

Bounded Memory and Biases in Information Processing*

Andrea Wilson[†]

University of Chicago

May 04, 2004

Abstract

This paper explores the connection between bounded memory and biases in information processing. An infinitely-lived decision maker receives a sequence of signals which provide partial information about the true “state of the world”; when this process terminates, which happens with probability η after each signal, he makes a decision and earns a payoff which depends on the true state. The agent has a bounded memory: this consists of a finite set of available memory states, and a memory rule which specifies the transition between states as new information is received. We show that the optimal memory rule may perform very poorly in the short run, and can explain several biases that psychologists have observed. In particular, we show that when η is small, the optimal memory process involves ignoring information with probability close to 1 in some memory states. As a result, the agent appears to display a confirmatory bias (tendency to ignore information which does not support his first impressions), and an overconfidence/underconfidence bias (tendency to infer too much from ambiguous information, too little from precise information).

*I am indebted to Wolfgang Pesendorfer for extremely valuable guidance and advice throughout this project. I also thank Faruk Gul for his help in developing some ideas in the paper, and Dilip Abreu, Ariel Rubinstein, and many seminar participants for useful comments.

[†]Correspondence address: University of Chicago, Economics Department; Chicago, IL 60637. Telephone number: (773) 702-8251. E-mail: awilson@uchicago.edu.

1 Introduction

Psychologists have observed many systematic biases in the way that individuals update their beliefs as new information is received. Many studies have suggested a first impressions matter bias: exchangeable signals are processed in a way that puts too much weight on the initial signals. In particular, people tend to pay too much attention to information which supports their initial hypotheses, while largely disregarding (or even misinterpreting) information which opposes these hypotheses.¹ As they become more convinced that their beliefs are correct, the problem becomes even more severe: many individuals seem to simply ignore all information once they reach a “confidence threshold”.

A related phenomenon is belief polarization. Several experiments have taken two individuals with opposing initial beliefs, then given them exactly the same sequence of information. In many cases, *both* individuals became even more convinced of their initial position. Obviously this is in contrast to Bayes’ rule, which says that the prior should not affect the way in which the new information is interpreted.²

A third bias is overconfidence/underconfidence: belief adjustments tend to be more extreme than those of a Bayesian after a relatively uninformative sequence of signals (overconfidence), but too conservative after a highly informative sequence of information.³

Several recent papers in behavioral economics have focused on identifying some of these biases, and exploring their implications for the standard economic models; see Rabin (1998) for a comprehensive survey. Mullainathan (1998) made a potential connection between memory and biased information processing, in a model which makes several explicit (psychology-based) assumptions on the memory process. In particular, he assumes that the agent’s ability to recall a past event depends on how similar it is to the current environment, how similar it is to a randomly drawn event, and how often the event has been recalled in the past.

The goal of this paper is to develop a model of memory based on more primitive assumptions, and to demonstrate how an *optimal* finite memory can explain many perceived

¹Rabin and Schrag (1999) model this behavior using a modified version of Bayes’ rule, which explicitly assumes that new evidence is weighted according to the prior.

²See Rabin (1999) for more examples and references.

³See Kahneman, Slovic, and Tversky (1982, pp.287-387) for many related experiments, and Rabin (1999) for a summary of the results.

biases.

The paper considers an infinitely lived decision maker, who receives a sequence of signals. Each signal provides partial information about the “state of the world”, which may be either H or L . At the start of each period, there is a probability η that this process will terminate, in which case the decision-maker must make a decision; the correct decision depends on the true state of the world. The paper will focus on the case where η is close to zero; this approximates a situation in which the agent expects to receive a long sequence of signals before making his decision, but is not exactly sure when the process will end.

A standard Bayesian decision-maker would be able to base his decision on the entire sequence of signals. In contrast, we study a decision-maker with a bounded memory: he is restricted to a finite set $\mathcal{N} = \{1, \dots, N\}$ of available memory states. A memory process on \mathcal{N} consists of an initial distribution g^0 (which tells the agent where to start), a transition rule σ (which tells him which state to go to when he receives new information, as a function of the current state), and an action rule d (which tells him which decision to make when the process terminates, as a function of the current memory state). This model of memory bears some resemblance to several others proposed in the literature. Dow (1991) studies an agent who searches sequentially for the lowest price, but can only remember each price as being in one of a finite number of categories. Lipman (1995) and Rubinstein (1998) also discuss related models.⁴

The basic idea of the model is that the decision-maker cannot recall all of the information that he receives, and cannot perform (and recall) an exact Bayesian update after each new signal. The finite-state memory system just describes the heuristic that he uses to process and store information. For example, suppose that there are two possible signals in the world: a low signal, and a high signal. Then each state might correspond to a set of signals that the decision-maker can recall (e.g., i high signals, $N - i$ low signals in state i).⁵ The N memory states might also correspond to a set of N different beliefs that the decision-maker can have; in this case, the transition rule σ describes how the decision-maker updates his beliefs after

⁴The memory process also resembles standard finite automata models of decision-making. See, for instance, Piccione and Rubinstein (1993), and Rubinstein (1998).

⁵Note that if we interpret each possible sequence of signal realizations as a memory state, then a standard Bayesian agent must have an infinite number of available states.

new information.

It is assumed throughout the paper that in each state i , the decision-maker's beliefs are correct (Bayesian) given the set of histories which could have occurred, conditional on i and the rule σ . However, since he must have the same beliefs every time he reaches state i , the change in his beliefs from one period to the next will typically appear non-Bayesian. Moreover, since he must behave the same way every time he reaches state i , the memory rule which is optimal in the long run may perform poorly in the short run. This accounts for many of the "biases" described above, which can simply result from the optimal long-run behavior when memory is finite.

Section 2 introduces the basic model, and defines the optimality concept that will be used. In general decision problems with imperfect recall, the behavior which is optimal ex ante may not be incentive compatible; therefore, there is some ambiguity over what is the "correct" notion of optimality. However, Theorem 1 states that for the problem considered here, there is no conflict between the two solution concepts: every ex-ante optimal strategy is also incentive compatible.

Section 3 develops some results on the optimal finite memory. Theorem 2 provides an upper bound on the decision-maker's expected payoff, achievable only in the limit as $\eta \rightarrow 0$, which depends on the number of memory states N and the likelihood ratios of the two most informative signals. One implication of the result is that for any finite N , the decision-maker chooses the wrong state with strictly positive probability - even after an infinite sequence of signals. However, the decision-maker's payoff does converge to that of a Bayesian agent as the number of memory states goes to infinity. Theorem 3 contains the main result of the paper, a precise characterization of the optimal memory when η is close to zero. First, the states $1, \dots, N$ are ordered according to the beliefs induced by (g^0, σ) , such that higher memory states assign higher probabilities to the event that H is the true state of the world; in particular, the decision-maker is most strongly convinced that state H is true when in memory state N , and that L is true when in memory state 1. We show that if $N \geq 3$ and η is close to zero, then optimality requires leaving states $1, N$ with probability close to zero. In other words, once the decision-maker reaches one of the two memory states with the most extreme beliefs, he optimally ignores all information with probability near 1.

An intuition for this result is as follows: the DM prefers to make all of his decisions in the two extreme states, where he has the best information (and hence obtains the highest expected payoff). This creates an incentive to avoid switching out of the extreme states. However, it is not optimal to actually get stuck in these states: ignoring information is costly, making each memory state less informative. In fact, if states $1, N$ were absorbing, the decision-maker’s payoff would fall to a level attainable with half as many memory states. The optimal solution therefore requires randomization: the decision-maker leaves states $1, N$ with a positive probability that goes to zero as $\eta \rightarrow 0$, but at a much slower rate. For small η , this strategy implies that the decision-maker is almost always in one of the two extreme states, and that over finite signal sequences it may appear as though he is completely unresponsive to information. However, the probability of eventually leaving the extreme states goes to 1 as $\eta \rightarrow 0$.

We then relate the optimal finite-state memory to some more descriptive, psychology-based models. A corollary to Theorem 3 provides a simple expression for the decision-maker’s beliefs in each memory state i . For the special case with two symmetric signals, high and low, the decision-maker’s beliefs in each state i correspond to those of a Bayesian who recalls a sequence of $(N - 1)$ signals; when he observes a high signal, his beliefs change as though he both added the new high signal to memory, and forgot one of the previously stored low signals; and similarly when he observes a low signal, his beliefs change as though the low signal replaced an old high signal. The expression is similar for a general signal space, and implies that the optimal finite-state memory closely resembles psychologists’ “attention and memory models”.⁶ These models are based on the idea that memory is an optimal storage system with limited capacity: people can control which facts to remember by paying attention to those that seem the most important, and knowledge which no longer seems (as) useful will be replaced by new information.

Section 4 shows how the result in Theorem 3 can create a first impressions matter/confirmatory bias. Consider a simple example with two symmetric signals, high and low, and three memory states, $1, 2, 3$. Suppose that the individual begins in state 2: then the optimal transition rule is to move to state 1 with probability 1 after a low signal, and to state 3 with prob-

⁶See Cowan (1995).

ability 1 after a high signal. Once the decision-maker reaches one of these extreme states, he may stay for a long time - leaving with probability γ each period, where γ is close to zero. This means that in the short run, it will appear as though one signal was enough to make the decision-maker ignore all opposing information. More generally, Theorem 4 shows that the order in which information is received can matter significantly. Over short signal sequences, first impressions matter: relative to a Bayesian, the decision-maker puts too much weight on the early signals that he receives. In the long run, last impressions matter: the decision-maker puts too much weight on the most recently received information.

Section 4 also discusses an experiment (conducted by Lord, Ross, and Lepper in 1979) which demonstrated that the same sequence of information can polarize beliefs. In the experiment, two groups of people were given a sequence of studies on the merits of capital punishment as a deterrent to crime. Group 1 individuals initially favored capital punishment, while Group 2 individuals were opposed. After seeing exactly the same information, Group 1 individuals became even more strongly convinced that capital punishment deters crime, while Group 2 individuals became even more strongly opposed: that is, both groups seemed to view the evidence as supporting their initial beliefs. We show how this can result from finite memory, and argue in Theorem 5 that even in the long run, there is a strictly positive probability that two individuals will disagree after observing the same information.

Section 5 demonstrates that the decision-maker will display an over/underconfidence bias: his beliefs are typically too extreme (compared to a Bayesian) after short or uninformative signal sequences, and too conservative after highly informative sequences. Moreover, if the most extreme signal in favor of state H is more informative than the most extreme signal in favor of state L , then the decision-maker's beliefs will relatively overrespond to the less informative low signal, and underrespond to the more informative high signal.

Section 6 concludes; all proofs are in the appendix.

2 The Model

There is a single infinitely-lived decision-maker (DM), who must eventually make a decision. The optimal decision depends on the true state of the world, $S \in \{H, L\}$, which remains

fixed throughout the problem but is unknown to the DM. The ex ante probability of state H is π .

The DM has access to an information process which terminates each period with probability η , and otherwise generates a signal $k \in \{1, 2, \dots, K\} \equiv \mathcal{K}$. The signals are informative about the state of the world: conditional on the process not terminating and on true state $S \in \{H, L\}$, the probability of receiving signal k is μ_k^S in every period. Without loss of generality, assume that μ_k^H / μ_k^L is increasing in k ; thus higher signals provide stronger evidence in favor of state H . The DM knows both the termination probability η and the conditional probabilities μ_k^S .

At the start of each period, the DM first learns whether or not the information process has terminated. If it has, then he makes a decision $d \in \{H, L\}$ and the problem ends; if it has not, then he receives a signal and continues to the next period. He does not discount the future, and earns a payoff of 1 if his decision matches the state S , 0 otherwise.

The DM cannot keep track of all of the information that he receives. He has a finite set of $N \geq 2$ available memory states, $\mathcal{N} \equiv \{1, 2, \dots, N\}$, together with an initial distribution $g^0 \in \Delta(\mathcal{N})$ (which determines the probability g_i^0 of starting in state i), and a rule $\sigma : \mathcal{N} \times \mathcal{K} \rightarrow \Delta(\mathcal{N})$ which specifies the transition between states after new information is received. Each memory state therefore represents a set of information histories which are indistinguishable to the DM.⁷ At the start of each period, he knows only his current memory state i , and any information about the history which can be inferred from being in memory state i . Then, after receiving a signal $k \in \mathcal{K}$, he updates his memory by choosing one of the states $j \in \mathcal{N}$. For all $i, j \in \mathcal{N}$ and $k \in \mathcal{K}$, let $\sigma_{i,j}^k$ denote the probability that the DM moves from memory state i to j after observing signal k . The assumption that the DM updates his memory after every signal is without loss of generality: for example, if he could costlessly record two signals at a time, then the results in this paper would continue to go through if the signal space \mathcal{K} were enlarged to include all pairs of signals.

Note that σ is stationary: that is, the DM must follow the same transition rule every time he is in memory state i . This stationarity requirement follows from the definition of a

⁷If the rule σ were deterministic, then the memory states would partition the set of all possible information histories into N equivalence classes. However, the optimal rule σ will require randomization: this implies that there is some intersection between the sets of histories represented by different memory states.

memory state as representing *all* of the information that the DM can recall. In particular, any information about the time period t must be encoded into the memory states; and since N is finite, this implies that (at least in the long run) the DM cannot keep track of time.⁸

In addition to choosing a memory rule (g^0, σ) , the DM must choose a decision rule. Let d_i denote the probability that the DM decides H , conditional on a decision in memory state i .

The DM can associate a belief about $S \in \{H, L\}$ with each memory state i , by using his rule (g^0, σ) to infer the set of signal sequences which could have occurred given that he is in memory state i , and the relative likelihood of each sequence conditional on S . More precisely, let k_t, i_t be random variables describing (respectively) the DM's signal in period t , memory state at the end of period t . Define $X_t(i) \equiv \{(S, i_0, (k_t, i_t)_{t=1}^T) \mid i_T = i\}$ as the set of t -period histories ending in memory state i , and define $X(i) \equiv \cup_{t=0}^{\infty} X_t(i)$ as the DM's information set in memory state i . Then the probability that the true state of the world is H , conditional on information set $X(i)$, is given by

$$\pi_i \equiv \frac{\pi P(X(i)|H)}{\pi P(X(i)|H) + (1 - \pi)P(X(i)|L)}$$

where for any history z , $P(z|S)$ is the probability induced by (g^0, σ) of history z , conditional on S ; and $P(X(i)|S) \equiv \sum_{z \in X(i)} P(z|S)$.

It is assumed throughout the paper that the DM's beliefs are consistent with his strategy, in the sense that his beliefs about the world in memory state i are Bayesian given his information set:

Assumption 1: In any memory state i with $P(X(i)|S) > 0$, the DM believes that the probability of state H is π_i .

Note that this consistency requirement implies that a DM with finite memory cannot perform (and recall) an exact Bayesian update after each new signal. At the start of each

⁸It is possible to keep track of time in the short run, though this is not optimal for small η . For example, suppose $N = 4$, and consider the rule which starts in state 1, moves to state 2 after the first signal, moves to state 3 after the second signal, moves to state 4 after the third signal, and never leaves state 4. A DM with this memory rule can infer that he's received 0 signals if he's in state 1, 1 signal if he's in state 2, 2 signals if he's in state 3, and at least 3 signals in state 4. However, since σ does not distinguish between signals in \mathcal{K} , he does not update his beliefs about the probability of state H beyond the prior, π .

period he knows only his current memory state i , and any information that can be inferred from being in memory state i ; thus, his posterior about the state of the world is π_i . After receiving a new signal $k \in \mathcal{K}$, he uses the rule σ to choose a (possibly different) memory state j , where he adopts the beliefs π_j in the subsequent period. Hence, a DM with N available memory states can have at most N different posteriors; the transition rule σ both determines these posteriors, and explains how beliefs are updated after each new signal.

Throughout the paper, it will be assumed higher memory states correspond to (weakly) higher posteriors π_i :

Assumption 2: π_i is non-decreasing in i .

(This is without loss of generality: for any rule (g^0, σ, d) , it simply says that the states are (re)-labelled according to the beliefs induced by the rule).

2.1 Preliminaries

Define $\tau_{i,j}^S \equiv (1 - \eta) \cdot \sum_{k \in \mathcal{K}} \mu_k^S \sigma_{i,j}^k$ as the probability induced by σ of a transition from memory state i to j , conditional on $S \in \{H, L\}$; and define T^S to be the $N \times N$ matrix with (i, j) th entry $\tau_{i,j}^S$. Then the probability that the DM will be in memory state i at the start of period τ , conditional on $S \in \{H, L\}$, is given by the i th coordinate of the $(N \times 1)$ vector $g^{S,\tau} \equiv g^0 (T^S)^\tau$. Next, define

$$f^S \equiv \lim_{T \rightarrow \infty} \sum_{t=0}^T \eta g^{S,t} = \lim_{T \rightarrow \infty} \eta g^0 \sum_{t=0}^T (T^S)^t \quad (1)$$

Lemma 1 in the Appendix argues that the sum on the RHS of (1) converges to a unique limit, and that f^S is a probability distribution over \mathcal{N} .

This vector f^S can be used to calculate both the DM's expected payoff, and his beliefs about the world in memory state i .

To calculate his expected payoff, we need to figure out the probability that a decision will be made in memory state i , conditional on $S \in \{H, L\}$: this is given by f_i^S , the i th

coordinate of f^S . Hence, the DM's expected payoff if he follows the rule (g^0, σ, d) is

$$\Pi(g^0, \sigma, d) \equiv \sum_{i \in \mathcal{N}} [\pi f_i^H d_i + (1 - \pi) f_i^L (1 - d_i)] \quad (2)$$

To calculate his beliefs π_i , note that $g_i^{S,t}$ could equivalently be described as the probability (induced by (g^0, σ)) of a t -period history ending in memory state i , conditional on S . Then $f_i^S / \eta \equiv P(X(i)|S)$: so Assumption 1 implies that in any state i with $P(X(i)|S) > 0$, the DM's beliefs about the state of the world are given by

$$\pi_i = \frac{\pi f_i^H}{\pi f_i^H + (1 - \pi) f_i^L} \quad (3)$$

Then for any state i with $\pi f_i^H + (1 - \pi) f_i^L > 0$, Assumption 2 says that states are labelled such that f_i^H / f_i^L is non-decreasing in i .

The paper will characterize optimal rules, defined as strategies (g^0, σ, d) which maximizes the DM's ex ante expected payoff:

Definition 1: (g^0, σ, d) is optimal if $\Pi(g^0, \sigma, d) \geq \Pi(g^{0'}, \sigma', d')$ for all $(g^{0'}, \sigma', d')$.

Piccione and Rubinstein (1997a) and Lipman (1997) argue that in decision problems with imperfect recall, it is not clear that this is the “correct” definition of an optimal strategy. In particular, a DM with imperfect recall may have an incentive to revise the ex ante optimal strategy as he acquires more information.⁹ As one alternative, Piccione and Rubinstein propose the following incentive compatibility concept¹⁰ :

Definition 2: (g^0, σ, d) is incentive compatible if the following two conditions hold:

1. d_i maximizes $\pi_i d_i + (1 - \pi_i)(1 - d_i)$, $\forall i \in \mathcal{N}$
2. $\sigma_{i,j}^k > 0 \Rightarrow j \in \arg \max_{j' \in \mathcal{N}} [\pi_i \mu_k^H v_j^H + (1 - \pi_i) \mu_k^L v_j^L]$, $\forall i, j \in \mathcal{N}$ and $k \in \mathcal{K}$

⁹In decision problems with perfect recall, there is no ambiguity: the ex ante optimal strategy is also time consistent (there is no information set where the DM would like to change his strategy, assuming that he can also change his behavior at any point in the future) and incentive compatible (there is no information set where the DM would like to deviate, assuming that future behavior is fixed according to the original strategy).

¹⁰Piccione and Rubinstein refer to this solution concept as “modified multi-self consistency”.

where π_i is given by (3), and v_i^S is the DM's expected continuation payoff starting in memory state i , conditional on true state S and on the strategy (g^0, σ, d) :

$$v_i^H = \eta d_i + \sum_{j \in \mathcal{N}} \tau_{i,j}^H v_j^H$$

$$v_i^L = \eta(1 - d_i) + \sum_{j \in \mathcal{N}} \tau_{i,j}^L v_j^L$$

This says that the DM has no incentive for one-shot deviations from his strategy, given the payoffs and beliefs induced by his strategy. In particular, part (1) says that conditional on stopping in memory state i , the DM chooses the decision which maximizes his expected payoff at the beliefs π_i ; part (2) says that (for all $i \in \mathcal{N}, k \in \mathcal{K}$) if he receives signal k in memory state i , he will choose a memory state which maximizes his expected continuation payoff at the revised beliefs $\Pr\{H\} = \frac{\pi_i \mu_k^H}{\pi_i \mu_k^H + (1 - \pi_i) \mu_k^L}$.

Aumann, Hart, and Perry (1997a) argue that this is the correct solution concept, regardless of whether or not it is ex ante optimal. They argue that decision problems with imperfect recall should be viewed as one-person games, and the DM as a collection of independent agents: so every time the DM reaches a particular information set he controls only his current behavior, viewing his future behavior as fixed. Then, any pre-planning stage must select a strategy such that the DM has no incentive to deviate, assuming that he does not deviate in the future. Gilboa (1997) also argues in favor of the multi-selves approach, but against a pre-planning state; hence a solution is simply any equilibrium of the one-person, multiselves game.

