

BOUNDED MEMORY AND BIASES IN INFORMATION PROCESSING

BY ANDREA WILSON¹

Before choosing among two actions with state-dependent payoffs, a Bayesian decision-maker with a finite memory sees a sequence of informative signals, ending each period with fixed chance. He summarizes information observed with a finite-state automaton.

I characterize the optimal protocol as an equilibrium of a dynamic game of imperfect recall; a new player runs each memory state each period. Players act as if maximizing expected payoffs in a common finite action decision problem. I characterize equilibrium play with many multinomial signals. The optimal protocol rationalizes many behavioral phenomena, like “stickiness,” salience, confirmation bias, and belief polarization.

KEYWORDS: Imperfect recall, bounded rationality, multiselves, bounded memory, absent-minded, biases.

1. INTRODUCTION

THIS PAPER DEVELOPS A SIMPLE MODEL of Bayesian learning with a finite memory. A decision-maker (DM) must choose between two actions with state-dependent payoffs. Ideally, he wishes to select the low and high actions, respectively, in the low and high states of the world. Before doing so, he is afforded the opportunity to observe a sequence of informative signals, concluding each period with a constant *termination chance*. A standard Bayesian agent would at each stage update the current posterior based on the newest signal and thereby eventually choose the action with highest expected payoff. By contrast, I explore what happens when the DM is constrained by a finite memory capacity. Specifically, he must summarize all signals so far into one of finitely many *memory states*, or information sets. He must optimally coarsen Bayes rule into a finite transition rule that specifies precisely how he will update his memory state in response to new information.

Formulating this as a decision problem naturally leads one to ponder an “absent-minded” decision-maker with imperfect recall of the past. Piccione and Rubinstein (1997) explored this strategic notion in a simple context, thinking of it as a dynamic game of “multiselves.” I adapt their strategic formulation to my infinite horizon problem, and then explore the role of incentive conditions and optimal beliefs.

This paper characterizes the optimal *memory protocol* for this context; specifically, a triple of an initial memory state, a decision rule should the game end in

¹This is a radical revision of my job market paper that I presented on the 2003 Review of Economic Studies Tour. I am indebted to Wolfgang Pesendorfer for extremely valuable guidance and advice throughout this project. I also thank Faruk Gul for his help in developing some ideas in the paper, and Ariel Rubinstein, the initial and final co-editors, five anonymous referees of this journal, and many seminar participants for useful comments. Throughout this editing, I have profited from the continued support of senior colleagues as well as my wonderful dog Mavi.

the current memory state, and a transition rule. But motivated by the absent-mindedness link, I immediately pursue a strategic reformulation. I assume that each memory state is run by a separate player (or “self”) each period, and that he enjoys a two-dimensional action space—decision and transition rules. Using methods from Markov chains on directed graphs to adapt the posterior belief construction of Piccione and Rubinstein to my infinite horizon setting, I develop a *team equilibrium* notion. I then argue in Proposition 1 that the DM’s optimal plan is a team equilibrium. Yet in the spirit of a few other dynamic games with a coordination flavor, the converse fails, as not all team equilibria yield optimal plans for the decision-maker.

I next turn to a few key results that are valid for all memory protocols and signal distributions. In the spirit of mechanism design, Lemma 1 argues that incentive compatibility alone—namely, that no self can profit from fully mimicking another’s strategy—implies that the *proportional* payoff gains from shifting to higher memory states must fall. Proposition 2 then leverages this to produce a simple characterization of how selves behave. Namely, each memory state is mapped onto a distinct interval of beliefs, and each self—after observing the signal—transitions to the state whose corresponding interval contains his posterior. Intuitively, the selves solve an optimal encoding problem, aware that no one later learns of their signal draw.

Next, restricting to my multinomial signals, I address whether all memory states are hit with positive probability in equilibrium. This requires that signals not be too asymmetric, and/or that the termination chance not be too large. Under an even stronger condition, I show in Lemma 2 that no memory state is absorbing and, thus, new signals can forever impact the memory state. Next, Proposition 3 derives limits on how well the DM does. Naturally, his payoff improves with a larger expected number of signals, but less obviously, I show that payoffs are forever strictly bounded below the unconstrained Bayesian full information payoff. This gap vanishes as the number of memory states explodes or any signal realization becomes perfectly informative. This underscores that the finite memory capacity in and of itself limits information accumulation, and cannot be circumvented by sufficiently many private signals.

My paper finally turns to the characterization of the optimal memory protocol with a vanishing termination chance, that is, an exploding expected number of signals. Proposition 4 then nearly fully characterizes the resulting decision and transition strategies. A key feature that emerges is that with enough signals, the optimal belief thresholds separating memory states spread apart: Only the two most extreme signal realizations are strong enough to push the DM out of his current state, and no observation is informative enough to push the DM past an adjacent state. Thus, the DM probabilistically shifts up or down by exactly one state after the two most powerful opposing signal realizations, and disregards all other observations. This speaks to the role of *salience* in learning, because nonextreme signals simply do not matter. In the limit with a vanishing termination chance (Corollary 3), the highest and lowest signals each leave the

DM precisely indifferent about switching memory states. Extremal memory states also play a special role. When maximally convinced of either high or low states of the world, the DM exhibits much inertia: He reacts to the strongest signal opposing his current belief with a probability vanishing as the square root of the termination chance.

Proposition 4 also offers a simple foundation for *fluid* versus *sticky* behavior—namely, whether the DM either always reacts to new extreme signals or only sometimes does. First of all, discussed in Section 6, Proposition 4 illustrates a *confirmation bias*: When the DM entertains a strong prior bias in favor of one state of the world, transitions in the opposite direction are sticky in some interior states. For a different perspective, observe how this updating behavior embodies *belief polarization*: Two individuals with different prior views may optimally move in opposite directions upon seeing the same information, each growing more convinced that his prior view is correct. Next, the stickiness in the extremal states implies that the order in which signals are seen is pivotal and, in particular, first impressions matter. If the DM initially sees enough high signals to reach his highest memory state, he will ignore a subsequent sequence of low signals with high probability, whereas if he had instead started out with a sequence of low signals, he might have reached his lowest state and disregarded the subsequent high signals. By contrast, standard Bayesian updating is obviously commutative in the observations. The optimal protocol also mimics a simpler bounded rationality in which the DM acts *as if* he could only remember a finite number of supporting facts. His beliefs adjust up or down as if he were mechanically replacing a previously stored opposing signal with the latest one.

Given the coarse way that past information can be stored, my paper reflects several themes of the longstanding informational herding model. For instance, the famous finding of incorrect herds with bounded likelihood ratios corresponds to my positive chance of incomplete learning with such signals. But unlike their setting, here one revealing extreme likelihood ratio leads to full learning in *both* states. And just as bad herds persist even with forward-looking behavior, so too here is full learning not approached even with enough signals. Technically, my finite memory states offer a simpler history summary, but my transition rule embeds a harder forward-looking optimization.

Hellman and Cover (1970) were the first to model learning with finitely many memory states. They assumed an unbounded number of signals, pursuing an approximate solution to the special question of estimating the correct state of the world. By contrast, I analyze the standard binary action decision problem, and reformulate it as an infinite dynamic game of imperfect recall. *I characterize optimal play in all memory states for any signal distribution and any termination chance.* I also derive sharper results for the limit model with a small termination chance and show that it captures many well known behavioral phenomena.

2. THE MODEL OF A LONG-LIVED DECISION-MAKER

A single infinitely-lived decision-maker (DM) is uncertain of the *true state* of the world θ , fixed for all time at either H or L . He must eventually choose a once-and-for-all action, either 1 or 0. The DM wishes to match the action to the state: the low action 0 is a safe action, yielding zero payoff in both states, and the high action 1 is a risky action paying $\pi^H > 0$ in state $\theta = H$ and $\pi^L < 0$ in state $\theta = L$. Call the indifference likelihood ratio $\ell^* = |\pi^L/\pi^H|$. Ex ante, the DM assigns probability p_0 to true state H and, thus, the *full information payoff* is $p_0\pi^H$. Further define the prior likelihood ratio $\ell_0 = p_0/(1 - p_0)$. Without loss of generality (WLOG), we assume throughout that the *prior bias* $\beta \equiv \ell_0\pi^H/|\pi^L|$ satisfies $\beta \geq 1$, so that the risky action is a priori optimal, yielding *myopic payoff* $p_0\pi^H + (1 - p_0)\pi^L > 0$.

A process of imperfectly informative independent and identically distributed (i.i.d.) signals ends each period with *termination chance* $\eta > 0$. In true state θ , the signal realization is $s \in \{1, 2, \dots, S\} \equiv \mathcal{S}$ with chance μ_s^θ .² The DM knows the termination chance η and the probabilities μ_s^θ . Signals are labeled such that their *likelihood ratios* $\xi(s) \equiv \mu_s^H/\mu_s^L$ are increasing in s ,³ so that higher realizations provide stronger evidence in favor of state H . I also make a standard full-support assumption that no signal is perfectly informative, so that $\mu_s^\theta > 0$ for all $s \in \mathcal{S}$ and $\theta \in \{H, L\}$; otherwise the model trivializes.

The DM cannot keep track of all of the information he receives. He has a finite set of available *memory states*, $\mathcal{M} \equiv \{1, 2, \dots, M\}$, and chooses an initial chance g_i^0 of starting in each state $i \in \mathcal{M}$,⁴ a *decision rule* d , which specifies the action choice $d_i \in \{0, 1\}$ if the problem terminates in memory state $i \in \mathcal{M}$,⁵ and a *transition rule* $\sigma: \mathcal{M} \times \mathcal{S} \rightarrow \Delta(\mathcal{M})$, which specifies the (possibly randomized) transition between memory states as new information is received.⁶ The DM wishes to design a *protocol* (g^0, σ, d) to maximize his expected payoff.⁷

The timing is as follows: at the start of a period, the DM learns whether the information process has ended. If it has, he chooses an action according to his *decision rule* d ; if it has not, he observes the signal and updates his memory according to the transition rule σ .

²For simplicity, I shall loosely refer to the signal realizations simply as signals.

³For beliefs about the true state of world θ , the information set is parenthesized. For θ -contingent chances of reaching (or payoffs starting in) a particular information set, the information set is a subscript and θ is a superscript.

⁴The mixture distribution $g^0 \in \Delta(\mathcal{N})$ is typically a point mass on one state.

⁵I could allow mixed action choices, with $d_i \in [0, 1]$ denoting the probability of choosing action 1 in memory state i . But since the payoff is linear in d_i (see (3)), it is clear that randomized actions can never help the DM.

⁶That the DM updates his memory after each signal is WLOG. If, for example, he can freely record two signals at a time, then all results go through, redefining the signal space \mathcal{S} as the set of pairs of signal realizations.

⁷I will suppress the functional dependence on the protocol (g^0, σ, d) throughout the paper.

That σ is stationary is an essential assumption; namely, the DM must follow the same transition rule every time he is in memory state i . For example, if the DM could condition his transition rule on the number of observed signals, then he could discriminate between infinitely many histories. But this would entirely circumvent the finite memory restriction.

3. PRELIMINARY ANALYSIS AND A STRATEGIC REFORMULATION

Let $\sigma_{i,j}^s$ be the chance that the DM moves from memory state i to j if he sees signal realization s . I shall denote $\sigma_{i,j}^s > 0$ more simply by $i \rightarrow_s j$. Let $\tau_{i,j}^\theta = (1 - \eta) \sum_{s \in S} \mu_s^\theta \sigma_{i,j}^s$ be the probability that the DM observes a signal and then moves from memory state i to j , given the true state of the world $\theta \in \{H, L\}$.⁸ Conditional on $\theta \in \{H, L\}$, the probability that the DM will be in memory state $i \in \mathcal{M}$ when the information process terminates is given by the *terminal probability*

$$(1) \quad f_i^\theta \equiv \sum_{t=0}^{\infty} \eta (1 - \eta)^t g_i^{\theta,t},$$

where $g_i^{\theta,t}$ is the probability of memory state $i \in \mathcal{M}$ in true state θ after t signal observations. Appendix A.2 derives these distributions, and proves that the terminal distribution $f^\theta = (f_i^\theta)_{i \in \mathcal{M}}$ is precisely the steady-state distribution of the following perturbed Markov process.⁹ With chance η , transitions are governed by the initial distribution g^0 , namely jumping to memory state i with chance g_i^0 ; with the remaining chance $1 - \eta$, transitions follow σ . Defining $\omega_{j,i}^\theta \equiv \eta g_i^0 + \tau_{j,i}^\theta$ as this *perturbed transition chance* from j to i , f^θ obeys the system of steady-state equations¹⁰

$$(2) \quad f_i^\theta = \sum_{j \in \mathcal{M}} f_j^\theta \omega_{j,i}^\theta \quad \forall i \in \mathcal{M}.$$

⁸The transition probabilities out of any state $(\sum_j \tau_{i,j}^\theta)$ sum to $(1 - \eta)$; the problem terminates with chance η .

⁹The proof simply calculates the infinite sum in (1), obtaining $f^\theta = \eta g^0 + (1 - \eta) T^\theta f^\theta$, where T is the matrix of transition chances specified by the DM, and then uses $\eta g^0 = \sum_{j=1}^M f_j^\theta \eta g^0$ (probabilities sum to 1) to write f^θ as the steady-state distribution of this perturbed Markov process.

¹⁰Following Freidlin and Wentzell (1984), the solutions f_i^θ to this system of steady-state equations can be written as $f_i^\theta = y_i^\theta / \sum_{j=1}^M y_j^\theta$, where y_j^θ is the θ -contingent probability of a "path" (connected acyclic graph) ending in state j . More precisely, define a j -tree as a directed graph on \mathcal{M} with no closed loops and such that each state $i \neq j$ has a unique successor k , indicated by $i \rightarrow k$. To calculate the *path chance* y_j^θ , take the sum, over all j -trees, of the product of the transition probabilities $\omega_{i,k}^\theta$ along the tree. For further explanation, see Kandori, Mailath, and Rob (1993).

Since the DM makes his decision in memory state $i \in \mathcal{M}$ with probability f_i^θ when the true state of the world is θ , whereupon the payoff is $d_i \pi^\theta$, his expected payoff is given by

$$(3) \quad p_0 \pi^H \sum_{i=1}^M f_i^H d_i + (1 - p_0) \pi^L \sum_{i=1}^M f_i^L d_i.$$

Appendix A.1 proves that the problem of maximizing this objective is isomorphic to a repeated memory constrained decision problem with flow payoffs rather than a terminal payoff.

A protocol (g^0, σ, d) maximizing (3) is *optimal*. For any memory size $M < \infty$ and termination chance η , let $\Pi^*(M, \eta)$ be the DM's *value*, that is, his expected payoff following an optimal protocol. The DM's objective is to choose transition chances and thereby via (1) to choose the optimal long-run distribution f^θ , subject to the constraint that he move from memory state i to j with chance at least ηg_j^0 after any signal observation. Since by (2) and (3), the expected payoff is continuous in σ and η when $\eta > 0$, the DM maximizes a continuous objective function over a compact constraint set; an optimal memory therefore exists. By continuity of Π in η , Berge's theorem of the maximum asserts upper hemicontinuity in η of the space of optimal protocols and continuity of the value $\Pi^*(M, \eta)$ in η . Since the DM can secure the myopic payoff $p_0 \pi^H + (1 - p_0) \pi^L$ by ignoring information, we have $\Pi^*(M, \eta) \geq p_0 \pi^H + (1 - p_0) \pi^L$. *Information matters* when this inequality is strict.¹¹

It will also be convenient to calculate the DM's beliefs and expected continuation payoffs in each memory state. I assume that the DM is Bayesian: he believes that if the true state of the world were θ , then he would have reached memory state $i \in \mathcal{M}$ after t signal observations with probability $g_i^{\theta, t}$, and so the *prospective* chance f_i^θ in (1) of ending in memory state $i \in \mathcal{M}$ is also the probability with which the DM believes he would reach memory state $i \in \mathcal{M}$ in any (unknown) period.¹² Then by Bayes' rule, letting $p(i)$ denote his updated belief that H is the true state conditional on memory state i , and letting $p(i, s)$ denote the posterior if he additionally observes signal realization s in memory

¹¹Define $\bar{\xi}(\eta) \equiv \max_{s^* \in S} (\eta + (1 - \eta) \sum_{s \geq s^*} \mu_s^H) / (\eta + (1 - \eta) \sum_{s \geq s^*} \mu_s^L)$. I show in Appendix D in the Supplemental Material (Wilson (2014)) that information matters if and only if $\beta < (\bar{\xi}(\eta) / \xi(1))^{M-1}$.

