

2021-09-07

Risky Learning with Imperfect Recall

Olson, Alaina

Olson, A. (2021). Risky Learning with Imperfect Recall (Master's thesis, University of Calgary, Calgary, Canada). Retrieved from <https://prism.ucalgary.ca>.

<http://hdl.handle.net/1880/113866>

Downloaded from PRISM Repository, University of Calgary

UNIVERSITY OF CALGARY

Risky Learning with Imperfect Recall

by

Alaina Olson

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE
DEGREE OF MASTER OF ARTS

GRADUATE PROGRAM IN ECONOMICS

CALGARY, ALBERTA

SEPTEMBER, 2021

© Alaina Olson 2021

Abstract

This paper proposes a model of endogenous memory recall in a general dynamic setting using the framework of rational inattention. I introduce the general model setup in chapter 2 through the lens of a rational Bayesian agent who faces no memory constraints. Building on this framework, I define a memory process in chapter 3. I then study how an exogenous memory process impacts the agent's optimal actions, and I discuss how the agent's choice of actions given her memory constraints can be used to approximate her risk preferences. Finally, endogenous memory costs enter into the agent's welfare maximization problem in chapter 4. These memory costs depend on the agent's memory process, and are amplified by an exogenously specified ability parameter. I show that in the presence of endogenous memory costs, the agent selects her memory process according to her given ability: higher ability agents choose sharper memory processes, make riskier decisions in the context of learning, and ultimately expect greater welfare.

Preface

This thesis is an original work by the author. No part of this thesis has been previously published.

Acknowledgements

I am deeply indebted to Dr. Alexander Jakobsen for his invaluable advice, patience and unwavering support over the course of the last three years. Dr. Jakobsen has afforded me every opportunity to pursue my theoretical interests and develop my skills along the way, and this project would not have been possible without his expert guidance. I would also like to thank my thesis committee members, Dr. Kenneth McKenzie and Dr. Dimitri Migrow, for insightful feedback and helpful discussions.

Finally, I would like to acknowledge the overwhelming amount of support I have had from my family and friends throughout my academic journey. I know haven't always made it easy for them, and I owe them a tremendous debt of gratitude for their continuing patience, support and encouragement. I am so incredibly fortunate to have such a solid support network (both inside and outside of academia) going into the next phase of my journey at the University of Toronto this fall.

Table of Contents

Abstract	ii
Preface	iii
Acknowledgements	iv
Table of Contents	v
List of Symbols, Abbreviations, and Nomenclature	vi
Epigraph	vii
1 Introduction	1
1.1 Background	1
1.2 Related Literature	3
2 The Rational Bayesian Agent	7
3 The Exogenous Memory Case	11
4 Endogenous Memory	13
5 Conclusion	23
Bibliography	25

List of Symbols, Abbreviations, and Nomenclature

Symbol	Definition
a	Vector representing actions a_1, a_2
$\ \cdot\ $	Euclidean norm
μ	Vector representing prior μ_1, μ_2
$v(\mu)$	Value of prior
$\mu \cdot a$	Dot product, representing one-shot expected payoff
δ	Discount factor on the open interval $(0, 1)$
d	Direct component of expected utility
$U(\cdot)$	Expected utility function
$p, 1 - p$	Memory process
$c(\cdot)$	Memory cost function from $[0, 1]$ to \mathbb{R}
$U(d^*(p)) - c(p)$	Welfare function
k	Ability parameter

Epigraph

I consider that a man's brain originally is like a little empty attic, and you have to stock it with such furniture as you choose. A fool takes in all the lumber of every sort that he comes across, so that the knowledge which might be useful to him gets crowded out, or at best is jumbled up with a lot of other things, so that he has a difficulty in laying his hands upon it. Now the skillful workman is very careful indeed as to what he takes into his brain-attic. He will have nothing but the tools which may help him in doing his work, but of these he has a large assortment, and all in the most perfect order. It is a mistake to think that that little room has elastic walls and can distend to any extent. Depend upon it there comes a time when for every addition of knowledge you forget something that you knew before. It is of the highest importance, therefore, not to have useless facts elbowing out the useful ones.

- Sir Arthur Conan Doyle, *A Study in Scarlet*

Chapter 1

Introduction

1.1 Background

Memory plays a crucial role in every aspect of our daily lives. We rely on our memory for everything from remembering where we left our keys, to forming our beliefs about ourselves and the world around us. In essence, memory helps shape who we are. However, a wide body of evidence from the psychology literature asserts that memory is highly imperfect; increasingly, this fact is supported by research in economics. As much as we rely on our memory to function, it can be a surprisingly fallible process. Nearly all of us are familiar with the experience of forgetting something important: perhaps we forget a friend's birthday on occasion, and many of us cannot recall events in our lives from before a certain age. Moreover, there has long been a thriving market for reminder apps, sticky notes, daily planners, and other memory-enhancing tools (Brenner and Brenner 1982).

While many instances of imperfect memory are relatively benign, sometimes the consequences may be more serious. Memory deficiencies can lead to faulty eyewitness testimony (and even erroneous convictions and acquittals) in criminal trials (Fisher, Brewer, and Mitchell 2009). Furthermore, experimental evidence suggests that limited memory impacts risk preferences (Muda, Niszczoła, and Augustynowicz 2020): better memory is associated with lower risk preferences. This suggests that poorer memory is associated with riskier choices, and potentially worse outcomes.

Economic theory is also concerned with the study of memory, from how we process information given memory constraints (Wilson 2014) to how we deal with complicated information such as complex contracts (Jakobsen 2020). Given its important role in making decisions, memory is clearly important to the study of economics.

This thesis adds a novel dimension to the existing theoretical literature in economics by exploring the impact of memory constraints on learning in a dynamic decision environment. We know that memory is fallible by nature, and that imperfect memory can be associated with suboptimal outcomes. Thus, it is natural to wonder how memory might be “managed” – in other words, where does memory fit into the problem of studying incentives?

Before we begin to explore this question further, it is important to note that “sharp” memory requires some degree of cognitive effort. To illustrate the trade-off between the benefits of sharper memory and the costs of cognitive resource use, consider a university student who is studying for a sequence of exams. She knows she wants to perform well, but she also knows that preparing for each of the tests will require significant cognitive effort. If she passes the exams, she gets what she wants – in other words, the “good” outcome. However, if she fails the exam – the “bad” outcome – she is worse off; in this way, the outcome of the test serves as a signal about the student’s ability. She does not know the outcome before taking her first action, which yields information about the true state of the world. Therefore, the student needs to decide how to allocate her cognitive resources in order to optimally prepare for the exams ahead of time. This simple example illustrates costly cognition in the context of memory and learning: the student stands to gain more when she retains more information (that is, with better memory), but she also faces costs of cognitive effort that constrain her overall payoff.

On a basic level, this example has three core properties: (1) an agent faces a decision problem multiple times; (2) making a decision in each period yields information about the true state of the world by way of observed payoffs; and (3) the agent has a memory problem – in other words, she only retains some of the information from her decisions in previous periods.

Accordingly, this model can be interpreted as an agent who chooses how to use her memory of past decisions to learn how to make better decisions in the future. Instead of treating memory as a signal about the state of the world, I interpret memory as a signal about the history of past choices and their consequences. The notion of a *memory process* helps to further conceptualize the problem. A memory process can be thought of as an information manipulator: an exogenous memory process acts like an information filter, while an endogenous process distorts the filter itself. In the case of an exogenous memory process, an agent loses certain kinds of information while retaining others; for example, she may remember information that supports her beliefs, and forget information that contradicts them. However, it seems unrealistic to assume that memory processes are uniformly exogenous: for example, it seems plausible that people might put more effort into recalling informative events when stakes are higher (such as in the example of the student preparing for an exam).

