Bounded Memory and Biases in Information Processing*

Andrea Wilson†
Princeton University

April 29, 2003

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 must make a decision whose payoff depends on the true state. We consider an agent with a bounded memory: he has a finite set of available memory states, and a memory process specifies which memory state to go to after new information is received (as a function of the current memory state). We show that with more than three memory states, the optimal memory process involves ignoring information with probability close to one, once the decision-maker is as convinced as his finite memory will allow. As a result, the behavior which is optimal in the long run creates many of the observed information processing biases in the short run. In particular, the agent appears to display a confirmatory bias (tendency to interpret information as supporting his current beliefs), and an overconfidence/underconfidence bias (tendency to read too much into ambiguous information, too little into highly diagnostic information).

---

*I am indebted to Wolfgang Pesendorfer for extremely valuable suggestions at all stages of this project, and for the unlimited amount of time that he is willing to spend providing encouragement and advice. I also thank Faruk Gul for many helpful suggestions in the development of this paper, and Dilip Abreu and Ariel Rubinstein for helpful conversations.

†Correspondence address: Department of Economics, 001 Fisher Hall, Princeton University, Princeton, NJ 08544. Telephone number: (609) 258-5409. E-mail: awilson@princeton.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 which puts too much weight on the initial signals.[1] 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.[2] 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.[3]

A third bias is overconfidence/underconfidence: belief adjustments tend to be more extreme than those of a Bayesian after a sequence of relatively uninformative signals (overconfidence), but too conservative after a highly informative sequence of information.[4]

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 the 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 he has recalled the event in the past. The goal of this paper is to develop

[1] There is also a lot of "popular evidence" for this phenomenon; a quick internet search will locate many businesses devoted entirely to teaching people to make a good first impression.

[2] Rabin 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.

[3] See Rabin (1999) for more examples and references.

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

a more primitive model of limited memory, and to demonstrate how these biases can arise from optimal behavior.

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 take an action; the correct action 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, but is not exactly sure when the process will end (and assigns an almost equal probability to all possibilities).

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 $\mathcal{N} = \{1, .., N\}$ 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 a (which tells him what to do in the terminal period, as a function of the current 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.[5]

The basic idea of the model is that the decision-maker cannot recall all of the information that he receives, and he 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).[6] The N memory states might also correspond to a set of N different beliefs that the decision-maker can have; in this sense, the transition rule

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

[6] Note 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.

σ describes his version of Bayesian updating. Since the decision-maker must behave the same way every time he reaches a particular memory state, 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 may simply result from the optimal long-run behavior when memory is bounded.

The paper characterizes optimal behavior for agents with bounded memory. As Rubinstein and Piccione (1997a,b) point out, the correct notion of optimality is not necessarily unambiguous in decision problems with imperfect recall. In particular, there may be an incentive to deviate from the rule which would be optimal under full commitment. Rubinstein and Piccione propose an alternative solution concept - *modified multi-self consistency* - which says that the agent cannot have any incentive for single deviations from his strategy, assuming that he follows the strategy at all other information sets. For the problem considered in this paper, there is no conflict between the two solution concepts: every ex-ante optimal strategy is modified multi-self consistent, so the agent has no incentive to deviate.

Section 2 introduces the basic model, and the simple binary signal structure that will be used in Sections 3-5: There are two possible signals, $\{l, h\}$, with the probabilites of receiving signal l in state L, h in state H both equal to $\rho > \frac{1}{2}$. (Section 6 will allow for more general signal structures). We define optimality and modified multi-self consistency for the decision problem considered, and show in Theorem 1 that the ex ante optimal solution is modified multi-self consistent.

Section 3 characterizes optimal strategies when η is small. First, the states $1, ..., N$ are ordered 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. Theorem 3 contains the main result of the paper: we show that if $N \geq 3$ and η is close to zero, then optimality requires leaving the states $1, N$ with probability close to zero. In other words, once the decision-maker reaches one of his two extreme states, he ignores all information with probability near 1. (In all other memory states, he moves to the next highest state with probability 1 after receiving a high signal, h, and to the next lowest state after receiving signal l).

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, and implies that a pure strategy (switching out of the extreme states with probability 1) is not optimal. On the other hand, it is not optimal to actually get stuck in the extreme states: ignoring information is costly, in the sense that it makes each state less informative. The optimal solution is to leave states $1, N$ with a positive probability γ which goes to zero as $\eta \to 0$, 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 sequences of signals, it may appear as though he is completely unresponsive to information. However, the probability of leaving states $1, N$ before the decision problem terminates is in fact strictly positive; this makes the extreme states considerably more informative than if they were absorbing states.

In a corollary to Theorem 3, we obtain a simple formula for the payoff: In the limit as $\eta \to 0$, the optimal strategy yields an expected payoff of $\left(1 + \left(\frac{1-\rho}{\rho}\right)^{N-1}\right)^{-1}$. For comparison, a strategy which never leaves the extreme states would yield a payoff of only $\left(1 + \left(\frac{1-\rho}{\rho}\right)^{\frac{N-1}{2}}\right)^{-1}$, as if there were half as many states. Note that the expected payoff is strictly increasing in N, but remains bounded below 1 as long as N is finite, and the signals are not perfectly informative ($\rho \neq 1$). This implies that even in the limit as $\eta \to 0$, corresponding to an infinite sequence of signals, the probability of ultimately choosing the correct action is bounded below 1 for an agent with bounded memory.

We then relate the optimal finite-state memory system to some more descriptive (and psychology-based) models. One implication of Theorem 3 is that beliefs in memory state i are as if the decision-maker recalled a sequence of $(i-1)$ high signals, and $(N-i)$ low signals. Moreover, the optimal process does not contain any "jumps": after receiving an $h-$signal, the decision-maker moves from state i to state $(i+1)$ - where his beliefs are as if he replaced one of the low signals in memory with the new high signal. This is a primitive version of psychologists' "attention and memory models"[7] - which 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 get replaced by new information.

Section 4 shows how the result in Theorem 3 can create a first impressions matter/confirmatory bias. Consider an individual with three memory states, who begins in the middle state 2. The

[7] See Cowan (1995).

optimal transition rule in state 2 is to move to state 1 with probability 1 after an l-signal, and to state 3 with probability 1 after an h-signal. Once the agent 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 were even more strongly in favor of capital punishment, while Group 2 individuals became more strongly opposed: that is, both groups viewed the evidence as support for their prior beliefs. Theorem 5 shows how belief polarization can occur when memories are bounded. As an example, consider two individuals, each with four memory states. Agent 1 starts in memory state 2, and Agent 2 starts in memory state 3. Suppose that they receive the signal sequence $\{l, h, h\}$. Agent 1 will immediately move to state 1 after the l-signal: with high probability he will remain here, ignoring the subsequent h-signals. Agent 2 will move down from state 3 to state 2 after the initial l-signal, but will then move up to state 4 after the subsequent two h-signals. That is, both individuals become even more convinced of their initial beliefs after the same sequence of information.

Section 5 demonstrates that the bounded-memory agent will display an overconfidence/underconfidence bias: For short sequences of signals, his beliefs are typically too extreme, relative to a Bayesian; while after long sequences of information, his beliefs are typically too conservative. Section 6 will discuss another type of overconfidence/underconfidence, showing that if one signal is more informative than another, then the decision-maker's beliefs will appear to respond too much to the less informative signal, and too little to the more informative signal.

Finally, Section 6 provides some results for a general signal structure $\mathcal{K} = \{1, 2, .., K\}$, where signal K provides the strongest support for state H, and signal 1 provides the strongest support for state L. We show that for sufficiently small η, the agent ignores all but the two most extreme signals, 1 and K. Moreover, regardless of the asymmetry between signals 1 and K (i.e., even if signal K provides much more compelling evidence for state H than signal 1 provides for state L), the optimal memory system becomes almost symmetric as $\eta \to 0$. Theorem 7 shows that in the limit as $\eta \to 0$, the DM attains a payoff of $\left(1+\left(\frac{1-\widetilde{\rho}}{\widetilde{\rho}}\right)^{N-1}\right)^{-1}$, where $\frac{\widetilde{\rho}}{1-\widetilde{\rho}}$ is a geometric mean of the likelihood ratios of the two extreme signals, 1 and K.

Section 7 concludes; proofs are in the appendix.

2 The Model

The model considers a decision-maker who receives a signal in every period. The signal is informative about the state of the world, $S \in \{L, H\}$ which remains fixed throughout the problem but is unknown to the decision maker. We assume that the two states are ex ante equally likely. At the beginning of every period, the process terminates with probability η and the decision-maker (DM) must take an action, H or L. He obtains a payoff of 1 if his action matches the state, and 0 otherwise. With probability $(1-\eta)$, the decision problem continues. In this case, the DM receives a binary signal $s \in \{l, h\}$ which provides information about the state S. The signal is i.i.d. and symmetric, and the conditional probabilities satisfy

$$\Pr(l|L) = \Pr(h|H) = \rho > \frac{1}{2}$$

Signal h is therefore more likely in state H and signal l is more likely in state L. Section 6 will provide some results for more general signal structures.

The DM's memory is described by a set $\mathcal{N} = \{1, 2, ..., N\}$ of memory states. For simplicity, we assume that N is odd.[8] In period zero, he chooses an initial state; let $g_0(i)$ denote the probability of starting in state i ($\forall i \in \mathcal{N}$), and let $g_0 \in \Delta(\mathcal{N})$ be the vector with ith component $g_0(i)$. When the decision problem ends, the DM must choose an action as a function of his current memory state. Let $a : \mathcal{N} \to [0, 1]$ denote this action rule, where $a(i)$ is the probability of choosing action H

[8] An even number of states would require slight adjustments to most of the proofs, but would not change the main results.

in state i. Finally, whenever the DM receives a signal, he must choose which memory state to go to. This transition is described by the stationary signal processing rule $\sigma : \mathcal{N} \times K \to \Delta(\mathcal{N})$. We write $\sigma(i,s)(j)$ to denote the probability that the agent moves from memory state i to memory state j after receiving signal s. Note that the restriction to a finite number of memory states implies that the DM cannot keep track of time, nor can he perform (and recall) an exact Bayesian update after each new piece of information.

For a given transition rule σ, denote the transition probability between memory states conditional on $S \in \{L, H\}$ by $\tau_{i,j}^S$, where

$$\tau_{i,j}^H = (1-\rho)\sigma(i,l)(j) + \rho\sigma(i,h)(j)$$
$$\tau_{i,j}^L = \rho\sigma(i,l)(j) + (1-\rho)\sigma(i,h)(j)$$

Define T^s to be the $N \times N$ transition matrix with (i,j)th entry $\tau_{i,j}^S$.

Let $g_t(i|S)$ denote the probability that the decision-maker will be in memory state i at the start of period t, conditional on S. Note that this notation suppresses dependence on the choice of (σ, g_0). Let $g_t^S \in \Delta(\mathcal{N})$ be the vector with ith component $g_t(i|S)$. Then,

$$g_t^S = g_{t-1}^S T^S = g_0 \left(T^S\right)^t$$

Finally, let f_i^S denote the probability of ending the decision problem in memory state i, conditional on $S \in \{L, H\}$, and let $f^S \in \Delta(\mathcal{N})$ be the vector with ith component f_i^S. This is given by

$$f^S \equiv \sum_{t=0}^{\infty} \eta(1-\eta)^t g_t^S = \lim_{T\to\infty} \left(\eta g_0 \sum_{t=0}^{T-1} \left((1-\eta)T^S\right)^t \right) \tag{1}$$

Note that f^S also describes the probability distribution over memory states when the DM does not know how many periods have been played. Lemma 1 in the appendix argues that the sum on the RHS converges, yielding a unique distribution f^S.

The payoff of the decision-maker, as a function of the strategy (g_0, σ, a), is given by

$$\Pi(g_0, \sigma, a) = \frac{1}{2} \sum_{i \in \mathcal{N}} \left(f_i^H a(i) + f_i^L (1 - a(i)) \right)$$

Definition 1: (g_0, σ, a) *is an optimal memory if* $\Pi(g_0, \sigma, a) \geq \Pi(g_0', \sigma', a')$ for all (g_0', σ', a').

The decision problem defined above has imperfect recall. As Rubinstein and Piccione (1997a,b) show, such decision problems can be viewed as one-person games. Moreover, they point out that optimal strategies (in the sense defined here) are not necessarily incentive compatible: that is, the DM may have an incentive to deviate from the ex-ante optimal strategy. This does not happen in the decision problem considered here. We next define incentive compatibility, and show in Theorem 1 that all optimal strategies are indeed incentive compatible.

For a given rule (g_0, σ), we can determine the beliefs that the DM should associate with each memory state. Let $\pi(i)$ denote the probability that the DM assigns to state H in memory state i. Since the states are ex ante equally likely, we have

$$\pi(i) = \frac{f_i^H}{f_i^H + f_i^L}$$

Given the beliefs π, we can determine the agent's expected continuation payoff in any period, as follows. For all $i \in \mathcal{N}$, let v_i^S denote the expected continuation payoff starting in memory state i, and conditional on true state (of the world) S. This is given by the following recursion formula:

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

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

Then, starting in state i, the DM's expected payoff is

$$\pi(i) v_i^H + \left(1 - \pi(i)\right) v_i^L \tag{2}$$

Definition 2: (σ, a) *is incentive compatible if for all* $i, j \in \mathcal{N}$ *and* $s \in \{l, h\}$,

1. $a(i) > 0 \Rightarrow \pi(i) \geq \frac{1}{2}$, *and* $a(i) = 1$ *if strict inequality holds*
2. $\sigma(i,h)(j) > 0$ *implies that for all* $k \in \mathcal{N}$,

$$\pi(i)\rho\left(v_j^H - v_k^H\right) + \left(1 - \pi(i)\right)(1-\rho)\left(v_j^L - v_k^L\right) \geq 0$$

3. $\sigma(i,l)(j) > 0$ *implies that for all* $k \in \mathcal{N}$,

$$\pi(i)(1-\rho)\left(v_j^H - v_k^H\right) + \left(1 - \pi(i)\right)\rho\left(v_j^L - v_k^L\right) \geq 0$$

This says that the DM has no incentive to deviate from his strategy, given the payoffs and beliefs induced by the strategy. In particular, part (2) says that if $\sigma(i,h)(j) > 0$, then choosing state j must maximize the expected continuation payoff when the probability of state H is $\frac{\pi(i)\rho}{\pi(i)\rho+(1-\pi(i))(1-\rho)}$. Part (3) is the analogous incentive compatibility condition for $\sigma(i,l)(j)$, and part (1) says that the action rule must be optimal, given the probabilities $\pi(\cdot)$. This definition corresponds exactly to the definition of modified multi-self consistency in Piccione-Rubinstein (1997), which requires that the agent cannot gain by deviating at a single information set, assuming that he follows his strategy at all other information sets.

Theorem 1 states that for the problem considered in this paper, there is no conflict between the ex ante optimality and incentive compatibility.

Theorem 1: *If (g_0, σ, a) is an optimal memory rule, then (σ, a) is incentive compatible.*

This result relies on the constant termination probability, η. We interpret η as measuring the amount of information that the decision-maker expects to acquire before having to make a decision. As $\eta \to 0$, the expected number of signals goes to infinity; while if η is close to 1, then he expects to make a decision almost immediately. The effect of the termination probability is that the DM cannot figure out how long he has been in the 'game': all periods look the same to him, and his expected continuation payoff does not depend on how long he has already been playing. Rubinstein and Piccione (1997a, Proposition 3) obtained the same result, under similar conditions: they showed that if a decision problem is repeated infinitely often, and the agent has no memory of how many periods have already been played, then optimality and incentive compatibility coincide.

In a finite-horizon game, with no probability of termination before the last period, the optimal rule may depend on the number of periods remaining. This means that optimal rules are generally not incentive compatible: information sets may arise in which the DM can infer something about the number of periods left, in which case he may want to deviate from the best full-commitment strategy.

The result also relies on the fact that the decision-maker does not discount his payoffs: he only cares about the probability distribution over terminal nodes, not *when* these nodes are reached. In models with perfect recall, there is no need for a discount factor on top of a termination probability - any discounting can be incorporated into η. In the problem here, on the other hand, η is used both to "discount" the future, and to figure out the beliefs that should be associated with each memory state (i.e., to form expectations about the possible histories that led to each information set). In particular, note that if payoffs were discounted at rate δ, then the expected payoff would change to

$$\frac{1}{2}\sum_{i\in\mathcal{N}}\left[f_i^{H,\delta}a(i)+f_i^{L,\delta}(1-a(i))\right],\text{ with } f^{S,\delta}\equiv\sum_{t=0}^{\infty}\eta\delta^t(1-\eta)^t g_t^S$$

Since the decision-maker's beliefs are unchanged at $\pi(i)=\frac{f_i^H}{f_i^H+f_i^L}$, an incentive compatible solution does not maximize this expression when $\delta<1$. Essentially, payoff discounting requires using a different η in the past than in the future, which creates a distortion between optimality and incentive compatibility.

Convention: In what follows, order the memory states such $j<i$ implies $\pi(j)<\pi(i)$. Thus the decision-maker is most strongly convinced that L is true in state 1, and that H is true in state N.

3 Optimal Behavior

Theorem 2 establishes the existence of an optimal strategy, and says that the decision-maker's payoff is increasing in the number of memory states N, and decreasing in the termination probability η. We denote with $\Pi^*(\eta,N)$ the DM's optimal payoff as a function of η,N.

Theorem 2: *An optimal memory process exists. The decision-maker's expected payoff* $\Pi^*(\eta,N)$ *is continuous and strictly decreasing in* η, *increasing in* N.

Theorem 3 contains the main result of the paper. In it, we characterize behavior for small η. We show that for a memory process to be optimal, it must satisfy the following conditions: (i) the probability of leaving the two extreme states, 1 and N, goes to zero as $\eta\to 0$; (ii) at any memory state $i\in\mathcal{N}\backslash\{1,N\}$, the decision-maker moves with probability 1 to the next highest state $(i+1)$ after an h-signal, and to the next lowest state $(i-1)$ after an l-signal; (iii)the DM starts in the middle state $\frac{N+1}{2}$ with probability 1; (iv) the action rule is deterministic, choosing action L or H in the middle state, action L in all lower states, and action H in all higher states.

Let $(\widehat{g}_0, \widehat{\sigma}, \widehat{a})$denote an optimal strategy, and let $\widehat{\tau}^S_{i,j}$ be the transition probability $(i \to j)$ induced by $\widehat{\sigma}$. For $N \geq 3$ memory states, optimal behavior for small η can be described as follows:

Theorem 3: *Let $N \geq 3$. For every $\varepsilon > 0$, there exists $\overline{\eta}$ such that whenever $0 < \eta < \overline{\eta}$,*

(i) $1 > \widehat{\tau}^S_{1,1} > 1 - \varepsilon$, $1 > \widehat{\tau}^S_{N,N} > 1 - \varepsilon$;

(ii) *For all $i \in \mathcal{N} \backslash \{1, N\}$, $\widehat{\sigma}^S(i,h)(i+1) = \widehat{\sigma}^S(i,l)(i-1) = 1$;*

(iii) $\widehat{g}_0 \left(\frac{N+1}{2}\right) = 1$;

(iv) $\widehat{a}(i) = \begin{cases} 0 & \text{if } i \leq \frac{N-1}{2} \\ 1 & \text{if } i \geq \frac{N+3}{2} \end{cases}$, and $\widehat{a}\left(\frac{N+1}{2}\right) \in \{0,1\}$.

Part (i) implies that the probability of leaving states $1, N$ is positive for $\eta > 0$, but goes to zero as $\eta \to 0$. One implication is that for small η, the decision-maker will spend most of his time in these two extreme memory states. Of course, an outsider cannot observe the decision-maker's memory state: this means that from his point of view, the decision-maker is simply unresponsive to information most of the time. This result will generate the confirmatory and overconfidence/underconfidence biases discussed in the subsequent sections. However, it is important to note that the probability of leaving the extreme states is strictly positive for $\eta > 0$; in fact, the probability of moving from state 1 to N (or vice versa) before the end of the decison problem goes to 1 as $\eta \to 0$. If the decision-maker actually got stuck in the extreme states, then they would become much less informative - reducing the expected payoff to as if there were only $\frac{N-1}{2}$ memory states. To see this, note that if states $1, N$ are both absorbing, then the limiting payoff conditional on H is simply the probability of reaching state N before state 1. This can be calculated by comparison with a classical "gambler's ruin" problem: Suppose that the gambler starts in state $\frac{N+1}{2}$, and state 1 represents "ruin"; the probability of winning each round is ρ, and the probability of losing is $(1-\rho)$. Then the probability of attaining goal N from $\frac{N+1}{2}$ before ruin is $\left(1 + \left(\frac{1-p}{p}\right)^{\frac{N-1}{2}}\right)^{-1}$ (see, for example, Billingsly p.93).

The proof of Theorem 3 also calculates bounds on the maximized expected payoff $\Pi^*(\eta, N)$, and on the beliefs associated with an optimal strategy. We obtain the following corollary:

Corollary: *For any $\varepsilon > 0$, there exists $\overline{\eta}$ such that whenever $0 < \eta < \overline{\eta}$,*

(i) $\left(1 + \left(\frac{1-\rho}{\rho}\right)^{(N-1)}\right)^{-1} - \varepsilon < \Pi^*(\eta, N) < \left(1 + \left(\frac{1-\rho}{\rho}\right)^{(N-1)}\right)^{-1}$

(ii) $\left| \frac{f_i^H}{f_i^L} - \left(\frac{\rho}{1-\rho}\right)^{(i-1)} \left(\frac{1-\rho}{\rho}\right)^{(N-i)} \right| < \varepsilon$.