¹²This posterior belief assignment rule adapts the consistency notion of Piccione and Rubinstein (1997) for an infinite horizon setting: It assumes that the DM assigns chance $\eta(1 - \eta)^t$ to the event that he has observed t signal realizations so far. To understand this assumption, it may be helpful to think of the decision problem as "resetting" every time the information process terminates; then, since the problem resets with probability η each period, the chance that it last reset t periods ago is equal to $\eta(1 - \eta)^t$.

state i , we have

$$(4) \quad p(i) = \frac{p_0 f_i^H}{p_0 f_i^H + (1 - p_0) f_i^L} \quad \text{and} \\ p(i, s) = \frac{p_0 f_i^H \mu_s^H}{p_0 f_i^H \mu_s^H + (1 - p_0) f_i^L \mu_s^L}.$$

This yields the *likelihood ratios* $\ell(i) \equiv p(i)/(1 - p(i)) = \ell_0 f_i^H / f_i^L$ and $\ell(i, s) \equiv \ell(i) \xi(s)$.

Let v_i^θ denote the DM's expected *continuation payoff* starting in memory state $i \in \mathcal{M}$, given true state of the world $\theta \in \{H, L\}$. These obey the recursive system of equations

$$(5) \quad v_i^\theta = \eta \pi^\theta d_i + \sum_{j \in \mathcal{M}} \tau_{i,j}^\theta v_j^\theta \quad \forall i \in \mathcal{M} \text{ and for } \theta = H, L.$$

Let us now re-envision the DM's problem as an equilibrium outcome of a dynamic game. When the DM chooses an action or a transition, he knows only his current memory state; he can infer the set of possible histories, but cannot recall precisely which history occurred. He knows that any deviations today from the plan (g^0, σ, d) will not be later recalled and so cannot influence future behavior. I thus formulate this decision problem with imperfect recall as a dynamic game played by infinitely many *selves*, each controlling the transitions and decisions at only one memory state in one period, and taking all future and past behavior as given. For this game, define a *team equilibrium* as a tuple $((g^0, \sigma, d), (v_i^\theta), (p(i)))$ satisfying the following three conditions.

(i) If the DM moves from memory state i to j after signal realization s , then the new memory state j maximizes his expected continuation payoff $p(i, s) v_j^H + (1 - p(i, s)) v_j^L$.

(ii) For all memory states $i \in \mathcal{M}$, the decision rule d_i maximizes $d_i(p(i) \pi^H + (1 - p(i)) \pi^L)$.

(iii) Continuation payoffs (v_i^θ) and beliefs $(p(i))$ are induced by (g^0, σ, d) via (4) and (5).

Conditions (i) and (ii) define an incentive compatibility notion respecting the finite-state measurability and so require that there be no profitable (one-shot) deviations from (g^0, σ, d) .

PROPOSITION 1: If (g^0, σ, d) maximizes (3), $((g^0, \sigma, d), (v_i^\theta), (p(i)))$ is a team equilibrium.

This result—true for all signal distributions—is proved in Section A.4. It implies that the search for an optimal protocol (g^0, σ, d) is equivalent to the

search for a payoff-maximizing team equilibrium.¹³ *Its converse fails.*¹⁴ There is a trivial equilibrium in which the transition rule σ specifies always (in all states, after all observations) moving to each memory state with equal probability. This renders all memory states completely uninformative, with identical continuation payoffs.

Either Proposition 1 or, more simply, condition (i) has the following useful implication.

COROLLARY 1—Revelation Principle (RP): *In all memory states $i \in \mathcal{M}$ hit with positive probability, no self thinks he can gain by mimicking the strategy of any other self: $i \in \arg \max_{j \in \mathcal{M}} (p(i)v_j^H + (1 - p(i))v_j^L)$.*

This corollary is so labeled as it has the flavor of the revelation principle, insofar as each memory state (“type”) wishes to choose the action and transitions designated for his state.

4. GENERAL PROPERTIES OF AN OPTIMAL MEMORY

Label the memory states in \mathcal{M} with $p(i)$ weakly increasing in i : higher memory states assign (weakly) higher probabilities to true state H . So the DM is most strongly convinced of the true state of the world L in memory state 1 and most strongly convinced of H in memory state M . I call memory states 1 and M the *extremal states*, and call states in $\mathcal{M} \setminus \{1, M\}$ the *interior states*.

Call two memory states *equivalent* if they yield the same continuation payoffs: $v_i^\theta - v_j^\theta = 0$ for $\theta = H, L$. Easily, if memory states i, j are equivalent, then the DM could earn the same payoff with a modified rule $(\tilde{g}^0, \tilde{\sigma}, d)$ that never uses memory state j .¹⁵ For ease of exposition, I therefore restrict attention to optimal protocols in which no two memory states are equivalent. In equation (7) below, I will provide a condition under which equivalent or unused memory states are strictly suboptimal.

Lemma 1 derives monotonicity properties on the memory state values. Let $\Delta_{j,i}^\theta \equiv v_j^\theta - v_i^\theta$ be the *payoff differential* of memory state j over i in true state θ . So $\Delta_{j,i}^\theta \equiv -\Delta_{i,j}^\theta$. The next three results are general findings for any optimal memory, irrespective of the signal distribution.

¹³This “second welfare theorem,” shown in Section A.4, is a feature of common interest games. Crawford and Haller (1990) is an early example of this decentralization for a common interest game.

¹⁴This failure of the converse is familiar in team theory contexts; the classic reference is Radner (1962).

¹⁵Namely, obtain \tilde{g}^0 from g^0 by moving any initial probability of memory state j onto memory state i , and obtain $\tilde{\sigma}$ from σ by moving all transitions into memory state j over to memory state i . An easy calculation yields $f_j^H = f_j^L = 0$, so that memory state j is never used, and continuation payoffs are unchanged.

LEMMA 1—Payoff Monotonicity: *Fix an optimal protocol (g^0, σ, d) .*

(a) *Higher memory states are better in state H and worse in state L: $\Delta_{j,i}^H > 0 > \Delta_{j,i}^L$ for $j > i$.*

(b) *Gains proportionately decline in the memory state: $\Delta_{j,i}^H/\Delta_{i,j}^L \geq \Delta_{k,j}^H/\Delta_{j,k}^L$ if $i < j < k$.*

PROOF: Take memory states $i < j$, that is, $p(i) \leq p(j)$. By optimality of (g^0, σ, d) and Corollary 1, the DM prefers memory state j to i at belief $p(j)$, and prefers i to j at belief $p(i)$. Incentive conditions are

$$(6) \quad p(j)\Delta_{j,i}^H + (1 - p(j))\Delta_{j,i}^L \geq 0 \geq p(i)\Delta_{j,i}^H + (1 - p(i))\Delta_{j,i}^L.$$

Adding yields $(p(j) - p(i))(\Delta_{j,i}^H - \Delta_{j,i}^L) \geq 0$. So $\Delta_{j,i}^H \geq \Delta_{j,i}^L$ if $p(j) > p(i)$, while if $p(i) = p(j)$, we can order states so that this inequality holds. But $\Delta_{j,i}^\theta \neq 0$ for some $\theta \in \{H, L\}$, as we have ruled out equivalent states. Then $\Delta_{j,i}^H$ and $\Delta_{j,i}^L$ are not both weakly negative by the first inequality in (6) or both weakly positive by the second inequality in (6). Since $\Delta_{j,i}^H \geq \Delta_{j,i}^L$, this implies (a). For (b), take memory states $i < j < k$. Use (a) to rewrite the left-hand side (LHS) of (6) as $\ell(j)(\Delta_{j,i}^H/\Delta_{i,j}^L) \geq 1$ and the right-hand side (RHS) of (6) (replacing i, j by j, k) as $\ell(j)(\Delta_{k,j}^H/\Delta_{j,k}^L) \leq 1$. Q.E.D.

I now reformulate the choice facing a self at a memory state in terms of posterior beliefs. The action choice $d_i \in \{0, 1\}$ is straightforward, but the transition choice problem is less obvious. Define the *indifference likelihood ratio* $\bar{\ell}_i = \Delta_{i+1,i}^L/\Delta_{i+1,i}^H$ as the belief likelihood that leaves any self indifferent between memory states $i + 1$ and i . We next argue that a standard finite decision problem is induced, with the DM choosing higher memory states for higher likelihood ratios, given the proportionately falling gains to choosing action 1 versus action 0.

PROPOSITION 2—Optimal Transitions: *There are cutoffs $0 < \bar{\ell}_1 \leq \bar{\ell}_2 \leq \dots \leq \bar{\ell}_{M-1} < \infty$ so that a self transitions to memory state i if and only if (iff) his posterior likelihood ratio ℓ lies in $[\bar{\ell}_{i-1}, \bar{\ell}_i]$.*

PROOF: Having arrived at a posterior likelihood ℓ , a self prefers memory state $i + 1$ to i iff he expects a higher continuation payoff in state $i + 1$, that is, when $\ell\Delta_{i+1,i}^H \geq \Delta_{i+1,i}^L$. By Lemma 1(a), this happens if and only if $\ell \geq \bar{\ell}_i$, and by Lemma 1(b) with $j = i + 1$ and $k = i + 2$, the indifference likelihood ratios are monotone: $\bar{\ell}_i \leq \bar{\ell}_{i+1}$. So a self prefers memory state i to other memory states if his posterior likelihood lies in $[\bar{\ell}_{i-1}, \bar{\ell}_i]$; his transitions follow, by Proposition 1. Q.E.D.

The DM entertains a prior belief p_0 before entering this memory protocol. Later, however, all selves are endowed with priors by the protocol; by Corollary 1, they obey the following corollary.

COROLLARY 2—Bayesian Rational Prior Beliefs: *Before seeing a signal, the likelihood ratio $\ell(i)$ of the self in memory state i obeys $\bar{\ell}_{i-1} \leq \ell(i) \leq \bar{\ell}_i$.*

In light of these corollaries, no self ever “misinterprets” evidence of one true state as favoring the alternative. If the state i self sees a signal realization favoring true state H , that is, with $\xi(s) > 1$, then his posterior rises to $\ell(i, s) > \ell(i)$. His likelihood ratio then strictly exceeds $\bar{\ell}_{i-1}$, the largest cutoff for moving to a state below i . So he either ignores the signal s or moves up. Symmetrically, a self who sees evidence in favor of true state L either ignores it or moves down.

Proposition 2 greatly simplifies the problem in some cases: *equilibrium is fully characterized by a belief vector in \mathbb{R}^{M-1} with continuous signal distributions*. But I have assumed multinomial signals, requiring mixed strategies. I now restrict focus to this signal class. Under inequality (7), the DM’s expected payoff is strictly increasing in M (see Appendix B in the Supplemental Material); thus, equivalent memory states are not optimal, and all are optimally employed, so all selves obtain with positive chance:

$$(7) \quad \frac{\mu_S^L}{\mu_S^H} \left(\frac{\eta + (1 - \eta)\mu_1^H}{\eta + (1 - \eta)\mu_1^L} \right) \leq \frac{\mu_1^H}{\mu_1^L} \frac{\mu_S^H}{\mu_S^L} \leq \frac{\mu_1^L}{\mu_1^H} \left(\frac{\eta + (1 - \eta)\mu_S^H}{\eta + (1 - \eta)\mu_S^L} \right).$$

This sandwich inequality on the likelihood ratio product $\xi(1)\xi(S)$ holds for η near zero, or when the likelihood ratios of the extreme signal realizations 1 and S are not too asymmetric. I defer an intuition for it until I present a related inequality (8) below. If (7) fails, then there are priors and memory sizes M for which the DM could not improve his payoff with just one extra memory state.¹⁶

Call memory state $i \in \mathcal{M}$ *absorbing* if the DM either stays in state i or moves to an equivalent absorbing state, after all signal realizations. We next explore whether absorbing memory states can be optimal. To this end, say that there are *no dominant signals* whenever

$$(8) \quad \frac{\eta + (1 - \eta)\mu_1^H}{\eta + (1 - \eta)\mu_1^L} < \frac{\mu_1^H}{\mu_1^L} \frac{\mu_S^H}{\mu_S^L} < \frac{\eta + (1 - \eta)\mu_S^H}{\eta + (1 - \eta)\mu_S^L}.$$

This implies inequality (7), and so equivalent memory states are not optimal given no dominant signals. As seen in Appendix A.3, the first inequality in (8) guarantees that signal realization S is sometimes strong enough to push the DM out of memory state 1 and thus absorption in memory state 1 is not optimal. The second inequality likewise precludes absorption in memory state M .

¹⁶This is trivially the case if the prior bias β violates the “information matters” condition with $M + 1$ memory states, but can occur even when information matters; see Appendix B in the Supplemental Material for an example with both absorbing and equivalent states.

For an intuition, consider a binary signal, $\mu_1^H + \mu_S^H = 1$. Then the first inequality in (8) becomes

$$1 < \frac{\mu_1^H}{\mu_1^L} \frac{\mu_S^H}{\mu_S^L} \frac{1 - (1 - \eta)\mu_S^L}{1 - (1 - \eta)\mu_S^H} = \frac{\mu_1^H}{\mu_1^L} \left(\frac{\mu_S^H + (1 - \eta)(\mu_S^H)^2 + \dots}{\mu_S^L + (1 - \eta)(\mu_S^L)^2 + \dots} \right).$$

The final RHS term is the posterior belief (likelihood) adjustment for a DM who believes that he has seen S -signals in every period until termination. This is the strongest possible evidence in favor of true state H . If it cannot overwhelm an initial 1-signal, then the latter is a dominant signal: a DM who initially sees a 1-signal will always believe that the evidence favors true state L , and so absorption in memory state 1 could optimally occur for some priors. Conversely:

LEMMA 2—Absorbing Memory States: *If information matters and there are no dominant signals, that is, (8) holds, then there are no absorbing memory states in an optimal memory.*

For an instructive contrast, consider the informational herding model. Like this finite memory world, private signals in that well explored setting are observationally filtered in every period through a finite mesh—in Smith and Sørensen's (2000) case, via the observation of one of finitely many actions. Like here, a different decision-maker acts each period, unaware of previously viewed signals. Unlike here, their decision-makers can see the entire action history. And like Bikhchandani, Hirshleifer, and Welch (1992), I assume multinomial private signals. But they conclude that “cascades” obtain, namely, where additional signals are ignored. On the other hand, absorbing states are never optimal for small η —the proper parallel to the social learning setting with an infinite number of agents. The logic is that a Markov process eventually lands in the set of absorbing states. This renders useless the multitude of additional signals that would arise with vanishing termination chance η .¹⁷

We next bound equilibrium payoffs. A rational Bayesian almost surely learns the true state in the limit $\eta \rightarrow 0$ and so secures the full information payoff $p_0\pi^H$. My DM earns strictly less than this, by an amount depending on the prior bias $\beta = \ell_0\pi^H/|\pi^L|$ and the *information quality* $1/\rho^* = \lambda^{M-1}$, where $\lambda = \xi(S)/\xi(1) > 1$ is the quotient of the extreme likelihood ratios favoring state H . So $1/\rho^*$ is the relevant measure of the information quality afforded the agent: As either the signal informativeness or M explodes, so does $1/\rho^*$, and the maximized expected payoff tends to the full information payoff. If information matters, then $\rho^* < 1/\beta$.¹⁸

¹⁷One might think this is due to the forward-looking character of my agent. In fact, Smith, Sørensen, and Tian (2012) show that cascades remain possible even when a very patient social planner makes action choices.

¹⁸This follows immediately from the condition in footnote 11, noting that $\bar{\xi}(\eta) \leq \xi(S) \forall \eta$, with $\lim_{\eta \rightarrow 0} (\bar{\xi}(\eta)/\xi(1))^{M-1} = \lambda^{M-1} \equiv 1/\rho^*$.

PROPOSITION 3—A Payoff Upper Bound: *If $\rho^* \geq 1/\beta$, then the DM earns the myopic payoff $p_0\pi^H + (1 - p_0)\pi^L$. If $\rho^* < 1/\beta$ and if information matters, then $\Pi^*(M, \eta)$ is strictly falling in the termination chance η , and its supremum and limit as $\eta \rightarrow 0$ is*

$$(9) \quad \bar{\Pi}^* \equiv \lim_{\eta \rightarrow 0} \Pi^*(\eta, M) = p_0\pi^H \cdot \frac{(1 - \sqrt{\rho^*/\beta})^2}{1 - \rho^*}.$$

Hellman and Cover (1970) derived a similar bound,¹⁹ while exploring the optimal design of an automaton to distinguish between two hypotheses after an infinite sequence of signals; I discuss their result in more detail after Proposition 4.

