Rational Memory with Decay

Nathaniel Neligh*

July 20, 2022

Abstract

Previous economic models of memory have failed to incorporate one of its critical features: decay. It has long been established that memories fade over time, losing fidelity. In this paper we show that a framework of rational memory with decay can produce the recency effect and other economically interesting phenomena across a wide range of economic contexts. We apply the framework to models of elections, insurance purchasing, and advertising. In these settings, the framework reproduces empirically established behaviors and produces additional insights.

1 Introduction

There is substantial evidence in the experimental literature that people have imperfect memory. One significant imperfection is the tendency of memories to decay over time, becoming noisier and less reliable. The decay of memories is one of the earliest and most robust findings of empirical psychology. First studied by Ebbinghaus (1885), it has been supported consistently in experiments on both animals and humans since that time (Averell and Heathcote, 2009; Averell and Heathcote, 2011; Murre and Dros, 2015). However, most existing economic models of imperfect memory do not directly incorporate decay, instead focusing on other memory features such as limited capacity, context dependence, retrieval, and motivated forgetting.[1]

*University of Tennessee, Knoxville Haslam College of Business Economics Department. Thanks to the audiences at the Chapman University Brown Bag Seminar and the 2019 Bay Area Behavioral and Experimental Economics Workshop for their attention and feedback. Special thanks to Mark Dean, David Rojo-Arjona, Ambuj Dewan, Andrew Hanson, John Quah, Andrew Kosenko, Adam Sanjurjo, Andrea Wilson, and Michael Woodford for their insights and advice.

[1]Limited Capacity: Kocer (2010); Drakopoulos et al. (2013); Wilson (2014); Sanjurjo (2015)
Context Dependence: Bordalo et al. (2020); Wachter and Kahana (2020)
Motivated forgetting: Bénabou and Tirole (2002), Zimmermann (2020)
Retrieval: Afrouzi et al. (2020)

Research has also found that memory consistently responds to incentives, with memories linked to high rewards decaying more slowly (Weiner, 1966; Loftus and Wickens, 1970; Spaniol et al., 2014). Performance on memory tasks increases as the incentives for a correct response increase, a phenomenon which would not arise if memory decay were a purely exogenous process (Wieth and Burns, 2006; Spaniol et al., 2014). This implies that memory retention is at least partially determined by a rational optimizing process.

In this paper we propose an optimizing framework of imperfect memory with decay. This framework provides a tractable and intuitive way to model the link between accuracy, time, and effort in the memory system. Using this framework we can derive temporal comparative statics and estimate the welfare impact of imperfect memory. For example, we can predict how the likelihood of a good election outcome changes as campaign seasons get longer and estimate how much value is lost to imperfect memory in insurance markets.

Importantly, this framework also generates the "recency effect" as a natural result of decay and Bayesian updating. This result is appealing in an Occam's Razor sense, because it combines two important and well studied psychological effects into the same phenomenon. The recency effect is a widely known psychological phenomenon whereby information that has arrived recently has a much larger impact on beliefs and behavior than older information (Ciccarelli and White, 2007).[2] The recency effect has been found to have significant impacts on economic behavior, as discussed by Kahneman and Tversky (2000). For example, the recency effect can influence voter turnout (Panagopoulos, 2011) and insurance purchasing (Kunreuther et al., 2013), and legislators often take advantage of the effect when deciding what legislation to pass (Amacher and Boyes, 1978). The recency effect can also account for persistent non-equilibrium play and the impact of different feedback regimes in games (Fudenberg and Peysakhovich, 2016).

To our knowledge, ours is the first formal memory framework which directly explores the natural connection between memory decay and the recency effect. Previous economic and psychology literatures have treated the recency effect as an exogenous process separate from memory decay or as the result of salience, but these approaches are not fully satisfying. Without some explanation of why recent memories are more salient, salience based explanations of the recency effect are essentially just re-phrasings.

In Section 2 we begin by presenting our costly memory framework inspired by a simple data storage problem with information loss. In this framework, the eventual

[2]The recency effect is also often referred to as recency bias, but we will avoid that term because it is not entirely clear what constitutes a bias when agents are rational and Bayesian.

accuracy of a memory is determined by how much effort was devoted to storing it and how long it has been stored.

For example, consider what happens when a person parks in an unfamiliar parking structure. They could simply glance around and go about their business, hoping they will be lucky enough to find their way back. Alternatively, they could repeat the spot row number and floor to themselves several times, rehearsing and enhancing the memory. Doing this would slow the memory's decay and improve odds of later success. If they were very concerned, they could develop a mnemonic for their row and floor or they could pull out their phone and take a picture. People have many tools, both mental and physical, that help them store information. Different tools have different rates of decay and different costs associated with them.[3]

For readers who believe that players will default to their most secure external memory media in most cases, we note that most people record very little of the information that they are exposed to in a day. Also information being learned is often subjective and cannot be effectively conveyed in an explicit way, leaving only mental tools available. For example, in learning about probabilities one may intellectually "know" the probability of disaster abstractly from exogenous information, but without experience memories one does not know what that probability really means.[4]

In Section 3 we provide the general results which will be applied in later sections. Section 3.1 looks at simple environments with only one exogenous signal and one action and shows how effort and behavior respond to changes in delay, incentives, and prior. Section 3.2 looks at environments with multiple signals and one action. These results quantify the recency effect in such settings and examine how it changes with the rate of decay. There is also a result showing how total recalled information increases with the length of the information gathering period. Section 3.3 is about environments with multiple signals and actions. The result here provides conditions for memory effort to converge to a stable level. There is a table which summarizes these results in Section 3.4.

In Section 4. We present solved examples for the general results in order to build intuition and demonstrate potential portability.

In Section 5 we apply our framework to a voting setting where every voter receives a sequence of informative exogenous signals and then votes between two alternatives. This setting could plausibly represent jury voting or undecided voters in a political election. There is evidence that recency effects play a strong role in jury trials (Costabile and

[3]For a review of the recent literature on memory consolidation see Squire et al. (2015).

[4]Abdellaoui et al. (2011) discuss the difference in behavior between revealed and experienced utility contexts.

Klein, 2008; Kerr and Jung, 2018). Studies have also found that the timing of political messages and campaigns can have a significant impact on voter behavior with events immediately preceding an election having a much stronger effect than those that happen earlier in the campaign season (e.g. Panagopoulos, 2011; Invernizzi, 2020).

For simplicity, we assume that all voters have aligned preferences and majority rule. This model is designed to isolate the effects of memory. The impact of things like unaligned preferences, skewed priors, and non-majority voting will introduce strategic voting, but will not interact with memory except to decrease the value of information. The value of information is discussed in Section A.11.

The voter wants to remember many exogenous signals in order to improve his or her chance of voting correctly in the case that he or she is pivotal. Due to decay, it is more difficult to remember earlier signals long enough to vote based on them. We find that the recency effect increases as signals become more informative, because more informative signals are more substitutable. Rational memory can reduce the impact of imperfect memory on utility relative to purely exogenous decay, but later events are still given a greater weight.

The model also generates a somewhat counter-intuitive result: when memory resources are rationally allocated, the performance of the electorate can worsen as the size of the electorate increases. As the number of voters increases, the probability of pivotality decreases, which in turn lowers player incentives and efforts. Because memory effort is a public good there is a free rider effect which means that new players may not bring in enough information to compensate the loss in existing players' information. This provides a stark contrast with the predictions of non-optimizing approaches to imperfect memory, where the Condorcet Jury Theorem guarantees that more participants will improve performance.

In Section 6 we apply our framework to make predictions in an insurance buying model. There is a phenomenon sometimes called "insurance cycles" where insurance demand spikes immediately after a disaster, only to drop back down when no new disasters are forthcoming.[5] This type of behavior can lead to serious inefficiencies, with substantial under- or over-purchasing of insurance.

In this game individuals live in either a high-risk or a low-risk world. Both worlds produce shocks, with the high-risk world producing larger shocks on average. Insurance which fully protects against shocks is available. The price of this insurance is set in such a way that the player wants to buy insurance in a high-risk world and not in a low-risk

[5] Browne and Hoyt (2000); Dumm et al. (2017); Dumm et al. (2020); Kunreuther et al. (2013); Volkman-Wise (2015)

world, which means that the player wants to learn the state of the world. As such, the player has an incentive to remember past shocks. However, information about past shocks cannot be reliably stored indefinitely, so as in the voting model, players primarily use more recent information.

In our model, the diminishing returns on information guarantee that the true risk level is never learned. This leads to beliefs evolving in a way that resembles a bounded random walk with no asymptotic convergence, which in turn leads to insurance cycles. We find that even when memory costs are very low, if decay is substantial insurance cycles can lead to substantial under- or over-purchasing of insurance. The impermanence of information suggests that informational interventions can only be effective if they provide consistent access to information. Providing one-time access to historical data is likely to improve efficiency only in the short term, and cannot provide experiential information. The results could also be used to support state-sponsored insurance or insurance mandates as long as the insurance policy rules are set by institutions that are not as cognitively constrained as private individuals.

We devote the main body of this paper to primarily decision theoretic environments to better focus the discussion, but we do briefly explore the more game theoretic setting of manipulation in Appendix D.

Note that ours is not a complete model of memory and even more incomplete as a model of information processing. Our model is meant to capture several first order economically important features of memory while remaining simple and portable. We do not claim that this is an exhaustive model of what is occurring in the real world scenarios of the applications, but the effects we describe are likely significant based on existing research and policy. Our model focuses on using two robustly replicable and broadly applicable features of memory (decay and response to incentives) to construct a model which can transparently and simply explain the recency effect and other memory related behaviors.

1.1 Memory Literature

There are many memory effects which may be considered as being of first order importance in some component of economic decision making including limited capacity, context dependence, motivated forgetting, primacy effects, repetition effects, spacing effects, effortful memory, and decay.[6] The requirement that effort is selected before the

[6]Limited capacity: Cowan (2001)
Context dependence: Gruneberg (1994)
Motivated forgetting: Weiner (1968)

exogenous signal is realized is innocuous in our examples, since every exogenous signal realization is essentially equally informative. In less symmetric signal structures, however, this will not always be the case. Players are assumed to remember the structure of the game, so they can deduce their past effort allocations and thus the precision of their memory signal.

There are potential modeling targets in memory encoding, maintenance, and retrieval. The degree to which various effects are separate or caused jointly by underlying features of the system is unclear and in some cases hotly contested, but all of these effects have been consistently observed to significantly impact behavior. This can make it extremely tempting to try to add more memory features to a model, but doing so can result in models that are intractable or only useful in specific highly memory focused settings.

Previous economic models have focused primarily on context dependence, limited capacity, or motivated forgetting.[7] Our model focuses primarily on decay and effortful memory with particular focus on encoding (although it can also provide insight into repetition effects and spacing effects).

Several existing models impose a cost or limit on the amount of information stored in memory (Kocer, 2010; Drakopoulos et al., 2013; Wilson, 2014; Sanjurjo, 2015; Afrouzi et al. (2020)). These models provide useful descriptions of short-term working memory, which is characterized by hard limits on how much information can be available at one time.[8] However, unintentional forgetting does not occur in these models, nor do memories decay over time. In addition, many of these models[9] use hard limits on memory rather than making memory costly, so they do not allow memory to respond to incentives unless an outside use for memory resources is explicitly included.

One economic model proposed by Bénabou and Tirole (2002) has memories vanish as a result of motivated forgetting. In this model, memories are forgotten to maintain positive beliefs about oneself. Some of the newest economic papers on memory, Bordalo et al. (2020) and Wachter and Kahana (2020), focus on context dependence which is a property whereby information is easier to retrieve from memory if it was initially encoded in a similar environment.

Repetition and spacing effects: Greene (2008)
Effortful memory: Hasher and Zagks (1979)
Decay: Ebbinghaus (1885)

[7]Limited capacity: Kocer (2010); Drakopoulos et al. (2013); Wilson (2014); Sanjurjo (2015)
Context Dependence: Bordalo et al. (2020); Wachter and Kahana (2020)
Motivated forgetting: Bénabou and Tirole (2002)

[8]Miller (1994)

[9]Specifically Kocer (2010); Drakopoulos et al. (2013); and Wilson (2014).

Both motivated forgetting and context dependence can contribute to memory decay. This can lead to identification issues. It can be difficult to tell the difference between memory decay, context dependence, or a shifting state of the world rendering older information obsolete. In some real world applications it will not be possible to rule out a shifting state, although in cases like earthquake insurance all evidence points to the state remaining largely consistent over time (barring fracking or other interventions).

Single-system psychological models of memory like the one discussed in Howard and Kahana (1999) often attribute decay and recency entirely to contextual drift and context dependence. Memories are context-dependent, and as more time passes the recall context generally becomes further removed from the context of the initial exposure. In the lab it is generally possible to rule out state changes, but context dependence cannot be fully controlled. That said, laboratory experiments do generally control context as much as is practically possible and observe very regular rates of memory decay.[10] There is also direct neurological evidence of decay unrelated to context dependence which supports decay as a separate phenomenon (Jonides et al., 2008).

This separation makes sense, since decay and context dependence are fundamentally different types of mechanisms. Using a computer based metaphor, there are two ways to fail in retrieving a memory: addressing issues and content decay. Addressing issues are when the location of a memory is lost. Content decay is when the information contained in the memory has been destroyed or damaged. Context dependence is an addressing issue while decay is a content problem. Note when addressing issues are present, the memory is effectively nonexistent as in Mullainathan (2002), discussed below. Content damage has a more continuous impact, reducing the reliability and emotional intensity of the memory. Evidence suggests that content decay and addressing decays are different effects with different underlying mechanisms. Cooper et al. (2019) find that content decay is mostly a mechanical effect of time, while addressing decay can be more affected by other factors like the emotional salience of an event.

Competition for attention can also partially explain the recency effect in settings where participants quickly memorize lists of words. In the psychology literature, the recency effect is often discussed along with a primacy effect where items that arrive earlier are recalled more accurately than those in the middle (Brodie and Murdock Jr, 1977; Glenberg et al., 1980). This effect cannot be explained by decay; it is generally attributed to the brain's limited ability to encode multiple memories within a specific time frame (or to "salience"). Items near the beginning or end of a list do not have to compete for attention as fiercely as those near the middle. However, the primacy effect

[10] Averell and Heathcote (2009); Averell and Heathcote (2011); Murre and Dros (2015)

is often found to be smaller than the recency effect even in experimental settings where information arrives rapidly, so decay likely still plays a role in such settings (Garbinsky et al., 2014). In the systems we model, information is arriving infrequently enough that encoding interference should not be important, so the primacy effect is unlikely to be substantial. There is evidence that the primacy effect may not be present in some investment-type decisions while the recency effect is highly significant (Pinsker, 2011).

While the connection between memory decay and the recency effect is fairly direct and may at first seem obvious and even mechanical, thus far the connection has not been highlighted in the literature. Two-system models like those described in Glenberg et al. (1980) generally attribute the recency effect to the fact that it is easier to retrieve memories from the short-term memory system than from the long-term memory system. However, this type of model is not able to account for recency effects on long-term memories, because all memories are presumably being retrieved from the long-term system in such cases (Bjork and Whitten, 1974).

Economic models with a recency effect generally do not seek to explain it and instead treat it as an exogenous weighting of evidence (Fudenberg and Levine, 2014; Fudenberg and Peysakhovich, 2016).

Mullainathan's (2002) model includes both a recency effect and decay, but the two features are not directly linked. Rather, newer information is given higher weight due to a dynamically changing state of the world. Older events have a chance of being forgotten every period, but any information that still exists is equally reliable. His model also assumes that players are naive regarding the fact that their memories decay, so agents are not fully Bayesian. Afrouzi et al. (2020) similarly present a model where recency is generated by a dynamic state of the world. Their model focuses on costly information retrieval.

Brocas and Carrillo (2016) present a noisy memory model with Bayesian updating which allows for unintentional forgetting, but their model is not dynamic and does not include time, so it cannot be used to explore the recency effect or decay.

The most similar framework to our own, which includes both decay and the recency effect, is being developed contemporaneously by Azeredo da Silveira and Woodford (2019). Their work focuses on the question of which elements of current beliefs will be stored in memory and how precisely they will be stored, when there is no *a priori* constraint on the structure of memories. Our paper instead takes as given an "episodic" structure where memories of different events are stored separately, and focuses on reasons for endogenous variation in the precision with which different events are stored – an issue that is not relevant to the applications that they consider. We compare the models more thoroughly and technically in Section 2.2.

1.2 Identification and Model Similarities

While time plays a critical role in our theory, in single action theoretical environments it only influences the difficulty of retaining information. This means that time could be replaced by other difficulty influencing factors without substantially changing results. Therefore, many of the results we present could also be used in models of single or multi-dimensional rational inattention with non-temporal "difficulty" factors. This is unsurprising as both memory and attention are essentially just noisy information channels. With the single noted exception of Remark 1 our results are all also novel within the rational inattention space to our knowledge, potentially providing a broader use for results.

Note, however, that our model is not isomorphic to a model with rational inattention and exogenous memory decay, because delay and effort interact. Effort influences the marginal impact of delay and vice versa. Effort is specifically impacting decay.

That said, there are some difficulties for someone trying to observationally separate the role of memory and attention. In decision problems with only one action, it is impossible to know whether information was forgotten or never acquired in the first place. In general, it is only possible to definitively identify memory decay directly in decision problems with multiple actions (by noting that information once had has been lost).

Indirect identification of decay is possible by comparing responses to information presented with different delays between exposure and action. If different delays lead to different responses in otherwise identical decision contexts, this implies the presence of decay. This is only an indirect observation of decay, however, because the difference in behavior observed may be partially due to direct information decay, and partially due to a change in attention strategy in response to the presence of decay. If a person knows they will forget something after a while, they won't bother trying to remember it if it isn't immediately useful. Observing a reduction in information after a longer delay in this case implies that decay exists but does not imply that decay is responsible for the full decrease in information.

In the text we will point out contexts where attention and memory may be observationally confounded and where they can be more fully disentangled. When a result could be generated in a Caplin and Dean (2015) style rational inattention framework with no decay, we say that there is an attention analog for that result.[11]

Note that it is ambiguous both observationally and within our model whether effort directly influences decay or merely influences attention gathering in response to decay. Without further assumptions on the decay process, any non-adjustable rational memory model, including our own, will be isomorphic to a rational attention model with exogenous memory decay. We do not concern ourselves overly with this distinction as the predictions are the same.

2 Rational Memory Framework

2.1 The Setting

We begin by explaining the framework conceptually. The player receives a sequence of signals over time informing him about a static state of the world. In simpler settings, this signal may be fully informative, leaving imperfect memory as the only source of uncertainty, but often this will not be the case. In elections, for example, the voter will rarely be presented with definitive evidence that one candidate is the better choice.

When each period begins, the player decides how much effort to devote to remembering the information he will receive during the period. This level of effort determines how well he will remember the current period's events in the future. The longer a memory of a past event is stored, the more it decays. The more effort a player devotes to memory storage during a specific period, the better he can preserve his memory of that period's event.

Next, in every period, the player will consult his memory of past events or signals. He will combine his memories with his prior to determine his posterior beliefs over states. After he has formed posterior beliefs, the player will choose an action that maximizes his expected utility based on those posterior beliefs. Aside from restricted memory, agents are assumed to be fully rational and sophisticated.