The limiting payoff expression is the probability that the DM will correctly identify the state after an infinite sequence of signals. Part (i) implies that this is strictly below 1 when N is finite, but goes to 1 in the limit as $N \to \infty$. Also note that the expected payoff is strictly increasing in the information content of the signal: a completely uninformative signal ($\rho = \frac{1}{2}$) yields an expected payoff of $\frac{1}{2}$ for all N, while the expected payoff goes to 1 as the signal becomes perfectly informative ($\rho \to 1$).

Part (ii) implies that the relative likelihood of state H is finite and bounded away from 0 in all memory states (provided that N is finite, and that ρ is bounded away from one). This means that although the DM ignores information with high probability in states $1, N$, his probability assessments $\pi(1), \pi(N)$ are bounded away from 0,1 (respectively). Therefore, the key difference between the bounded memory agent and a Bayesian is that the bounded memory DM starts to ignore information more quickly. In particular, note that a Bayesian will also eventually become unresponsive to information, but only when he is virtually certain that his beliefs are correct. In contrast, the bounded memory agent ignores information with high probability once he is as confident as his memory will allow - even though his beliefs are far from certain.

Part (ii) also implies that in the limit as $\eta \to 0$, the beliefs in state i satisfy $\frac{f_i^H}{f_i^L} = \left(\frac{\rho}{1-\rho}\right)^{(i-1)} \left(\frac{1-\rho}{\rho}\right)^{(N-i)}$. Note that these are exactly the Bayesian beliefs for a sequence containing $(i-1)$ h-signals, and $(N-i)$ l-signals. Moreover, recall from Theorem 3 part (ii) that the optimal transition rule $\widehat{\sigma}$ does not have any jumps: the decision-maker moves to the next highest memory state after an h-signal, and to the next lowest memory state after an l-signal. This means that in memory state i, the DM behaves as if he remembers $(i-1)$ h-signals, and $(N-i)$ l-signals; after receiving a new h-signal, he moves up to state $(i+1)$ - where his new beliefs are as if he replaced one of the l-signals in memory with the new h-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 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.[9] This seems like a very intuitive notion. For example, consider a student studying for an exam. He will typically go over material repeatedly, while trying to block out all other distractions: that is, choose which facts to store in memory. If he continues to use the material after the exam, then he will continue to remember it. If not, then he will quickly forget everything - replacing the information with something more important. There is also considerable evidence that people are much better at recalling events that they have actually experienced, and that they typically store the event as a whole - rather than extracting a particular piece of information to remember. For example, Anderson et al (1976) describe an experiment in which one group of individuals memorized a list of words while on land, and the other group memorized the list while under water. Then they were asked to try to recall the list - both while on land, and while in water. The first group performed much better when on land again, while the second group did much better in water. For the model studied here, this would suggest that individuals are largely restricted to remembering signals as they are received - rather than just extracting the useful part of the information (e.g., rather than performing and recalling an exact Bayesian update).

To provide an intuition for Theorem 3, we now sketch the proof using a simple 3-state example, with $\mathcal{N} = \{1, 2, 3\}$.

Recall that $g_t(\cdot|S)$ is the probability distribution over memory states at the start of period t. This can be described recursively by

$$g_t(i|S) = (1-\eta) \sum_{j \in \mathcal{N}} g_{t-1}(j|S) \tau^S_{j,i} \tag{3}$$

Lemma 1 in the appendix shows that f^S is equal to the steady-state distribution for this process, $g(\cdot|S)$, and therefore solves $f^S_i = (1-\eta) \sum_j f^S_j \tau^S_{j,i}$.

[9] See Cowan (1995) for a discussion of "attention and memory" models.

Suppose first that $\tau^S_{1,3} = \tau^S_{3,1} = 0$; then this yields

$$\frac{f_3^H}{f_3^L} = \frac{\tau^H_{2,3}}{\tau^L_{2,3}} \left(\frac{\eta + (1-\eta)\tau^L_{3,2}}{\eta + (1-\eta)\tau^H_{3,2}} \right) \frac{f_2^H}{f_2^L} \tag{4}$$

Note that for fixed $\frac{f_2^H}{f_2^L}$, the likelihood ratio $\frac{f_3^H}{f_3^L}$ is maximized if $\frac{\tau^H_{2,3}}{\tau^L_{2,3}} = \frac{\tau^L_{3,2}}{\tau^H_{3,2}} = \frac{\rho}{1-\rho}$; this suggests that to make the states as informative as possible, the DM should only switch to higher states after an h-signal, and to lower states after an l-signal. Lemma 2 argues that this must be true of any optimal rule $\widehat{\sigma}$.

Since the states are ordered such that $\frac{f_1^H}{f_1^L} < \frac{f_2^H}{f_2^L} < \frac{f_3^H}{f_3^L}$, the probability of error is the highest in state 2 (where the DM has the least information). This implies that the DM would like to minimize the probability of ending the decision problem in state 2, while making states 1,3 as informative as possible. Note first that if the only objective was to minimize the probability of ending in state 2, then the DM would set $\tau^S_{1,2} = \tau^S_{3,2} = 0$ - that is, never leave the extreme states once he arrives. However, ignoring all information is very costly, in that it substantially reduces the informativeness of states 1 and 3. For example, consider equation (4); if $\tau^S_{3,2} = 0$, then $\frac{f_3^H}{f_3^L} = \frac{\rho}{1-\rho}\frac{f_2^H}{f_2^L}$ $\forall$ η; while any positive value for $\tau^S_{3,2}$ would yield $\lim_{\eta \to 0} \frac{f_3^H}{f_3^L} = \left(\frac{\rho}{1-\rho}\right)^2 \frac{f_2^H}{f_2^L}$, clearly the upper bound on $\frac{f_3^H}{f_3^L}$ relative to $\frac{f_2^H}{f_2^L}$. This suggests that the optimal solution requires randomization 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 $\eta \to 0$, the DM does not expect to make a decision for a long time: this makes it *almost* costless to ignore information *almost* all the time. In particular, suppose that $\lim_{\eta \to 0} \tau^S_{3,2} = \lim_{\eta \to 0} \tau^S_{1,2} = 0$, but $\lim_{\eta \to 0} \frac{\eta}{\tau^S_{3,2}} = \lim_{\eta \to 0} \frac{\eta}{\sigma^S_{1,2}} = 0$; so the probability of leaving the extreme states 1,3 goes to zero as $\eta \to 0$, but at a much slower rate. Then $f_2^S \to 0$ as $\eta \to 0$, while at the same time, equation (4) shows that the upper bound $\lim_{\eta \to 0} \frac{f_3^H}{f_3^L} = \frac{f_2^H}{f_2^L}\left(\frac{\rho}{1-\rho}\right)^2$ is still achieved in the limit. Using a similar argument for $\tau^S_{1,2}$, we conclude that the optimal solution is to switch out of the extreme states 1,3 with a positive probability that goes to zero as $\eta \to 0$, but at a slower rate; this reduces the probability of making a decision in state 2 to almost zero, without affecting the limiting informativeness of the states.

This is argued more generally in the appendix. Lemma 5 calculates an upper bound on the expected payoff, and Lemma 8 shows that this bound can only be achieved if the probability of leaving states $1, N$ goes to zero as $\eta \to 0$, but at a slower rate. Lemma 7 then shows that the optimal solution does not contain any "jumps"; that is, each state moves only to adjacent states after new information is received. This maximizes the number of h-signals, l-signals (respectively) required to reach states $N, 1$, making them as informative as possible. (And since the probability of making a decision in state 1 or N is close to 1 for small η, this maximizes the expected payoff).

To see that the optimal solution is state-symmetric, solve (3) for the limiting distribution f^S, to obtain

$$\lim_{\eta \to 0} \frac{f_3^S}{f_1^S} = \lim_{\eta \to 0} \left(\frac{\tau_{2,3}^S}{\tau_{2,1}^S} \right) \left(\frac{\eta + (1-\eta)\tau_{1,2}^S}{\eta + (1-\eta)\tau_{3,2}^S} \right) = \begin{cases} \frac{\sigma(1,h)(2)}{\sigma(3,l)(2)} \left(\frac{\rho}{1-\rho} \right)^2 & \text{if } S = H \\ \frac{\sigma(1,h)(2)}{\sigma(3,l)(2)} \left(\frac{1-\rho}{\rho} \right)^2 & \text{if } S = L \end{cases} \tag{5}$$

Define $\alpha^* \equiv \frac{\sigma(1,h)(2)}{\sigma(3,l)(2)}$, and recall from above that $\lim_{\eta \to 0} f_2^S = 0$ in an optimal solution. Then the DM's expected payoff is

$$\frac{1}{2}\left[f_3^H + f_1^L\right] = \frac{1}{2}\left[\frac{f_3^H}{f_1^H + f_3^H} + \frac{f_1^L}{f_1^L + f_3^L}\right] = \frac{1}{2}\left[\frac{1}{1 + \frac{1}{\alpha^*}\left(\frac{1-\rho}{\rho}\right)^2} + \frac{1}{1 + \alpha^*\left(\frac{1-\rho}{\rho}\right)^2}\right]$$

So as $\alpha^* \to 0$, the DM obtains an expected payoff (in the limit as $\eta \to 0$) of 0 in state H, and 1 in state L.[10] Similarly, a solution with $\alpha^* \to \infty$ would yield a limiting payoff of 1 in state H, 0 in state L. However, given that the prior is $\frac{1}{2}$, it is easily verified that the expression is maximized at $\alpha^* = 1$; this is shown more generally in Lemma 4.

The intuition for part (iii) is quite straightforward: Since the optimal strategy is symmetric in the limit as $\eta \to 0$, beliefs in the middle state, $\pi\left(\frac{N+1}{2}\right)$, are very close to the prior $\frac{1}{2}$. Incentive compatibility requires that state $\frac{N+1}{2}$ maximize the expected continuation payoff at the beliefs $\pi\left(\frac{N+1}{2}\right)$, implying that state $\frac{N+1}{2}$ also maximizes the ex ante expected payoff when H, L are equally likely.

For part (iv), it is clear that if the states are ordered such $\pi(i) < \frac{1}{2}$ for $i < \frac{N+1}{2}$, and $\pi(i) > \frac{1}{2}$ for $i > \frac{N+1}{2}$, then the DM must set $a(i) = 1$ whenever $i > \frac{N+1}{2}$, and $a(i) = 0$ whenever $i < \frac{N+1}{2}$. Lemma 9 shows that randomization in the middle state is not optimal, implying that there are two

[10] Note that this does not imply a perfectly informative memory state: $\frac{f_3^H}{f_3^L}, \frac{f_1^H}{f_1^L}$ remain bounded away from 0 and ∞. However, if the probability of leaving state 1, relative to the probability of leaving state 3, is almost zero, then the DM ends up in memory state 1 with probability close to 1. This implies that conditional on L, he is almost always right; conditional on H, he is almost always wrong.

optimal solutions: one with $a\left(\frac{N+1}{2}\right)=1$, and one with $a\left(\frac{N+1}{2}\right)=0$. In the limit as $\eta\to 0$, the associated optimal transition rules converge; moreover, any asymmetry in the prior or the signals would yield a unique equilibrium.

Finally, note that the Theorem assumes $N\geq 3$. For $N=2$ there is no bad middle state to avoid, so the optimal solution switches out of both states with probability 1.

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:[11]

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.

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 (i) 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 he does not ignore evidence in the intermediate states $\mathcal{N}\backslash\{1,N\}$, this implies that an initial string of h-signals can make it appear (in the short run) as though he's convinced that state H is true. Moreover, as his initial position gets closer to state N (so his initial impression favors H more strongly), it becomes more likely that he will receive a sequence of information which contains enough h-signals to take him to state N.

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

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. Fix an optimal strategy $(\widehat{g}_0, \widehat{\sigma}, \widehat{a})$, and choose $\widehat{\eta}$ small enough that the characterization in Theorem 3 holds for $\eta < \widehat{\eta}$. Let i_t be a random variable describing the DM's memory state in period t, so $\Pr\{i_t = i|S\} \equiv g_t(i|S)$ (where $g_t(\cdot|S)$ is as defined in Section 2). Let s_t be a random variable describing the signal received in period t, and let $s^t \in \{l, h\}^t$ denote a t-period sequence of signals. For any sequence s^t, define $g_t(i|s^t) \equiv \Pr\{i_t = i|s^t\}$.

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 h^τ of τ consecutive h-signals, and $[\tau, T]$ as the set of all T-period sequences which begin with a block of τ consecutive h-signals. More precisely,

$$[T, \tau] \equiv \{s^T \in \{l, h\}^T \mid s_t = h \text{ for } t = T - \tau + 1, ..., T\}$$
$$[\tau, T] \equiv \{s^T \in \{l, h\}^T \mid s_t = h \text{ for } t = 1, 2, ..., \tau\}$$

For any sequence s^T and collection of signal sequences A, define $\Pr\{s^T|A\}$ as the probability of sequence s^T, conditional on A. Recall that $\pi(i)$ 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

$$E(\pi(i_T) \mid A) = \sum_{i \in \mathcal{N}} \pi(i) \left(\sum_{s^T \in A} g_T\left(i|s^T\right) \Pr\left\{s^T \mid A\right\} \right)$$

Clearly, there is a 1-1 map from elements of $[T, \tau]$ to elements of $[\tau, T]$ which leaves the number of h- and l-signals unchanged. Therefore, for a standard Bayesian decision-maker, $E(\pi(i_T) \mid [T, \tau]) = E(\pi(i_T) \mid [\tau, T])$. However, this is not true for the bounded-memory DM considered here. Theorem 4 says that if $\eta < \widehat{\eta}$ (where $\widehat{\eta}$, as defined above, is such that Theorem 3 holds for $\eta < \widehat{\eta}$), then for any τ, T, (i) First impressions matter in the short run: if η is small enough relative to T, then a block of τ consecutive h-signals has a larger effect on expected period T beliefs if it occurs in period 0, rather than in period $T - \tau + 1$; (ii) Last impressions matter in the long run: if T is large enough relative to η, then the block of h-signals has a larger effect on beliefs if it occurs at the end.

Theorem 4: (i) *For any $T \geq \frac{N+1}{2}$, there exists η^* such that whenever $\eta < \eta^*$, $E(\pi(i_T) \mid [\tau, T]) > E(\pi(i_T) \mid [T, \tau])$*

(ii) *For any* $\eta < \widehat{\eta}$, *there exists* T^* *such that whenever* $T > T^*$, $E(\pi(i_T) \mid [T, \tau]) > E(\pi(i_T) \mid [\tau, T])$.

The intuition for this result follows directly from Theorem 3. First consider part (i), which says that for fixed T, τ, it is possible to choose η small enough that first impressions matter. This follows from part (i) of Theorem 3, which says that for small enough η, the probability of leaving states $1, N$ can be made arbitrarily close to zero. Then for finite sequences of information, we can choose η small enough that once the DM reaches an extreme state, the probability of leaving this state before time T is arbitrarily close to zero. Note that if states $1, N$ are essentially absorbing, then after any signal sequence the DM will either get stuck in state N (if he reaches state N before 1), or get stuck in state 1 (if he reaches 1 before N), or never reach an extreme state. The proof argues that if the DM reaches state N before state 1 after a sequence $(s^{T-\tau}, h^\tau)$, where h^τ is a block of τ consecutive h-signals, then he will also reach state N before 1 if the sequence starts with the block h^τ: moving the initial state closer to N cannot reduce the probability of reaching N first. However, there are sequences $s^{T-\tau}$ such that the DM will get stuck in state N following $(h^\tau, s^{T-\tau})$, but not following $(s^{T-\tau}, h^\tau)$. For example, if $\tau \geq \frac{N-1}{2}$ (implying that any sequence which starts with the block h^τ will cause the DM to get stuck in state N, starting in state $\frac{N+1}{2}$), then this will happen for any sequence $s^{T-\tau}$ which causes the DM to get stuck in state 1, starting in state $\frac{N+1}{2}$. Therefore, changing the order of each sequence $(s^{T-\tau}, h^\tau) \in [T, \tau]$ to $(h^\tau, s^{T-\tau})$ cannot reduce the posterior memory state, but may cause it to increase; this implies an increase in the expected posterior beliefs.

Next, consider part (ii): this says that for fixed η, it is possible to choose T large enough that last impressions dominate first impressions. This is true by the ergodic theorem: for T sufficiently large, the probability distribution g_T over memory states is independent of the initial distribution. In particular, this means that an initial block of τ consecutive h-signals, which simply increases the initial state by τ, does not affect the expected beliefs $(T - \tau)$ periods later. However, for any fixed probability distribution over $\mathcal{N}$, adding a sequence of h-signals at the end will cause an increase in the posterior memory state (for fixed $\eta > 0$). In other words, the first impressions effect wears off after sufficiently long sequences; but the last impressions effect is independent of the sequence length.

Theorem 4 matches the stylized facts. The phrase “first impressions matter” appears everywhere, and is supported by both experimental evidence and “popular evidence”. For example, a job candidate will typically pay a lot of attention to the small details of his appearance, resume, etc., and is repeatedly told that his chances of being hired are very small if he does not make a good first impression. Rabin describes the first impressions matter/confirmatory bias in the context of a school teacher, trying to assess the competence of a student: “a teacher can often interpret a question or answer by a student as either creative or just plain stupid; he will often interpret according to his previous hypothesis about the student’s aptitude”. Since the previous hypothesis is likely based on the student’s initial performance, this creates a confirmatory bias: a few bad initial exams make it more likely that subsequent exams will be judged unfavorably, which strengthens the teacher’s beliefs in the student’s incompetence, which makes it even more likely that subsequent exams will be judged unfavorably.

However, these are short-run examples. 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 third result shows how the same sequence of information can polarize the beliefs of two different individuals, as described in the third fact 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). 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".

To illustrate how this can occur with bounded memories, consider the following experiment. There are two states of the world: state H corresponds to "capital punishment deters crime", and state L to "it does not". There are two inviduals, each with $N = 5$ memory states: state 1 contains the strongest evidence for state L, and state 5, the strongest evidence for state H. Agent 1 is currently in state $i_0^1 = 2$ (weakly opposed to capital punishment), and Agent 2 is currently in state $i_0^2 = 4$ (weakly in favor of capital punishment). The experimentor has a collection of studies on capital punishment which fall into two categories, H and L. Conditional on state H being true (capital punishment does deter crime), any study will yield a type H report with probability $\rho > \frac{1}{2}$, and a type L report with probability $1 - \rho$. Conditional on state L, any study will yield a type L report with probability ρ. Both individuals agree on the probability ρ, and can correctly identify a report as type H or L.

Each individual is given the collection of studies; he can take as much time as needed to review the reports, and will then be asked whether they favor state L or state H. The analysis below will assume that η is close to zero, as the individual can reasonably expect to collect information for a long time (perhaps forever) before being asked to make a payoff-relevant decision.

In fact, there is an equal number τ of each type of report - so the overall evidence is completely ambiguous. The individuals, of course, are unaware of this; they only know that they will receive a total of 2τ reports. The experimenter shuffles the reports to obtain a random ordering, and sends the reports to each individual in the same order.

To see how divergence can occur, suppose that $\tau = 4$, so there are $\binom{4}{2} = 6$ possible orderings of the reports. Consider the sequences $\{L, H, H, L\}$ and $\{H, L, L, H\}$. In the first sequence, Agent 1 will immediately get stuck in state 1 and stay there; Agent 2 will initially move down to state 3, but then the subsequent two high signals will cause him to end up at state 5. This means that *both* individuals view the evidence as confirming their prior beliefs. The sequence $\{H, L, L, H\}$ will similarly cause beliefs to diverge. (Actually for this small sample size, every sequence will result in partial divergence. There are two other sequences starting with L : both will move agent 1 to state 1, while agent 2 will remain at state 4 - correctly identifying the sequence as uninformative. Similarly, both sequences starting with H will move agent 2 up to state 5, while agent 1's beliefs will not change).

Although divergence becomes less likely as memory and sample sizes increase, it continues to occur with positive probability. To see this, let i_t^j be a random variable representing the "posterior" memory state after t periods, for an agent who starts in state j. As in the model, assume that reports are randomly generated according to the distribution $\Pr\{\text{type } H \text{ report}|\text{ state } H\} = \Pr\{\text{type } L \text{ report}|\text{state } L\} = \rho$; assume further that the two agents observe exactly the same sequence of reports. The following theorem says that (i)for any pair of priors which are distinct and not on the memory boundaries, there is a positive probability of divergence; (ii)the probability of divergence strictly increases as either prior moves closer to the boundary.

Theorem 5: (i)If $1 < j < k < N$ and $t \geq N - 1 - (k - j)$, then $\Pr\{i_t^j < j < k < i_t^k\} > 0$; (ii)$\Pr\{i_t^j < j < k < i_t^k\}$ is strictly decreasing in $j, (N - k)$.

Proof. For (i), note that a sequence of $(j-1)$ consecutive low signals, followed by $(j-1+N-k)$ consecutive high signals, will result in $i_t^j = 1$ and $i_t^k = N$. For (ii), note that any sequence which causes agent 1 to reach state 1 from j, will also take him to state 1 if he starts in $j - 1$; at this point, the probability of leaving state 1 is independent of the initial position. However, he is strictly more likely to reach state 1 if he starts at $j - 1$: for example, an initial sequence of $(j - 2)$ low signals, followed by long sequence of high signals, will only cause divergence starting at $j - 1$. ■

Finally, note that from an ex ante point of view, these biases becomes very unlikely as $N \to \infty$. If an agent starts with basically no information and N is very large, then he requires an almost infinite sequence of high signals to reach the extreme high memory state, N. This is infinitely more likely to occur in state H than state L. Thus, with probability near 1, any agent will eventually reach the extreme high state. For large N, remaining here is close to the unconstrained optimal (i.e. Bayesian) behavior, and beliefs are almost Bayesian.

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 this bias works in both ways: while people tend to be overconfident of relatively ambiguous information, their beliefs are typically too conservative after receiving highly diagnostic information.