As the transition chance η vanishes, the payoff upper bound $\Pi^*(M, \eta)$ increases in the information quality $1/\rho^*$. Learning is complete as $1/\rho^*$ explodes, for instance, as the signal becomes perfectly informative of either state ($\xi(1) \downarrow 0$ or $\xi(S) \uparrow \infty$). For a useful contrast to the informational herding literature, note that if either likelihood ratio explodes, learning is complete in both states. By contrast, in Smith and Sørensen (2000), when the support of the likelihood ratio density is boundedly positive but not boundedly finite (or conversely), there is incomplete learning in the low state (or high state). This holds even though the action history is observed in their setting but hidden in mine, which frustrates learning in my setting.

To understand the difference, consider an example with two memory states, representing the low action and the high action. Assume there is a nearly perfect signal for the high state, observed with chance ε near zero in the low state, but no such signal for the low state. In Smith and Sørensen (2000), the threshold $\hat{\ell}$ for switching from the high to the low action is exogenously specified: thus, if $\hat{\ell}$ is very low and if current beliefs (mistakenly) indicate the high state strongly enough, there may be no private signal strong enough to shift posteriors below $\hat{\ell}$. Thus, even in the low state of the world, public beliefs may converge to a bad cascade on the high action. My model, in contrast, has endogenously chosen thresholds. As the informativeness of the high signal explodes, so does the optimal threshold $\hat{\ell}$: if the DM anticipates being reluctant to switch out of the high action, he is correspondingly more careful about taking it in the first place, which in turn makes it easier for his posterior to drop below $\hat{\ell}$. In particular, by switching from the high action to the low action with a probability that is very small, but large compared to ε , the DM can guarantee ending up at the right action with chance near 1.

¹⁹Since they assume a continuous signal, their payoff bound replaces ξ_1 and ξ_S with $\inf_{\mathcal{A}} \frac{P^H(\mathcal{A})}{P^L(\mathcal{A})}$ and $\sup_{\mathcal{A}} \frac{P^H(\mathcal{A})}{P^L(\mathcal{A})}$, where P^H and P^L are the conditional probability measures, and \inf and \sup are over measurable sets \mathcal{A} .

5. OPTIMAL PROTOCOLS WITH LOW TERMINATION CHANCES

I next flesh out some crucial aspects of the optimal protocol with small termination chances. Among other things, we will see that only extreme signals matter and that the transition rule only takes unit steps.²⁰ Along with our finding in Corollary 1 that the DM can only move up (or down) after evidence in favor of true state H (or L), it then follows that he moves up by at most one state after signal realization S , down by at most one state after signal realization 1, and remains in his current memory state in all other cases. We need only determine the optimal transition chances. We find that transitions out of the extreme memory states 1 and M are very rare when the termination probability is small, while transitions in the interior memory states are deterministic, except to balance out asymmetries in either the prior bias or the signal.

To describe the optimal rule, define two critical cutoff memory states: the DM's initial state i_0 , obeying $\bar{\ell}_{i_0-1} \leq \ell_0 \leq \bar{\ell}_{i_0}$, and the lowest memory state i^* to choose action 1, satisfying $\bar{\ell}_{i^*-1} \leq \ell^* \leq \bar{\ell}_{i^*}$. In other words, the DM starts in the memory state i_0 yielding the highest continuation payoff for prior likelihood ratio ℓ_0 , while i^* is the lowest memory state with a belief above the threshold ℓ^* to prefer action 1 (by Corollary 1). Appendix A.8 reveals that if information matters, then

$$(10) \quad i^* \equiv \left\lfloor \frac{M+1}{2} - \frac{\log \sqrt{\beta}}{\log \lambda} \right\rfloor \quad \text{and} \quad i_0 \equiv \left\lfloor \frac{M+1}{2} + \frac{\log \sqrt{\beta}}{\log \lambda} - \frac{\log \xi(S)}{\log \lambda} \right\rfloor.$$

Observe that both i_0 and i^* move closer to the middle memory state as the signal grows more informative (so that λ increases), while an increase in the prior bias β —indicating a stronger prior in favor of H —pushes i^* further below the midpoint and pushes i_0 further above. My assumed prior bias $\beta > 1$ implies that generally $i^* < i_0$; define $I^* = \{i^*, i^* + 1, \dots, i_0\}$. Also note that $\beta < 1/\rho^* \equiv \lambda^{M-1}$. Then $i^* \geq 2$. In Supplemental Material Appendix C (see below equation (C.8)), I construct a related interior set of states \hat{I} that depends on the asymmetry between signals 1 and S .²¹

Call transitions in memory state i *sticky up* if $\sigma_{i,i+1}^S < 1$, *sticky down* if $\sigma_{i,i-1}^1 < 1$, and *fluid up* (or down) if not sticky up (not sticky down). Call signal

²⁰In Appendix E in the Supplemental Material, I offer a fully solved three-state example with a symmetric signal and general termination chance $\eta \in (0, 1)$. The example shows how team equilibria are constructed, indicates *how* small η must be to ignore intermediate signals, and shows that some features of the optimal protocol described in Proposition 4 persist even for larger values of η : namely, transitions out of the extremal states are always sticky, and the DM never jumps between states 1 and 3. The example also demonstrates the impossibility of a converse to Proposition 1 by showing that while there is an asymptotically efficient symmetric team equilibrium, it is Pareto-dominated by an asymmetric memory protocol.

²¹The final section of the Supplemental Material shows that for symmetric signals, $\hat{I} \supseteq \{i^* + 1, \dots, i_0 - 1\}$. As signal 1 grows weaker compared to S , the set I typically expands.

S a geometrically more rare signal 1 if $\mu_S^H \mu_S^L \leq \mu_1^L \mu_1^H$. This holds for symmetric signals, with $\mu_i^H \equiv \mu_{S+1-i}^L$. For a binary signal, with $\mu_1^H + \mu_S^H = 1$, it asserts that signal S more strongly indicates state H than signal 1 indicates state L .²² The Supplemental Material example exhibits sticky transitions into the middle state.

PROPOSITION 4—Optimal Protocols: Fix a memory size M and memory ratio $r^* < 1/\beta$. There exists $\eta^* > 0$ such that when $\eta \in (0, \eta^*)$, the DM begins in memory state i_0 . Also, the following statements hold:

(a) *Only extreme signals matter:* In all memory states $i \in \mathcal{M}$, the DM reacts only to signal realizations $s = 1, S$ and remains in memory state i after any signal $s \in S \setminus \{1, S\}$.

(b) *No jumps:* In all memory states $i \in \mathcal{M}$, the DM either stays in i or moves up one state after signal realization $s = S$, and stays or moves down one state after $s = 1$.

(c) *No interior memory state is both sticky up and sticky down, and interior transition chances are bounded away from zero as $\eta \rightarrow 0$.* If signal S is geometrically more rare than signal 1, then all interior states are fluid up, all interior states outside \hat{I} are fluid down, and all interior states in \hat{I} are sticky down. If signal 1 is geometrically more rare than signal S , then interior states outside of I^* are fluid down.

(d) Let $(g_\eta^0, \sigma_\eta, d_\eta)$ be a sequence of optimal protocols with $\eta \rightarrow 0$. If $M \geq 3$, then the transition chances out of extreme memory states 1 and M vanish at a rate asymptotic to $\sqrt{\eta}$, while transition chances in the interior memory states are bounded away from 0.

Recall the convention that higher memory states correspond to higher posteriors in favor of the true state H . Parts (b) and (c) then say that the DM moves up by at most one step if he observes the strongest available evidence in favor of state H (an S -signal), down by at most one step if he observes the strongest evidence for L (a 1-signal), and ignores all other signal realizations. For an intuition, consider the world with an arbitrary but not necessarily small termination chance $\eta > 0$. In any state, the DM can in principle choose any one of M transitions. Since $p(i)$ is weakly increasing in i , the DM clearly should rank signals, with higher signals leading to higher transitions. Now, as the termination chance η falls, the DM secures a higher expected payoff precisely because the beliefs $p(i)$ incur a “spread,” with $p(1)$ falling and $p(M)$ rising. As a result, the DM grows more discriminating in his willingness to take larger memory transitions, and eventually, for small enough η , no longer jumps over memory states. At the same time, only the extreme signals are able to push

²²Indeed, $\mu_S^H/\mu_S^L \geq \mu_1^L/\mu_1^H \Leftrightarrow \mu_S^H(1 - \mu_S^H) \geq \mu_S^L(1 - \mu_S^L) \Leftrightarrow \mu_S^H - \mu_S^L \geq (\mu_S^H)^2 - (\mu_S^L)^2 \Leftrightarrow 1 \geq \mu_S^H + \mu_S^L \Leftrightarrow \mu_1^H \mu_1^L = (1 - \mu_S^H)(1 - \mu_S^L) \geq \mu_S^H \mu_S^L$.

the DM high enough to move up to the next belief $p(i + 1)$ or move down to $p(i - 1)$.²³

Part (c) only partially characterizes the transitions at interior memory states. Outside of I^* , the DM moves deterministically after the more powerful of the two signals $s \in \{1, S\}$. That is, transitions are fluid up if signal S is stronger, and fluid down if signal 1 is stronger; transitions may be sticky in response to the weaker signal, but a precise characterization is difficult and depends on the signal asymmetry. But when the prior bias $\beta > 1$ *strongly* favors the true state H , there is a nonempty block of interior memory states I^* in which transitions are sticky down even if the signal realization 1 is stronger than S .²⁴ In other words, inside I^* , the DM reacts only probabilistically to evidence against his prior bias. In the special case of a binary symmetric signal, transitions are fluid both up and down in all states outside of I^* , while all states in I^* are sticky down.

For insight into part (d), consider memory state 1, in which the DM is maximally convinced of true state L . He is thus unable to react to any additional evidence in favor of true state L and wishes to leave memory state 1 only when the net observed information favors H . A rule that *never* leaves memory state 1 cannot be optimal, as this makes it so uninformative that the DM will find it optimal to leave, for we have seen in Lemma 2 that absorbing states are sub-optimal for small η . Conversely, a rule that *always* leaves memory state 1 after high signals cannot be optimal, for then whenever the DM is in memory state 1, he can infer that he has never seen a high signal, yet has likely seen a large number of low signals; thus, he will not leave memory state 1 after just one high signal. The optimal rule therefore requires randomization. The chance of leaving memory state 1 after high signals is chosen so that in expectation, by the time the DM actually leaves, he believes that he has likely seen equal amounts of evidence in favor of true states H and L . With small termination chance η , this requires an exit chance asymptotically proportional to $\sqrt{\eta}$.

Observe the important role of the strategic aspects of this decision problem. No Bayesian decision-maker ever passes up free information since it is always valuable. But here, the DM is willing to end the problem (and decide early) rather than randomize in states 1 and M ; indeed, he does so, by Proposi-

²³Another perspective is that since (by Proposition 4(d)) the probability of ending up in state 1 or M tends to 1 as $\eta \rightarrow 0$, the goal for small η is to maximize the relative likelihood of ending up in state M (instead of state 1) in the high state of the world and to minimize this likelihood in the low state. This is achieved by ensuring that the DM only moves from state 1 to M (or from M to 1) after the most extreme available sequences of evidence for true state H (or L). In particular, by moving only one state at a time, only after extreme signals, the ratio $f_M^H f_1^L / f_M^L f_1^H$ tends to $(\xi(S)/\xi(1))^{M-1}$ as $\eta \rightarrow 0$; this ratio falls if the DM jumps states or reacts to intermediate signals, as it becomes easier to mistakenly move between states 1 and M . See also the illustration in Appendix E (Step 1 of the example) in the Supplemental Material.

²⁴This asymmetry reflects the assumed prior bias $\beta > 1$. Had the prior bias strongly favored the true state L , transitions would have been sticky up at an interior block of memory states.

tion 4(d). So the DM in these memory states must be indifferent about moving and yet the DM can only profit from further signals by moving.²⁵

Hellman and Cover (1970) characterized a class of ε -optimal automata, satisfying analogs to Proposition 4(a) and (b), and a modification of (d). They found that transition chances out of memory states 1 and M must be small to generate payoffs within any $\varepsilon > 0$ of $\lim_{\eta \rightarrow 0} \Pi^*(\eta, M)$. They require ε -optimality, since transition chances out of memory states 1 and M are forever positive, but vanish as $\eta \rightarrow 0$: a payoff discontinuity arises when states 1 and M are absorbing.

Proposition 4 yields a simple expression for Bayesian beliefs and cutoffs in Proposition 2.

COROLLARY 3—Limit Beliefs: *As $\eta \rightarrow 0$, the likelihood ratio $\ell(i)$ with an optimal memory tends to $\sqrt{\ell_0|\pi^L|/\pi^H} \lambda^{i-(M+1)/2}$ and the indifference ratio $\bar{\ell}_i$ tends to $\xi(S)\ell(i) = \xi(1)\ell(i+1)$.*

So each state- i self starts out with the prior $\ell(i)$, then observes the signal and moves to memory state j if his posterior likelihood ratio belongs to $[\bar{\ell}_{j-1}, \bar{\ell}_j]$. As the termination chance $\eta \rightarrow 0$, these intervals tend to $[\bar{\ell}_{j-1}, \bar{\ell}_j] \rightarrow [\xi(1)\ell(j), \xi(S)\ell(j)]$. Critically observe that in this limit, there arises indifference after the extremal signals 1, S , consistent with the thrust of Proposition 4(a) and (c). This is the reason why no memory state jumps arise *and* why no nonextremal signals matter. Notice also that the belief likelihoods $\ell(j)$ grow geometrically in j , with $\ell(j+1)/\ell(j) = \xi(S)/\xi(1)$.

This likelihood equation in particular implies that if η is near zero and the DM advances from memory state i to $i+1$ after observing an S -signal, then his beliefs adjust as if he were a Bayesian agent who observed the S -signal, but also *forgot* a previously observed 1-signal. The optimal memory thereby symmetrically treats the signals $s = 1, S$ even if their informativeness radically differs. So the optimal finite-state memory behaves like a limited capacity memory: the DM stores $M - 1$ signals, and when the memory is full, can only learn a new one by replacing a previously stored one. To this end, observe that the likelihood ratios $\ell(i)$ are uniformly bounded and finite for all memory states i , for fixed M . In other words, expecting to observe an arbitrarily large number of signals (with vanishing η) is not an avenue toward complete learning, which can only be secured by increasing the memory size M .

Next, observe how the prior and the payoffs affect beliefs across states: When action 1 is riskier, namely the payoff ratio $|\pi^L|/\pi^H$ is higher, all memory states are associated to higher probabilities and, thus, can better identify

²⁵Precisely, it turns out that instead of moving from memory state 1 to 2 with probability $\sigma_{1,2}^S$, the DM could earn exactly the same payoff by modifying this transition chance to $\hat{\sigma}_{1,2}^S$ and also terminating the problem in memory state 1 with a probability ϕ_1 satisfying $\frac{\eta(1-\phi_1)}{\eta+(1-\eta)\phi_1} \hat{\sigma}_{1,2}^S = \sigma_{1,2}^S$.

state H . Similarly, when the DM's prior on state H rises with higher ℓ_0 , associated memory state probabilities shift up.

6. BEHAVIORAL PREDICTIONS

An optimized finite-state memory sheds light on some well known behavioral biases. First, *people tend to ignore all but the most salient and striking information events*. For instance, Proposition 4 predicts that when the DM's memory is optimized for a small termination chance $\eta > 0$, he optimally reacts only to the two most extreme signal realizations. In essence, the DM avoids wasting limited memory resources on all but the two most informative observations.

Next, *the order in which information is observed matters and, in particular, "first impressions matter."* By Proposition 4, if the highest and lowest signal realizations $s = 1, S$ similarly favor the respective true states L and H , then a DM who starts in an interior memory state will move up one state after each high observation and move down one state after each low observation. Also, once he reaches one of the extreme memory states 1 or M , he exits each period with only a small probability proportional to $\sqrt{\eta}$. Now, fix a time $t < \infty$ and $\varepsilon > 0$, and assume that the signal sequence begins with enough consecutive high observations $s = S$ so as to land the DM in memory state M after at most $M - 1$ steps. Then for $\eta > 0$ small enough, the DM ignores any subsequent low signals $s = 1$ with high probability, and thus remains there until time t with chance at least $1 - \varepsilon$. Similarly, if the sequence instead begins with $M - 1$ consecutive low signal observations, then the DM ends up in memory state 1 with probability at least $1 - \varepsilon$. While this contrasts with a Bayesian DM, for whom the order of signal observations does not matter, it agrees with the predictions of the Bayesian social learning models. In both cases, a coarse observed historical signal is critical.