Chapter 2 introduces the model by examining the behaviour of a rational Bayesian agent who faces her

intertemporal decision problem in two periods, and in the absence of memory constraints. The rational Bayesian case captures the agent’s optimal strategy from her perspective of not having learned yet, but understanding that she will process information and choosing her actions accordingly. As the agent is unconstrained by her memory in this context, the rational Bayesian case provides insight into how the agent would behave if she did not face such constraints.

Chapter 3 introduces an exogenous memory process in which the agent is either a rational Bayesian with some probability, or does not update her beliefs after observing information with some other probability. With exogenously constrained memory, the agent again faces her decision problem in two periods; however, she potentially has different information on which to base her decisions (depending on her exogenously specified memory process). In this way, we can infer that, depending on her exogenous memory process, the agent may be learning from her past actions and adjusting her behaviour accordingly.

Finally, endogenous memory formation is addressed in Chapter 4. In the endogenous case, I model memory formation as a form of rational inattention in which “finer” (more detailed) memory requires greater cognitive resources and, by extension, incurs greater cognitive costs. When memory is costly and endogenous, I derive five main results. First, when memory is costless, the agent will choose perfect recall; however, if memory costs are sufficiently high, the agent will bypass these costs by forgetting information from previous periods. Moreover, when memory costs are restricted to a range of positive values, the agent will trade off between improving her memory and exhausting her cognitive resources. Ultimately, higher memory costs lead the agent to choose less risky actions, which reduces her welfare by limiting her potential benefit from information.

1.2 Related Literature

Once we receive new information, we can store it for later use; this feature is fundamental to the current understanding of memory. Moreover, our cumulative experience of the world around us is guided by our beliefs about our observations and experiences. Given that our perspectives are profoundly shaped by our beliefs, which depend on how we process and store the information we receive, it is worthwhile to wonder how our memories shape our beliefs.

This model can be interpreted as an addition to the broad theoretical literature on non-Bayesian updating in economics. It is well known that belief updating patterns are well-approximated by Bayes’ rule: when faced with new information, we generally incorporate our new knowledge into what we already believe about the world. However, examples of belief updating that deviate from the standard Bayesian method are well documented in economics. From a theoretical perspective, Rabin and Schrag (1999) propose a model in

which an agent misinterprets novel information as confirming beliefs he already holds. In this model, there are two states of the world, and while the agent initially believes that either state is the true state with equal probability, the distribution of signals the agent receives (and perceives) influence the beliefs at which he ultimately arrives.

Examples of non-Bayesian updating have also been explored in the experimental economics literature. Eil and Rao (2011) study the interplay between preferences and beliefs, and find that individuals tend to incorporate positive signals into their beliefs in line with the standard Bayesian model, but they also downplay the strength of negative signals (thereby acting in a non-Bayesian manner). In other words, people respond differently to information that supports their beliefs than they do to information that challenges their preconceptions.

How we use our memory to shape our beliefs is governed by the process of memory retrieval. Psychologists recognize two forms of memory retrieval: recognition, and recall. Recognition involves awareness of a stimulus item (something that one has seen or experienced before). Recall, on the other hand, requires us to reproduce a stimulus item (Baddeley 2004). Clearly, recall forces us to think in some detail about what we have experienced or seen.

However, thinking requires significant cognitive effort; this should not be a controversial statement. In economics, the study of costly cognition is facilitated using the framework of rational inattention. Sims (2003) develops the canonical model of rational inattention in economics with the goal of capturing the idea that people have limited capacity for processing information; many subsequent papers in recent years have applied the rational inattention framework to study various dynamic behavioural problems. For example, Gossner, Steiner and Stewart (2021) study how attention can be manipulated when an agent receives sequential information before making a decision, and find that focusing attention on a particular item increases its favourability regardless of its actual value. In other words, people tend to favour things that grab their attention, perhaps because such things reduce their cognitive workload. As is the case with attention, memory requires some degree of cognitive effort; as a costly cognitive process, memory fits neatly into the rational inattention framework.

Perhaps (at least in part) because of its potential costs, the process of memory recall is highly imperfect. The limitations of memory have important implications for decisions. Piccione and Rubinstein (1997) illustrate this fact through the paradox of the absent-minded driver: an agent who does not distinguish between repeated occurrences of their decision problem exhibits imperfect recall, which constrains his payoff at the end of the game. I adapt this formulation of imperfect memory recall to include the possibility that an agent can be aware of her own absent-mindedness, and subsequently seek to remedy it.

The assumption that an agent can be aware of (and modify) her memory capacity is rooted in a phe-

nomenon known in psychology as *metamemory* (Flavell and Wellman 1975). Metamemory encompasses our own practical awareness of our memory capacities, the extent to which they are malleable, and what we must do to facilitate the act of remembering. Accordingly, metamemory can be broken down into three broad categories: knowledge of one’s own characteristics related to the act of remembering, knowledge about differences in tasks that are important for how they are remembered, and knowledge of strategies that are useful in remembering (Flavell and Wellman 1975). For the purposes of my model, I assume the agent is aware of her limited memory capacity, and can employ memory-enhancing strategies to improve her memory.

Several strands of the literature in psychology view memory as an ongoing record of processes occurring during learning (Wingfield and Byrnes 2013), and there is broad agreement that memory is inextricably linked to learning. As such, an agent who is aware of her own limited memory and wants to improve it over time can perhaps be thought of as wanting to “learn how to learn”. The theory of multi-armed bandit problems is highly applicable to the study of optimal sequential learning (Gittins and Jones (1974); (1979)). More specifically, multi-armed bandit problems can be thought of in terms of an agent choosing how to allocate a fixed set of resources across competing options so as to maximize her expected outcome over time. This is commonly referred to as the “exploration-exploitation” trade-off: exploration involves trying different actions that could potentially provide different information and lead to the discovery of the optimal action (since payoffs depend on the true state), while exploitation involves repeatedly choosing the action with the best outcome in order to maximize payoffs. In particular, incentives for exploration depend on the agent’s memory of past information.

Similar to other models of limited memory in the theoretical economics literature, this model incorporates the notion that rehearsal facilitates memory (and, by extension, learning): recalling an event once makes it easier to retrieve the memory in the future. Mullainathan (2002) proposes a model in which imperfect memory distorts actions by way of influencing beliefs. The agent operates in a finite number of periods, receives belief shocks that are either transitory or permanent, and information received is either “hard” (regarding past events) or “soft” (concerning new events). Modeling memory as a stochastic map from the true history to the agent’s perceived history, this paper finds that agents with limited memory tend to overreact to memorable events and underreact (or forget) when there is a significant information lag.

Wilson (2014) proposes an alternative conception of memory, wherein an agent is faced with an infinite-horizon problem but a finite set of memory states. The agent in this context observes a sequence of informative signals, and combines these signals into a memory state that is subject to updating when presented with new information. From this perspective, limited memory impacts incentives, leading to common biases in information processing such as confirmation bias and belief polarization.

I depart from these formulations of memory in several important ways. First, rather than treating

memory as uniformly exogenous, my model is based on the idea that incentives and memory can jointly influence each other. In essence, I focus on memory as an endogenous process: memory affects outcomes, but if outcomes can be improved with better memory, then outcomes can also motivate incentives to affect the memory process. Moreover, I do not restrict attention to a finite set of memory states. Rather, I address memory as a continuous process in which the agent either learns to work with the memory states she is given, as in the exogenous memory case, or optimally influences her memory in light of her cognitive constraints, as in the endogenous memory case.