The decision-maker with bounded memory will display two types of overconfidence/underconfidence: one based on sample size, and one based on the quality of the information that he receives. In this section we illustrate the first type, showing that the decision-maker will typically be overconfident after short sequences of information, and underconfident after long sequences. This matches the empirical findings: Griffin and Tversky 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 the overconfidence/underconfidence bias, consider the set of T-period sequences of signals. For any sequence $s^T \in \{l, h\}^T$, define $\Delta(s^T)$ as the number of h-signals in s^T, less the number of l-signals in s^T. Fix $\delta \geq 0$, and define $|T, \delta|$ as the set of all T-period signal sequences s^T such that $\left|\Delta(s^T)\right| = \delta$. Note that this definition only depends on the "net" information content of the sequence, without specifying whether there are more h- or l-signals.

For any s^T, define $i_T(s^T)$ as the DM's posterior memory state after the sequence s^T (which may be random); then his probability assessment to state H is $\pi\left(i_T(s^T)\right)$. For any probability assessment π, define the confidence of the assessment as

$$\left|\pi - \frac{1}{2}\right|$$

That is, confidence is increasing in π if the DM believes that H is the more likely state, and decreasing in π (increasing in $\Pr\{L\} = 1 - \pi$) if the DM believes that L is more likely. Note that after any sequence $s^T \in |T, \delta|$, the correct probability assessment depends only on $\Delta(s^T)$, and is given by

$$\pi^*\left(\Delta(s^T)\right) = \begin{cases} \frac{1}{1+\left(\frac{1-\rho}{\rho}\right)^\delta} & \text{if } \Delta(s^T) = \delta \\ \frac{1}{1+\left(\frac{\rho}{1-\rho}\right)^\delta} & \text{if } \Delta(s^T) = -\delta \end{cases}$$

Then define overconfidence/underconfidence as follows:

Definition: The DM is overconfident after $s^T \in |T, \delta|$ if $\left|\pi\left(i_T(s^T)\right) - \frac{1}{2}\right| > \left|\pi^*(\Delta\left(s^T\right)) - \frac{1}{2}\right|$.

Choose $\widehat{\eta}$ small enough that Theorem 3 holds for $\eta < \widehat{\eta}$. We have the following result for $\eta < \widehat{\eta}$:

Theorem 6: *Fix T, δ, and let $\eta < \widehat{\eta}$. For any sequence $s^T \in |T, \delta|$:*

(i) *If $T < \frac{N-1}{2}$ and $\delta > 0$, then the decision-maker is overconfident after s^T with probability 1.*

(ii) *If $T > \frac{N-1}{2}$ and $0 < \delta < N-1$, then for every $\varepsilon > 0$ there exists $\overline{\eta}$ such that whenever $\eta < \overline{\eta}$, the decision-maker is overconfident after s^T with probability $1-\varepsilon$.*

(iii) *If $\delta > N-1$, then the decision-maker is underconfident after s^T with probability 1.*

Proof: Recall from the Corollary to Theorem 3 that in the limit as $\eta \to 0$, beliefs in state i are as if the DM recalled a sequence of $(i-1)$ h-signals, and $(N-i)$ l-signals. This implies that in moving from state i to $i+1$ after an h-signal, beliefs adjust as if he had received two h-signals; and in moving from state i to $(i-1)$, beliefs adjust as if he had received two l-signals. Also note that beliefs in state N are equivalent to the beliefs of a Bayesian who observes a sequence s^T with $\Delta(s^T) = (N-1)$, and beliefs in state 1 are equivalent to those of a Bayesian who observes s^T with $-\Delta(s^T) = N-1$.

For part (i): if $T < \frac{N-1}{2}$, then the sequence contains fewer than $\frac{N-1}{2}$ of each type of signal, so it is impossible for the DM to reach either state 1 or state N. Then he ends up in state $\frac{N+1}{2} + \Delta(s^T)$ with probability 1, with the overconfident beliefs $\pi^*\left(2\Delta(s^T)\right)$. For part (iii): if $\delta > N-1$, then the correct beliefs are more confident than the beliefs in either state 1 or state N, implying underconfidence with probability 1. For part (ii): for η sufficiently small, with probability close to 1 the DM will end up in state 1, state N, or state $\frac{N+1}{2} + \Delta(s^T)$. Then for $\left|\Delta(s^T)\right| < N-1$, the confidence of his

probability assessment is at least $\left|\pi^*\left(2\Delta(s^T)\right) - \frac{1}{2}\right| > \left|\pi^*\left(\Delta(s^T)\right) - \frac{1}{2}\right|$, implying overconfidence.
■

This result states that for sequences in which the difference between the number of h- and l-signals is small, the DM is almost always overconfident (parts (i) and (ii)); while if the difference is large, then he is almost always underconfident (part (iii)). The underconfidence simply results from the fact that bounded memory implies bounded probability assessments; if the number of memory states is small relative to the informativeness of the signal sequence, then the correct beliefs are more extreme than the DM is able to accomodate. The overconfidence after short (uninformative) sequences results from the fact that beliefs adjust by too much after each signal - in particular, as if two signals had been received (in the limit as $\eta \to 0$). This overadjustment is created by the optimal long-run behavior. More precisely, recall the sketch of the proof of Theorem 3. For a 3-state example, we showed that long-run beliefs satisfy $\frac{f_3^H}{f_3^L} = \left(\frac{\rho}{1-\rho}\right)\left(\frac{\eta+(1-\eta)\tau_{3,2}^L}{\eta+(1-\eta)\tau_{3,2}^H}\right)\frac{f_2^H}{f_2^L}$, with $\frac{\tau_{3,2}^L}{\tau_{3,2}^H} = \frac{\rho}{1-\rho}$. If $\sigma(3,l)(2)$ is large relative to η, then in the limit as $\eta \to 0$ this becomes $\frac{f_3^H}{f_3^L} = \left(\frac{\rho}{1-\rho}\right)^2\frac{f_2^H}{f_2^L}$. This implies that if the DM begins in state 2, an h-signal will take him to state 3, where his beliefs adjust by a factor of $\left(\frac{\rho}{1-\rho}\right)^2$ - as if he had received two h-signals. The only way to avoid this overadjustment is to set $\tau_{3,2}^S = 0$; but this substantially reduces the long-run information content in state 3, and therefore reduces the decision-maker's expected payoff. For large N, the argument is a bit more complicated, but the intuition is the same: maximizing the long-run probability of a correct action, requires that beliefs overadjust in the short run.

Note also that when T is large, $\Delta(s^T)$ is expected to be large whenever $\rho \neq \frac{1}{2}$; therefore, the underconfidence result in part (iii) is likely to arise for large sample sizes.

Moreover, small values of $\Delta(s^T)$ are expected when T is small, or ρ is close to $\frac{1}{2}$. Theorem 6 then implies that we are most likely to see overconfidence when either the sample size is small, or the signal is not very informative. This addresses both types of overconfidence/underconfidence mentioned at the beginning of the section, and closely matches the empirical evidence: many experiments (see Edwards (1965), Lichtenstein and Fischhoff (1977)) have suggested that overconfidence is most likely to arise after small samples, in which each signal provides very little diagnostic information (in this context, where ρ is close to $\frac{1}{2}$). Lichtenstein and Fischhoff, for example, investigated the effect of task difficulty on confidence: as summarized by Griffen and Tversky (1992), "Their "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". Griffen 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). With this interpretation, the empirical results closely resemble the predictions in Theorem 6.

Section 6, which introduces a more general structure, will provide a further explanation for this second type of overconfidence/underconfidence: when signals are asymmetric, beliefs appear to adjust by too much after the relatively uninformative signals, and by too little after the more informative signals.

6 Generalized Signal Structures

This section shows that the results derived in Sections 3-5 are not special to the symmetric binary signal structure. Consider the model described in Section 2, but now assume that the set of possible signals is $\mathcal{K} = \{1, 2, ..., K\}$. Denote by μ_k^S the probability of receiving signal k in state $S \in \{L, H\}$, and order the signals such that $i < j \Rightarrow \frac{\mu_i^H}{\mu_i^L} < \frac{\mu_j^H}{\mu_j^L}$. Let $\widehat{\sigma}$ be an optimal transition rule, and $\widehat{\tau}_{i,j}^S$ the associated transition probabilities $i \to j$. As in Section 2, $\widehat{\sigma}$ is also incentive compatible.

It continues to hold that for small η, the probability of switching out of the extreme states is close to zero. The new result is that for η sufficiently small, the DM completely ignores all but the two most extreme signals. More precisely:

Theorem 7: *For every $\varepsilon > 0$, there exists $\overline{\eta}$ such that whenever $0 < \eta < \overline{\eta}$,*

(i) $1 > \widehat{\tau}_{1,1}^S > 1 - \varepsilon$, $1 > \widehat{\tau}_{N,N}^S > 1 - \varepsilon$;

(ii) *For all* $i \in \mathcal{N} \backslash \{1, N\}$ and $s \in \mathcal{K} \backslash \{1, K\}$, $\widehat{\sigma}(i, s)(i) = 1$;

(iii) *For all* $i \in \mathcal{N} \backslash \{1, N\}$, $\widehat{\sigma}(i, K)(i + 1) > 0$, $\widehat{\sigma}(i, 1)(i - 1) > 0$, and $\widehat{\tau}_{i,i+\Delta}^S = 0$ whenever $\Delta \geq 2$ or $\Delta \leq -2$;

(iv) $\widehat{g}_0\left(\frac{N+1}{2}\right) = 1$;

(v) $\widehat{a}(i) = \begin{cases} 0 & \text{if } i \leq \frac{N-1}{2} \\ 1 & \text{if } i \geq \frac{N+3}{2} \end{cases}$, and $\widehat{a}\left(\frac{N+1}{2}\right) \in \{0, 1\}$.

Moveover, in the limit as $\eta \to 0$, the optimal rule $\widehat{\sigma}$ yields a symmetric set of memory states, and a state-symmetric payoff. Defining $\widetilde{\rho}$ by $\frac{\widetilde{\rho}}{1-\widetilde{\rho}} \equiv \sqrt{\frac{\mu_1^L}{\mu_1^H}\frac{\mu_K^H}{\mu_K^L}}$, we obtain the following corollary:

Corollary: *For any* $\varepsilon > 0$, *there exists* $\overline{\eta}$ *such that whenever* $0 < \eta < \overline{\eta}$,

(i) $\left(1+\left(\frac{1-\widetilde{\rho}}{\widetilde{\rho}}\right)^{N-1}\right)^{-1} - \varepsilon < \Pi^*(\eta, N) < \left(1+\left(\frac{1-\widetilde{\rho}}{\widetilde{\rho}}\right)^{N-1}\right)^{-1}$;

(ii) *For all* $i \in \mathcal{N}$, $\left|\frac{f^H(i)}{f^L(i)} - \left(\frac{\widetilde{\rho}}{1-\widetilde{\rho}}\right)^{i-1}\left(\frac{1-\widetilde{\rho}}{\widetilde{\rho}}\right)^{N-i}\right| < \varepsilon$, *where* $\frac{\widetilde{\rho}}{1-\widetilde{\rho}} \equiv \sqrt{\frac{\mu_1^L}{\mu_1^H}\frac{\mu_K^H}{\mu_K^L}}$

The intuition for part (i) of the Theorem is exactly the same as in Theorem 3: the DM prefers to avoid intermediate states, and can do so almost costlessly when η is small (by leaving the extreme states with a probability that goes to zero, but at a slower rate than η). Part (ii) says that for small η, the DM ignores all signals other than 1 and K. To see why this is optimal, note that as the probability of being in one of the extreme states (in the terminal period) goes to 1, the optimization problem reduces to maximizing the payoff in the two extreme states. This, in turn, is accomplished by ignoring all but the two most extreme signals: this makes the extreme states as informative as possible, by requiring the most informative possible sequence of signals to reach them. (Of course, this rule would not be optimal for large η; if the DM is almost sure that the problem will end tomorrow, then he should pay attention to any information that he can get today - rather than trying to maximize his probability of being correct 10000 periods later). The intuition for parts (iii)-(v) is as in Theorem 3.

The corollary is analagous to the Corollary to Theorem 3, but replacing $\frac{\rho}{1-\rho}$ with $\frac{\widetilde{\rho}}{1-\widetilde{\rho}} \equiv \sqrt{\frac{\mu_1^L}{\mu_1^H}\frac{\mu_K^H}{\mu_K^L}}$. The proofs of Lemmas 4 and 5 basically show that $\left(\frac{\mu_1^L}{\mu_1^H}\frac{\mu_K^H}{\mu_K^L}\right)^{N-1}$ is the maximum amount of "information" that can be stored in the memory system, and that it is optimal to split this information equally among states H, L when they are ex ante equally likely; that is, the optimal payoff is state-symmetric. As an intuition for why this is implied by incentive compatibility (and hence optimality, by Theorem 1), start with the symmetric case: $\frac{\widetilde{\rho}}{1-\widetilde{\rho}} = \frac{\mu_1^L}{\mu_1^H} = \frac{\mu_K^H}{\mu_K^L}$. Since $\widehat{\tau}_{1,2}^S \in (0,1)$ and $\widehat{\tau}_{N,N-1}^S \in (0,1)$, the DM must be indifferent about moving from memory state 1 to 2 after signal K, and from memory state N to $N-1$ after signal 1. If we now increase $\frac{\mu_1^L}{\mu_1^H}$ and decrease $\frac{\mu_K^H}{\mu_K^L}$, then he will no longer be willing to move from state 1 to 2 after signal K, but will strictly prefer moving from state N to $N-1$ after signal 1. To make him indifferent, we need to make the beliefs in state N slightly stronger (so that he is more convinced that H is true), and the beliefs in state 1 slightly weaker (less convinced that L is true). This is accomplished by increasing $\sigma(N,1)(N-1)$, and decreasing $\sigma(1, K)(2)$ - which implies that on average, it requires more $K-$signals to reach state N, than $1-$signals to reach state 1. (For example, consider the extreme rule $\sigma(N, 1)(N-1) = 1$, and $\sigma(1, K)(2) = 0$. Then in state 1, the DM knows only that he received enough 1-signals to reach state 1 *once*; after this, anything could have happened, making the state quite uninformative. In state N, on the other hand, he knows that he never received enough $1-$signals to reach state 1, and that he did he did not receive a 1-signal in the most recent period; on average, this implies a relatively large number of $K-$signals).

In the limit as $\eta \to 0$, incentive compatibility requires choosing these probabilities such that the memory states are symmetric, $\pi(i) = \pi(N-i)$, and such that beliefs adjust by the same amount after a $K-$signal as after a $1-$signal. In particular, part (ii) of the Corollary implies that the increase in the relative likelihood that the DM assigns to state H (vs L) after a $K-$signal, and the increase in the relative likelihood that he assigns to state L (vs H) after a $1-$signal, are both equal to $\left(\frac{\widetilde{\rho}}{1-\widetilde{\rho}}\right)^2$. This provides another, and possibly more intuitive, explanation for the overconfidence/underconfidence bias. In the short run, it appears that the DM's beliefs adjust by relatively too much after the less informative signal, say 1 (overconfidence), and by too little after the more informative signal K (underconfidence). This is because the DM correctly infers that in the long run, it requires a larger number of $1-$signals to reach state i, than $K-$signals to reach state $(N-i)$.

7 Conclusion

This paper has demonstrated that decision-makers with limited memory will optimally display many of the observed 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 taken a very fundamental definition of bounded memory, simply that the decision-maker is restricted to a finite set of memory states, and made several specific predictions about the biases that will result. 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 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 maintains almost the maximum possible information content in the extreme memory states.

Sections 4 and 5 discuss the empirical implications of Theorem 3. The main result in Theorem 3 states 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 matters 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: the most recently received signals have the largest impact on beliefs. This matches the empirically observed behavior: first impressions 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 system 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 incorportate it into an optimal long-run memory system.

Section 6 extends the results to a more general signal structure. Theorem 7 shows that with an arbitrary set of K signals, the decision-maker will ignore all but the two most extreme signals when η is close to zero; again, this makes the two extreme states as informative as possible. It is also shown that regardless of the asymmetry in the signals, the optimal set of memory states is almost symmetric when η is close to zero. This is accomplished by switching out of the memory states which are based on weak signals (after receiving an opposing strong signal) with high probability, and out of the memory states based on strong signals with a slightly lower probability. This means that memories based on weak signals will appear to be overconfident; in fact, the decision-maker is correctly inferring that if he is still *in* the weak memory state, he must not have received any strong opposing information. A further implication is that a strong negative signal will have a much greater impact on the decison-maker's beliefs than a sequence of moderately negative signals. This may provide an explanation for why politicians generally delay disseminating bad information: they may expect that one huge bad signal (e.g., a press conference announcing every error that has been made in the last year) will have a greater impact on beliefs than a sequence of less severe bad signals (e.g., only exposing one error at a time).

The results of this paper focused on the case when η is small. As η increases, the probability of leaving the two extreme states increases; this means that the biases described will arise less frequently, but still with positive probability as long as $\eta < 1$. This also seems to match the empirical findings. For example, most of the studies which demonstrate extreme polarization and confirmatory biases have involved experiments in which individuals probably do not expect to make a payoff-relevant decision for a long time, if ever; this corresponds to η close to zero. In contrast, the "herding behavior" in financial markets suggests that individuals respond to information in basically the same way; this is a situation where η should be large, as decisions must be made very quickly, and information quickly becomes obsolete. 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 as the decision problem becomes more important.

A

A.1 Proofs of Theorems 1,4

Proof of Theorem 1: This is an almost exact adaptation of the argument of Proposition 3 in Piccione-Rubinstein (1997a). Let $Z^{\tau} = \left\{ S, i_0, (k_t, i_{t+1})_{t=0}^{\tau-1} \mid k_t \in \mathcal{K} \text{ for } t \geq 0, i_t \in \mathcal{N} \right\}$ denote the set of all $\tau-$period histories; so a typical $\tau-$period history is a sequence $z = (k_t, i_{t+1})_{t=0}^{\tau-1}$, where $k_t \in \mathcal{K}$ denotes the signal received at the start of period t, and i_t denotes the DM's memory state at the start of period t. For $i \in \mathcal{N}$, let $X_{\tau}(i) = \{z \in Z_{\tau} | i_{\tau} = i\}$ represent the set of τ-period histories which end in memory state $i_{\tau} = i$. Let ϕ denote the event that the game is terminated, and for $k \in \{\mathcal{K}, \phi\}$ and $i \in \mathcal{N}$, define $X_{\tau}(i,k) \equiv \{(z,k) | z \in X_{\tau}(i)\}$ 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 two histories z, z' and a transition rule σ, define $p_{\sigma}(z|z')$ as the probability of history z given z', according to σ. Now, fix a strategy $(\widehat{g}_0, \widehat{\sigma}, \widehat{a})$, and pick any $i^*, j^* \in \mathcal{N}$ and $k^* \in \mathcal{K}$ such that $\widehat{\sigma}(i^*, k^*)(j^*) > 0$.