Third, *people tend to see what they want to see*. Proposition 4(c) asserts that for large enough $\beta > 1$, there is a nonempty set of states I^* in which transitions are sticky down, while the DM reacts with probability 1 to observations supporting his initial bias. A symmetric result obtains for low enough $\beta < 1$, yielding an interior block of states in which transitions are sticky up, but not sticky down. In other words, given a strong prior bias toward one of the two action choices, the DM optimally (probabilistically) ignores evidence favoring the alternative choice. All told, we see that a *confirmation bias* optimally emerges as a best response to bounded memory.

In a related finding, this memory model predicts *belief polarization*: People with conflicting initial views can observe exactly the same information and then move in *opposite directions*. Indeed, consider two DM's, one with large $\beta > 1$ and one with small $\beta < 1$. After seeing the same two extreme signal realizations $s = 1, S$, the first agent moves up in expectation (surely reacting to the high observation $s = S$, but sometimes ignoring $s = 1$), while the second agent moves down in expectation, reacting for sure only to the low observation. In

other words, each reads the evidence as supporting his hypothesis and contradicting the alternative.²⁶ Note that polarization occurs even though the signal distribution μ^θ is known; there is no ambiguity in the interpretation of evidence. In the extreme case, with sufficiently opposed prior biases, agreement is essentially impossible between the two agents: one only moves up and the other only moves down.

7. RELATED LITERATURE AND CONCLUSION

7.1. *Computer Science*

Hellman and Cover (1970) introduced the problem of the optimal design of a finite-state automaton—in their case, maximizing the chance of correctly identifying a state $\theta \in \{H, L\}$ after an infinite sequence of informative signals. They first derived a payoff upper bound and then showed that one can achieve a payoff arbitrarily close to this number. Their paper corresponds to the limit as $\eta \downarrow 0$ of my model with identical payoffs $\pi^H = |\pi^L|$. As I explained after Proposition 4, no optimal solution exists at this limit, thus explaining why Hellman and Cover focused exclusively on ε -optimal rules. Despite having continuous rather than multinomial signals, as I do, their results are comparable to the payoff bound in my Proposition 3, and to the transition rule characterizations in parts (a), (b), and (d) of Proposition 4. In particular, my finding that the DM only reacts to the two most informative signal realizations when η is small corresponds to their result that only signal realizations with sufficiently extreme likelihood ratios matter.

My paper makes many technical advances over Hellman and Cover. I show that a robust class of perturbations of their environment yields existence, and essentially uniqueness, of an optimal solution, and not just an ε -optimal one. With a small termination chance $\eta > 0$, transition chances out of the extremal memory states are asymptotically proportional to $\sqrt{\eta}$. I also find that transitions are fluid at most interior memory states, whereas Hellman and Cover note only that ε -optimality requires transition chances in the interior states are large compared to those in extremal states.

Technical robustness aside, my primary novel contribution is the strategic formulation and characterization of the learning problem. I have recast it as a dynamic game with imperfect recall, in which selves must optimally encode new signals by wisely passing the baton to the next agent or by retaining it. I have computed the optimal decision rule of each realized self: In the equivalent team equilibrium, all selves act as if in a time invariant M -action decision problem with fixed but endogenous payoffs (the baton toss and the memory state values).

²⁶For instance, in a famous experiment, Lord, Ross, and Lepper (1979) allowed two groups of people—one in favor of capital punishment and one opposed—to read identical studies for and against. Many of the initial proponents moved even more strongly in favor of capital punishment, while many opponents became even more opposed.

7.2. *Psychology*

My model shares some structure of two memory agendas in the psychology literature, and also realizes some of their goals. Cowan (1995) explores finite-capacity memory models of *short-term memory*. My paper provides a novel rational foundation for this assumption. I argued that the DM optimally behaves as if he has a finite-capacity memory, and must forget a previously stored extremal observation whenever he learns a new one.

I also relate to the *long-term working memory* (LTWM) literature (see Chase and Ericsson (1982)). Ericsson and Kintsch (1995) propose an LTWM model consisting of a set of *nodes*, each representing a block of information stored in long-term memory; only the memory node can be recalled, and not the original information sequence. This resembles my assumed structure with finitely many memory states and the Markovian transition rule restriction. Ericsson and Staszewski (1989) pointedly argue that “to meet the particular demands for working memory in a given skilled activity, subjects must acquire encoding methods and retrieval structures that allow efficient storage and retrieval.” Indeed, my optimization precisely solves for the *optimal* such encoding. And consistent with my optimization of the memory for the *specific decision problem*, they observe that “LTWM is therefore closely tailored to the demands of a specific activity.”

7.3. *Economics*

Economists have explored models in which information histories must be bundled, such as the discussed informational herding literature. As seen, my model offers an infinite horizon application of the solution concept in Piccione and Rubinstein’s (1997) model of imperfect recall in the absent-minded driver’s paradox. I have applied their solution concept and shown that a steady-state Markov model captures their belief restriction. Dow (1991) introduced a simple sequential price search model in which the DM could only recall whether he had categorized past prices as low or high, and must optimally design the categories.²⁷ Finally, Mullainathan (1998) and Rabin and Schrag (1999) explicitly build models around biased inferences; my paper shows that many systematic biases are consistent with Bayesian rationality subject to a constraint and, hence, do not necessarily imply that people are fundamentally unable to make probability judgments.

²⁷For another approach to decision-making through categories, see the case-based and analogy-based models in Gilboa and Schmeidler (1995) and Jehiel (2005).

APPENDIX A: OMITTED PROOFS

A.1. An Equivalent Model With Flow Payoffs

Consider a repeated problem: the DM chooses among actions 0 and 1 every period, earning payoffs 0 and π^θ , respectively. In true state θ , let X_t^θ be the expected payoff if the problem terminates after t signals, and let x_τ^θ denote the expected payoff increment earned in period $\tau \leq t$. By linearity of expectations, $X_t^\theta = \sum_{\tau=0}^t x_\tau^\theta$. Recalling from (1) that $g_i^{\theta,\tau}$ is the chance of memory state i after τ observations, so $x_\tau^\theta = \sum_{i=1}^M d_i \pi^\theta g_i^{\theta,\tau}$, the expected payoff in true state θ is

$$\begin{aligned} (11) \quad \sum_{t=0}^{\infty} \eta(1-\eta)^t X_t^\theta &= \sum_{t=0}^{\infty} \eta(1-\eta)^t \sum_{\tau=0}^t \left(\sum_{i=1}^M d_i \pi^\theta g_i^{\theta,\tau} \right) \\ &= \sum_{i=1}^M d_i \pi^\theta \sum_{t=0}^{\infty} \sum_{\tau=0}^t \eta(1-\eta)^t g_i^{\theta,\tau}. \end{aligned}$$

Changing the order of the summation,

$$\begin{aligned} \sum_{t=0}^{\infty} \sum_{\tau=0}^t \eta(1-\eta)^t g_i^{\theta,\tau} &= \sum_{\tau=0}^{\infty} \left(\sum_{t=\tau}^{\infty} \eta(1-\eta)^t \right) g_i^{\theta,\tau} \\ &= \sum_{\tau=0}^{\infty} (1-\eta)^\tau g_i^{\theta,\tau} \equiv f_i^\theta / \eta, \end{aligned}$$

final equality by (1). Substituting into (11), the DM's payoff in each state θ is then $\frac{1}{\eta}$ times his one-shot payoff $\sum_{i=1}^M f_i^\theta d_i \pi^\theta$, establishing the desired equivalence.

A.2. The Terminal Distribution

LEMMA 3: *The sum in (1) converges for any (g^0, σ, d) , yielding a unique distribution f^θ . This distribution f^θ equals the steady-state distribution of a perturbed Markov process with transition probabilities $\omega_{i,j}^\theta \equiv \eta g_{i,j}^\theta + \tau_{i,j}^\theta$ from i to j .*

PROOF: Fix a memory size M , termination chance η , and protocol (g^0, σ, d) . Let T^θ be the induced transition matrix with (j, i) th entry $\tau_{i,j}^\theta = (1-\eta) \times \sum_{s \in \mathcal{S}} \mu_s^\theta \sigma_{i,j}^s$. Then the distribution $g^{\theta,t}$ after t observations satisfies $g^{\theta,t} = (T^\theta / (1-\eta))^t g^0$, so by (1), $f^\theta = \sum_{t=0}^{\infty} \eta(1-\eta)^t g^{\theta,t} = \sum_{t=0}^{\infty} (T^\theta)^t \eta g^0$. Now, transform T^θ into an irreducible matrix by deleting row/column j whenever $\tau_{i,j}^\theta = 0 \forall i$, and redefine this as T^θ . Since the column entries in T sum to $(1-\eta) < 1$, T^θ is an irreducible substochastic matrix, and hence invertible, with inverse $(I - T^\theta)^{-1} = \sum_{t=0}^{\infty} (T^\theta)^t$. Substituting into our expression for f^θ , we obtain $f^\theta = (I - T^\theta)^{-1}(\eta g^0) \Leftrightarrow f^\theta = \eta g^0 + T^\theta f^\theta$. For state $i \in \mathcal{M}$, this

yields $f_i^\theta = \eta g_i^\theta + \sum_j f_j^\theta \tau_{j,i}^\theta$, which rearranges using $\sum_j f_j^\theta = 1$ to the desired expression: $f_i^\theta = \sum_j f_j^\theta (\eta g_i^\theta + \tau_{j,i}^\theta) \equiv \sum_j f_j^\theta \omega_{j,i}^\theta$. *Q.E.D.*

A.3. Absorbing Memory States: Proof of Lemma 2

Fix (M, η) , assume that information matters, and let (g^θ, σ, d) be a protocol satisfying team equilibrium conditions (i) and (ii). I will prove that if state 1 is absorbing, then the first inequality in (8) is violated; together with a symmetric argument for state M , (8) thus rules out team equilibria with absorbing states.²⁸ Lemma 2 then follows immediately by Proposition 1.

STEP 1: Define $j^* = \arg \min_j (f_j^H \tau_{j,1}^H) / (f_j^L \tau_{j,1}^L)$. If memory state 1 is absorbing, then

$$(12) \quad 1 \geq \left(\frac{\tau_{j^*,1}^H \mu_S^H}{\tau_{j^*,1}^L \mu_S^L} \frac{\eta + \tau_{j^*,1}^L}{\eta + \tau_{j^*,1}^H} \right) \left(\frac{p_0 \frac{f_{j^*}^H}{f_{j^*}^L} \frac{\eta d_{j^*} \pi^H + \sum_{i \geq 2} \tau_{j^*,i}^H \Delta_{i,j^*}^H}{\eta d_{j^*} |\pi^L| + \sum_{i \geq 2} \tau_{j^*,i}^L \Delta_{i,j^*}^L}}{1 - p_0} \right).$$

PROOF: If memory state 1 is absorbing, then it must be optimal to stay after an S -signal, requiring $1 \geq \ell(1, S) \Delta_{j,1}^H / \Delta_{1,j}^L \forall j \in \mathcal{M}$ (by Lemma 1(a)). To compute $\Delta_{j,1}^\theta$, note that $v_1^\theta = 0$ (since state 1 is absorbing and chooses action 0), while for any other memory state j , (5) yields $v_j^\theta = \eta d_j \pi^\theta + \sum_i \tau_{j,i}^\theta v_i^\theta$. Subtracting $(\sum_{i \geq 2} \tau_{j,i}^\theta) v_j^\theta = (1 - \eta - \tau_{j,1}^\theta) v_j^\theta$ from both sides, we obtain

$$(13) \quad \Delta_{j,1}^\theta \equiv v_j^\theta - v_1^\theta = \left(\eta d_j \pi^\theta + \sum_{i \geq 2} \tau_{j,i}^\theta \Delta_{i,j}^\theta \right) / (\eta + \tau_{j,1}^\theta).$$

To compute $\ell(1, S)$, recall from (1) that $f_1^\theta \sum_{j \geq 2} \omega_{1,j}^\theta = \sum_{j \geq 2} f_j^\theta \omega_{j,1}^\theta$. Since memory state 1 is absorbing but not an initial state, this becomes $f_1^\theta \eta = \sum_{j \geq 2} f_j^\theta \tau_{j,1}^\theta$. Taking ratios and using (4) for the first expression below, we deduce

$$(14) \quad \ell(1, S) \equiv \ell_0 \xi(S) \frac{f_1^H}{f_1^L} = \frac{p_0 \mu_S^H}{1 - p_0 \mu_S^L} \frac{\sum_{j \geq 2} f_j^H \tau_{j,1}^H}{\sum_{j \geq 2} f_j^L \tau_{j,1}^L} \geq \frac{p_0 \frac{f_{j^*}^H}{f_{j^*}^L} \tau_{j^*,1}^H \mu_S^H}{1 - p_0 \frac{f_{j^*}^L}{f_{j^*}^L} \tau_{j^*,1}^L \mu_S^L}.$$

Together with (13) (evaluated at $j = j^*$), we obtain the RHS in (12) as a lower bound on $\ell(1, S) \Delta_{j^*,1}^H / \Delta_{1,j^*}^L$, which then cannot exceed 1. *Q.E.D.*

²⁸Note that it suffices to consider absorbing extremal states. If some state $1 < j < M$ is absorbing, then so too must be state 1 (if $d_j = 0$) or M (if $d_j = 1$): If the state- j self is unwilling to ever change his action, then so are all selves with more extreme posteriors in favor of action d_j .

STEP 2: *The second term in parentheses in (12) is at least 1.*

PROOF: By condition (ii) for a team equilibrium, $d_j > 0$ implies $\ell(j)\pi^H/|\pi^L| \geq 1$. Using $\ell(j) \equiv (p_0 f_j^H / (1 - p_0) f_j^L)$ (by (4)), this rearranges as

$$(15) \quad d_j p_0 f_j^H \pi^H \geq d_j (1 - p_0) f_j^L |\pi^L|.$$

And by condition (i) for a team equilibrium, a transition $j \rightarrow_s i$ requires that payoffs be weakly higher in state i than j at posterior $p(j, s)$. Thus

$$\begin{aligned} \sigma_{j,i}^s > 0 \quad \Rightarrow \quad 0 &\leq p(j, s) \Delta_{i,j}^H + (1 - p(j, s)) \Delta_{i,j}^L \\ &\propto p_0 f_j^H \mu_s^H \Delta_{i,j}^H + (1 - p_0) f_j^L \mu_s^L \Delta_{i,j}^L. \end{aligned}$$

Summing over all signal realizations s and recalling $\tau_{j,i}^\theta \equiv (1 - \eta) \sum_s \mu_s^\theta \sigma_{j,i}^s$, we obtain $p_0 f_j^H \tau_{j,i}^H \Delta_{i,j}^H \geq (1 - p_0) f_j^L \tau_{j,i}^L \Delta_{i,j}^L$. Finally, summing this inequality over all states $i \geq 2$ and then adding (15) yields

$$(16) \quad p_0 f_j^H \left(\eta d_j \pi^H + \sum_{i \geq 2} \tau_{j,i}^H \Delta_{i,j}^H \right) \geq (1 - p_0) f_j^L \left(\eta d_j |\pi^L| + \sum_{i \geq 2} \tau_{j,i}^L \Delta_{i,j}^L \right).$$

At $j = j^*$, this rearranges to the desired inequality when the RHS is positive, which follows from (13) (the RHS term in parentheses is $\Delta_{1,j}^L(\eta + \tau_{j,1}^L)$) and Lemma 1(a) ($\Delta_{1,j}^L > 0$). Q.E.D.