We now present the framework more formally. There is a set of time periods $t \in \{1, 2, ..., T\}$ which may or may not be finite. If T is not finite, the game ends every period with a probability $\beta \in (0, 1)$. There is also a static state of the world θ, which comes from a state space $\Theta \subseteq \mathbb{R}$. The player has a prior over states of the world $\pi \in \Delta(\Theta)$. There is a space of actions $A \subseteq \mathbb{R}$ with elements a.

There are two important types of signal spaces in this framework. First, there is a

[11] Sometimes the implicit rational inattention model must be supplemented by introducing multiple information gathering periods for the analog to make sense.

space of exogenous signals $S \subseteq \mathbb{R}$ with elements denoted s. The exogenous signals can be thought of as the events that the player wishes to remember. The probability of receiving a given exogenous signal depends on the state of the world and is exogenous. In order to maintain tractability when multiple memories are involved, we will be using only normally distributed signals and uninformative signals, so $\forall t$ we have

$$s_t \sim N\left(\theta, \sigma_t^2\right)$$

or

$$s_t = \emptyset$$

Exogenous signals are all drawn independently. In most of the applications we consider, there will either be only one informative exogenous signal, or informative signals will be iid, but in a few cases there will be changes in signal quality, and in one noted case the player's previous actions will influence whether they receive an informative signal.

2.1.1 Memory Signals

Second, there is a space of memory signals $M \subseteq \mathbb{R}$ with elements denoted m. We define m_t^τ as the memory of the event that took place at time τ, recalled at time t. Memory signals inform the player about past exogenous signals. The memory signals are retrieved when the player recalls a past event. The probability of receiving a given memory signal depends on the corresponding exogenous signal, the effort devoted to preserving the memory, and how long it has been since the memory was laid down.

We assume that the memory signal m_t^τ is equal to s_τ plus some normal noise.

$$m_t^\tau \sim N\left(s_\tau, \tfrac{g(t-\tau,\delta)}{n_\tau}\right) \tag{2.1}$$

$g(t-\tau,\delta)$ is a decay function. δ is an exogenous persistence parameter and n_τ is the amount of effort devoted to the memory of s_τ. We refer to $t-\tau$ as the delay, because it is the number of periods between when an event occurs and when it is recalled.

Assume $g(t-\tau,\delta)$ is increasing in $t-\tau$ and decreasing in δ. Assume also that $g(t-\tau,\delta)$ goes to infinity as $t-\tau$ goes to infinity and $g(0,\delta)=0$, so variance goes to zero as the delay goes to zero. Memory signal variance is decreasing in n_τ for all $t-\tau$, so more effort always yields a better signal. Because δ is taken as exogenous and fixed in most of our applications, we will suppress it and refer to $g(t-\tau,\delta)$ as $g_{t-\tau}$ going forward.

Where needed in the applications, we use the following functional form

$$g(t-\tau,\delta)=\frac{1}{\delta^{t-\tau}}$$

We also use this form in some propositions where specified.

Note that the framework does not make assumptions about the inter-temporal correlations between memory errors (whether $m_{t-1}^{\tau}-s_{\tau}$ is correlated with $m_{t}^{\tau}-s_{\tau}$). Due to the linear nature of expected discounted utility, this assumption does not matter for the player's attention selection problem, but we do need to make a choice for simulations. In simulating the data, we introduced correlation rather than drawing these errors independently. We decayed memories by adding normal errors of the appropriate variance to the previous period's memory signal, such that

$$m_t^{\tau}=m_{t-1}^{\tau}+\epsilon_t^{\tau} \tag{2.2}$$

where

$$\epsilon_t^{\tau}\sim N\left(0,\frac{g_{t-\tau}-g_{t-1-\tau}}{n_{\tau}}\right) \tag{2.3}$$

This means that errors are correlated across time. Allowing the errors to be correlated over time smooths beliefs, which makes it easier to see the evolution of beliefs.[12] This approach also reflects the natural evolution of memory better than uncorrelated errors, because memory errors are generally persistent (Mullet and Marsh, 2015).

2.1.2 The Game

The player has a sequence of utility functions, $u_t(a_t,\theta,\bullet)$, which depend on their action and the state. These utilities can depend on other factors, such as the actions of others (in the voting application) or the exogenous signal realization (in the insurance application), but a_t and θ will be the most important determinants. Throughout the game, players will select levels of costly effort n_t devoted to preserving specific memories. We refer to the vector of all n_ts as $\boldsymbol{n}$, which has $T-1$ dimensions in finite games and a variable number of dimensions in infinite horizon games. At the end of the game, the player receives a payoff

$$\sum_{t\in\mathcal{T}}u_t(a_t,\theta,\bullet)-c(\boldsymbol{n})$$

[12] Example Figure 6.3

where $c(\boldsymbol{n})$ is a cost that depends on how much effort was devoted to memories throughout the game. Throughout the paper we will assume that $c(\boldsymbol{n})$ is increasing in each element. $\mathcal{T}$ is the set of time periods where actions are payoff relevant. In our finite time horizon applications, the player will take action in the final period, so $\mathcal{T} = \{T\}$. In our infinite horizon game in Section 6, the player takes action every period, so $\mathcal{T}$ includes all periods before the end game's stochastic end.

Every period t has two sub-periods. Information does not persist between the sub periods.

Sub-period $t.a$ Recall and Action:

1. The player receives memory signals about every past exogenous signal. Call the set of received memory signals $\boldsymbol{m}_t = \{m_t^1, m_t^2, ..., m_t^{t-1}\}$. The probability of receiving a particular signal m_t^τ at time t regarding each previous exogenous signal s_τ depends on n_τ and on the original s_τ as previously discussed.

2. The player updates his prior based on signals received this period, using Bayes' Theorem to generate a posterior about the state of the world, $\gamma(\boldsymbol{m}_t, \boldsymbol{n}_t, \pi) \in \Delta(\Theta)$. Here $\boldsymbol{n}_t = \{n_1, n_2, ..., n_{t-1}\}$ is the history of memory efforts. Note past n_τs are deduced in equilibrium from the rules of the game rather than freely remembered, so they cannot be used to send signals.

3. The player chooses an action, a_t.

Sub-period $t.b$ Memory Encoding:

1. The player receives an exogenous signal about the state of the world, s_t, and stores it. Note that s_ts are not necessarily produced by the same signal generating process every period.

2. The player selects a costly effort level, n_t, which determines how well this period's exogenous signal, s_t, will be remembered in future periods by influencing the quality of future memory signals.

The order of the sub-periods is generally unimportant, although reversing the order would not make sense with the insurance application, because it does not make sense to buy insurance for a disaster after it has occurred.

Note that players do not remember past actions in our model and only recall past exogenous signals through memory signals. Signal precision is deduced in equilibrium. Only the game rules are freely remembered at all times. This prevents players from potentially abusing actions to send signals to their future selves. A more complex model might allow for actions to also be remembered with decay, but this would not add significant insight.

Each period the player chooses a_t and n_t optimally. The action, a_t, is chosen to maximize the player's expected payoff for the current period given the player's posterior beliefs after receiving his memory signals. Formally,

$$a_t \in \arg\max_a E(u_t(a, \theta) | \gamma(\boldsymbol{m}_t, \boldsymbol{n}_t, \pi)) \tag{2.4}$$

Memory effort, n_t, is chosen to maximize the player's expected payoff stream net of effort costs, conditional on the current history of memory efforts. Formally,

$$n_t \in \arg\max_n E\left(\sum_{t \in \mathcal{T}} u_t(a_t, \theta, \bullet) - c(\boldsymbol{n}) | n, \boldsymbol{n}_{t-1}\right) \tag{2.5}$$

In all games we assume that there is common knowledge of rationality and memory decay mechanisms. A strategy in this game is vector of memory efforts $\boldsymbol{n}$ and a mapping from received memories $\boldsymbol{m_t}$ signals to actions a_t for each non-trivial period. For our solution concept we use subgame perfect equilibrium (SPE). In most cases this coincides with basic Nash Equilibrium (NE).

2.2 Discussion of Framework

There are several features of this framework that deserve discussion. The state of the world is kept static in order to isolate decay as a source of recency. In an environment with a dynamic state, it is optimal to weight more recent information more. Arguably, this is not a recency effect since information is being used optimally.

Due to the normality of signals, all realizations of s_t for a given period are equally informative and are therefore given equal effort. However, if information could bleed over between sub-periods the player might be able to deduce how accurate of a realization s_t is through their memory signals, which could lead them to change their effort allocation. While this could lead to interesting effects like confirmation bias, it impacts tractability, and similar mechanisms are already covered by Wilson (2014).

Players are assumed to remember specific past events rather than summary statistics or beliefs, which may seem inefficient. Neuro-scientific evidence suggests that individuals more often store memories of past events separately and in whole, rather than remembering only summaries of past events. d'Acremont et al. (2013) find evidence that there are distinct regions encoding prior beliefs and memories of past events which both activate when forming beliefs. Shadlen and Shohamy (2016) also find evidence that individuals form posterior beliefs and make decisions by sampling their memories rather than by accessing some stored belief distribution.

The episodic nature of memory storage helps explain a number of cognitive effects. For example there is the peak-end effect described by Kahneman and Tversky (2000) where people tend to judge experiences based on the most intense part and most recent part rather than the average. This would not be possible in a memory framework which stored updated belief distributions.[13] Ratcliff (1978) finds that a theory where memories are stored and retrieved episodically explains a broad range of experimental data. Similarly, Murty et al. (2016) find that subjects who failed to form episodic memories associating lottery choices with specific payouts did not update their behavior and learn in an ambiguous lottery selection task.

In addition to being realistic, episodic memory can be rational. Kumaran et al. (2016) posit that human brains use this approach, because storing episodic memories is critical for meeting the flexible needs of human existence. By storing memories as episodes, new information can easily be integrated and old information can easily be re-interpreted in the face of new developments without the risk of correlation neglect. We are not arguing that players do not recall information summaries, merely that episodic memory often plays a critical role in determining beliefs.[14]

Episodic memory also has expositional advantages: by assuming episodic memory, we shut down the optimal coding dimension of the memory problem, which examines the way that different features of the player's beliefs are stored. This component of the memory problem is explored extensively by Kocer (2010), Drakopoulos et al. (2013), Sanjurjo (2015), and Wilson (2014). Assuming that each informative event is stored independently makes it easier to isolate the impacts of effort and decay.

There are a number of different types of memory and memory systems discussed by psychologists and neuroscientists. Squire (2004) discusses the differences between these different types of memory: semantic memories pertain to facts and important events while procedural memories relate to skills and processes.[15] Semantic memories can be further subdivided into declarative memories and episodic memories. Declarative memories involve general facts like the capital of Denmark while episodic memories relate to experienced events like what you had for breakfast. Episodic memories in particular are helpful for the formation of beliefs about risk, utility, and complex states of the world

[13]This is a somewhat technical point. One could have a mechanism where individuals updated strongly after outliers and distributions decayed, leading to something like a peak-end effect, but outliers would have to be judged relative to the current distribution, so it would overweight local peaks rather than the global peak.

[14]Sanjurjo (2014) discusses situations where memories are likely to be combined and summarized.

[15]See Schacter et al. (2009) for further discussion of different memory types and systems.

because they provide a sample of events. Our framework is specifically designed to represent episodic memories, although it can also reasonably cover declarative memories. Procedural memories do not map well onto the framework.

Another important feature of the model is that the effort allocated to a memory is determined entirely when the memory is first laid down without later opportunity for adjustment. This is in contrast with the Azeredo da Silveira and Woodford (2019) framework which is fully adjustable, meaning that the amount of effort devoted to remembering each component of memory is re-optimized every period.

In Appendix C we show how frameworks reflecting either of these assumptions produce counterintuitive predictions in different scenarios. Non-adjustable models predict that memory retention will not respond to increases in the expected value of previously acquired information. Fully adjustable models predict that even well memorized information will be lost unrealistically quickly when the expected value of information drops, including the complete instant forgetting of any information that has lost its expected future value.

Intuition and evidence suggest a somewhat richer model where, after memorization effort has been expended, decay can be further slowed but not generally worsened. Most memorized information does not require constant effort or attention to maintain, but regularly rehearsing information will make memories last much longer.

However, this more complex model is not required in the settings we examine, because the expected value of information does not change, leaving us to choose primarily based on expositional considerations. In fact, as we discuss in Appendix C, the adjustable memory effort selection problem and the non-adjustable memory effort selection problem are essentially identical in many of the scenarios we discuss.[16] We use a non-adjustable framework for greater tractability and notational simplicity.

The normality of signals is assumed solely for tractability in the case of multiple signals. Even simple environments like Bernouli signals and binary actions become intractable with only a few signals. Generally normality is not as important in single signal cases, but the assumption is maintained there for consistency.

Note that the prior does not decay like memories do in this model. This can be thought of as representing a default belief state which a person tends to return to. For example, an optimistic person might default to a position that the person they are talking to has good intentions. However, in most cases discussed the prior is chosen as a

[16] Note that this is not always true more generally. In particular the adjustable model allows players to change their effort in response to different signal realizations when these realizations have different informational value as is often the case in less symmetric setups. Whether this responsiveness is more realistic is unclear and likely situation dependent.

position of no information in as far as such a thing is possible in a Bayesian framework.

Throughout the paper, we often assume that the cost of memory effort is convex. We justify the use of cost convexity on the idea that a person has a wide range of possible uses for their effort. When they reallocate their effort to memory the player should generally first reallocate the effort being spent on lower return operations. In this way the marginal cost of effort should be increasing as effort increases. It may be that the function is not truly perfectly convex but is instead a step function which is approximately convex, but this should not dramatically impact results. If costs are not convex, the results may still hold as long as the benefits of information are sufficiently concave.

3 General Results

Before exploring the economic applications of the framework, we first present some general results from the model. These results are organized based on the types of environments where they apply starting from the simplest. Major results are labeled as propositions while minor results are included as remarks or corollaries.

3.1 One Signal One Action

We begin by considering the simplest environment with only one exogenous signal and one choice. Say that one exogenous signal s_t arrives at time t, $\mathcal{T} = \{T\}$, and $u_T(a_T, \theta, \bullet)$ depends only on a_T and θ. Because we are only considering one exogenous signal to begin, only one period's memory effort will be potentially greater than zero. As such we say $c(\boldsymbol{n}) = c(n_t)$. We refer to optimal n_ts as n_t^*. We call this the One Signal One Action (OSOA) environment.

For the first result, we need the following definition.

Definition 1. The strong set order $\geq_{SSO}$, is defined such that if $A \geq_{SSO} B$, then $\forall a \in A, b \in B$, $\max\{a, b\} \in A$ and $\min\{a, b\} \in B$.

This says that all elements of A not in B are greater than all elements of B and all elements of B not in A are less than all elements of A. Given these preliminaries, we have the following result.

Proposition 1. *In the OSOA environment, if $c(n_t)$ is convex, then the optimal $E(u_T(a_T, \theta))$ is decreasing in the delay $T - t$ in a strong set order sense.*

See Appendix A.1 for proof. Here we use a set order because there might be multiple optimal performance levels in some cases. Due to the presence of decay, adding a delay between information and action (stimulus and response in psychology) will decrease performance. Intuitively, the effort selection problem can be rewritten as a precision selection problem with a cost of precision that depends on $T - t$. Under the given conditions, that cost has increasing differences in precision and delay, so monotone comparative statics gives this result.

Note that this proposition does not preclude longer delay leading to greater effort as long as the greater effort is not sufficient to lead to higher performance. We will show a non-monotone relationship between effort and delay in Section 5. Also, this proposition does not have a rational memory analog due to its dynamic nature.

We will refer to the quantity $E(u_T(a_T, \theta))$ as a player's gross expected utility since it does not include cost. Note that a player's gross expected utility is a function of memory effort and delay in this setup. This proposition tells us that under cost convexity, increasing the delay between information exposure and the resulting decision will decrease the player's performance in making that decision. This is intuitive, as longer delay means more opportunity to forget and more effort needed to effectively preserve a memory.

Rationality means that effort will also respond to incentive level.

Remark 1. In the OSOA environment, if $u_T(a_T, \theta) = r \times f(a_T, \theta)$, where $r \geq 0$, and $f(\bullet)$ is an arbitrary function then n_t^* is weakly increasing in r.

See Appendix A.2 for proof. Increasing the incentive level leads to an increase in effort and a resulting increase in $E(f(a_T, \theta))$. This result will be unsurprising to those familiar with the rational inattention literature, because it is a direct analog of NIAC condition from Caplin and Dean (2015). The single exogenous signal, single action framework is isomorphic to a rational attention problem if the delay is held fixed. Unlike Proposition 1, this result does guarantee a monotonicity in effort.

Similarly, when the prior is near degenerate, there is no reason to devote effort to remembering signals.

Remark 2. In the OSOA environment, let $u_T(a_T, \theta)$ be continuous and $E(u_T(a_T, \theta)|n_t)$ be concave in n_t. Let $c'(0) = 0$, $c(n_t)$ be convex and $\pi = (1 - \lambda)\pi' + \lambda\pi_{\theta'}$ where $\pi_{\theta'}$ is a distribution with all its mass on θ'. Then as $\lambda \to 1$, $n_t^* \to 0$.

See Appendix A.3 for proof. This result can also be applied to attention problems. In the one signal, one action framework, only changing the delay allows for the identification of imperfect memory separate from imperfect attention.

3.2 Many Signals One Action

We now consider an environment with multiple exogenous signals but still only one choice. Say that an exogenous signal arrives in each of some subset of time periods, $\mathcal{T}$ =$\{T\}$, and $u_T(a_T, \theta, \bullet)$ depends only on a_T and θ. We call this the Multiple Signals One Action (MSOA) environment.

The first proposition in this section shows how more recent events are remembered with greater precision and are therefore given more weight. Define $h_t^T = \left(\sigma_t^2 + \frac{g_{T-t}}{n_t}\right)^{-1}$ as the precision of the memory signal m_T^t as a signal of θ. Note that h has an inverted scripting convention from m. This is done to keep the notation in line with other sections while minimizing the overall number of potentially confusing superscripts in the paper.

This is how the recency effect manifests in our model.

Proposition 2. *In the MSOA environment, if i.i.d. exogenous signals arrive every period from 1 till $T-1$, and $c(\boldsymbol{n})$ is Schur-Convex*[17] *then the player will pick a sequence of h_t^Ts that is increasing in t.*

For proof, see Appendix A.4. This result does not have an attention analogue.

Note that receiving several normal signals with the same mean is informationally equivalent to receiving one normal signal equal to the precision weighted average of the component signals.

$$M_T = \frac{\sum_{t=1}^{T-1} h_t m_T^t}{\sum_{t=1}^{T-1} h_t} = \frac{\sum_{t=1}^{T-1} h_t s_t}{\sum_{t=1}^{T-1} h_t} + \epsilon$$

DefineM_T as the composite signal. Here ϵ is some mean 0 normal noise. Note all symmetric convex functions are Schur-Convex (Roberts and Varberg, 1973).

Exogenous signals that are more precisely remembered are given greater weight in the composite signal and therefore have a higher impact on decision making. If newer exogenous signals are recalled with higher precision, they will be given more weight, hence recency.

If we put more structure on the problem, we can get a specific form for the sequence of precisions.

Define total precision of memory signals $H_T = \sum_{\tau=1}^{T} h_t^T$.

Proposition 3. *In the MSOA environment, if i.i.d. exogenous signals arrive every period from 1 till $T-1$, and $c(\boldsymbol{n})$ is a strictly convex differentiable function of $\sum n_t$, and $E(u_T(a_T, \theta)|H_T)$ is differentiable with respect to H_T then*

[17] Roberts and Varberg (1973)

$$h_t^T = \max\left(\frac{1}{\sigma^2}\left(1 - z\sqrt{g_{T-t}}\right), 0\right)$$

where z is a fixed value that does not change with t.