For any history z, let $\delta(z)$ denote the number of occurrances of the sequence (i^*, k^*, j^*) in z. Note that for $S \in \{L, H\}$ and $i_0 \in \mathcal{N}$, $p_{\sigma}(z|S, i_0)$ can be written as $(\sigma(i^*, k^*)(j^*))^{\delta(z)}\ C_{\sigma}^{S}(z)$,where

$C^S_\sigma(z) \equiv \dfrac{p_\sigma(z|S,i_0)}{(\sigma(i^*,k^*)(j^*))^{\delta(z)}}$ is a constant which does not depend on $\sigma(i^*,k^*)(j^*)$. Then

$$\left.\frac{d\sum\limits_{z\in X(i,\phi)} p_\sigma(z|S,i_0)}{d\sigma(i^*,k^*)(j^*)}\right|_{\sigma=\widehat{\sigma}} = \sum_{z\in X(i,\phi)} \delta(z)\frac{(\widehat{\sigma}(i^*,k^*)(j))^{\delta(z)}}{\widehat{\sigma}(i^*,k^*)(j)} C^S_{\widehat{\sigma}}(z) = \sum_{z\in X(i,\phi)} \frac{\delta(z)p_{\widehat{\sigma}}(z|S,i_0)}{\widehat{\sigma}(i^*,k^*)(j)} \tag{A1}$$

Note that for any history z,

$$\delta(z)p_{\widehat{\sigma}}(z|S,i_0) = \sum_{z'\in X(i^*,k^*)} p_{\widehat{\sigma}}(z'|S,i_0)\widehat{\sigma}(i^*,k^*)(j^*)p_{\widehat{\sigma}}(z|S,(z',j^*))$$

(The summand on the RHS is zero if z' is not a subhistory of z, and otherwise is equal to $p_{\widehat{\sigma}}(z|S,i_0)$; $\delta(z)$ is the number of subhistories z' of z which end in (i^*,k^*,j^*)). Using this and summing over all histories in $X(i,\phi)$,

$$\begin{aligned}\sum_{z\in X(i,\phi)} \frac{\delta(z)p_{\widehat{\sigma}}(z|S,i_0)}{\widehat{\sigma}(i^*,k^*)(j^*)} &= \sum_{z\in X(i,\phi)}\sum_{z'\in X(i^*,k^*)} p_{\widehat{\sigma}}(z'|S,i_0)p_{\widehat{\sigma}}(z|S,(z',j^*)) \\ &= \sum_{z'\in X(i^*,k^*)} p_{\widehat{\sigma}}(z'|S,i_0) \sum_{z\in X(i,\phi)} p_{\widehat{\sigma}}(z|S,(z',j^*))\end{aligned}$$

Next, note that $\sum\limits_{z\in X(i,\phi)} p_{\widehat{\sigma}}(z|S,(z',j^*))$ is just the probability of ending in state i, conditional on $(S,(z',j^*))$; the stationarity of $\widehat{\sigma}$ and constant termination probability η imply that this does not depend on z'. Then the above expression simplifies to

$$\sum_{z\in X(i,\phi)} \frac{\delta(z)p_{\widehat{\sigma}}(z|S,i_0)}{\widehat{\sigma}(i^*,k^*)(j^*)} = \left(\sum_{z'\in X(i^*,k^*)} p_{\widehat{\sigma}}(z'|S,i_0)\right)\left(\sum_{z\in X(i,\phi)} p_{\widehat{\sigma}}(z|S,j^*)\right) \tag{A2}$$

Assume that $\widehat{g}_0$ is deterministic, with initial state i_0 (the argument is easily modified without the assumption). Then since $\sum_{z\in X(i,\phi)} p_{\widehat{\sigma}}(z|S,i_0) \equiv f^S_i$, the DM's expected payoff is

$$\Pi(\widehat{g}_0,\widehat{\sigma},\widehat{a}) = \frac{1}{2}\sum_{i\in\mathcal{N}}\left[a(i)\sum_{z\in X(i,\phi)} p_{\widehat{\sigma}}(z|H,i_0) + (1-a(i))\sum_{z\in X(i,\phi)} p_{\widehat{\sigma}}(z|L,i_0)\right]$$

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

$$\frac{1}{2}\sum_{i\in\mathcal{N}}\left[\begin{array}{c} a(i)\left(\sum_{z'\in X(i^*,k^*)} p_{\widehat{\sigma}}(z'|H,i_0)\right)\left(\sum_{z\in X(i,\phi)} p_{\widehat{\sigma}}(z|H,j^*)\right) \\ +(1-a(i))\left(\sum_{z'\in X(i^*,k^*)} p_{\widehat{\sigma}}(z'|L,i_0)\right)\left(\sum_{z\in X(i,\phi)} p_{\widehat{\sigma}}(z|L,j^*)\right)\end{array}\right] \tag{A3}$$

Recall from the text that

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

So the expression in (A3) can be written as

$$\begin{aligned}
&\frac{1}{\eta} \cdot \frac{1}{2} \sum_{i \in \mathcal{N}} \left[a(i) f_{i^*}^H \mu_{k^*}^H \left(\sum_{z \in X(i,\phi)} p_{\widehat{\sigma}}(z|H, j^*) \right) + (1-a(i)) f_{i^*}^L \mu_{k^*}^L \left(\sum_{z \in X(i,\phi)} p_{\widehat{\sigma}}(z|L, j^*) \right) \right] \\
&= \frac{1}{2\eta} \left[f_{i^*}^H \mu_{k^*}^H \left(\sum_{i \in \mathcal{N}} a(i) \sum_{z \in X(i,\phi)} p_{\widehat{\sigma}}(z|H, j^*) \right) + f_{i^*}^L \mu_{k^*}^L \left(\sum_{i \in \mathcal{N}} (1-a(i)) \sum_{z \in X(i,\phi)} p_{\widehat{\sigma}}(z|L, j^*) \right) \right] \\
&= \frac{1}{2\eta} \left[f_{i^*}^H \mu_{k^*}^H v_{j^*}^H + f_{i^*}^L \mu_{k^*}^L v_{j^*}^L \right] \qquad \text{(A4)}
\end{aligned}$$

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

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

By (A4), this says

$$f_{i^*}^H \mu_{k^*}^H \left(v_{j^*}^H - v_{j'}^H \right) + f_{i^*}^L \mu_{k^*}^L \left(v_{j^*}^L - v_{j'}^L \right) \geq 0$$

which is exactly the condition for incentive compatibility of $\widehat{\sigma}(i^*, k^*)(j^*)$, by Definition 2. Since i^*, k^*, j^* were chosen arbitrarily, this implies that any optimal strategy is incentive compatible. ■

Proof of Theorem 4: Let h^τ denote a block of τ consecutive h-signals; decompose the set $[T, \tau]$ as $\left\{ \left(s^{T-\tau}, h^\tau \right) | s^{T-\tau} \in \{l, h\}^{T-\tau} \right\}$, and the set $[\tau, T]$ as $\left\{ \left(h^\tau, s^{T-\tau} \right) | s^{T-\tau} \in \{l, h\}^{T-\tau} \right\}$.

For part (i), choose any sequence $s^T = \left(s^{T-\tau}, h^\tau \right)$. Recall that in the limit as $\eta \to 0$, $\widehat{\tau}_{1,1}^S = \widehat{\tau}_{N,N}^S = 1$. Then i_T depends deterministically on s^T, with three possible outcomes: $i_T(s^T) = 1$, $i_T(s^T) = N$, or $i_T(s^T) = \frac{N+1}{2} + \Delta(s^T)$, where $\Delta(s^T)$ is the difference between the number of h- and l-signals in sequence s^T.

Suppose first that $i_T(s^T) = N$: then either $i_T(s^{T-\tau}) = N$ (the DM moved from state $\frac{N+1}{2}$ to N after $s^{T-\tau}$ and got stuck), or $i_T(s^{T-\tau}) = \frac{N+1}{2} + \Delta(s^{T-\tau}) < N$, and $\frac{N+1}{2} + \Delta(s^{T-\tau}) + \tau \geq N$ (the DM never reached state 1 or N after $s^{T-\tau}$, kept track of $\Delta(s^{T-\tau})$, then moved to N after h^τ). Now change the order to $(h^\tau, s^{T-\tau})$, so the DM observes $s^{T-\tau}$ starting in state $\max\{\frac{N+1}{2} + \tau, N\}$. If the DM reaches state N before state 1 after $s^{T-\tau}$ starting at $\frac{N+1}{2}$, then he also reaches state N before

1 starting at $\frac{N+1}{2}+\tau$; so if the first case above holds, $i_T(s^{T-\tau})=N$, then also $i_T(h^\tau,s^{T-\tau})=N$. If the DM does not reach state 1 after $s^{T-\tau}$ starting in state $\frac{N+1}{2}$, then he also will not reach state 1 starting in state $\frac{N+1}{2}+\tau$. So if the second case above holds, then the DM never reaches state 1, and therefore again ends up in state $i_T=N$ (since $\frac{N+1}{2}+\tau+\Delta(s^{T-\tau})\geq N$). So in the limit as $\eta\to 0$, $i_T(s^{T-\tau},h^\tau)=N\Rightarrow i_T(h^\tau,s^{T-\tau})=N$.

Now, suppose that $i_T(s^{T-\tau},h^\tau)=\max\left\{1,\frac{N+1}{2}+\Delta(s^{T-\tau})+\tau\right\}<N$. Then changing the order to $(h^\tau,s^{T-\tau})$ cannot reduce i_T : this is trivially true if $i_T(s^{T-\tau},h^\tau)=1$, and if $i_T(s^{T-\tau},h^\tau)>1$, implying that $s^{T-\tau}$ does not take the DM from state $\frac{N+1}{2}$ to 1, then $s^{T-\tau}$ also does not take him from $\frac{N+1}{2}+\tau$ to 1. However, there are sequences $s^{T-\tau}$ such that $i_T(s^{T-\tau},h^\tau)<N$, but $i_T(h^\tau,s^{T-\tau})=N$: for instance, any $s^{T-\tau}$ which starts with $\max\{0,\frac{N-1}{2}-\tau\}$ consecutive h-signals, then ends with $T-\tau-\max\{0,\frac{N-1}{2}-\tau\}$ consecutive l-signals (this number is positive by the assumption $T>\frac{N-1}{2}$).

Since $\Pr[T,\tau]=\Pr[\tau,T]$ (by the assumption of conditionally independent signals), we have

$$E\left(\pi(i_T)\mid[\tau,T]\right)-E\left(\pi(i_T)\mid[T,\tau]\right)=\sum_{i\in\mathcal{N}}\pi(i)\left(\frac{\sum\limits_{s^{T-\tau}}\Pr\{s^{T-\tau}\}\cdot\left(g_T\left(i|(h^\tau,s^{T-\tau})\right)-g_T\left(i|\left(s^{T-\tau},h^\tau\right)\right)\right)}{\Pr[T,\tau]}\right)$$

In the limit as $\eta\to 0$, $g_T(\cdot|s^T)$ is deterministic; we showed above that $g_T\left(i|\left(s^{T-\tau},h^\tau\right)\right)=1$ implies that for some $j\geq i$, $g_T(j|(h^\tau,s^{T-\tau}))=1$; while there exist sequences $s^{T-\tau}\in[T,\tau]$ such that $g_T(i|(s^{T-\tau},h^\tau))=1$ for $i<N$, but $g_T\left(N|(h^\tau,s^{T-\tau})\right)=1$. Since $\pi\left(N\right)>\pi(i)$, this implies that the above expression is strictly positive. For any ε', part (i) of Theorem 3 implies that there exists η' such that for all $\eta<\eta'$, and for all T-period sequences s^T and $i\in\mathcal{N}$, $\left|g_T(i|S^T)-\lim\limits_{\eta\to 0}g_T\left(i|s^T\right)\right|<\varepsilon'$; hence, the above expression is also strictly positive for sufficiently small η.

For part (ii): fix an optimal transition rule σ, and let S be the true state of the world; define $g_t(i|S,i_0)\equiv\Pr\{i_t=i|S,i_0\}$ as the probability of memory state i after t periods, conditional on S,i_0 (as in Section 2, with i_0 instead of g_0). Fix $\eta<\widehat{\eta}$, so that Theorem 3 holds. Then an initial sequence of τ h-signals changes the initial state to $\min\left\{N,\frac{N+1}{2}+\tau\right\}$, so $E\left[\pi(i_T)\mid[\tau,T]\right]=\sum_{i\in\mathcal{N}}\pi(i)g_{T-\tau}\left(i\mid S,\min\left\{N,\frac{N+1}{2}+\tau\right\}\right)$. For sequences in $[T,\tau]$, note that if the DM ignored the final block h^τ, then his expected belief would be $E\left[\pi(i_T)\mid[T,\tau]\right]=\sum_{i\in\mathcal{N}}\pi(i)g_{T-\tau}\left(i\mid S,\frac{N+1}{2}\right)$. Since h^τ will in fact increase the final memory state with strictly positive probability for $\eta>0$, there exists some constant c^* such that $E\left[\pi(i_T)\mid[T,\tau]\right]=\sum_{i\in\mathcal{N}}\pi(i)g_{T-\tau}\left(i\mid S,\frac{N+1}{2}\right)+c^*$. Lemma

5 argues that all memory states are ergodic. Then by Theorem 11.4 of Stokey-Lucas, for any initial state i_0, the period-T distribution $g_T(\cdot|i_0)$ 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 $T > T'$, $\left|g_{T-\tau}\left(i \mid S, \max\left\{N, \frac{N+1}{2}+\tau\right\}\right) - g_{T-\tau}\left(i \mid S, \frac{N+1}{2}\right)\right| < \varepsilon'$. Therefore, for any $\varepsilon^* > 0$, there exists T^* such that whenever $T > T^*$,

$$\sum_{i\in\mathcal{N}} \pi(i) g_{T-\tau}\left(i \mid S, \frac{N+1}{2}\right) > \sum_{i\in\mathcal{N}} \pi(i) g_{T-\tau}\left(i \mid S, \max\left\{N, \frac{N+1}{2}+\tau\right\}\right) - \varepsilon^*$$

For $\varepsilon^* < c^*$, this implies $E\left[\pi(i_T) \mid S, [T,\tau]\right] > E\left[\pi(i_T) \mid S, [\tau, T]\right]$, as desired. ■

A.2 Lemmas 1-8, for Theorems 2,3 and 7

Lemmas 1-8 provide the results which will be used to prove Theorem 2, and all parts of Theorems 3,7 except part (v) of Theorem 7, and part (ii) of Theorem 3. Lemma 1 shows that the distribution f^S (defined as an infinite sum in equation (1) of the text) is well-defined and unique, by constructing an auxiliary Markov process with stationary distribution f^S; most of the subsequent proofs will use this auxiliary process. Lemmas 2 and 3 prove some basic monotonicity, concavity, and continuity properties of an optimal solution, which are used in subsequent proofs. Lemma 4 shows that the maximized payoff is state-symmetric; Lemma 5 (using Lemma 4) calculates an upper bound on the expected payoff, and constructs a strategy which achieves this bound in the limit as $\eta \to 0$. Lemmas 6, 7, and 8 provide some basic characterizations of an optimal strategy: they consider sequences of transition rules σ^η with $\eta \to 0$, and show that in order for σ^η to obtain the upper bound in Lemma 5, it must be that (Lemma 6) each state moves to both higher and lower states with a probability that is strictly positive, and goes to zero at a slower rate than η (if at all); (Lemma 7) each state moves only to adjacent states, and ignores all signals $k \notin \{1, K\}$; and (Lemma 8) the probability of leaving states $1, N$ must go to zero as $\eta \to 0$. We then prove Theorem 7 and its corollary, using the continuity established in Lemma 3 to conclude that an optimal sequence must satisfy the conditions in Lemmas 6-8.

Let Ω^S be an $N \times N$ Markov matrix, with (i,j)th element

$$\omega_{ij}^S = \eta g_0(j) + (1-\eta)\tau_{ij}^S \tag{A5}$$

Recall that T^S is the matrix of transition probabilities τ_{ij}^S. Let y_i^S denote the determinant of the matrix obtained from $\left(I-(1-\eta)T^S\right)$ by replacing the ith row with the initial distribution vector, g_0 :

$$y_i^S = \left|I-(1-\eta)T^S\right|_i \tag{A6}$$

Lemma 1: *For any (g_0, σ), the sum $f^S \equiv \lim\limits_{T\to\infty} \eta g_0 \left(\sum_{t=0}^{T-1}\left((1-\eta)T^S\right)^t\right)$ converges to a unique limit. This limit f^S is equal to the stationary distribution of a Markov process with the transition probabilities ω_{ij}^S defined in (A5), and solves (with y_i^S as defined in (A6))*

$$f_i^S = \frac{y_i^S}{\sum\limits_{j\in\mathcal{N}} y_j^S} \tag{A7}$$

Proof: From equation (1) in the text, f^S is given by

$$f^S = \lim_{T\to\infty} \eta g_0 \left(\sum_{t=0}^{T-1}\left((1-\eta)T^S\right)^t\right)$$

Now consider the Markov process with initial distribution ηg_0, and transition probabilities ω_{ij}^S. Then for a state i_0 with $g_0(i_0) > 0$, $\omega_{i,i_0}^S \geq \eta g_0(i_0)\ \forall i$; so by Theorem 11.4 in Stokey and Lucas, there exists a unique ergodic set and long-run distribution. For this process, the distribution $\widetilde{g}_t$ over states in period t can be described by the following recursion formula:

$$\widetilde{g}_0 = \eta g_0; \quad \widetilde{g}_t = \eta g_0 + \widetilde{g}_{t-1}\left((1-\eta)T^S\right) \tag{A8}$$
$$\Rightarrow \widetilde{g}_T = \eta g_0 \left(\sum_{t=0}^{T-1}\left((1-\eta)T^S\right)^t\right) + \eta g_0 \left((1-\eta)T^S\right)^T$$

Since the process converges to a stationary distribution, it must be that $\lim_{T\to\infty}\left((1-\eta)T^S\right)^T = 0$. This implies

$$\lim_{T\to\infty} \widetilde{g}_T = \eta g_0 \left(\lim_{T\to\infty}\left(\sum_{t=0}^{T-1}\left((1-\eta)T^S\right)^t\right)\right) \equiv f^S$$

Using the first line of (A8), the stationary distribution f^S solves

$$f^S = \eta g_0 \left(I-(1-\eta)T^S\right)^{-1}$$

Then by Cramer's rule,

$$f_i^S = \frac{y_i^S}{\left|I-(1-\eta)T^S\right|}$$

Summing this equation over all $i \in \mathcal{N}$, and noting that the probabilities must sum to 1, we obtain $\left|I - (1-\eta)T^S\right| = \sum_{i \in \mathcal{N}} y_i^S$, as desired. ■

For any transition rule σ, let T^S be the induced transition matrix in state S (as defined in the text), and define Ω^S as in equation (A5). Define $\widetilde{\mathcal{N}}(T)$ as the set of ergodic states for the Markov process with transition probabilities τ_{ij}^S, and $\widetilde{\mathcal{N}}(\Omega)$ as the set of ergodic states for the Markov process with transition probabilities ω_{ij}^S (as in (A5)).

Lemma 2: *If (σ, a) is incentive compatible, then for all $i, j \in \widetilde{\mathcal{N}}(\Omega)$:* [12]

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 $j < i$, then $v_i^H \geq v_j^H$, and $v_i^L \leq v_j^L$*
3. *If $j < i$, then for any signal k, $\sigma(i,k)(j) > 0 \Rightarrow \frac{\mu_k^H}{\mu_k^L} \leq 1$,* and $\sigma(j,s)(i) > 0 \Rightarrow \frac{\mu_k^H}{\mu_k^L} \geq 1$
4. $\forall i, (i-1), (i+1) \in \widetilde{\mathcal{N}}(T)$, $\dfrac{v_{i+1}^H - v_i^H}{v_i^L - v_{i+1}^L} \leq \dfrac{v_i^H - v_{i-1}^H}{v_{i-1}^L - v_i^L}$.

(For the symmetric 2-signal case, part 3 states that the DM will only switch to higher memory states after signal h, and to lower states after l. Parts 2 and 4 (respectively, monotonicity and concavity conditions on the optimal v_i's) will be used in subsequent lemmas).

Proof: Pick any two states i and j. For $k \in \mathcal{K}$, define

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

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

$$\begin{aligned} V(i) &= \eta\left(\pi(i)a(i) + (1-\pi(i))(1-a(i))\right) + (1-\eta)\sum_{k \in \mathcal{K}}\left[\pi(i)\mu_k^H \cdot v_{i_k^*}^H + (1-\pi(i))\mu_k^L \cdot v_{i_k^*}^L\right] \\ &\geq \eta\left(\pi(i)a(j) + (1-\pi(i))(1-a(j))\right) + (1-\eta)\sum_{k \in \mathcal{K}}\left[\pi(i)\mu_k^H \cdot v_{j_k^*}^H + (1-\pi(i))\mu_k^L \cdot v_{j_k^*}^L\right] \\ &= \pi(i)v_j^H + (1-\pi(i))v_j^L \end{aligned}$$

[12] If i, j are not part of this ergodic set, then they are never chosen with positive probability; hence, any behavior is incentive compatible.

This proves part (1). For part (2), order i, j such that $j < i$. By part (1),

$$\pi(i)\left(v_i^H - v_j^H\right) + (1-\pi(i))\left(v_i^L - v_j^L\right) \geq 0 \tag{A10}$$

$$-\pi(j)\left(v_i^H - v_j^H\right) - (1-\pi(j))\left(v_i^L - v_j^L\right) \geq 0 \tag{A11}$$

Adding the inequalities,

$$(\pi(i) - \pi(j)) \cdot \left[\left(v_i^H - v_j^H\right) - \left(v_i^L - v_j^L\right)\right] \geq 0$$

Since $j < i \Rightarrow \pi(j) < \pi(i)$, this implies $\left(v_i^H - v_j^H\right) \geq \left(v_i^L - v_j^L\right)$. Thus the LHS of (A10) is at most $\left(v_i^H - v_j^H\right)$, so the inequality requires $v_i^H \geq v_j^H$. Similarly the LHS of (A11) is at most $-\left(v_i^L - v_j^L\right)$, 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 $k \in \mathcal{K}$, $\sigma(i,k)(j) > 0$. By part (1),

$$\pi(i)\left(v_i^H - v_j^H\right) + (1-\pi(i))\left(v_i^L - v_j^L\right) \geq 0$$

$$\Leftrightarrow \frac{\pi(i)}{(1-\pi(i))} \frac{\left(v_i^H - v_j^H\right)}{\left(v_j^L - v_i^L\right)} \geq 1 \text{ (by part (2)}$$

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

For part (4): suppose not, so for some i, $\frac{v_{i+1}^H - v_i^H}{v_i^L - v_{i+1}^L} > \frac{v_i^H - v_{i-1}^H}{v_{i-1}^L - v_i^L}$. Since $i \in \widetilde{\mathcal{N}}(T)$, there must exist $j \in \mathcal{N}$ and $k \in \mathcal{K}$ such that $\sigma(j,k)(i) > 0$. Suppose first that $j > i$; then by part (3), incentive compatibility requires $\frac{\pi(j)}{1-\pi(j)} \frac{\mu_k^H}{\mu_k^L} \left(\frac{v_{i+1}^H - v_i^H}{v_i^L - v_{i+1}^L}\right) \leq 1$. But then the above inequality implies $\frac{\pi(j)}{1-\pi(j)} \frac{\mu_k^H}{\mu_k^L} \left(\frac{v_i^H - v_{i-1}^H}{v_{i-1}^L - v_i^L}\right) < 1$; again by part (3), this contradicts incentive compatibility of $\sigma(j,k)(i) > 0$ (the DM would prefer state $(i-1)$ to i). So, it must be that $j < i$. Then incentive compatibiltiy requires $\frac{\pi(j)}{1-\pi(j)} \frac{\mu_k^H}{\mu_k^L} \left(\frac{v_i^H - v_{i-1}^H}{v_{i-1}^L - v_i^L}\right) \geq 1$; by the above inequality, this implies $\frac{\pi(j)}{1-\pi(j)} \frac{\mu_k^H}{\mu_k^L} \left(\frac{v_{i+1}^H - v_i^H}{v_i^L - v_{i+1}^L}\right) > 1$; but then the DM would rather switch to state $(i+1)$, contradicting optimality of $\sigma(j,k)(i) > 0$. ■

Using Lemma 2 and the fact that $\frac{\pi(i)}{(1-\pi(i))} = \frac{f_i^H}{f_i^L}$, incentive compatibility (hence optimality) of

$\sigma(i,k)(j) > 0$ requires

$$i < j : \frac{f_i^H}{f_i^L}\frac{\mu_k^H}{\mu_k^L}\left(\frac{v_{\widehat{i}}^H - v_{\widehat{i}-1}^H}{v_{\widehat{i}-1}^L - v_{\widehat{i}}^L}\right) \geq 1 \ \forall \widehat{i} \in \{i+1, i+2, ..., j\} \tag{A12}$$

$$i > j : \frac{f_i^H}{f_i^L}\frac{\mu_k^H}{\mu_k^L}\left(\frac{v_{\widehat{i}+1}^H - v_{\widehat{i}}^H}{v_{\widehat{i}}^L - v_{\widehat{i}+1}^L}\right) \leq 1 \ \forall \widehat{i} \in \{j, j+1, ..., i-1\}$$

Lemma 3: *For all η, an optimal rule σ^η exists. The set of optimal rules σ^η is upper hemi-continuous in η, and the DM's maximized expected payoff $\Pi^*(\eta, N)$ is continuous and strictly decreasing in η.*

Proof: Given an action rule a, an optimal transition rule σ^η maximizes $\sum_{i\in\mathcal{N}}(1-a(i))\, f_i^L + \sum_{i\in\mathcal{N}} a(i) f_i^H$. By Lemma 1, the probability distribution f^S is equal to the limiting distribution of a Markov process with transition probabilities given by Ω^S. Thus, choosing σ (hence τ_{ij}^S) is equivalent to choosing ω_{ij}^S, subject to the constraint $\sigma(i,k)(j) \geq \eta g_0(j)$ for all signals k and states j. Note that f_i^S is a polynomial (and hence continuous) in ω_{ij}^S; 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. Finally, Lemma 2 part (3) implies that $\sigma(i,k)(j) > 0$ is not incentive compatible for all $k \in \mathcal{K}$. Then by Theorem 1, the constraint $\sigma(i,s)(j) \geq \eta g_0(j)$ $\forall s$ must be binding, establishing that payoffs are decreasing in η. ■

Lemma 4: *The DM's expected payoff is at most $\frac{1}{1+r^*}$, where r^* is a lower bound on*

$$r(\sigma) \equiv \sqrt{\frac{\sum_{i<i^*} y_i^H}{\sum_{i<i^*} y_i^L} \cdot \frac{\sum_{i\geq i^*} y_i^L}{\sum_{i\geq i^*} y_i^H}}, \text{ with } y_i^S \text{ as defined in (A6)}$$

The payoff $\frac{1}{1+r(\sigma)}$ is attained iff $\frac{\sum_{i<i^*} y_i^H}{\sum_{i\geq i^*} y_i^H} = \frac{\sum_{i\geq i^*} y_i^L}{\sum_{i<i^*} y_i^L}$.