Chapter 2

The Rational Bayesian Agent

Consider a rational Bayesian agent who faces a decision problem at two distinct points in time. There are two states of the world, and the agent can observe two possible outcomes: a “good” outcome and a “bad” outcome. The agent has a utility of 1 from the good outcome, and 0 from the bad outcome. Moreover, the agent chooses an action, which determines the probability of observing the good outcome in different states. Framed from this perspective, actions are Anscombe-Aumann acts (Anscombe and Aumann 1963), which are described by a coordinate pair (a_1, a_2) , where a_i denotes the probability that the agent receives the good outcome in state i . In selecting an action, the agent is effectively choosing an information structure: a_1 and a_2 represent probabilities that, along with her realized outcome, provide a signal about the state of the world.

In general, this setup can be applied to a compact, convex set of such pairs; for ease of analysis, we restrict attention to pairs of action contained in the unit circle. More succinctly, $a_1, a_2 \in A = \{(a_1, a_2) : \|a\|^2 = 1\}$, where the notation $\|\cdot\|$ refers to the Euclidean norm.

While the agent faces her decision problem exactly twice, we consider three periods of time, $t \in \{0, 1, 2\}$. The agent chooses an action in period 0, and observes the payoff from that action at the start of period 1. After choosing an action, the agent subsequently updates her prior $\mu = (\mu_1, \mu_2)$ in period 1 according to Bayes’ rule. After revising her beliefs, the agent takes another action based on her posterior at the end of the same period. Finally, the game concludes when the agent observes her payoff from the second action at the start of period 2.

In a one-shot setting, the agent maximizes her expected utility $\mu \cdot a$ subject to $\|a\|^2 = 1$. Straightforward calculations show that her optimal strategy is given by $a^*(\mu) = \frac{\mu}{\|\mu\|}$. In light of her prior beliefs, the value of the agent’s optimal strategy is given by $v(\mu) = a^*(\mu) \cdot \mu$. This simply reduces to $v(\mu) = \|\mu\|$.

It is worth emphasizing that a rational Bayesian agent chooses her strategy with the knowledge that she will update her prior after observing her first payoff. Her posteriors then serve as informative signals that help her choose an action in the following period. If the agent observes the good payoff in period 1, she updates her beliefs according to $\hat{\mu}(1) = \frac{\mu a}{\mu \cdot a}$ with probability $\mu \cdot a$; here, we draw attention to the fact that μa refers to the component-wise vector multiplication of beliefs and actions in each period; that is, $\mu a = (\mu_1 a_1, \mu_2 a_2)$. Analogously, observing the bad payoff causes the agent to update her beliefs according to $\hat{\mu}(0) = \frac{\mu(e-a)}{\mu \cdot (e-a)}$ with probability $\mu \cdot (e - a)$ (where $e = (1, 1)$ denotes the uninformative signal). To keep things simple, I analyze the model for the case where the agent has a uniform prior; that is, when $\mu = (\frac{1}{2}, \frac{1}{2})$.

There are two ways we could determine the agent's choice in period 0 and study her resulting payoff. One method is to start at the end of the game, and then proceed via backward induction to determine what she will choose at the start. Alternatively, we can view the agent as "forward looking", and construct her reasoning and preferences from the perspective of period 0. The two methods produce identical results when the agent is a rational Bayesian, but the "forward looking" approach allows us to study a non-Bayesian agent's behaviour.

The agent holds beliefs in every period, and standard Bayesian equilibrium concepts require those beliefs to be correct given the history of play up to that time. A non-Bayesian agent, however, may hold beliefs that are incorrect at some histories. To bypass this problem, we assume that the agent correctly forecasts her own belief updating and decision behaviour. Accordingly, the agent is still able to perform expected utility assessments.

Definition 1. An agent is *forward-looking* if her period-0 preferences can be captured by an expected utility function of the form

$$U(a) = \mu \cdot a + \delta \left((\mu \cdot a) v \left(\frac{\mu a}{\mu \cdot a} \right) + (\mu \cdot (e - a)) v \left(\frac{\mu(e - a)}{\mu \cdot (e - a)} \right) \right).$$

The forward-looking expected utility function can be broken down into two parts: $\mu \cdot a$ represents the agent's expected payoff at time 1, and $(\mu \cdot a) v \left(\frac{\mu a}{\mu \cdot a} \right) + (\mu \cdot (e - a)) v \left(\frac{\mu(e - a)}{\mu \cdot (e - a)} \right)$ is the expected payoff at time 2. The time-2 payoff further depends on two factors. First, $\mu \cdot a$ denotes the probability that the agent observes the good outcome, and updates her prior to $\frac{\mu a}{\mu \cdot a}$ according to Bayes' rule. Similarly, $\mu \cdot (e - a)$ corresponds to the probability that the agent observes the bad outcome, and updates her prior to $\frac{\mu(e - a)}{\mu \cdot (e - a)}$. With these considerations in mind, $U(a)$ specifies the agent's expected utility from choosing an action in period 0, given that her period-1 decisions will maximize her expected utility conditional on information generated by her period-1 choice. The future-period expected payoff is discounted by $\delta \in (0, 1)$.

Definition 2. Let $U(a) = \mu \cdot a + \delta(\|\mu a\| + \|\mu(e - a)\|)$.

1. The *direct component* $d := \mu \cdot a$ is the agent's expected payoff at $t = 1$.
2. The *information component* $\|\mu a\| + \|\mu(e - a)\|$ is the agent's expected payoff at $t = 2$.

The following result demonstrates that, given a uniform prior, choosing an optimal strategy is analogous to choosing an optimal direct component.

Proposition 1. Consider a forward-looking agent with a uniform prior. The agent's problem of choosing an optimal strategy a can be transformed into a problem of choosing an optimal direct component d .

Proof. Suppose the agent is forward looking and has a uniform prior. Then the agent's preferences are captured by $U(a) = \mu \cdot a + \delta(\|\mu a\| + \|\mu(e - a)\|)$; assuming a uniform prior, this becomes

$$\begin{aligned}
U(a) &= \mu \cdot a + \delta(\|\mu a\| + \|\mu(e - a)\|) \\
&= \frac{1}{2}(a_1 + a_2) + \delta \left(\sqrt{\left(\frac{a_1}{2}\right)^2 + \left(\frac{a_2}{2}\right)^2} + \sqrt{\left(\frac{1-a_1}{2}\right)^2 + \left(\frac{1-a_2}{2}\right)^2} \right) \\
&= \frac{1}{2}(a_1 + a_2) + \delta \left(\sqrt{\left(\frac{1}{2}\right)^2 (a_1^2 + a_2^2)} + \sqrt{\left(\frac{1}{2}\right)^2 ((1-a_1)^2 + (1-a_2)^2)} \right) \\
&= \frac{1}{2}(a_1 + a_2) + \delta \left(\frac{1}{2}\sqrt{a_1^2 + a_2^2} + \frac{1}{2}\sqrt{(1-a_1)^2 + (1-a_2)^2} \right) \\
&= \frac{1}{2}(a_1 + a_2) + \frac{\delta}{2}(\|a\| + \|e - a\|) \\
&= \frac{1}{2}(a_1 + a_2) + \frac{\delta}{2} \left(\|a\| + \sqrt{(1-a_1)^2 + (1-a_2)^2} \right) \\
&= \frac{1}{2}(a_1 + a_2) + \frac{\delta}{2} \left(\|a\| + \sqrt{(1-2a_1+a_1^2) + (1-2a_2+a_2^2)} \right) \\
&= \frac{1}{2}(a_1 + a_2) + \frac{\delta}{2} \left(\|a\| + \sqrt{(a_1^2 + a_2^2) - 2(a_1 + a_2) + 2} \right) \\
U(a) &= \frac{1}{2}(a_1 + a_2) + \frac{\delta}{2} \left(\|a\| + \sqrt{\|a\| - 2(a_1 + a_2) + 2} \right).
\end{aligned}$$