For proof and an expression for z see Appendix A.5. This result does not have an attention analogue.

Note that as t increases, g_{T-t} decreases and h_t increases. This means memories of more recent signals have a higher precision, which in turn gives them a greater weight when determining posterior beliefs. This is how the recency effect manifests in this model. Exogenous signals that are sufficiently old receive no memory effort and are therefore immediately forgotten and have 0 precision.

It is interesting to note that in our formulation, the relationship between precision and time is fixed for a given z. While the number of players and $c(\bullet)$ can influence the amount of effort the players devote to memory, they do not affect the distribution of that attention across memories.[18] Many commonly used utility functions will lead to expected utility being differentiable with respect to H_T. For example, the setting in Section 5 and the quadratic loss utility function both satisfy this requirement.

The next proposition shows how giving players more time to gather information will always weakly improve performance.

Proposition 4. *In the MSOA environment, if i.i.d. exogenous signals arrive every period from 1 till $T-1$, , and $c(\boldsymbol{n}) = f\left(\sum g(n_t)\right)$ with $f(\bullet)$ and $g(\bullet)$ increasing and convex functions, then $E(u_T(a_T, \theta))$ is weakly increasing in T.*

For proof see Appendix A.6. This result requires a bit more structure than one might expect, because while it is straightforward that increasing T makes it cheaper to achieve any given precision, it is not trivial to demonstrate that the effect is stronger for higher precisions as the proposition, which is required to guarantee the result. This Proposition does have a rational attention analogue, since more exposures give more information to learn.

As we discuss later in Section 5, more time will always weakly improve performance, but the gains can rapidly drop off after a certain length of time. In the Section 5 example, it is possible to reach a point where additional time periods provide no benefit, because it is inefficient to give any attention to memories of events that are sufficiently far back.

[18]In setups without normal signals, this will not generally be the case.

The final result for this environment says that, if decay is not too significant, the more accurate exogenous signal will receive more memory effort. This is fairly intuitive, since remembering more informative events should be more useful.

Remark 3. In the MSOA environment,if $c(\boldsymbol{n})$ is symmetric and δ is sufficiently close to 1, then if exogenous signals arrive in period t and t' with standard deviations σ_t and $\sigma_{t'}$ such that $\sigma_{t'} > \sigma_t$, the optimal memory efforts satisfy $n_t^* \geq n_{t'}^*$.

For proof see Appendix A.7. This result says that, all else being equal, the player will generally devote more effort to remembering more informative exogenous signals. This result has a rational attention analogue: more effort will be spent in attending to more informative signals.

Note that the inequality is often strict here including when $\delta = 1$, so it will also often be the case that $h_t^* \geq h_{t'}^*$ for δ sufficiently large. Importantly, this shows us that the recency effect is not always dominant. When decay is low and signals are heterogeneous, signal quality may be more important than signal timing.

3.3 Multiple Signals Multiple Actions

Finally we look at the dynamic programming scenario. In this environment there is no set end; instead the game continues with probability β. Exogenous signals are i.i.d. with variance σ^2 and arrive every period. Utility is the same every period. Assume $c(\boldsymbol{n}) = \sum c(n_t)$ in order to deal with $\boldsymbol{n}$ having a potentially variable number of dimensions. This defines the Dynamic Programming (DP) environment.

In the DP environment we can formulate the consumer's memory effort selection problem as a dynamic programming problem:

$$\max_{\{n_1, n_2, \ldots\}} \sum_{t=1}^{\infty} \beta^{t-1} \left(E(u|\boldsymbol{n}_{t-1}) - c(n_t) \right) \tag{3.1}$$

Where $\boldsymbol{n}_t = (\boldsymbol{n}_{t-1}, n_t)$ and $\boldsymbol{n}_0$ is a 0 dimensional vector. If we assume a convenient functional form for g_t we can guarantee convergence to a steady state

Proposition 5. *In the RDP environment, if $E(u(a_t, \theta, \bullet)|H_t)$ is differential and concave in H_t and $g_t = \frac{1}{\delta^t}$, then H_t converges to a steady state defined by the solution, $\bar{H}$, to*

$$u'(H) = \left(\frac{1}{\delta\beta} - 1 \right) c' \left(\left(\frac{1}{\delta} - 1 \right) H \right)$$

If $c(\bullet)$ and is convex, we can interpret this result cleanly. The LHS is decreasing in H and the RHS is increasing in H. Increasing δ or β decreases the RHS for a given H

meaning $\bar{H}$ increases. This makes sense, since increasing the future value of information increases its overall value and therefore the amount gathered. A higher decay parameter means slower decay, increasing the future usefulness and value of information gathered now. A higher discount rate will also increase the future value of information directly.

It is natural to ask whether n_t and H_t converge given other decay functions. The answer is yes, but some simplifying assumptions must be made to avoid infinite dimensional state space. For the sake of space, we do not cover this consideration in detail.

3.4 Summary Tables

Major Results Table

#	Description	Attention Analog
One Signal One Action (OSOA) Environment		
1	Performance in memory tasks decreases with delay	No
Multiple Signals One Action (MSOA) Environment		
2	Players will weight more recent events more heavily	No
3	Gives a functional form to the weighting	No
4	When players have more periods to gather information performance improves	Yes
Multiple Signals Multiple Actions (MSMA) Environment		
5	Convergence to steady state	No

Minor Results Table

#	Description	Attention Analog
One Signal One Action (OSOA) Environment		
1	Effort in memory tasks increases with incentive	Yes
2	If priors are near degenerate, memory effort approaches 0	Yes
Multiple Signals One Action (MSOA) Environment		
3	When decay is low, more effort is spent remembering more informative events	Yes

4 Solved Examples

In this section we assume a convenient functional form for utility and costs and use it to find closed form solutions in the various environments discussed above. In the case of the MSMA framework, a closed form has not been found, but we do characterize the solution.

For convenient forms we assume $u_t(a_t, \theta) = -(a_t - \theta)^2$ and $c(\boldsymbol{n}_t) = \kappa \sum n_t^2$. This utility function has the useful property that $E(u_t(a_t, \theta)|H_t) = -\frac{1}{H_t}$.

We begin with the OSOA setting

Corollary 1. *Assuming the convenient functional forms, in the OSOA the optimal memory effort* $n^* = \left(\frac{g_{T-t}}{2\kappa}\right)^{1/3}$ *and the optimal gross expected utility is given by* $-\left(\left(2g_{T-t}^2\kappa\right)^{1/3} + \sigma^2\right)$.

For proof see Appendix B.1. Players put in more effort when decay is worse and when costs are low.

Next we look at the MSOA setting with an additional assumption on exogenous signal variance to make results more readable.

Corollary 2. *Assuming the convenient functional forms and* $\sigma^2 = 0$, *in the MSOA the optimal memory effort* $n_t^* = \frac{1}{(G^2 2\kappa)^{1/3} g_{T-t}}$ *and the optimal gross expected utility is given by* $-\left(\frac{2\kappa}{G}\right)^{1/3}$*where* $G = \sum_{t=1}^{T-1} \frac{1}{g_t^2}$.

For proof see Appendix B.2. One interesting result is that the expected utility loss from inaccurate guessing only goes to 0 as T goes to infinity if G is unbounded. This means that performance will be asymptotically imperfect for many natural decay functions.

Corollary 3. *Assuming the convenient functional forms,* $\sigma^2 = 0$, $c(n_t) = \kappa n_t^2$, $g_t = \frac{1}{\delta^t}$, *in the MSMA the optimal solution to the resulting dynamic programming problem converges to a unique steady state* $\bar{H} = \left(2\kappa\left(\frac{1}{\delta} - 1\right)\left(\frac{1}{\delta\beta} - 1\right)\right)^{-1/3}$.

For proof see Appendix B.3. To our knowledge this difference equation does not have a known closed form solution, but it is amendable to linear approximation, and we can construct a flow field to graphically analyze the solution. The difference equation and the evolution of precision give us the following relationships

$$\begin{gathered} \Delta H_t = (\delta - 1) H_t + \delta n_t \\ \Delta n_t = \left(\frac{1}{\delta\beta} - 1\right) n_t - \frac{1}{2\kappa(\delta H_t + \delta n_t)^2} \end{gathered} \tag{4.1}$$

Figure 4.1 shows the result. Note that the true solution will approach the stable point along the saddle path from the upper left, since there is no incentive to overshoot the steady state H_t when costs are non-concave. If costs were linear, there would be one

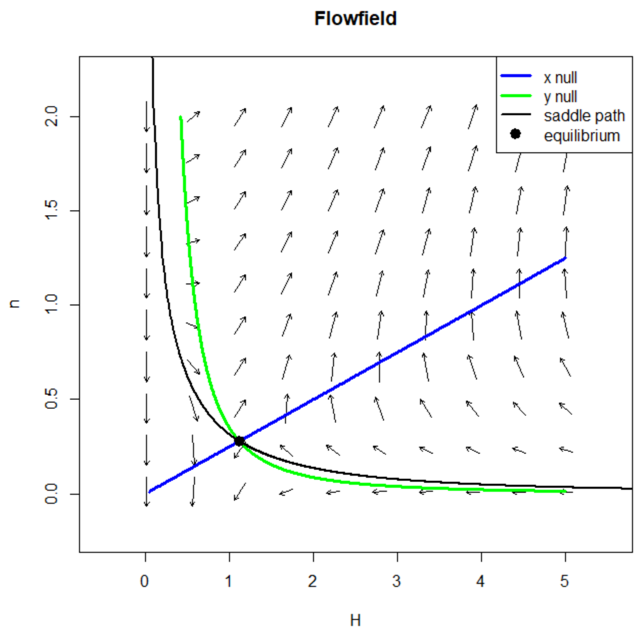

Figure 4.1: Flow field for equations 4.1 with $\delta = \beta = 0.8$ and $\kappa = 1$.

period of extremely high n_t followed by steady state n_t forever. The convexity of costs incentivizes the player to spread out the above-steady-state n_t over multiple periods.

Note that if we use the cost function $c(n_t) = \kappa n_t$ instead, the dynamics become trivial with H_t immediately jumping to its steady state.

5 Voting and The Recency Effect

In court cases, the jurors must remember evidence over the course of a trial in order to reach a verdict. Often, voters seem to display a recency effect whereby more recent events and signals have disproportionate importance in determining voting behavior. Similarly, in most political elections, an individual must collate information gathered over a campaign season in order to make an informed decision about how to vote.

There is substantial evidence from the social psychology literature that the recency effect plays a strong role in jury behavior with evidence presented more recently having a stronger impact on jury voting (Kerr and Jung, 2018). Furthermore, there is evidence that this recency bias is mediated by how well specific pieces of evidence are remembered, lending credence to the forgetting hypothesis (Costabile and Klein, 2008).

For political elections, it is commonly noted that voters respond much more to events like scandals and recessions that appear near an election than they do to events that occur earlier in the election cycle (Healy and Lenz, 2013; Pereira and Waterbury, 2019). While it is somewhat difficult to precisely measure the recency effect in real voting behavior due to the subjective nature of exogenous signal quality in the real world, the recency effect has been found in voting games in the lab (Invernizzi, 2020). There is also a spike in the efficacy of voter mobilization efforts immediately before elections, with earlier voter mobilization being significantly less effective (Panagopoulos, 2011).

While primarily kept simple to highlight the impact of memory, the model can be applied directly to grand juries or demographically similar subsets of swing voters. Grand juries are generally majority rule in the US, and we expect them to have aligned preferences and unbiased priors. In elections, many individuals will eventually be or become confident enough in their opinion of a candidate that they stop gathering new information. This can also happen in juries but is likely less common due to less biased initial states. At this point they have essentially entered an absorbing state and are unlikely to be swayed, but swing voters do not enter such states by definition. Demographically different voters may have poorly aligned interests, but similar voters are more likely to have similar interests.

5.1 Voting Game Model

In this section, we consider what happens when voters must aggregate information over multiple time periods in order to make a decision at a later time. We use our model to make quantifiable predictions about the recency effect, welfare, and voter performance. The model is based on a very simple election framework. There is a set of J *ex ante* identical players indexed $j \in \{1, 2, 3, ..., J\}$ who will vote in an election. There are $T - 1$ evidence gathering periods indexed $t \in \{1, 2, 3, ..., T - 1\}$ and all players must cast their votes during period T, so $\mathcal{T} = \{T\}$. There are only two choices, 1 and -1, and all players have identical preferences, so one option is better than the other for the electorate. When all players have cast their votes, the choice with the majority of votes is implemented. Assume J is odd and abstention is impossible, so no tie breaking is needed. We normalize payoffs such that players receive a payoff of 1 if the good option is selected and 0 otherwise.

The underlying state of the world $\theta \in \{-1, 1\}$ represents which option is better. If $\theta = 1$, then players receive the payoff if option 1 is implemented, and if $\theta = -1$ then they receive a payoff if -1 is implemented. Both states are equally probable *a priori*. In each information gathering period, each player receives a private signal about the state of the world. The exogenous signals are i.i.d. across players as well as time.

Note that each player has their own $\boldsymbol{n}$ and $\boldsymbol{m_T}$, so we could index these by j.

However, we will be examining this problem from the point of view of an individual player. As such we suppress dependence on j to avoid awkward triple indices.

Before continuing we need to introduce a piece of voting terminology.

Definition 2. We say a player votes *sincerely* if the voter always chooses whichever option maximizes their expected payoff given their posterior after updating their beliefs based on their own private signals. In our environment this means

$$a_T \in \arg\max \gamma(\boldsymbol{m}_T, \boldsymbol{n}_T, \pi) \tag{5.1}$$

Critically, the beliefs which define sincere voting are not conditioned on pivotality. This is the standard definition used in the voting literature (Austen-Smith and Banks, 1996).

We also must define one additional piece of game terminology.

Definition 3. A symmetric responsive equilibrium is an equilibrium in which all players adopt the same strategy and player actions depend on signals in a non-trivial way.

This definition eliminates pathological and trivial equilibria like the one where all players always choose $a_T = 1$.

5.2 Perfect Learning Benchmark

The following Lemma is useful when establishing benchmarks for comparison to the rational memory model.

Lemma 1. *In our voting game, if all players receive the same distribution of memory signals, there exists a symmetric responsive equilibrium in which each player will vote sincerely.*

For proof see Appendix A.9.

We begin by considering the perfect memory case, where all memory signals have zero variance, as one benchmark. Because all of the signals are independently and identically normally distributed, in the symmetric responsive equilibrium the players will select option 1 if the average of their signals is above 0 and -1 if it is below. Note that all of the signals in this benchmark have the same weight, because they all have the same variance, σ^2. This means that there is no recency effect, since there is no loss of memory fidelity over time.

The average of the T normal signals a player receives is distributed $N\left(\theta, \frac{\sigma^2}{T}\right)$. The only task that agents must perform in this game is to infer the state based on their signals and then vote accordingly. Therefore, under perfect memory, a player's probability of voting correctly is $1-\Phi\left(\frac{-\sqrt{T}}{\sigma}\right)$. The probability of the majority selecting the correct option is $1-\mathcal{B}\left(\frac{J-1}{2}, J, 1-\Phi\left(\frac{-\sqrt{T}}{\sigma}\right)\right)$, where $\mathcal{B}$ is the binomial cumulative distribution function, and J is the number of players. Note that in this benchmark more players always lead to a higher probability of the good outcome being realized, in accordance with the Condorcet Jury Theorem.

5.3 Solution

Now consider what happens when we add rational memory and decay. Assume that the cost of memory effort $c(\boldsymbol{n})$ is continuous, convex, increasing, and satisfies the Inada-like conditions $\frac{d}{dn_t}c(0) = 0$ and $\lim_{n_t\to\infty} \frac{d}{dn_t}c(\boldsymbol{n}) = \infty$.

Before we discuss the recency effect for voters in this game, it is important to establish that a symmetric responsive equilibrium exists.

Proposition 6. *There exists an equilibrium of the voting game with optimizing memory in which all players select the same $\boldsymbol{n}$'s and vote sincerely.*

For proof see Appendix A.10.

Define $h_t = \left(\sigma^2 + \frac{g_{T-t}}{n_t}\right)^{-1}$ as the precision of memory signal m_T^t, as a signal on the state θ. Note that this means h_t incorporates the exogenous variance σ^2. We can formally represent the recency effect through the following proposition.

Corollary 4. *If $c(\boldsymbol{n})$ is a function of $\sum n_t$, optimal precision in equilibrium takes the form given by Proposition 2.*

For proof see Appendix A.11. Now we use an example to show how this recency can manifest. Figure 5.1 shows how the optimal effort devoted to the memory of the current period changes over the campaign season, as well as the precision of each resulting memory signal. We call the relationship between the time period and the optimal precision the precision curve. The figure demonstrates a pronounced recency effect, with signals that are closer to voting time having much higher precisions and hence a larger impact on results. Any event that occurs prior to around $t = 20$ has no impact.

There is no recency effect in the perfect memory benchmark, but we can also compare our results to another benchmark in which players devote the same effort to remembering each memory. Given fixed memory effort we have

$$h_t = \frac{\bar{n}}{\sigma^2\bar{n}+g_{T-t}}$$

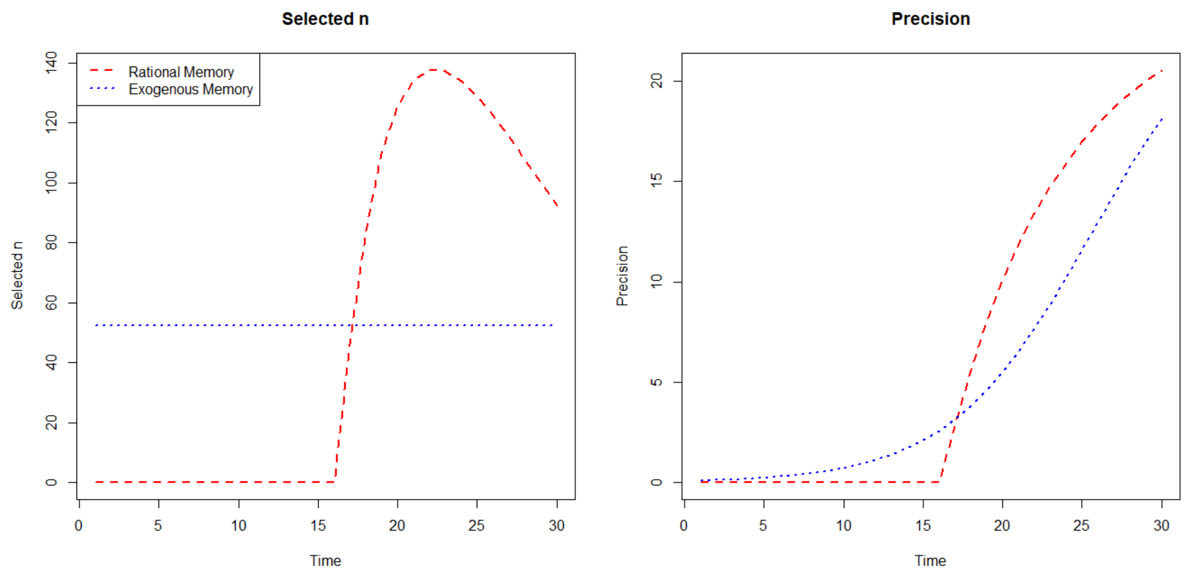

Figure 5.1: Selected n_t (left) and resulting memory signal precision (right) for a rational memory voter and an exogenous memory voter. Parameters $n_{T-1} = 100$, $T = 50$, $\sigma = 0.04$, $\delta = 0.75$. Note that the specific cost function is irrelevant, since any cost function which induces the same n_{T-1} will generate the same graph.

where $\bar{n}$ is some fixed level of effort. For comparison, Figure 5.1 also includes a fixed $n_t = \bar{n}$ exogenous memory benchmark (set at $\bar{n} = 52.4$, the mean from the rational memory solution). Rational memory leads to higher effort and higher precision for the newer memories, while older memories have lower precision. The total precision achieved by the rational memory is higher than that achieved by exogenous memory as expected. Thus, rational memory can mitigate the losses caused by imperfect memory, but it can actually increase the intensity of the recency effect beyond the purely mechanical effect.