Proof: Assume that there exists $i^* \in \mathcal{N}$ such that $a(i) = 1$ for $i \geq i^*$, and $a(i) = 0$ for $i < i^*$. (The argument does not rely on this, aside from the notation). Then the expected payoff is $\frac{1}{2}\left[\sum_{i\geq i^*} f_i^H + \sum_{i<i^*} f_i^L\right]$; write the bracketed term as

$$\frac{\sum\limits_{i\geq i^*} y_i^H}{\sum\limits_j y_j^H} + \frac{\sum\limits_{i<i} y_i^L}{\sum\limits_j y_j^L} = \frac{1}{1+\frac{\sum_{i<i^*} y_i^H}{\sum_{i\geq i^*} y_i^H}} + \frac{1}{1+\frac{\sum_{i\geq i^*} y_i^L}{\sum_{i<i^*} y_i^L}}$$

Let $x \equiv \frac{\sum_{i<i^*} y_i^H}{\sum_{i\geq i^*} y_i^H}$. Then using the definition of $r(\sigma)$, the expected payoff is $\frac{1}{2}\left(\frac{1}{1+x} + \frac{1}{1+\frac{(r(\sigma))^2}{x}}\right)$. Then observe that

$$\begin{aligned}
&(1+r(\sigma))(1+x)\left(1+\frac{(r(\sigma))^2}{x}\right)\left[\frac{2}{(1+r(\sigma))} - \left(\frac{1}{(1+x)} + \frac{1}{\left(1+\frac{(r(\sigma))^2}{x}\right)}\right)\right] \\
&= \frac{1-r(\sigma)}{x}\cdot(x-r(\sigma))^2
\end{aligned}$$

Since $\frac{f_i^H}{f_i^L}$ (therefore also $\frac{y_i^H}{y_i^L}$) is increasing in i, it follows easily from the definition that $r(\sigma) < 1$. Then the above expression is non-negative, and equal to zero iff $x = r(\sigma) \Leftrightarrow \frac{\sum_{i<i^*} y_i^H}{\sum_{i\geq i^*} y_i^H} = \frac{\sum_{i\geq i^*} y_i^L}{\sum_{i<i^*} y_i^L}$. This implies that $\frac{1}{(1+x)} + \frac{1}{\left(1+\frac{(r(\sigma))^2}{x}\right)} \leq \frac{2}{1+r(\sigma)}$, with equality iff the condition in the Lemma holds. ■

Lemma 5: *There exists a sequence of strategies* σ^η *in which* $\lim_{\eta\to 0} r(\sigma^\eta) = r^* \equiv \left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H}\right)^{\frac{N-1}{2}}$, *and* $\lim_{\eta\to 0} \Pi^*(\eta, N) = \frac{1}{1+r^*}$. *This is an upper bound on the expected payoff, which is attained only if* (i) $\lim_{\eta\to 0} \frac{y_i^{S,\eta}}{y_1^S} = \lim_{\eta\to 0} \frac{y_i^S}{y_N^S} = 0$ $\forall i \in \mathcal{N}\backslash\{1, N\}$; (ii) $\lim_{\eta\to 0} \frac{y_1^{H,\eta}}{y_1^{L,\eta}} = \left(\frac{\mu_1^H}{\mu_1^L}\right)^{N-1}$, $\lim_{\eta\to 0} \frac{y_N^{H,\eta}}{y_N^{L,\eta}} = \left(\frac{\mu_K^H}{\mu_K^L}\right)^{(N-1)}$; (iii) *there exists* $\widehat{\eta}$ *such that for all* $\eta < \widehat{\eta}$, *all states are ergodic.*

Proof: Lemma 3 established that the set of optimal transition rules is upper hemi-continuous in η; this implies that there exists a convergent sequence of transition rules σ^η with $\eta \to 0$. Fix such a sequence, and define $\omega_{ij}^{S,\eta}$ as in (A5), now making explicit the dependence on η; let $\Omega^{S,\eta}$ be the associated transition matrix. Since σ^η converges, there exists $\widehat{\eta}$ such that whenever $\eta \in (0, \widehat{\eta})$, the ergodic set $\widetilde{\mathcal{N}}(\Omega^\eta)$ is constant. Fix such an $\widehat{\eta}$, and let $\widetilde{\mathcal{N}}$ be the ergodic set.

Using the characterization of Markov processes given in Freidlin and Wentzell (see also Kandori, Mailath, Rob (1993)), $y_j^{S,\eta}$ is given by

$$y_j^{S,\eta} = \sum_{q\in Q_j} \prod_{(i\to i')\in q} \omega_{i,i'}^{S,\eta} \tag{A13}$$

where Q_j is the set of all $j-$trees, and a $j-$tree is a directed graph on $\widetilde{\mathcal{N}}$ such that (i)each state except j has a unique successor; (ii)there are no closed loops. Suppose that $\widetilde{\mathcal{N}} = \{1, 2, ..., N\}$ (without loss of generality; if $\widetilde{\mathcal{N}}$ is strictly contained in $\mathcal{N}$, renumber the states such that $\widetilde{\mathcal{N}} =$

$\{1,2,...,N'\}$, where $N' \equiv \#\widetilde{\mathcal{N}}'$). Now, consider the following strategy:

$$\tau_{12}^{S,\eta} = \sqrt{\eta}\mu_K^S;\ \tau_{N,N-1}^{S,\eta} = \sqrt{\eta}\left(\frac{\mu_K^H\mu_K^L}{\mu_1^H\mu_1^L}\right)^{\frac{N-1}{2}}\mu_1^S;\ \text{for } i \in \widetilde{\mathcal{N}}\backslash\{1,N\},\ \tau_{i,i+\Delta}^{S,\eta} = \begin{cases} \mu_K^S & \text{if } \Delta = 1 \\ \mu_1^S & \text{if } \Delta = -1 \\ 0 & \text{otherwise} \end{cases}$$

Under this process,

$$\lim_{\eta\to 0}\frac{\omega_{12}^{S,\eta}}{\sqrt{\eta}} = \mu_K^S;\ \lim_{\eta\to 0}\frac{\omega_{N,N-1}^{S,\eta}}{\sqrt{\eta}} = \left(\frac{\mu_K^H\mu_K^L}{\mu_1^H\mu_1^L}\right)^{\frac{N-1}{2}}\mu_1^S;\ \text{for } i \in \widetilde{\mathcal{N}}\backslash\{1,N\}, \lim_{\eta\to 0}\omega_{i,j}^{S,\eta} = \begin{cases} \mu_K^S & \text{if } j = i+1 \\ \mu_1^S & \text{if } j = i-1 \\ 0 & \text{otherwise} \end{cases}$$

Note that if $\omega_{ij}^{S,\eta} = 0$ whenever $j \notin \{i-1, i+1\}$, while $\omega_{i,i+1}^{S,\eta} > 0$ and $\omega_{i,i-1}^{S,\eta} > 0$, then $\forall j \in \widetilde{\mathcal{N}}$ there is exactly one j-tree - namely, in which each state $i < j$ goes to $i+1$, and each state $i > j$ goes to $i-1$. Then in the process just described,

$$\lim_{\eta\to 0}\frac{y_1^{S,\eta}}{\sqrt{\eta}} = \lim_{\eta\to 0}\frac{\omega_{N,N-1}^{S,\eta}}{\sqrt{\eta}}\lim_{\eta\to 0}\prod_{i\in N\backslash\{1,N\}}\omega_{i,i-1}^{S,\eta} = \left(\mu_1^S\right)^{N-1}\left(\frac{\mu_K^H\mu_K^L}{\mu_1^H\mu_1^L}\right)^{\frac{N-1}{2}}$$

Similarly, $\lim_{\eta\to 0}\frac{y_N^{S,\eta}}{\sqrt{\eta}} = \left(\mu_K^S\right)^{N-1}$; and for $i \in \widetilde{\mathcal{N}}\backslash\{1,N\}$, $\lim_{\eta\to 0}\dfrac{y_i^{S,\eta}}{\sqrt{\eta}} = 0$ (since both $\omega_{1,2}^{S,\eta}$ and $\omega_{N,N-1}^{S,\eta}$ go to zero at rate $\sqrt{\eta}$). Therefore

$$\lim_{\eta\to 0}\frac{\sum_{i<i^*} y_i^{H,\eta}}{\sum_{i<i^*} y_i^{L,\eta}} = \lim_{\eta\to 0}\frac{\sum_{i<i^*}\frac{y_i^{H,\eta}}{\sqrt{\eta}}}{\sum_{i<i^*}\frac{y_i^{L,\eta}}{\sqrt{\eta}}} = \lim_{\eta\to 0}\frac{\frac{y_1^{H,\eta}}{\sqrt{\eta}}}{\frac{y_1^{L,\eta}}{\sqrt{\eta}}} = \left(\frac{\mu_1^H}{\mu_1^L}\right)^{N-1},\ \text{and}\ \lim_{\eta\to 0}\frac{\sum_{i\geq i^*} y_i^{L,\eta}}{\sum_{i\geq i^*} y_i^{H,\eta}} = \left(\frac{\mu_K^L}{\mu_K^H}\right)^{N-1}$$

$$\Rightarrow \lim_{\eta\to 0} r(\sigma) = \lim_{\eta\to 0}\sqrt{\left(\frac{\sum_{i<i^*} y_i^{H,\eta}}{\sum_{i<i^*} y_i^{L,\eta}}\right)\left(\frac{\sum_{i\geq i^*} y_i^{L,\eta}}{\sum_{i\geq i^*} y_i^{H,\eta}}\right)} = r^*$$

Furthermore, note that $\lim_{\eta\to 0}\frac{\sum_{i<i^*} y_i^{H,\eta}}{\sum_{i\geq i^*} y_i^{H,\eta}} = \lim_{\eta\to 0}\frac{y_1^{H,\eta}/\sqrt{\eta}}{y_N^{H,\eta}/\sqrt{\eta}} = \left(\frac{\mu_1^H\mu_K^L}{\mu_K^H\mu_1^L}\right)^{\frac{N-1}{2}} = \lim_{\eta\to 0}\frac{\sum_{i\geq i^*} y_i^{L,\eta}}{\sum_{i<i^*} y_i^{L,\eta}}$. So by Lemma 4, this implies that σ^η attains the limiting payoff $\frac{1}{1+r^*}$.

To see that this is the upper payoff bound, we just need to show that r^* is a lower bound on $r(\sigma)$. To see this, note that each k−tree is a product of $(N-1)$ entries $\omega_{i,j}^{S,\eta} \equiv \eta g_0(j) + (1-\eta)\tau_{i,j}^{S,\eta}$. Since $\frac{\omega_{i,j}^{H,\eta}}{\omega_{i,j}^{L,\eta}} \geq \frac{\mu_1^H}{\mu_1^L}$ and $\frac{\omega_{i,j}^{L,\eta}}{\omega_{i,j}^{H,\eta}} \geq \frac{\mu_K^L}{\mu_K^H}$ and likelihood ratios $\frac{y_i^{H,\eta}}{y_i^{L,\eta}}$ are increasing in i, it is then clear from (A13) that $\forall j \in \widetilde{\mathcal{N}}$, $\frac{y_j^{H,\eta}}{y_j^{L,\eta}} \geq \frac{y_1^{H,\eta}}{y_1^{L,\eta}} \geq \left(\frac{\mu_1^H}{\mu_1^L}\right)^{N-1}$, $\frac{y_j^{L,\eta}}{y_j^{H,\eta}} \geq \frac{y_N^{L,\eta}}{y_N^{H,\eta}} \geq \left(\frac{\mu_K^L}{\mu_K^H}\right)^{N-1}$; this implies $\lim_{\eta\to 0} r(\sigma^\eta) \geq \sqrt{\frac{y_1^H}{y_1^L}\frac{y_N^L}{y_N^H}} \geq r^*$. This inequality is strict unless (i) for all $i \in \widetilde{\mathcal{N}}\backslash\{1,N\}$, $\lim_{\eta\to 0}\frac{y_i^{S,\eta}}{y_1^S} = \lim_{\eta\to 0}\frac{y_i^S}{y_N^S} = 0$;

(ii) $\lim_{\eta\to 0} \frac{y_1^H}{y_1^L} = \left(\frac{\mu_1^H}{\mu_1^L}\right)^{N-1}$ and $\lim_{\eta\to 0} \frac{y_N^H}{y_N^L} = \left(\frac{\mu_K^H}{\mu_K^L}\right)^{N-1}$; this establishes conditions (i) and (ii) of the Lemma. Finally, note that if $\widetilde{\mathcal{N}}$ is strictly contained in $\mathcal{N}$, then the upper bound for $\lim_{\eta\to 0} r(\sigma^\eta)$ is as above, but replacing N with $N' < N$. Since $\left(\frac{\mu_1^H \mu_K^L}{\mu_1^L \mu_K^H}\right) < 1$, this implies a strictly lower payoff. So r^* can only be attained if $\widetilde{\mathcal{N}} = \mathcal{N}$, establishing (iii). ■

Lemma 6: *Fix a convergent sequence σ^η. The bound $\lim_{\eta\to 0} r(\sigma^\eta) = r^*$ is attained only if* $\lim_{\eta\to 0} \frac{\eta}{\sum_{j>i} \omega_{i,j}^{S,\eta}} = \frac{\eta}{\sum_{j<i} \omega_{i,j}^{S,\eta}} = 0 \;\forall i \in \mathcal{N}$.

Proof: Suppose that the condition does not hold: there exist $i \in \mathcal{N}$ and $\delta^S > 0$ such that $\lim_{\eta\to 0} \frac{\eta}{\sum_{j>i} \omega_{i,j}^{S,\eta}} = \delta^S$. By the definition of $\omega_{ij}^{S,\eta}$, this is equivalent to requiring that the limit $\lim_{\eta\to 0} \frac{\tau_{ij}^{S,\eta}}{\eta}$ exists for all $j \in \mathcal{N}$. If there is no $N-$tree q_N^* with $\lim_{\eta\to 0} \prod_{(i\to j)\in q_N} \frac{\omega_{i,j}^{L,\eta}}{\omega_{i,j}^{H,\eta}} = \left(\frac{\mu_K^L}{\mu_K^H}\right)^{N-1}$, then (A13) implies that $\frac{y_N^L}{y_N^H}$ is bounded above $\left(\frac{\mu_K^L}{\mu_K^H}\right)^{N-1}$, so by Lemma 5 (ii) we are done. So, assume that there is such q_N^*; if there are several, then choose one which goes to zero at the slowest rate. Note that $\lim_{\eta\to 0} \prod_{(i\to j)\in q_N^*} \frac{\omega_{i,j}^{L,\eta}}{\omega_{i,j}^{H,\eta}} = \left(\frac{\mu_K^L}{\mu_K^H}\right)^{N-1}$ requires $\lim_{\eta\to 0} \omega_{i,j}^{S,\eta} = \mu_K^S$ for all edges $(i \to j) \in q_N^*$; by Lemma 2 part (3), this implies that all states in q_N^* switch up to higher states.

Choose an initial state i_0 (such that $g_0(i_0) > 0$). The proof will argue that there is another tree $\widehat{q}$ which goes to zero at least as slowly as q_N^*, with $\lim_{\eta\to 0} \prod_{(i\to j)\in \widehat{q}} \frac{\omega_{i,j}^{L,\eta}}{\omega_{i,j}^{H,\eta}}$ bounded above the desired limit $\left(\frac{\mu_K^L}{\mu_K^H}\right)^{N-1}$; this implies $\lim_{\eta\to 0} \frac{y_N^H}{y_N^L} > \left(\frac{\mu_K^L}{\mu_K^H}\right)^{N-1}$, so $\lim_{\eta\to 0} r(\sigma^\eta) > r^*$.

Suppose first that $i < i_0$, and let j be $i's$ successor in q_N^*. Consider the tree $\widehat{q}_N$ which is identical to q_N^*, but has the edge $(i \to i_0)$ instead of $(i \to j)$; since i_0 switches up to a higher state, this does not create any closed loops, and hence $\widehat{q}_N$ is also an $N-$tree. By hypothesis (see paragraph 1), $\lim_{\eta\to 0} \frac{\tau_{i,i_0}^{S,\eta}}{\eta}$ is finite: this implies that there exists δ' such that $\lim_{\eta\to 0} \frac{\omega_{i,i_0}^{L,\eta}}{\omega_{i,i_0}^{H,\eta}} \equiv \lim_{\eta\to 0} \frac{\eta+(1-\eta)\tau_{i,i_0}^{L,\eta}}{\eta+(1-\eta)\tau_{i,i_0}^{H,\eta}} > \frac{\mu_K^L}{\mu_K^H} + \delta'$. Thus, $\lim_{\eta\to 0} \prod_{(i\to j)\in \widehat{q}_N} \frac{\omega_{i,j}^{L,\eta}}{\omega_{i,j}^{H,\eta}}$ is bounded above $\left(\frac{\mu_K^L}{\mu_K^H}\right)^{N-1}$. Since $\omega_{i,i_0}^{S,\eta} \equiv \eta + (1-\eta)\tau_{i,i_0}^{S,\eta}$ obviously does not go to zero faster than η, while by assumption, $\omega_{i,j}^{S,\eta}$ goes to zero at least as quickly as η, it must be that $\widehat{q}_N$ goes to zero at least as slowly as q_N^*. By construction, this rate is at least as slow as any other N-tree which achieves the ratio $\left(\frac{\mu_K^L}{\mu_K^H}\right)^{N-1}$. Then by (A13), $\lim_{\eta\to 0} \frac{y_N^{L,\eta}}{y_N^{H,\eta}}$ is bounded above $\left(\frac{\mu_K^L}{\mu_K^H}\right)^{N-1}$.

Now suppose that $i > i_0$, and remove the edge $(i \to j)$ from q_N^*. There are two cases to consider: (i) if the remaining graph contains a path from i_0 to N, then it must be that q_N^* does not contain a path from i_0 to i (all states in q_N^* go up to higher states; so if there is a path from i_0 to i, then removing $(i \to j)$ breaks the path to N). In this case, there is a tree $\widehat{q}_N$ which is identical to q_N^*, but with the edge $(i \to i_0)$ instead of $(i \to j)$; at this point, the analysis is identical to the previous paragraph. (ii) If the remaining graph does not contain a path from i_0 to N, then consider the tree which is identical to q_N^*, but removes the edge $(i \to j)$, and adds the edge $(N \to i_0)$. By construction, adding the edge $(N \to i_0)$ does not create any closed loops, so we now have an i_0-tree. Since $\omega_{N,i_0}^{S,\eta}$ goes to zero at rate η or slower (as above), this tree cannot go to zero faster than q_N^*. Since $\lim_{\eta\to 0} \frac{y_{i_0}^L}{y_{i_0}^H}$ is bounded above $\left(\frac{\mu_K^L}{\mu_K^H}\right)^{N-1}$, this implies (by Lemma 4 (ii)) that $\lim_{\eta\to 0} r(\sigma^\eta)$ is bounded above r^*. The argument for $\lim_{\eta\to 0} \frac{\eta}{\sum_{j<i} \omega_{i,j}^{S,\eta}} = 0$ is identical. ∎

Lemma 7: *Let σ^η be a convergent sequence of incentive compatible transition rules, and $\tau_{i,j}^{S,\eta}$ the associated sequence of transition probabilities. Then $\lim_{\eta\to 0} r(\sigma^\eta) = r^*$ only if there exists $\overline{\eta}$ such that for all $i \in \mathcal{N}$, $\eta < \overline{\eta}$ implies* (i) $\sigma^\eta(i,k)(i) = 1$ *whenever* $k \in \mathcal{K}\backslash\{1, K\}$; (ii) $\tau_{i,i+\Delta}^{S,\eta} = \tau_{i,i-\Delta}^{S,\eta} = 0$ *whenever* $\Delta \geq 2$.