The direct component becomes $d = \frac{1}{2}(a_1 + a_2)$ with a uniform prior; moreover, the agent's utility is constrained by $\|a\|^2 = 1$. Making the appropriate substitutions allows us to reduce the dimension of $U(\cdot)$, so that it can be expressed as a single-variable problem in terms of d :

$$\begin{aligned}
U(d) &= d + \frac{\delta}{2} \left(1 + \sqrt{1 - 2(2d) + 2} \right) \\
U(d) &= d + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d} \right).
\end{aligned}$$

□

The optimal direct component d^* is found by maximizing $U(d)$ with respect to d ; that is,

$$\max_d U(d) = 1 + \frac{\delta}{2} \left(\frac{1}{2} (3 - 4d)^{-1/2} (-4) \right).$$

Setting the derivative equal to zero, we have

$$0 = 1 - \delta (3 - 4d)^{-1/2}$$

$$1 = \delta (3 - 4d)^{-1/2}$$

$$\delta = (3 - 4d)^{-1/2}$$

$$\delta^2 = 3 - 4d$$

$$4d = 3 - \delta^2$$

$$d^* = \frac{3 - \delta^2}{4}.$$

Substituting d^* into the expected utility function yields the agent's expected payoff from maximizing the direct component.

$$\begin{aligned} U(d^*) &= d^* + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*} \right) \\ &= \frac{3 - \delta^2}{4} + \frac{\delta}{2} \left(1 + \sqrt{3 - 4 \left(\frac{3 - \delta^2}{4} \right)} \right) \\ &= \frac{3 - \delta^2}{4} + \frac{\delta(1 + \delta)}{2}. \end{aligned}$$

Framed this way, it is easy to study the agent's payoff from each component of her expected utility function. First, we notice that a lower value of δ indicates that the agent more heavily discounts her future payoff. Therefore, agent's expected payoff in period 1 is decreasing in δ , and a higher value of δ means that she cares more about her period-2 payoff than for lower values of δ .

Note that the optimal direct component cannot be larger than $d^* = \frac{1}{\sqrt{2}}$, given the agent's constraint $\|a\|^2 = 1$. Thus, setting $d^* = \frac{1}{\sqrt{2}}$, we see that the agent must discount her future payoff to such an extent that $\delta \geq \frac{1}{\sqrt{2}}$. Moreover, when $\delta = 1$, the information component dominates the direct component, so the smallest feasible direct component is optimal; in other words, when the agent does not discount her future payoff at all, she always chooses to acquire the most information that she can. As such, we assume that $\delta \in \left(\frac{1}{\sqrt{2}}, 1 \right)$ to ensure that the agent can choose a direct component that falls between the two extremes.

Chapter 3

The Exogenous Memory Case

Now that we have established how a rational Bayesian agent would behave, we are ready to introduce the idea of a memory process.

Definition 3. Let $p \in [0, 1]$ denote the probability that the agent updates her beliefs according to Bayes' rule; similarly, let $1 - p$ denote the probability that the agent does not update her prior. The parameter p captures a *memory process*.

When the agent updates her prior according to Bayes' rule, she benefits from information via the information component of her expected utility function. Similarly, when the agent does not update her beliefs, she does not acquire any extra information. Defining memory this way allows us to incorporate memory into the agent's expected utility function. That is, by definition of a memory process, the agent observes

$$U(d) = d + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d}\right) \text{ with probability } p, \text{ and}$$
$$U(d) = (1 - p)(1 + \delta)d \text{ with probability } (1 - p).$$

Across memory processes, the agent's expected payoff is

$$U(d) = p \left(d + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d}\right) \right) + (1 - p)(1 + \delta)d$$

As in the rational Bayesian case, maximizing $U(d)$ allows us to find the optimal direct component:

$$\max_d U(d) = p + \frac{\delta}{2} \left(\frac{1}{2} (3 - 4d)^{-1/2} (-4) \right) - p(1 + \delta) + (1 + \delta)$$

Setting the derivative equal to zero, we have

$$\begin{aligned}
0 &= p - \delta p (3 - 4d)^{-1/2} + (1 + \delta)(1 - p) \\
\delta p (3 - 4d)^{-1/2} &= p - (1 + \delta)(1 - p) \\
\frac{\delta p}{1 + \delta(1 - p)} &= (3 - 4d)^{-1/2} \quad (\text{note: } p - (1 + \delta)(1 - p) = 1 + \delta(1 - p)) \\
\left(\frac{\delta p}{1 + \delta(1 - p)} \right)^2 &= 3 - 4d \\
4d &= 3 - \left(\frac{\delta p}{1 + \delta(1 - p)} \right)^2 \\
d^* &= \frac{1}{4} \left(3 - \left(\frac{\delta p}{1 + \delta(1 - p)} \right)^2 \right)
\end{aligned}$$

Substituting d^* into the expected utility function, we get

$$U(d^*) = p \left(d^* + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*} \right) \right) + (1 - p)(1 + \delta) d^*,$$

which simplifies to

$$\begin{aligned}
U(d^*) &= \frac{p}{4} \left(3 - \left(\frac{\delta p}{1 + \delta(1 - p)} \right)^2 \right) + \frac{\delta p}{2} \left(\frac{1 + \delta}{1 + \delta(1 - p)} \right) + \\
&\quad \frac{(1 - p)(1 + \delta)}{4} \left(3 - \left(\frac{\delta p}{1 + \delta(1 - p)} \right)^2 \right)
\end{aligned}$$

Taking p as given, we can look at the comparative statics for d^* and $U(d^*)$. Clearly, d^* is decreasing in p , and $U(d^*)$ is increasing in p .

The intuition behind these observations is straightforward: better memory induces the agent to sacrifice more of her period-1 payoff in favour of a potentially higher payoff in period 2. In this way, the agent's choice of optimal direct component captures her risk preferences: choosing a lower direct component is akin to engaging in riskier behaviour, since the agent is willing to trade off more of her immediate expected payoff in favour of potentially doing better in the future.

The agent is more likely to treat her memory process as exogenous if she feels she cannot do anything to influence her memory. However, when the stakes are higher – perhaps when she can potentially improve her outcome by doing so – the agent might be concerned with choosing her memory process more carefully to enhance her learning. The following section endogenizes the memory process so that we may study such behaviour.

Chapter 4

Endogenous Memory

Finally, we consider the case of endogenous memory recall. To endogenize the memory process, we proceed via backward induction. First, recall that in the exogenous memory case, we found d^* in terms of p ; in the endogenous memory case, d^* is a function of p . In other words,

$$d^*(p) = \frac{1}{4} \left(3 - \left(\frac{\delta p}{1 + \delta(1-p)} \right)^2 \right)$$

Substituting $d^*(p)$ into the utility function, we end up with the value function

$$U(d^*(p)) = p \left(d^*(p) + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*(p)} \right) \right) + (1 + \delta)(1 - p)d^*(p).$$

However, we know that improving cognitive processes, including memory, generally comes with some cost; at this point, we are ready to introduce endogenous memory costs into the model. In this setting, memory costs constrain the agent's acquisition of information.