For the problem considered here, the “correct” optimality concept is perhaps even less clear: it depends whether one believes that the memory is based on conscious decisions (in which case “optimality” would rule out strategies which violate incentive compatibility) or subconscious decisions (in which case, evolutionary arguments would favor ex ante optimal rules).

Fortunately for the problem considered here, there is no conflict between the two concepts:

Theorem 1: *If (g^0, σ, d) is an optimal memory rule, then it is also incentive compatible.*

Thus, finding an ex ante optimal rule is equivalent to finding the incentive compatible

rule with the highest expected payoff.

This result relies on the fact that the DM lives in a stationary environment: his beliefs about the number of periods played so far, and the number of periods remaining, do not change over time. It also relies on the fact that he does not discount the future. (In terms of finding the ex ante optimal rule, any impatience about the future can simply be incorporated into the termination probability - namely, by changing the termination probability to $\eta' = 1 - \delta(1 - \eta)$. However, Theorem 1 implies that the ex ante optimal rule for this problem would be incentive compatible iff the DM's beliefs were given by $\pi_i = \frac{\pi f_i^{H,\delta}}{\pi f_i^{H,\delta} + (1-\pi)f_i^{L,\delta}}$, where $\frac{f_i^{S,\delta}}{\eta} = \sum_{t=0}^{\infty} \delta^t g^{t,S}$; but these beliefs are incorrect when $\delta \neq 1$.)

The optimal rule is not time consistent in the usual sense - there are information sets (memory states) at which, in light of the new information, the DM would prefer to revise his entire strategy. This paper takes the view that in decision problems with imperfect recall, a time consistency requirement does not make sense. It would rule out the ex ante optimal strategy (and hence any evolutionary arguments), and also would implicitly assume that the DM will choose deviations which would only be profitable if they triggered a sequence of future deviations: but the DM will not recall the initial deviation, and hence cannot plan a sequence of future deviations without adding more memory states.¹¹

3 Optimal Memory Rules

This section provides some results on the optimal finite memory. Theorem 2 establishes existence of an optimal rule whenever $\eta \neq 0$, and provides an upper bound on the payoff $\Pi^*(\eta, N)$. Theorem 3 provides a precise characterization of the (generically unique) optimal rule when η is small. The paper focuses on optimal behavior when η is small - in other words, where the DM expects to receive a lot of information relative to the size of his memory. This is partly to obtain cleaner proofs and results, but is also the more interesting case: this is where the finite memory is very restrictive for the DM, generating a sharp divergence between his optimal behavior and that of a Bayesian.

¹¹This corresponds to the view of Battigali (1997).

Theorem 2: *An optimal memory process exists for all $\eta > 0$. If $\max\left\{\frac{\pi}{1-\pi}, \frac{1-\pi}{\pi}\right\} \geq \left(\frac{\mu_1^L \mu_K^H}{\mu_1^H \mu_K^L}\right)^{N-1}$, then the maximum payoff is $\max\{\pi, 1-\pi\}$. If $\max\left\{\frac{\pi}{1-\pi}, \frac{1-\pi}{\pi}\right\} < \left(\frac{\mu_1^L \mu_K^H}{\mu_1^H \mu_K^L}\right)^{N-1}$, then the maximum payoff $\Pi^*(\eta, N)$ is continuous and strictly decreasing in η , and is bounded above by*

$$\lim_{\eta \rightarrow 0} \Pi^*(\eta, N) \equiv \frac{1 - 2\sqrt{\pi(1-\pi)}r^*}{1 - (r^*)^2}, \text{ where } \frac{1}{r^*} \equiv \left(\frac{\mu_1^L \mu_K^H}{\mu_1^H \mu_K^L}\right)^{\frac{N-1}{2}}$$

This result says that an optimal rule exists provided that $\eta \neq 0$, and gives an upper bound on the maximized payoff $\Pi^*(\eta, N)$. The upper bound depends on the prior π , and on a parameter $\frac{1}{r^*} \equiv \left(\frac{\mu_1^L \mu_K^H}{\mu_1^H \mu_K^L}\right)^{\frac{N-1}{2}}$ which is an increasing function of the memory size N , the informativeness $\frac{\mu_1^L}{\mu_1^H}$ of signal 1, and the informativeness $\frac{\mu_K^H}{\mu_K^L}$ of signal K . The first part says that if this term $\frac{1}{r^*}$ is sufficiently small compared to the prior, then the DM can never store enough information to outweigh the prior: he will therefore decide H in all memory states if $\pi > \frac{1}{2}$, and L in all memory states if $\pi < \frac{1}{2}$. The second part says that if the prior is not too extreme, then the expected payoff is strictly increasing in the expected amount of information (as measured by $\frac{1}{\eta}$). The upper bound is therefore $\lim_{\eta \rightarrow 0} \Pi^*(\eta, N) = \frac{1-2\sqrt{\pi(1-\pi)}r^*}{1-(r^*)^2}$, which is the (maximized) probability of making a correct decision in the limit as the expected number of signals goes to infinity.

Observe that $\lim_{\eta \rightarrow 0} \Pi^*(\eta, N)$ is strictly increasing in $\frac{1}{r^*}$, but is bounded below 1 as long as r^* is positive. This means that as long as N is finite and no signal is perfectly informative, the DM will choose the wrong state S with strictly positive probability even in the long run. This is in contrast to a Bayesian, who would correctly identify the state with probability 1 after a sufficiently large number of signals. It is true, however, that $\lim_{\eta \rightarrow 0} \Pi^*(\eta, N)$ is strictly increasing in N , and converges to the Bayesian payoff 1 as $N \rightarrow \infty$. It is also easy to show that $\lim_{\eta \rightarrow 0} \Pi^*(\eta, N)$ is concave in N ; this implies that if the DM could choose the number of memory states at some cost, there would be a unique optimal choice of N for small η .

Cover and Hellman (1970) study the optimal design of a hypothesis-testing machine with finite memory.¹² Their definition of finite memory is identical to the notion considered in this paper. The information structure is also similar, except that the signal space is not

¹²I thank Ken Arrow, who drew my attention to this paper after noticing several similar results.

discrete, and there is no termination probability; therefore the objective is to maximize the probability of being correct after an infinite sequence of signals.

For their problem, no optimal solution exists. This agrees with Theorem 2, which establishes existence of an optimum only when $\eta \neq 0$; the failure of existence at the limit $\eta = 0$ will be explained following Theorem 3. Cover and Hellman instead analyze ε -optimal rules, and obtain an upper bound on the set of achievable payoffs which corresponds exactly to $\lim_{\eta \rightarrow 0} \Pi^*(\eta, N)$.¹³

3.1 Characterization of the Optimum for small η

The following theorem describes the optimal rule for small η , when $r^* < \max \left\{ \sqrt{\frac{\pi}{1-\pi}}, \sqrt{\frac{1-\pi}{\pi}} \right\}$. Denote this optimal rule by $(\widehat{g}^0, \widehat{\sigma}, \widehat{d})$, and the induced transition probabilities by $\widehat{\tau}_{i,j}^S$.

Theorem 3: *Let $r^* < \max \left\{ \sqrt{\frac{\pi}{1-\pi}}, \sqrt{\frac{1-\pi}{\pi}} \right\}$, and $N \geq 3$. There exists η^* such that whenever $\eta \in (0, \eta^*)$, the optimal rule $(\widehat{g}^{\eta,0}, \widehat{\sigma}^\eta, \widehat{d}^\eta)$ satisfies the following conditions:*

- i. For all $k \in K$, $\sigma_{i,j}^{k,\eta} = 0$ whenever $j \geq i + 2$ or $j \leq i - 2$.
- ii. For all $i \in N$, $\sigma_{i,i+1}^{k,\eta} > 0$ iff $k = K$, and $\sigma_{i,i-1}^{k,\eta} = 0$ iff $k = 1$.
- iii. The sequences $\sigma_{1,2}^{K,\eta}, \sigma_{N,N-1}^{1,\eta}$ satisfy $\lim_{\eta \rightarrow 0} \widehat{\sigma}_{1,2}^{K,\eta} = \lim_{\eta \rightarrow 0} \widehat{\sigma}_{N,N-1}^{1,\eta} = 0$, and $\lim_{\eta \rightarrow 0} \eta / \widehat{\sigma}_{1,2}^{K,\eta} = \lim_{\eta \rightarrow 0} \eta / \widehat{\sigma}_{N,N-1}^{1,\eta} = 0$.
- iv. For $i \in N \setminus \{1, N\}$, there exist $\Delta_1, \Delta_K > 0$ such that $\sigma_{i,i+1}^{K,\eta} \geq \Delta_K$, and $\sigma_{i,i-1}^{1,\eta} \geq \Delta_1$. Moreover, $\Delta_K = 1$ if $\frac{\mu_K^H \mu_K^L}{\mu_1^H \mu_1^L} \leq 1$, and $\Delta_1 = 1$ if $\frac{\mu_1^H \mu_1^L}{\mu_K^H \mu_K^L} \geq 1$.
- v. Define $i_0 \equiv \left\lceil \frac{N-1}{2} + \log \left(\sqrt{\frac{\pi}{1-\pi}} \frac{\mu_1^L}{\mu_1^H} \right) / \log \frac{\mu_1^L \mu_K^H}{\mu_1^H \mu_K^L} \right\rceil$. The optimal initial distribution g^0 chooses state i_0 with probability 1.
- vi. The optimal action rule sets $d_i = 0$ if $i < i^*$, and $d_i = 1$ if $i > i^*$, where $i^* \equiv \frac{N+1}{2} - \log \sqrt{\frac{\pi}{1-\pi}} / \log \frac{\mu_1^L}{\mu_1^H} \frac{\mu_K^H}{\mu_K^L}$.

¹³Their model considers a continuous signal space. For the payoff bound, $\frac{\mu_1^H}{\mu_1^L}$ is replaced by $\inf \frac{P^H(A)}{P^L(A)}$ and $\frac{\mu_K^H}{\mu_K^L}$ is replaced by $\sup \frac{P^H(A)}{P^L(A)}$, where P^H, P^L are the conditional probability measures and the inf, sup are taken over all measurable sets A .

Parts (i) and (ii) describe some basic features of the optimal transition rule. Part (i) says that if η is sufficiently small, then the optimal rule does not contain any jumps: the DM will never move up or down by more than one memory state. Part (ii) says that he will move to the next highest memory state with positive probability if and only if he observes signal K , the most extreme evidence in favor of state H ; and similarly he moves down by one state if and only if he observes signal 1, the most extreme evidence in favor of state L .

Parts (iii) and (iv) describe the likelihood that the DM will move to a different state after observing signal 1 or K . Part (iii) says that at the boundary memory states 1 and N , the optimal transition probabilities $\sigma_{1,2}^K$ and $\sigma_{N,N-1}^1$ go to zero as $\eta \rightarrow 0$, but at a slower rate. This means that if the DM is expecting a sufficiently large amount of information, then he spends most of his time in states 1 and N , remaining there for an arbitrarily long period of time once he arrives. To an outside observer, it will appear as though he is simply unresponsive to new information. However, it is important to note that the probability of leaving the extreme states is strictly positive for any $\eta > 0$; in fact, the probability of leaving before the information process terminates goes to 1 as $\eta \rightarrow 0$. If states 1 and N were absorbing, then the limit payoff would fall to the level attainable with half as many states. This discontinuity at zero is the reason why no optimal solution exists at $\eta = 0$: for any positive $\tau_{1,2}^S, \tau_{N,N-1}^S$, the DM could increase his payoff by reducing the transition probabilities; but the payoff drops dramatically once they actually reach zero.

Part (iv) says that the interior transition probabilities, $\sigma_{i,i+1}^K, \sigma_{i,i-1}^1$ for $i \notin \{1, N\}$, are bounded away from zero for all η . The result also implies that in the special case where the signals are symmetric, so $\mu_K^H = \mu_1^L$ and $\mu_K^L = \mu_1^H$, the DM moves up one state with probability 1 after signal K , and down one state with probability 1 after signal 1.

Finally, parts (v) and (vi) describe the optimal action rule and initial distribution.

The proof of Theorem 3 also determines the limiting beliefs $\lim_{\eta \rightarrow 0} \pi_i$ associated with an optimal strategy. Note that $\lim_{\eta \rightarrow 0} \pi_i$ is the long-run probability of state H conditional on memory state i , for a DM who believes that he has observed an infinite number of signals.

Corollary: *The DM's beliefs in state i satisfy*
$$\lim_{\eta \rightarrow 0} \frac{\pi_i}{1 - \pi_i} = \sqrt{\frac{\pi}{1 - \pi}} \left(\frac{\mu_1^L \mu_K^H}{\mu_1^H \mu_K^L} \right)^{i - \frac{N+1}{2}}.$$

This implies that the ex ante likelihood ratio $\frac{\pi}{1-\pi}$ does affect the long-run beliefs, but only by a factor of $\sqrt{\frac{\pi}{1-\pi}}$. Note that as long as the signals 1, K are not perfectly informative, the likelihood ratios $\frac{\pi_i}{1-\pi_i}$ are bounded away from both zero and infinity for all i . In particular, this means that although the DM ignores information with probability close to 1 in states 1 and N , it is not because he is certain that his beliefs are correct. This is one main difference between the finite-memory DM and a Bayesian agent, and will be interpreted more extensively in Section 4 as a confirmatory bias.

Also note that $\lim_{\eta \rightarrow 0} \frac{\pi_{i+1}}{1-\pi_{i+1}} \bigg/ \frac{\pi_i}{1-\pi_i} = \frac{\mu_K^H}{\mu_K^L} \bigg/ \frac{\mu_1^H}{\mu_1^L}$: this implies that if η is close to zero and the DM moves from memory state i to $i+1$ after observing a K -signal, his beliefs adjust as if he were a Bayesian agent who observed one additional K -signal, and one less 1-signal. Similarly, if he observes signal 1 and moves from state $i+1$ to i , his beliefs adjust as though he replaced one K -signal with the new 1-signal. This suggests that the optimal finite-state memory behaves like a “fact-based” memory with limited capacity, in which the DM stores $(N-1)$ signals (facts) at a time. If he receives new information when his memory is full, he can either ignore it, or replace a previously stored fact with the new signal. This relation is especially transparent in the symmetric 2-signal case with $\pi = \frac{1}{2}, \mu_K^H = \mu_1^L = \rho, \mu_K^L = \mu_1^H = 1 - \rho$, where the belief formula can be written as $\lim_{\eta \rightarrow 0} \frac{\pi_i}{1-\pi_i} = \left(\frac{\rho}{1-\rho}\right)^{i-1} \left(\frac{1-\rho}{\rho}\right)^{N-i}$: that is, the DM’s beliefs in state i are as if he is a Bayesian agent who recalls a sequence containing $(i-1)$ K -signals, and $(N-i)$ 1-signals.

This closely resembles standard psychological theories of attention and memory. These theories are based on the idea that memory is an optimal information storage/retrieval system with limited capacity. Individuals can (to some extent) control what they remember, by focusing their attention on what seems the most important. The main features of these models are that in order for a stimulus to enter short-term memory, individuals must pay attention to it - i.e., try to remember it. From here, the brain decides whether to store the information in the long-term memory. This depends on how useful it is; information which no longer seems important will eventually be replaced by more relevant knowledge.¹⁴

For an intuition behind Theorem 3, consider first part (iii). Note that the probability of error is the lowest in states 1 and N , where the DM has the best information; this

¹⁴See Cowan (1995) for a discussion of “attention and memory” models.

suggests that his payoff would be the highest if he could guarantee making all decisions in the extreme memory states. One way to achieve this, would be to simply never leave states $1, N$ once he arrives: of course, the problem with this rule is that it is very costly to ignore all information, and makes the states much less informative. This suggests that the optimal solution requires randomizing in the extreme states - switching out with a probability that is positive (to maintain informativeness), but below 1 (to reduce the likelihood of ending in a middle state). In the limit as $\eta \rightarrow 0$, the DM does not expect to make a decision for a long time: this makes it *almost* costless to ignore information for very long, but finite, periods of time. In particular, the rule described in part (iii) guarantees that as $\eta \rightarrow 0$, the DM is almost always in states $1, N$, yet they are almost as informative as if no information was ignored.

Next consider parts (i) and (ii). If the DM is almost always in states $1, N$, then the objective is to make these states as informative as possible. Lemma 5 in the appendix argues that this is achieved by ignoring all by the most extreme signals, 1 and K , and by only moving one state at a time after observing one of these signals. Intuitively, the extreme states will be the most informative if they can only be reached after the most informative signal sequences.

Part (iv) says that any randomizing in the interior memory states is due only to asymmetry between signals $1, K$. This is because it is still costly to ignore information, but it may be necessary to balance the relative amount of time spent in states $1, N$. For example, if signal 1 almost never occurs in either state of the world, then the DM may sometimes randomize after signal K to avoid spending all of his time in state N . Parts (v) and (vi) essentially just say that the DM chooses action H in state i iff he believes that state H is more likely, and begins in the state with the beliefs closest to the prior π .

Finally, note that the Theorem assumes $N \geq 3$. For $N = 2$ there is no bad middle state to avoid, so the DM does not avoid switching out of the “extreme” states: any randomizing will again arise only due to asymmetry between the signals.

As η increases, the optimal transition rules will obviously be less extreme than those described in Theorem 3. One difference is that the probabilities of leaving states $1, N$ increase: the DM knows that he probably has not received many signals, and therefore his beliefs are

not extreme enough to justify ignoring new information. He will also start paying attention to less extreme signals, and will move by more than one memory state after the most extreme signals; the only straightforward generalization (see Lemma 2 in the appendix) is that he will only move to higher memory states after signals which support state H (i.e., those with $\mu_k^H > \mu_k^L$), and to lower memory states after signals which support state L . Note also that the optimal rule in Theorem 3 implies that there are no transient states, which by definition implies that it is not optimal to use any of the states to keep track of time. For sufficiently high η this result will also fail; some memory states will only be reached once, and in fact the maximized payoff is not always monotonically increasing in N .¹⁵

Therefore, the biases described in the following sections are the most severe when η is close to zero; they still exist for larger η , but (continuously) become less severe. The intuition is fairly clear: if η is high, then the DM does not expect to receive many signals - and if he is not expecting to run out of memory, then there is no reason for any non-Bayesian behavior. He will therefore appear almost Bayesian in the short run, but will also obtain a very low payoff - as none of the memory states will be very informative.

4 Belief Perseverance and Confirmatory Bias

After forming sufficiently strong initial beliefs, people tend to pay too little attention to opposing evidence; they may simply ignore it, or even interpret it as supporting evidence. Experiments have suggested the following stylized facts:¹⁶

1. People tend to display a confirmatory bias: as they become more convinced of their initial hypotheses, it becomes more likely that they will disregard any information which contradicts these hypotheses.

¹⁵The range of values of η for which the Theorem holds depends significantly on \mathcal{K} . For instance, if there exists a signal j such that $\frac{\mu_j^H}{\mu_j^L}$ is very close to the most extreme likelihood ratio $\frac{\mu_K^H}{\mu_K^L}$, then part (ii) of the result fails very quickly: η needs to be very close to zero for the DM to optimally ignore signal j . On the other hand if \mathcal{K} consists of two symmetric signals, then most of the results do generalize as η increases; the only effect is an increase in the transition probabilities out of state N (which are proportional to $\sqrt{\eta}$, as can be seen from the proof of Theorem 3 (iv) in the appendix).

¹⁶This closely follows the description in Rabin (1999) and Rabin-Schrag (1999); see also Kahneman, Slovic, and Tversky (1982, pp.144-149) for a summary.

2. First impressions matter: exchangeable information is processed in a way that puts too much weight on early signals.
3. Providing the same evidence to people with different initial beliefs can move their beliefs even further apart.

The confirmatory bias described in Fact 1 follows directly from Theorem 3. Part (iii) of the theorem states that once the decision-maker reaches a “threshold of confidence” (memory state 1 or N), he ignores any opposing evidence with probability close to 1. Since (by part (iv)) he does not ignore evidence in the intermediate states $\mathcal{N} \setminus \{1, N\}$, this implies that an initial string of K -signals can make it appear, in the short run, as though he’s convinced that state H is true. Moreover, the closer he is to state N - that is, the more strongly his initial beliefs favor state H - the more likely it is that he will receive a sequence of information containing enough K -signals to take him to memory state N .

The next two results show how the biases described in Facts 2 and 3 can arise as a result of the optimal behavior described in Theorem 3. As in Section 2, let i_t, k_t be random variables describing (respectively) the DM’s memory state, signal in period t . Let $k^t \in \mathcal{K}^t$ denote a t -period sequence of signals; and for any sequence k^t , define $g^t(i|k^t) \equiv \Pr\{i_t = i \mid k^t\}$.