STEP 3—Completing the Proof: *If state 1 is absorbing, then the following violation of (8) obtains:*

$$(17) \quad 1 \geq \frac{\tau_{j^*,1}^H \mu_s^H}{\tau_{j^*,1}^L \mu_s^L} \frac{\eta + \tau_{j^*,1}^L}{\eta + \tau_{j^*,1}^H} > \frac{\mu_1^H \mu_s^H}{\mu_1^L \mu_s^L} \frac{\eta + (1 - \eta) \mu_1^L}{\eta + (1 - \eta) \mu_1^H}.$$

PROOF: The first inequality in (17) is by Steps 1 and 2. So it remains to establish the second inequality, which rearranges as

$$(18) \quad \begin{aligned} &\tau_{j^*,1}^H \mu_1^L \left(\frac{\eta}{1 - \eta} + \tau_{j^*,1}^L \right) \left(\frac{\eta}{1 - \eta} + \mu_1^H \right) \\ &\quad - \tau_{j^*,1}^L \mu_1^H \left(\frac{\eta}{1 - \eta} + \tau_{j^*,1}^H \right) \left(\frac{\eta}{1 - \eta} + \mu_1^L \right) > 0. \end{aligned}$$

For this, I first show that $\sigma_{j^*,1}^1 = 1$. Since state 1 is absorbing, we need $\ell(1, S) \leq \bar{\ell}_1$ by Proposition 2; but then by (14), using $\mu_s^H / \mu_s^L > 1$ and $\tau_{j,1}^H / \tau_{j,1}^L \geq \mu_1^H / \mu_1^L$ for the strict inequality below, we obtain

$$\bar{\ell}_1 \geq \ell(1, S) \geq \ell_0 \frac{f_{j^*}^H}{f_{j^*}^L} \frac{\tau_{j^*,1}^H}{\tau_{j^*,1}^L} \frac{\mu_s^H}{\mu_s^L} > \ell_0 \frac{f_{j^*}^H}{f_{j^*}^L} \frac{\mu_1^H}{\mu_1^L} \equiv \ell(j^*, 1).$$

Thus $\ell(j^*, 1) < \bar{\ell}_1$, so $\sigma_{j^*,1}^1 = 1$ by Proposition 2. Then defining $x^\theta \equiv \sum_{s \geq 2} \mu_s^\theta \times \sigma_{j^*,1}^s$, we have $\tau_{j^*,1}^\theta = (1 - \eta)(\mu_1^\theta + x^\theta)$. With this, (18) factors as

$$\eta \left(\frac{\eta}{1 - \eta} (\mu_1^L x^H - \mu_1^H x^L) + (x^H \mu_1^{2L} - x^L \mu_1^{2H}) + x^H x^L (\mu_1^L - \mu_1^H) \right) > 0,$$

which follows at once by our signal ordering, $\mu_1^L > \mu_1^H$ and $\mu_1^L / \mu_1^H > \mu_2^L / \mu_2^H \geq x^L / x^H$. Q.E.D.

A.4. Equilibrium: Proof of Proposition 1 and Corollary 1

PROOF OF PROPOSITION 1: Fix an optimal protocol (g^0, σ, d) . We wish to prove that under (4) and (5), conditions (i) and (ii) for a team equilibrium are satisfied. For (i), we follow Piccione and Rubinstein's (1997) Proposition 3 proof. Let $\Pi(g^0, \sigma, d)$ denote the payoff in (3). For any $i, j \in \mathcal{M}$ and $s \in \mathcal{S}$, optimality of $\sigma_{i,j}^s > 0$ requires

$$(19) \quad \frac{\partial}{\partial \sigma_{i,j}^s} \Pi(g^0, \sigma, d) \geq \frac{\partial}{\partial \sigma_{i,j'}^s} \Pi(g^0, \sigma, d) \quad \text{for all } j'.$$

Now, define $X(j', \zeta)$ as the set of all terminal histories ending in state j' , $X(j', s)$ as the set of all (nonterminal) histories ending with observation $s \in \mathcal{S}$ in memory state j' , and define $P^\sigma(z|z', \theta)$ as the probability of history z , conditional on history z' and true state θ . So $f_{j'}^\theta = \sum_{z \in X(j', \zeta)} P^\sigma(z|\theta)$. Observe that for any two memory states i, j , $P^\sigma(z|\theta)$ can be written as $(\sigma_{i,j}^s)^{\delta(z)}$ times a term independent of $\sigma_{i,j}^s$, where $\delta(z)$ denotes the number of occurrences of the transition $i \rightarrow_s j$ along history z . Thus $\frac{\partial}{\partial \sigma_{i,j}^s} P^\sigma(z|\theta) = \delta(z) P^\sigma(z|\theta) / \sigma_{i,j}^s$. Defining $\mathcal{H}(z|i, s)$ as the set of all subhistories of z ending with observation s in state i then yields

$$\begin{aligned} \frac{\partial}{\partial \sigma_{i,j}^s} P^\sigma(z|\theta) &= \sum_{z' \in \mathcal{H}(z|i, s)} \frac{P^\sigma(z|\theta)}{\sigma_{i,j}^s} = \sum_{z' \in \mathcal{H}(z|i, s)} P^\sigma(z'|\theta) P^\sigma(z|z', j, \theta) \\ &= \sum_{z' \in X(i, s)} P^\sigma(z'|\theta) P^\sigma(z|z', j, \theta). \end{aligned}$$

Then using this for the equality below, we get the following expression for $\frac{\partial}{\partial \sigma_{i,j}^s} \sum_{j' \in \mathcal{M}} d_{j'} \pi^\theta f_{j'}^\theta$:

$$\begin{aligned} (20) \quad & \sum_{j' \in \mathcal{M}} \sum_{z \in X(j', \zeta)} d_{j'} \pi^\theta \frac{\partial}{\partial \sigma_{i,j}^s} P^\sigma(z|\theta) \\ &= \sum_{z' \in X(i, s)} P^\sigma(z'|\theta) \left(\sum_{j' \in \mathcal{M}} \sum_{z \in X(j', \zeta)} d_{j'} \pi^\theta P^\sigma(z|z', j, \theta) \right). \end{aligned}$$

By stationarity, the expression in parentheses in (20) is independent of z' and is precisely the continuation payoff v_j^θ . And by (1), $\sum_{z' \in X(i,s)} P^\sigma(z'|\theta) \equiv f_i^\theta \mu_s^\theta$. Thus, $\frac{\partial}{\partial \sigma_{i,j}^s} \sum_j d_j \pi^\theta f_j^\theta = f_i^\theta \mu_s^\theta v_j^\theta$ by (20), and so (taking expectations with respect to true state θ)

$$\begin{aligned} \frac{\partial}{\partial \sigma_{i,j}^s} \Pi(g^0, \sigma, d) &= p_0 f_i^H \mu_s^H v_j^H + (1 - p_0) f_i^L \mu_s^L v_j^L \\ &\propto p(i, s) v_j^H + (1 - p(i, j)) v_j^L. \end{aligned}$$

Substituting into (19) yields condition (i) for a team equilibrium. Condition (ii) follows immediately from the linearity of (3) in d_i , noting from (1) that f^θ is independent of d . Q.E.D.

PROOF OF COROLLARY 1: Consider an optimal protocol satisfying equilibrium conditions (i) and (ii). Choose a memory state i with $f_i^\theta > 0$ and suppose, by contradiction, that the DM prefers some state $j \neq i$ at belief $p(i)$, say (WLOG) $j < i$. So $\ell(i) \Delta_{i,j}^H / \Delta_{j,i}^L < 1$, using Lemma 1(a) implication $\Delta_{i,j}^H > 0 > \Delta_{i,j}^L$. Also, if $g_i^0 > 0$, then state i must yield a weakly higher continuation payoff than j at the prior p_0 , so $\ell_0 \Delta_{i,j}^H / \Delta_{j,i}^L \geq 1$. Combining inequalities, we deduce that $\ell_0 > \ell(i) \Leftrightarrow 1 > (f_i^H / f_i^L)$ whenever $g_i^0 > 0$. Then since $f_i^\theta = \eta g_i^0 + \sum_{k \in \mathcal{M}} f_k^\theta \tau_{k,i}^\theta$ by (2), it follows that

$$(21) \quad \frac{f_i^H}{f_i^L} \equiv \frac{\eta g_i^0 + \sum_{k \in \mathcal{M}} f_k^H \tau_{k,i}^H}{\eta g_i^0 + \sum_{k \in \mathcal{M}} f_k^L \tau_{k,i}^L} \geq \frac{\sum_{k \in \mathcal{M}} f_k^H \tau_{k,i}^H}{\sum_{k \in \mathcal{M}} f_k^L \tau_{k,i}^L} \geq \min_{\{k \in \mathcal{M} : \tau_{k,i}^\theta > 0\}} \frac{f_k^H \tau_{k,i}^H}{f_k^L \tau_{k,i}^L}.$$

But by condition (i) for a team equilibrium, the DM only puts positive probability on a transition $k \rightarrow_s i$ if state i maximizes his payoff at posterior likelihood $\ell(k, s)$. In particular, state i must be weakly preferred to $j < i$, so $\ell(k, s) \Delta_{i,j}^H / \Delta_{j,i}^L \geq 1$. Together with our assumed inequality $\ell(i) \Delta_{i,j}^H / \Delta_{j,i}^L < 1$, we deduce that $\sigma_{k,i}^s > 0 \Rightarrow \ell(i) < \ell(k, s)$, so $(f_i^H / f_i^L) < (f_k^H / f_k^L) \xi(s)$. Summing over signal realizations, recalling that $\tau_{k,i}^\theta \equiv (1 - \eta) \sum_{s \in \mathcal{S}} \sigma_{k,i}^s \mu_s^\theta$, we contradict (21):

$$\tau_{k,i}^\theta > 0 \quad \Rightarrow \quad \frac{f_k^H \tau_{k,i}^H}{f_k^L \tau_{k,i}^L} \equiv \frac{f_k^H \sum_{s \in \mathcal{S}} \sigma_{k,i}^s \mu_s^H}{f_k^L \sum_{s \in \mathcal{S}} \sigma_{k,i}^s \mu_s^L} > \frac{f_i^H}{f_i^L}. \quad \text{Q.E.D.}$$

A.5. Optimal Rules: Proof of Propositions 3 and 4(a), (b)

STEP 1: $\Pi^*(\eta, M)$ is strictly decreasing in η whenever information matters.

PROOF: Fix (g^0, d) . Recall that the DM's objective is to then choose σ , and hence f^θ , to maximize (3). And that this is equivalent to choosing the best steady-state distribution, subject to the constraint that for any pair of states i, j with $g_j^0 > 0$, the DM puts probability at least ηg_j^0 on the transition $i \rightarrow_s j$ after all signal realizations s . But by the Proposition 2 lesson, transitions $i \rightarrow_s j$ are strictly suboptimal if $\xi(s) < 1$ and $j > i$ or $\xi(s) > 1$ and $j < i$. Therefore, the constraint $\sigma_{i,j}^s \geq \eta g_j^0 \forall s$ is binding for any $\eta > 0$, with higher η forcing a higher probability on suboptimal transitions, and thus $\Pi^*(\eta, M)$ is strictly decreasing in η . Q.E.D.

STEP 2: Fix (g^0, d) and (M, η) , and let $i^* = \min\{i | d_i = 1\}$. For any transition rule σ and the induced distributions f^θ , the DM's expected payoff is $(1 - p_0)|\pi^L| \cdot \Pi(x, r)$, where

$$(22) \quad \Pi(x, r) = \frac{\beta}{1+x} - \frac{1}{1+\frac{x}{r^2}},$$

$$x \equiv \frac{\sum_{i=1}^{i^*-1} f_i^H}{\sum_{i=i^*}^M f_i^H}, \quad \text{and} \quad r = \sqrt{\frac{\sum_{i=i^*}^M f_i^L \sum_{i=1}^{i^*-1} f_i^H}{\sum_{i=i^*}^M f_i^H \sum_{i=1}^{i^*-1} f_i^L}}.$$

PROOF: By team equilibrium condition (ii) and the memory state ordering, the DM chooses action 1 in all memory states $j \geq i^*$ and action 0 below i^* . Using (3), his expected payoff is then

$$\begin{aligned} & p_0 \pi^H \sum_{i=i^*}^M f_i^H + (1 - p_0) \pi^L \sum_{i=i^*}^M f_i^L \\ &= (1 - p_0) |\pi^L| \left(\beta \sum_{i=i^*}^M f_i^H - \sum_{i=i^*}^M f_i^L \right) \\ &= (1 - p_0) |\pi^L| \left(\frac{\beta}{1 + \sum_{i=1}^{i^*-1} f_i^H / \sum_{i=i^*}^M f_i^H} - \frac{1}{1 + \sum_{i=1}^{i^*-1} f_i^L / \sum_{i=i^*}^M f_i^L} \right) \end{aligned}$$

(final equality by dividing each sum $\sum_{i=i^*}^M f_i^H$ by $\sum_{i=1}^M f_i^\theta = 1$). This is found in (22). Q.E.D.

STEP 3: Fix a value $r < 1$. If $r\sqrt{\beta} > 1$, then $\Pi(x, r)$ reaches a maximum value at $x = 0$ of $\beta - 1$. If $r\sqrt{\beta} < 1$, then $\Pi(x, r)$ is maximized at $x^*(r) =$

$r(1 - \sqrt{\beta}r)/(\sqrt{\beta} - r)$, with corresponding value $\bar{\Pi}(r) \equiv \arg \max_{x \geq 0} \Pi(x, r) = (\sqrt{\beta} - r)^2/(1 - r^2)$.

PROOF: Differentiating $\Pi(x, r)$ with respect to x yields

$$(23) \quad \frac{\partial \Pi(x, r)}{\partial r} = -\left(\frac{\sqrt{\beta}}{1+x}\right)^2 + \left(\frac{r}{r^2+x}\right)^2.$$

This is weakly positive iff $x(\sqrt{\beta} - r) \leq r(1 - r\sqrt{\beta})$. If $r\sqrt{\beta} > 1$, then this inequality can never hold, as the RHS is strictly negative, while the LHS is weakly positive by $x \geq 0$ and $r < 1$ (recalling our assumption $\beta \geq 1$). Thus $\frac{\partial}{\partial r} \Pi(x, r) < 0$ over this range, and so $\bar{\Pi}(x, r)$ is maximized at $x = 0$. If $r\sqrt{\beta} < 1$, then $\Pi(x, r)$ increases in x when $x < x^*(r)$ and decreases when $x > x^*(r)$. It is maximized at $x = x^*(r)$. Evaluating $\Pi(x, r)$ at $x = x^*(r)$ yields the desired $\bar{\Pi}(r)$ expression. *Q.E.D.*

STEP 4: Define $r^* = (\xi(1)/\xi(S))^{(M-1)/2}$. Then r^* is a lower bound on r (from (22)), and the value $\bar{\Pi}(r)$ from Step 3 is decreasing in r , with upper bound $\bar{\Pi}(r^*)$.

PROOF: Assume $f_i^\theta > 0$ for all memory states $i \in \mathcal{M}$ (otherwise, define \mathcal{M} as the set of states with $f_i^\theta > 0$). By Appendix A.2, f^θ is the steady-state distribution of a Markov process with perturbed transition chances $\omega_{i,j}^\theta$, implying the following steady-state relationship: for any block of memory states A , the probability of exiting A must equal the probability coming into A : $\sum_{j \in A} f_j^\theta (\sum_{l \notin A} \omega_{j,l}^\theta) = \sum_{l \notin A} f_l^\theta (\sum_{j \in A} \omega_{l,j}^\theta)$. Setting $A = \{1, 2, \dots, i\}$ and taking ratios, this implies that for any memory state i and transition rule σ ,

$$(24) \quad \frac{\sum_{j=1}^i f_j^H}{\sum_{j=1}^i f_j^L} = \left(\frac{\sum_{l=i+1}^M \omega_{j,l}^L}{\sum_{l=i+1}^M \omega_{j,l}^H} \right) \left(\frac{\sum_{j=1}^i \omega_{l,j}^H}{\sum_{j=1}^i \omega_{l,j}^L} \right) \left(\frac{\sum_{l=i+1}^M f_l^H}{\sum_{l=i+1}^M f_l^L} \right).$$

By our ordering of the memory states, the LHS of (24) is at most f_i^H/f_i^L and the final RHS term is at least f_{i+1}^H/f_{i+1}^L . And for a bound on the first two RHS terms, recall that the perturbed transition probability from any state i to j is given by $\omega_{i,j}^\theta = \eta g_j^0 + (1 - \eta) \sum_s \mu_s^\theta \sigma_{i,j}^s$. Taking ratios and using $\xi(1) \leq \mu_s^H/\mu_s^L \leq \xi(S)$ and $\xi(1) < 1 < \xi(S)$, we deduce $\xi(1) \leq \omega_{i,j}^H/\omega_{i,j}^L \leq \xi(S)$, so that the RHS of (24) is at least $(\xi(1)/\xi(S))(f_{i+1}^H/f_{i+1}^L)$. It then follows from (24) that for any transition rule σ , the induced distributions f^θ satisfy (25) below, and iterating then yields (26):

$$(25) \quad f_i^H/f_i^L \geq (f_{i+1}^H/f_{i+1}^L)(\xi(1)/\xi(S)) \quad \text{for all } i \in \mathcal{M}$$

$$(26) \quad \Rightarrow \quad (f_1^H/f_1^L) \geq (f_2^H/f_2^L)/\lambda \geq \dots \geq (f_M^H/f_M^L)/\lambda^{M-1}.$$

Now, by (22) and our ordering of the memory states (in particular, $\ell(1) \leq \ell(i) \leq \ell(M)$), we have

$$(27) \quad r = \sqrt{\frac{\sum_{i=i^*}^M f_i^L \sum_{i=1}^{i^*-1} f_i^H}{\sum_{i=i^*}^M f_i^H \sum_{i=1}^{i^*-1} f_i^L}} \geq \sqrt{\frac{f_M^L f_1^H}{f_M^H f_1^L}}.$$

Substituting (26) into this inequality yields $r \geq r^*$. Finally, an immediate calculation shows that $r\sqrt{\beta} < 1 \Rightarrow \frac{d\bar{\Pi}(r)}{dr} < 0$ and thus $r \geq r^* \Rightarrow \bar{\Pi}(r) \leq \bar{\Pi}(r^*)$. *Q.E.D.*

STEP 5: Let (g^0, σ, d) be an equilibrium protocol that violates (at least) one of parts (a) or (b) of Proposition 4. Then there exists $\delta > 0$ such that the DM earns at most $(1 - p_0)|\pi^L|\bar{\Pi}(r^* + \delta)$.