Definition 4. A *memory cost* is a function $c : [0, 1] \rightarrow \mathbb{R}$ that depends on a memory process p .

To study the effects of memory costs explicitly, we will consider the case where $c(p) = kp^2$. In this context, k is interpreted as a measure of cognitive ability: as k increases, the agent's cognitive ability decreases. Intuitively, if the agent has a higher level of cognitive ability, she will not have to exert as much effort to improve her memory, and will face lower costs in doing so. Taken together, the value function $U(d^*(p))$ and memory costs capture the tradeoff between the agent's benefits of sharper memory versus the costs of achieving it. Thus, the agent must decide how to optimally select p to balance these two competing forces.

It is helpful to think of the optimal selection of p as occurring before the remaining steps in the game

are carried out; that is, when the agent selects her optimal memory process, she knows exactly how she will select her future-period actions based on this choice. The following definition specifies the objective function that weighs the costs and benefits of acquiring a memory process.

Definition 5. A forward-looking agent has *rationally inattentive memory recall* if her preferences take the form

$$U(d^*(p)) - c(p) = p \left(d^*(p) + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*(p)} \right) \right) + (1 + \delta) (1 - p) d^*(p) - c(p).$$

Furthermore, $U(d^*(p)) - c(p)$ captures the agent's welfare; that is, her overall expected payoff from her choice of memory process and the actions she subsequently takes.

In the case where $c(p) = kp^2$,

$$U(d^*(p)) - c(p) = p \left(d^*(p) + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*(p)} \right) \right) + (1 + \delta) (1 - p) d^*(p) - kp^2.$$

The next set of results outline the circumstances in which it is optimal for the agent to choose either perfect memory recall, or no recall at all.

Proposition 2. A forward looking agent with rationally inattentive memory recall will choose $p = 1$ when $k = 0$.

Proof. Suppose $k = 0$. Then the agent simply maximizes $U(d^*(p))$; that is,

$$U(d^*(p)) = p \left(d^*(p) + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*(p)} \right) \right) + (1 + \delta) (1 - p) d^*(p).$$

To establish that perfect information is optimal in this case, we only need to look at the components of the utility function. The direct component, $d^*(p) = \frac{1}{4} \left(3 - \left(\frac{\delta p}{1 + \delta(1 - p)} \right)^2 \right)$, is decreasing in p ; however, the information component is increasing in p (by virtue of also being decreasing in $d^*(p)$).

When $d^*(p)$ is maximized, $d^*(p) = \frac{1}{\sqrt{2}}$ and

$$\begin{aligned} \frac{1}{2} \left(1 + \sqrt{3 - 4d^*(p)} \right) &= \frac{1}{2} \left(1 + \sqrt{3 - 4 \left(\frac{1}{\sqrt{2}} \right)} \right) \\ &= \frac{1}{2} \left(1 + \sqrt{3 - 2\sqrt{2}} \right) \\ &= \frac{\sqrt{2}}{2} \\ &= \frac{1}{\sqrt{2}}. \end{aligned}$$

In other words, the direct and information components of the utility function contribute equally to the agent's payoff when the direct component is maximized. However, for higher values of p , the information component contributes more to the agent's payoff than the direct component (by virtue of their relationships with p). This implies that, in the absence of any costs associated with increasing p , the agent benefits from choosing the highest possible p , since a higher p ensures a larger payoff from increasing information. \square

Proposition 3. A forward-looking agent with rationally inattentive memory recall will choose $p = 0$ for all $p \in (0, p^*)$; that is, for some $p^* \in (0, 1)$, $U(d^*(0)) - c(0) > U(d^*(p)) - c(p)$ for all $p < p^*$.

Proof. Suppose the agent is forward-looking with rationally inattentive memory recall. Let $p = 0$. Then

$$\begin{aligned} d^*(0) &= \frac{1}{4} \left(3 - \left(\frac{\delta p}{1 + \delta(1-p)} \right)^2 \right) \\ &= \frac{1}{4} \left(3 - \left(\frac{0}{1 + \delta(1-0)} \right)^2 \right) \\ &= \frac{3}{4}. \end{aligned}$$

However, the direct component is maximized at $d^*(p) = \frac{1}{\sqrt{2}}$; therefore,

$$d^*(0) = \frac{1}{\sqrt{2}}$$

When $p = 0$, the agent's expected payoff is

$$\begin{aligned} U(d^*(0)) - c(0) &= 0 \left(\frac{1}{\sqrt{2}} + \frac{\delta}{2} \left(1 + \sqrt{3 - 4 \left(\frac{1}{\sqrt{2}} \right)} \right) \right) + \frac{(1 + \delta)(1 - 0)}{\sqrt{2}} - k(0)^2 \\ &= \frac{1 + \delta}{\sqrt{2}}. \end{aligned}$$

However, we also have $d^*(p) = \frac{1}{\sqrt{2}}$ when

$$\begin{aligned} d^*(p) &= \frac{1}{4} \left(3 - \left(\frac{\delta p}{1 + \delta(1-p)} \right)^2 \right) \\ \frac{1}{\sqrt{2}} &= \frac{1}{4} \left(3 - \left(\frac{\delta p}{1 + \delta(1-p)} \right)^2 \right) \\ \sqrt{2} &= 3 - \left(\frac{\delta p}{1 + \delta(1-p)} \right)^2 \\ \frac{\delta p}{1 + \delta(1-p)} &= \sqrt{3 - 2\sqrt{2}} \\ \delta p &= (1 + \delta(1-p)) \left(\sqrt{3 - 2\sqrt{2}} \right) \end{aligned}$$

$$p = \left(\frac{1+\delta}{\delta} \right) \left(\frac{3-2\sqrt{(2)}}{2} \right)^{1/2}.$$

In other words,

$$d^* \left(\left(\frac{1+\delta}{\delta} \right) \left(\frac{3-2\sqrt{(2)}}{2} \right)^{1/2} \right) = \frac{1}{\sqrt{2}}.$$

Let $p^* = \left(\frac{1+\delta}{\delta} \right) \left(\frac{3-2\sqrt{(2)}}{2} \right)^{1/2}$. The agent's expected payoff from choosing p^* is given by

$$\begin{aligned} U(d^*(p^*)) - c(p^*) &= p^* \left(\frac{1}{\sqrt{2}} + \frac{\delta}{2} \left(1 + \sqrt{3 - 4 \left(\frac{1}{\sqrt{2}} \right)} \right) \right) + \frac{(1+\delta)(1-p^*)}{\sqrt{2}} - kp^{*2} \\ &= p^* \left(\frac{1}{\sqrt{2}} + \frac{\delta}{2} \left(1 + \sqrt{3 - 2\sqrt{2}} \right) - \frac{(1+\delta)}{\sqrt{2}} \right) + \frac{(1+\delta)}{\sqrt{2}} - kp^{*2} \\ &= p^* \left(\frac{1}{\sqrt{2}} + \frac{\sqrt{2}\delta}{2} - \frac{(1+\delta)}{\sqrt{2}} \right) + \frac{(1+\delta)}{\sqrt{2}} - kp^{*2} \\ &= p^* \left(\frac{(1+\delta) - (1+\delta)}{\sqrt{2}} \right) + \frac{(1+\delta)}{\sqrt{2}} - kp^{*2} \\ &= p^*(0) + \frac{(1+\delta)}{\sqrt{2}} - kp^{*2} \\ &= \frac{1+\delta}{\sqrt{2}} - kp^{*2}. \end{aligned}$$

Proposition 2 establishes that the agent will choose $p = 1$ if $k = 0$, so we must have $k > 0$; therefore,

$$\frac{1+\delta}{\sqrt{2}} > \frac{1+\delta}{\sqrt{2}} - kp^{*2}, \text{ which implies}$$

$$U(d^*(0)) - c(0) > U(d^*(p)) - c(p) \text{ for all } p < p^*.$$

□

Proposition 3 provides us with some preliminary insight into the properties of the welfare function. In particular, if the agent chooses to maximize the direct component for some range of p , we can infer how her preferences behave for such a range of p close to 0.