The first result considers T -period sequences of signals, which differ only in the order in which the signals are received. Fix positive integers T, τ . Define $[T, \tau]$ as the set of all T -period sequences which end in a block of τ consecutive K -signals, and $[\tau, T]$ as the set of all T -period sequences which begin with a block of τ consecutive K -signals. More precisely,

$$[T, \tau] \equiv \{k^T \in \mathcal{K}^T \mid k_t = K \text{ for } t = T - \tau + 1, \dots, T\}$$

$$[\tau, T] \equiv \{k^T \in \mathcal{K}^T \mid k_t = K \text{ for } t = 1, 2, \dots, \tau\}$$

For any sequence k^T and collection of signal sequences A , define $\Pr\{k^T|A\}$ as the probability of sequence k^T , conditional on A . Recall that π_i is the probability of H in memory state i . Therefore, the individual’s expected probability assessment to state H in period T ,

conditional on A , is

$$E(\pi_{i_T} | A) = \sum_{i \in \mathcal{N}} \pi_i \left(\sum_{k^T \in A} g^T(i | k^T) \Pr\{k^T | A\} \right)$$

Clearly, there is a 1-1 map from elements of $[T, \tau]$ to elements of $[\tau, T]$ which leaves unchanged the total number of each type of signal. Therefore, for a standard Bayesian decision-maker, $E(\pi_{i_T} | [T, \tau]) = E(\pi_{i_T} | [\tau, T])$. However, this is not true for the finite-memory DM:

Theorem 4: *Choose $\hat{\eta}$ small enough that Theorem 3 holds for $\eta < \hat{\eta}$, and fix an optimal strategy sequence σ^n . Then:*

- i. *For any $\eta < \hat{\eta}$, there exists T^* such that whenever $T > T^*$, $E(\pi_{i_T} | [T, \tau]) > E(\pi_{i_T} | [\tau, T])$.*
- ii. *For any $T \geq \frac{N+1}{2}$, there exists (η^*, τ^*) such that whenever $\eta < \eta^*$ and $\tau \geq \tau^*$, $E(\pi_{i_T} | [\tau, T]) > E(\pi_{i_T} | [T, \tau])$; moreover $\tau^* = 1$ if $\mu_K^H \mu_K^L = \mu_1^H \mu_1^L$.*

Part (i) of this result says roughly that last impressions matter in the long run: if the sequence length T is sufficiently large compared to η , then the block of K -signals has a larger effect on period T beliefs if it occurs at the end. This is true by the ergodic theorem: putting the K -signals at the start just changes the initial state, which has a negligible long-run effect; however putting the K -signals at the end will increase the memory state, since $\sigma_{i,i+1}^K > 0 \forall i$. In other words, the first impressions effect eventually wears off, while the last impression will always have some effect on beliefs.

Part (ii) says that first impressions matter in the short run: if η is sufficiently small compared to T (equivalently, if the sequence is sufficiently short compared to the expected amount of information $\frac{1}{\eta}$), then provided that τ is not too small, a block of τ consecutive K -signals has a larger effect on beliefs if it occurs at the beginning. This is a consequence of Theorem 3 (iii), which implies that for any fixed T , we can choose η small enough that the probability of leaving states 1, N within T periods is arbitrarily close to zero. The proof basically argues that putting the K -signals at the start (hence increasing the DM's initial state) makes it more likely that the DM will eventually reach state N and get stuck, and less

likely that he will eventually reach state 1 and get stuck. It then follows that if he does not randomize in the interior states, which follows from Theorem 3 (iv) when $\mu_1^H \mu_1^L = \mu_K^H \mu_K^L$, then changing the order of a sequence $(k^{T-\tau}, K^\tau)$ to $(K^\tau, k^{T-\tau})$ has zero probability of reducing the final memory state in the limit as $\eta \rightarrow 0$, but a strictly positive probability of increasing it.

This result matches the stylized facts. There is a lot of evidence, both experimental and from casual observation, suggesting that it is difficult to change a first impression in the short run. However, there is also substantial evidence that people tend to remember recent events the most vividly, while information which was learned a long time ago eventually fades from memory: this suggests that in the long run, there should be a bias towards last impressions.

The next result shows how the same sequence of information can polarize the beliefs of two different individuals, as described in Fact 3 above. An experiment by Lord, Ross, and Lepper (1979) provided an example of this bias. They asked a group of 151 students about their attitudes toward capital punishment, then selected 24 opponents, 24 proponents. The students were then given exactly the same sequence of studies on capital punishment, and asked again about their beliefs. Nearly all of the proponents became even more convinced that capital punishment deters crime, while the opponents became even more convinced that it does not. (Additionally, the graph on p.146 of Kahneman, Slovic, and Tversky (1982) shows that the overall change in beliefs for both types was the largest when they *first* received the supporting evidence - which is consistent with Theorem 4). The study concluded that “there is considerable evidence that people tend to interpret subsequent evidence so as to maintain their initial beliefs...Indeed, they may even come to regard the ambiguities and conceptual flaws in the data opposing their hypotheses as somehow suggestive of the fundamental correctness of those hypotheses. Thus, completely inconsistent or even random data - when “processed” in a suitably biased fashion - can maintain or even reinforce one’s preconceptions”.

As a simple illustration of how polarization can result from the behavior described in Theorem 3, suppose that \mathcal{K} consists of two symmetric signals, 1 and K . Identify state H with the state “capital punishment deters crime”, state L with “capital punishment does not

deter crime”; and interpret signal K as a piece of evidence in favor of capital punishment, signal 1 as an equally compelling piece of evidence against. Now suppose that two individuals, each with $N = 4$ memory states, observe the sequence $(K, 1, 1)$. Suppose further that the two individuals agree on the symmetry of signals 1, K , but have different beliefs prior to observing the sequence: specifically, let agent 2 start in state 3, and agent 1 in state 2. Then Theorem 3 implies that agent 2 will move to state 4 after the initial K -signal, then with high probability will ignore the subsequent 1-signals and remain in state 4; thus he becomes even more convinced that capital punishment deters crime. Agent 1, on the other hand, will end up in memory state 1 with probability 1 - where he is even more convinced that capital punishment does not deter crime. In other words, even though the two agents observe exactly the same sequence of information, and interpret each individual piece of evidence in exactly the same way, their beliefs end up moving in opposite directions.

More generally, polarization after a sequence occurs if and only if at least one of the agents reaches a partially absorbing state, 1 or N , and gets stuck for the duration of the sequence. Therefore, it is most likely to be observed in the short run, between two agents who are expecting to receive a large number of signals.

In the long run, the beliefs of two individuals will repeatedly polarize, then converge, then polarize again. Thus, it does not really make sense to refer to polarization after a specific sequence. Note, however, that the beliefs of two agents are *not* guaranteed to converge in the long run: the randomization in states 1, N implies that even after arbitrarily long sequences of signals, there is a strictly positive probability that two agents will reach different conclusions. In general, we can obtain the following result:

Theorem 5: *Choose $\hat{\eta}$ small enough that Theorem 3 holds for $\eta < \hat{\eta}$, and consider two agents with initial states $1 < i_0^1 < i_0^2 < N$ who observe the same signals every period.*

Then:

- i. *Any collection of T signals satisfying $\#1 > \#K \geq N - i_0^2$ and/or $\#K > \#1 \geq i_0^1 - 1$ can be ordered such that $\Pr\{i_T^1 < i_0^1 < i_0^2 < i_T^2\} > 0$.*
- ii. *The probability of long-run disagreement, $P^* \equiv \lim_{T \rightarrow \infty} \Pr\{i_T^1 = 1, i_T^2 = N\}$, satisfies $\lim_{\eta \rightarrow 0} P^* = 2\sqrt{\pi(1-\pi)}r^* \left(\frac{1 - \frac{1}{\sqrt{\pi(1-\pi)}}r^* + (r^*)^2}{(1-(r^*)^2)^2} \right)$, with $r^* \equiv \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H} \right)^{\frac{N-1}{2}}$.*

Part (i) is fairly trivial. Suppose that the first condition holds, $\#1 > \#K \geq N - i_0^2$: then the statement just says that if a sequence contains more 1-signals than K -signals, but enough K -signals to move agent 2 from state i_0^2 to N , then the sequence can be re-ordered such that the agents polarize with positive probability. To see this, consider a sequence which starts with $(N - i_0^2)$ K -signals: so Theorem 3 implies that at this point, agent 2 will be in state N , while agent 1 will be in state $i_0^1 + N - i_0^2 < N$.¹⁷ For any finite T , Theorem 3 (iii) implies that η can be chosen small enough that agent 2 will stay in state N for the remainder of the sequence with probability arbitrarily close to 1. Then since $\#1 > \#K$, and since (by Theorem 3 (ii)) all signals other than 1, K are ignored, the remaining signals can be ordered such that agent 1 never hits state N , and ends up below his initial state. If signals 1, K are symmetric, then the probability of polarization after this (re-ordered) sequence will in fact go to 1 as $\eta \rightarrow 0$.

Note that this result implies that if two agents observe the same signal generating process and follow the same memory rule, then the probability of polarization after t periods is strictly positive whenever $t \geq \min\{i_0^1 - 1, N - i_0^2\}$. Moreover, polarization becomes more likely as the agents' initial beliefs become more extreme (i.e., as their initial states i_0^1, i_0^2 move closer to the boundary states 1, N).

Part (ii) quantifies the long-run probability of disagreement, defined as the probability that agent 1 will eventually be in state 1, and agent 2 will eventually be in state N . The result assumes that the agents observe the same sequence of information, but does not specify the sequence; that is, the probability is calculated over all possible signal sequences. Note that the limiting expression is independent of the initial states; this means that the disagreement described could represent any of belief polarization (if $i_0^1 < i_0^2$), the opposite of polarization (if $i_0^1 > i_0^2$, so the agents switch positions), or belief divergence from the same prior. Finally, note that the probability of disagreement is increasing in r^* , hence decreasing in N and the informativeness of signals 1, K ; as $N \rightarrow \infty$, the probability of long-run disagreement goes to zero.

¹⁷This statement assumes that both agents move up one state with probability 1 after a K -signal, which is true provided $\mu_1^H \mu_1^L \geq \mu_K^H \mu_K^L$; if this inequality does not hold, then the statement holds with positive probability, which is sufficient for the argument.

5 Overconfidence/Underconfidence

People tend to be overconfident after receiving weak information, and underconfident after receiving strong information. Rabin (1996) states that “there is a mass of psychological research that finds that people are prone towards overconfidence in their judgements”. (A related phenomenon is the “law of small numbers” - people infer too much from too little evidence). A series of experiments in Kahneman, Slovic, and Tversky (1982) demonstrate that the bias works both ways: while people tend to be overconfident after relatively ambiguous information, their beliefs are typically too conservative after receiving highly diagnostic information.

The DM with finite memory displays two types of overconfidence/underconfidence. The first type, overconfidence/underconfidence after a single piece of information, relies on asymmetry between the signals and follows immediately from the Corollary to Theorem 3. Specifically, suppose that $\frac{\mu_K^H}{\mu_K^L} > \frac{\mu_1^L}{\mu_1^H}$ - so that signal K is more compelling evidence (in favor of state H) than signal 1 (in favor of state L). The corollary implies that if the DM observes a K -signal and moves up to the next state, then he will become $\left(\frac{\mu_K^H \mu_1^L}{\mu_K^L \mu_1^H}\right)$ times more convinced that the true state is H . On the other hand, if he observes a 1-signal and moves down by one state, then his beliefs will adjust in favor of state L by the same factor. In other words, even though he is aware that signal K is more informative than signal 1, his beliefs adjust by the same magnitude after both signals.

In the case where also $\frac{\mu_K^H \mu_K^L}{\mu_1^H \mu_1^L} < 1$, this is because the DM more frequently ignores signal 1.¹⁸ Thus, when he does move to a lower memory state after observing a 1-signal, his beliefs must relatively overadjust to reflect the fact that he has probably already seen and ignored some 1-signals. If $\frac{\mu_K^H \mu_K^L}{\mu_1^H \mu_1^L} > 1$, then K can only be the more informative signal if $\frac{1-\mu_1^L-\mu_K^L}{1-\mu_1^H-\mu_K^H} > 1$ - so that the signals $i \notin \{1, K\}$ which the DM ignores favor (on average) state L . In this case, he again overresponds to the less informative 1-signal to compensate for the fact that he likely has already observed (and ignored) some evidence in favor of state L .

The second type of overconfidence/underconfidence considers the change in the DM’s beliefs after observing a sequence of information. The experimental evidence suggests that

¹⁸This is somewhat implied by Theorem 3 (iv), and is explained in the intuition for Theorem 3.

people are typically underconfident (compared to a Bayesian) after long, informative signal sequences, and overconfident after sequences which contain very little information. Griffin and Tversky (1992) state that “Edwards and his colleagues, who used a sequential updating paradigm, argued that people are conservative in the sense that they do not extract enough information from sample data. On the other hand, Tversky and Kahneman (1971), who investigated the role of sample size in researchers’ confidence...concluded that people ...make radical inferences on the basis of small samples. In some updating experiments conducted by Edwards, people were exposed to large samples of data..This is the context in which we expect underconfidence or conservatism. The situations studied by Tversky and Kahneman, on the other hand, involve..fairly small samples. This is the context in which overconfidence is likely to prevail”.

To illustrate how this bias can result from finite memory, suppose that the DM is initially in memory state i_0 , with belief π_{i_0} . For any T -period sequence of signals k^T , let $i(k^T)$ denote the DM’s (realized) memory state after observing k^T , and define $\Delta(k^T)$ as the number of K -signals in k^T , less the number of 1-signals. Recall that in memory state $i(k^T)$, the DM believes that state H is $\left(\frac{\pi_{i(k^T)}}{1 - \pi_{i(k^T)}}\right)$ times more likely than state L : therefore if he started in state $i_0 < i(k^T)$, an outside observer would conclude that he responded *as if* the observed evidence was $\frac{\pi_{i(k^T)}}{1 - \pi_{i(k^T)}} / \frac{\pi_{i_0}}{1 - \pi_{i_0}}$ times more likely to occur in state H than state L . Similarly if $i_0 > i(k^T)$, then the DM’s beliefs adjust as if the observed evidence was $\frac{1 - \pi_{i(k^T)}}{\pi_{i(k^T)}} / \frac{1 - \pi_{i_0}}{\pi_{i_0}}$ times more likely to occur in state L . Let $\Pr(k^T|S)$ denote the true probability of sequence k^T occurring, conditional on state S . We say that the DM is overconfident if the change in his beliefs is more extreme than it should be, in *either* direction:

Definition: The DM is *overconfident after* k^T if

$$\max \left\{ \frac{\pi_{i(k^T)}}{1 - \pi_{i(k^T)}} / \frac{\pi_{i_0}}{1 - \pi_{i_0}}, \frac{1 - \pi_{i(k^T)}}{\pi_{i(k^T)}} / \frac{1 - \pi_{i_0}}{\pi_{i_0}} \right\} > \max \left\{ \frac{\Pr(k^T|H)}{\Pr(k^T|L)}, \frac{\Pr(k^T|L)}{\Pr(k^T|H)} \right\}$$

So for example, if the DM receives a sequence of information which is two times more likely to occur in state H than L , then according to this definition he is overconfident if he becomes more than twice as convinced that the true state is H , OR that the true state is L .

The cleanest overconfidence/underconfidence result obtains in the special case where \mathcal{K} consists of two symmetric signals - which was in fact the case considered in the experiments discussed by Griffin and Tversky (1992):

Theorem 6: *Suppose that $\mathcal{K} = \{1, K\}$, $\mu_1^L = \mu_K^H \equiv \rho$, and η is small enough that Theorem 3 holds. Then:*

- i. *After any sequence k^T with $T < \min\{i_0 - 1, N - i_0\}$ and $\Delta(k^T) \neq 0$, the DM is overconfident with probability 1.*
- ii. *For any $\varepsilon > 0$, there exists T^* such that in any period $T \geq T^*$, the DM is underconfident with probability greater than $1 - \varepsilon$; moreover T^* is decreasing in ρ .*

Part (i) of the result says that the DM will be overconfident with probability 1 after any sequence which is sufficiently short, but not completely uninformative. This is due to Theorem 3 (iv) and the Corollary, which imply (for the symmetric two-signal case) that the DM's beliefs overadjust to each individual signal: in particular, if η is close to zero, then a K -signal (1-signal) has almost the same effect on his beliefs as if he were a Bayesian who observed two K -signals (1-signals). Then as long as he does not reach a boundary memory state, and therefore correctly keeps track of the difference between the number of K -signals and 1-signals, his beliefs will adjust by almost twice as much as they should.

Part (ii) implies that in the long run, the DM is almost certain to be underconfident (that is, for sufficiently large T , it is almost certain that the DM will observe a sequence k^T to which his beliefs underrespond). The second part of statement says that as the signals become more informative (the probability of a correct signal, ρ , increases), the DM will become underconfident more quickly. The reason is simply that since the DM has a finite number of memory states, his probability assessments are bounded. Therefore in the long run, the correct inference will, with probability 1, be more extreme than he can accommodate.

19

¹⁹Theorem 6 in fact holds for all η . However, the magnitude of the overconfidence in part (i) decreases as η increases: the DM's overreaction to individual signals is the most dramatic when η is near zero.

This suggests that overconfidence is most likely after short sequences, while underconfidence is most likely after long sequences - particularly when the individual signals are very informative. This matches the experimental evidence: in Edwards (1965), the sample size was large and individuals were underconfident; in Kahneman and Tversky (1971), the sample size was small and individuals were overconfident; and Lichtenstein and Fischhoff investigated the effect of task difficulty on overconfidence, finding, as summarized by Griffen and Tversky (1992), that “easy items produced underconfidence through much of the confidence range,...and their “impossible” task (discriminating European from American handwriting, accuracy = 51%) showed dramatic overconfidence through the entire range”.²⁰

For the general signal space \mathcal{K} , the DM ignores some information; this makes overconfidence less likely.

6 Conclusion

This paper has demonstrated that decision-makers with finite memory will optimally display several biases in information processing. There are other papers (Dow (1991), Lipman (1995), Piccione-Rubinstein (1993, 1997a,b) which discuss similar models; however, these papers make only a limited connection between bounded memory and errors in judgement. There are also papers (Mullainathan (1998), Rabin-Schrag (1998)) which focus on making empirical predictions for errors in judgement; however, these papers also make very specific assumptions on the way that information is stored and processed. In this paper, we have attempted to combine the two approaches. We have started with a very fundamental notion of finite memory, and made several specific predictions about the biases that will arise. These predictions match the empirically observed behavior, and do not require any assumptions beyond optimality.

Most of the biases are driven by Theorem 3, which states that if the termination probability η is close to zero (so the decision is not expected to occur for a long time), then the decision-maker will ignore almost all information once he reaches one of his two extreme

²⁰Griffen and Tversky interpret “difficult tasks” as inference problems in which the information is of relatively poor quality (that is, when the signals provide very little evidence for one hypothesis over the other).

memory states. The intuition behind this result is quite straightforward. Since the decision-maker's expected payoff in the terminal period is the highest when he is in an extreme memory state (where he has the best information), there is an incentive to avoid leaving these states. In general, ignoring information makes each state less informative; however, when η is close to zero, the decision-maker is able to do this almost costlessly. By leaving the extreme states with a probability that is close to zero, but still much higher than η , he both avoids the middle states with high probability, and makes the extreme memory states almost as informative as possible.

Sections 4 and 5 discuss the empirical implications of Theorem 3. An implication of Theorem 3 is that when the decision-maker is as convinced as his finite memory will allow, he behaves as if he is virtually certain of the true state - even though his beliefs are far from certain. This is a confirmatory bias, and implies that the order in which information is received can matter significantly. More precisely, Theorem 4 shows that for relatively short sequences of information, there is a typically a first impressions bias: early signals have the largest effect on expected beliefs. For longer sequences of information, the first impressions effect wears off, and is dominated by a last impressions bias. This is consistent with the experimental literature: first impressions indeed matter for most people, but in the long run, it is the most recent events which are remembered the most vividly. Theorem 5 shows how the first impressions bias can lead to polarization: two agents with opposing initial beliefs may move even further apart after seeing exactly the same sequence of information. Theorem 6 shows that optimal behavior involves an overconfidence/underconfidence bias: after relatively short or uninformative sequences of signals, beliefs are typically more extreme than those of a Bayesian (overconfident); while after longer or more informative sequences, beliefs are typically too conservative.

All of these biases result from the fact that when the decision-maker is restricted to a finite number of memory states, and behaves the same way every time he reaches a particular memory state, the optimal long-run behavior may perform poorly in the short run. In particular, in order to maximize the probability of ultimately making a correct decision, the decision-maker ignores information with high probability once he reaches an extreme memory state, 1 or N ; in all other memory states, his beliefs overadjust to the information received.

Our interpretation is that people typically are not sure of exactly when information will cease to be useful, or exactly when decisions will be made. Therefore, the optimal memory should immediately interpret and store new information according to a rule which is optimal in the long run, but which may appear biased in the short run. This interpretation is supported by several psychological experiments, which have demonstrated that even information which is completely discredited will typically affect beliefs; this suggests that individuals do not simply memorize information as given, but rather incorporate it into an optimal long-run memory system.