PROOF: Since we showed in Step 3 that an optimal choice of x yields payoff $\bar{\Pi}(r)$, decreasing in r , it suffices to prove that a violation of (a) or (b) yields $r > r^* + \delta$. Suppose first that part (a) is violated, so there are states i, j and a nonextreme signal realization s with $\sigma_{i,j}^s > 0$. Assume WLOG that $j > i$ (a symmetric argument applies if $j < i$). Then by Proposition 2 with $j \geq i + 1$, the equilibrium condition for transition $i \rightarrow_s j$ implies $\ell(i, s) \geq \bar{\ell}_{j-1} \geq \bar{\ell}_i$. Also, $\ell(i + 1, 1) \leq \bar{\ell}_i$; otherwise, $\ell(j, s) > \bar{\ell}_i$ for all states $j \geq i + 1$ and all signal realizations s , implying via Proposition 2 that no state $j \geq i + 1$ moves below $i + 1$; but then the block of states $\{i + 1, \dots, M\}$ is absorbing, contradicting optimality via Lemma 2. Combining inequalities, we deduce

$$\ell_0(f_i^H/f_i^L)\xi(s) \equiv \ell(i, s) \geq \bar{\ell}_i \geq \ell(i + 1, 1) \equiv \ell_0(f_{i+1}^H/f_{i+1}^L)\xi(1).$$

So $\ell(i) \geq \ell(i + 1)(\xi(1)/\xi(s))$. We have $\ell(1)/\ell(M) \geq (1/\lambda)^{M-2}(\xi(1)/\xi(s))$ for $s < S$, given (25) and (26). Substituting into (27), we conclude $r \geq r^* + \delta_1$, where

$$(28) \quad \delta_1 \equiv (1/\lambda)^{(M-2)/2}(\sqrt{\xi(1)/\xi(s)} - \sqrt{\xi(1)/\xi(S)}).$$

Suppose next that part (b) is violated, so there are states i, j with $|j - i| \geq 2$ such that the DM jumps $i \rightarrow_s j$ with positive probability. WLOG, assume $j \geq i + 2$. By Proposition 2, the DM must then find it incentive compatible to jump from i up to j after the highest possible signal realization, S , requiring $\ell(i, S) \geq \bar{\ell}_{j-1} \geq \bar{\ell}_{i+1}$. But by Proposition 2, we also have $\ell(i + 1) \leq \bar{\ell}_{i+1}$. Combining inequalities, we deduce $\ell(i, S) \geq \bar{\ell}_{i+1} \geq \ell(i + 1)$, and so (using (4))

$$\ell_0(f_i^H/f_i^L)\xi(S) \geq \ell_0(f_{i+1}^H/f_{i+1}^L) \Rightarrow f_i^H/f_i^L \geq (f_{i+1}^H/f_{i+1}^L)/\xi(S).$$

Together with (25) and (26), this implies $\ell(1)/\ell(M) \geq (1/\lambda)^{M-2}(1/\xi(S))$. Substituting into (27), we conclude $r \geq r^* + \delta_2$, where

(29) $\delta_2 \equiv (\xi(1)/\xi(S))^{(M-2)/2}(\sqrt{1/\xi(S)} - \sqrt{\xi(1)/\xi(S)}).$

Finally, since $\xi(1) < \xi(s) < \xi(S)$ for any $s < S$, the expressions δ_1 and δ_2 defined in (28) and (29) are both strictly positive. Thus, defining $\delta = \min\{\delta_1, \delta_2\}$, we have shown that a violation of part (a) or (b) of Proposition 4 yields $r \geq r^* + \delta$ for some $\delta > 0$, as desired. Q.E.D.

STEP 6: *There exists a sequence of memory protocols (g^0, σ, d) , with associated payoffs Π^η , such that $\Pi^\eta \rightarrow \bar{\Pi}(r^*)$ as $\eta \rightarrow 0$.*

PROOF: Choose any initial state, and any action rule with $d_1 = 0$ and $d_M = 1$ (fixed for all protocols along the sequence). Now define a sequence of transition rules σ^η as follows: (i) $\sigma_{i,i}^{s,\eta} = 1$ if $s \notin \{1, S\}$, (ii) $\sigma_{1,2}^{S,\eta} = \sqrt{\eta}$, (iii) $\sigma_{M,M-1}^{1,\eta} = \sqrt{\eta}(\frac{\mu_S^H \mu_S^L}{\mu_1^H \mu_1^L})^{(M-1)/2}(1 - r^* \sqrt{\beta})/(\sqrt{\beta} - r^*)$, and (iv) for all interior states i , $\sigma_{i,i+1}^S = 1$ and $\sigma_{i,i-1}^1 = 1$. Putting these into the steady-state system of equations in (2), this rule yields $f_i^\theta \rightarrow 0$ for all interior memory states i , while

$$\begin{aligned} \frac{f_1^\theta}{f_M^\theta} &\rightarrow \frac{\sigma_{M,M-1}^1}{\sigma_{1,2}^S} \left(\frac{\mu_1^\theta}{\mu_S^\theta} \right)^{M-1} \\ &\Rightarrow \frac{f_1^H}{f_M^H} \rightarrow r^* \frac{(1 - r^* \sqrt{\beta})}{(\sqrt{\beta} - r^*)}, \quad \frac{f_1^L}{f_M^L} \rightarrow \frac{1}{r^*} \frac{(1 - r^* \sqrt{\beta})}{(\sqrt{\beta} - r^*)}. \end{aligned}$$

The limiting payoff is then $p_0 \pi^H f_M^H + (1 - p_0) \pi^L f_M^L$. Dividing each term f_M^θ by $1 = \sum f_j^\theta \rightarrow f_1^\theta + f_M^\theta$ and substituting the above expressions for f_1^θ/f_M^θ yields limit payoff $\bar{\Pi}(r^*)$. Q.E.D.

STEP 7—Completing the Proof: Step 2 rewrote the payoff as $(1 - p_0)|\pi^L| \times \Pi(x, r)$, Step 3 derived the value $\bar{\Pi}(r) \equiv \max_x \Pi(x, r)$, and Step 4 proved that $\bar{\Pi}(r^*)$ is an upper bound on $\bar{\Pi}(r)$. Together with the Step 1 result that the payoff is strictly decreasing in η and noting that $(1 - p_0)|\pi^L| \cdot \bar{\Pi}(r^*)$ is precisely the payoff bound in Proposition 3, with $r^* = \sqrt{\rho^*}$, this completes the proof of Proposition 3. Next, Step 5 proved that a protocol that reacts to noisy signal realizations $s \notin \{1, S\}$ or that jumps to nonadjacent states implies that the value $\bar{\Pi}(r)$ is bounded below $\bar{\Pi}(r^*)$. Thus, there exists $\varepsilon > 0$ such that a violation of Proposition 4(a) or (b) earns payoff at most $(1 - p_0)|\pi^L| \cdot \bar{\Pi}(r^*) - \varepsilon$. But by Step 6, there exists a sequence of protocols with limit payoff $(1 - p_0)|\pi^L| \cdot \bar{\Pi}(r^*)$; then by continuity, for any $\varepsilon > 0$, there exists η_ε such that an optimal protocol earns a payoff above $(1 - p_0)|\pi^L| \cdot \bar{\Pi}(r^*) - \varepsilon$ whenever $\eta < \eta_\varepsilon$. Thus,

for $\eta < \eta_\varepsilon$, a rule that jumps or reacts to signal observations $s \notin \{1, S\}$ is not optimal, proving Proposition 4(a) and (b). Q.E.D.

A.6. Extremal States: Transition Chances Vanish as $\eta \rightarrow 0$

Fix a memory size M , assume information matters, and consider a sequence of protocols (g^0, σ, d) with $\eta \rightarrow 0$ (suppressing dependence on η). Let Π^n , r^n , and x^n denote (respectively) the sequence of associated payoffs and values of r , x (from (22)). Steps 1 and 2 are technical lemmas; Step 3 proves that transition chances out of the extreme states vanish as $\eta \rightarrow 0$, and Step 3(iv) in Section A.9 completes the proof of Proposition 4(d), proving that $\sigma_{1,2}^S$ and $\sigma_{M,M-1}^1$ are in fact asymptotic to $\sqrt{\eta}$.

STEP 1: If $\Pi^n \rightarrow (1 - p_0)|\pi^L|\bar{\Pi}(r^*)$, then $x^n \rightarrow x^*(r^*)$, $r^n \rightarrow r^*$. Also,

$$(30) \quad \ell(i)/\ell(i+1) \rightarrow \xi(1)/\xi(S) \quad \text{and} \quad r^n \rightarrow \sqrt{\ell(1)/\ell(M)} \rightarrow (1/\lambda)^{(M-1)/2}.$$

PROOF: The first assertion is immediate from Steps 3 and 4 in Section A.5. And by (27), $r^n \rightarrow r^*$ if and only if (25), (26), and (27) all hold with equality in the limit, yielding conditions in (30). Q.E.D.

STEP 2: If $\Pi^n \rightarrow (1 - p_0)|\pi^L| \cdot \bar{\Pi}(r^*)$, then $f_j^\theta/f_1^\theta \rightarrow 0$ for all interior states j .

PROOF: Assume that $x^n \rightarrow x^*(r^*)$ and that the first limit in (30) holds. Rewrite (27) as

$$(31) \quad r = \frac{\sqrt{1 + \sum_{j=i^*}^{M-1} f_j^L / f_M^L}}{\sqrt{1 + \sum_{j=i^*}^{M-1} f_j^H / f_M^H}} \frac{\sqrt{1 + \sum_{j=2}^{i^*-1} f_j^H / f_1^H}}{\sqrt{1 + \sum_{j=2}^{i^*-1} f_j^L / f_1^L}} \sqrt{\frac{\ell(1)}{\ell(M)}}.$$

By the ordering of the memory states, the RHS of (31) is at least $\sqrt{\ell(1)/\ell(M)}$; thus, for it to tend to exactly $\sqrt{\ell(1)/\ell(M)}$, as required by the second limit in (30), the first and second terms must both tend to 1 as $\eta \rightarrow 0$. But by $\xi(1) < 1 < \xi(S)$ and the first limit in (30), there exist $\eta^* > 0$ and $\hat{\lambda} > 1$ such that for all $\eta < \eta^*$, $\ell(M)/\ell(j) > \hat{\lambda}$ and $\ell(j)/\ell(1) \geq \hat{\lambda}$. Rearranging yields

$$(32) \quad (f_j^L/f_M^L) \geq \hat{\lambda}(f_j^H/f_M^H) \quad \text{and} \quad (f_j^H/f_1^H) \geq \hat{\lambda}(f_j^L/f_1^L).$$

But then the first RHS term in (31) can only tend to 1 if $f_j^\theta/f_M^\theta \rightarrow 0 \forall i^* \leq j \leq M-1$ (by the first inequality in (32)); similarly the second RHS term can

only tend to 1 if $f_j^\theta/f_1^\theta \rightarrow 0 \forall 2 \leq j \leq i^* - 1$ (by the second inequality in (32)). But these two limits imply via (22) that $x^n \rightarrow f_1^\theta/f_M^\theta$, and so $x^n \rightarrow x^*(r^*) > 0$ implies $f_1^\theta/f_M^\theta \rightarrow 0$. Thus, for all $i^* \leq j \leq M - 1$, we have $f_j^\theta/f_M^\theta \rightarrow 0 \Rightarrow f_j^\theta/f_1^\theta = (f_j^\theta/f_M^\theta)(f_M^\theta/f_1^\theta) \rightarrow 0$, and since we have also just shown that $f_j^\theta/f_1^\theta \rightarrow 0 \forall 2 \leq j \leq i^* - 1$, this completes the proof. Q.E.D.

STEP 3: If $\Pi^n \rightarrow (1 - p_0)|\pi^L| \cdot \bar{\Pi}(r^*)$, then transition chances $\sigma_{1,2}^S, \sigma_{M,M-1}^1$ vanish as $\eta \rightarrow 0$.

PROOF: Suppose, by contradiction, that $\Pi^n \rightarrow (1 - p_0)|\pi^L| \cdot \bar{\Pi}(r^*)$ but that transition chances out of an extremal state do not vanish, say $\sigma_{1,2}^S \not\rightarrow 0$. Then there exists $\Delta > 0$ and a convergent subsequence of protocols with limit payoff $(1 - p_0)|\pi^L| \cdot \bar{\Pi}(r^*)$ and with $\sigma_{1,2}^S \geq \Delta$ for all subsequences along the protocol. For a contradiction, it suffices by Step 2 to find an interior state j with f_j^θ/f_1^θ boundedly positive as $\eta \rightarrow 0$. Now, consider state $j = 2$, which (by (2)) obeys the steady-state equation

$$(33) \quad f_2^\theta = \sum_j f_j^\theta \omega_{j,2}^\theta \geq f_1^\theta \omega_{1,2}^\theta \geq f_1^\theta (1 - \eta) \mu_S^\theta \sigma_{1,2}^S.$$

So $f_2^\theta/f_1^\theta \geq (1 - \eta) \mu_S^\theta \sigma_{1,2}^S$, which stays boundedly positive, as desired, by $\sigma_{1,2}^S \geq \Delta > 0$. Q.E.D.

A.7. Beliefs: Proof of Corollary 3

By Steps 1 and 2 in Section A.6 (respectively), an optimal sequence of protocols satisfies $x^n \rightarrow x^*(r^*)$ and $x^n \rightarrow f_1^H/f_M^H$, with $x^*(r^*)$ as defined in Step 3 of Section A.5. Thus,

$$(34) \quad x^n \rightarrow f_1^H/f_M^H \rightarrow r^*(1 - \sqrt{\beta}r^*)/(\sqrt{\beta} - r^*).$$

But since f^H is a probability distribution and $f_j^\theta/f_1^\theta \rightarrow 0$ for all interior j (again by Step 2 of Section A.6), we have $f_1^\theta + f_M^\theta \rightarrow \sum_j f_j^\theta = 1$. Substituting $f_M^H = 1 - f_1^H$ into (34), we deduce

$$(35) \quad f_1^H \rightarrow r^*(1 - \sqrt{\beta}r^*)/\sqrt{\beta}(1 - (r^*)^2).$$

By Step 1 in Section A.6, an optimal sequence of protocols also satisfies (30); thus (from the second limit therein) $\sqrt{(f_1^H/f_M^H)(f_M^L/f_1^L)} \rightarrow r^*$. Substituting (34) into this expression and using $f_1^L + f_M^L = 1$, we obtain $f_1^L \rightarrow (1 - \sqrt{\beta}r^*)/(1 - (r^*)^2)$. Together with (35), we then get $f_1^H/f_1^L \rightarrow r^*/\sqrt{\beta}$, so (by (4)), $\ell(1) = \ell_0(f_1^H/f_1^L) = \sqrt{\ell_0|\pi^L|/\pi^H}(1/\lambda)^{(M-1)/2}$. Then by the first limit in (30), $\ell(i) \rightarrow \lambda^{i-1}\ell(1) = \sqrt{\ell_0|\pi^L|/\pi^H}\lambda^{i-(M+1)/2}$, the desired expression. For the indifference

likelihood ratios, recall from Lemma 2 that there are no absorbing (blocks of) states in an optimal protocol. Then by Proposition 2, a self in state i must find it optimal to move up after the highest observation $s = S$, so $\ell(i, S) \geq \bar{\ell}_i$ (otherwise the block $\{1, 2, \dots, i\}$ is absorbing), and similarly we need $\ell(i + 1, 1) \leq \bar{\ell}_i$; thus $\ell(i, S) \geq \bar{\ell}_i \geq \ell(i + 1, 1)$. But since $\ell(i, S) \rightarrow \ell(i + 1, 1)$ by the first limit in (30), this yields the desired $\bar{\ell}_i = \ell(i, S)$.