Corollary 3.1. $U(d^*(p)) - c(p)$ is constant on the interval $[0, p^*]$.

This implication is important to note for further analysis. An expected payoff that is constant on a given interval does not achieve a maximum on that interval; in other words, the agent would never choose a memory process that would make her worse off than if she chose to avoid information altogether.

Proposition 4. A forward-looking agent with rationally inattentive memory recall will choose $p = 0$ for sufficiently large values of k .

Proof. Rearranging the expected payoff function slightly, we see that

$$U(d^*(p)) - c(p) = p \left(d^*(p) + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*(p)} \right) - kp \right) + (1 + \delta)(1 - p)d^*(p).$$

Suppose k is sufficiently large that $kp > d^*(p) + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*(p)} \right)$ for all non-zero values of p . Then

$$p \left(d^*(p) + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*(p)} \right) - kp \right) < 0,$$

which implies that $U(d^*(p)) - c(p) < U(d^*(0)) - c(0) = \frac{1+\delta}{\sqrt{2}}$ for any strictly positive value of p . Thus, the agent receives the highest expected payoff from choosing $p = 0$ for sufficiently large k . \square

This result has a simple interpretation. When the agent chooses $p = 0$, she never updates her beliefs, so she does not benefit from information and effectively repeats her one-shot problem in both decision periods (much like the absent-minded driver in Piccione and Rubinstein (1997)) Proposition 4 essentially tells us that when the agent chooses $p = 0$, she behaves as if she is avoiding riskier actions for which the costs outweigh the benefits.

Now that we have examined the behaviour of the agent's preferences at the endpoints, the next definition helps us pave the way to studying memory that is not "all-or-nothing".¹

Definition 6. Any $p \in (p^*, 1)$ that maximizes $U(d^*(p)) - c(p)$ is an *interior solution*.

When we think about memory and its limitations, we understand that it is not reasonable to restrict ourselves to the boundary cases (where the agent either remembers everything, or forgets everything). Memory is clearly not an all-or-nothing process: for example, we do not forget where we left our keys every time we want to leave the house, but we also do not remember every time either. The next result captures the relationship between memory and ability in the context of learning.

Proposition 5. Let $p(k)$ denote the optimal memory process for a given k . A forward looking agent with rationally inattentive memory recall will choose $p(k) \in (p^*, 1)$ for a range of positive k .

Proof. To show that an interior solution exists for some k , we first need to establish the first-order condition

¹It is worth mentioning that, in the real world, we do not directly choose our memory processes; rather, the selection of p arises from a gradual process in which we learn from our previous actions. We can think of the cost function $c(p)$ and the ability parameter k as a sort of evolutionary constraint, where nature acts as a "planner" that determines the agent's long-run outcome and the agent herself acts as a "doer" who is primarily concerned with her immediate payoff.

for an optimal solution on $(0,1)$. First, we notice that

$$U(d^*(p)) - c(p) = p \left(d^*(p) + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*(p)} \right) \right) + (1 + \delta)(1 - p)d^*(p) - kp^2$$

is continuous on $[0, 1]$ and differentiable on $(0, 1)$ in p and k . Thus, for an optimal p on $(0, 1)$,

$$\begin{aligned} \frac{d}{dp}U(d^*(p)) - c'(p) &= 0, \text{ which implies} \\ c'(p) &= \frac{d}{dp}U(d^*(p)). \end{aligned}$$

By the envelope theorem, $\frac{d}{dp}U(d^*(p)) = \frac{d}{dp}U(d)$ evaluated at $d = d^*(p)$. Therefore,

$$\begin{aligned} 2kp &= d^*(p) + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*(p)} \right) - (1 + \delta)d^*(p) \\ &= (1 - 1 - \delta)d^*(p) + \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*(p)} \right) \\ &= \frac{\delta}{2} \left(1 + \sqrt{3 - 4d^*(p)} \right) - \delta d^*(p). \end{aligned}$$

Substituting for $d^*(p)$ gives us our first-order condition for an interior solution in terms of k and p :

$$\begin{aligned} 2kp &= \frac{\delta}{2} \left(1 + \sqrt{3 - 4 \left(\frac{1}{4} \left(3 - \left(\frac{\delta p}{1 + \delta(1 - p)} \right)^2 \right) \right)} \right) - \frac{\delta}{4} \left(3 - \left(\frac{\delta p}{1 + \delta(1 - p)} \right)^2 \right) \\ &= \frac{\delta}{2} \left(1 + \frac{\delta p}{1 + \delta(1 - p)} \right) - \frac{\delta}{4} \left(3 - \left(\frac{\delta p}{1 + \delta(1 - p)} \right)^2 \right) \\ &= \frac{\delta}{2} \left(\frac{1 + \delta}{1 + \delta(1 - p)} \right) - \frac{\delta}{4} \left(3 - \left(\frac{\delta p}{1 + \delta(1 - p)} \right)^2 \right). \end{aligned}$$

Let $K = \{k : \frac{d}{dp}U(d^*(p)) - c'(p) \leq 0\}$. Suppose $p(k) \in (p^*, 1)$ maximizes $U(d^*(p)) - c(p)$ for a fixed value of k . The first-order condition for an interior solution tells us that, at $p(k)$,

$$\begin{aligned} \frac{d}{dp}U(d^*(p)) - c'(p) &= \frac{\delta}{2} \left(\frac{1 + \delta}{1 + \delta(1 - p(k))} \right) - \frac{\delta}{4} \left(3 - \left(\frac{\delta p(k)}{1 + \delta(1 - p(k))} \right)^2 \right) - 2kp(k) \\ &= 0. \end{aligned}$$

By Corollary 3.1, and since $U(d^*(p)) - c(p)$ is continuous, choose $p', p'' \in (p^*, 1)$ such that $p' < p(k) < p''$.

Since $p(k)$ maximizes $U(d^*(p)) - c(p)$ by assumption, then at $p' < p(k)$,

$$0 < \frac{\delta}{2} \left(\frac{1 + \delta}{1 + \delta(1 - p')} \right) - \frac{\delta}{4} \left(3 - \left(\frac{\delta p'}{1 + \delta(1 - p')} \right)^2 \right) - 2kp'.$$

Similarly, at $p'' > p(k)$,

$$0 > \frac{\delta}{2} \left(\frac{1 + \delta}{1 + \delta(1 - p'')} \right) - \frac{\delta}{4} \left(3 - \left(\frac{\delta p''}{1 + \delta(1 - p'')} \right)^2 \right) - 2kp''.$$

Rearranging these conditions to isolate k , we have

$$k < \frac{\delta}{4p'} \left(\frac{1 + \delta}{1 + \delta(1 - p')} \right) - \frac{\delta}{8p'} \left(3 - \left(\frac{\delta p'}{1 + \delta(1 - p')} \right)^2 \right), \text{ and}$$

$$k > \frac{\delta}{4p''} \left(\frac{1 + \delta}{1 + \delta(1 - p'')} \right) - \frac{\delta}{8p''} \left(3 - \left(\frac{\delta p''}{1 + \delta(1 - p'')} \right)^2 \right).$$

Given assumptions about δ and p' , it is clear that $U(d^*(p)) - c(p)$ is increasing for some positive k . Then, by the Intermediate Value Theorem, there will be at least one value $p(k) \in (p^*, 1)$ such that $\frac{d}{dp}U(d^*(p)) - c'(p) = 0$ whenever

$$k \in \left(\frac{\delta}{4p''} \left(\frac{1 + \delta}{1 + \delta(1 - p'')} \right) - \frac{\delta}{8p''} \left(3 - \left(\frac{\delta p''}{1 + \delta(1 - p'')} \right)^2 \right), \frac{\delta}{4p'} \left(\frac{1 + \delta}{1 + \delta(1 - p')} \right) - \frac{\delta}{8p'} \left(3 - \left(\frac{\delta p'}{1 + \delta(1 - p')} \right)^2 \right) \right). \quad \square$$

While Proposition 5 establishes the conditions under which an interior solution can exist, we do not know if it is unique for a given ability level. The next result demonstrates that any interior solution $p(k)$ is, in fact, unique for some k within the appropriate bounds.