Proof: Let $\eta < \widehat{\eta}$, with $\widehat{\eta}$ as defined in Lemma 5 (iii). Then all states are ergodic, so incentive compatibility requires that (A12) hold for all $i, j \in \mathcal{N}$ and $k \in \mathcal{K}$ with $\sigma(i,k)(j) > 0$. For all $i \in \mathcal{N}\backslash\{N\}$, Lemma 6 implies that for $\eta < \widehat{\eta}$, (1) there exists $j > i$ with $\tau_{i,j}^{s,\eta} > 0$: by (A12), this requires $\frac{f_i^{H,\eta}}{f_i^{L,\eta}} \frac{\mu_K^H}{\mu_K^L} \frac{v_{i+1}^{H,\eta} - v_i^{H,\eta}}{v_i^{L,\eta} - v_{i+1}^{L,\eta}} \geq 1$; (2) there exists $j' < i+1$ such that $\tau_{i+1,j'}^{S,\eta} > 0$; by (A12), this requires $\frac{f_{i+1}^{H,\eta}}{f_{i+1}^{L,\eta}} \frac{\mu_1^H}{\mu_1^L} \frac{v_{i+1}^{H,\eta} - v_i^{H,\eta}}{v_i^{L,\eta} - v_{i+1}^{L,\eta}} \leq 1$. Combining inequalities, it must be that for all i,

$$\frac{f_i^{H,\eta}}{f_i^{L,\eta}} \frac{\mu_K^H}{\mu_K^L} \geq \frac{f_{i+1}^{H,\eta}}{f_{i+1}^{L,\eta}} \frac{\mu_1^H}{\mu_1^L} \Rightarrow \frac{y_i^{H,\eta}}{y_i^{L,\eta}} \geq \frac{y_{i+1}^{H,\eta}}{y_{i+1}^{L,\eta}} \left(\frac{\mu_1^H}{\mu_1^L} \frac{\mu_K^L}{\mu_K^H}\right) \tag{A14}$$

Fix a state i. Applying (A14) to every pair $(j, j+1)$ with $j < i$, and then to every pair $(j, j+1)$ with $j \geq i+1$, yields

$$\frac{y_1^{H,\eta}}{y_1^{L,\eta}} \geq \left(\frac{\mu_1^H}{\mu_1^L} \frac{\mu_K^L}{\mu_K^H}\right)^{i-1} \frac{y_i^{H,\eta}}{y_i^{L,\eta}}; \quad \frac{y_{i+1}^{H,\eta}}{y_{i+1}^{L,\eta}} \geq \left(\frac{\mu_1^H}{\mu_1^L} \frac{\mu_K^L}{\mu_K^H}\right)^{(N-i-1)} \frac{y_N^{H,\eta}}{y_N^{L,\eta}} \tag{A15}$$

For part (i), suppose not. Then there exist $i \in \mathcal{N}$, $k \in \mathcal{K}\backslash\{1, K\}$, and $j > i$ such that we can construct a subsequence with $\sigma^\eta(i,k)(j) > 0$. By (A12), this requires $\frac{f_i^{H,\eta}}{f_i^{L,\eta}} \frac{\mu_k^H}{\mu_k^L} \frac{v_{i+1}^{H,\eta} - v_i^{H,\eta}}{v_i^{L,\eta} - v_{i+1}^{L,\eta}} \geq 1$. By (2) above, we also need $\frac{f_{i+1}^{H,\eta}}{f_{i+1}^{L,\eta}}\frac{\mu_1^H}{\mu_1^L}\frac{v_{i+1}^{H,\eta}-v_i^{H,\eta}}{v_i^{L,\eta}-v_{i+1}^{L,\eta}} \leq 1$. Combining inequalities,

$$\frac{f_i^{H,\eta}}{f_i^{L,\eta}} \geq \frac{f_{i+1}^{H,\eta}}{f_{i+1}^{L,\eta}}\left(\frac{\mu_1^H}{\mu_1^L}\frac{\mu_k^L}{\mu_k^H}\right) \Rightarrow \frac{y_i^{H,\eta}}{y_i^{L,\eta}} \geq \frac{y_{i+1}^{H,\eta}}{y_{i+1}^{L,\eta}}\left(\frac{\mu_1^H}{\mu_1^L}\frac{\mu_k^L}{\mu_k^H}\right)$$

Using this in (A15), we obtain

$$\frac{y_1^{H,\eta}}{y_1^{L,\eta}} \geq \left(\frac{\mu_1^H}{\mu_1^L}\frac{\mu_K^L}{\mu_K^H}\right)^{N-2}\left(\frac{\mu_1^H}{\mu_1^L}\frac{\mu_k^L}{\mu_k^H}\right)\frac{y_N^{H,\eta}}{y_N^{L,\eta}} \tag{A16}$$

Recall (established in the final paragraph of the proof of Lemma 5) that $r(\sigma) \geq \sqrt{\frac{y_1^H}{y_1^L}\frac{y_N^L}{y_N^H}}$. Since $\frac{\mu_k^L}{\mu_k^H} > \frac{\mu_K^L}{\mu_K^H}$, (A16) then implies

$$\lim_{\eta\to 0} r(\sigma^\eta) \geq \lim_{\eta\to 0}\sqrt{\frac{y_1^{H,\eta}}{y_1^{L,\eta}}\frac{y_N^{L,\eta}}{y_N^{H,\eta}}} \geq \sqrt{\left(\frac{\mu_1^H}{\mu_1^L}\frac{\mu_K^L}{\mu_K^H}\right)^{N-2}\left(\frac{\mu_1^H}{\mu_1^L}\frac{\mu_k^L}{\mu_k^H}\right)} > r^*$$

So the limit r^* cannot be attained if $\sigma(i,k)(j) > 0$ for some $k \notin \{1, K\}$; this proves part (i) for $j > i$, and the argument for $j < i$ is symmetric.

For part (ii): suppose that for some $i \in \mathcal{N}$ and $j \geq i+2$, $\tau_{i,j}^{S,\eta} > 0$. By (A12), this requires $\frac{f_i^{H,\eta}}{f_i^{L,\eta}}\frac{\mu_K^H}{\mu_K^L}\frac{v_j^{H,\eta}-v_{j-1}^{H,\eta}}{v_{j-1}^{L,\eta}-v_j^{L,\eta}} \geq 1$. By Lemma 6, there exists $j' < j$ such that $\tau_{j,j'}^{S,\eta} > 0$; by (A12), this requires $\frac{f_j^{H,\eta}}{f_j^{L,\eta}}\frac{\mu_1^H}{\mu_1^L}\frac{v_j^{H,\eta}-v_{j-1}^{H,\eta}}{v_{j-1}^{L,\eta}-v_j^{L,\eta}} \leq 1$. Combining inequalities, $\frac{f_i^{H,\eta}}{f_i^{L,\eta}} \geq \frac{f_j^{H,\eta}}{f_j^{L,\eta}}\left(\frac{\mu_1^H}{\mu_1^L}\frac{\mu_K^L}{\mu_K^H}\right)$. Since $j \geq i+2$, this implies

$$\frac{y_i^{H,\eta}}{y_i^{L,\eta}} \geq \frac{y_{i+2}^{H,\eta}}{y_{i+2}^{L,\eta}}\left(\frac{\mu_1^H}{\mu_1^L}\frac{\mu_K^L}{\mu_K^H}\right)$$

Substituting this into (A16) implies $\frac{y_1^{H,\eta}}{y_1^{L,\eta}}\frac{y_N^{L,\eta}}{y_N^{H,\eta}} \geq \left(\frac{\mu_1^H}{\mu_1^L}\frac{\mu_K^L}{\mu_K^H}\right)^{N-2}$, which implies $\lim_{\eta\to 0} r(\sigma^\eta) \geq \sqrt{\frac{y_1^{H,\eta}}{y_1^{L,\eta}}\frac{y_N^{L,\eta}}{y_N^{H,\eta}}} > r^*$. This proves part (ii) for $j \geq i+2$, and the argument for $j \leq i-2$ is symmetric. ■

Lemma 8: *Fix a sequence σ^η. Then $\lim_{\eta\to 0} r(\sigma^\eta) = r^*$ only if $\lim_{\eta\to 0}\tau_{1,2}^{S,\eta} = \lim_{\eta\to 0}\tau_{N,N-1}^{S,\eta} = 0$.*

Proof: Suppose that $\lim_{\eta\to 0}\tau_{1,2}^{S,\eta} \neq 0$. By Lemma 6, $\lim_{\eta\to 0}\frac{\eta}{\tau_{i,j}^{S,\eta}} = 0\ \forall i,j \in \mathcal{N}$; then to calculate $\lim_{\eta\to 0}\frac{y_i^{S,\eta}}{y_1^{S,\eta}}$ ($i \in \mathcal{N}$) using (A13), we can replace $\omega_{i,j}^{S,\eta}$ with $\tau_{i,j}^{S,\eta}$. By Lemma 7, $\tau_{i,j}^{S,\eta} = 0$ (for small η) whenever i,j are not adjacent states. As noted in the proof of Lemma 5, there is only one $i-$tree ($\forall i \in \mathcal{N}$) when all states switch only to adjacent states. Hence,

$$\lim_{\eta\to 0}\frac{y_2^{S,\eta}}{y_1^{S,\eta}} = \lim_{\eta\to 0}\frac{\tau_{1,2}^{S,\eta}\cdot\prod_{i>2}\tau_{i,i-1}^{S,\eta}}{\tau_{2,1}^{S,\eta}\cdot\prod_{i>2}\tau_{i,i-1}^{S,\eta}} = \lim_{\eta\to 0}\frac{\tau_{1,2}^{S,\eta}}{\tau_{2,1}^{S,\eta}}$$

The assumption $\lim_{\eta\to 0}\tau_{1,2}^{S,\eta}\neq 0$ implies that the that the final ratio is positive. By Lemma 5 (i), this implies that the bound r^* is not attained. Similarly if $\lim\limits_{\eta\to 0}\tau_{N,N-1}^{S,\eta}\neq 0$, then $\lim\limits_{\eta\to 0}\frac{y_{N-1}^{S,\eta}}{y_N^{S,\eta}}>0$, so $\lim\limits_{\eta\to 0} r(\sigma^\eta)>r^*$. ■

A.3 Proof of Theorems 2,7 and Corollary to Theorem 7

Proof of Theorem 2: Lemma 3 contains all of the results except for the final part, $\Pi^*(\eta,N)$ increasing in N. Weak monotonicity is clear by revealed preference (the DM could always choose to not use an extra state), and strict monotonicity for small η is established in Lemma 5 (iii) - which proves that all states must be used to achieve the optimal payoff (hence, it is increasing in the number of available states). ■

Proof of Theorem 7, parts (i),(ii),(iii): Consider a sequence σ^η of transition rules, with $\eta\to 0$. Lemma 5 calculated an upper bound on the expected payoff, $\frac{1}{1+r^*}$, and showed that this limit is attainable in the limit as $\eta\to 0$; since the set of optimal strategies is upper hemi-continuous in η (Lemma 3), this implies that σ^η can only be optimal if it achieves the limiting payoff, $\frac{1}{1+r^*}$. By Lemma 8, this requires $\lim\limits_{\eta\to 0}\tau_{1,2}^{S,\eta}=\lim\limits_{\eta\to 0}\tau_{N,N-1}^{S,\eta}=0$; this establishes part (i) of Theorem 7. By Lemma 6, there exist $j'<i<j$ such that $\tau_{i,j'}^S>0$ and $\tau_{i,j}^S>0$, $\forall i\in\mathcal{N}$. Then by Lemma 7, if the sequence σ^η is incentive compatible (which is a necessary condition for optimality, by Theorem 1), it can only attain the limiting payoff $\frac{1}{1+r^*}$ if parts (ii) and (iii) of Theorem 7 hold. ■

Proof of the Corollary to Theorem 7: For part (i): Lemma 5 established that the payoff $\frac{1}{1+r^*}=\left(1+\left(\frac{\mu_1^H}{\mu_1^L}\frac{\mu_K^L}{\mu_K^H}\right)^{\frac{N-1}{2}}\right)^{-1}=\left(1+\left(\frac{1-\widetilde{\rho}}{\widetilde{\rho}}\right)^{(N-1)}\right)^{-1}$ is attainable in the limit as $\eta\to 0$. Then the result follows from Theorem 2, which established that $\Pi^*(\eta,N)$ is continuous and strictly increasing in η. For part (ii): as argued in the above proof of Theorem 7, a sequence σ^η can only be optimal if it attains the limiting payoff $\frac{1}{1+r^*}$. By Lemma 4, this requires

$$\lim_{\eta\to 0}\frac{\sum_{i<i^*}y_i^{H,\eta}}{\sum_{i\geq i^*}y_i^{H,\eta}}=\lim_{\eta\to 0}\frac{\sum_{i\geq i^*}y_i^{L,\eta}}{\sum_{i<i^*}y_i^{L,\eta}}\tag{A17}$$

By Lemma 5, we need $\lim\limits_{\eta\to 0}\frac{y_i^S}{y_1^S}=\lim\limits_{\eta\to 0}\frac{y_i^S}{y_N^S}=0$ for $i\in\mathcal{N}\backslash\{1,N\}$, which also implies $\lim\limits_{\eta\to 0}\frac{f_i^S}{f_1^S}=\lim\limits_{\eta\to 0}\frac{f_i^S}{f_N^S}=0$. Then the LHS of (A17) is

$$\lim_{\eta\to 0}\frac{\sum_{i<i^*}y_i^{H,\eta}}{\sum_{i\geq i^*}y_i^{H,\eta}}=\lim_{\eta\to 0}\frac{y_1^{H,\eta}}{y_N^{H,\eta}}=\lim_{\eta\to 0}\frac{y_1^{H,\eta}/\sum y_i^{H,\eta}}{y_N^{H,\eta}/\sum y_i^{H,\eta}}=\lim_{\eta\to 0}\frac{f_1^{H,\eta}}{f_N^{H,\eta}}=\lim_{\eta\to 0}\frac{f_1^{H,\eta}}{1-f_1^{H,\eta}}$$

Similarly, the RHS of (A17) is equal to $\lim_{\eta\to 0}\frac{f_N^{L,\eta}}{1-f_N^{L,\eta}}$. Then (A17) requires $\lim_{\eta\to 0}\frac{f_1^{H,\eta}}{1-f_1^{H,\eta}}=\lim_{\eta\to 0}\frac{f_N^{L,\eta}}{1-f_N^{L,\eta}}$, which implies $\lim_{\eta\to 0} f_1^{H,\eta}=\lim_{\eta\to 0} f_N^{L,\eta}$, and $f_1^{L,\eta}=1-f_N^{L,\eta}=1-f_1^{H,\eta}=f_N^{H,\eta}$. This yields

$$1=\lim_{\eta\to 0}\frac{f_1^{H,\eta}}{f_1^{L,\eta}}\frac{f_N^{H,\eta}}{f_N^{L,\eta}}=\lim_{\eta\to 0}\frac{y_1^{H,\eta}y_N^{H,\eta}}{y_1^{L,\eta}y_N^{L,\eta}}\left(\frac{y_1^{L,\eta}+y_N^{L,\eta}}{y_1^{H,\eta}+y_N^{H,\eta}}\right)^2$$

By (A15),

$$\frac{y_1^{H,\eta}}{y_1^{L,\eta}}\left(\frac{\mu_1^L\mu_K^H}{\mu_1^H\mu_K^L}\right)^{i-1}\geq\frac{y_i^{H,\eta}}{y_i^{L,\eta}}\geq\left(\frac{\mu_1^H}{\mu_1^L}\frac{\mu_K^L}{\mu_K^H}\right)^{(N-i)}\frac{y_N^{H,\eta}}{y_N^{L,\eta}}$$

By Lemma 5, an optimal sequence σ^η satisfies $\lim_{\eta\to 0}\frac{y_1^{H,\eta}}{y_1^{L,\eta}}=\left(\frac{\mu_1^H}{\mu_1^L}\right)^{N-1}$, $\lim_{\eta\to 0}\frac{y_N^{H,\eta}}{y_N^{L,\eta}}=\left(\frac{\mu_K^H}{\mu_K^L}\right)^{N-1}$. Using this and taking limits in the above equation, we obtain $\lim_{\eta\to 0}\frac{y_i^{H,\eta}}{y_i^{L,\eta}}=\left(\frac{\mu_1^H}{\mu_1^L}\right)^{N-i}\left(\frac{\mu_K^H}{\mu_K^L}\right)^{(i-1)}$. So

$$\lim_{\eta\to 0}\frac{f_i^{H,\eta}}{f_i^{L,\eta}}=\lim_{\eta\to 0}\frac{y_i^{H,\eta}}{y_i^{L,\eta}}\left(\frac{y_1^{L,\eta}+y_N^{L,\eta}}{y_1^{H,\eta}+y_N^{H,\eta}}\right)=\left(\frac{\mu_K^H}{\mu_K^L}\right)^{i-1}\left(\frac{\mu_1^H}{\mu_1^L}\right)^{N-i}\left(\frac{\mu_1^L}{\mu_1^H}\frac{\mu_K^L}{\mu_K^H}\right)^{\frac{N-1}{2}}$$

This simplifies to $\left(\frac{\widetilde{\rho}}{1-\widetilde{\rho}}\right)^{i-1}\left(\frac{1-\widetilde{\rho}}{\widetilde{\rho}}\right)^{N-i}$, as desired. ■

Proof of Theorem 7, part (iv): Choose i with $g_0(i)>0$. By the above Corollary to Theorem 7, an optimal sequence of strategies σ^η has $\lim_{\eta\to 0}\frac{f^{H,\eta}(i)}{f^{L,\eta}(i)}=\left(\frac{\widetilde{\rho}}{1-\widetilde{\rho}}\right)^{(i-1)}\left(\frac{1-\widetilde{\rho}}{\widetilde{\rho}}\right)^{(N-i)}$; this is equal to 1 at $i=\frac{N+1}{2}$, is at least $\left(\frac{\widetilde{\rho}}{1-\widetilde{\rho}}\right)^2$ for $i\geq\frac{N+3}{2}$, and is at most $\left(\frac{1-\widetilde{\rho}}{\widetilde{\rho}}\right)^2$ for $i\leq\frac{N-1}{2}$. If $g_0(i)>0$,then it must be that $v_i^H+v_i^L\geq v_j^H+v_j^L$ $\forall j\in\mathcal{N}$ (otherwise, starting in state j instead of i would increase the expected payoff, at prior $\frac{1}{2}$). Applying this inequality to $j\in\{i-1,i+1\}$ and using the monotonicity condition in Lemma 2 part (2), we obtain

$$\frac{v_{i+1}^H-v_i^H}{v_i^L-v_{i+1}^L}\leq 1\leq\frac{v_i^H-v_{i-1}^H}{v_{i-1}^L-v_i^L}$$

If $i\leq\frac{N-1}{2}$, then the LHS of this inequality implies that

$$\frac{f_i^H}{f_i^L}\frac{\mu_K^H}{\mu_K^L}\frac{v_{i+1}^H-v_i^H}{v_i^L-v_{i+1}^L}\leq\left(\frac{1-\widetilde{\rho}}{\widetilde{\rho}}\right)^2\frac{\mu_K^H}{\mu_K^L}(1)=\frac{\mu_1^H}{\mu_1^L}<1$$

So incentive compatibility (hence optimality) requires $\widehat{\sigma}(i,K)(i+1)=0$, contradicting Lemmas 7 and 8 (which together imply $\widehat{\tau}_{i,i+1}^{S,\eta}>0$ $\forall\eta$). Similarly if $i\geq\frac{N+3}{2}$, then the RHS of the inequality implies $\frac{f_i^H}{f_i^L}\frac{\mu_1^H}{\mu_1^L}\frac{v_i^H-v_{i-1}^H}{v_{i-1}^L-v_i^L}>1$; so $\widehat{\sigma}(i,l)(i-1)>0$ is not optimal, contradicting Lemmas 7 and 8. So, it must be that $i=\frac{N+1}{2}$. ■

Proof of Theorem 7, part (v): Since $\frac{f_i^H}{f_i^L} > 1$ for $i > \frac{N+1}{2}$ and $\frac{f_i^H}{f_i^L} < 1$ for $i < \frac{N+1}{2}$ (by the corollary to Theorem 7), it is clear that incentive compatibility, and hence optimality, requires $a_i = 0$ if $i \leq \frac{N-1}{2}$, and $a_i = 1$ if $i \geq \frac{N+3}{2}$. So, we just need to show that optimality requires $a_{\frac{N+1}{2}} \in \{0,1\}$. The DM's expected payoff is half of

$$\left[\sum_{i<\frac{N+1}{2}} f_i^L + \left(1-\widehat{a}_{\frac{N+1}{2}}\right) f^L_{\frac{N+1}{2}}\right] + \left[\sum_{i>\frac{N+1}{2}} f_i^H + \widehat{a}_{\frac{N+1}{2}} f^H_{\frac{N+1}{2}}\right]$$

Using $\sum_i f_i^S = 1$, rewrite this as

$$f_1^L\left[1+\sum_{i=2}^{\frac{N-1}{2}}\frac{f_i^L}{f_1^L} + \left(1-\widehat{a}_{\frac{N+1}{2}}\right)\frac{f^L_{\frac{N+1}{2}}}{f_1^L}\right] - f_1^H\left[1+\sum_{i=2}^{\frac{N-1}{2}}\frac{f_i^H}{f_1^H} + \left(1-\widehat{a}_{\frac{N+1}{2}}\right)\frac{f^H_{\frac{N+1}{2}}}{f_1^H}\right] + 1 \tag{A22}$$

Suppose, by contradiction, that there is an optimal solution $(\widehat{g}_0, \widehat{\sigma}, \widehat{a})$ with $\widehat{a}_{\frac{N+1}{2}} \in (0,1)$; this requires $\frac{d\Pi(\widehat{g}_0,\widehat{\sigma},\widehat{a})}{d\widehat{a}_{\frac{N+1}{2}}} = 0$. Since $\Pi\left(\widehat{g}_0, \widehat{\sigma}, \widehat{a}\right)$ is linear in $\widehat{a}$, this implies that $\Pi\left(\widehat{g}_0, \widehat{\sigma}, \widehat{a}\right)$ would not change if we increased $\widehat{a}\left(\frac{N+1}{2}\right)$ to 1. Then to prove that $(\widehat{g}_0, \widehat{\sigma}, \widehat{a})$ is not optimal, it is sufficient to show that $\widehat{\sigma}$ does not maximize the expected payoff when $\widehat{a}\left(\frac{N+1}{2}\right) = 1$.