Theorem 3 also shows that with an arbitrary set of K signals, the decision-maker will ignore all but the two signals which provide the most extreme evidence in favor of states H, L . One implication of this result is that a strong negative signal will have a much greater impact on the decision-maker's beliefs than a sequence of moderately negative signals. This suggests that an individual who must give out bad news would probably want to break the news into pieces - so as to induce a smaller negative reaction to the information - while an individual with good news would prefer to tell everything at once.

The results of this paper focused on the case when η is small. As η increases, the optimal behavior and associated biases become less extreme. The probability of leaving the extreme memory states increases, reducing the confirmatory bias. The decision-maker will start to pay attention to more moderate signals, and will move by more than one memory state after the most extreme signals. One implication is that the decision-maker's beliefs will not overreact so dramatically after each signal; this reduces the tendency toward overconfidence, while making underconfidence more likely as the beliefs in states 1, N become less extreme. The intuition is clear: if η is large, then the DM does not expect to receive enough signals to run out of memory; this means that there is no reason to react differently to information than a Bayesian.

Finally, the paper assumed that the number of memory states, N , is exogenous. A model which went one step further back would consider how many memory states should be allocated to each particular decision problem. Presumably, the optimal N would be increasing in the importance of decisions. Since the results predicted more extreme biases for small N , this suggests they will arise less frequently for more important decision problems.

A Appendix

A.1 Proof of Theorem 1

This is an almost exact adaptation of the argument of Proposition 3 in Piccione-Rubinstein (1997a). As in Section 2, define $X_t(i) \equiv \{(S, i_0, (k_t, i_t)_{t=1}^T) \mid i_T = i\}$ as the set of t -period histories ending in memory state i . Let ϕ denote the event that the game is terminated, and for $k \in \{\mathcal{K}, \phi\}$ and $i \in \mathcal{N}$, define $X_\tau(i, k) \equiv \{(z, k) \mid z \in X_\tau(i)\}$ as the set obtained by adding nature's action k to each history in $X_\tau(i)$. Finally, define $X(i, k) \equiv \bigcup_{\tau=0}^{\infty} X_\tau(i, k)$ as the DM's information set when he is in memory state i , then observes signal (or termination) k .

For any strategy (g^0, σ) , define $P^\sigma(z|z')$ as the probability of history z given z' , according to σ . Fix a strategy $\hat{\sigma}$, and choose $i^*, j^* \in \mathcal{N}$ and $k^* \in \mathcal{K}$ such that $\hat{\sigma}_{i^*, j^*}^{k^*} > 0$. For any history z , let $\delta(z)$ denote the number of occurrences of the sequence (i^*, k^*, j^*) in z , and define $C_\sigma^S(z) \equiv \frac{P^\sigma(z|S)}{(\sigma_{i^*, j^*}^{k^*})^{\delta(z)}}$. Note that $C^\sigma(z|S)$ is independent of $\sigma_{i^*, j^*}^{k^*}$; then we obtain

$$\left. \frac{d \sum_{z \in X(i, \phi)} P^\sigma(z|S, i_0)}{d\sigma_{i^*, j^*}^{k^*}} \right|_{\sigma = \hat{\sigma}} = \sum_{z \in X(i, \phi)} \delta(z) \frac{(\hat{\sigma}_{i^*, j^*}^{k^*})^{\delta(z)}}{\hat{\sigma}_{i^*, j^*}^{k^*}} C^{\hat{\sigma}}(z|S) = \sum_{z \in X(i, \phi)} \frac{\delta(z) P^{\hat{\sigma}}(z|S)}{\hat{\sigma}_{i^*, j^*}^{k^*}} \quad (\text{A1})$$

Note that for any history z ,

$$\delta(z) P^{\hat{\sigma}}(z|S) = \sum_{z' \in X(i^*, k^*)} P^{\hat{\sigma}}(z'|S) \hat{\sigma}_{i^*, j^*}^{k^*} P^{\hat{\sigma}}(z|S, (z', j^*))$$

(The summand on the RHS is zero if z' is not a subhistory of z , and otherwise is equal to $P^{\hat{\sigma}}(z'|S)$; $\delta(z)$ is the number of subhistories z' of z which end in (i^*, k^*, j^*)). Using this and summing over all histories in $X(i, \phi)$,

$$\begin{aligned} \sum_{z \in X(i, \phi)} \frac{\delta(z) P^{\hat{\sigma}}(z|S)}{\hat{\sigma}_{i^*, j^*}^{k^*}} &= \sum_{z \in X(i, \phi)} \sum_{z' \in X(i^*, k^*)} P^{\hat{\sigma}}(z'|S) P^{\hat{\sigma}}(z|S, (z', j^*)) \\ &= \sum_{z' \in X(i^*, k^*)} P^{\hat{\sigma}}(z'|S) \sum_{z \in X(i, \phi)} P^{\hat{\sigma}}(z|S, (z', j^*)) \end{aligned}$$

Next, note that $\sum_{z \in X(i, \phi)} P^{\hat{\sigma}}(z|S, (z', j^*))$ is just the probability of ending in state i , conditional on $(S, (z', j^*))$; moreover the stationarity of $\hat{\sigma}$ implies that this does not depend on z' . Then the above expression simplifies to

$$\sum_{z \in X(i, \phi)} \frac{\delta(z) P^{\hat{\sigma}}(z|S)}{\hat{\sigma}_{i^*, j^*}^{k^*}} = \left(\sum_{z' \in X(i^*, k^*)} P^{\hat{\sigma}}(z'|S) \right) \left(\sum_{z \in X(i, \phi)} P^{\hat{\sigma}}(z|S, j^*) \right) \quad (\text{A2})$$

Assume that \hat{g}^0 is deterministic, with initial state i_0 (the argument is easily modified without the assumption). Then since $\sum_{z \in X(i, \phi)} P^{\hat{\sigma}}(z|S, i_0) \equiv f_i^S$, the DM's expected payoff is

$$\Pi(\hat{g}_0, \hat{\sigma}, \hat{d}) = \sum_{i \in \mathcal{N}} \left[\pi d_i \sum_{z \in X(i, \phi)} P^{\hat{\sigma}}(z|H, i_0) + (1 - \pi)(1 - d_i) \sum_{z \in X(i, \phi)} p_{\hat{\sigma}}(z|L, i_0) \right]$$

By (A1) and (A2), the derivative of $\Pi(\hat{g}_0, \sigma, \hat{d})$ w.r.t. $\sigma_{i^*, j^*}^{k^*}$, evaluated at $\sigma = \hat{\sigma}$, is

$$\sum_{i \in \mathcal{N}} \left[\begin{aligned} & \pi d_i \left(\sum_{z' \in X(i^*, k^*)} P^{\hat{\sigma}}(z'|H, i_0) \right) \left(\sum_{z \in X(i, \phi)} P^{\hat{\sigma}}(z|H, j^*) \right) \\ & + (1 - \pi)(1 - d_i) \left(\sum_{z' \in X(i^*, k^*)} P^{\hat{\sigma}}(z'|L, i_0) \right) \left(\sum_{z \in X(i, \phi)} P^{\hat{\sigma}}(z|L, j^*) \right) \end{aligned} \right] \quad (\text{A3})$$

Recall from the text (under equation (2)) that

$$f_{i^*}^S \mu_{k^*}^S = \sum_{\tau=0}^{\infty} \eta (1 - \eta)^\tau g_{i^*}^{\tau, S} \mu_{k^*}^S = \sum_{\tau=0}^{\infty} \sum_{z' \in X_\tau(i^*, k^*)} \eta P^{\hat{\sigma}}(z'|S, i_0) = \sum_{z' \in X(i^*, k^*)} \eta P^{\hat{\sigma}}(z'|S, i_0)$$

So the expression in (A3) can be written as

$$\begin{aligned} & \frac{1}{\eta} \sum_{i \in \mathcal{N}} \left[\pi d_i f_{i^*}^H \mu_{k^*}^H \sum_{z \in X(i, \phi)} P^{\hat{\sigma}}(z|H, j^*) + (1 - \pi)(1 - d_i) f_{i^*}^L \mu_{k^*}^L \sum_{z \in X(i, \phi)} P^{\hat{\sigma}}(z|L, j^*) \right] \\ & = \frac{1}{\eta} \left[\pi f_{i^*}^H \mu_{k^*}^H \left(\sum_{i \in \mathcal{N}} d_i \sum_{z \in X(i, \phi)} P^{\hat{\sigma}}(z|H, j^*) \right) \right. \\ & \quad \left. + (1 - \pi) f_{i^*}^L \mu_{k^*}^L \left(\sum_{i \in \mathcal{N}} (1 - d_i) \sum_{z \in X(i, \phi)} P^{\hat{\sigma}}(z|L, j^*) \right) \right] \\ & \equiv \frac{1}{\eta} \left[\pi f_{i^*}^H \mu_{k^*}^H v_{j^*}^H + (1 - \pi) f_{i^*}^L \mu_{k^*}^L v_{j^*}^L \right] \quad (\text{A4}) \end{aligned}$$

(recall that $v_{j^*}^S$ is the expected payoff conditional on S and initial state j^*). Since $\hat{\sigma}$ is optimal and $\hat{\sigma}_{i^*,j^*}^{k^*} > 0$, it must be that for any $j' \neq j^*$,

$$\left. \frac{d\Pi}{d\sigma(i^*, k^*)(j^*)} \right|_{\hat{\sigma}} \geq \left. \frac{d\Pi}{d\sigma(i^*, k^*)(j')} \right|_{\hat{\sigma}}, \text{ with equality if } \hat{\sigma}(i^*, k^*)(j') > 0$$

By (A4), this says $\pi f_{i^*}^H \mu_{k^*}^H (v_{j^*}^H - v_{j'}^H) + (1 - \pi) f_{i^*}^L \mu_{k^*}^L (v_{j^*}^L - v_{j'}^L) \geq 0$, which is exactly the condition in Definition 2 for incentive compatibility of $\hat{\sigma}_{i^*,j^*}^{k^*}$. Since i^*, k^*, j^* were chosen arbitrarily, this implies that any optimal strategy is incentive compatible. ■

A.2 Proofs of Theorems 2, 3 i-iii,v,vi; and Corollary to 3

The proof will proceed through five lemmas. Lemma 1 shows that sum (1) in the text is well-defined, and obtains a useful expression for f^S . Lemma 2 is used in the subsequent proofs, and also contains some results which hold for general η : in particular, part (3) implies that the DM only moves to higher memory states after signals k which satisfy $\mu_k^H / \mu_k^L > 1$, and to lower memory states after signals satisfying the reverse inequality. Lemma 3 proves two of the results in Theorem 2: existence of an optimal strategy, and monotonicity in η . It also proves upper hemi-continuity of the set of optimal strategies, which is useful in the subsequent proofs. Lemma 4 obtains an upper bound on the expected payoff, which nearly completes the proof of Theorem 2 - it just remains to be shown that the upper bound is attained in the limit as $\eta \rightarrow 0$. Lemma 5 shows that the upper payoff bound can only be attained by a transition rule in which there are no jumps between states, and in which all but the most extreme signals are ignored; this result is the main ingredient in the proof of Theorem 3 parts (i) and (ii). The proofs of Lemmas 4 and 5 have been improved from the original version, using arguments from Cover and Hellman (1970).

Lemma 1: For any (g^0, σ) , the sum $f^S \equiv \lim_{T \rightarrow \infty} \eta g_0 \left(\sum_{t=0}^{T-1} ((1 - \eta)T^S)^t \right)$ converges to the steady-state distribution of a Markov process with transition probabilities $\omega_{i,j}^S = \eta g_j^0 + \tau_{i,j}^S$.

Proof: Observe that $\frac{1}{1-\eta} T^S$ is a stochastic matrix (row sum $\sum_j \frac{\tau_{i,j}^S}{1-\eta} = \sum_j \sum_k \sigma_{i,j}^k \mu_k^S = 1$), therefore all of its eigenvalues are less than or equal to 1 in absolute value. This implies that

$\left| \frac{I}{1-\eta} - \frac{T^S}{1-\eta} \right| \neq 0$, so $(I - T^S)$ is invertible with inverse $(I - T^S)^{-1} = \sum_{t=0}^T (T^S)^t$. Therefore by (A1),

$$f^S = \eta g^0 (I - T^S)^{-1} \Rightarrow f^S = \eta g^0 + f^S T^S$$

Since f^S is a probability distribution over \mathcal{N} , the vector ηg^0 can be written as $f^S [\eta g^0]$, where $[\eta g^0]$ is a matrix with N identical rows, each equal to the row vector ηg^0 . Thus $f^S = f^S ([\eta g^0] + T^S)$; this implies the result, since the (i, j) th entry of the matrix $([\eta g^0] + T^S)$ is $\omega_{i,j}^S$. ■

Lemma 2: *If (σ, d) is optimal, then for all i, j with $f_i^S, f_j^S > 0$:*

1. $\pi_i v_i^H + (1 - \pi_i) v_i^L \geq \pi_i v_j^H + (1 - \pi_i) v_j^L$
2. *If $\pi_j < \pi_i$, then $v_i^H \geq v_j^H$, and $v_i^L \leq v_j^L$*
3. *If $\pi_j < \pi_i$ and $v_j^S \neq v_i^S$, then $\sigma_{i,j}^k > 0 \Rightarrow \frac{\mu_k^H}{\mu_k^L} \leq 1$, and $\sigma_{j,i}^k > 0 \Rightarrow \frac{\mu_k^H}{\mu_k^L} \geq 1$*

Proof: Pick any two states i and j with $f_i^S, f_j^S > 0$. For $k \in \mathcal{K}$, define

$$i(k) = \arg \max_{i'} [\pi_i \mu_k^H \cdot v_{i'}^H + (1 - \pi_i) \mu_k^L \cdot v_{i'}^L]$$

and define $j(k)$ similarly. For all $k \in \mathcal{K}$, pick $i_k^* \in i(k)$ and $j_k^* \in j(k)$.

By the definition of v^H, v^L , and incentive compatibility, the expected payoff $V_i \equiv \pi_i v_i^H + (1 - \pi_i) v_i^L$ is equal to

$$\begin{aligned} V_i &= \eta (\pi_i d_i + (1 - \pi_i)(1 - d_i)) + (1 - \eta) \sum_{k \in \mathcal{K}} \left[\pi_i \mu_k^H \cdot v_{i_k^*}^H + (1 - \pi_i) \mu_k^L \cdot v_{i_k^*}^L \right] \\ &\geq \eta (\pi_i d_j + (1 - \pi_i)(1 - d_j)) + (1 - \eta) \sum_{k \in \mathcal{K}} \left[\pi_i \mu_k^H \cdot v_{j_k^*}^H + (1 - \pi_i) \mu_k^L \cdot v_{j_k^*}^L \right] \\ &= \pi_i v_j^H + (1 - \pi_i) v_j^L \end{aligned}$$

This proves part (1). For part (2): part (1) implies that

$$\pi_i (v_i^H - v_j^H) + (1 - \pi_i) (v_i^L - v_j^L) \geq 0 \quad (\text{A5})$$

$$-\pi_j (v_i^H - v_j^H) - (1 - \pi_j) (v_i^L - v_j^L) \geq 0 \quad (\text{A6})$$

Adding the inequalities,

$$(\pi_i - \pi_j) \cdot [(v_i^H - v_j^H) - (v_i^L - v_j^L)] \geq 0$$

So if $\pi_i > \pi_j$, this implies $(v_i^H - v_j^H) \geq (v_i^L - v_j^L)$. Thus the LHS of (A5) is at most $(v_i^H - v_j^H)$, so the inequality requires $v_i^H \geq v_j^H$. Similarly the LHS of (A6) is at most $-(v_i^L - v_j^L)$, requiring $v_i^L \leq v_j^L$; this proves part (2).

For part (3): pick $i, j \in \mathcal{N}$ with $j < i$ such that for some $k \in \mathcal{K}$, $\sigma_{i,j}^k > 0$. By part (1),

$$\begin{aligned} \pi_i (v_i^H - v_j^H) + (1 - \pi_i) (v_i^L - v_j^L) &\geq 0 \\ \Leftrightarrow \frac{\pi_i}{1 - \pi_i} \frac{(v_i^H - v_j^H)}{(v_j^L - v_i^L)} &\geq 1 \text{ (by part (2), provided } v_j^L \neq v_i^L) \end{aligned}$$

So $\frac{\mu_k^H}{\mu_k^L} > 1$ implies $\frac{\pi_i}{1 - \pi_i} \frac{\mu_k^H}{\mu_k^L} \frac{(v_i^H - v_j^H)}{(v_j^L - v_i^L)} > 1$. Thus after receiving signal k , a DM in memory state i strictly prefers state i to state j ; so incentive compatibility requires $\sigma_{i,j}^k = 0$, a contradiction; therefore it must be that $\frac{\mu_k^H}{\mu_k^L} \leq 1$. The proof that $\sigma_{j,i}^k > 0 \Rightarrow \frac{\mu_k^H}{\mu_k^L} \geq 1$ is identical. ■

Lemma 3: *For any $\eta > 0$, an optimal rule σ^η exists. The set of optimal rules σ^η is upper hemi-continuous in η , and the maximized expected payoff $\Pi^*(\eta, N)$ is continuous in η . Moreover, provided that d_i is not identically 0 or 1 for all i , $\Pi^*(\eta, N)$ is strictly decreasing in η .*

Proof: Given an action rule d , an optimal transition rule σ^η maximizes the expression in (2). By Lemma 1, the probability distribution f^S is equal to the long-run distribution of a Markov process with transition probabilities $\omega_{i,j}^S = \eta g_j^0 + \tau_{i,j}^S$. Thus, choosing σ to achieve the optimal distribution f^S in the terminal period, subject to termination probability η , is equivalent to choosing σ to achieve the optimal long-run distribution, subject to the constraint that $\sigma_{i,j}^k \geq \eta g_j^0$ for all $i, j \in \mathcal{N}$ and $k \in \mathcal{K}$.

To prove existence and continuity of $\Pi^*(\eta, N)$ in η : By Lemma 1, $f^S = \eta g^0 (I - T^S)^{-1}$; therefore $f_i^S = \frac{|I - T^S|_i}{|I - T^S|}$, where $|I - T^S|_i$ is the determinant of the matrix obtained by replacing the i th row of $(I - T^S)$ with g^0 . Since the numerator and denominator are both polynomials in the transition probabilities $\sigma_{i,j}^k$ and the denominator is non-zero for all $\eta > 0$

(again by Lemma 1), it follows that f_i^S , and hence the DM's expected payoff, is continuous in σ for all $\eta \in (0, 1)$. So the problem is to maximize a continuous objective function over a compact constraint set which is continuous in η . Then the theorem of the maximum implies that an optimal solution exists, the set of optimal solutions is upper hemi-continuous in η , and the maximized payoff is continuous in η .

To prove that $\Pi^*(\eta, N)$ is strictly decreasing in η (unless d_i is constant across i): choose a state j with $g_j^0 > 0$. Lemma 2 part (2) implies that $\sigma_{i,j}^k$ can only be optimal if $i < j$ and $\frac{\mu_k^H}{\mu_k^L} > 1$, or $i > j$ and $\frac{\mu_k^H}{\mu_k^L} < 1$. Therefore the constraint $\sigma_{i,j}^k \geq \eta g_j^0 \forall i \in \mathcal{N}, k \in \mathcal{K}$ must be binding when $\eta > 0$, establishing that $\Pi^*(\eta, N)$ is strictly decreasing in η . ■

Lemma 4: Define $r^* = \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H} \right)^{\frac{N-1}{2}}$. Then the DM's expected payoff $\Pi^*(\eta, N)$ is bounded above by

$$\overline{\Pi^*(N)} \equiv \begin{cases} \max\{\pi, 1 - \pi\} & \text{if } r^* \geq \min\left\{\sqrt{\frac{1-\pi}{\pi}}, \sqrt{\frac{\pi}{1-\pi}}\right\} \\ \frac{1-2\sqrt{\pi(1-\pi)}r^*}{1-(r^*)^2} & \text{if } r^* < \min\left\{\sqrt{\frac{1-\pi}{\pi}}, \sqrt{\frac{\pi}{1-\pi}}\right\} \end{cases}$$

Moreover, $\Pi(g^0, \sigma, d) = \frac{1-2\sqrt{\pi(1-\pi)}r^*}{1-(r^*)^2}$ iff $r \equiv \sqrt{\frac{\sum_{i \in \mathcal{N}} f_i^L d_i \sum_{i \in \mathcal{N}} f_i^H (1-d_i)}{\sum_{i \in \mathcal{N}} f_i^H d_i \sum_{i \in \mathcal{N}} f_i^L (1-d_i)}} = r^*$ and $x \equiv \frac{\sum_{i \in \mathcal{N}} f_i^H (1-d_i)}{\sum_{i \in \mathcal{N}} f_i^H d_i} = \frac{r^*(1-\sqrt{\frac{\pi}{1-\pi}}r^*)}{(\sqrt{\frac{\pi}{1-\pi}}-r^*)}$.