A.8. Initial State and Action Rule

The expressions for i_0, i^* in (10) follow at once by substituting the Corollary 3 expressions for $\bar{\ell}_{i_0-1}, \bar{\ell}_{i_0}, \ell(i)$ into the inequality $\bar{\ell}_{i_0-1} \leq \ell_0 \leq \bar{\ell}_{i_0}$, which says that the DM must earn a higher payoff in state i_0 than either $i_0 - 1$ or $i_0 + 1$ at his prior likelihood ℓ_0 , and the inequality $\ell(i^*)\pi^H/|\pi^L| \geq 1 \geq \ell(i^* - 1)\pi^H/|\pi^L|$, which says that i^* is the smallest state to prefer action 1. Note that $i_0 \geq i^*$ whenever M is odd or M is even and $\beta \geq \xi(1)\xi(S)$. Indeed, recall that $\lambda \equiv \xi(S)/\xi(1) > \xi(S)$, so $\log \xi(S)/\log \lambda < 1$. If M is odd, it then follows that i_0 is increasing in $\log \sqrt{\beta}/\log \lambda$ with $i_0 \geq \frac{M+1}{2}$, which is the maximum value for i^* by $\beta > 1$. If M is even, then $\beta > \xi(1)\xi(S)$ implies $\log \sqrt{\beta} - \log \xi(S) > -\frac{1}{2} \log \lambda$, implying $i_0 \geq \frac{M}{2} + 1$, which strictly exceeds i^* by $\beta > 1$.

A.9. Interior States: Proof of Proposition 4(c)

Steps 1 and 2 in this proof are technical; Step 1 derives recursive equations on the steady-state distribution and the payoff differentials, and Step 2 uses these to derive necessary conditions of an optimal protocol (for η near zero). Step 3 proves the first assertion in Proposition 4(c), namely, that no state can be sticky both up and down, and that interior transition chances stay bounded away from zero as $\eta \rightarrow 0$. Step 3 also completes the proof of Proposition 4(d), proving that $\sigma_{1,2}^S$ and $\sigma_{M,M-1}^1$ are asymptotic to $\sqrt{\eta}$ (Section A.6 proved only that they are vanishing).

STEP 1: *Fix a convergent sequence of optimal protocols as $\eta \rightarrow 0$. There exists $\eta^* > 0$ with fluid transitions ($\sigma_{2,1}^1 = \sigma_{N-1,N}^S = 1$) into extreme states for all $\eta < \eta^*$. Define $\alpha_l \equiv \prod_{j=2}^{l-1} (\sigma_{j+1,j}^1/\sigma_{j,j+1}^S)$ and $\beta_l = \prod_{j=l+1}^{M-1} (\sigma_{j-1,j}^S/\sigma_{j,j-1}^1)$.²⁹ Terminal distributions f^θ and payoff differentials $\Delta_{i,i-1}^\theta$ obey*

$$(36) \quad i \leq i_0: \quad \alpha_i f_i^\theta / f_1^\theta = \eta \sum_{l=2}^i (\alpha_l / \sigma_{l,l-1}^1 \mu_1^\theta) (\mu_S^\theta / \mu_1^\theta)^{i-l} \\ + \sigma_{1,2}^S (\mu_S^\theta / \mu_1^\theta)^{i-1} + o(\eta),$$

²⁹As usual, I use the convention that $\alpha_1 = \sigma_{1,2}^S$, $\alpha_2 = \beta_{M-1} = 1$, and $\beta_M = \sigma_{M,M-1}^1$.

$$(37) \quad i \geq i_0: \quad \frac{f_i^\theta}{f_M^\theta} \beta_i = \eta \sum_{l=i}^{M-1} (\beta_l / \sigma_{l,l+1}^S \mu_S^\theta) (\mu_1^\theta / \mu_S^\theta)^{l-i} \\ + \sigma_{M,M-1}^1 (\mu_1^\theta / \mu_S^\theta)^{M-i} + o(\eta),$$

$$(38) \quad i \leq i^*: \quad \frac{\sigma_{i-1,i}^1}{\alpha_i} \Delta_{i,i-1}^\theta / \Delta_{2,1}^\theta = \sigma_{1,2}^S \sum_{l=3}^i (1/\alpha_{l-1}) (\mu_1^\theta / \mu_S^\theta)^{i-l} \\ + (\mu_1^\theta / \mu_S^\theta)^{i-2} + O(\eta),$$

$$(39) \quad i \geq i^*: \quad \frac{\sigma_{i-1,i}^S}{\beta_{i-1}} \frac{\Delta_{i,i-1}^\theta}{\Delta_{M,M-1}^\theta} = \sigma_{M,M-1}^1 \sum_{l=i-1}^{M-2} \left(\frac{1}{\beta_{l+1}} \right) \left(\frac{\mu_S^\theta}{\mu_1^\theta} \right)^{l+1-i} \\ + \left(\frac{\mu_S^\theta}{\mu_1^\theta} \right)^{M-i} + O(\eta).$$

PROOF: Choose η^* so that Proposition 4(a), (b), and (d) hold for $\eta < \eta^*$. We first prove (36) by induction and prove $\sigma_{2,1}^1 = 1$. For $i = 2 \leq i_0$, evaluate (2) at $i = 1$ and $g_i^0 = 0$; this yields

$$(40) \quad (\eta + (1 - \eta) \sigma_{1,2}^S \mu_S^\theta) f_1^\theta = f_2^\theta (1 - \eta) \mu_1^\theta \sigma_{2,1}^1 \\ \Rightarrow \quad \frac{f_2^\theta}{f_1^\theta} = \frac{1}{\sigma_{2,1}^1} \left(\eta \frac{1}{\mu_1^\theta} + \sigma_{1,2}^S \frac{\mu_S^\theta}{\mu_1^\theta} + o(\eta) \right).$$

This is precisely (36) evaluated at $i = 2$, as long as $\sigma_{2,1}^1 = 1$. To prove that $\sigma_{2,1}^1 = 1$, take ratios in the first expression in (40) to deduce $\ell(1, S) > \ell(2, 1)$; but then since $\ell(1, S) = \bar{\ell}_1$ by Proposition 2 and Proposition 4(d)—which requires that the DM be indifferent between states 1 and 2 after observing an S -signal in state 1—we have $\ell(2, 1) < \bar{\ell}_1$; by Proposition 2, optimality then requires $\sigma_{2,1}^1 = 1$, as desired. Now (inductive hypothesis) assume (36) holds at $i - 1$ and i for $i \leq i_0 - 1$. Now, for $i + 1$, by (2) in state i , using Proposition 4(a) and (b) together with $g_i^0 = 0$ yields

$$(\eta + (1 - \eta) (\mu_1^\theta \sigma_{i,i-1}^1 + \mu_S^\theta \sigma_{i,i+1}^S)) f_i^\theta \\ = f_{i-1}^\theta (1 - \eta) \mu_S^\theta \sigma_{i-1,i}^S + f_{i+1}^\theta (1 - \eta) \mu_1^\theta \sigma_{i+1,i}^1.$$

Taking limits and multiplying by $\alpha_{i+1}/(\sigma_{i+1,i}^1 \mu_1^\theta f_1^\theta)$, noting that $\alpha_{i+1}/\alpha_i = \sigma_{i+1,i}^1/\sigma_{i,i+1}^S$, this yields

$$\alpha_{i+1} \frac{f_{i+1}^\theta}{f_1^\theta} = \left(\frac{\sigma_{i,i-1}^1}{\sigma_{i,i+1}^S} + \frac{\mu_S^\theta}{\mu_1^\theta} \right) \left(\alpha_i \frac{f_i^\theta}{f_1^\theta} \right) - \frac{\sigma_{i,i-1}^1}{\sigma_{i,i+1}^S} \left(\alpha_{i-1} \frac{\mu_S^\theta}{\mu_1^\theta} \frac{f_{i-1}^\theta}{f_1^\theta} \right) + o(\eta).$$

Substituting expressions for f_i^θ/f_1^θ and $f_{i-1}^\theta/f_1^\theta$ from (36), noting that $(\alpha_{i-1}f_{i-1}^\theta/f_1^\theta)(\mu_S^\theta/\mu_1^\theta) = \alpha_i(f_i^\theta/f_1^\theta - \eta/\sigma_{i,i-1}^1\mu_1^\theta) + o(\eta)$, this simplifies precisely to the desired expression.

To prove (38), for $2 \leq i \leq i^* - 1$, (5) yields the next expression for v_i^θ by Proposition 4(a) and (b) (no jumps or transitions after $s \neq 1, S$) and the fact that $i \leq i^* - 1 \Rightarrow d_i = 0$:

$$\begin{aligned} (41) \quad v_i^\theta &= (1 - \eta)\mu_1^\theta\sigma_{i,i-1}^1v_{i-1}^\theta + (1 - \eta)\mu_S^\theta\sigma_{i,i+1}^Sv_{i+1}^\theta \\ &\quad + (1 - \eta)(1 - \sigma_{i,i-1}^1\mu_1^\theta - \sigma_{i,i+1}^S\mu_S^\theta)v_i^\theta \\ &\Rightarrow (\eta + (1 - \eta)\mu_1^\theta\sigma_{i,i-1}^1)\Delta_{i,i-1}^\theta + \eta v_{i-1}^\theta = (1 - \eta)\mu_S^\theta\sigma_{i,i+1}^S\Delta_{i+1,i}^\theta. \end{aligned}$$

But from (5) at $i = 1$, using $\tau_{1,1}^\theta = (1 - \eta)(1 - \sigma_{1,2}^S\mu_S^\theta)$, we obtain $\eta v_1^\theta = (1 - \eta)\sigma_{1,2}^S\mu_S^\theta\Delta_{2,1}^\theta$. Along with the identity $\Delta_{i-1,i}^\theta = \Delta_{i-1,1}^\theta + v_1^\theta$, I may rewrite (41) as (for $i \leq i^* - 1$)

$$\begin{aligned} (42) \quad &(\eta + (1 - \eta)\mu_1^\theta\sigma_{i,i-1}^1)\Delta_{i,i-1}^\theta + \eta\Delta_{i-1,1}^\theta + (1 - \eta)\sigma_{1,2}^S\mu_S^\theta\Delta_{2,1}^\theta \\ &= (1 - \eta)\mu_S^\theta\sigma_{i,i+1}^S\Delta_{i+1,i}^\theta. \end{aligned}$$

Then taking limits in (42) and rearranging gives

$$(43) \quad \mu_S^\theta\sigma_{i,i+1}^S\Delta_{i+1,i}^\theta = \mu_1^\theta\sigma_{i,i-1}^1\Delta_{i,i-1}^\theta + \sigma_{1,2}^S\mu_S^\theta\Delta_{2,1}^\theta + O(\eta).$$

Solving this recursion yields the expression in (38) as is easily verified by induction. Then (37) and (39) are found symmetrically, indexing states by their distance from M , not 1. Q.E.D.

STEP 2: Assume $i^* \geq i_0$, choose η small enough that the Step 1 expressions hold, and, letting $\lambda_1 \equiv \mu_S^H/\mu_1^H$ and $\lambda_2 \equiv \mu_1^L/\mu_S^L$, define the expressions

$$\begin{aligned} (44) \quad x_l &\equiv \left(\frac{\xi(1)\lambda^{l-1} - 1}{\mu_1^H\lambda_1^{l-1}} \right), \quad y_l \equiv \left(\frac{\lambda^{l-1} - 1}{\lambda_2^{l-1}} \right), \\ v_l &\equiv \left(\frac{\lambda^{M-l-1} \frac{1}{\xi(1)} - 1}{\mu_S^L\lambda_2^{M-l}} \right), \quad u_l \equiv \frac{\lambda^{M-l} - 1}{\lambda_1^{M-l}}. \end{aligned}$$

In an optimal protocol, the left-hand sides of (45), (46), and (47) hold, with equality if $\sigma_{i,i+1}^S \in (0, 1)$, and the right-hand sides hold, with equality if $\sigma_{i,i-1}^1 \in (0, 1)$:

$$\begin{aligned} (45) \quad i \leq i^* - 1: \quad &\sum_{l=2}^i \frac{\alpha_l x_l}{\sigma_{l,l-1}^1} \bigg/ \sum_{l=2}^i \frac{y_l}{\alpha_l} \leq (\sigma_{1,2}^S)^2 / \eta + o(\eta) / \eta \\ &\leq \sum_{l=2}^i \frac{\alpha_l x_l}{\sigma_{l,l-1}^1} \bigg/ \sum_{l=2}^{i-1} \frac{y_l}{\alpha_l}, \end{aligned}$$

(46)
$$i^* \leq i \leq i_0: \frac{\sum_{l=2}^i \frac{\alpha_l}{\sigma_{l,l-1}^1} x_l}{\sum_{l=2}^{i^*-1} \frac{y_l}{\alpha_l} + \frac{\alpha_M}{\sigma_{1,2}^S} \sum_{l=i^*}^i \frac{u_l}{\alpha_l}} \leq \frac{(\sigma_{1,2}^S)^2 + o(\eta)}{\eta}$$
$$\leq \frac{\sum_{l=2}^i \frac{\alpha_l}{\sigma_{l,l-1}^1} x_l}{\sum_{l=2}^{i^*-1} \frac{y_l}{\alpha_l} + \frac{\alpha_M}{\sigma_{1,2}^S} \sum_{l=i^*}^{i-1} \frac{u_l}{\alpha_l}},$$

(47)
$$i \geq i_0: \sum_{l=i}^{M-1} \frac{\beta_l v_l}{\sigma_{l,l+1}^S} \bigg/ \sum_{l=i+1}^{M-1} \frac{u_l}{\beta_l} \geq \frac{(\sigma_{M,M-1}^1)^2 + o(\eta)}{\eta}$$
$$\geq \sum_{l=i}^{M-1} \frac{\beta_l v_l}{\sigma_{l,l+1}^S} \bigg/ \sum_{l=i}^{M-1} \frac{u_l}{\beta_l},$$

(48)
$$\frac{(\sigma_{1,2}^S)^2}{\eta} \rightarrow \frac{\sum_{l=2}^{i_0} \frac{\alpha_l}{\sigma_{l,l-1}^1} x_l + \frac{\sigma_{1,2}^S}{\alpha_M} \sum_{l=i_0}^{M-1} \frac{\alpha_l}{\sigma_{l,l+1}^S} v_l}{\sum_{l=2}^{i^*-1} \frac{y_l}{\alpha_l} + \frac{\alpha_M}{\sigma_{1,2}^S} \sum_{l=i^*}^{M-1} \frac{u_l}{\alpha_l}},$$
$$\frac{a_M}{\sigma_{1,2}^S} \rightarrow \left(\frac{\lambda_1}{\lambda_2}\right)^{(M-1)/2} \left(\frac{\lambda^{(M-1)/2} - \sqrt{\beta}}{\sqrt{\beta} \lambda^{(M-1)/2} - 1}\right).$$

PROOF: Consider first a state $i \leq i^* - 1 \leq i_0$. By Lemma 2, moving up a state must be optimal after an S -signal, and by Proposition 4(d), $\sigma_{1,2}^S \in (0, 1)$. Using Proposition 2, the respective optimality conditions are $\ell(i, S)/\bar{\ell}_i \geq 1$, with equality if $\sigma_{i,i+1}^S \in (0, 1)$, and $\ell(1, S)/\bar{\ell}_1 = 1$. Thus, using (4) and $\bar{\ell}_i \equiv \Delta_{i,i+1}^L/\Delta_{i+1,i}^H$, optimality implies the expression, with both inequalities tight if $\sigma_{i,i+1}^S \in (0, 1)$,

(49)
$$\ell(i, S)/\bar{\ell}_i \geq (1, S)/\bar{\ell}_1$$
$$\Leftrightarrow (f_i^H/f_1^H)(\Delta_{i+1,i}^H/\Delta_{2,1}^H) \geq (f_i^L/f_1^L)(\Delta_{i,i+1}^L/\Delta_{1,2}^L).$$

By (36) and (38), $(f_i^\theta/f_1^\theta)(\Delta_{i+1,i}^\theta/\Delta_{2,1}^\theta)$ is proportional to the next expression, plus an $o(\eta)$ term:

$$\eta \sum_{l=2}^i \frac{\alpha_l}{\sigma_{l,l-1}^1} \frac{1}{\mu_1^\theta} \left(\frac{\mu_1^\theta}{\mu_S^\theta}\right)^{l-1} + (\sigma_{1,2}^S)^2 \sum_{l=2}^i \frac{1}{\alpha_l} \left(\frac{\mu_S^\theta}{\mu_1^\theta}\right)^{l-1} + \sigma_{1,2}^S.$$

Substituting into (49) and solving for $(\sigma_{1,2}^S)^2/\eta$ yields the first inequality in (45), from the definitions of x_l, y_l in (44). Likewise, the optimality condition for $\sigma_{i,i-1}^1 > 0$ can be written in terms of the condition for $\sigma_{1,2}^S \in (0, 1)$ as $\ell(i, 1)/\bar{\ell}_{i-1} \leq \ell(1, S)/\bar{\ell}_1$, with equality if $\sigma_{i,i-1}^1 \in (0, 1)$; putting (36) and (38) into this inequality and solving for $(\sigma_{1,2}^S)^2/\eta$ yields the second inequality in (45).