From Corollary 4 we know that only two parameters, σ and δ, impact the distribution of effort across memories for a given z. Therefore only these parameters impact the severity of the recency effect. Changing δ can be thought of as essentially changing the units of the time axis. Lower δ compresses the time axis by speeding up decay, while higher δ stretches out the time axis by slowing it. The effect of σ is somewhat more interesting. Higher σ reduces the slope of the precision curve and therefore reduces the intensity of the recency effect. Recall that σ governs the variance of exogenous signals and therefore provides a bound on the maximum achievable precision for each memory signal.

When σ is higher, the bound on precision is lower, guaranteeing that returns for memory effort devoted to a specific exogenous signal diminish more rapidly. Put another way, when exogenous signals are more accurate they are more substitutable, so later signals get more attention. When exogenous signals are less accurate, they become less substitutable and attention is spread more evenly. The fact that more accurate signals are more substitutable, because they contain more similar information is a fairly general property and does not depend strongly on the setup we use.

To develop intuition about the impact of signal precision under rational memory, consider the extreme case of $\sigma \to 0$ where all exogenous signals are perfectly informative and therefore contain the same information; e.g., memories are perfect substitutes and $\frac{dh_t}{dn_t} = \frac{1}{g_{T-t}}$ is constant. In this situation devoting bits to the most recent memory is strictly better, because those bits have a lower chance of decaying before the voting period. There is no reason to devote any cognitive energy to older exogenous signals, so the recency effect is maximized in the rational memory framework. However, in the fixed memory effort benchmark, older memories still have positive precision and positive weight even as $\sigma \to 0$.

The presence of a recency effect in jury and electoral systems may explain the existence of policies which police when information can be provided to voters and jurors. For example, electioneering laws restrict the expression of support for candidates within some radius of polling places throughout the United States. In France campaigning and some types of political reporting are restricted on the day of elections or preceding days. Arguments have been made that defendants should be given the last word in closing arguments, partially due to recency.[19]

For a brief discussion of manipulation using information timing, see Appendix D. This discussion is general, but makes some reference to electoral contexts.

5.4 Deliberation Length and Number of Voters

Campaign seasons have been getting longer and more expensive in the United States (Nichols and McChesney, 2013). It is important, then, to understand what impact this lengthening might have on the welfare of the electorate. There is evidence that campaign seasons educate voters (Arceneaux, 2006), so reasonably one might expect longer campaigns to educate voters better. While causation is difficult to establish, there is evidence that certain jury biases may be more pronounced in shorter trials (Lemley et al., 2013). This change in performance is consistent with our model.

From Proposition 4 we know that lengthening the campaign season or trial will

[19] Mitchell (2000)

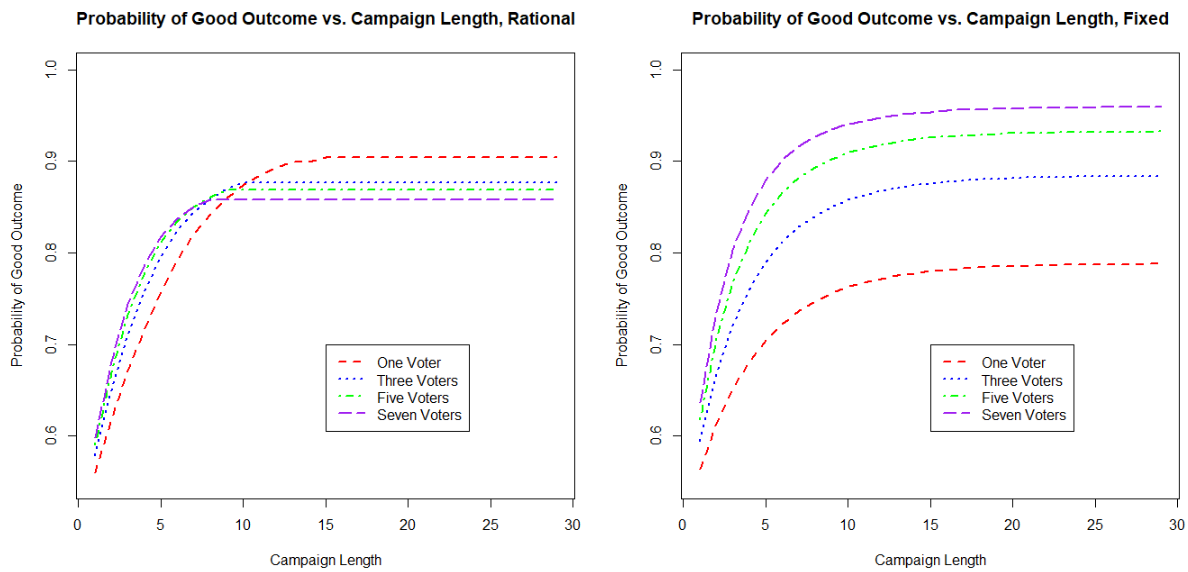

Figure 5.2: Equilibrium probability of good election outcome. $\delta = 0.8$, $\sigma^2 = 5$. Rational framework $c(\boldsymbol{n}) = 0.02 \cdot \sum n_t$ (left) and fixed exogenous effort benchmark $n = 0.1$ (right).

always weakly improve the performance of the voters, but the gains can rapidly drop off once a certain length is reached. In fact, as shown in Figure 5.2 when lengthening the campaign, it is possible to reach a point where additional time periods provide no benefit, because it is inefficient to give any attention to memories of events that are sufficiently far back.

The drop-off in the marginal benefit from longer campaigns is actually more significant when the electorate is larger to the extent that when campaign seasons are sufficiently long, having more voters can actually decrease election performance. This result is interesting, because it contrasts with the Condorcet Jury Theorem which guarantees that adding people will always improve the probability of a good election outcome. Note that the Condorcet Jury Theorem does hold in the fixed memory effort benchmark, where additional players always improve electoral performance.

In our rational memory model more players lowers the probability of pivotality. The lower probability of pivotality substantially lowers voter investment, because voters can free ride on the memory investments of others. This, combined with the disaggregation of information inherent in larger voting groups, leads to the reduction in relative performance for larger voting bodies.

There does not appear to be a parsimonious way to express exactly when additional players will cause a drop in equilibrium performance.

These results suggest that, while long campaigns and trials can be costly, they at least do not reduce system performance. Larger electorates and juries, however, might. The fact that high stakes felony trials have more jurors is at least partially based on the logic that larger juries will produce better results, but we have shown this is not inherently the case. This problem could be even more pronounced for elections. While this concept is a bit unusual, future research could explore a system where instead of all people voting, votes are allocated to a random subset of the population at the beginning of the campaign season. This random subset would have more voting power and therefore more incentive to collect information and vote conscientiously. Conceptually, modern jury selection is already based on similar logic.

6 Insurance Cycles

In this section we apply our framework to study insurance purchasing. People seem to exhibit a strong recency effect when making their insurance purchasing decisions leading to a phenomenon called insurance cycles.[20] When disasters arise, individuals increase their demand for insurance. Over time, however, they seem to forget about past disasters and allow their insurance to lapse. This is problematic because states (local risk) are often very stable over time, which means the recency effect leads to inefficient behavior.

While one would expect a perfect memory individual to respond to disasters and calm periods to some degree through Bayes updating, these responses should asymptotically vanish as information is accrued, but this does not seem to be the case. Dumm et al. (2020) estimate that consumers who have not experienced disasters recently underestimate risk by 30% while those who have experienced disaster temporarily overestimate risk by 50%. This means disaster leads to a doubling of believed disaster risk. While a direct comparison is not made, evidence suggests that recent disasters have an effect on insurance demand that is less than but comparable to the impact of income and insurance price both at the individual and state levels.[21]

While insurance cycles and similar effects have sometimes been attributed to salience or the availability heuristic[22], these explanations do not reveal why more recent memories are more available or more heavily weighted. The temporal nature of the drop-off in demand suggests that it may be more deeply described as a recency effect. In our

[20] Browne and Hoyt (2000); Dumm et al. (2017); Dumm et al. (2020); Kunreuther et al. (2013); Volkman-Wise (2015)

[21] Browne and Hoyt (2000); Dumm et al. (2020)

[22] Dumm et al. (2017); Dumm et al. (2020); McCoy and Walsh, 2018

model, recency could be considered a cause of salience with specific predictions and mechanisms. Time decays both the factual and emotional content of memories. Newer memories would be more informative factually and emotionally and therefore more salient.[23]

6.1 Insurance Purchasing Model

This application uses a repeated game with a continuation probability of β, and the player's utility depends on their actions in every period played. Exogenous signals arrive every period which also serve as shocks to the player's income. The state is again binary, so $\theta \in \{\theta_B, \theta_G\}$ for Bad and Good with $\theta_B > \theta_G$. The two underlying states are *a priori* equally likely. The player has two possible actions every period, $a = 1$ (insure) and $a = 0$ (do not insure).

The player's uninsured per period utility, $u(0, s_t)$, is a static, strictly monotone decreasing concave function of the shock he receives. We refer to the expected uninsured utility corresponding to average shocks θ_B and θ_G as $u_B = E(u(0, s_t)|\theta = \theta_B)$ and $u_G = E(u(0, s_t)|\theta = \theta_G)$. Shocks are normally distributed and i.i.d. with constant variance σ^2. Note that $u_G > u_B$ because the shocks in state θ_G are smaller. For simplicity assume that the insurer offers complete insurance, which guarantees that the insuree receives a payoff $u(1, \bullet) = u_k = \frac{u_B + u_G}{2}$ for the insured period.

Note that due to the concavity of $u(0, \bullet)$, it will always be possible for a risk neutral insurance company to offer such a contract and profit on average at the prior. Since the insurance company is insuring against both uncertainty about the state and uncertainty about the shock size given the state, if σ^2 is high enough an insurance company that believes that the state is B with certainty could still profitably insure an individual to u_k on average.

Because we are dealing with a non-finite game, we assume $c(\boldsymbol{n})$ is a function that can take a countable number of inputs and satisfies $c(\boldsymbol{n}) = \sum c(n_t)$.

We assume that the shock is fully observable even if insurance is purchased and that $g_t = \frac{1}{\delta^t}$. As before, memory signals are distributed normally around the corresponding s_τ with a variance of $\frac{1}{\delta^t n_\tau}$. The player does not have the ability to remember his history of insurance purchases. This assumption does not fit well with rationality, but it is realistic. Most people would not consult their purchase history when considering whether to purchase insurance. Assume that $c(n_t)$ is continuous, convex, and satisfies $\frac{d}{dn_t} c(0) = 0$ and $\lim_{n_t \to \infty} \frac{d}{dn_t} c(n_t) = \infty$.

[23] Cooper et al. (2019)

The player's strategy consists of an infinite sequence of memory efforts n_t and a mapping for each period from memory effort history and received set of memory signals to insurance purchasing decisions. Due to the way in which information arrives, a player will have no reason to change his or her chosen memory effort level in response to past events. As there is no strategic interaction in this game, an equilibrium is simply an optimal strategy for the player.

6.2 Solution

We solve by first considering the insurance purchasing decision. The agent will purchase insurance as long as

$$u_k \geq \gamma_B u_B + (1-\gamma_B)u_G \tag{6.1}$$

where γ_B is the posterior probability that $\theta = \theta_B$. After observing the set of memory signals $\boldsymbol{m}_t$,

$$\gamma_B(\boldsymbol{m}_t) = \frac{e^{-\sum h_\tau^t (m_t^\tau - \theta_B)^2}}{e^{-\sum h_\tau^t (m_t^\tau - \theta_B)^2} + e^{-\sum h_\tau^t (m_t^\tau - \theta_L)^2}} \tag{6.2}$$

Here h_τ^t is the precision of the signal m_t^τ which depends on the effort n_τ and the delay $t-\tau$. The condition for the player to purchase insurance is

$$\sum h_\tau^t m_t^\tau \geq \tfrac{\theta_B+\theta_G}{2} H(\boldsymbol{n}_t) \tag{6.3}$$

where $\boldsymbol{n}_t$ is the vector of n_i's up through and including the current n_t and $H(\boldsymbol{n}_t) = \sum_{\tau=0}^{t} h_\tau^t(n_\tau)$.

This means that the player will buy insurance with a probability of

$$1 - \Phi\left(\left(\tfrac{\theta_B+\theta_G}{2} - \theta\right)\sqrt{H(\boldsymbol{n}_t)}\right)$$

For a given $\boldsymbol{n}_t$, ex ante expected gross utility for a given period is

$$\tfrac{u_B+u_G}{2} + \tfrac{1}{4}(u_G - u_B)\left(2\Phi\left(\left(\tfrac{\theta_H-\theta_G}{2}\right)\sqrt{H(\boldsymbol{n}_t)}\right) - 1\right)$$

We invoke Proposition 5.

Corollary 5. *H_t will converge to a steady state value in the insurance purchasing game to the value defined by*

$$\tfrac{1}{4}(u_G - u_B)\phi\left(\left(\tfrac{\theta_H-\theta_G}{2}\right)\sqrt{H}\right)\tfrac{1}{\sqrt{H}} = \left(\tfrac{1}{\delta\beta} - 1\right)c'\left(\left(\tfrac{1}{\delta} - 1\right)H\right)$$

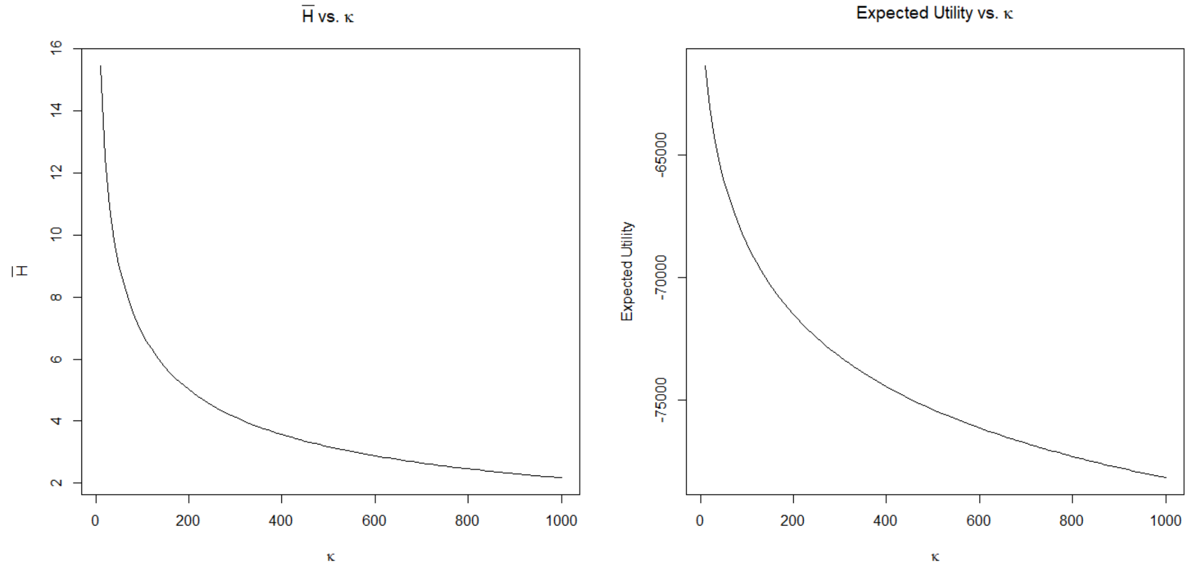

Figure 6.1: Optimal $\bar{H}$ and resulting ex ante expected utility by κ

For proof, see Appendix A.12. To illustrate the phenomenon of insurance cycles, we consider an example with $u(0, s_t) = -e^{2s_t}$ and $c(n_t) = \kappa \times n_t^2$. For parameters we choose $\sigma^2 = 2$, $\beta = 0.75$, $\theta_B = 2$, $\theta_G = 1$, $\delta = 0.5$, $L = 30$.

Figure 6.1 shows how the optimal $\bar{n}$ and resulting ex ante expected utility vary as we change the cost of memory effort, κ.

Figure 6.2 shows how the probability of insurance purchase changes with κ in both states. As we increase κ, the resulting drop in $\bar{H}$ causes insurance buying behavior in the high- and low-risk states to converge. This convergence does become quite slow as it approaches the asymptote, however, due to the quadratic cost function, which makes $n_t < 1$ very cheap. Even small amounts of effort devoted to memory can allow for some distinction between states.

In order to demonstrate the insurance cycles that motivated this analysis, we simulate shock and memory data with $\kappa = 1$. Figure 6.3 shows the evolution of beliefs over time.

Players purchase insurance whenever their beliefs satisfy

$$\gamma_B \geq \tfrac{u_G - u_k}{u_G - u_B} = 0.5 \tag{6.4}$$

which we represent with a dotted line in the figure. Recall $u_k = \frac{u_B + u_G}{2}$.

The simulation generates pronounced insurance cycles with periods of insurance purchase frequently linked to large shocks or "disasters," represented by the black x's

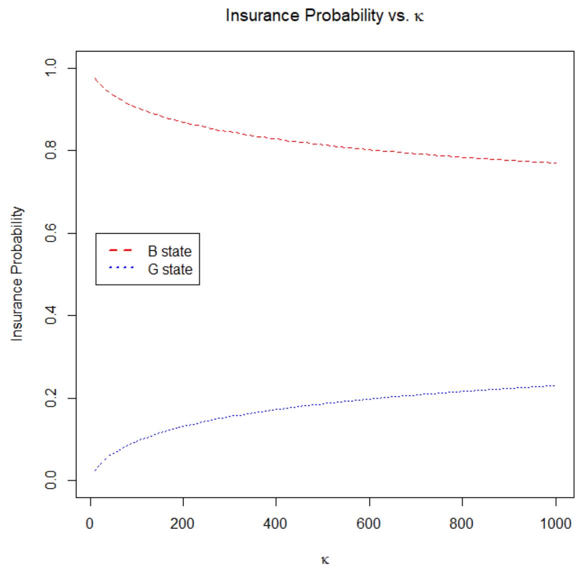

Figure 6.2: Insurance purchase probability by κ in different states

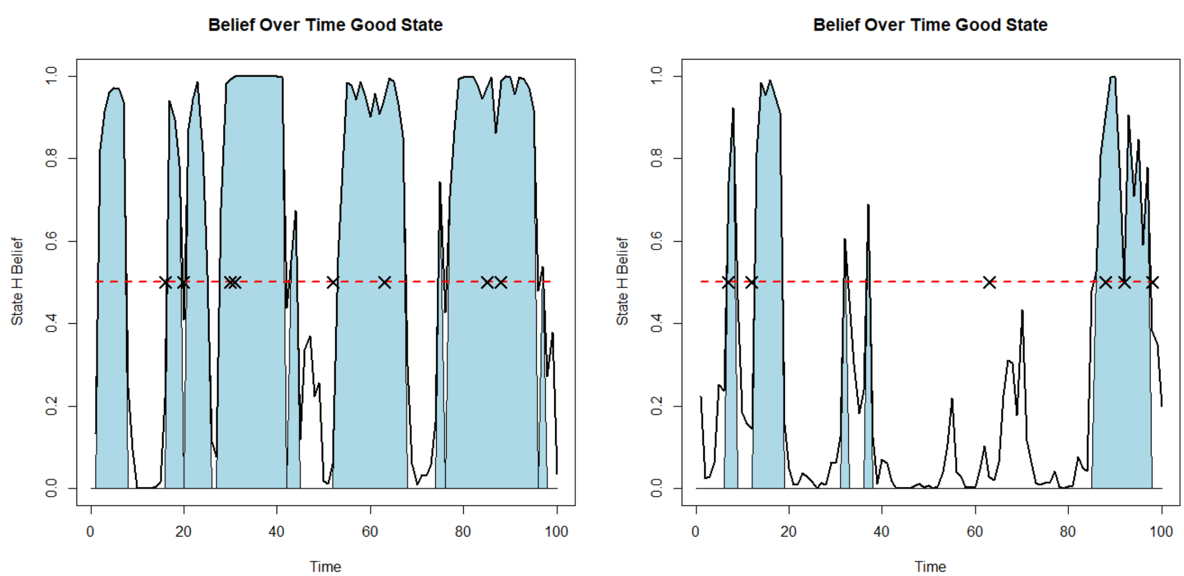

Figure 6.3: Belief that $\theta = \theta_B$ over time when $\theta = \theta_B$(left) and $\theta = \theta_G$ (right). Red dotted line indicates purchase threshold. Blue shading indicates times with insurance purchase. x indicates a shock at least 1.5 standard deviations larger than the mean.

in Figure 6.3. These insurance cycles arise under stable levels of memory investment, meaning that they are an effectively permanent feature of the system.