By Theorem 7 (i), $\widehat{\sigma}_{N,N-1} \in (0,1)$, so the derivative of (A22) w.r.t. $\sigma_{N,N-1}$ must equal zero. Noting that $\frac{f_i^S}{f_1^S}$ is independent of $\sigma_{N,N-1}$ (easily established by induction, noting that f_i^S solves the recursion $f_i^S\left(\eta + \tau^S_{i,i-1} + \tau^S_{i,i+1}\right) = f^S_{i-1}\tau^S_{i-1,i} + f^S_{i+1}\tau^S_{i+1,i}$), this FOC can be written as

$$\frac{\frac{df_1^L}{d\sigma^1_{N,N-1}}}{\frac{df_1^H}{d\sigma^1_{N,N-1}}} = \frac{1+\sum_{i=2}^{\frac{N-1}{2}}\frac{f_i^H}{f_1^H} + \left(1-\widehat{a}_{\frac{N+1}{2}}\right)\frac{f^H_{\frac{N+1}{2}}}{f_1^H}}{1+\sum_{i=2}^{\frac{N-1}{2}}\frac{f_i^L}{f_1^L} + \left(1-\widehat{a}_{\frac{N+1}{2}}\right)\frac{f^L_{\frac{N+1}{2}}}{f_1^L}} = \frac{f_1^L}{f_1^H}\frac{\left(\sum_{i=1}^{\frac{N-1}{2}} f_i^H + \left(1-\widehat{a}_{\frac{N+1}{2}}\right) f^H_{\frac{N+1}{2}}\right)}{\left(\sum_{i=1}^{\frac{N-1}{2}} f_i^L + \left(1-\widehat{a}_{\frac{N+1}{2}}\right) f^L_{\frac{N+1}{2}}\right)}$$

The LHS is independent of $\widehat{a}_{\frac{N+1}{2}}$, while $\frac{f_i^H}{f_i^L}$ strictly increasing in i implies that the RHS is strictly increasing in $\left(1-\widehat{a}_{\frac{N+1}{2}}\right)$. This implies that if $\widehat{\sigma}$ is optimal at $\widehat{a}_{\frac{N+1}{2}} \in (0,1)$, it cannot satisfy the F.O.C. at $\widehat{a}_{\frac{N+1}{2}} = 1$. ■

A.4 Proof of Theorem 3 (ii)

For this section, let $\sigma_{i,i+1} \equiv \sigma(i,h)(i+1)$, and $\sigma_{i,i-1} \equiv \sigma(i,l)(i-1)$. Note that all results in Theorem 3 (and the corollary) are special cases of the results in Theorem 7, except for part (ii): Theorem 7 only implies that $\sigma_{i,i+1}$ and $\sigma_{i,i-1}$ are strictly positive, not that they are equal to 1.

Fix an optimal strategy (σ, g_0, a). For any i with $a_i = 0$, define $\widetilde{\Delta^S_{i,i+1}} \equiv \left(\prod_{j=2}^{i} \tau^S_{j,j+1}\right) \frac{v^S_{i+1} - v^S_i}{v^S_2 - v^S_1}$; and for any $i \leq \frac{N+1}{2}$, define $\widetilde{f^S_i} \equiv \left(\prod_{j=2}^{i} \tau^S_{j,j-1}\right) \frac{f^S_i}{f^S_1}$.

Claim 0: In an optimal solution, the following conditions must hold:

(i) For any i with $a_i = 0$: $\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}} \frac{\widetilde{\Delta^H_{i,i+1}}}{\widetilde{\Delta^L_{i,i+1}}} \geq 1$, with equality if $\sigma_{i,i+1} \in (0,1)$

(ii) For any $i \leq \frac{N+1}{2}$: $\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}} \frac{1-\rho}{\rho} \frac{\widetilde{\Delta^H_{i-1,i}}}{\widetilde{\Delta^L_{i-1,i}}} \leq 1$, with equality if $\sigma_{i,i-1} \in (0,1)$.

Proof: The inequality $\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}} \frac{\widetilde{\Delta^H_{i,i+1}}}{\widetilde{\Delta^L_{i,i+1}}} \geq 1$ is exactly equivalent to the optimality condition for $\sigma_{i,i+1}$, noting that (i) in a symmetric solution, $\frac{\tau^H_{i,i+1}}{\tau^L_{i,i+1}} = \frac{\rho}{1-\rho} = \frac{\tau^L_{i,i-1}}{\tau^H_{i,i-1}}$; (ii) optimality for $\sigma_{1,2}$ requires $\frac{f^H_1}{f^L_1} \frac{\rho}{1-\rho} \frac{v^H_2 - v^H_1}{v^L_1 - v^L_2} = 1$. Similarly, (ii) is equivalent to the optimality condition for $\sigma_{i,i-1}$, given that $\sigma_{1,2} \in (0,1)$ is optimal. ■

Claim 1: $\widetilde{f^S_i}, \widetilde{\Delta^S_{i,i+1}}$ satisfy the following recursion formulas:

$$\widetilde{\Delta^S_{i,i+1}} = \left(\eta + \tau^S_{i,i-1} + \tau^S_{i-1,i}\right) \widetilde{\Delta^S_{i-1,i}} - \tau^S_{i-1,i-2} \tau^S_{i-1,i} \widetilde{\Delta^S_{i-2,i-1}} \tag{A23}$$

$$\widetilde{f^S_i} = \left(\eta + \tau^S_{i-1,i-2} + \tau^S_{i-1,i}\right) \widetilde{f^S_{i-1}} - \tau^S_{i-2,i-1} \tau^S_{i-1,i-2} \widetilde{f^S_{i-2}} \tag{A24}$$

$$\widetilde{f^S_i} = \widetilde{\Delta^S_{i,i+1}} - \tau^S_{i,i-1} \widetilde{\Delta^S_{i-1,i}}, \text{ and } \eta \widetilde{\Delta^S_{i,i+1}} = \widetilde{f^S_{i+1}} - \tau^S_{i,i+1} \widetilde{f^S_i} \tag{A25}$$

Proof: For equation (A23): recall that if $a_i = 0$, then v^H_i satisfies

$$v^H_i = \frac{\tau^H_{i,i-1} v^H_{i-1} + \tau^H_{i,i+1} v^H_{i+1}}{\eta + \tau^H_{i,i-1} + \tau^H_{i,i+1}}$$

Subtracting this from v^H_{i+1}, and solving for $(v^H_{i+1} - v^H_i)$, yields

$$\tau^H_{i,i+1} \left(v^H_{i+1} - v^H_i\right) = \eta v^H_i + \tau^H_{i,i-1} \left(v^H_i - v^H_{i-1}\right) \tag{A26}$$

$$= \eta v^H_{i-1} + \left(\eta + \tau^H_{i,i-1}\right) \left(v^H_i - v^H_{i-1}\right) \tag{A27}$$

Now lag (A26) by one, to obtain $\eta v^H_{i-1} = \tau^H_{i-1,i}(v^H_i - v^H_{i-1}) - \tau^H_{i-1,i-2}(v^H_{i-1} - v^H_{i-2})$; substituting this into (A27) yields

$$\tau^H_{i,i+1} \left(v^H_{i+1} - v^H_i\right) = \left(\eta + \tau^H_{i,i-1} + \tau^H_{i-1,i}\right) \left(v^H_i - v^H_{i-1}\right) - \tau^H_{i-1,i-2}(v^H_{i-1} - v^H_{i-2})$$

Now multiply both sides by $\frac{\prod_{j=2}^{i-1}\tau_{j,j+1}^H}{v_2^H - v_1^H}$, to obtain the desired expression for H; the calculation for L is identical.

For equation (A24): for all $i < \frac{N+1}{2}$, f_i^S solves

$$f_i^S = \frac{\tau_{i+1,i}^S f_{i+1}^S + \tau_{i-1,i}^S f_{i-1}^S}{\eta + \tau_{i,i-1}^S + \tau_{i,i+1}^S}$$

Solving for f_{i+1}^S yields $\tau_{i+1,i}^S f_{i+1}^S = f_i^S\left(\eta + \tau_{i,i-1}^S + \tau_{i,i+1}^S\right) - \tau_{i-1,i}^S f_{i-1}^S$; then multiplying both sides by $\frac{\prod_{j=2}^{i}\tau_{j,j-1}^S}{f_1^S}$ yields the desired expression.[13]

For equation (A25) (by induction from (A23) and (A24)): Since $\widetilde{f_1^S} = \widetilde{\Delta_{1,2}^S} = 1$ (by definition), and evaluating equation (2) at $i = 2$ yields $\widetilde{f_2^S} = \eta + \tau_{1,2}^S$, it follows that

$$\widetilde{f_1^S} = \widetilde{\Delta_{1,2}^S}, \text{ and } \eta\widetilde{\Delta_{1,2}^S} = \widetilde{f_2^S} - \tau_{1,2}^S\widetilde{f_1^S}$$

verifying equation (A25) for $i = 1$. Now suppose (A25) holds for all $j \leq i-1$. To show that it also holds for $j = i$, rewrite (A23) as

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

By the first equation in (A25) (at $i-1$), the second bracketed term is equal to $\widetilde{f_{i-1}^S}$; by the second equation in (A25) (at $i-1$), $\tau_{i-1,i}^S\widetilde{f_{i-1}^S} = \widetilde{f_i^S} - \eta\widetilde{\Delta_{i-1,i}^S}$; so, we obtain $\widetilde{\Delta_{i,i+1}^S} = \left(\eta + \tau_{i,i-1}^S\right)\widetilde{\Delta_{i-1,i}^S} + \widetilde{f_i^S} - \eta\widetilde{\Delta_{i-1,i}^S}$, which verifies the first equation in (A25) for $j = i$. A similar calculation verifies the second equation for $j = i$. ∎

Claim 2: Suppose that for all $j < i$, $\sigma_{j,j-1} = \sigma_{j,j+1} = 1$. Then the following inequalities hold: (i)$\frac{\widetilde{\Delta_{i,i+1}^H}}{\widetilde{\Delta_{i,i+1}^L}} \leq \frac{\rho}{1-\rho}\frac{\widetilde{\Delta_{i-1,i}^H}}{\widetilde{\Delta_{i-1,i}^L}}$; (ii)$\frac{\widetilde{f_i^H}}{\widetilde{f_{il}^L}} > 1$; (iii)$\frac{\widetilde{\Delta_{i-1,i}^H}}{\widetilde{\Delta_{i-1,i}^L}} < 1$; (iv)$\frac{\widetilde{\Delta_{i,i+1}^H}}{\widetilde{\Delta_{i,i+1}^L}} > \frac{1-\rho}{\rho}\frac{\widetilde{\Delta_{i,i-1}^H}}{\widetilde{\Delta_{i,i-1}^L}}$; (v)$\frac{\widetilde{f_{i+1}^H}}{\widetilde{f_{i+1}^L}} < \frac{\rho}{1-\rho}\frac{\widetilde{f_i^H}}{\widetilde{f_i^L}}$.

Proof: Part (i) is equivalent to Lemma 2 (iv). For (ii): Equation (A24) implies that $\widetilde{f_{j+1}^S} = \widetilde{f_j^S} - (1-\eta)^2\rho(1-\rho)\widetilde{f_{j-1}^S}$ for all $2 \leq j \leq i-1$, so $\frac{\widetilde{f_{j+1}^H}}{\widetilde{f_{j+1}^L}} > \frac{\widetilde{f_j^H}}{\widetilde{f_j^L}} \Leftrightarrow \frac{\widetilde{f_j^H}}{\widetilde{f_j^L}} > \frac{\widetilde{f_{j-1}^H}}{\widetilde{f_{j-1}^L}}$. Then since $\frac{\widetilde{f_2^H}}{\widetilde{f_2^L}} = \frac{\eta+\tau_{1,2}^H}{\eta+\tau_{1,2}^L} > 1$ and $\frac{\widetilde{f_1^H}}{\widetilde{f_1^L}} = 1$, it follows that $\frac{\widetilde{f_i^H}}{\widetilde{f_i^L}} > 1$. For (iii): Equation (A23) implies that $\widetilde{\Delta_{j,j+1}^S} = \widetilde{\Delta_{j-1,j}^S} - (1-\eta)^2\rho(1-\rho)\widetilde{\Delta_{j-2,j-1}^S}$ for $2 \leq j \leq i-1$, so $\frac{\widetilde{\Delta_{j,j+1}^H}}{\widetilde{\Delta_{j,j+1}^L}} < \frac{\widetilde{\Delta_{j-1,j}^H}}{\widetilde{\Delta_{j-1,j}^L}} \Leftrightarrow \frac{\widetilde{\Delta_{j-1,j}^H}}{\widetilde{\Delta_{j-1,j}^L}} < \frac{\widetilde{\Delta_{j-2,j-1}^H}}{\widetilde{\Delta_{j-2,j-1}^L}}$; since

[13] Note that the recursion equation for $\widetilde{\Delta_{i,i+1}^S}$ relies on $a_i = 0$, and the equation for $\widetilde{f_{i+1}^S}$ relies on $i \leq \frac{N-1}{2}$.

$\frac{\widetilde{\Delta^H_{2,3}}}{\widetilde{\Delta^L_{2,3}}} = \frac{\eta+\tau^H_{2,1}+\tau^H_{1,2}}{\eta+\tau^L_{2,1}+\tau^L_{1,2}} < 1 = \frac{\widetilde{\Delta^H_{1,2}}}{\widetilde{\Delta^L_{1,2}}} = 1$, we then obtain $\frac{\widetilde{\Delta^H_{j,j+1}}}{\widetilde{\Delta^L_{j,j+1}}} < 1\ \forall j \le i-1$. For (iv), suppose the inequality is false. Then by the first equation in (A25), $\widetilde{\Delta^S_{i,i+1}} = \widetilde{f^S_i} + \tau^S_{i,i-1}\widetilde{\Delta^S_{i-1,i}}$, it must be that $\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}} \le \frac{1-\rho}{\rho}\frac{\widetilde{\Delta^H_{i,i-1}}}{\widetilde{\Delta^L_{i,i-1}}}$; this implies $\left(\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}}\right)^2 \le \frac{\widetilde{f^H_i}}{\widetilde{f^L_i}}\frac{1-\rho}{\rho}\frac{\widetilde{\Delta^H_{i,i-1}}}{\widetilde{\Delta^L_{i,i-1}}}$. The RHS must be at most 1 by the optimality condition for $\tau^S_{i,i-1} > 0$, so we need $\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}} \le 1$; this contradicts (ii). For (v): the second equation in (A25) yields $\frac{\widetilde{f^H_{i+1}}}{\widetilde{f^L_{i+1}}} = \frac{\tau^H_{i,i+1}\widetilde{f^H_i}+\eta\widetilde{\Delta^H_{i,i+1}}}{\tau^L_{i,i+1}\widetilde{f^L_i}+\eta\widetilde{\Delta^L_{i,i+1}}}$, so we need to show that $\frac{\widetilde{\Delta^H_{i,i+1}}}{\widetilde{\Delta^L_{i,i+1}}} < \frac{\rho}{1-\rho}\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}}$; this follows from (i) and (iii), which imply that the LHS is below $\frac{\rho}{1-\rho}$, and (ii), which implies that the RHS is greater than $\frac{\rho}{1-\rho}$. ■

Claim 3: Suppose that for all $j < i$, $\sigma_{j,j-1} = \sigma_{j,j+1} = 1$. Then:

(i) if $\sigma_{i,i+1} \in (0,1)$ is optimal, then $\sigma_{i,i-1} = \sigma_{i+1,i} = 1$, and $\frac{\widetilde{f^H_{i+1}}}{\widetilde{f^L_{i+1}}}, \frac{\widetilde{\Delta^H_{i+1,i+2}}}{\widetilde{\Delta^L_{i+1,i+2}}}$ are both increasing in $\sigma_{i,i+1}$;

(ii) if $\sigma_{i,i-1} \in (0,1)$ is optimal, then $\sigma_{i-1,i} = \sigma_{i,i+1} = 1$, and $\frac{\widetilde{f^H_{i+1}}}{\widetilde{f^L_{i+1}}}, \frac{\widetilde{\Delta^H_{i,i+1}}}{\widetilde{\Delta^L_{i,i+1}}}$ are decreasing in $\sigma_{i,i-1}$.

Proof: For (i): suppose that $\sigma_{i,i+1} \in (0,1)$ is optimal, requiring $\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}}\frac{\widetilde{\Delta^H_{i,i+1}}}{\widetilde{\Delta^L_{i,i+1}}} = 1$; then parts (iv) and (v) of Claim 2 imply that the IC conditions for $\tau^S_{i,i-1} > 0, \tau^S_{i+1,i} > 0$ cannot hold with equality, so $\sigma_{i,i-1} = \sigma_{i+1,i} = 1$. To show that $\frac{\widetilde{f^H_{i+1}}}{\widetilde{f^L_{i+1}}}$ is increasing in $\sigma_{i,i+1}$, we need to show that $\frac{d\widetilde{f^H_{i+1}}/d\sigma_{i,i+1}}{d\widetilde{f^L_{i+1}}/d\sigma_{i,i+1}} > \frac{\widetilde{f^H_{i+1}}}{\widetilde{f^L_{i+1}}}$. By (A24), using $\sigma_{i,i-1} = 1$ and noting that $\widetilde{f^H_j}$ is independent of $\sigma_{i,i+1}$ for $j \le i$, we obtain $\frac{d\widetilde{f^H_{i+1}}/d\sigma_{i,i+1}}{d\widetilde{f^L_{i+1}}/d\sigma_{i,i+1}} = \frac{\rho}{1-\rho}\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}}$; this is greater than $\frac{\widetilde{f^H_{i+1}}}{\widetilde{f^L_{i+1}}}$ by Claim 2 (v). To show that $\frac{\widetilde{\Delta^H_{i+1,i+2}}}{\widetilde{\Delta^L_{i+1,i+2}}}$ is increasing in $\sigma_{i,i+1}$, we need to show that $\frac{d\widetilde{\Delta^H_{i+1,i+2}}/d\sigma_{i,i+1}}{d\widetilde{\Delta^L_{i+1,i+2}}/d\sigma_{i,i+1}} > \frac{\widetilde{\Delta^H_{i+1,i+2}}}{\widetilde{\Delta^L_{i+1,i+2}}}$. By (A23), using $\sigma_{i+1,i} = \sigma_{i,i-1} = 1$ and noting that $\widetilde{\Delta^H_{i,i+1}}, \widetilde{\Delta^H_{i-1,i}}$ are independent of $\sigma_{i,i+1}$, we obtain $\frac{d\widetilde{\Delta^H_{i+1,i+2}}/d\sigma_{i,i+1}}{d\widetilde{\Delta^L_{i+1,i+2}}/d\sigma_{i,i+1}} = \frac{\rho}{1-\rho}\left(\frac{\widetilde{\Delta^H_{i,i+1}}-(1-\eta)(1-\rho)\widetilde{\Delta^H_{i,i-1}}}{\widetilde{\Delta^L_{i,i+1}}-(1-\eta)\rho\widetilde{\Delta^L_{i,i-1}}}\right)$. This is greater than $\frac{\rho}{1-\rho}\frac{\widetilde{\Delta^H_{i,i+1}}}{\widetilde{\Delta^L_{i,i+1}}}$ by Claim 2 (iv), which is at least $\frac{\widetilde{\Delta^H_{i+1,i+2}}}{\widetilde{\Delta^L_{i+1,i+2}}}$ by Claim 2 (i).