The argument for states $i \geq i_0 \geq i^*$ is symmetric: by Proposition 4(d), $\sigma_{M,M-1}^1 \in (0, 1)$; and by Lemma 2, $\sigma_{i,i-1}^1 > 0$ and $\sigma_{i,i+1}^S > 0$. By Proposition 2, optimality then demands $\ell(i, S)/\bar{\ell}_i \geq \ell(M, 1)/\bar{\ell}_{M-1} \geq \ell(i, 1)/\bar{\ell}_{i-1}$: the first inequality is tight if $\sigma_{i,i+1}^S \in (0, 1)$ and the second is tight if $\sigma_{i,i-1}^1 \in (0, 1)$. Substituting (37) and (39), and solving for $(\sigma_{M,M-1}^1)^2/\eta$ yields (47).

To obtain (46), we again need $\ell(i, S)/\bar{\ell}_i \geq \ell(1, S)/\bar{\ell}_1$, with equality if $\sigma_{i,i+1}^S \in (0, 1)$, and $\ell(i, 1)/\bar{\ell}_{i-1} \leq \ell(1, S)/\bar{\ell}_1$, with equality if $\sigma_{i,i-1}^1 \in (0, 1)$. Write the first inequality (for $\sigma_{i,i+1}^S > 0$) as

$$(50) \quad \frac{f_i^H}{f_1^H} \frac{\Delta_{i+1,i}^H}{\Delta_{1,2}^H} \geq \frac{f_i^L}{f_1^L} \frac{\Delta_{i,i+1}^L}{\Delta_{1,2}^L}, \quad \text{where} \quad \frac{\Delta_{i+1,i}^\theta}{\Delta_{2,1}^\theta} = \frac{\Delta_{i+1,i}^\theta}{\Delta_{M,M-1}^\theta} \frac{\Delta_{M,M-1}^\theta}{\Delta_{i^*,i^*-1}^\theta} \frac{\Delta_{i^*,i^*-1}^\theta}{\Delta_{2,1}^\theta}.$$

For $i^* \leq i \leq i_0$, use (36) to get an expression for f_i^θ/f_1^θ , and use (38), (39), and the identity in (50) to get an expression for $\Delta_{i+1,i}^\theta/\Delta_{2,1}^\theta$. Substituting into the inequality in (50) and solving for $(\sigma_{1,2}^S)^2/\eta$ yields the first inequality in (46), using $\sigma_{M,M-1}^1/\beta_l = \alpha_M/\alpha_l$. The second inequality is similar.

To obtain the first expression in (48), since an optimal protocol has $\sigma_{1,2}^S \in (0, 1)$ and $\sigma_{M,M-1}^1 \in (0, 1)$, Proposition 2 demands $\ell(1, S)/\bar{\ell}_1 = \ell(M, 1)/\bar{\ell}_{M-1} = 1$. Recalling (4) and the definition $\bar{\ell}_i = \Delta_{i,i+1}^L/\Delta_{i+1,i}^H$, this requires that $(f_1^\theta/f_M^\theta)(\mu_S^\theta/\mu_1^\theta)(\Delta_{2,1}^\theta/\Delta_{M,M-1}^\theta)$ be state symmetric (i.e., equal in states $\theta = H, L$), which rearranges to the condition

$$\left(\frac{f_{i_0}^H \mu_S^H \Delta_{i^*,i^*-1}^H}{f_M^H \mu_1^H \Delta_{M,M-1}^H} \right) \left(\frac{f_{i_0}^L \Delta_{i^*-1,i^*}^L}{f_1^L \Delta_{1,2}^L} \right) = \left(\frac{f_{i_0}^L \mu_S^L \Delta_{i^*,i^*-1}^L}{f_M^L \mu_1^L \Delta_{M,M-1}^L} \right) \left(\frac{f_{i_0}^H \Delta_{i^*-1,i^*}^H}{f_1^H \Delta_{1,2}^H} \right).$$

Substitute (36), (37), (38), and (39) into this expression and solve for $(\sigma_{1,2}^S)^2/\eta$ to obtain the desired expression.

For the next expression in (48), by footnote 10, $f_1^H/f_M^H \rightarrow y_1^H/y_M^H$, and by Proposition 4(a), (b), and (d), $y_1^H/y_M^H \rightarrow (\prod_{j=1}^{M-1} \tau_{j+1,j}^H)/(\prod_{j=1}^{M-1} \tau_{j,j+1}^H) \rightarrow (\alpha_M/\sigma_{1,2}^S)(\mu_1^H/\mu_S^H)^{M-1}$. Put this into (34) and solve for $\prod_{j=1}^{M-1} \sigma_{j+1,j}^1/\sigma_{j,j+1}^S \equiv \alpha_M/\sigma_{1,2}^S$ to get the desired expression. Q.E.D.

STEP 3: *There exists η^* such that in an optimal protocol with $\eta < \eta^*$, (i) no memory state is both sticky up and sticky down; (ii) if i is sticky up, then $i + 1$ is fluid down; (iii) transition chances in the interior states are boundedly positive as*

$\eta \rightarrow 0$; (iv) *transition chances in the extreme states, $\sigma_{1,2}^S$ and $\sigma_{M,M-1}^1$, are asymptotic to $\sqrt{\eta}$.*

PROOF: Parts (i) and (ii) are immediate from (45), (46), and (47). For example, if a state $i \leq i^* - 1$ is sticky up, then the first expression in (45) holds with equality: $(\sigma_{1,2}^S)^2/\eta + o(\eta) = \sum_{l=2}^i x_l / \sum_{l=2}^i y_l$. But the RHS of this expression is strictly below both $\sum_{l=2}^i x_l / \sum_{l=2}^{i-1} y_l$ and $\sum_{l=2}^{i+1} x_l / \sum_{l=2}^i y_l$, implying that the second expression in (45) cannot hold with equality for either state i or $i + 1$; therefore, neither i nor $i + 1$ is sticky down in an optimal protocol. The proof for states $i \geq i^*$ is similar, given the corresponding expression (46) or (47).

For part (iv), since $\beta > 1$ implies (via (10)) that $i_0 \geq 2$ and since $i^* \geq 2$ if information matters, optimality of $\sigma_{2,3}^S > 0$ yields the first inequality below, and optimality of $\sigma_{3,2}^1 > 0$ yields the second inequality. Also, if $\sigma_{2,3}^S \in (0, 1)$, the LHS expression holds with equality:

$$\frac{x_2}{z_2} \leq \frac{(\sigma_{1,2}^S)^2 + o(\eta)}{\eta} \leq \frac{x_2 + x_3/\sigma_{2,3}^S}{z_2}, \quad \text{where}$$

$$z_2 = \begin{cases} y_2, & \text{if } i^* \geq 3, \\ \frac{\alpha_M}{\sigma_{1,2}^S} u_2, & \text{if } i^* = 2. \end{cases}$$

By (44), the LHS and RHS expressions are both positive and finite, thus immediately implying that $(\sigma_{1,2}^S)^2$ must be asymptotically proportional to η , as desired. The argument for state M is symmetric as long as $i_0 < M$. If $i_0 = M$, see the final paragraph of this proof.

For part (iii), suppose, by contradiction, that there is some state i with $\sigma_{i,i-1}^1 \rightarrow 0$ or $\sigma_{i,i+1}^S \rightarrow 0$. If $i \leq i_0 - 1$, choose the smallest such i , and if $i \geq i_0$, choose the largest such i . Suppose first that $i \leq i_0$ and $\sigma_{i,i-1}^1 \rightarrow 0$; then $\alpha_i \rightarrow 0$. But then if $i \leq i^* - 2$, the RHS expression in (45) tends to zero for state $i + 1$, since the final term in the denominator, y_i/α_i , explodes as $\alpha_i \rightarrow 0$, while all other terms in the expression are boundedly positive by construction. While if $i \geq i^* - 1$, then the RHS expression in (46) tends to zero in state $i + 1$, since the final denominator term u_i/α_i explodes, while all other terms are again positive. Either way, the optimality condition for $\sigma_{i+1,i}^1 > 0$ implies that $(\sigma_{1,2}^S)^2/\eta$ must tend to zero as $\eta \rightarrow 0$; but this violates the LHS expression in (45) at $i = 2$, which requires $(\sigma_{1,2}^S)^2/\eta \rightarrow x_2/y_2 > 0$ —a contradiction.

Suppose next that $i \geq i_0 + 1$ and that $\sigma_{i,i-1}^1 \rightarrow 0$, so $\beta_{i-1} \rightarrow \infty$ (using the fact that transition chances in interior states above i are boundedly positive by construction, and the part (ii) implication that $\sigma_{i,i-1}^1 \in (0, 1) \Rightarrow \sigma_{i-1,i}^S = 1$). Then the RHS expression in (47), evaluated at $i - 1$ rather than i , explodes as $\eta \rightarrow 0$, implying that optimality of $\sigma_{i-1,i-2}^1 > 0$ requires $(\sigma_{M,M-1}^1)^2/\eta \rightarrow \infty$. But if $i = M - 1$, this contradicts the second expression in (47) at $i = M - 1$ (which must hold with equality if $\sigma_{M-1,M-2}^1 \rightarrow 0$, so $(\sigma_{M,M-1}^1)^2/\eta \rightarrow v_{M-1}/u_{M-1}$,

which is finite by (44)), and if $i \leq M - 2$, then we have a contradiction to the first expression in (47) at $i = M - 2$, which requires $(\sigma_{M,M-1}^1)^2/\eta \rightarrow (v_{M-2}/\sigma_{M-1,M-2}^1 + v_{M-1})/u_{M-1}$ (which is finite by (44) and the fact that $\sigma_{M-1,M-2}^1$ is boundedly positive by construction).

And finally, if $\sigma_{i_0,i_0-1}^1 \rightarrow 0$, while transition chances in states above and below i_0 are boundedly positive, then $\alpha_M/\sigma_{1,2}^S = \prod_{j=2}^M \sigma_{j,j-1}^1/\sigma_{j-1,j}^S \rightarrow 0$ (using the result from the last paragraph that $\sigma_{M,M-1}^1$ and $\sigma_{1,2}^S$ go to zero at the same rate if i_0 is interior), contradicting the second expression in (48). This completes the proof that $\sigma_{i,i-1}^1$ must be boundedly positive as $\eta \rightarrow 0$ for all interior states i and the proof that $\sigma_{i,i+1}^S \rightarrow 0$ is symmetric.

Finally, to complete the proof of (iv) if $i_0 = M$, since transition chances in all interior states are boundedly positive as $\eta \rightarrow 0$ by part (iii), and since we showed that $\sigma_{1,2}^S$ is asymptotic to $\sqrt{\eta}$, the second expression in (48) can only hold if also $\sigma_{M,M-1}^1$ is asymptotic to $\sqrt{\eta}$. Q.E.D.

STEP 4: *There exists η^* such that in any optimal protocol with $\eta < \eta^*$, (i) all interior states are fluid up if $\lambda_1 \leq \lambda_2$, while all interior states outside of I^* are fluid down if $\lambda_1 \geq \lambda_2$; (ii) if $\lambda_1 \leq \lambda_2$, then there is a set of interior states I such that all states $i \notin I$ are fluid down, while all states $i \in I$ are sticky down. Moreover, for symmetric signals, $I \supseteq \{i^* + 1, \dots, i_0 - 1\}$.*

The proof is straightforward, but algebraically intensive; the proof is in Appendix C in the Supplemental Material.

REFERENCES

- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy*, 100, 992–1026. [2267]
- CHASE, W. G., AND K. A. ERICSSON (1982): "Skill and Working Memory," in *The Psychology of Learning and Motivation*, Vol. 16, ed. by G. H. Bower. New York: Academic Press. [2275]
- COWAN, N. (1995): *Attention and Memory: An Integrated Framework*. London: Oxford University Press. [2275]
- CRAWFORD, V., AND H. HALLER (1990): "Learning How to Cooperate: Optimal Play in Repeated Coordination Games," *Econometrica*, 58 (3), 571–595. [2264]
- DOW, J. (1991): "Search Decisions With Limited Memory," *Review of Economic Studies*, 58, 1–14. [2275]
- ERICSSON, K. A., AND W. KINTSCH (1995): "Long-Term Working Memory," *Psychological Review*, 102, 211–245. [2275]
- ERICSSON, K. A., AND J. J. STASZEWSKI (1989): "Skilled Memory and Expertise: Mechanisms of Exceptional Performance," in *Complex Information Processing: The Impact of Herbert A. Simon*, ed. by D. Klahr and K. Kovotsky. Hillsdale, NJ: Lawrence Erlbaum. [2275]
- FREIDLIN, M., AND A. D. WENTZELL (1984): *Random Perturbations of Dynamical Systems*. New York: Springer-Verlag. [2261]
- GILBOA, I., AND D. SCHMEIDLER (1995): "Case-Based Decision Theory," *Quarterly Journal of Economics*, 110 (3), 605–639. [2275]

- HELLMAN, M., AND T. COVER (1970): "Learning With Finite Memory," *The Annals of Mathematical Statistics*, 41 (3), 765–782. [2259,2268,2272,2274]
- JEHIEL, P. (2005): "Analogy-Based Expectation Equilibrium," *Journal of Economic Theory*, 123, 81–104. [2275]
- KANDORI, M., G. MAILATH, AND R. ROB (1993): "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, 61 (1), 29–56. [2261]
- LORD, C., L. ROSS, AND M. R. LEPPER (1979): "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence," *Journal of Personality and Social Psychology*, 37, 2098–2110. [2274]
- MULLAINATHAN, S. (1998): *A Memory Based Model of Bounded Rationality*. Cambridge: MIT. [2275]
- PICCIONE, M., AND A. RUBINSTEIN (1997): "On the Interpretation of Decision Problems With Imperfect Recall," *Games and Economic Behavior*, 20, 3–24. [2257,2262,2275,2279]
- RABIN, M., AND J. SCHRAG (1999): "First Impressions Matter: A Model of Confirmatory Bias," *Quarterly Journal of Economics*, 114 (1), 37–82. [2275]
- RADNER, R. (1962): "Team Decision Problems," *The Annals of Mathematical Statistics*, 33 (3), 857–881. [2264]
- SMITH, L., AND P. SØRENSEN (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398. [2267,2268]
- SMITH, L., P. SØRENSEN, AND J. TIAN (2012): "Informational Herding, Optimal Experimentation, and Contrarianism," Working Paper, UW Madison. [2267]
- WILSON, A. (2014): "Supplement to 'Bounded Memory and Biases in Information Processing'," *Econometrica Supplemental Material*, 82, <http://dx.doi.org/10.3982/ECTA12188>. [2262]

Dept. of Economics, University of Wisconsin–Madison, 1180 Observatory Dr., Madison, WI 53706, U.S.A.; awilson@ssc.wisc.edu.

Manuscript received May, 2004; final revision received June, 2014.