The insurance cycles represent a large source of welfare loss for the agent, because they will frequently purchase insurance when it is not optimal to do so (the low-risk state) and fail to purchase insurance when doing so is optimal (the high-risk state). To see the impact of these cycles, compare the expected utility values in Figure 6.1 to the perfect memory expected utility of -57208.55.

Interestingly, due to concave benefits and convex costs of effort, overall memory costs can be quite small. In the specification shown in Figure 6.3, for example, the overall cost of memory is only 2.1% of the overall gain from memory (as calculated by subtracting the no-information expected utility from the equilibrium expected utility).

6.3 Discussion

In terms of policy, the predictions of this model provide additional support, beyond the standard adverse selection arguments, for insurance mandates. If properly implemented by an agency with near-perfect memory, such mandates could avoid inefficient insurance under-provision. The model also provides evidence that it may be efficiency-improving to forbid the sale of certain types of insurance. Mandates could also reduce the amount of cognitive effort the average consumer spends on learning about disaster frequency, although the direct welfare impact of this reduction may be small.

Weaker interventions such as continual information campaigns may suffice, but such campaigns would entail ongoing direct and attentional costs. These campaigns may also fail to fully convey risk information as effectively as remembered experience.

In a real market, the insurance companies themselves would set prices optimally and potentially respond to shocks as well, generally by changing u_k. Setting $u_k \neq \frac{u_B+u_G}{2}$ introduces local non concavities in the value of information when H_t is small, and this prevents the use of standard analysis methods. While some recent work has been devoted to dynamic programming problems with local non-concavity, the tools are not yet available to make solving this type of problem a simple task (Pennanen et al., 2017). Based on simulations, the non-concavities can lead to situations where no information is gathered, but otherwise results should be relatively similar.

It is natural to consider how one might embed this decision theoretic model in a more complete market environment. If insurance price dynamics contained information about risk, the model would become extremely complex, but in several realistic scenarios this would not be the case. As mentioned, since the insurance protects buyers against both state based uncertainty and state conditional uncertainty, it is possible that the insurance company would be able to profitably offer complete insurance regardless of buyer and insurer beliefs. In such cases, insurance prices would only reflect the insurer's beliefs about the buyer's beliefs, since the insurance company would only care about extracting as much money as the buyer was willing to pay. If the seller had minimal ability to learn buyer beliefs or price discriminate, there would be little difference from the decision theoretic environment of this section.

In the case where the insurance seller has perfect knowledge about buyer beliefs and perfect price discrimination power, the result is degenerate. The seller would always set the price such that the buyer would be indifferent between buying and not. This means that the utility of buying insurance is equal to the utility of not buying which is linear in belief.

$$E(u(1, s_t)) = E(u(0, s_t)) = \gamma_B u_B + (1 - \gamma_B) u_G$$

The buyer would then have no incentive to gather information, since it would not influence actions. It would also have no impact on expected utility conditional on buying insurance, because insurance buying utility is linear in beliefs and the expectation of beliefs is equal to the prior.

$$E\left(\gamma_B u_B + (1 - \gamma_B) u_G | H_t\right) = \pi_B u_B + (1 - \pi_B) u_G \forall H_t$$

As such the buyer would always buy insurance and never put effort into remembering events.

In many types of disaster insurance, actuarial risk for an area is very stable over time, implying that price dynamics do not contain substantial information about risk.[24] In such cases changes to the price of insurance should primarily reflect the company's beliefs about the willingness to pay of buyers rather than underlying risk. In these cases insurance companies would raise and lower their prices in response to disaster solely to take advantage of demand fluctuations. It is often illegal to raise premiums on existing home insurance customers in response to weather or natural disaster related claims, but there is evidence that insurance companies raise rates on new customers in states where disasters have recently occurred.[25] Anecdotal evidence also suggests that filing car insurance claims related to earthquake damage can also raise premiums.

[24] Clark (2010) indicates that only one major update to USGS earthquake risk estimates has occurred since 1974.

[25] Born and Viscusi (2006)

7 Conclusion

In this paper, we propose a novel framework of rational memory with decay which can produce a recency effect and other empirically validated predictions in several contexts. In voting games, memory decay produces a pronounced recency effect, which can explain the tendency for voter behavior to be excessively influenced by events that occur close to elections. Our framework also can reproduce the phenomenon of insurance cycles, where consumers demand more insurance after a disaster occurs even in conditions where disasters are not serially correlated. We hope this framework allows more researchers to incorporate memory decay into their work.

A number of avenues for future research present themselves. As we discuss in Appendix C, both fully adjustable and non-adjustable frameworks of rational memory present problematic predictions in some scenarios. This suggests that developing frameworks with partially adjustable memory may be valuable, particularly when studying environments where the expected value of information fluctuates, such as the tech industry, where obsolescence can rapidly decrease the value of specialized knowledge.

This paper is also designed as a step towards more complete frameworks of rational cognition in economics. By using a framework compatible with those used in the rational inattention literature, we can potentially provide a platform for further exploration of phenomena involving both attention and memory. This joint framework could be very valuable for understanding the effect of exposure timing on memory generally and the impact of cognitive overload on learning.

References

Abdellaoui, M., L'Haridon, O., and Paraschiv, C. (2011). Experienced vs. described uncertainty: Do we need two prospect theory specifications? *Management Science*, 57(10):1879–1895.

Afrouzi, H., Kwon, S. Y., Landier, A., Ma, Y., and Thesmar, D. (2020). Overreaction and working memory. *NBER Working Paper No. 27947.*

Amacher, R. and Boyes, W. (1978). Cycles in senatorial voting behavior: Implications for the optimal frequency of elections. *Public Choice*, 33(3):5–13.

Arceneaux, K. (2006). Do campaigns help voters learn? a cross-national analysis. *British Journal of Political Science*, 36(1):159–173.

Austen-Smith, D. and Banks, J. S. (1996). Information aggregation, rationality, and the condorcet jury theorem. *The American Political Science Review*, 90(1):34–45.

Averell, L. and Heathcote, A. (2009). Long term implicit and explicit memory for briefly studied words. In Taatgen, A. and van Rijn, H., editors, *Proceedings of the 31st annual conference of the cognitive science society, Cognitive Science Society, Austin, TX*, pages 267–281.

Averell, L. and Heathcote, A. (2011). The form of the forgetting curve and the fate of memories. *Journal of Mathematical Psychology*, 55(1):25–35.

Azeredo da Silveira, R. and Woodford, M. (2019). Noisy memory and over-reaction to news. *NBER Working Paper Series.*

Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics*, 117(3):871–915.

Bjork, R. and Whitten, W. (1974). Recency sensitive retrieval processes in long-term free recall. *Cognitive Psychology*, 6(2):173–189.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2020). Memory, attention, and choice. *The Quarterly Journal of Economics*, 135(3):1399–1442.

Born, P. and Viscusi, W. K. (2006). The catastrophic effects of natural disasters on insurance markets. *Journal of Risk and Uncertainty*, 33:55–72.

Brocas, I. and Carrillo, J. (2016). A neuroeconomic theory of memory retrieval. *Journal of Economic Behavior and Organization*, 130:198–205.

Brodie, D. and Murdock Jr, B. (1977). Effect of presentation time on nominal and functional serial-position curves of free recall. *Journal of Verbal Learning and Verbal Behavior*, 16(2):185–200.

Browne, M. and Hoyt, R. (2000). The demand for flood insurance: Empirical evidence. *Jounral of Risk and Uncertainty*, 20(3):291–306.

Caplin, A. and Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203.

Ciccarelli, S. and White, J. N. (2007). *Psychology (3rd Edition)*. Pearson Education.

Clark, K. (2010). The new u.s. earthquake models: A wake-up call to actuaries? *Society of Actuaries Newsletters*.

Cooper, R., Kensinger, E., and Ritchey, M. (2019). Memories fade: The relationship between memory vividness and remembered visual salience. *Psychological Science*, 30(5):657–668.

Costabile, K. and Klein, S. (2008). Finishing strong: Recency effects in juror judgments. *Basic and Applied Social Psychology*, 27(1):47–58.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(1):97–185.

d'Acremont, M., Schultz, W., and Bossaerts, P. (2013). The human brain encodes event frequencies while forming subjective beliefs. *Journal of Neuroscience*, 33(26):10887–10897.

Drakopoulos, K., Ozdaglar, A., and Tsitsiklis, J. (2013). On learning with finite memory. *IEEE Transactions on Information Theory*, 59(10).

Dumm, R., Eckles, D., Nyce, C., and Volkman-Wise (2017). Demand for windstorm insurance coverage and the representative heuristic. *The Geneva Risk and Insurance Review*, 42:117–139.

Dumm, R., Eckles, D., Nyce, C., and Volkman-Wise, J. (2020). The representative heuristic and catastrophe-related risk behaviors. *Journal of Risk and Uncertainty*, 60:157–185.

Ebbinghaus, H. (1885). Memory: A contribution to experimental psychology. *Classics in the History of Psychology*. An internet resource developed by Christopher D. Green. York University, Toronto, Ontario.

Francis, J., Pagach, D., and Stephan, J. (1992). The stock market response to earnings announcements released during trading versus nontrading periods. *Jounral of Accounting Research*, 30(2):165–184.

Fudenberg, D. and Levine, D. (2014). Learning with recency bias. *Proceedings of the National Academy of Sciences*, (111).

Fudenberg, D. and Peysakhovich, A. (2016). Recency, records and recaps: Learning and non-equilibrium behavior in a simple decision problem. *ACM Transactions on Economics and Computation (TEAC) - Special Issue on EC' 14*, 4(4).

Garbinsky, E., Morewedge, C., and Shiv, B. (2014). Interference of the end: Why recency bias in memory determines when a food is consumed again. *Psychological Science*, 25(7):1466–1474.

Glenberg, A., Bradley, M., Stevenson, J., Kraus, T., Tkachuk, M., Gretz, A., Fish, J., and Turpin, B. (1980). A two-process account of long-term serial position effects. *Journal of Experimental Psychology: Human Learning and Memory*.

Greene, R. (2008). Repetition and spacing effects. In III, H. R., editor, *Learning and memory: A comprehensive reference*, volume 2 of *Cognitive Psychology of Memory*, pages 65–78. Elsevier.

Gruneberg, M. (1994). *Theoretical Aspects of Memory*. Routledge, 2nd edition.

Hasher, L. and Zagks, R. (1979). Automatic and effortful processes in memory. *Journal of Experimental Psychology: General*, 108(3):356–388.

Healy, A. and Lenz, G. (2013). Substituting the end for the whole: Why voters respond primarily to the election-year economy. *American Journal of Political Science*, 58(1):31–47.

Howard, M. and Kahana, M. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 25(4):923–941.

Invernizzi, G. (2020). Public information: Relevance or salience? *Games*, 11(1).

Jonides, J., Lewis, R., Nee, D. E., Lustig, C., and Berman, M. (2008). The mind and brain of short-term memory. *Annual Review of Psychology*, 59:193–224.

Kahneman, D. and Tversky, A. (2000). *Choices, Values and Frames*, chapter Evaluation by moments, past and future, page 693. Cambridge University Press.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Kerr, N. and Jung, J. (2018). Should jurors be allowed to discuss trial evidence before deliberation?: New research evidence. *Law and Human Behavior*, 42(5):413–426.

Kocer, Y. (2010). Endogenous learning with bounded memory. *Princeton University - Economic Theory Center Research Paper.*

Kumaran, D., Hassabis, D., and McClelland, J. (2016). What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in Cognitive Sciences*, 20(7):512–534.

Kunreuther, H., Meyer, R., and Michel-Kerjan, E. (2013). *Overcoming decision biases to reduce losses from natural catastrophes*, chapter 23, pages 398–413. Princeton University Press.

Lemley, M., Kendall, J., and Martin, C. (2013). Rush to judgment: Trial length and outcomes in patent cases. *AIPLA Quarterly Journal*, 41(2):169–204.

Lipnowski, E. and Mathevet, L. (2018). Disclosure to a psychological audience. *AEJ: Microeconomics*, 10(4):67–93.

Loftus, G. and Wickens, T. (1970). Effect of incentive on storage and retrieval processes. *Journal of Experimental Psychology*, 85(1):141–147.

McCoy, S. and Walsh, R. (2018). Wildfire risk, salience & housing demand. *Journal of Environmental Economics and Management*, 91:203–228.

Milgrom, P. and Shannon, C. (1994). Monotone comparative statics. *Econometrica*, 62(1):157–180.

Miller, G. (1994). The magical number seven, plus or minus two some limits on our capacity for processing information. *Psychological Review*, 101(2):343–352. Reprint.

Mitchell, J. (2000). Why should the prosecutor get the last word? *American Journal of Criminal Law*, 27(2):139–216.

Mullainathan, S. (2002). A memory-based model of bounded rationality. *The Quarterly Journal of Economics*, 117(3):735–774.

Mullet, H. and Marsh, E. (2015). Correcting false memories: Errors must be noticed and replaced. *Memory & Cognition*, 44(3):403–412.

Murre, J. M. J. and Dros, J. (2015). Replication and analysis of ebbinghaus forgetting curve. *PLOS ONE*, 10(7).

Murty, V., FeldmanHall, O., Hunter, L. E., Phelps, E. A., and Davachia, L. (2016). Episodic memories predict adaptive value-based decision-making. *Journal of Experimental Psychology: General*, 145(5):548–558.

Nichols, J. and McChesney, R. (2013). *Dollarocracy: How the Money-and media Election Complex Is Destroying America*. Nation Books.

Panagopoulos, C. (2011). Timing is everything? primacy and recency effects in voter mobilization campaigns. *Political Behavior*, 33(1):71–93.

Pennanen, T., Perkkiö, A.-P., and Rásonyi, M. (2017). Existence of solutions in a nonconvex dynamic programming and optimal investment. *Mathematical and Financial Economics*, 11(2):173–188.

Pereira, M. and Waterbury, N. (2019). Do voters discount political scandals over time? *Political Research Quarterly*, 72(3):584–595.

Pinsker, R. (2011). Primacy or recency? a study of order effects when nonprofessional investors are provided a long series of disclosure. *Behavioral Research in Accounting*, 23(1).

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2):59–108.

Roberts, W. and Varberg, D. (1973). *Convex Functions*. Academic Press, New York.

Sanjurjo, A. (2014). The role of memory in search and choice. Unpublished.

Sanjurjo, A. (2015). Search, memory, and choice error: An experiment. *PLOS ONE*, 10(6).

Schacter, D., Gilbert, D., and Wegner, D. (2009). *Psychology*, chapter Semantic and episodic memory, pages 185–186. Worth Publishers.

Shadlen, M. and Shohamy, D. (2016). Decision making and sequential sampling from memory. *Neuron*, (90):927–938.

Spaniol, J., Schain, C., and Bowen, H. (2014). Reward-enhanced memory in younger and older adults. *Journal of Gerontology: Series B*, 69(5):730–740.

Squire, L. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory*, 82(3):171–177.

Squire, L., Genzal, L., Wixted, J., and Morris, R. (2015). Memory consolidation. *Cold Springs Harbor Perspectives in Biology*, 7(8).

Stokey, N. and Lucas, R. (1989). *Recursive Methods in Economics Dynamics*, chapter Dynamic Programming Under Certainty, pages 66–102. Harvard University Press.

Volkman-Wise, J. (2015). Representativeness and managing catastrophe risk. *Journal of Risk and Uncertainty*, 51:267–290.

Wachter, J. A. and Kahana, M. J. (2020). A retrieved-context theory of financial decisions. In *10th Miami Behavioral Finance Conference*.

Weiner, B. (1966). Motivation and memory. *Psychological Monographs: General and Applied*, 80(5):1–22.

Weiner, B. (1968). Motivated forgetting and the study of repression. *Journal of Personality*, 36(2):213–234.

Wieth, M. and Burns, B. (2006). Incentives improve performance on both incremental and insight problem solving. *The Quarterly Journal of Experimental Psychology*, 59(8).

Wilson, A. (2014). Bounded memory and biases in information processing. *Econometrica*, 82(6):2257–2294.

Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review*, 110(2):337–363.

Appendix

A Proofs

A.1 Proof of Proposition 1

Due to the way normal signals compound, from the point of view of the player in Period T, the memory signal m_T^t behaves as a single normal signal with a mean θ and a precision

$$h_t = \left(\sigma^2 + \tfrac{g_{T-t}}{n}\right)^{-1}$$

Because normal signals with higher precision dominate those with lower precision in the Blackwell (1953) sense $E(u_T(a_T, \theta)|h_t)$, we can define as the player's gross expected utility as a function of precision $U(h_t) = E(u_T(a_T, \theta)|h_t)$. Say $C(h_t, T-t) = c\left(\left(\frac{1}{h_t} - \sigma^2\right)^{-1} g_{T-t}\right)$ is the cost of achieving that precision given a specific delay. We can rewrite the player's effort selection problem as a precision selection problem.

$$h_t^* = \max_{h_t} U(h_t) - C(h_t, T-t)$$

We want to show that optimal $U(h_t)$ is decreasing in $T-t$, which is equivalent to saying h_t^* is decreasing in $T-t$.