Proposition 6. Any interior solution $p(k)$ is unique for a given k .

Proof. Let $k > 0$ such that an interior solution exists. Suppose $\hat{p}(k)$ and $\bar{p}(k)$ are both interior solutions to the agent's problem for a given k . Then $\hat{p}(k), \bar{p}(k) \in (p^*, 1)$,

$$\frac{d}{dp}U(d^*(p)) - c'(p) = \frac{\delta}{2} \left(\frac{1 + \delta}{1 + \delta(1 - \hat{p}(k))} \right) - \frac{\delta}{4} \left(3 - \left(\frac{\delta \hat{p}(k)}{1 + \delta(1 - \hat{p}(k))} \right)^2 \right) - 2k\hat{p}(k),$$

and

$$\frac{d}{dp}U(d^*(p)) - c'(p) = \frac{\delta}{2} \left(\frac{1 + \delta}{1 + \delta(1 - \bar{p}(k))} \right) - \frac{\delta}{4} \left(3 - \left(\frac{\delta \bar{p}(k)}{1 + \delta(1 - \bar{p}(k))} \right)^2 \right) - 2k\bar{p}(k).$$

Setting both derivatives equal to zero, this further implies that

$$\begin{aligned} \frac{\delta}{2} \left(\frac{1+\delta}{1+\delta(1-\hat{p}(k))} \right) - \frac{\delta}{4} \left(3 - \left(\frac{\delta\hat{p}(k)}{1+\delta(1-\hat{p}(k))} \right)^2 \right) - 2k\hat{p}(k) = \\ \frac{\delta}{2} \left(\frac{1+\delta}{1+\delta(1-\bar{p}(k))} \right) - \frac{\delta}{4} \left(3 - \left(\frac{\delta\bar{p}(k)}{1+\delta(1-\bar{p}(k))} \right)^2 \right) - 2k\bar{p}(k) \end{aligned}$$

$$\begin{aligned} \frac{\delta}{2} \left(\frac{1+\delta}{1+\delta(1-\hat{p}(k))} \right) + \frac{\delta}{4} \left(\frac{\delta\hat{p}(k)}{1+\delta(1-\hat{p}(k))} \right)^2 - 2k\hat{p}(k) = \\ \frac{\delta}{2} \left(\frac{1+\delta}{1+\delta(1-\bar{p}(k))} \right) - \frac{\delta}{4} \left(\frac{\delta\bar{p}(k)}{1+\delta(1-\bar{p}(k))} \right)^2 - 2k\bar{p}(k). \end{aligned}$$

Rearranging and simplifying, we notice that

$$\begin{aligned} \frac{\delta(1+\delta)}{2} \left(\frac{\bar{p}(k) - \hat{p}(k)}{(1+\delta(1-\hat{p}(k)))(1+\delta(1-\bar{p}(k)))} \right) + \\ \frac{\delta^3}{2} \left(\frac{(\bar{p}(k))^2 - (\hat{p}(k))^2}{((1+\delta(1-\hat{p}(k)))(1+\delta(1-\bar{p}(k))))^2} \right) - 2k(\hat{p}(k) - \bar{p}(k)) = 0. \end{aligned}$$

Since $\hat{p}(k), \bar{p}(k) \in (p^*, 1)$, this condition only holds if $\hat{p}(k) = \bar{p}(k)$. Hence, $p(k)$ is unique for a given k . \square

Once an optimal memory process is chosen, we can analyze the agent's actions in terms of her choice of d^* . Note that when the agent chooses either $p = 0$ or $p = 1$, her choice of d^* is trivial; she will either maximize the direct component (assume no risk) or the information component (assume full risk), respectively. The agent's choice of optimal direct component when $p(k)$ is an interior solution is more subtle; the next result demonstrates that the agent's optimal direct component depends on her ability.

Proposition 7. Let $p(k)$ be an interior solution. The optimal direct component $d^*(p(k))$ is increasing in ability k .

Proof. Suppose $p(k)$ is an interior solution. Then $p(k)$ satisfies the first-order condition for an optimal solution, which is given by

$$2kp(k) = \frac{\delta}{2} \left(\frac{1+\delta}{1+\delta(1-p(k))} \right) - \frac{\delta}{4} \left(3 - \left(\frac{\delta p(k)}{1+\delta(1-p(k))} \right)^2 \right).$$

Rearranging this equation, we see that

$$\begin{aligned} p(k) &= \frac{\delta}{2k} \left(\frac{1}{2} \left(\frac{1+\delta}{1+\delta(1-p(k))} \right) - \frac{1}{4} \left(3 - \left(\frac{\delta p(k)}{1+\delta(1-p(k))} \right)^2 \right) \right) \\ &= \frac{\delta}{k} \left(\frac{1}{4} \left(\frac{1+\delta}{1+\delta(1-p(k))} \right) - \frac{1}{8} \left(3 - \left(\frac{\delta p(k)}{1+\delta(1-p(k))} \right)^2 \right) \right). \end{aligned}$$

Evaluated at $p(k)$, the optimal direct component becomes

$$\begin{aligned}
d^*(p(k)) &= \frac{1}{4} \left(3 - \left(\frac{\delta p(k)}{1 - \delta(1 - p(k))} \right)^2 \right) \\
&= \frac{3}{4} - \frac{\delta \left(\frac{\delta}{k} \left(\frac{1}{4} \left(\frac{1+\delta}{1+\delta(1-p(k))} \right) - \frac{1}{8} \left(3 - \left(\frac{\delta p(k)}{1+\delta(1-p(k))} \right)^2 \right) \right) \right)}{4 \left(1 + \delta \left(\frac{\delta}{4k} \left(\frac{1+\delta}{1+\delta(1-p(k))} \right) - \frac{\delta}{8k} \left(3 - \left(\frac{\delta p(k)}{1+\delta(1-p(k))} \right)^2 \right) \right) \right)} \\
&= \frac{3}{4} - \frac{\delta^2 \left(\frac{1}{4} \left(\frac{1+\delta}{1+\delta(1-p(k))} \right) - \frac{1}{8} \left(3 - \left(\frac{\delta p(k)}{1+\delta(1-p(k))} \right)^2 \right) \right)}{4k \left(1 + \delta - \frac{\delta^2}{k} \left(\frac{1}{4} \left(\frac{1+\delta}{1+\delta(1-p(k))} \right) - \frac{1}{8} \left(3 - \left(\frac{\delta p(k)}{1+\delta(1-p(k))} \right)^2 \right) \right) \right)} \\
&= \frac{3}{4} - \frac{\delta^2 \left(\frac{1}{4} \left(\frac{1+\delta}{1+\delta(1-p(k))} \right) - \frac{1}{8} \left(3 - \left(\frac{\delta p(k)}{1+\delta(1-p(k))} \right)^2 \right) \right)}{4k + 4\delta k - \delta^2 \left(\frac{1}{4} \left(\frac{1+\delta}{1+\delta(1-p(k))} \right) - \frac{1}{8} \left(3 - \left(\frac{\delta p(k)}{1+\delta(1-p(k))} \right)^2 \right) \right)} \\
&= \frac{3}{4} - \frac{\delta^2 \left(\frac{1}{4} \left(\frac{1+\delta}{1+\delta(1-p(k))} \right) - \frac{1}{8} \left(3 - \left(\frac{\delta p(k)}{1+\delta(1-p(k))} \right)^2 \right) \right)}{4k(1 + \delta) - \delta^2 \left(\frac{1}{4} \left(\frac{1+\delta}{1+\delta(1-p(k))} \right) - \frac{1}{8} \left(3 - \left(\frac{\delta p(k)}{1+\delta(1-p(k))} \right)^2 \right) \right)}.
\end{aligned}$$