Proof: Define $x \equiv \frac{\sum_{i \in \mathcal{N}} f_i^H (1-d_i)}{\sum_{i \in \mathcal{N}} f_i^H d_i}$, and $r \equiv \sqrt{\frac{\sum_{i \in \mathcal{N}} f_i^L d_i \sum_{i \in \mathcal{N}} f_i^H (1-d_i)}{\sum_{i \in \mathcal{N}} f_i^H d_i \sum_{i \in \mathcal{N}} f_i^L (1-d_i)}}$. If d_i is not identically zero for all i (and therefore x, r are positive and finite), then by equation (2) the DM's expected payoff can be written as

$$\frac{\pi \sum_i f_i^H d_i}{\sum_i f_i^H d_i + \sum_i f_i^H (1-d_i)} + \frac{(1-\pi) \sum_i f_i^L (1-d_i)}{\sum_i f_i^L (1-d_i) + \sum_i f_i^L d_i} = \frac{\pi}{1+x} + \frac{1-\pi}{1+\frac{r^2}{x}} \quad (\text{A7})$$

If $r > \sqrt{\frac{1-\pi}{\pi}}$, then this is below π whenever $x > 0$; and if $r > \sqrt{\frac{\pi}{1-\pi}}$, then this is below $(1-\pi)$ for finite x . Therefore it follows that if $r \geq \min\left\{\sqrt{\frac{1-\pi}{\pi}}, \sqrt{\frac{\pi}{1-\pi}}\right\}$, then the highest achievable payoff is $\max\{\pi, 1-\pi\}$, attained by setting $d_i = 1 \forall i$ if $\pi > \frac{1}{2}$, and $d_i = 0 \forall i$ if $\pi < \frac{1}{2}$.

If $r < \left\{\sqrt{\frac{1-\pi}{\pi}}, \sqrt{\frac{\pi}{1-\pi}}\right\}$, then the RHS of (A7) is increasing in x whenever $x > \frac{r(1-\sqrt{\frac{\pi}{1-\pi}}r)}{(\sqrt{\frac{\pi}{1-\pi}}-r)}$, and decreasing in x whenever the reverse inequality holds; therefore the payoff is at most

$$\max_r \left(\left[\frac{\pi}{1+x} + \frac{1-\pi}{1+\frac{r^2}{x}} \right]_{x=\frac{r(1-\sqrt{\frac{\pi}{1-\pi}}r)}{(\sqrt{\frac{\pi}{1-\pi}}-r)}} \right) = \max_r \frac{1-2r\sqrt{\pi(1-\pi)}}{(1-r^2)}$$

This is decreasing in r whenever $r < \left\{ \sqrt{\frac{1-\pi}{\pi}}, \sqrt{\frac{\pi}{1-\pi}} \right\}$: so to complete the proof, we need to show that if d_i is not identically 0 or 1 for all i , then r^* is a lower bound on r . This yields the desired upper bound on the payoff, which is attained iff $r = r^*$ and $x = \frac{r^*(1-\sqrt{\frac{\pi}{1-\pi}}r^*)}{(\sqrt{\frac{\pi}{1-\pi}}-r^*)}$.

To see this, let \mathcal{M} be the set of states i with $f_i^S > 0$. Recall from Lemma 1 that f^S is the steady-state distribution of a Markov process with transition probabilities $\omega_{i,j}^S = \eta g_j^0 + \tau_{i,j}^S$; and note also that if $f_i^S > 0$ and $\omega_{i,j}^S > 0$, then $f_j^S > 0$ so $j \in \mathcal{M}$. This implies that for all $i \in \mathcal{M}$ and for $S = H, L$,

$$\sum_{\{j \in \mathcal{M} | j \leq i\}} f_j^S = \sum_{\{j \in \mathcal{M} | j \leq i\}} f_j^S \left(1 - \sum_{\{l \in \mathcal{M} | l \geq i+1\}} \omega_{j,l}^S \right) + \sum_{\{l \in \mathcal{M} | l \geq i+1\}} f_l^S \sum_{\{j \in \mathcal{M} | j \leq i\}} \omega_{l,j}^S$$

Since the signal ordering implies that $\frac{\mu_1^H}{\mu_1^L} \leq \frac{\omega_{i,j}^H}{\omega_{i,j}^L} \leq \frac{\mu_K^H}{\mu_K^L}$ for any i, j with $\omega_{i,j}^S > 0$, and since $\frac{f_i^H}{f_i^L}$ is non-decreasing, this implies that for any $i \in \mathcal{M}$

$$\frac{f_i^H}{f_i^L} \frac{\mu_K^H}{\mu_K^L} \geq \frac{\sum_{\{j \in \mathcal{M} | j \leq i\}} f_j^H \sum_{\{l \in \mathcal{M} | l \geq i+1\}} \omega_{j,l}^H}{\sum_{\{j \in \mathcal{M} | j \leq i\}} f_j^L \sum_{\{l \in \mathcal{M} | l \geq i+1\}} \omega_{j,l}^L} = \frac{\sum_{\{l \in \mathcal{M} | l \geq i+1\}} f_l^H \sum_{\{j \in \mathcal{M} | j \leq i\}} \omega_{l,j}^H}{\sum_{\{l \in \mathcal{M} | l \geq i+1\}} f_l^L \sum_{\{j \in \mathcal{M} | j \leq i\}} \omega_{l,j}^L} \geq \min_{\{l \in \mathcal{M} | l \geq i+1\}} \frac{f_l^H}{f_l^L} \frac{\mu_1^H}{\mu_1^L} \quad (\text{A8})$$

Iterating this argument,

$$\min_{i \in \mathcal{M}} \frac{f_i^H}{f_i^L} \geq \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H} \right)^{\#\mathcal{M}-1} \max_{l \in \mathcal{M}} \frac{f_l^H}{f_l^L} \quad (\text{A9})$$

Finally, again using the fact that $\frac{f_i^H}{f_i^L}$ is non-decreasing for $i \in \mathcal{M}$, this implies

$$r^2 \equiv \frac{\sum_{i \in \mathcal{M}} f_i^H (1-d_i) \sum_i f_i^L d_i}{\sum_{i \in \mathcal{M}} f_i^L (1-d_i) \sum_i f_i^L d_i} \geq \min_{i \in \mathcal{M}} \frac{f_i^H}{f_i^L} / \max_{l \in \mathcal{M}} \frac{f_l^H}{f_l^L} \geq \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H} \right)^{\#\mathcal{M}-1} \quad (\text{A10})$$

which completes the proof since $\#\mathcal{M} \leq N$, with equality iff $\mathcal{M} = \mathcal{N}$. ■

Lemma 5: Let σ^η be a convergent sequence of strategies with $\eta \rightarrow 0$, and let $f^{S,\eta}, r^\eta$ be the f^S, r induced by σ^η . Then $\lim_{\eta \rightarrow 0} r^\eta \geq r^*$, with equality iff the following conditions hold for all η sufficiently small: (i) $\sigma_{i,j}^{k,\eta} = 0$ whenever $j \geq i + 2$ or $j \leq i - 2$; (ii) $\sigma_{i,i}^{k,\eta} = 1$ whenever $k \notin \{1, K\}$; (iii) $\lim_{\eta \rightarrow 0} \frac{\eta}{\tau_{i,i+1}^{S,\eta}} = \lim_{\eta \rightarrow 0} \frac{\eta}{\tau_{i,i-1}^{S,\eta}} = 0 \forall i$; (iv) $\lim_{\eta \rightarrow 0} \frac{f_i^{S,\eta}}{f_1^{S,\eta}} = \lim_{\eta \rightarrow 0} \frac{f_i^{S,\eta}}{f_N^{S,\eta}} = 0 \forall i \in \mathcal{N} \setminus \{1, N\}$; (v) there exists $\bar{\eta}$ s.t. $f_i^{S,\eta} > 0 \forall \eta < \bar{\eta}$.

Proof: For part (v): if the statement is false, then for all $\eta < \bar{\eta}$, the set \mathcal{M}^η of states with $f_i^{S,\eta} > 0$ satisfies $\#\mathcal{M}^\eta \leq N - 1$; then (A10) implies that $r^\eta \geq \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H}\right)^{\frac{N-2}{2}}$, hence $\lim_{\eta \rightarrow 0} r^\eta \geq \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H}\right)^{\frac{N-2}{2}} > r^*$. Thus $\lim_{\eta \rightarrow 0} r^\eta = r^*$ can only hold if $f_i^{S,\eta} > 0$ for sufficiently small η ; assume for the remainder of the proof that $f_i^{S,\eta} > 0 \forall i \in \mathcal{N}$. For parts (i),(ii), and (iii), note by (A10) that the bound r^* can only be attained if (A9) holds with equality. By (A8), and since $f_i^{S,\eta} > 0$, this can only hold in the limit if $\lim_{\eta \rightarrow 0} \frac{f_i^{H,\eta} \mu_K^H}{f_i^{L,\eta} \mu_K^L} = \lim_{\eta \rightarrow 0} \frac{f_{i+1}^{H,\eta} \mu_1^H}{f_{i+1}^{L,\eta} \mu_1^L} \forall i \in \mathcal{N}$; by (A8), this requires

$$\lim_{\eta \rightarrow 0} \frac{f_i^{H,\eta} \mu_K^H}{f_i^{L,\eta} \mu_K^L} = \lim_{\eta \rightarrow 0} \frac{\sum_{j \leq i} f_j^{H,\eta} \sum_{l \geq i+1} \omega_{j,l}^{H,\eta}}{\sum_{j \leq i} f_j^{L,\eta} \sum_{l \geq i+1} \omega_{j,l}^{L,\eta}} \quad (\text{A11})$$

$$\lim_{\eta \rightarrow 0} \frac{f_{i+1}^{H,\eta} \mu_1^H}{f_{i+1}^{L,\eta} \mu_1^L} = \lim_{\eta \rightarrow 0} \frac{\sum_{l \geq i+1} f_l^{H,\eta} \sum_{j \leq i} \omega_{l,j}^{H,\eta}}{\sum_{l \geq i+1} f_l^{L,\eta} \sum_{j \leq i} \omega_{l,j}^{L,\eta}} \quad (\text{A12})$$

Since $\frac{\omega_{j,l}^{H,\eta}}{\omega_{j,l}^{L,\eta}} \leq \frac{\mu_K^H}{\mu_K^L}$, (A11) can only hold if $\lim_{\eta \rightarrow 0} \omega_{j,l}^{S,\eta} = 0$ whenever $\frac{f_j^{H,\eta}}{f_j^{L,\eta}} < \frac{f_i^{H,\eta}}{f_i^{L,\eta}}$; so if $\frac{f_j^{H,\eta}}{f_j^{L,\eta}} = \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H}\right) \frac{f_{j+1}^{H,\eta}}{f_{j+1}^{L,\eta}} < \frac{f_{j+1}^{H,\eta}}{f_{j+1}^{L,\eta}}$ for all pairs $(j, j+1)$, this requires $\lim_{\eta \rightarrow 0} \sum_{l \geq i+1} \omega_{j,l}^{S,\eta} = 0 \forall j \leq i+1$. By an identical argument, (A12) can only hold if $\lim_{\eta \rightarrow 0} \sum_{j \leq i} \omega_{l,j}^{S,\eta} = 0$ whenever $l \geq i+2$. This establishes part (i). For part (ii), assume condition (i) holds: then (A11) reduces to the requirement

$$\frac{\mu_K^H}{\mu_K^L} = \lim_{\eta \rightarrow 0} \frac{\omega_{i,i+1}^{H,\eta}}{\omega_{i,i+1}^{L,\eta}} \equiv \lim_{\eta \rightarrow 0} \frac{\eta g_{i+1}^0 + \tau_{i,i+1}^{H,\eta}}{\eta g_{i+1}^0 + \tau_{i,i+1}^{L,\eta}}$$

This holds only if for all i , $\lim_{\eta \rightarrow 0} \frac{\eta}{\tau_{i,i+1}^{S,\eta}} = 0$, and if $\frac{\tau_{i,i+1}^{H,\eta}}{\tau_{i,i+1}^{L,\eta}} \equiv \frac{\sum_{k \in \mathcal{K}} \mu_k^H \sigma_{i,i+1}^{k,\eta}}{\sum_{k \in \mathcal{K}} \mu_k^L \sigma_{i,i+1}^{k,\eta}} = \frac{\mu_K^H}{\mu_K^L} \Leftrightarrow \sigma_{i,i+1}^{k,\eta} = 0$ whenever $k \neq K$. By an identical argument, (A12) holds only if for all i , $\lim_{\eta \rightarrow 0} \eta / \tau_{i,i-1}^{S,\eta} = 0$,

and $\sigma_{i,i-1}^{k,\eta} = 0$ whenever $k \neq 1$. This establishes parts (ii) and (iii). Finally for part (iv), observe (since $\frac{f_i^H}{f_i^L}$ is non-decreasing) that the second last inequality in (A10) can only hold

with equality if $\lim_{\eta \rightarrow 0} \frac{\sum_{\{i|d_i < 1\}} f_i^{H,\eta}(1-d_i)}{\sum_{\{i|d_i < 1\}} f_i^{L,\eta}(1-d_i)} = \lim_{\eta \rightarrow 0} \frac{f_1^{H,\eta}}{f_1^{L,\eta}}$, and $\lim_{\eta \rightarrow 0} \frac{\sum_{\{i|d_i > 0\}} f_i^{L,\eta} d_i}{\sum_{\{i|d_i > 0\}} f_i^{H,\eta} d_i} = \lim_{\eta \rightarrow 0} \frac{f_N^{L,\eta}}{f_N^{H,\eta}}$.

If (A8) holds with equality for all i , implying that $\frac{f_i^{H,\eta}}{f_i^{L,\eta}}$ is in fact strictly increasing in i , then

this can only hold if $\lim_{\eta \rightarrow 0} \frac{f_i^{S,\eta}}{f_1^{S,\eta}} = 0 \forall i$ with $d_i < 1$, and $\lim_{\eta \rightarrow 0} \frac{f_i^{S,\eta}}{f_N^{S,\eta}} = 0 \forall i$ with $d_i > 0$. ■

A.2.1 Proof of Theorem 2

The existence, continuity, and monotonicity results were established in Lemma 3, and the upper bound on the payoff was established in Lemma 4. To complete the proof, it remains to be shown that this upper bound is achieved by a sequence of optimal strategies in the limit as $\eta \rightarrow 0$. For this, it is sufficient by Lemma 4 to construct a sequence of strategies $(g^{0,\eta}, \sigma^\eta, d^\eta)$ which satisfies $\lim_{\eta \rightarrow 0} r^\eta = r^*$, and $\lim_{x \rightarrow 0} x^\eta = \frac{r^*(1 - \sqrt{\frac{\pi}{1-\pi}} r^*)}{(\sqrt{\frac{\pi}{1-\pi}} - r^*)}$ (with r^η, x^η as defined in Lemma 4, for the strategy $(g^{0,\eta}, \sigma^\eta, d^\eta)$).

Choose any g^0 with $g_1^0 = g_N^0 = 0$, and any d with $d_1 = 0, d_N = 1$; and for all η set $d^\eta = d, g^{0,\eta} = g^0$. Now consider a sequence σ^η of transition rules satisfying the following conditions: (i) $\sigma_{i,i}^k = 1 \forall k \in \mathcal{K} \setminus \{1, K\}, i \in \mathcal{N}$; (ii) for $i \in \mathcal{N} \setminus \{1, N\}$, $\sigma_{i,i+1}^{K,\eta} = 1$ (so $\sigma_{i,j}^{K,\eta} = 0$ if $j \neq i+1$) and $\sigma_{i,i-1}^{1,\eta} = 1$ (so $\sigma_{i,j}^{1,\eta} = 0 \forall j \neq i-1$); (iii) $\sigma_{1,2}^{K,\eta} = \sqrt{\eta}, \sigma_{N,N-1}^{1,\eta} = \sqrt{\eta} \left(\frac{\mu_K^H \mu_K^L}{\mu_1^H \mu_1^L} \right)^{\frac{N-1}{2}} \frac{(1 - \sqrt{\frac{\pi}{1-\pi}} r^*)}{(\sqrt{\frac{\pi}{1-\pi}} - r^*)}$.

Recall from Lemma 1 that f^S is the steady-state distribution of a Markov process with transition probabilities $\omega_{i,j}^S$. This implies that for any i , $f_i^S \sum_{j \neq i} \omega_{i,j}^S = \sum_j f_j^S \omega_{j,i}^S$. Under the rule (g^0, σ^η) , this implies for $i = 1$ that $f_1^{S,\eta} \left(\eta + (1-\eta) \sigma_{1,2}^{K,\eta} \mu_K^S \right) = f_2^{S,\eta} (1-\eta) \mu_1^S \Rightarrow \lim_{\eta \rightarrow 0} \frac{f_2^{S,\eta}}{\sqrt{\eta} f_1^{S,\eta}} = \frac{\mu_K^S}{\mu_1^S}$. Similarly, the expression for $i = N$ yields $\lim_{\eta \rightarrow 0} \frac{f_N^S \sigma_{N,N-1}^1}{f_{N-1}^S} = \frac{\mu_K^S}{\mu_1^S}$. For any $2 \leq i \leq N-2$, f_i^S solves

$$\begin{aligned} f_i^S &= f_i^S (1 - (1-\eta) \mu_1^S - (1-\eta) \mu_K^S - \eta(1-g_i^S)) \\ &\quad + f_{i-1}^S (\eta g_i^S + (1-\eta) \sigma_{i-1,i}^K \mu_K^S) + f_{i+1}^S (\eta g_i^S + (1-\eta) \mu_1^S) \end{aligned}$$

which implies

$$\lim_{\eta \rightarrow 0} \frac{f_{i+1}^{S,\eta}}{\sqrt{\eta} f_1^{S,\eta}} = \left(1 + \frac{\mu_K^S}{\mu_1^S}\right) \lim_{\eta \rightarrow 0} \frac{f_i^{S,\eta}}{\sqrt{\eta} f_1^{S,\eta}} - \frac{\mu_K^S}{\mu_1^S} \lim_{\eta \rightarrow 0} \frac{f_{i-1}^{S,\eta} \sigma_{i-1,i}^K}{\sqrt{\eta} f_1^{S,\eta}} \quad (\text{A13})$$

Solving (A13) recursively, using the above expression $\lim_{\eta \rightarrow 0} \frac{f_2^{S,\eta}}{\sqrt{\eta} f_1^{S,\eta}} = \frac{\mu_K^S}{\mu_1^S}$, we obtain

$$\lim_{\eta \rightarrow 0} \frac{f_{N-1}^{S,\eta}}{\sqrt{\eta} f_1^{S,\eta}} = \left(\frac{\mu_K^H}{\mu_1^H}\right)^{N-2}; \text{ then since } \lim_{\eta \rightarrow 0} \frac{f_N^{H,\eta} \sigma_{N,N-1}}{f_{N-1}^{H,\eta}} = \frac{\mu_K^H}{\mu_1^H}, \text{ this implies that}$$

$$\left(\frac{\mu_K^H \mu_K^L}{\mu_1^H \mu_1^L}\right)^{\frac{N-1}{2}} \frac{\left(1 - \sqrt{\frac{\pi}{1-\pi}} r^*\right)}{\left(\sqrt{\frac{\pi}{1-\pi}} - r^*\right)} \lim_{\eta \rightarrow 0} \frac{f_N^{H,\eta}}{f_1^{H,\eta}} \equiv \lim_{\eta \rightarrow 0} \frac{f_{N-1}^{H,\eta}}{\sqrt{\eta} f_1^{H,\eta}} \frac{f_N^{H,\eta} \sigma_{N,N-1}}{f_{N-1}^{H,\eta}} = \left(\frac{\mu_K^H}{\mu_1^H}\right)^{N-1}$$

Solving this yields $\lim_{\eta \rightarrow 0} x^\eta \equiv \lim_{\eta \rightarrow 0} \frac{f_1^H}{f_N^H} = \frac{r^* \left(1 - \sqrt{\frac{\pi}{1-\pi}} r^*\right)}{\left(\sqrt{\frac{\pi}{1-\pi}} - r^*\right)}$, as desired; moreover, note that σ^η satisfies all conditions in Lemma 5, implying that $\lim_{\eta \rightarrow 0} r^\eta = r^*$. ■

A.2.2 Proof of Theorem 3 (i)-(iii)

Consider a sequence $(g^{0,\eta} \sigma^\eta, d^\eta)$ of rules, with $\eta \rightarrow 0$. Theorem 2 calculated an upper bound on the expected payoff, and showed that this bound $\overline{\Pi^*(N)}$ is achievable in the limit as $\eta \rightarrow 0$; since the set of optimal strategies is upper hemi-continuous in η (by Lemma 3), this implies that $(g^{0,\eta}, \sigma^\eta, d^\eta)$ can only be optimal if $\lim_{\eta \rightarrow 0} \Pi(d^\eta, g^{0,\eta}, \sigma^\eta) = \overline{\Pi^*(N)}$. By Lemma 4, this requires $\lim_{\eta \rightarrow 0} r^\eta = r^*$; then by Lemma 5 parts (i) and (ii) imply parts (i) and (ii) of the Theorem. Part (iii) follows almost immediately from Lemma 5 part (iv), which implies that an optimal sequence $(g^{0,\eta}, \sigma^\eta, d^\eta)$ must satisfy $\lim_{\eta \rightarrow 0} \frac{f_2^{S,\eta}}{f_1^{S,\eta}} = \lim_{\eta \rightarrow 0} \frac{f_{N-1}^{S,\eta}}{f_N^{S,\eta}} = 0$: This is only possible if $\lim_{\eta \rightarrow 0} \tau_{1,2}^S = \lim_{\eta \rightarrow 0} \tau_{N,N-1}^S = 0$, since Lemma 1 implies that for a rule satisfying Theorem 3 (i), $\frac{f_2^S}{f_1^S} = \frac{\omega_{1,2}^S}{\omega_{2,1}^S}$ and $\frac{f_{N-1}^S}{f_N^S} = \frac{\omega_{N,N-1}^S}{\omega_{N-1,N}^S}$. ■