In order to apply Theorem 5 of Milgrom and Shannon (1994) directly, we need to modify the problem slightly. Say $\tilde{h} = -h$ and for convenience rewrite $T - t = \tilde{T}$. The player's problem is then

$$\max_{\tilde{h}} U(-\tilde{h}) - C(-\tilde{h}, \tilde{T})$$

Theorem 5 of Milgrom and Shannon (1994) guarantees that the optimal $\tilde{h}$ will be weakly increasing in the strong set order as long as $U(-\tilde{h}) - C(-\tilde{h}, \tilde{T})$ has increasing differences in $(\tilde{h}, \tilde{T})$. This is equivalent to saying h_t^* and therefore the player's gross expected utility is decreasing in $T-t$. Therefore we need to show

$$U(-\tilde{h}) - C(-\tilde{h}, \tilde{T}') - \left(U(-\tilde{h}') - C(-\tilde{h}', \tilde{T}')\right) \geq U(-\tilde{h}) - C(-\tilde{h}, \tilde{T}) - \left(U(-\tilde{h}') - C(-\tilde{h}', \tilde{T})\right)$$

$$\forall \tilde{h} > \tilde{h}', \tilde{T}' > \tilde{T}$$

This can be simplified as

$$C(-\tilde{h}', \tilde{T}') - C(-\tilde{h}, \tilde{T}') \geq C(-\tilde{h}', \tilde{T}) - C(-\tilde{h}, \tilde{T}) \forall \tilde{h} > \tilde{h}', \tilde{T}' > \tilde{T}$$

or

$$C(h_t', \tilde{T}') - C(h_t, \tilde{T}') \geq C(h_t', \tilde{T}) - C(h_t, \tilde{T}) \forall h_t' < h_t, \tilde{T}' > \tilde{T}$$

or

$$c\left(\left(\frac{1}{h_t'} - \sigma^2\right)^{-1} g_{\tilde{T}'}\right) - c\left(\left(\frac{1}{h_t} - \sigma^2\right)^{-1} g_{\tilde{T}'}\right) \geq c\left(\left(\frac{1}{h_t'} - \sigma^2\right)^{-1} g_{\tilde{T}}\right) - c\left(\left(\frac{1}{h_t} - \sigma^2\right)^{-1} g_{\tilde{T}}\right)$$

$$\forall h_t' < h_t, \tilde{T}' > \tilde{T}$$

which is true since $c(\bullet)$ is increasing and convex, $h_t, g_t \geq 0$, g_t is increasing in t, and

$$\left(\frac{1}{h_t'} - \sigma^2\right)^{-1} > \left(\frac{1}{h_t} - \sigma^2\right)^{-1}$$

□

A.2 Proof of Remark 1

Consider two reward levels, $r' > r$. Call the resulting optimal effort levels $n_t^{*\prime}$ and n_t^* respectively. Assume for the sake of contradiction that $n_t^{*\prime} < n_t^*$. By optimality of n_t^* we have

$$r \times E(f(a_T, \theta)|n_t^*) - c(n_t^*) \geq r \times E(f(a_T, \theta)|n_t^{*\prime}) - c(n_t^{*\prime})$$

which we rewrite as

$$r \times (E(f(a_T, \theta)|n_t^*) - E(f(a_T, \theta)|n_t^{*\prime})) \geq c(n_t^*) - c(n_t^{*\prime})$$

Note that $E(f(a, \theta)|n_t^*) - E(f(a, \theta)|n_t^{*\prime}) \geq 0$ since $n_t^* > n_t^{*\prime}$ and since a normal signal with lower variance Blackwell dominates one with higher variance.[26]

By optimality of $n_t^{*\prime}$ we have

$$r' \times E(f(a_T, \theta)|n_t^{*\prime}) - c(n_t^{*\prime}) \geq r' \times E(f(a_T, \theta)|n_t^*) - c(n_t^*)$$

[26]Blackwell, 1953

which we rewrite as

$$c(n_t^*) - c(n_t^{*\prime}) \geq r' \times (E(f(a_T, \theta)|n_t^*) - E(f(a_T, \theta)|n_t^{*\prime}))$$

Together these imply

$$r \times (E(f(a_T, \theta)|n_t^*) - E(f(a_T, \theta)|n_t^{*\prime})) \geq r' \times (E(f(a_T, \theta)|n_t^*) - E(f(a_T, \theta)|n_t^{*\prime}))$$

which can only be possible if

$$E(f(a_T, \theta)|n_t^*) - E(f(a_T, \theta)|n_t^{*\prime}) = 0$$

However, as established previously, we must have

$$E(f(a_T, \theta)|n_t^*) - E(f(a_T, \theta)|n_t^{*\prime}) \geq \tfrac{c(n_t^*) - c(n_t^{*\prime})}{r} > 0$$

Hence we have a contradiction. $\square$

A.3 Proof of Remark 2

This result is a fairly direct application of the theorem of the maximum. The player chooses

$$n_t^* \in \max_{n_t} E(u_T(a_T, \theta)|n_t, \pi(\lambda)) - c(n_t)$$

$$s.t.\ n_t \in \mathbb{R}^+$$

The objective function is continuous in λ and quasi-concave, so by the theorem of the maximum $n_t^*(\lambda)$ is a continuous function of λ. In addition $n^*(0) = 0$.

A.4 Proof of Proposition 2

We prove this by contradiction. Due to the normality of signals and the prior, player beliefs depend on the vector of precisions for memory signals arriving in period T, $\boldsymbol{h}^T$, only through $\sum h_t^T$ and do so in a monotone increasing fashion. Say $t' > t$ and $h_t^T > h_{t'}^T$. Note achieving these precisions requires $n_t = g_{T-t} h_t^T$ and $n_{t'} = g_{T-t'} h_{t'}^T$. Note $g_{T-t} > g_{T-t'}$, so $n_t > n_{t'}$. One could achieve the same total precision by setting $\tilde{n}_t = g_{T-t} h_{t'}^T$ and $\tilde{n}_{t'} = g_{T-t'} h_t^T$.

Note $\tilde{n}_t + \tilde{n}_{t'} < n_t + n_{t'}$ and $n_t > \max(\tilde{n}_t, \tilde{n}_{t'})$, so $\boldsymbol{n} \succ_w \tilde{\boldsymbol{n}}$ where $\succ_w$ denotes weak majorization. By the definition of Schur-Convexity this means $c(\boldsymbol{n}) > c(\tilde{\boldsymbol{n}})$ and both vectors produce the same total precision, contradicting the optimality of $\boldsymbol{n}$.

A.5 Proof of Proposition 3

We begin by showing that the FOC must apply for some optimal n_t. Note that $E(u_T(a_T, \theta))$ depends on h_t^T only through H_T.

The first derivative of gross expected utility with respect to n_t is

$$\tfrac{d}{dH_T} E(u_T(a_T, \theta)|H_T) \tfrac{dH_T}{dn_t}$$

This expression is positive, because $\frac{dH_T}{dn_t} = \frac{dh_t}{dn_t} = \left(\sigma^2 n_t + g_{T-t}\right)^{-2} g_{T-t} > 0$. Note that this value is strictly positive as n_t goes to 0. In addition, by inspection we can see that as n_t goes to infinity, the derivative goes to zero, so it is asymptotically bounded.

The second derivative of utility is

$$\tfrac{d^2}{dH_T^2} E(u_T(a_T, \theta)|H_T) \left(\tfrac{dh_t}{dn_t}\right)^2 - \tfrac{d}{dH_T} E(u_T(a_T, \theta)|H_T) \tfrac{d^2 h_t}{dn_t^2}$$

which is negative in every term, since $\frac{d^2 h_t^T}{dn_t^2} = -2\sigma^2 \left(\sigma^2 n_t + g_{T-t}\right)^{-3} g_{T-t} < 0$. This suffices to show that any interior optimal n_t must satisfy the following FOC:

$$\tfrac{d}{dH_T} E(u_T(a_T, \theta)|H_T) \tfrac{dh_t}{dn_t} = \tfrac{d}{dn_t} c(\boldsymbol{n})$$

Since gross expected utility is bounded and cost grows arbitrarily, any non-interior solution must have $n_t = 0$.

Recall that $c(\boldsymbol{n})$ is a function of $\sum_{t=1}^{T-1} n_t$. This guarantees that $\frac{d}{dn_t} c(\boldsymbol{n}) = \frac{d}{dn_\tau} c(\boldsymbol{n})$ $\forall t, \tau$. This in turn means we can find a relationship between any interior n_ts or h_ts by taking the ratio of the FOCs, which gives the following result.

$$\tfrac{d}{dn_t} h_t^T = \tfrac{d}{dn_\tau} h_\tau^T$$

which can be written out and transformed to

$$n_t = \tfrac{1}{\sigma^2} \left(\left(n_\tau \sigma^2 + g_{T-t}\right) \sqrt{\tfrac{g_{T-t}}{g_{T-\tau}}} - g_{T-t} \right)$$

Note that $n_\tau = \frac{g_{T-\tau}}{h_\tau^{-1}-\sigma^2}$, which we can substitute into the above to get

$$n_t = \left(\frac{g_{T-\tau}}{h_\tau^{-1}-\sigma^2} + \frac{g_{T-\tau}}{\sigma^2}\right)\sqrt{\frac{g_{T-t}}{g_{T-\tau}}} - \frac{g_{T-t}}{\sigma^2}$$

Rearrange and multiply by g_t to get

$$\frac{g_{T-t}}{n_t} = \frac{\sigma^2}{\left(\frac{1}{(h_\tau^T)^{-1}-\sigma^2}+1\right)\sqrt{\frac{g_{T-\tau}}{g_{T-t}}}-1}$$

Then add σ^2 and substitute in for h_t to get

$$\left(h_\tau^T\right)^{-1} = \sigma^2\left(1 + \frac{1}{\left(\frac{1}{h_\tau^{-1}-\sigma^2}+1\right)\sqrt{\frac{g_{T-\tau}}{g_{T-t}}}-1}\right)$$

Rearrange to get

$$h_\tau^T = \frac{1}{\sigma^2}\left(1 - \frac{\left(h_\tau^T\right)^{-1}-\sigma^2}{\left((h_\tau^T)^{-1}-\sigma^2+1\right)g_{T-\tau}}\sqrt{g_{T-t}}\right)$$

Note that if any n_ts are interior, n_{T-1} must be as well, because if not, one could set n_{T-1} equal to $n_t > 0$ and set that n_t equal to zero and improve overall total precision without changing cost. Therefore, we can set $\tau = T-1$ and define $\frac{\left(h_{T-1}^T\right)^{-1}-\sigma^2}{\left(\left(h_{T-1}^T\right)^{-1}-\sigma^2+1\right)g_1} = z$ which gives us

$$h_t^T = \frac{1}{\sigma^2}\left(1 - z\sqrt{g_{T-t}}\right)$$

for any interior h_t^T. Otherwise, $h_t^T = 0$.

□

A.6 Proof of Proposition 4

We prove this by contradiction. Consider two election games which are identical except game one ends at time T_1 and game two ends at time T_2. Assume $T_2 > T_1$.

Note that $E(u_T(a_T, \theta))$ is a function of the total memory signal precision H_T. Call this function $U(H)$.

Call the total precision of players in the game ending at T_1, H_1^* and the total precision selected by players in the game ending at time T_2, H_2^*.

Assume for a contradiction that the probability of a good election is higher in game one. This implies $H_1^* > H_2^*$. The individual's total precision selection problem in game i can be written as

$$\max_H U(H) - C(H,T)$$

where $C(H,T)$ is the cost incurred for getting a total memory precision of H given a T period campaign season. Note that $C(H,T)$ is decreasing in T because any vector of n_t's that is feasible when T is smaller is also feasible when T is larger for at most the same cost.

We need to prove one lemma on $C(H,T)$ before continuing the proof.

Lemma 2. *$C(H,T)$ has decreasing differences in (H,T).*

Proof. We want to show that

$$C(H_2,T_2) - C(H_1,T_2) < C(H_1,T_1) - C(H_1,T_1)$$

$\forall H_2 > H_1, T_2 > T_1$.

We can write $C(H,T)$ as

$$C(H,T) = \min_{\boldsymbol{h}} c\left(\boldsymbol{n}\left(\boldsymbol{h}\right)\right)$$

$$s.t. \sum_{t=1}^{T} h_t \geq H$$

where $\boldsymbol{h}$ is the vector of h_t's and the tth element of $\boldsymbol{n}\left(\boldsymbol{h}\right)$ is given by

$$n_t(h_t) = g_t \left(\tfrac{1}{h_t} - \sigma^2\right)^{-1}$$

We also define the reverse transformation so the tth element of $\boldsymbol{h}(\boldsymbol{n})$ is given by

$$h_t(n_t) = \left(\tfrac{g_t}{n_t} + \sigma^2\right)^{-1}$$

Finally define $\boldsymbol{h}_i(T_j)$ optimal vector of precision for achieving H_i given the game length T_j with a corresponding effort vector $\boldsymbol{n}_i(T_j)$.

We want to show

$$c(\boldsymbol{h_1}(T_2)) - c(\boldsymbol{h}_1(T_1)) \geq c(\boldsymbol{h}_2(T_2)) - c(\boldsymbol{h}_2(T_1))$$

We begin by showing $\boldsymbol{h}_2 > \boldsymbol{h}_1$ element-wise. We do this by looking at the constrained optimization problem which determines $\boldsymbol{h}$ in any given problem

$$\min_{\boldsymbol{n}} f\left(\sum g\left(n_i\right)\right)$$

$$s.t. \sum \boldsymbol{h}(\boldsymbol{n}) = H_i$$

which has a Lagrangian

$$f\left(\sum g\left(n_i\right)\right) - \lambda(H(\boldsymbol{n}) - H)$$

Take the FOC w.r.t. n_i to get

$$g'(n_i)/\tfrac{dH}{dn_i} = \lambda/f'(G)$$

for all i where it is possible. For all other i, $n_i = 0$. Note the LHS is increasing in n_i and does not depend on T and the RHS is identical for all i in a given problem. This going from H_1 to H_2 increases $\lambda/f'(G)$, and therefore increases all n_i. Therefore, $\boldsymbol{h_2} > \boldsymbol{h_1}$ element-wise.

Next define$\boldsymbol{\Delta h}_1 = \boldsymbol{h_1}(T_2) - \boldsymbol{h_1}(T_1)$.

Note that $\boldsymbol{\Delta h}_1$ has two components, $\boldsymbol{\Delta h}_1[1; T_1] \leq 0$ and $\boldsymbol{\Delta h}_1[T_1 + 1; T_2] \geq 0$. This comes from the previous result, since if $\boldsymbol{\Delta h}_1[1; T_1] > 0$ in any element it must also be greater in every element which contradicts the fact that $\sum \Delta \boldsymbol{h}_1 = 0$.

Further define $\boldsymbol{h}' = \boldsymbol{h_2}(T_1) + \boldsymbol{\Delta h}(T_1)$

By optimality $c(\boldsymbol{h}') > c(\boldsymbol{h_2}(T_2))$, so

$$c(\boldsymbol{h_2}(T_2)) - c(\boldsymbol{h_1}(T_2)) \leq c(\boldsymbol{h}') - c(\boldsymbol{h_1}(T_2))$$

Therefore it suffices to show

$$c(\boldsymbol{h_1}(T_2)) - c(\boldsymbol{h}_1(T_1)) \geq c(\boldsymbol{h}') - c(\boldsymbol{h}_2(T_1))$$

Define $G(\boldsymbol{h}) = \sum g(n(h_i))$.

Note that $n(h_i)$ is also a convex increasing function, so $G(\bullet)$ is increasing and convex.

By the convexity of f it is sufficient to show

$$G(\boldsymbol{h}_1(T_1)) \leq G(\boldsymbol{h}_2(T_1))$$

which is true trivially

and

$$G(\boldsymbol{h}_1(T_1)) - G(\boldsymbol{h_1}(T_2)) \leq G(\boldsymbol{h}_2(T_1)) - G(\boldsymbol{h}')$$

Note that components $\boldsymbol{\Delta h}_1[T_1 + 1; T_2]$ contribute equally to both sides of the equation while $\boldsymbol{\Delta h}_1[1;T_1]$ contributes more to the RHS since $g(\bullet)$ is convex and $\boldsymbol{h}_2(T_1) \geq \boldsymbol{h}_1(T_1)$ element-wise.

□

By optimality we must have

$$U(H_1^*) - C(H_1^*, T_1) \geq U(H_2^*) - C(H_2^*, T_1)$$

and

$$U(H_2^*) - C(H_2^*, T_2) \geq U(H_1^*) - C(H_1^*, T_2)$$

Add together to get

$$U(H_1^*) - C(H_1^*, T_1) + U(H_2^*) - C(H_2^*, T_2)$$

$$\geq U(H_2^*) - C(H_2^*, T_1) + U(H_1^*) - C(H_1^*, T_2)$$

And rearrange to get

$$C(H_2^*, T_1) - C(H_1^*, T_1) \geq C(H_2^*, T_2) - C(H_1^*, T_2)$$

Which, given $T_2 > T_1$ only satisfies Lemma 2 if $H_2^* \geq H_1^*$.

□

A.7 Proof of Remark3

We prove this by contradiction. Say that the optimal $n_{t'}^* > n_t^*$. Due to the normality of signals and the prior, player beliefs depend on $\boldsymbol{h}^T$only through $\sum h_t^T$ and do so in a monotone increasing fashion. By symmetry, we know we can switch the memory effort levels without changing the cost of memory.

Therefore all that remains to be shown is that, given a sufficiently small δ, switching memory efforts would weakly increase $h_t^T + h_{t'}^T$. We can show this by proving $\frac{d}{dn}h_t^T(n) \geq \frac{d}{dn}h_{t'}^T(n)$.

Note that

$$\tfrac{d}{dn_t}h_t^T = \tfrac{d}{dn_t}\left(\sigma_t + \tfrac{g_{T-t}}{n_t^*}\right)^{-1} = -\left(\sigma_t + \tfrac{g_{T-t}}{n_t^*}\right)^{-2}\tfrac{d}{dn_t}\tfrac{g_{T-t}}{n_t^*}$$

$$-\left(\sigma_t + \tfrac{g_{T-t}}{n_t^*}\right)^{-2}\tfrac{d}{dn_t}\tfrac{g_{T-t}}{n_t^*} = -\left(\sigma_t + \tfrac{g_{T-t}}{n_t^*}\right)^{-2} - \tfrac{g_{T-t}}{{n_t^*}^2}$$

$$= \left(\frac{n_t^*}{\sigma_t n_t^* + g_{T-t}}\right)^2 \frac{g_{T-t}}{n_t^{*2}}$$

$$= g_{T-t} \left(\frac{1}{\sigma_t n_t^* + g_{T-t}}\right)^2$$

which is decreasing in σ_t and continuous in $g_{T-t} \forall g_{T-t} > 0$.

$\square$

A.8 Proof of Proposition 5

To prove convergence, we need Assumptions 4.3, 4.4, 4.7, and 4.8 of Stokey and Lucas (1989).

Assumption 4.3 requires that n_t be bounded. Define $\boldsymbol{n}_t^{-\tau}$ as the history of memory efforts at time t with n_τ omitted. Say $\bar{u}(t, \boldsymbol{n}_t^{-\tau}) = \lim_{n_\tau \to \infty} E(u|\boldsymbol{n}_t^{-\tau}, n_\tau)$. The marginal benefit of n_τ is bounded above by $\frac{1}{1-\beta} \max_{t, \boldsymbol{n}_t^{-\tau}} \left(\bar{u}(t, \boldsymbol{n}_t^{-\tau}) - E(u|\boldsymbol{n}_t^{-\tau}, n_\tau)\right)$, by monotonicity. By the continuity of $E(u|\boldsymbol{n}_t)$, the quantity $\bar{u}(t, \boldsymbol{n}_t^{-\tau}) - E(u|\boldsymbol{n}_t^{-\tau}, n_\tau)$ goes to 0 as $n_\tau \to \infty$ for all $t, \boldsymbol{n}_t^{-\tau}$. Therefore, the marginal benefit of n_τ goes to 0 as $n_\tau \to \infty$. Since the marginal cost of n_τ is non-decreasing and positive somewhere, costs will eventually outstrip benefits. We simply assume that n_t is bounded by some $\bar{n}$ beyond that point.