For (ii): the first statement is equivalent to the first statement in (i). To show that $\frac{\widetilde{\Delta^H_{i,i+1}}}{\widetilde{\Delta^L_{i,i+1}}}$ is decreasing in $\sigma_{i,i-1}$, we need to show that $\frac{d\widetilde{\Delta^H_{i,i+1}}/d\sigma_{i,i-1}}{d\widetilde{\Delta^L_{i,i+1}}/d\sigma_{i,i-1}} < \frac{\widetilde{\Delta^H_{i,i+1}}}{\widetilde{\Delta^L_{i,i+1}}}$. By (A23), using the fact that

$\widetilde{\Delta^S_{i-1,i}}, \widetilde{\Delta^S_{i-1,i-2}}$ are independent of $\sigma_{i,i-1}$, we obtain $\frac{d\widetilde{\Delta^H_{i,i+1}}/d\sigma_{i,i-1}}{d\widetilde{\Delta^L_{i,i+1}}/d\sigma_{i,i-1}} = \frac{1-\rho}{\rho}\frac{\widetilde{\Delta^H_{i-1,i}}}{\widetilde{\Delta^L_{i-1,i}}}$; this is below $\frac{\widetilde{\Delta^H_{i,i+1}}}{\widetilde{\Delta^L_{i,i+1}}}$ by Claim 2 (iv). To show that $\frac{\widetilde{f^H_{i+1}}}{\widetilde{f^L_{i+1}}}$ is decreasing in $\sigma_{i,i-1}$, we need to show that $\frac{d\widetilde{f^H_{i+1}}/d\sigma_{i,i-1}}{d\widetilde{f^L_{i+1}}/d\sigma_{i,i-1}} < \frac{\widetilde{f^H_{i+1}}}{\widetilde{f^L_{i+1}}}$. By (A24), using $\sigma_{i,i+1} = 1$ and the fact that $\widetilde{f^S_i}, \widetilde{f^S_{i-1}}$ are independent of $\sigma_{i,i-1}$,we obtain $\frac{d\widetilde{f^H_{i+1}}/d\sigma_{i,i-1}}{d\widetilde{f^L_{i+1}}/d\sigma_{i,i-1}} = \left(\frac{1-\rho}{\rho}\right)\left(\frac{\widetilde{f^H_i}-(1-\eta)\rho\widetilde{f^H_{i-1}}}{\widetilde{f^L_i}-(1-\eta)(1-\rho)\widetilde{f^L_{i-1}}}\right)$; this is below $\frac{1-\rho}{\rho}\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}}$ by Claim 2 (v), which is below $\frac{\widetilde{f^H_{i+1}}}{\widetilde{f^L_{i+1}}}$ iff $\frac{\widetilde{f^H_{i+1}}}{\widetilde{f^L_{i+1}}} > \frac{\widetilde{f^H_i}}{\widetilde{f^L_i}}$, true by construction. ■

Claim 4: Define $\alpha \equiv (1-\eta)^2\rho(1-\rho)$, and define coefficients γ_j recursively by $\gamma_1 = 0, \gamma_2 = 1$, and $\gamma_{j+1} = \gamma_j - \alpha\gamma_{j-2}$ for $j \geq 2$.[14] Then:

(i) If $\sigma_{j,j-1} = \sigma_{j,j+1} = 1 \ \forall\ 2 \leq j < i$, then $\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}}\frac{\widetilde{\Delta^H_{i,i+1}}}{\widetilde{\Delta^L_{i,i+1}}} \geq 1 \Leftrightarrow \gamma_i^2\left(2\eta\sigma_{1,2} + (1-\eta)\sigma_{1,2}^2 - \eta\right) - \alpha\gamma_{i-1}\gamma_i\sigma_{1,2}^2 \geq 0$.

(ii) If $\sigma_{j,j-1} = \sigma_{j,j+1} = 1 \ \forall\ 2 \leq j < i$, then $\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}}\frac{1-\rho}{\rho}\frac{\widetilde{\Delta^H_{i-1,i}}}{\widetilde{\Delta^L_{i-1,i}}} \Leftrightarrow \gamma_i^2\eta - \alpha\gamma_{i-1}\gamma_i\sigma_{1,2}^2 \geq 0$.

Proof: If $\sigma_{j,j-1} = \sigma_{j,j+1} = 1 \ \forall\ 2 \leq j < i$, then solving the recursions in (A23) and (A24) yields the following expression:

$$\widetilde{f^S_i} = \gamma_i\widetilde{f^S_2} - \alpha\gamma_{i-1}\sigma_{1,2} \tag{A28}$$

$$\widetilde{\Delta^S_{i,i+1}} = \gamma_i\widetilde{\Delta^S_{2,3}} - \alpha\gamma_{i-1} - (1-\sigma_{i,i-1})(1-\eta)\mu^S_l\widetilde{\Delta^S_{i-1,i}} \tag{A29}$$

$$(1-\eta)\mu^S_l\widetilde{\Delta^S_{i-1,i}} = \gamma_i(1-\eta)\mu^S_l - \alpha\gamma_{i-1}(1-\sigma_{1,2}) \tag{A30}$$

For expression (ii), rewrite the condition $\frac{\widetilde{f^H_i}}{\widetilde{f^L_i}}\frac{1-\rho}{\rho}\frac{\widetilde{\Delta^H_{i-1,i}}}{\widetilde{\Delta^L_{i-1,i}}} \leq 1$ as $\widetilde{f^H_i}(1-\eta)\mu^H_l\widetilde{\Delta^H_{i-1,i}} - \widetilde{f^L_i}(1-\eta)\mu^L_l\widetilde{\Delta^L_{i-1,i}} \leq 0$. Using (A28) and (A30), this is

$$\gamma_i^2(1-\eta)\left[\widetilde{f^H_2}(1-\rho) - \widetilde{f^L_2}\rho\right] - \alpha\gamma_{i-1}\gamma_i\left[\sigma_{1,2}(1-\eta)(1-2\rho) + (1-\sigma_{1,2})\left(\widetilde{f^H_2} - \widetilde{f^L_2}\right)\right] \leq 0$$

[14]Note that $\alpha < \frac{1}{4}$, and that the recursion for γ_j is solved by $\gamma_j \equiv \frac{1}{\sqrt{1-4\alpha}}\left(\left(\frac{2\alpha}{1-\sqrt{1-4\alpha}}\right)^{j-1} - \left(\frac{2\alpha}{1+\sqrt{1-4\alpha}}\right)^{j-1}\right)$. This is positive for $j \geq 2$, which also implies that $\gamma_{j+1} = \gamma_j - \alpha\gamma_{j-1} > 0$ for $j \geq 2$; so $\sqrt{\frac{\gamma_i}{\alpha\gamma_{i-1}}}$, $\sqrt{1-\frac{\alpha\gamma_{i-1}}{\alpha\gamma_i}}$ are both real numbers. Also note that γ_j satisfies $\gamma_{j+1}\gamma_{j-1} - \gamma_j^2 = \alpha\left(\gamma_j\gamma_{j-2} - \gamma_{j-1}^2\right)$. This implies that $\frac{\gamma_i}{\gamma_{i+1}} > \frac{\gamma_{i-1}}{\gamma_i}$ iff $\frac{\gamma_{i-1}}{\gamma_i} > \frac{\gamma_{i-2}}{\gamma_{i-1}}$; since $\frac{\gamma_2}{\gamma_3} = 1 > 0 = \frac{\gamma_1}{\gamma_2}$, it then follows that $\frac{\gamma_i}{\gamma_{i+1}}$ is strictly increasing in i. Finally, a straightforward calculation shows that $\lim_{j\to\infty}\left(\lim_{\eta\to 0}\frac{\gamma_j}{\alpha\gamma_{j-1}}\right) = \lim_{j\to\infty}\left(\frac{\left(\rho^{j-1}-(1-\rho)^{j-1}\right)}{\rho(1-\rho)\left(\rho^{j-2}-(1-\rho)^{j-2}\right)}\right) = \frac{1}{(1-\rho)}$. Since $\frac{\gamma_j}{\alpha\gamma_{j-1}}$ is decreasing in j $\forall\eta$, this implies that $\lim_{\eta\to 0}\frac{\gamma_j}{\alpha\gamma_{j-1}}$ is strictly greater than $\frac{1}{(1-\rho)}$ for finite N.

Using (A23) and (A24) to evaluate $\widetilde{f_2^H}$ and $\widetilde{f_2^L}$, and dividing the expression by $(1-\eta)(2\rho-1)$,yields the desired inequality.

For (i), rewrite $\frac{\widetilde{f_i^H}}{\widetilde{f_i^L}} \frac{\widetilde{\Delta_{i,i+1}^H}}{\widetilde{\Delta_{i,i+1}^L}} \geq 1$ as $\widetilde{f_i^H}\widetilde{\Delta_{i,i+1}^H} - \widetilde{f_i^L}\widetilde{\Delta_{i,i+1}^L} \geq 0$; using (A28) and (A29), this is

$$0 \leq \left(\gamma_i \widetilde{f_2^H} - \alpha\gamma_{i-1}\sigma_{1,2}\right)\left(\gamma_i \widetilde{\Delta_{2,3}^H} - \alpha\gamma_{i-1}\right) - \left(\gamma_i \widetilde{f_2^L} - \alpha\gamma_{i-1}\sigma_{1,2}\right)\left(\gamma_i \widetilde{\Delta_{2,3}^L} - \alpha\gamma_{i-1}\right)$$
$$- (1-\eta)(1-\sigma_{i,i-1})\left(\widetilde{f_i^H}(1-\rho)\widetilde{\Delta_{i-1,i}^H} - \widetilde{f_i^L}\rho\widetilde{\Delta_{i-1,i}^L}\right)$$

The second line must equal zero, since optimality for $\sigma_{i,i-1}$ requires that the term in square brackets be non-positive, with $\sigma_{i,i-1} = 1$ whenever it is strictly negative. Using (A23) and (A24) to evaluate $\widetilde{f_2^S}$ and $\widetilde{\Delta_{2,3}^S}$, then dividing the first line by $(1-\eta)(2\rho-1)$, yields the desired inequality. ■

Proof of Theorem: Assume that $\widehat{a}_{\frac{N+1}{2}} = 0$; the proof if $\widehat{a}_{\frac{N+1}{2}} = 1$ is symmetric.

Suppose first that $\sigma_{i,i+1} \in (0,1)$ for some $2 \leq i < \frac{N+1}{2}$, and choose the smallest such i. By Claim 3 (i) and our choice of i, $\frac{\widetilde{f_{i+1}^H}}{\widetilde{f_{i+1}^L}} \frac{\widetilde{\Delta_{i+1,i+2}^H}}{\widetilde{\Delta_{i+1,i+2}^L}}$ is bounded above by the value which would be attained at $\sigma_{j,j-1} = \sigma_{j,j+1} = 1 \ \forall \ j \leq i$. So by Claim 0 and Claim 4 (i), optimality of $\sigma_{i,i+1} \in (0,1)$ requires $\frac{\sigma_{1,2}}{\sqrt{\eta}} = \frac{1}{\sqrt{1-\frac{\alpha\gamma_{i-1}}{\gamma_i}+\sqrt{\eta}}}$, and optimality of $\sigma_{i+1,i+2} > 0$ requires $\frac{\sigma_{1,2}}{\sqrt{\eta}} > \frac{1}{\sqrt{1-\frac{\alpha\gamma_{i-1}}{\gamma_i}+\sqrt{\eta}}}$. But since $\frac{\gamma_{i-1}}{\gamma_i} < \frac{\gamma_i}{\gamma_{i+1}}$ (established in footnote 14), it is impossible for both of these inequalities to be satisfied.

Next, suppose $\sigma_{i,i-1} \in (0,1)$ for some $3 \leq i < \frac{N+1}{2}$, and choose the smallest such i. By Claim 3 (ii) and our choice of i, $\frac{\widetilde{f_{i+1}^H}}{\widetilde{f_{i+1}^L}} \frac{1-\rho}{\rho} \frac{\widetilde{\Delta_{i,i+1}^H}}{\widetilde{\Delta_{i,i+1}^L}}$ is bounded below by the value which would be attained at $\sigma_{j,j-1} = \sigma_{j,j+1} = 1 \ \forall \ j \leq i$. So by Claim 0 and Claim 4 (ii), optimality of $\sigma_{i,i-1} \in (0,1)$ requires $\frac{\sigma_{1,2}}{\sqrt{\eta}} = \sqrt{\frac{\gamma_i}{\alpha\gamma_{i-1}}}$, and optimality of $\sigma_{i+1,i} \in (0,1)$ requires $\frac{\sigma_{1,2}}{\sqrt{\eta}} < \sqrt{\frac{\gamma_{i+1}}{\alpha\gamma_i}}$. But since $\frac{\gamma_i}{\gamma_{i-1}} > \frac{\gamma_{i+1}}{\gamma_i}$, it is impossible for both inequalities to be satisfied. Thus, it must be that $\sigma_{i,i-1} = \sigma_{i,i+1} = 1 \ \forall \ 2 \leq i \leq \frac{N-1}{2}$.

A symmetric argument rules out $\sigma_{i,i-1} \in (0,1)$ for $i \geq \frac{N+5}{2}$, and $\sigma_{i,i+1} \in (0,1) \in (0,1)$ for $i \geq \frac{N+3}{2}$. [15] So, we just need to verify that $\sigma_{\frac{N+1}{2},\frac{N-1}{2}} = \sigma_{\frac{N+1}{2},\frac{N+3}{2}} = \sigma_{\frac{N+3}{2},\frac{N+1}{2}} = 1$. For future

[15] More precisely, Claims 1-4 developed optimality conditions for $\sigma_{i,i+1}$ when $a_i = 0$, written in terms of the optimality condition for $\sigma_{1,2} \in (0,1)$ and the behavior of states $j < i$. By counting states from N, rather than from 1, we could obtain an analogous optimality condition for $\sigma_{N+1-i,N+1-(i+1)}$, for any i s.t. $a_{N+1-i} = 1$, in terms of the optimality condition for $\sigma_{N,N-1} \in (0,1)$ and the behavior of states $j < N+1-i$. Similarly, the optimality condition for $\sigma_{N+1-i,N+1-(i-1)}$ (in terms of the behavior of higher states) is analogous to the optimality condition for $\sigma_{i,i-1}$

reference: from the proof of the Corollary to Theorem 7, an optimal solution requires

$$1=\lim_{\eta\to 0}\frac{y_1^H+y_N^H}{y_1^L+y_N^L}=\frac{\left(\prod_{j=2}^{N}\sigma_{j,j-1}\right)(1-\rho)^N+\left(\prod_{j=1}^{N-1}\sigma_{j,j+1}\right)\rho^N}{\left(\prod_{j=2}^{N}\sigma_{j,j-1}\right)\rho^N+\left(\prod_{j=1}^{N-1}\sigma_{j,j+1}\right)(1-\rho)^N}\Leftrightarrow\lim_{\eta\to 0}\frac{\prod_{j=1}^{N-1}\sigma_{j,j+1}}{\prod_{j=2}^{N}\sigma_{j,j-1}}=1 \quad \text{(A31)}$$

Suppose first that $\sigma_{\frac{N+3}{2},\frac{N+1}{2}}\in(0,1)$. By (analogy to) Claim 4 (i), evaluated at $i=\frac{N-1}{2}$ and replacing $\sigma_{1,2}$ with $\sigma_{N,N-1}$, this requires $\frac{\sigma_{N,N-1}}{\sqrt{\eta}}=\frac{1}{\sqrt{1-\frac{\alpha\gamma_{\frac{N-3}{2}}}{\gamma_{\frac{N-1}{2}}}+\sqrt{\eta}}}$, which is strictly below $\frac{1}{\sqrt{1-\frac{\alpha\gamma_{\frac{N-1}{2}}}{\gamma_{\frac{N+1}{2}}}+\sqrt{\eta}}}$. By Claim 4 (i), optimality of $\sigma_{\frac{N+1}{2},\frac{N+3}{2}}>0$ requires $\frac{\sigma_{1,2}}{\sqrt{\eta}}\geq\frac{1}{\sqrt{1-\frac{\alpha\gamma_{\frac{N-1}{2}}}{\gamma_{\frac{N+1}{2}}}+\sqrt{\eta}}}$. So $\lim_{\eta\to 0}\frac{\sigma_{N,N-1}}{\sigma_{1,2}}<1<\frac{1}{\sigma_{\frac{N+3}{2},\frac{N+1}{2}}\sigma_{\frac{N+1}{2},\frac{N-1}{2}}}\Rightarrow\lim_{\eta\to 0}\frac{\prod_{j=1}^{N-1}\sigma_{j,j+1}}{\prod_{j=2}^{N}\sigma_{j,j-1}}>1$, contradicting (A31).

So, $\sigma_{\frac{N+3}{2},\frac{N+1}{2}}=1$. Suppose next that $\sigma_{\frac{N+1}{2},\frac{N+3}{2}}\in(0,1)$. By Claim 4 (i), evaluated at $i=\frac{N+1}{2}$, this requires

$$\gamma_{\frac{N+1}{2}}^2\left(2\eta\sigma_{1,2}+(1-\eta)\sigma_{1,2}^2-\eta\right)-\alpha\gamma_{\frac{N-1}{2}}\gamma_{\frac{N+1}{2}}\sigma_{1,2}^2=0 \quad \text{(A32)}$$

By analogy to Claim 4 (ii), evaluated at $i=\frac{N+1}{2}$ and replacing $\sigma_{1,2}$ with $\sigma_{N,N-1}$, the optimality condition for $\sigma_{\frac{N+1}{2},\frac{N+3}{2}}\in(0,1)$ can also be written as

$$\gamma_{\frac{N+1}{2}}^2\eta-\alpha\gamma_{\frac{N-1}{2}}\gamma_{\frac{N+1}{2}}\sigma_{N,N-1}^2=0 \quad \text{(A33)}$$

By (A32), $\lim_{\eta\to 0}\frac{\sigma_{1,2}^2}{\eta}=\lim_{\eta\to 0}\frac{\gamma_{\frac{N+1}{2}}}{\gamma_{\frac{N+3}{2}}}$; by (A33), $\lim_{\eta\to 0}\frac{\sigma_{N,N-1}^2}{\eta}=\lim_{\eta\to 0}\frac{\gamma_{\frac{N+1}{2}}}{\alpha\gamma_{\frac{N-1}{2}}}$; thus,

$$\lim_{\eta\to 0}\frac{\sigma_{N,N-1}}{\sigma_{1,2}}=\lim_{\eta\to 0}\frac{\gamma_{\frac{N+3}{2}}}{\alpha\gamma_{\frac{N-1}{2}}}=\lim_{\eta\to 0}\frac{\gamma_{\frac{N+1}{2}}}{\alpha\gamma_{\frac{N-1}{2}}}-1>\frac{\rho}{1-\rho}$$

Where the final inequality follows from footnote (14). This is greater than 1 for $\rho>\frac{1}{2}$, implying that $\lim_{\eta\to 0}\frac{\sigma_{N,N-1}}{\sigma_{1,2}\sigma_{\frac{N+1}{2},\frac{N+3}{2}}}$ is strictly greater than 1, contradicting (A31).

Finally, suppose $\sigma_{\frac{N+1}{2},\frac{N-1}{2}}\in(0,1)$. First, since (A32) and (A33) simply give two different ways to write the optimality condition for $\sigma_{\frac{N+!}{2},\frac{N+3}{2}}>0$, they must be equivalent; setting the LHS's equal and taking limits as $\eta\to 0$, this requires

$$2=\lim_{\eta\to 0}\left(\frac{\sigma_{1,2}^2}{\eta}+\frac{\alpha\gamma_{\frac{N-1}{2}}}{\gamma_{\frac{N+1}{2}}}\left(\frac{\sigma_{N,N-1}^2}{\eta}-\frac{\sigma_{1,2}^2}{\eta}\right)\right) \quad \text{(A34)}$$

with $i\leq\frac{N+1}{2}$. Then an argument symmetric to the first paragraph above rules out $\sigma_{i,i-1}\in(0,1)$ for any state i other than the lowest with $a_i=1$ (here $\frac{N+3}{2}$), and an argument symmetric to the second paragraph above rules out $\sigma_{i,i+1}\in(0,1)$ for any state other than $i=\frac{N+1}{2}$.

By Claim 4 (ii), $\sigma_{\frac{N+1}{2},\frac{N-1}{2}} \in (0,1)$ requiers $\frac{\sigma_{1,2}^2}{\eta} = \frac{\gamma_{\frac{N+1}{2}}}{\alpha\gamma_{\frac{N-1}{2}}}$. By (A31), we need $\lim_{\eta\to 0} \frac{\sigma_{N,N-1}\sigma_{\frac{N+1}{2},\frac{N-1}{2}}}{\sigma_{1,2}} = 1$, which implies that $\lim_{\eta\to 0} \frac{\sigma_{N,N-1}}{\sigma_{1,2}} > 1$, or $\lim_{\eta\to 0} \left(\frac{\sigma_{N,N-1}^2}{\eta} - \frac{\sigma_{1,2}^2}{\eta} \right) > 0$. Then the RHS of (A34) is strictly greater than $\lim_{\eta\to 0} \frac{\sigma_{1,2}^2}{\eta}$, which is strictly above 2 by footnote (14), a contradiction. ■

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] Cowan, Nelson (1995) "Attention and Memory: an Integrated Framework." New York, NY: Oxford University Press.

[3] Dow, J. (1991) "Search Decisions with Limited Memory." *Review of Economic Studies* 58, 1-14.

[4] Edwards, W. (1968) "Conservatism in Human Information Processing." Reprinted in Kahneman, Slovic, and Tversky (1982), 359-369.

[5] Freidlin, M., and A.D. Wentzell (1984), "*Random Perturbations of Dynamical Systems"*. New York: Springer Verlag.

[6] Griffin, D. and A. Tversky (1992) "The Weighing of Evidence and the Determinants of Confidence," *Cognitive Psychology* 24, 411-435.

[7] Kandori,M., Mailath, G., and R. Rob (1993), "Learning, Mutations, and Long Run Equilibria in Games", *Econometrica*, Vol.61 29-56.

[8] Kahneman, Slovic, and Tversky (1982) "Judgement under Uncertainty: Heuristics and Biases." New York, NY: Cambridge University Press.

[9] 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.

[10] Lipman, B. (1995) "Information Processing and Bounded Rationality: A Survey." *Canadian Journal of Economics 28*, 42-67

[11] 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

[12] Mullainathan, S. (1998) "A Memory Based Model of Bounded Rationality", mimeo, M.I.T.

[13] Piccione, M. and A. Rubinstein (1993) "Finite Automata Play a Repeated Extensive Game." *Journal of Economic Theory* 61, 160-168

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

[15] Piccione, M. and A. Rubinstein (1997b) "The Absent-Minded Driver's Paradox: Synthesis and Reponses." *Games and Economic Behavior* 20, 121-130.

[16] Rabin, M. (1998) "Psychology and Economics." *Journal of Economic Literature* XXXVI, 11-46

[17] Rabin, M. and J. Schrag (1999), "First Impressions Matter: A Model of Confirmatory Bias," *Quarterly Journal of Economics* 114(1), 37-82.

[18] Rubinstein, A. (1998) "Modeling Bounded Rationality." M.I.T. Press.

[19] Stokey, N. and R. Lucas (1989), "*Recursive Methods in Economic Dynamics"*, Cambridge, MA: Harvard University Press.