A.2.3 Proof of the Corollary to Theorem 3

As argued above in the proof of Theorem 3 (i)-(iii), an optimal sequence of strategies $(g^{0,\eta}, \sigma^\eta, d^\eta)$ must satisfy $\lim_{\eta \rightarrow 0} \Pi(g^{0,\eta}, \sigma^\eta, d^\eta) = \overline{\Pi^*(N)}$. By Lemma 4, this requires $\lim_{\eta \rightarrow 0} r^\eta = r^*$,

and

$$\lim_{\eta \rightarrow 0} \frac{\sum_i f_i^{H,\eta} (1 - d_i)}{\sum_i f_i^{H,\eta} d_i} \equiv \lim_{\eta \rightarrow 0} x^\eta = \lim_{\eta \rightarrow 0} \frac{r^* \left(1 - \sqrt{\frac{\pi}{1-\pi}} r^*\right)}{\left(\sqrt{\frac{\pi}{1-\pi}} - r^*\right)} \quad (\text{A14})$$

By Lemma 5 part (iv), $\lim_{\eta \rightarrow 0} r^\eta = r^*$ requires $\lim_{\eta \rightarrow 0} \sum_i f_i^{S,\eta} = \lim_{\eta \rightarrow 0} (f_1^{S,\eta} + f_N^{S,\eta})$; since $f_i^{S,\eta}$ is a probability distribution, this implies that $\lim_{\eta \rightarrow 0} f_1^{S,\eta} = 1 - \lim_{\eta \rightarrow 0} f_N^{S,\eta}$, and that $\lim_{\eta \rightarrow 0} \frac{\sum_i f_i^H (1 - d_i)}{\sum_i f_i^H d_i} = \lim_{\eta \rightarrow 0} \frac{f_1^{H,\eta}}{f_N^{H,\eta}}$. Therefore (A14) implies

$$\begin{aligned} \lim_{\eta \rightarrow 0} f_1^{H,\eta} &= \frac{r^* \left(1 - \sqrt{\frac{\pi}{1-\pi}} r^*\right)}{\left(\sqrt{\frac{\pi}{1-\pi}} - r^*\right)} \left(1 - \lim_{\eta \rightarrow 0} f_1^{H,\eta}\right) \\ \Rightarrow \lim_{\eta \rightarrow 0} f_1^{H,\eta} &= \frac{\sqrt{\frac{1-\pi}{\pi}} r^* - (r^*)^2}{1 - (r^*)^2}, \quad \lim_{\eta \rightarrow 0} f_N^{H,\eta} = \frac{1 - \sqrt{\frac{1-\pi}{\pi}} r^*}{1 - (r^*)^2} \end{aligned} \quad (\text{A15})$$

Substituting these expressions into the condition $\lim_{\eta \rightarrow 0} r^\eta = \lim_{\eta \rightarrow 0} \sqrt{\frac{f_1^H f_N^L}{f_1^L f_N^H}} = r^*$ yields an expression for $\lim_{\eta \rightarrow 0} \frac{f_1^{L,\eta}}{f_N^{L,\eta}}$, which can be solved using $\lim_{\eta \rightarrow 0} (f_1^{L,\eta} + f_N^{L,\eta}) = 1$ to obtain

$$\lim_{\eta \rightarrow 0} f_N^{L,\eta} = \frac{\sqrt{\frac{\pi}{1-\pi}} r^* - (r^*)^2}{1 - (r^*)^2}, \quad \lim_{\eta \rightarrow 0} f_1^L = 1 - \lim_{\eta \rightarrow 0} f_N^{L,\eta} = \frac{1 - \sqrt{\frac{\pi}{1-\pi}} r^*}{1 - (r^*)^2} \quad (\text{A16})$$

Then

$$\lim_{\eta \rightarrow 0} \frac{f_N^{H,\eta}}{f_N^{L,\eta}} = \frac{1 - \sqrt{\frac{1-\pi}{\pi}} r^*}{\sqrt{\frac{\pi}{1-\pi}} r^* - (r^*)^2} = \sqrt{\frac{1-\pi}{\pi}} \frac{1}{r^*} \equiv \sqrt{\frac{1-\pi}{\pi}} \left(\frac{\mu_K^H \mu_1^L}{\mu_K^L \mu_1^H}\right)^{\frac{N-1}{2}}$$

This yields the desired expression for $i = N$. Then the proof follows from (A11) and Lemma 5 (i)-(iii), which imply that $\lim_{\eta \rightarrow 0} \frac{f_i^{H,\eta}}{f_i^{L,\eta}} = \lim_{\eta \rightarrow 0} \frac{f_{i+1}^{H,\eta} \mu_1^H \mu_K^L}{f_{i+1}^{L,\eta} \mu_1^L \mu_K^H} \forall i$; this gives

$$\lim_{\eta \rightarrow 0} \frac{f_i^{H,\eta}}{f_i^{L,\eta}} = \lim_{\eta \rightarrow 0} \frac{f_N^{H,\eta}}{f_N^{L,\eta}} \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H}\right)^{N-i} = \sqrt{\frac{1-\pi}{\pi}} \left(\frac{\mu_K^H \mu_1^L}{\mu_K^L \mu_1^H}\right)^{i - \frac{N+1}{2}}$$

■

A.2.4 Proof of Theorem 3 parts (v) and (vi)

Since Theorem 1 implies that an optimal rule must be incentive compatible, the DM will start in the state i which maximizes his expected continuation payoff at the prior belief, π . In particular, starting in state i must yield a (weakly) higher continuation payoff than starting in state $i + 1$, or state $i - 1$: using Lemma 2, this condition requires $\frac{\pi}{1 - \pi} \left(\frac{v_{i+1}^H - v_i^H}{v_i^L - v_{i+1}^L} \right) \leq 1 \leq \frac{\pi}{1 - \pi} \left(\frac{v_i^H - v_{i-1}^H}{v_{i-1}^L - v_i^L} \right)$. By Lemma 5 (iii), $\sigma_{i+1,i}^1 < 0$ and $\sigma_{i,i+1}^K > 0$: and by Theorem 1, this can only be optimal if

$$\frac{\pi}{1 - \pi} \frac{f_{i+1}^H \mu_1^H}{f_{i+1}^L \mu_1^L} \left(\frac{v_{i+1}^H - v_i^H}{v_i^L - v_{i+1}^L} \right) \leq 1 \leq \frac{\pi}{1 - \pi} \frac{f_i^H \mu_K^H}{f_i^L \mu_K^L} \left(\frac{v_{i+1}^H - v_i^H}{v_i^L - v_{i+1}^L} \right) \quad (\text{A17})$$

Taking limits as $\eta \rightarrow 0$, using the Corollary to Theorem 3, and noting (by the Corollary) that the LHS,RHS converge to the same limit, this implies

$$\lim_{\eta \rightarrow 0} \left(\frac{v_{i+1}^H - v_i^H}{v_i^L - v_{i+1}^L} \right) = \sqrt{\frac{1 - \pi}{\pi}} \frac{\mu_K^L}{\mu_K^H} \left(\frac{\mu_K^L \mu_1^H}{\mu_K^H \mu_1^L} \right)^{i - \frac{N+1}{2}}$$

Substituting this into the optimality condition for starting in i rather than $i - 1$ or $i + 1$, we conclude that starting in i is optimal only if

$$\sqrt{\frac{\pi}{1 - \pi}} \frac{\mu_K^L}{\mu_K^H} \left(\frac{\mu_K^L \mu_1^H}{\mu_K^H \mu_1^L} \right)^{i - \frac{N+1}{2}} \leq 1 \leq \sqrt{\frac{\pi}{1 - \pi}} \frac{\mu_1^L}{\mu_1^H} \left(\frac{\mu_K^L \mu_1^H}{\mu_K^H \mu_1^L} \right)^{i - \frac{N+1}{2}}$$

Generically this will only hold for one integer i - namely, the smallest integer i for which $\sqrt{\frac{\pi}{1 - \pi}} \frac{\mu_K^L}{\mu_K^H} \left(\frac{\mu_K^L \mu_1^H}{\mu_K^H \mu_1^L} \right)^{i - \frac{N+1}{2}} \leq 1$; this yields the expression in (v).

For part (vi): by Theorem 1, $d_i = 0$ is optimal iff $\frac{\pi_i}{1 - \pi_i} \equiv \frac{\pi}{1 - \pi} \frac{f_i^H}{f_i^L} \leq 1$, and $d_i = 1$ is optimal iff the reverse inequality holds. Using the Corollary to Theorem 3, this says that $d_i = 0$ is optimal iff

$$\sqrt{\frac{\pi}{1 - \pi}} \left(\frac{\mu_K^H \mu_1^L}{\mu_K^L \mu_1^H} \right)^{i - \frac{N+1}{2}} \leq 1 \Leftrightarrow i \leq \frac{N + 1}{2} + \frac{\log \sqrt{\frac{1 - \pi}{\pi}}}{\log \left(\frac{\mu_K^H \mu_1^L}{\mu_K^L \mu_1^H} \right)}$$

as desired. ■

A.3 Proof of Theorem 3 (iv)

Define i_0 as the initial state, and j^* as the highest state j with $d_j = 0$. For any $j \leq j^*$, define

$$\widetilde{\Delta}_{j,j+1}^S \equiv \left(\prod_{l=2}^j \tau_{l,l+1}^S \right) \left(\frac{v_{j+1}^S - v_j^S}{v_2^S - v_1^S} \right); \text{ and for } j \geq j^* \text{ define}$$

$$\widehat{\Delta}_{j,j+1}^S \equiv \left(\prod_{l=j+1}^{N-1} \tau_{l,l-1}^S \right) \left(\frac{v_{j+1}^S - v_j^S}{v_N^S - v_{N-1}^S} \right). \text{ For any } j \leq i_0 \text{ define } \widetilde{f}_j^S \equiv \left(\prod_{l=2}^j \tau_{l,l-1}^S \right) \frac{f_j^S}{f_1^S}; \text{ and for}$$

any $j \geq i_0$ define $\widehat{f}_j^S = \left(\prod_{l=j}^{N-1} \tau_{l,l+1}^S \right) \frac{f_j^S}{f_N^S}$. The proof will proceed through four claims which calculate useful recursion formulas for these expressions.

Claim 1: $\widetilde{f}_i^S, \widetilde{\Delta}_{i,i+1}^S$ satisfy the following recursion formulas:

$$\widetilde{f}_i^S = \widetilde{\Delta}_{i,i+1}^S - \omega_{i,i-1}^S \widetilde{\Delta}_{i-1,i}^S \quad (\text{A18})$$

$$\eta \widetilde{\Delta}_{i,i+1}^S = \widetilde{f}_{i+1}^S - \omega_{i,i+1}^S \widetilde{f}_i^S \quad (\text{A19})$$

Proof (by induction): First observe that f_i^S, v_i^S solve the following recursion formulas when $i < i_0$ and $d_i = 0$:

$$f_i^S (\eta + \omega_{i,i-1}^S + \omega_{i,i+1}^S) = \omega_{i-1,i}^S f_{i-1}^S + \omega_{i+1,i}^S f_{i+1}^S \quad (\text{A20})$$

$$v_i^S (\eta + \omega_{i,i-1}^S + \omega_{i,i+1}^S) = \omega_{i,i-1}^S v_{i-1}^S + \omega_{i,i+1}^S v_{i+1}^S + \eta \cdot 1_{S=L} \quad (\text{A21})$$

Evaluating (A20) at $i = 1$ yields $f_1^S (\eta + \omega_{1,2}^S) = \omega_{2,1}^S f_2^S$, which implies $\widetilde{f}_2^S \equiv \omega_{2,1}^S f_2^S / f_1^S = \eta + \omega_{1,2}^S$. Then since $\widetilde{\Delta}_{1,2}^S = \widetilde{f}_1^S = 1$ and there is no state zero, an immediate calculation yields (A18),(A19) at $i = 1$.

Now suppose (A18),(A19) both hold at $i - 1$. To show that (A18) holds also at i , solve (A21) for $(v_{i+1}^S - v_i^S)$, obtaining

$$\begin{aligned} \omega_{i,i+1}^S (v_{i+1}^S - v_i^S) &= (\eta + \omega_{i,i-1}^S) (v_i^S - v_{i-1}^S) + \eta (v_{i-1}^S - 1_{S=L}) \\ &= (\eta + \omega_{i,i-1}^S) (v_i^S - v_{i-1}^S) + \eta (v_{i-1}^S - v_{i-2}^S) + \eta (v_{i-2}^S - 1_{S=L}) \end{aligned}$$

Now lag the first line of this equation by one, solve for $\eta(v_{i-1}^S - 1_{S=L})$, and substitute the resulting expression into the second equation: this yields

$$\omega_{i,i+1}^S (v_{i+1}^S - v_i^S) = (\eta + \omega_{i,i-1}^S) (v_i^S - v_{i-1}^S) + \omega_{i-1,i}^S (v_i^S - v_{i-1}^S) - \omega_{i-1,i-2}^S (v_{i-1}^S - v_{i-2}^S)$$

Multiplying both sides of this equation by $\frac{\prod_{j=2}^{i-1} \omega_{j,j+1}^H}{v_2^H - v_1^H}$, we obtain

$$\widetilde{\Delta_{i,i+1}^S} = (\eta + \omega_{i,i-1}^S) \widetilde{\Delta_{i-1,i}^S} + \omega_{i-1,i}^S \left(\widetilde{\Delta_{i-1,i}^S} - \omega_{i-1,i-2}^S \widetilde{\Delta_{i-2,i-1}^S} \right)$$

By (A18) at $i-1$, the last bracketed term is equal to $\widetilde{f_{i-1}^S}$; by (A19) at $i-1$, $\eta \widetilde{\Delta_{i-1,i}^S}$ is equal to $\widetilde{f_i^S} - \omega_{i-1,i}^S \widetilde{f_{i-1}^S}$; substituting both into the above expression, we obtain (A18) for i .

To show that (A19) holds at i , multiply both sides of (A20) by $\left(\prod_{j=2}^i \tau_{j,j-1}^S \right) / f_1^S$, to obtain

$$\begin{aligned} \widetilde{f_{i+1}^S} &= \omega_{i,i+1}^S \widetilde{f_i^S} + \eta \widetilde{f_i^S} + \omega_{i,i-1}^S \left(\widetilde{f_i^S} - \omega_{i-1,i}^S \widetilde{f_{i-1}^S} \right) \\ &= \omega_{i,i+1}^S \widetilde{f_i^S} + \eta \left(\widetilde{f_i^S} + \omega_{i,i-1}^S \widetilde{\Delta_{i-1,i}^S} \right) \text{ by (A19) at } i-1 \\ &= \omega_{i,i+1}^S \widetilde{f_i^S} + \eta \widetilde{\Delta_{i,i+1}^S} \text{ by (A18) at } i \end{aligned}$$

This is equivalent to (A19) at i . ■

Claim 2: $\widetilde{f_j^S}, \widehat{f_j^S}, \widetilde{\Delta_{j,j+1}^S}, \widehat{\Delta_{j,j+1}^S}$ satisfy the following recursion formulas:

$$\begin{aligned} \frac{\widetilde{f_j^S}}{\left(\prod_{l=1}^{j-1} \sigma_{l,l+1} \right) (\mu_K^S)^{j-1}} &= \eta \left(\sum_{l=1}^{j-1} \left(\prod_{m=2}^l \frac{\sigma_{m,m-1}^1}{\sigma_{m,m+1}^K} \right) \frac{1}{\mu_1^S} \left(\frac{\mu_1^S}{\mu_K^S} \right)^l \right) + \sigma_{1,2}^K + o(\eta) \\ \frac{\widetilde{\Delta_{j,j+1}^S}}{\left(\prod_{l=2}^{j-1} \sigma_{l+1,l} \right) (\mu_1^S)^{j-1}} &= \sigma_{2,1}^1 + \sigma_{1,2}^K \sum_{l=1}^{j-1} \left(\prod_{m=2}^l \frac{\sigma_{m,m+1}^K}{\sigma_{m+1,m}^1} \right) \left(\frac{\mu_K^S}{\mu_1^S} \right)^l + O(\eta) \\ \frac{\widehat{f_j^S} / (\mu_1^S)^{N-j}}{\left(\prod_{l=1}^{N-1-j} \sigma_{N-l,N-l-1}^1 \right)} &= \eta \left(\sum_{l=1}^{N-j} \left(\prod_{m=1}^{l-1} \frac{\sigma_{N-m,N-m+1}^K}{\sigma_{N-m,N-m-1}^1} \right) \frac{1}{\mu_K^S} \left(\frac{\mu_K^S}{\mu_1^S} \right)^l \right) + \sigma_{N,N-1}^1 + o(\eta) \\ \frac{\widehat{\Delta_{j,j+1}^S} / \sigma_{N-1,N}^K (\mu_K^S)^{N-j-1}}{\left(\prod_{l=2}^{N-1-j} \sigma_{N-l,N-l+1}^K \right)} &= 1 + \frac{\sigma_{N,N-1}^1}{\sigma_{N-1,N}^K} \sum_{l=1}^{N-1-j} \left(\prod_{m=2}^l \frac{\sigma_{N-m+1,N-m}^1}{\sigma_{N-m,N-m+1}^1} \right) \left(\frac{\mu_1^S}{\mu_K^S} \right)^l + O(\eta) \end{aligned}$$

Proof: Solving (A18),(A19) recursively yields

$$\begin{aligned}\widetilde{f}_j^S &= \left(\prod_{l=2}^{j-1} \sigma_{l,l+1} \right) \left(1 + \sum_{l=1}^{j-2} \prod_{m=1}^l \frac{\tau_{m+1,m}^S}{\tau_{m+1,m+2}^S} \right) (\mu_K^S)^{j-2} \eta + \left(\prod_{l=1}^{j-1} \sigma_{l,l+1} \right) (\mu_K^S)^{j-1} + o(\eta) \\ \widetilde{\Delta}_{j,j+1}^S &= \left(\prod_{l=1}^{j-1} \sigma_{l+1,l} \right) (\mu_1^S)^{j-1} + \left(\prod_{l=1}^{j-1} \sigma_{l,l+1} \right) (\mu_K^S)^{j-1} \left(1 + \sum_{l=1}^{j-2} \prod_{m=1}^l \frac{\tau_{j-(m-1),j-m}^S}{\tau_{j-m,j-(m-1)}^S} \right) + O(\eta)\end{aligned}$$

which can be arranged to the expressions in Claim 2 for $\widetilde{f}_j^S, \widetilde{\Delta}_{j,j+1}^S$. The expressions for $\widehat{f}_j^S, \widehat{\Delta}_{j,j+1}^S$ can then be obtained by symmetry. ■

Claim 3: Let σ^η be a sequence of transition rules with $\eta \rightarrow 0$, and for any integer l define

$$\begin{aligned}a_l^\eta &\equiv \prod_{m=2}^l \frac{\sigma_{m,m-1}^{1,\eta}}{\sigma_{m,m+1}^{K,\eta}} \left(\frac{1}{\mu_1^L} \left(\frac{\mu_1^L}{\mu_K^L} \right)^l - \frac{1}{\mu_1^H} \left(\frac{\mu_1^H}{\mu_K^H} \right)^l \right) & d_l^\eta &\equiv \prod_{m=2}^l \frac{\sigma_{N-m+1,N-m}^{1,\eta}}{\sigma_{N-m,N-m+1}^{K,\eta}} \left(\left(\frac{\mu_1^L}{\mu_K^L} \right)^l - \left(\frac{\mu_1^H}{\mu_K^H} \right)^l \right) \\ c_l^\eta &\equiv \prod_{m=2}^l \frac{\sigma_{N-m+1,N-m+2}^{K,\eta}}{\sigma_{N-m+1,N-m}^{1,\eta}} \left(\frac{1}{\mu_K^H} \left(\frac{\mu_K^H}{\mu_1^H} \right)^l - \frac{1}{\mu_K^L} \left(\frac{\mu_K^L}{\mu_1^L} \right)^l \right) & b_l^\eta &\equiv \prod_{m=2}^l \frac{\sigma_{m+1,m}^{K,\eta}}{\sigma_{m+1,m}^{1,\eta}} \left(\left(\frac{\mu_K^H}{\mu_1^H} \right)^l - \left(\frac{\mu_K^L}{\mu_1^L} \right)^l \right)\end{aligned}$$

If the optimality conditions for 1 and N are satisfied, then the optimality conditions for $\sigma_{j,j+1}^{K,\eta} > 0$ imply the following:

$$\begin{aligned}j \leq \min\{i_0, j^*\} : \quad & \lim_{\eta \rightarrow 0} \left(\frac{(\sigma_{1,2}^{K,\eta})^2}{\eta \sigma_{2,1}^{1,\eta}} - \frac{\sum_{l=1}^{j-1} a_l^\eta}{\sum_{l=1}^{j-1} b_l^\eta} \right) \geq 0 \\ j^* < j \leq i_0 : \quad & \lim_{\eta \rightarrow 0} \left(\frac{(\sigma_{1,2}^{K,\eta})^2}{\eta \sigma_{2,1}^{1,\eta}} - \frac{\sum_{l=1}^{j-1} a_l^\eta}{\sum_{l=1}^{j^*-1} b_l^\eta + \frac{\sigma_{2,1}^1 \sigma_{N,N-1}^{1,\eta}}{\sigma_{1,2}^K \sigma_{N-1,N}^{K,\eta}} \sum_{N-j}^{N-1-j^*} d_l^\eta} \right) \geq 0 \\ j \geq i_0 : \quad & \lim_{\eta \rightarrow 0} \left(\frac{(\sigma_{N,N-1}^{1,\eta})^2}{\eta \sigma_{N-1,N}^{K,\eta}} - \frac{\sum_{l=1}^{N-j} c_l^\eta}{\sum_{l=1}^{N-j-1} d_l^\eta} \right) \leq 0\end{aligned}$$

Each condition must hold with equality if $\lim_{\eta \rightarrow 0} \sigma_{j,j+1}^{K,\eta} \in (0, 1)$. The conditions for $\sigma_{j,j-1}^1 > 0$ are identical except for the following modifications: for $j \leq \min\{i_0, j^*\}$, remove the term b_{j-1}^η from the denominator; for $j^* < j \leq i_0$, remove the term d_{N-j}^η from the denominator; for $j \geq i_0$, add the term d_{N-j}^η to the denominator; and in all three cases, reverse the inequality.