Since $1 \geq E(u|\boldsymbol{n}_{t-1}) \geq E(u|\boldsymbol{n}_0)$, we know that $U(\boldsymbol{n}_{t-1}, n_t)$ is also bounded. This fact gives us Assumption 4.4, and Assumption 4.8 is immediate from the structure of the Bellman Equation.

All that remains is Assumption 4.7. which requires that $E(u|\boldsymbol{n}_t)$ is increasing in all n_τ and concave in $\boldsymbol{n}_\tau$. Increasing comes from the fact that higher precision normal signals Blackwell (1953) dominate lower precision ones. Concavity is a condition of the proposition. We now move on to prove the convergence value.

Due to the chosen g_t we have $H(\boldsymbol{n}_{t-1}) = H_t = \sum_{\tau=1}^{t} \delta^{t-\tau} n_\tau$

This has the convenient property that

$H_{t+1} = \delta H_t + \delta n_t$

We can then write the Bellman Equation

$v(H) = \max_{n_t} u(H_t) - c(n_t) + \beta v(\delta H + \delta n_t)$

Which can be rewritten as

$v(H_t) = \max_{H_{t+1}} u(H_t) - c(\frac{1}{\delta} H_{t+1} - H_t) + \beta v(H_{t+1})$

Which has a FOC

$$\frac{d}{dH_{t+1}} v(H_t) = -\frac{1}{\delta} c'(\frac{1}{\delta} H_{t+1} - H_t) + \beta \frac{d}{dH_{t+1}} v(H_{t+1}) = 0$$

Note
$v(H_{t+1}) = u(H_{t+1}) - c(\frac{1}{\delta}H_{t+2} - H_{t+1}) + \beta v(H_{t+2})$
So by the enveloped theorem
$\frac{d}{dH_{t+1}}v(H_{t+1}) = u'(H_{t+1}) + c'(\frac{1}{\delta}H_{t+2} - H_{t+1}) + 0$
Which we can substitute in to get

$$-\frac{1}{\delta}c'(\frac{1}{\delta}H_{t+1} - H_t) + \beta u'(H_{t+1}) + \beta c'(\frac{1}{\delta}H_{t+2} - H_{t+1}) = 0$$

at the steady state

$$-\frac{1}{\delta}c'(\frac{1}{\delta}\bar{H} - \bar{H}) + \beta u'(H_{t+1}) + \beta c'(\frac{1}{\delta}\bar{H} - \bar{H}) = 0$$

$$u'(H_{t+1}) = \left(\frac{1}{\delta\beta} - 1\right) c'\left(\left(\frac{1}{\delta} - 1\right)\bar{H}\right)$$

A.9 Proof of Lemma 1

We keep this proof short and light on formalism, because it is only a minor variation on Theorem 1 of Austen-Smith and Banks (1996). As long as all other voters $j \neq i$ have incentives identical to Player i's, all other players have identically distributed signals, the prior is uniform, and the election is majority rule, then Player i's beliefs about the state of the world, conditional on pivotality, are also uniform. If Player i has uniform interim beliefs conditional on pivotality, it is optimal for him or her to vote sincerely based on his posterior.

□

A.10 Proof of Proposition 6

We want to show that there exists an $\boldsymbol{n}^*$ such that, conditional on other players choosing $\boldsymbol{n}^*$and voting sincerely, it is optimal for a given player to choose $\boldsymbol{n}^*$ and vote sincerely.

As established in the proof of Lemma 1, if all other players receive equally distributed signals and vote sincerely, the prior is uniform, and the election is majority rule, then it is optimal for a player to vote sincerely.

Given that all memory signals are independent normals, they are informationally equivalent to one normal signal equal to the precision weighted average $\frac{\sum_{t=1}^{T-1} h_t m_t^T}{\sum_{t=1}^{T-1} h_t}$. Here $h_t = \left(\sigma^2 + \frac{g_{T-t}}{n_t}\right)^{-1}$ is the precision of signal m_t^T. Recall that σ^2 and g_{T-t} are both constants determined by the parameters of the model, while n_t is memory effort, which is selected by the player.

The composite signal has a mean θ and a precision of $\sum_{t=1}^{T-1} h_t$. Due to the symmetric prior (conditional on pivotality), the player will pick option 1 as long as $\frac{\sum_{t=1}^{T-1} h_t m_t^T}{\sum_{t=1}^{T-1} h_t} > 0$. This means that the probability of a player voting correctly is $1 - \Phi(-\sqrt{\sum h_t})$. For notational convenience, we will sometimes write$\sum_{t=1}^{T-1} h_t$ as H. As such Player i's utility function is given by

$$K\left(\boldsymbol{n}^{-i}\right) * \left(1 - \mathcal{P}\left(\boldsymbol{n}^{-i}\right)\right) + \mathcal{P}\left(\boldsymbol{n}^{-i}\right)\left(1 - \Phi\left(-\sqrt{H(\boldsymbol{n}^i)}\right)\right) - c\left(\boldsymbol{n}^i\right)$$

Here $\boldsymbol{n}^i$ is the vector of n_t^is chosen by Player i, and $c(\bullet)$ is the cost of memory resources. Here $\mathcal{P}$ is the probability of pivotality, and K is the probability of the correct choice being implemented given that the player is not pivotal. Both K and $\mathcal{P}$ depend on $\boldsymbol{n}^{-i}$ which is the vector of n_t's chosen by all players other than Player i. Note that all players other than Player i have the same memory efforts, so we can represent $\boldsymbol{n}^{-i}$ with a vector that has the same number of dimensions as $\boldsymbol{n}^i$.

Therefore, to show the proposition, we simply need to show the existence of a vector $\boldsymbol{n}^*$ that satisfies

$$\boldsymbol{n}^* \in \varphi(\boldsymbol{n}^*)$$

$$\varphi(\boldsymbol{n}^*) := \arg\max_{\boldsymbol{n}^i} \mathcal{P}\left(\boldsymbol{n}^*\right)\left(1 - \Phi\left(-\sqrt{H(\boldsymbol{n}^i)}\right)\right) - c\left(\boldsymbol{n}^i\right)$$

By Kakutani's Fixed Point Theorem, such an $\boldsymbol{n}^*$ exists as long as $\boldsymbol{n}^*$ belongs to a non-empty, convex, compact subset of Euclidean space and $\varphi(\bullet)$ is closed graph, compact valued, and convex valued for all $\boldsymbol{n}^*$. We can easily bound $\boldsymbol{n}^*$ in a non-empty, convex, compact subset of Euclidean space. Simply say $\bar{n}_t$ is some value of n_t such that $c(0,0,0,...,\bar{n}_t,...,0,0,0) > 1$. It will never be optimal to pick $n_t > \bar{n}_t$, because doing so guarantees that the player would make less than zero utility while picking all zeroes will provide more than zero. Say that $\bar{n} = \max_t \bar{n}_t$. We can say WLOG that $\boldsymbol{n}^* \in [0, \bar{n}]^T$, which is a non-empty, convex, compact subset of Euclidean space.

To prove that $\varphi(\bullet)$ is closed graph, compact-valued, and convex valued for all $\boldsymbol{n}^*$, we need to use the theorem of the maximum.

Theorem of the maximum states that if

$$\mathcal{P}\left(\boldsymbol{n}^*\right)\left(1 - \Phi\left(-\sqrt{H(\boldsymbol{n}^i)}\right)\right) - c\left(\boldsymbol{n}^i\right)$$

is continuous in $\boldsymbol{n}^i$ for all $\boldsymbol{n}^*$ then $\varphi(\bullet)$ is non-empty, compact-valued, and upper hemicontinuous. This would immediately guarantee $\varphi(\bullet)$ is closed graph as well, because it is a correspondence between closed subsets of $\mathbb{R}^T$ which are metric spaces. Compact-valued and upper hemicontinuous correspondences are closed graph if they map into a metric space. We get continuity of the expression immediately from the fact that it is a composition of continuous functions.

To get $\varphi(\bullet)$ convex-valued, from the Theorem of the Maximum, we also need to show that

$$\mathcal{P}\left(\boldsymbol{n}^*\right)\left(1-\Phi\left(-\sqrt{H(\boldsymbol{n}^i)}\right)\right)-c\left(\boldsymbol{n}^i\right)$$

is quasiconcave. It suffices to show that $1-\Phi\left(-\sqrt{H(\boldsymbol{n}^i)}\right)$ is concave, since $\mathcal{P}\left(\boldsymbol{n}^*\right)>0$ and $c\left(\boldsymbol{n}^i\right)$ is concave. Recall that a player's probability of pivotality does not depend on their own memory effort.

If $H(\boldsymbol{n}^i)$ is concave and increasing in each element, then $\sqrt{H(\boldsymbol{n}^i)}$ is concave and increasing in each element, by properties of the composition of concave increasing functions. This implies $-\sqrt{H(\boldsymbol{n}^i)}$ is convex and decreasing. $\Phi(x)$ is convex and increasing for $x \leq 0$ and $-\sqrt{H(\boldsymbol{n}^i)}<0$, so again by composition properties, $\Phi\left(-\sqrt{H(\boldsymbol{n}^i)}\right)$ is convex, which means $-\Phi\left(-k\sqrt{H(\boldsymbol{n}^i)}\right)$ is concave.

Therefore, it suffices to show that $H(\boldsymbol{n}^i)$ is concave. $H(\boldsymbol{n}^i)$ is concave as long as its Hessian matrix is negative semi-definite. The Hessian is a diagonal matrix with elements $\frac{d^2h_t}{dn_t^2}$, because all the cross derivatives are zero, so if we can show $\frac{d^2h_t}{dn_t^2}\leq 0$ we are done.

$$\frac{dh_t}{dn_t}=\left(\sigma^2+\frac{g_{T-t}}{n_t}\right)^{-2}\frac{g_{T-t}}{n_t^2}$$

$$=\left(\sigma^2 n_t+g_{T-t}\right)^{-2}g_{T-t}$$

So

$$\frac{d^2h_t}{dn_t^2}=-2\sigma^2\left(\sigma^2 n_t+g_{T-t}\right)^{-3}g_{T-t}$$

Note that $g_{T-t}, n_t, \sigma^2>0$, so

$$-2\sigma^2\left(\sigma^2 n_t+g_{T-t}\right)^{-3}g_{T-t}\leq 0$$

□

A.11 Proof of Corollary 4

we simply need to show that expected utility is continuous in H_T. Expected utility is given by

$$\mathcal{P}\phi\left(-\sqrt{H_T}\right)\left(\frac{1}{2\sqrt{H_T}}\right)\frac{dh_t}{dn_t}$$

where $\mathcal{P}$ is the individual player's probability of pivotality. By inspections, we can see that this is a continuous function. □

A.12 Proof of Corollary 5

We know from the proof of Proposition 4 that $\sqrt{\sum h_\tau}$ is concave and increasing in $\boldsymbol{n}$. If we can show

$$f(x) = \frac{u_B+u_G}{2} + \frac{1}{4}(u_G - u_B)\left(2\Phi\left(\left(\frac{\theta_B-\theta_G}{2}\right)x\right) - 1\right)$$

is concave in x, then we will have the result from the property of compositions of concave increasing functions. We can again ignore the constant components and focus on

$$\frac{1}{4}(u_G - u_B)\left(2\Phi\left(\left(\frac{\theta_B-\theta_G}{2}\right)x\right)\right)$$

which is a concave increasing function of x, because $u_G > u_B$, $\theta_B > \theta_G$ and $\Phi(x)$ is an increasing concave function of x as long as $x \geq 0$. Recall that the precisions are always weakly positive.

B Closed Form Proofs

B.1 OSOA Closed Form Proof

$$E(u|H(\boldsymbol{n}_t)) = -\frac{1}{h_T(n_t)} = -\frac{g_{T-t}+n_t\sigma^2}{n_t} = -\left(\frac{g_{T-1}}{n_t} + \sigma^2\right)$$

So the player's objective function is$-\left(\frac{g_{T-1}}{n_t} + \sigma^2\right) - \kappa n_t^2$
Which has a FOC $\frac{g_{T-1}}{n_t^2} = 2\kappa n_t$
Which we can solve for $\left(\frac{g_{T-1}}{\kappa 2}\right)^{1/3} = n_t^*$
Plug in for gross expected utility to get $-\left(\left(2g_{T-1}^2\kappa\right)^{1/3} + \sigma^2\right)$

B.2 MSOA Closed Form

$E(u|H(\boldsymbol{n}_T)) = -\frac{1}{H(\boldsymbol{n}_T)}$ where $H(\boldsymbol{n}_T) = \sum_{t=1}^{T-1} \frac{n_t}{g_{T-t}}$

So the objective function is $-\frac{1}{H(\boldsymbol{n}_T)} - \kappa \sum_{t=1}^{T-1} n_t^2$

Which has a FOC $\frac{1}{H(\boldsymbol{n}_T)^2} \left(\frac{1}{g_{T-t}}\right) = \kappa n_t$

Take the ratio of the FOC for τ and τ' to get the ratio condition: $\frac{n_\tau}{n_{\tau'}} = \frac{g_{t-\tau'}}{g_{t-\tau}}$

Rewrite everything as a function of n_{T-1}using $n_t = \frac{g_1}{g_{T-t}} n_{T-1}$ and define $G = \sum_{i=1}^{t-1} \frac{1}{g_i^2} = G$

Objective function can be rewritten as $-\frac{1}{n_{T-1} g_1 G} - \kappa n_{T-1}^2 g_1^2 G$

Take the FOC $\frac{1}{n_{T-1}^2 g_1 G} - \kappa 2 n_{T-1} g_1^2 G = 0$

Solve to get$(g_1^3 G^2 2\kappa)^{-1/3} = n_{T-1}^*$

Plug into the ratio condition to get

$n_t^* = \frac{\left(G^2 2\kappa\right)^{-1/3}}{g_{T-t}}$

$H^* = \left(\frac{G}{2\kappa}\right)^{1/3}$

Optimal Gross Expected Utility $-\left(\frac{2\kappa}{G}\right)^{1/3}$

B.3 MSMA Closed Form

$\frac{1}{\delta}$

$\Delta H_t = \left(\frac{1}{\alpha} - 1\right) H_t + \frac{1}{\alpha} n_t$ from the evolution of H_t

Sub $c(\bullet)$ into Proposition 5 to get

$$\beta u'(H_{t+1}) + 2\kappa \frac{1}{\delta} H_t - 2\kappa \left(\frac{1}{\delta^2} + \beta\right) H_{t+1} + 2\kappa\beta \frac{1}{\delta} H_{t+2} = 0$$

or

$$\beta u'(H_{t+1}) - 2\kappa \frac{1}{\delta} n_t + 2\kappa\beta n_{t+1} = 0$$

Since $\alpha H_{t+1} - H_t = n_t$ we can rewrite as

$$\Delta n_t = \left(\frac{1}{\delta\beta} - 1\right) n_t - \frac{1}{2\kappa} u'(H_{t+1})$$

We can further sub in for H_{t+1} to get

$$\Delta n_t = \left(\frac{1}{\delta\beta} - 1\right) n_t - \frac{1}{2\kappa} u'\left(\delta H_t + \delta n_t\right)$$

Steady State solution

$$\frac{\beta}{2\kappa} u'(\bar{H}) = \left(\frac{1}{\delta} - 1\right) \left(\frac{1}{\delta} - \beta\right) \bar{H}$$

Now sub in for $u(\bullet)$

Find the flow field

$$0 = (\delta - 1) H_t + \delta n_t$$
$$0 = \left(\frac{1}{\beta} - 1\right) n_t - \frac{1}{2\kappa(\delta H_t + \delta n_t)^2}$$

Find the steady state

$$\frac{\beta}{2\kappa} \frac{-1}{(H)^2} = (\frac{1}{\delta} - 1)(\frac{1}{\delta} - \beta) H$$

$$H^* = \left(\frac{\beta}{2\kappa\left(\frac{1}{\delta} - \beta\right)\left(\frac{1}{\delta} - 1\right)}\right)^{1/3}$$

The quantity in the root is positive, so this has only one real solution. By the monotonicity of the H_t evolution equation there is only one $\bar{n}$ for any $\bar{H}$.

C Adjustability

As discussed in Section 2.2, the framework we are using does not allow for players to adjust their memory effort in response to new information about the value of their memories (by, for example, devoting fewer neurons to a memory they have realized is unimportant or rehearsing a memory they expect to need). On the other hand, allowing for full adjustability could also produce unrealistic predictions such as unrealistically quick forgetting. In this section, we construct two thought experiments, with the first showing the pitfalls of leaving adjustability out of the framework, and the second showing how allowing for full adjustability can lead to its own problematic predictions.

C.1 Testing Adjustability

We now construct a simple thought experiment for which our baseline framework with inflexible memory plans generates somewhat counter intuitive predictions. Consider a three-period game, which we will call the adjustability experiment game. We deviate slightly from the baseline framework by saying that the state θ has two components, a reward level $r \in \{r_h, r_l\}$ with $r_h > r_l$ and an action relevant component $\eta \in \mathbb{R}$. The order of play in the game goes as follows:

1. In Period 1, the player learns η and r.

2. Period 2 passes.

3. In Period 3, the player receives a memory signal m_3 about η and takes an action a_3.

After the game is complete, the player receives a reward of

$$u_3(a_3, \theta) = r * f(a_3, \eta)$$

Assume for simplicity that once learned, r is remembered with certainty for free.

Non-Adjustable Memory: If the player has non-adjustable memory, as is the case in the baseline framework used throughout the paper, he chooses one memory effort level at the beginning of Period 1 and pays a cost

$$c(n_1)$$

In the last period, he receives a memory signal

$$m_3 = \eta + \epsilon$$

$$\epsilon \sim N\left(0, \tfrac{g_2}{n_1}\right)$$

The baseline player will not be able to adjust the precision of his memory in response to r.

Adjustable Memory: Now consider a player who can adjust his memory effort. Say the adjustable memory player selects an initial memory effort, n_1, at the beginning of period 1 and then selects a second memory effort, n_2, at the beginning of Period 2 after learning η and, more importantly, r. His memory signal is then

$$m_3 = \eta + \epsilon_1 + \epsilon_2$$

where

$$\epsilon_1 \sim N\left(0, \tfrac{g_1}{n_1}\right)$$

$$\epsilon_2 \sim N\left(0, \tfrac{g_1}{n_2}\right)$$

and he pays a cost

$$\tilde{c}(n_1, n_2)$$

We use $\tilde{c}(\bullet)$ to differentiate the cost functions between frameworks. Note that $c(\bullet)$ takes one argument while $\tilde{c}(\bullet)$ takes two.

Proposition 7. *Within the adjustability experiment game, the adjustable memory player's equilibrium satisfies*

$$E(u(a_3,\eta)|r_h) \geq E(u(a_3,\eta)|r_l)$$

while the baseline player's equilibrium satisfies

$$E(u(a_3,\eta)|r_h) = E(u(a_3,\eta)|r_l)$$

Proof. That performance is identical in the non-adjustable framework is trivial, since the memory plan can only be conditioned on information known to the player when he receives the signal. We will prove the remaining result by contradiction. Take the player's choice of n_1 as given, since this choice cannot be conditioned on r. In the adjustable framework, the player in Period 2 must solve the problem

$$\max_{n_2} r \times E(u(a,\eta)|n_2,n_1) - \tilde{c}(n_2,n_1)$$

Conjecture that there are effort levels n_2^h and n_2^l, corresponding to r_h and r_l respectively, such that $E(u(a,\eta)|n_2^l,n_1) > E(u(a,\eta)|n_2^h,n_1)$.