As we can see from the equation above, $d^*(p(k))$ is influenced by k both directly and through $p(k)$. Clearly, $d^*(p(k))$ is increasing directly in k . To help us ascertain the impact of k on $d^*(p(k))$ by way of $p(k)$, we analyze how k impacts $p(k)$.

By Proposition 2, $p = 1$ when $k = 0$; similarly, $p = 0$ for sufficiently high k by Proposition 4. In the case of an interior solution, we return to the first-order condition for an optimal solution:

$$kp(k) = \frac{\delta}{4} \left(\frac{1 + \delta}{1 + \delta(1 - p(k))} \right) - \frac{\delta}{8} \left(3 - \left(\frac{\delta p(k)}{1 + \delta(1 - p(k))} \right)^2 \right).$$

Looking at this equation, it is apparent that if k increases, $p(k)$ must decrease accordingly so as to maintain equality. Since larger values of k are therefore associated with lower values of $p(k)$, and d^* is decreasing in p (as discussed in Section 3), we can infer that $d^*(p(k))$ is increasing in k . \square

Proposition 7 demonstrates that the agent's choice of direct component is increasing in k . As we know that the optimal direct component is inversely related to the information component of the agent's expected utility function, we can infer that the information component must be decreasing in the parameter k . In other words, if the agent is of higher ability, then she expects a larger payoff in period 2, and she will adjust her choice of direct component downward (thereby taking riskier actions) to reflect her expectations. The

final piece of this analysis deals with the impact of the agent's ability on her overall welfare.

Proposition 8. Welfare is decreasing in k .

Proof. Suppose $p(k)$ is an interior solution. By Proposition 7, $d^*(p(k))$ is decreasing in k . As discussed in Section 3, $U(d)$ is decreasing in $d^*(p)$. Thus, as k increases, $d^*(p(k))$ increases and $U(d^*(p(k)))$ decreases. Moreover, the memory cost function evaluated at $p(k)$, $c(p(k)) = kp(k)^2$, is clearly increasing in k .

Welfare is given by $U(d^*(p(k))) - kp(k)^2$. Since $U(d^*(p(k)))$ is decreasing in k , and $kp(k)^2$ is increasing in k , we conclude that welfare is decreasing in k . \square

Recall that higher values of k are associated with lower levels of cognitive ability. Accordingly, the agent's welfare is ultimately increasing in her ability. From an intuitive standpoint, this finding makes sense: a high-ability agent can improve her memory more easily than a low-ability agent, as higher ability is associated with smaller cognitive costs.

Chapter 5

Conclusion

This model I have proposed in this thesis was motivated by several key features of memory and learning: (1) learning and memory are, to some degree, skills that can be manipulated at the individual level (Wingfield and Byrnes 2013); (2) humans have some awareness of their own memory capacities (Flavell and Wellman 1975); (3) memory is highly imperfect; and (4) memory impacts decisions and outcomes (Wingfield and Byrnes 2013). Given that memory and learning are malleable processes, and since we have some general awareness of our memory capacities, it is clear that incentives to improve our memory might arise when outcomes depend on these processes.

I have constructed a simple model of rationally inattentive memory recall in a general dynamic setting to study a subset of “intentional remembering” problems (Wingfield and Byrnes 2013) that are characterized by three properties: (1) an agent faces the same decision problem in two periods; (2) an agent’s choice of action produces tangible information about the true state via the observed outcome; and (3) the agent has imperfect memory recall. By framing memory as a tool for manipulating information, I have considered how both exogenous and endogenous memory processes can impact an agent’s optimal choice of actions when she faces her decision problem in different periods. When memory constraints are exogenous, an agent’s choice of action is influenced by her memory, which is outside of her control. Rather than choosing her memory process to select the best course of action, the agent adapts to her cognitive constraints in this setting.

However, when memory is endogenous, a forward-looking agent who faces a rationally inattentive memory recall problem will choose a memory process according to the costs associated with doing so. When memory is costless, such an agent will choose perfect recall, as she stands to benefit the most from information she receives from her actions. On a similar note, a memory-constrained agent will choose not to exert cognitive effort when the costs of doing so are sufficiently high. Finally, the agent will trade off between the potential

benefits of risky information and the costs of acquiring it if her cognitive ability falls within a range of positive values. These results both agree with and add a new dimension to the established literature on imperfect memory recall.

Bibliography

- Anscombe, F. J. and Aumann (1963). A definition of subjective probability. *Annals of mathematical statistics* 34(1), 199–205.
- Baddeley, A. D. (2004). The psychology of memory. *The essential handbook of memory disorders for clinicians*, 1–13.
- Brenner, G. A. and R. Brenner (1982). Memory and markets, or why are you paying \$2.99 for a widget? *Journal of Business*, 147–158.
- Eil, D. and J. M. Rao (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics* 3(2), 114–38.
- Fisher, R. P., N. Brewer, and G. Mitchell (2009). The relation between consistency and accuracy of eyewitness testimony: Legal versus cognitive explanations. *R. Bull, T. Valentine, & T. Williamson (Eds.), Handbook of psychology of investigative interviewing: Current developments and future directions*, 121–136.
- Flavell, J. H. and H. M. Wellman (1975). Metamemory.
- Gittins, J. and D. Jones (1974). A dynamic allocation index for the sequential allocation of experiments, in (j. gani, et al, eds.) progress in statistics.
- Gittins, J. C. and D. M. Jones (1979). A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika* 66(3), 561–565.
- Gossner, O., J. Steiner, and C. Stewart (2021). Attention please! *Econometrica* 89(4), 1717–1751.
- Jakobsen, A. M. (2020). A model of complex contracts. *American Economic Review* 110(5), 1243–73.
- Muda, R., P. Niszczoła, and P. Augustynowicz (2020). The effect of imperfect memory recall on risk preferences. *Journal of Behavioral Decision Making* 33(5), 683–690.
- Mullainathan, S. (2002). A memory-based model of bounded rationality. *The Quarterly Journal of Economics* 117(3), 735–774.

- Piccione, M. and A. Rubinstein (1997). On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior* 20(1), 3–24.
- Rabin, M. and J. L. Schrag (1999). First impressions matter: A model of confirmatory bias. *The quarterly journal of economics* 114(1), 37–82.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics* 50(3), 665–690.
- Wilson, A. (2014). Bounded memory and biases in information processing. *Econometrica* 82(6), 2257–2294.
- Wingfield, A. and D. L. Byrnes (2013). *The psychology of human memory*. Academic Press.