Proof: First consider the condition for $j \leq \min\{i_0, j^*\}$. By Theorem 3 (iii) we know that $\sigma_{1,2}^K \in (0, 1)$, and by Theorem 1 the corresponding optimality condition is $\frac{\pi}{1-\pi} \frac{f_1^H \mu_K^H v_2^H - v_1^H}{f_1^L \mu_K^L v_1^L - v_2^L} = 1$; if this condition holds, then the optimality condition for $\sigma_{j,j+1}^K \geq 0$ can be written as $\frac{\widehat{f_j^H} \widehat{\Delta_{j,j+1}^H}}{(\mu_1^H \mu_K^H)^{j-1}} \geq \frac{\widehat{f_j^L} \widehat{\Delta_{j,j+1}^L}}{(\mu_1^L \mu_K^L)^{j-1}}$. Using the expressions in Claim 2 and the above definitions for a_l, b_l , this requires $(\sigma_{1,2}^K)^2 \sum_{l=1}^{j-1} b_l - \eta \sigma_{2,1}^1 \left(\sum_{l=1}^{j-1} a_l \right) + o(\eta) = 0$, which implies the desired expression in the limit. The calculation for $\sigma_{j,j-1}^1$ is similar, and the calculations when $j \geq i_0$ are symmetric - just measuring states by their distance from N rather than from 1, and interchanging signals 1 and K . For $j^* < j \leq i_0$: if the optimality conditions for $\sigma_{1,2}^K \in (0, 1), \sigma_{N,N-1}^1 \in (0, 1)$ both hold, then the optimality condition for $\sigma_{j,j+1} > 0$ can be written as $\frac{\widehat{f_j^H} \widehat{\Delta_{j,j+1}^H} \widehat{\Delta_{j^*,j^*+1}^H} \widehat{\Delta_{j^*,j^*+1}^L}}{\widehat{f_j^L} \widehat{\Delta_{j,j+1}^L} \widehat{\Delta_{j^*,j^*+1}^L} \widehat{\Delta_{j^*,j^*+1}^H}} \left(\frac{\mu_1^L \mu_K^L}{\mu_1^H \mu_K^H} \right)^{j^*-1} \geq 1$; using the expressions in Claim 2, this implies the desired expression. ■

Claim 4: If $\lim_{\eta \rightarrow 0} \sigma_{j,j+1}^{K,\eta} \in (0, 1)$, then $\lim_{\eta \rightarrow 0} \sigma_{j,j-1}^{1,\eta} = 1$ and $\lim_{\eta \rightarrow 0} \sigma_{j+1,j}^{1,\eta} = 1$.

Proof: It is clear from the equations in Claim 3 that it is impossible for the optimality conditions for $\sigma_{j,j+1}^K, \sigma_{j,j-1}^1$ to both hold with equality; and similarly for $\sigma_{j,j+1}^K$ and $\sigma_{j+1,j}^1$. ■

Proof of Theorem 3 (iv): We first prove the result for states $j \leq \min\{i_0, j^*\}$. For the first part, existence of $\Delta_1, \Delta_K > 0$, we need to show that for any state $2 \leq j \leq \min\{i_0, j^*\}$, both $\lim_{\eta \rightarrow 0} \sigma_{j,j-1}^{1,\eta}$ and $\lim_{\eta \rightarrow 0} \sigma_{j,j+1}^{K,\eta}$ are bounded away from zero. To see this, suppose the claim is false: then there exists an optimal sequence of strategies such that for some state $2 \leq j \leq \min\{i_0, j^*\}$, either $\lim_{\eta \rightarrow 0} \sigma_{j,j-1}^{1,\eta} = 0$ or $\lim_{\eta \rightarrow 0} \sigma_{j,j+1}^{K,\eta} = 0$; choose the smallest such j . Then the definitions in Claim 3 imply that $\lim_{\eta \rightarrow 0} a_{l+1}, \lim_{\eta \rightarrow 0} b_l$ are both positive and finite for any $l \leq j-2$; while if $\lim_{\eta \rightarrow 0} \sigma_{j,j-1}^{1,\eta} = 0$, then $b_{j-1} \rightarrow \infty$ as $\eta \rightarrow 0$; and if $\lim_{\eta \rightarrow 0} \sigma_{j,j+1}^{K,\eta} = 0$, then $a_j \rightarrow \infty$ as $\eta \rightarrow 0$. Suppose first that $j \geq 3$ $\lim_{\eta \rightarrow 0} \sigma_{j,j-1}^{1,\eta} = 0$. Then by Claim 3, optimality of $\sigma_{j+1,j}^{1,\eta}$ requires

$$\lim_{\eta \rightarrow 0} \frac{(\sigma_{1,2}^{K,\eta})^2}{\eta \sigma_{2,1}^{1,\eta}} = 0 < \lim_{\eta \rightarrow 0} \frac{\sum_{l=1}^{j-2} a_l}{\sum_{l=1}^{j-2} b_l}$$

which violates the optimality condition for $\sigma_{j-1,j}^{K,\eta}$. Therefore if the strategy sequence is optimal, then there exists $\bar{\eta}$ such that whenever $\eta < \bar{\eta}$, either $\sigma_{j-1,j}^{K,\eta} = 0$ or $\sigma_{j+1,j}^{1,\eta} = 0$; this

violates Lemma 5 (iii), contradicting optimality of σ^η . Next, note that the optimality condition in Claim 3 for $\sigma_{2,1}^{1,\eta}$ can never hold with equality, so an optimal strategy sequence must have $\lim_{\eta \rightarrow 0} \sigma_{2,1}^{1,\eta} = 0$. Finally, suppose that $\lim_{\eta \rightarrow 0} \sigma_{j,j+1}^{K,\eta} = 0$. Then optimality of $\sigma_{j+1,j+2}^{K,\eta}$ requires that $\frac{(\sigma_{1,2}^{K,\eta})^2}{\eta \sigma_{2,1}^{1,\eta}} \rightarrow \infty$ as $\eta \rightarrow 0$, which violates the optimality condition for $\sigma_{j,j-1}^{1,\eta}$; as above, this leads to a contradiction of optimality using Lemma 5 (iii).

For the second part of the result, we need to prove that for all $2 \leq j \leq \min\{i_0, j^*\}$, $\lim_{\eta \rightarrow 0} \sigma_{j,j+1}^{K,\eta} = 1$ whenever $\frac{\mu_1^L \mu_1^H}{\mu_K^L \mu_K^H} \geq 1$, and $\lim_{\eta \rightarrow 0} \sigma_{j,j-1}^{1,\eta} = 1$ whenever $\frac{\mu_1^L \mu_1^H}{\mu_K^L \mu_K^H} \leq 1$. First consider the case $\frac{\mu_1^L \mu_1^H}{\mu_K^L \mu_K^H} \geq 1$, and suppose the claim is false. Choose the smallest $2 \leq j \leq j^*$ with $\lim_{\eta \rightarrow 0} \sigma_{j,j+1}^{K,\eta} \in (0, 1)$, and let i be the highest state below j to randomize after one of signals 1, K . The optimality condition for $\sigma_{j,j+1}^{K,\eta} \in (0, 1)$ requires $\lim_{\eta \rightarrow 0} \frac{1}{\eta} \frac{(\sigma_{1,2}^{K,\eta})^2}{\sigma_{2,1}^{1,\eta}} = \lim_{\eta \rightarrow 0} \frac{\sum_{l=1}^{j-1} a_l}{\sum_{l=1}^{j-1} b_l}$; and if this equality holds, then the optimality condition for $\sigma_{j+1,j+2}^{K,\eta} > 0$ requires $\lim_{\eta \rightarrow 0} \frac{(\sigma_{1,2}^{K,\eta})^2}{\eta \sigma_{2,1}^{1,\eta}} \geq \lim_{\eta \rightarrow 0} \frac{a_j}{b_j}$. So to obtain a contradiction, it is sufficient to show that $\lim_{\eta \rightarrow 0} \frac{\sum_{l=1}^{j-1} a_l}{\sum_{l=1}^{j-1} b_l} < \lim_{\eta \rightarrow 0} \frac{a_j}{b_j}$. To see this, suppose first that $i = 1$. Then since optimality of $\lim_{\eta \rightarrow 0} \sigma_{j,j+1}^{K,\eta} \in (0, 1)$ implies $\lim_{\eta \rightarrow 0} \sigma_{j,j-1}^{1,\eta} = \lim_{\eta \rightarrow 0} \sigma_{j+1,j}^{K,\eta} = 1$ (by Claim 4) and no states between 1 and j randomize, it follows for any $l \leq j-1$, the expression for $\frac{\sigma_{2,1}^{1,\eta} a_l}{b_l^\eta}$ reduces to $\frac{1}{\mu_1^L} \left(\frac{\mu_1^H \mu_K^L}{\mu_K^H \mu_1^L} \right)^l \left[\left(1 - \frac{\mu_1^L}{\mu_1^H} \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H} \right)^l \right) / \left(1 - \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H} \right)^l \right) \right]$. Since $\frac{\mu_1^L}{\mu_1^H} > 1 > \frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H}$ (by the signal ordering), the term in square brackets is increasing in l ; then $\frac{\mu_1^L \mu_1^H}{\mu_K^L \mu_K^H} \geq 1$ implies that $\frac{a_l}{b_l^\eta}$ is strictly increasing in l for $l \leq j-1$. Therefore,

$$\lim_{\eta \rightarrow 0} \frac{\sum_{l=1}^{j-1} a_l}{\sum_{l=1}^{j-1} b_l^\eta} < \lim_{\eta \rightarrow 0} \frac{a_{j-1}}{b_{j-1}^\eta} < \lim_{\eta \rightarrow 0} \frac{\frac{1}{\mu_1^L} \left(\frac{\mu_1^H \mu_1^L}{\mu_K^H \mu_K^L} \right)^j \left(1 - \frac{\mu_1^L}{\mu_1^H} \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H} \right)^j \right)}{\sigma_{2,1}^{1,\eta} \left(\sigma_{j,j+1}^{K,\eta} \right)^2 \left(1 - \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H} \right)^j \right)} \equiv \lim_{\eta \rightarrow 0} \frac{a_j}{b_j^\eta}$$

as desired. Suppose next that $i > 1$. Then the optimality condition for $\sigma_{i,i-1}^{1,\eta} \in (0, 1)$ implies $\lim_{\eta \rightarrow 0} \frac{(\sigma_{1,2}^{K,\eta})^2}{\eta} = \frac{\sigma_{2,1}^{1,\eta} \sum_{l=1}^{i-1} a_l}{\sum_{l=1}^{i-2} b_l^\eta}$; so using this, the condition for $\lim_{\eta \rightarrow 0} \sigma_{j,j+1}^{K,\eta} \in (0, 1)$ reduces to $\lim_{\eta \rightarrow 0} \frac{(\sigma_{1,2}^{K,\eta})^2}{\eta} = \lim_{\eta \rightarrow 0} \frac{\sigma_{2,1}^{1,\eta} \sum_{l=i}^{j-1} a_l}{\sum_{l=i-1}^{j-1} b_l^\eta}$. Using the fact that no states between i and j randomize,

this can be written as

$$\lim_{\eta \rightarrow 0} \frac{(\sigma_{1,2}^{K,\eta})^2}{\eta \sigma_{2,1}^{1,\eta}} = \lim_{\eta \rightarrow 0} \frac{\sum_{l=i}^{j-1} a_l^\eta}{\sum_{l=i-1}^{j-1} b_l^\eta} < \lim_{\eta \rightarrow 0} \frac{\sum_{l=i}^{j-1} a_l^\eta}{\sum_{l=i}^{j-1} b_l^\eta} = \lim_{\eta \rightarrow 0} \left(\prod_{m=2}^i \frac{(\sigma_{m,m-1}^{1,\eta})^2}{(\sigma_{m,m+1}^{K,\eta})^2} \right) \frac{\sum_{l=i}^{j-1} \left(\frac{1}{\mu_1^L} \left(\frac{\mu_1^L}{\mu_K^L} \right)^l - \frac{1}{\mu_1^H} \left(\frac{\mu_1^H}{\mu_K^H} \right)^l \right)}{\sum_{l=i}^{j-1} \left(\left(\frac{\mu_K^H}{\mu_1^H} \right)^l - \left(\frac{\mu_K^L}{\mu_1^L} \right)^l \right)}$$

which is less than $\lim_{\eta \rightarrow 0} \frac{a_j^\eta}{b_j^\eta} = \lim_{\eta \rightarrow 0} \left(\frac{\prod_{m=2}^i (\sigma_{m,m-1}^{1,\eta})^2}{(\sigma_{j,j+1}^{K,\eta})^2 \prod_{m=2}^i \sigma_{m,m+1}^2} \right) \left(\frac{\frac{1}{\mu_1^L} \left(\frac{\mu_1^L}{\mu_K^L} \right)^j - \frac{1}{\mu_1^H} \left(\frac{\mu_1^H}{\mu_K^H} \right)^j}{\left(\frac{\mu_K^H}{\mu_1^H} \right)^j - \left(\frac{\mu_K^L}{\mu_1^L} \right)^j} \right)$ by

the same argument as above, proving the claim if $\frac{\mu_1^L \mu_1^H}{\mu_K^L \mu_K^H} \geq 1$.

Next, let $\frac{\mu_1^L \mu_1^H}{\mu_K^L \mu_K^H} \leq 1$ and suppose the claim is false. Let $2 \leq j \leq j^*$ be the smallest state with $\sigma_{j,j-1} \in (0, 1)$, and let i be the largest state below j to randomize. By construction of j , we know that $\sigma_{i,i+1}^{K,\eta} \in (0, 1)$; the corresponding optimality condition is $\lim_{\eta \rightarrow 0} \frac{(\sigma_{1,2}^{K,\eta})^2}{\eta \sigma_{2,1}^{1,\eta}} = \frac{\sum_{l=1}^{i-1} a_l}{\sum_{l=1}^{i-1} b_l}$; and if this condition holds then the optimality condition for $\sigma_{j,j-1}^{1,\eta} \in (0, 1)$ reduces to

$$\lim_{\eta \rightarrow 0} \frac{(\sigma_{1,2}^{K,\eta})^2}{\eta \sigma_{2,1}^{1,\eta}} = \frac{\sum_{l=i}^{j-1} a_l}{\sum_{l=i}^{j-2} b_l} = \left(\prod_{m=2}^i \frac{\sigma_{m,m-1}^2}{\sigma_{m,m+1}^2} \right) \frac{\sum_{l=i}^{j-1} \left(\frac{1}{\mu_1^L} \left(\frac{\mu_1^L}{\mu_K^L} \right)^l - \frac{1}{\mu_1^H} \left(\frac{\mu_1^H}{\mu_K^H} \right)^l \right)}{\sum_{l=i}^{j-2} \left(\left(\frac{\mu_K^H}{\mu_1^H} \right)^l - \left(\frac{\mu_K^L}{\mu_1^L} \right)^l \right)}$$

If this holds, then the condition for $\sigma_{j+1,j}^{1,\eta} > 0$ requires

$$\lim_{\eta \rightarrow 0} \frac{(\sigma_{1,2}^{K,\eta})^2}{\eta \sigma_{2,1}^{1,\eta}} \leq \lim_{\eta \rightarrow 0} \frac{a_j^\eta}{b_{j-1}^\eta} \equiv \lim_{\eta \rightarrow 0} \frac{(\sigma_{j,j-1}^{1,\eta})^2}{\sigma_{2,1}^{1,\eta}} \left(\prod_{m=2}^i \frac{\sigma_{m,m-1}^2}{\sigma_{m,m+1}^2} \right) \frac{\frac{1}{\mu_1^L} \left(\frac{\mu_1^L}{\mu_K^L} \right)^j - \frac{1}{\mu_1^H} \left(\frac{\mu_1^H}{\mu_K^H} \right)^j}{\left(\frac{\mu_K^H}{\mu_1^H} \right)^{j-1} - \left(\frac{\mu_K^L}{\mu_1^L} \right)^{j-1}}$$

So to obtain a contradiction, it is sufficient to show that $\lim_{\eta \rightarrow 0} \frac{a_i + \sum_{l=i}^{j-2} a_{l+1}}{\sum_{l=i}^{j-2} b_l} > \frac{a_j}{b_{j-1}}$. This

follows from the fact that when $\frac{\mu_1^L \mu_1^H}{\mu_K^L \mu_K^H} \leq 1$,

$$\prod_{m=2}^i \frac{(\sigma_{m,m+1}^{K,\eta})^2}{(\sigma_{m,m-1}^{1,\eta})^2} \frac{\sigma_{2,1}^{1,\eta} a_{l+1}^\eta}{b_l^\eta} = \frac{1}{\mu_K^H} \left(\frac{\mu_1^L \mu_1^H}{\mu_K^L \mu_K^H} \right)^l \left(\frac{\frac{\mu_K^H}{\mu_1^H} - \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H} \right)^l}{1 - \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H} \right)^l} \right)$$

is strictly decreasing in l .

This completes the proof for $j \leq \min\{i_0, j^*\}$. The proof for $j \geq \max\{i_0, j^*\}$ is symmetric, since the expression for $\sigma_{j,j+1}^K > 0$ ($\sigma_{j,j-1}^1 > 0$) with $j > \max\{i_0, j^*\}$ is symmetric to the expression for $\sigma_{j,j-1}^1 > 0$ ($\sigma_{j,j+1}^K > 0$) with $j < \min\{i_0, j^*\}$ - just renumbering states according to their distance from N rather than from 1, and interchanging signals 1, K . The proof for j between i_0, j^* can be obtained by a similar argument. ■

A.4 Proof of Theorem 4

Let K^τ denote a block of τ consecutive K -signals; decompose the set $[T, \tau]$ as $\{(s^{T-\tau}, K^\tau) \mid s^{T-\tau} \in \mathcal{K}^{T-\tau}\}$, and the set $[\tau, T]$ as $\{(K^\tau, s^{T-\tau}) \mid s^{T-\tau} \in \mathcal{K}^{T-\tau}\}$.

For part (i): fix an optimal transition rule σ , and let S be the true state of the world; define $g^t(i \mid S, i_0)$ as the probability of memory state i after t periods, conditional on state S and on starting in state i_0 . Fix $\eta < \hat{\eta}$, so that Theorem 3 holds. Then an initial sequence of τ K -signals changes the initial state to $\min\{N, i_0 + \tau\}$, so $E[\pi_{i_T} \mid [\tau, T]] = \sum_{i \in \mathcal{N}} \pi_i g^{T-\tau}(i \mid S, \min\{N, i_0 + \tau\})$. For sequences in $[T, \tau]$, note that if the DM ignored the final block K^τ , then his expected belief would be $E[\pi_{i_T} \mid [T, \tau]] = \sum_{i \in \mathcal{N}} \pi_i g^{T-\tau}(i \mid S, i_0)$. Since K^τ will in fact increase the final memory state with strictly positive probability for $\eta > 0$, there exists some constant c^* s.t. $E[\pi_{i_T} \mid [T, \tau]] = \sum_{i \in \mathcal{N}} \pi_i g^{T-\tau}(i \mid S, i_0) + c^*$. Lemma 5 argued that all memory states are ergodic. Then by Theorem 11.4 of Stokey-Lucas, for any initial state i_0 , the period- T distribution $g^T(\cdot \mid i_0)$ converges to a unique limit, and convergence is at a geometric rate that is independent of i_0 . This implies that for any $\varepsilon' > 0$, there exists T' such that whenever $T > T'$,

$$|g^{T-\tau}(i \mid S, \max\{N, i_0 + \tau\}) - g^{T-\tau}(i \mid S, i_0)| < \varepsilon'$$

Therefore, for any $\varepsilon^* > 0$, there exists T^* such that whenever $T > T^*$,

$$\sum_{i \in \mathcal{N}} \pi(i) g_{T-\tau}(i \mid S, i_0) > \sum_{i \in \mathcal{N}} \pi(i) g_{T-\tau}(i \mid S, \max\{N, i_0 + \tau\}) - \varepsilon^*$$

For $\varepsilon^* < c^*$, this implies $E[\pi_{i_T} \mid S, [T, \tau]] > E[\pi_{i_T} \mid S, [\tau, T]]$, as desired.