By incentive compatibility we know

$$r_l \times E(f(a,\eta)|n_2^l,n_1) - \tilde{C}(n_2^l,n_1) > r_l \times E(f(a,\eta)|n_2^h,n_1) - \tilde{c}(n_2^h,n_1)$$

and

$$r_h \times E(f(a,\eta)|n_2^h,n_1) - \tilde{c}(n_2^h,n_1) > r_h \times E(f(a,\eta)|n_2^l,n_1) - \tilde{c}(n_2^l,n_1)$$

which implies

$$\tilde{c}(n_2^l,n_1) - \tilde{c}(n_2^h,n_1) > r_h \times E(f(a,\eta)|n_2^l,n_1) - r_h \times E(f(a,\eta)|n_2^h,n_1)$$

and

$$\tilde{c}(n_2^l,n_1) - \tilde{c}(n_2^h,n_1) < r_l \times E(f(a,\eta)|n_2^l,n_1) - r_l \times E(f(a,\eta)|n_2^h,n_1)$$

But if $E(f(a,\eta)|n_2^l,n_1) > E(f(a,\eta)|n_2^h,n_1)$ then

$$r_h \times \left(E(f(a,\eta)|n_2^l, n_1) - E(f(a,\eta)|n_2^h, n_1)\right) >$$

$$r_l \times \left(E(f(a,\eta)|n_2^l, n_1) - E(f(a,\eta)|n_2^h, n_1)\right)$$

is impossible, since $r_h > r_l$. This means it is impossible for both IC constraints to hold. Thus by contradiction we have the result.

Note that this is essentially an Application of the NIAC from Caplin and Dean (2015). As such, it can also apply to more general settings without normal signals.

In this case, fully adjustable memory leads to a more reasonable prediction, because new information changes the expected value of old information. Learning that information will be more or less valuable should increase or decrease a person's tendency to remember that information.

Note that if information about r did not arrive after an effort level had already been chosen, both frameworks would be effectively identical. To see this note that, if no information on r is provided, the only behaviorally important feature of the cost function in the adjustable framework is the minimum cost required to achieve any particular precision for m_3. No other features are behaviorally relevant, because a player will never use any set of n_1 and n_2 to achieve their desired precision that costs more than the minimum. If $c(\bullet)$ and $\tilde{c}(\bullet)$ produce the same minimum costs of precision for each possible precision value, then we can rewrite both the adjustable effort selection problem and the non-adjustable effort selection problem as the same precision selection problem.

While it may seem like this example depends heavily on the exact timing we assume in our framework, we can explore the same phenomenon with different period timing just by slightly altering the form of the game. If, for example, we allow the player to choose his memory effort after receiving the exogenous signal in each period, then the same counter-intuitive result can be obtained if we modify the game by having the player learn r in period 2.

C.2 Pitfalls of Full Adjustability

Allowing for full adjustability of memory resources creates its own set of unrealistic predictions. In particular, it predicts an extremely rapid drop-off in memory quality once the future value of a memory is diminished.

Consider the following thought experiment. Student 1 is told to memorize a list of words. In 15 minutes, he will be asked to repeat the list for $100. After another 15 minutes he is asked to do so again, but this time the prize is only $1 for success. Student 2 is given the same list and asked to memorize it. After 15 minutes he will be asked to repeat the list for a $1 prize with no further tests.

Who will do better on their $1 test? Intuitively, Student 1 has already spent the memorization effort for the previous $100 prize, so their performance should be higher. However, in a fully adjustable memory model, their performance will be lower. We formalize this below.

Adjustability Pitfalls Game: First, consider a simple three period game, which we will call the adjustability pitfalls game. For simplicity, assume the signal is fully revealing of the state.

1. In Period 1, the player learns θ.
2. In Period 2, he receives a signal m_2 and takes an action a_2.
3. In Period 3, the player receives a signal m_3 and takes an action a_3.

After the game is complete, the player receives a reward of

$$r_2 \times u(a_2, \theta) + r_3 \times u(a_3, \theta)$$

A player with adjustable memory selects an initial memory effort, n_1, at the beginning of Period 1 and then selects a second memory effort, n_2, at the beginning of Period 2 and pays a cost

$$\tilde{c}(n_1, n_2)$$

In Period 2, he receives a memory signal

$$m_2 = \theta + \epsilon_1^2$$

$$\epsilon_1^2 \sim N\left(0, \frac{g_1}{n_1}\right)$$

And in Period 3, he receives a memory signal

$$m_3 = \theta + \epsilon_1^3 + \epsilon_2^3$$

$$\epsilon_1^3 \sim N\left(0, \frac{g_2}{n_1}\right)$$

$$\epsilon_2^3 \sim N\left(0, \frac{g_1}{n_2}\right)$$

Comparison Game: Also consider a two-period comparison game.

1. In Period 1, the player learns θ
2. In Period 2, he receives a signal m_2 and takes an action $\tilde{a}_2$

After the game is complete, the player receives a reward of

$$r_3 \times u(\tilde{a}_2, \theta) - \tilde{c}(\tilde{n}_1, 0)$$

Note that the comparison game has the same reward level as the one associated with the action chosen in the second period of the adjustability pitfalls game. In the comparison game, whether the player has adjustable memory is irrelevant. He will select memory effort n_1 in Period 1, and in Period 2 he will receive a memory signal

$$m_2 = \eta + \epsilon_1^2$$

$$\epsilon_1^2 \sim N\left(0, \frac{g_1}{n_1}\right)$$

Before we go on, we introduce the following definition.

Definition 4. A function $c(\bullet)$ is symmetric if $c(a,b) = c(b,a) \forall a, b$.

The issue with full adjustability comes from the following result.

Proposition 8. *Under the adjustable memory framework, if $c(\bullet)$ is strictly increasing and convex in each element, symmetric, and has increasing differences in (n_1, n_2) then the performance in Period 2 of the comparison game will always be greater than the performance in Period 3 of the adjustability pitfalls game ($E(u(\tilde{a}_2, \theta)) > E(u(a_3, \theta))$).*

Proof. In the adjustability pitfalls game, under the adjustable memory framework, the player must choose n_1 and n_2 to maximize

$$r_2 \times U(h(n_1)) + r_3 U(h(n_1, n_2)) - \tilde{c}(n_1, n_2)$$

where $h()$ is the precision of a memory signal, so

$$h_t(n_1, n_2, ...n_J) = \left(\sum_{j=1}^{t} \frac{g_{J-j+1}}{n_j}\right)^{-1}$$

In addition $U(n_1) = E(u(a,\theta)|h(n_1))$ and $U(n_1,n_2) = E(u(a,\theta)|h(n_1,n_2))$. One necessary condition for optimality is that n_2 maximizes

$$r_3 \times U(h(n_1,n_2)) - \tilde{c}(n_1,n_2)$$

which defines an optimal n_2^*.

In the comparison game, the player must choose $\tilde{n}_1$to maximize

$$r_3 \times U(h(\tilde{n}_1)) - \tilde{c}(\tilde{n}_1,0)$$

Note that $h_2(n_1,n_2) = \frac{1}{\left(\frac{g_2}{n_1}+\frac{g_1}{n_2}\right)}$ and $h_1(\tilde{n}_1) = \frac{\tilde{n}_1}{g_1}$.

We define $C_2(h,n_1) = \frac{1}{r_3}\tilde{c}\left(n_1, g_1\left(\frac{1}{h} - \frac{g_2}{n_1}\right)^{-1}\right)$ and $C_1(h) = \frac{1}{r_3}\tilde{c}\left(g_1 h, 0\right)$

so we can rewrite the second period optimand for the adjustment pitfalls game as

$$U(h) - C_2(h,n_1)$$

with an optimum h^*

We rewrite the optimand for the comparison game as

$$U(h) - C_1(h)$$

With an optimum $\tilde{h^*}$

Note that it is possible to have $\tilde{h^*}$ which is not attainable in the adjustment pitfalls game, in which case we have the result trivially. Assume that $\tilde{h^*}$ is feasibly attainable in the adjustment pitfalls game from here on.

The fact that each selected precision is incentive compatible gives us

$$U(h^*) - C_2(h^*,n_1) \geq U(\tilde{h^*}) - C_2(\tilde{h^*},n_1)$$

$$U(\tilde{h^*}) - C_1(\tilde{h^*}) \geq U(h^*) - C_1(h^*)$$

which sum to

$$-C_2(h^*,n_1) - C_1(\tilde{h^*}) \geq -C_2(\tilde{h^*},n_1) - C_1(h^*)$$

which we can rearrange to get

$$C_2(\tilde{h}^*,n_1) - C_1(\tilde{h}^*) \geq C_2(h^*,n_1) - C_1(h^*)$$

Substitute back in and rearrange to get

$$\tilde{c}\left(n_1, g_1\left(\frac{1}{\tilde{h}^*} - \frac{g_2}{n_1}\right)^{-1}\right) - \tilde{c}\left(n_1, g_1\left(\frac{1}{h^*} - \frac{g_2}{n_1}\right)^{-1}\right) \geq \tilde{c}\left(g_1\tilde{h}^*, 0\right) - \tilde{c}\left(g_1 h^*, 0\right)$$

By symmetry we have

$$\tilde{c}\left(g_1\left(\frac{1}{\tilde{h}^*} - \frac{g_2}{n_1}\right)^{-1}, n_1\right) - \tilde{c}\left(g_1\left(\frac{1}{h^*} - \frac{g_2}{n_1}\right)^{-1}, n_1\right) \geq \tilde{c}\left(g_1\tilde{h}^*, 0\right) - \tilde{c}\left(g_1 h^*, 0\right)$$

or

$$\tilde{c}\left(g_1\left(\frac{1}{h^*} - \frac{g_2}{n_1}\right)^{-1}, n_1\right) - \tilde{c}\left(g_1\left(\frac{1}{\tilde{h}^*} - \frac{g_2}{n_1}\right)^{-1}, n_1\right) \leq \tilde{c}\left(g_1 h^*, 0\right) - \tilde{c}\left(g_1\tilde{h}^*, 0\right)$$

which is a necessary condition for the optimality of h^* and $\tilde{h}^*$. Now for the purpose of contradiction assume $h^* > \tilde{h}^*$. By increasing differences we know

$$\tilde{c}\left(g_1 h^*, n_1\right) - \tilde{c}\left(g_1\tilde{h}^*, n_1\right) \geq \tilde{c}\left(g_1 h^*, 0\right) - \tilde{c}\left(g_1\tilde{h}^*, 0\right)$$

Note that $g_1\left(\frac{1}{h} - \frac{g_2}{n_1}\right)^{-1}$ is a strictly increasing monotone function which is greater than or equal to $g_1 h \forall h \geq 0$ and $\tilde{c}\left(\bullet\right)$ is strictly increasing and convex in each element. Therefore,

$$\tilde{c}\left(g_1\left(\frac{1}{h^*} - \frac{g_2}{n_1}\right)^{-1}, n_1\right) - \tilde{c}\left(g_1\left(\frac{1}{\tilde{h}^*} - \frac{g_2}{n_1}\right)^{-1}, n_1\right) > \tilde{c}\left(g_1 h^*, n_1\right) - \tilde{c}\left(g_1\tilde{h}^*, n_1\right)$$

This implies

$$\tilde{c}\left(g_1\left(\frac{1}{h^*} - \frac{g_2}{n_1}\right)^{-1}, n_1\right) - \tilde{c}\left(g_1\left(\frac{1}{\tilde{h}^*} - \frac{g_2}{n_1}\right)^{-1}, n_1\right) > \tilde{c}\left(g_1 h^*, 0\right) - \tilde{c}\left(g_1\tilde{h}^*, 0\right)$$

which contradicts optimality of h^* and $\tilde{h}^*$.

$\square$

This result seems particularly counter intuitive in situations where r_2 is much larger than r_3. One would expect that giving someone a large reward to memorize a fact and then a small reward to repeat it later would produce a better result than just offering the small reward for one recitation. Normally, the high memorization effort in the first period would be expected to produce good memory signals in later periods even if incentives did not remain high. Note that the non-adjustable framework does not make a firm prediction in this case, with the outcome depending on the relative impacts of effort and time.

As we can see, both full adjustability and non-adjustability present issues. Neither of these issues is likely to impact results earlier in the paper, however, as we do not consider any games where the expected value of previously learned information can change stochastically, and we do not include any problems where information value decreases.

D Manipulation Through Timing of Information Release

In this section we show that even when memory resources are allocated rationally, the timing of when information is released can be used to manipulate the behavior of even fully Bayesian but forgetful individuals. There are several examples from the real world where evidence indicates that there is the possibility of manipulating individuals through the timing of information release or that such manipulation is already occurring. Amacher and Boyes (1978) find evidence that legislator behavior changes over the course of a session, with legislative actions more closely aligning with the preferences of the constituency when an election is coming up soon. This pattern suggests that legislators are taking advantage of the tendency of voters to forget old information in order to boost their chance of reelection. Panagopoulos (2011) finds that voter mobilization campaigns are more effective closer to the date of the election, suggesting that an agency wanting to maximize voter turnout can take advantage of the voter's ability to remember new information more easily. Moving away from elections, Francis et al. (1992) find that companies often report worse than expected earnings outside of trading hours to minimize the impact on their stock price.

To explore this phenomenon, we embed our model in a two player manipulation. In our model there are two agents: a Sender and a Receiver. Both Sender and Receiver have an identical prior π over states of the world. Assume $\mathcal{T} = T$, so the Receiver selects an action a_T from an action space A at time T leading to the final payoffs: $u(a_T, \theta)$ for the Sender and $v(a_T, \theta)$ for the Receiver.

This model borrows some structure and notation from the canonical Bayesian Persuasion model of Kamenica and Gentzkow (2011), and is simlar in principal, but there are key differences. In Kamenica and Gentzkow's (2001) model, the Sender has full control over the signal structure which informs the Receiver about the state of the world. In our model, the Sender only has one dimension of manipulation. They control when information arrives which in turn controls signal structure variance. Furthermore, in our model, the Receiver does have some control over his own memory signal structure through the exertion of effort. This greater restriction on the Sender and less stringent restriction on the Receiver makes manipulation a more difficult task in our setting, but it is still very possible as we shall show.

Due to the similarity with existing Bayesian Persuasion frameworks, we focus on dynamic statements that are unique to our setting.

At the beginning of the game, the Sender determines a time $t < T$ when the Receiver will see it. At the time t, the Receiver gets the exogenous signal s_t and selects how much effort n_t to devote to storing that signal.

The timing of the game is as follows:

1. The Sender selects when the signal will arrive for the Receiver, $t \in \{1, 2, ..., T-1\}$.
2. At time t, the Receiver selects his memory effort n_t and pays the cost $c(n_t)$.
3. The Receiver receives his exogenous signal s_t and remembers it.
4. At time T, the Receiver receives his memory signal m_T^t.
5. The Receiver updates his prior and selects an action, a_T.
6. The Sender gets a payoff of $v(a_T, \theta)$ and the Receiver gets a payoff of $u(a_T, \theta)$.

The Sender's strategy is a distribution over time period. The Receiver's strategy has two parts. The first is a mapping from the timing of information arrival to memory effort. The second is a mapping from the tuple of memory signals and the timing of information arrival to actions. In equilibrium, the Receiver's actions must be optimal given his beliefs. For simplicity we assume that the sender cannot change timing of information release in response to signal realizations, although this assumption actually is not important for the specific results we present.

First we consider how the alignment of preferences between the Sender and the Receiver influences equilibrium.

Corollary 6. *1. If $c(n_t)$ is convex and $v(\bullet)$ is an increasing affine transformation of $u(\bullet)$, there exists an uninformative timing equilibrium where the Sender will release the information at time $T-1$.*

2. *If $c(n_t)$ is convex and $v(\bullet)$ is a decreasing affine transformation of $u(\bullet)$, there exists an uninformative timing equilibrium where the Sender will release the information at time* 1.

Proof. This result is largely trivial. Proposition then 1 guarantees that the Receiver will have a Blackwell less informative memory signal structure in period T if the exogenous signal arrives earlier. This means that for any realization of the exogenous signal, earlier information release leads to the memory signal being less Blackwell informative about the state and therefore to a weakly lower $E(u)$. If u and v are related by an affine transformation, then $E(u)$ and $E(v)$ are related by the same affine transformation.

□

This result tells us that if incentives are perfectly aligned, the Sender wants to send the information late in the game (near time T) in order to minimize forgetting by the Receiver.[27] If incentives are perfectly misaligned, she wants to send the information early to encourage maximum forgetfulness and reduce Receiver performance.[28]

It is also interesting to consider how the convexity or concavity of Sender preferences can influence the timing of information release. Before we discuss such a result, however, a definition is needed.

Definition 5. We say a memory signal structure is *straightforward* if it is optimal after observing any memory signal m_T^t to take the action $a_T = m_T^t$.

This definition is identical to the one used in Kamenica and Gentzkow (2011). It allows us to avoid considering complicated maps between the signal space and the action space.

Using this definition, we present the following result.

Corollary 7. *Assume $c(n_t)$ is convex. If memory signals are straightforward for all s_t, n_t and $T - t$, then*

1. *If v is convex in $a_T \forall \theta$, there exists an uninformative timing equilibrium where the Sender will release information at time* 1.

[27]This assumes that the Receiver is perfectly rational other than having imperfect memory. This may not always be the case, as discussed by Lipnowski and Mathevet (2018).

[28]Note that this result does not necessarily hold with non-affine monotone transformations. For example, consider a case where the utility of the Sender is an increasing transformation of the Receiver's utility. If the Sender is more risk averse than the Receiver, there are situations where the Sender might send information early to reduce the informativeness of the Receiver's memory signal which would encourage him to take a safe option.

2. *If v is concave in $a_T \forall \theta$, there exists an uninformative timing equilibrium where the Sender will release information at time T.*

Proof. The combination of the assumed memory signal structure and straightforward signals guarantees that the average action is fixed for a given exogenous signal. From the Proof of Corollary 6 we know that increasing $T-t$ will essentially introduce a garbling to the memory signal. Therefore conditional on the exogenous signal, the distribution of actions given release time t will always be a mean-preserving spread of the conditional distribution of actions for release time t' if $t < t'$. By Jensen's Inequality, Senders with concave utility will prefer less variance and Senders with convex utility will prefer more. □

Straightforwardness and form of memory signals together fix the mean action of the Receiver based on the exogenous signal. Given that, this Proposition is a natural consequence of Jensen's Inequality. When the Sender has convex utility, she wants to introduce more noise into the Receiver's actions. Since signals are straightforward, this means introducing more noise into the signal. She can introduce more uncertainty and noise by releasing information earlier and giving it more time to decay. When the Sender has concave utility, however, she wants to minimize uncertainty, so she releases information later to avoid decay.

The requirement that the memory signal structure always be straightforward is somewhat restrictive when assuming normal signals and generally requires that the prior contain no information or that the signal be very informative. However, the results underlying the manipulation discussion don't require normal signals. See Appendix **??** for more details.

Note that the proofs for Corollary 6 and Corollary 7 are based on weak optimality of information release timing. There can be other uninformative timing equilibria due to indifferences and trivial choices. For example, consider the case in which the Receiver only has a single action. The Sender can send information at any time in equilibrium without influencing behavior.