For part (ii): suppose first that $\mu_1^H \mu_1^L = \mu_K^H \mu_K^L$. Then by Theorem 3 (iv), the transition

rule is deterministic and identical for all states $i \neq 1, N$, with $\sigma_{i,i+1}^K = \sigma_{i,i-1}^1 = 1 \forall i \in \mathcal{N} \setminus \{1, N\}$; and by Theorem 3 (iii), for any fixed T and any $\varepsilon > 0$, we can choose η small enough that the probability of leaving states $1, N$ within $T - \tau$ periods is smaller than ε . Consider any sequence $s^{T-\tau} \in \mathcal{K}^{T-\tau}$, and define $\Delta(s^{T-\tau})$ as the number of K -signals, less the number of 1-signals, in $s^{T-\tau}$. Suppose first that $s^{T-\tau}$ is such that the DM would reach state 1 before N starting in state $i_0 + \tau$: then he would also reach state 1 before N starting in state i_0 , where he would then ignore a block of K -signals with probability near 1; this means that in the limit as $\eta \rightarrow 0$, the DM will end up in memory state 1 with probability 1 after 1 after either sequence $(s^{T-\tau}, K^\tau)$ or $(K^\tau, s^{T-\tau})$. Similarly, if $s^{T-\tau}$ is such that either (i) the DM would reach state N before 1 starting in state i_0 ; or (ii) the DM would never reach state 1 or N from either i_0 or $i_0 + \tau$; then in the limit as $\eta \rightarrow 0$, the DM will end in the same state with probability 1 regardless of whether he observes $(K^\tau, s^{T-\tau})$ or $(s^{T-\tau}, K^\tau)$. So, the position of the block K^τ can only affect the final memory state with positive probability if $s^{T-\tau}$ is such that the DM would reach N before 1 starting in $i_0 + \tau$ but not starting in i_0 ; or such that he would reach state 1 before N starting in state i_0 but not starting in state $i_0 + \tau$. In this case, the DM ends in a higher memory state after $(K^\tau, s^{T-\tau})$ than $(s^{T-\tau}, K^\tau)$; then the claim follows from the fact that there is a strictly positive probability of such a sequence $s^{T-\tau}$ occurring.

If $\mu_1^H \mu_1^L \neq \mu_K^H \mu_K^L$, then the fact that the DM may randomize after (one of) signals $1, K$ means that the claim may not be true for all τ, T . Of course it is true for sufficiently large τ : if the initial block contains enough consecutive K -signals to move the DM from state i_0 to N , then as $\eta \rightarrow 0$ he will end in the maximum state N with probability 1 after $(K^T, s^{T-\tau})$ for any $s^{T-\tau}$. ■

A.5 Proof of Theorem 5

The proof for part (i) is straightforward (exactly following the argument given in the text), and is therefore omitted. For part (ii): consider the Markov process over $\mathcal{N} \times \mathcal{N}$ with transition probabilities $\tau_{(i^1, i^2), (j^1, j^2)}^{S, \eta} = (1 - \eta) \sum_{k \in \mathcal{K}} \mu_k^S \sigma_{i^1, j^1}^k \sigma_{i^2, j^2}^k$: thus $\tau_{(i^1, i^2), (j^1, j^2)}^S$ is the probability (conditional on true state S) that agents (1,2) will move from memory states (i^1, i^2) to (j^1, j^2) , assuming that they observe the same signal and follow the same rule σ .

For all $(i, j) \in \mathcal{N} \times \mathcal{N}$, define $\widehat{f}_{(i,j)}^S$ as the probability of ending in memory states (i, j) for agents $(1, 2)$, conditional on S . Then the long-run disagreement probability, $\lim_{\eta \rightarrow 0} P^*$, is given by

$$\lim_{\eta \rightarrow 0} \left(\pi \widehat{f}_{(1,N)}^{H,\eta} + (1 - \pi) \widehat{f}_{(1,N)}^{L,\eta} \right)$$

A simple calculation using (A15),(A16) shows that $\lim_{\eta \rightarrow 0} \left(\pi f_1^{H,\eta} f_N^{H,\eta} + (1 - \pi) f_1^{L,\eta} f_N^{L,\eta} \right)$ is equal to the desired expression. Therefore, it is sufficient to show that

$$\lim_{\eta \rightarrow 0} \widehat{f}_{(1,N)}^{S,\eta} = \lim_{\eta \rightarrow 0} f_1^{S,\eta} f_N^{S,\eta} \quad (\text{A22})$$

By the same argument as was used in the proof of Lemma 1, $\widehat{f}_{(i,j)}^{S,\eta}$ is uniquely determined by σ and the agents' initial states, (i_0^1, i_0^2) , and this distribution over memory state pairs is equal to the steady-state distribution of a Markov process with transition probabilities $\omega_{(i^1, i^2), (j^1, j^2)}^{S,\eta} \equiv \tau_{(i^1, i^2), (j^1, j^2)}^{S,\eta} + \eta \cdot 1_{(j^1, j^2) = (i_0^1, i_0^2)}$. Then $\widehat{f}_{(i^1, i^2)}^{S,\eta}$ solves $\widehat{f}_{(i^1, i^2)}^{S,\eta} = \sum_{(j^1, j^2) \in \mathcal{N} \times \mathcal{N}} \widehat{f}_{(j^1, j^2)}^S \omega_{(j^1, j^2), (i^1, i^2)}^S$. If both agents are following an optimal rule σ^η , then Theorem 3 (i),(ii) implies that for state (i, N) , this simplifies to

$$\begin{aligned} \frac{\widehat{f}_{(i,N)}^{S,\eta}}{1 - \eta} &= \mu_K^S \left[\widehat{f}_{(i,N)}^{S,\eta} (1 - \sigma_{i,i+1}^K) + \widehat{f}_{(i-1,N-1)}^{S,\eta} \sigma_{i-1,i}^K \sigma_{N-1,N}^K + \widehat{f}_{(i-1,N)}^{S,\eta} \sigma_{i-1,i}^K \right] \\ &\quad + \mu_1^S \left[\widehat{f}_{(i,N)}^S (1 - \sigma_{i,i-1}^1) (1 - \sigma_{N,N-1}^1) + \widehat{f}_{(i+1,N)}^S \sigma_{i+1,i}^1 (1 - \sigma_{N,N-1}^1) \right] \\ &\quad + (1 - \mu_1^S - \mu_K^S) \widehat{f}_{(i,N)}^{S,\eta} + 1_{(i_0^1, i_0^2) = (i, N)} \cdot O(\eta) \end{aligned}$$

Solving for $\widehat{f}_{(i-1,N)}^{S,\eta}$ and dividing by $\sigma_{N,N-1}^1 \widehat{f}_{(N,N)}^{S,\eta}$, this yields

$$\begin{aligned} \frac{\widehat{f}_{(i-1,N)}^{S,\eta}}{\sigma_{N,N-1}^1 \widehat{f}_{(N,N)}^{S,\eta}} \sigma_{i-1,i}^{K,\eta} \mu_K^S &= \frac{\widehat{f}_{(i,N)}^{S,\eta}}{\sigma_{N,N-1}^1 \widehat{f}_{(N,N)}^{S,\eta}} \left(\frac{\eta}{1 - \eta} + \sigma_{i,i-1}^{1,\eta} \mu_1^S + \sigma_{i,i+1}^{K,\eta} \mu_K^S \right) \\ &\quad - \frac{\widehat{f}_{(i+1,N)}^{S,\eta}}{\sigma_{N,N-1}^1 \widehat{f}_{(N,N)}^{S,\eta}} \sigma_{i+1,i}^{1,\eta} (1 - \sigma_{N,N-1}^1) \mu_1^S \\ &\quad + \frac{\widehat{f}_{(i,N)}^{S,\eta}}{\widehat{f}_{(N,N)}^{S,\eta}} (1 - \sigma_{i,i-1}^{1,\eta}) \mu_1^S - \frac{\widehat{f}_{(i-1,N-1)}^{S,\eta}}{\sigma_{N,N-1}^1 \widehat{f}_{(N,N)}^S} \sigma_{i-1,i}^{K,\eta} \sigma_{N-1,N}^{K,\eta} \mu_K^S \\ &\quad + 1_{(i_0^1, i_0^2) = (i, N)} \cdot \frac{O(\eta)}{\sigma_{N,N-1}^1 \widehat{f}_{(N,N)}^{S,\eta}} \end{aligned}$$

Recall from Lemma 5 (iv) that $\lim_{\eta \rightarrow 0} \frac{f_i^{S,\eta}}{f_N^{S,\eta}} = 0 \forall i \notin \{1, N\}$: this implies that $\lim_{\eta \rightarrow 0} \frac{\widehat{f}_{(i,N)}^{S,\eta}}{\widehat{f}_{(N,N)}^{S,\eta}} = \lim_{\eta \rightarrow 0} \frac{\widehat{f}_{(i-1,N-1)}^{S,\eta}}{\widehat{f}_{(i-1,N)}^{S,\eta}} = 0$, and therefore that the third line of this expression goes to zero (as $\eta \rightarrow 0$) at a faster rate than the first two lines. Moreover, Theorem 3 (iii) implies that the last line goes to zero as $\eta \rightarrow 0$, and that $\lim_{\eta \rightarrow 0} \sigma_{N,N-1}^1 = 0$. Then taking limits as $\eta \rightarrow 0$, we obtain

$$\lim_{\eta \rightarrow 0} \frac{\widehat{f}_{(i-1,N)}^{S,\eta}}{\sigma_{N,N-1}^{1,\eta} \widehat{f}_{(N,N)}^{S,\eta}} = \lim_{\eta \rightarrow 0} \frac{1}{\sigma_{i-1,i}^{K,\eta}} \left[\frac{\widehat{f}_{(i,N)}^{S,\eta}}{\sigma_{N,N-1}^{1,\eta} \widehat{f}_{(N,N)}^{S,\eta}} \left(\sigma_{i,i-1}^{1,\eta} \frac{\mu_1^S}{\mu_K^S} + \sigma_{i,i+1}^{K,\eta} \right) - \frac{\widehat{f}_{(i+1,N)}^{S,\eta}}{\sigma_{N,N-1}^{1,\eta} \widehat{f}_{(N,N)}^{S,\eta}} \sigma_{i+1,i}^{1,\eta} \frac{\mu_1^S}{\mu_K^S} \right] \quad (\text{A23})$$

By Lemma 1 and Theorem 3 (i),(ii),(iii), the identical limiting expression holds replacing $\widehat{f}_{(j,N)}^S$ with f_j^S : that is, for any $i \geq 2$,

$$\lim_{\eta \rightarrow 0} \frac{f_{i-1}^{S,\eta}}{\sigma_{N,N-1}^{1,\eta} f_N^{S,\eta}} = \lim_{\eta \rightarrow 0} \frac{1}{\sigma_{i-1,i}^{K,\eta}} \left[\frac{f_i^{S,\eta}}{\sigma_{N,N-1}^{1,\eta} f_N^{S,\eta}} \left(\sigma_{i,i-1}^{1,\eta} \frac{\mu_1^S}{\mu_K^S} + \sigma_{i,i+1}^{K,\eta} \right) - \frac{\widehat{f}_{i+1}^{S,\eta}}{\sigma_{N,N-1}^{1,\eta} \widehat{f}_N^{S,\eta}} \sigma_{i+1,i}^{1,\eta} \frac{\mu_1^S}{\mu_K^S} \right] \quad (\text{A24})$$

It then follows that for any $i \geq 2$,

$$\lim_{\eta \rightarrow 0} \frac{\widehat{f}_{(i-1,N)}^{S,\eta}}{\sigma_{N,N-1}^{1,\eta} \widehat{f}_{(N,N)}^{S,\eta}} = \lim_{\eta \rightarrow 0} \frac{f_{i-1}^{S,\eta}}{\sigma_{N,N-1}^{1,\eta} f_N^{S,\eta}} \quad (\text{A25})$$

Finally, to show that (A22) holds: since Lemma 5 (iv) implies that $\lim_{\eta \rightarrow 0} \widehat{f}_{(i^1,i^2)}^{S,\eta} = 0$ whenever $i^1, i^2 \notin \{1, N\}$ and $\widehat{f}^{S,\eta}$ is a probability distribution over memory state pairs, we know that

$$1 = \lim_{\eta \rightarrow 0} \widehat{f}_{(1,N)}^{S,\eta} \left[1 + \frac{\widehat{f}_{(N,1)}^{S,\eta}}{\widehat{f}_{(1,N)}^{S,\eta}} + \frac{\widehat{f}_{(1,1)}^{S,\eta}}{\widehat{f}_{(1,N)}^{S,\eta}} + \frac{\widehat{f}_{(N,N)}^{S,\eta}}{\widehat{f}_{(1,N)}^{S,\eta}} \right] \quad (\text{A26})$$

Evaluating (A25) at $i = 2$ implies that $\lim_{\eta \rightarrow 0} \frac{\widehat{f}_{(N,N)}^{S,\eta}}{\widehat{f}_{(1,N)}^{S,\eta}} = \lim_{\eta \rightarrow 0} \frac{f_N^{S,\eta}}{f_1^{S,\eta}}$; by a symmetric ar-

gument, $\lim_{\eta \rightarrow 0} \frac{\widehat{f}_{(1,1)}^{S,\eta}}{\widehat{f}_{(N,1)}^{S,\eta}} = \lim_{\eta \rightarrow 0} \frac{f_1^{S,\eta}}{f_N^{S,\eta}}$; and since the limiting distribution $\lim_{\eta \rightarrow 0} \widehat{f}^{S,\eta}$ is in-

dependent of the initial state (i_0^1, i_0^2) , it must be that $\lim_{\eta \rightarrow 0} \frac{\widehat{f}_{(N,1)}^{S,\eta}}{\widehat{f}_{(1,N)}^{S,\eta}} = 1$. Substituting these

expressions into (A26) and solving for $\widehat{f}_{(1,N)}^{S,\eta}$, we obtain $\lim_{\eta \rightarrow 0} \widehat{f}_{(1,N)}^{S,\eta} = \lim_{\eta \rightarrow 0} \frac{f_1^{S,\eta} f_N^{S,\eta}}{(f_1^{S,\eta} + f_N^{S,\eta})^2}$. Then the desired expression (A22) follows from Lemma 5 (iv) and the fact that f^S is a probability distribution, which imply $\lim_{\eta \rightarrow 0} (f_1^{S,\eta} + f_N^{S,\eta}) = 1$. ■

A.6 Proof of Theorem 6

For part (i): after any sequence k^T with $T < \min\{i_0 - 1, N - i_0\}$, Theorem 3 (ii),(iv) imply (for the symmetric two-signal case) that the DM will end in state $i_0 + \Delta(k^T)$ with probability 1. By the Corollary to Theorem 3, this implies that his beliefs adjust by a factor of $\left(\frac{\rho}{1-\rho}\right)^{2\Delta(k^T)}$, which is overconfident compared to the correct adjustment $\left(\frac{\rho}{1-\rho}\right)^{\Delta(k^T)}$. For part (ii), note that since the DM's beliefs are most extreme in states 1, N , his beliefs can never adjust by a factor greater than $\left(\frac{\rho}{1-\rho}\right)^{2 \max\{i_0 - 1, N - i_0\}}$; this means that he will be underconfident in period T whenever he observes a sequence k^T with $|\Delta(k^T)| > 2 \max\{i_0 - 1, N - i_0\}$. Let K_T be a random variable representing the number of “correct” signals - i.e. K -signals if the true state is H , 1-signals if the true state is L - in a T -period sample. Then for any $T > \frac{\max\{i_0 - 1, N - i_0\}}{\rho - \frac{1}{2}}$,

$$\begin{aligned} \Pr\{\text{underconfidence at } T\} &\geq \Pr\left\{\left|K_T - \frac{T}{2}\right| > \max\{i_0 - 1, N - i_0\}\right\} \\ &= 1 - \Pr\left\{\left|\frac{K_T}{T} - \frac{1}{2}\right| < \frac{\max\{i_0 - 1, N - i_0\}}{T}\right\} \\ &> 1 - \Pr\left\{\left|\frac{K_T}{T} - \rho\right| > \left(\rho - \frac{1}{2} - \frac{\max\{i_0 - 1, N - i_0\}}{T}\right)\right\} \end{aligned}$$

Since K_T follows a binomial distribution with mean ρ , variance $N\rho(1 - \rho)$, Chebyshev's inequality then implies that this is greater than

$$1 - \frac{\rho(1 - \rho)}{T \left(\rho - \frac{1}{2} - \frac{\max\{i_0 - 1, N - i_0\}}{T}\right)^2}$$

The last term is decreasing in ρ and converges to zero as $T \rightarrow \infty$, implying the result. ■

References

- [1] Anderson, R.C., Pichert, J.W., Goetz, E.T., Schallert, D.L., Stevens, K.V., and Trollip, S.R. (1976) "Instantiation of General Terms," *Journal of Verbal Learning and Verbal Behavior*, 15, 667-679
- [2] Aumann, R., Hart, S., and Perry, M. (1997a) "The Absent-Minded Driver". *Games and Economic Behavior* 20, 102-116.
- [3] Aumann, R., Hart, S., and Perry, M. (1997a) "The Forgetful Passenter." *Games and Economic Behavior* 20, 117-120.
- [4] Battigali, P. (1997) "Dynamic Inconsistency and Imperfect Recall." *Games and Economic Behavior* 20, 31-50.
- [5] Cover, T. and Hellman, M. (1970) "Learning With Finite Memory." *Annals of Mathematical Statistics* 41, No.3, 765-782.
- [6] Cowan, Nelson (1995) "Attention and Memory: an Integrated Framework." New York, NY: Oxford University Press.
- [7] Dow, J. (1991) "Search Decisions with Limited Memory." *Review of Economic Studies* 58, 1-14.
- [8] Edwards, W. (1968) "Conservatism in Human Information Processing." Reprinted in Kahneman, Slovic, and Tversky (1982), 359-369.
- [9] Freidlin, M., and A.D. Wentzell (1984), "*Random Perturbations of Dynamical Systems*". New York: Springer Verlag.
- [10] Griffin, D. and A. Tversky (1992) "The Weighing of Evidence and the Determinants of Confidence," *Cognitive Psychology* 24, 411-435.
- [11] Gilboa, I. (1997) "A Comment on the Absent Minded Driver Paradox." *Games and Economic Behavior* 20, 25-30.

- [12] Kahneman, Slovic, and Tversky (1982) “Judgement under Uncertainty: Heuristics and Biases.” New York, NY: Cambridge University Press.
- [13] Lichtenstein, S. and Fischhoff, B. (1977) “Do those who know more also know more about how much they know? The calibration of probability judgements.” *Organizational Behavior and Human Performance*, 20, 159-183.
- [14] Lipman, B. (1995) “Information Processing and Bounded Rationality: A Survey.” *Canadian Journal of Economics* 28, 42-67
- [15] Lord, C., Lepper, M.R., and Ross, L. (1979) “Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence.” *Journal of Personality and Social Psychology*, 37, 2098-2110
- [16] Mullainathan, S. (1998) “A Memory Based Model of Bounded Rationality”, mimeo, M.I.T.
- [17] Piccione, M. and A. Rubinstein (1993) “Finite Automata Play a Repeated Extensive Game.” *Journal of Economic Theory* 61, 160-168
- [18] Piccione, M. and A. Rubinstein (1997a) “On the Interpretation of Decision Problems with Imperfect Recall.” *Games and Economic Behavior* 20, 3-24.
- [19] Piccione, M. and A. Rubinstein (1997b) “The Absent-Minded Driver’s Paradox: Synthesis and Responses.” *Games and Economic Behavior* 20, 121-130.
- [20] Rabin, M. (1998) “Psychology and Economics.” *Journal of Economic Literature* XXXVI, 11-46
- [21] Rabin, M. and J. Schrag (1999), “First Impressions Matter: A Model of Confirmatory Bias,” *Quarterly Journal of Economics* 114(1), 37-82.
- [22] Rubinstein, A. (1998) “Modeling Bounded Rationality.” M.I.T. Press.
- [23] Stokey, N. and R. Lucas (1989), “*Recursive Methods in Economic Dynamics*”, Cambridge, MA: Harvard University Press.