Bounded Memory in a Changing World: Biases in Behaviour and Belief*

Kalyan Chatterjee[†] Konstantin Guryev[‡] Tai-Wei Hu[§]

April 25, 2022

Abstract

A decision-maker faces a decision problem to choose an action, at a randomly determined time, to match an unknown state of nature. She has access to a sequence of signals partially informative of the current state of nature. The state of nature evolves according to a Markov chain. The decision-maker is subject to constraints on information-processing capacity, modelled here by a finite set of memory states. We characterize when optimal inference is possible with these constraints and, when it is not, what the optimal constrained inference is in two broad classes of environments. In the first class where the signals have similar strengths, optimal inference can be represented by simple rules corresponding to heuristics, like the "recency bias", which have been studied by experimental researchers. In the second class where one signal is very informative, the constrained optimal rule ignores the possibility of regime changes.

---

*We are grateful to Nageeb Ali, Yeon-Koo Che, Federico Echenique, Vijay Krishna, Daniel Monte, Stephen Morris, Ariel Rubinstein, Muhamet Yildiz, participants at the NSF-CEME Decentralisation Conference (2016), Pennsylvania Economic Theory Conference (2016), Cornell-PSU Macro Conference (2016) and the North American Summer Meetings of the Econometric Society (2016). We also thank Andrea Wilson and Yilmaz Kocer for reading an earlier draft, and thank an anonymous referee and the Editor, Marciano Siniscalchi, for suggestions that helped improve the paper.

[†]Pennsylvania State University. E-mail - kchatterjee@psu.edu

[‡]Pennsylvania State University. E-mail - kxg331@psu.edu

[§]University of Bristol. E-mail - taiwei.hu@bristol.ac.uk

1 Introduction

In this paper, we consider a model of learning by a single decision-maker, who is rational but constrained by a finite memory capacity. There are two states of the world, H and L, unknown to the decision-maker, and a potentially infinite sequence of signals that she can observe. Our model features "changing worlds", that is, the state of the world itself changes according to a Markov process, with transitions also being unobservable. This new feature allows us to investigate the interaction between the decision-maker's memory capacity and the persistence of the process that governs the transition of the states of the world. This interaction leads to several significant differences with the analysis of the same problem where the state of the world stays fixed forever. In particular, we are able to rationalize some well-known heuristics as optimal responses when the memory constraint is binding.

Some of these differences have analogues in the well-known discussion of systematic deviations from Bayesian inference. In particular, the "availability heuristic" described by Tversky and Kahnemann (1974) and the ignoring of regime change arise naturally in two different environments in our model. The first characterises decision-makers who rely only on the most recent signals while ignoring all past but relevant informative events. The latter is reminiscent of descriptions of otherwise well-informed observers ignoring the possibility of housing price decline in the 2008-9 financial crisis despite many informative signals about that possibility. Both results can be illustrated in environments with two signals only, h and ℓ. The first feature is prominent in environments where the two signals indicate different states of the world but have similar strengths. The second feature is prominent in environments where the two signals have very asymmetric strengths and one of them is very informative.

Following Wilson (2014), we consider a stylized model in which there is an exogenous probability in any given period (independently across periods and independent of the signals received) that the decision-maker is called to take a terminal action. The inverse of that probability is then the expected horizon of the decision problem. Without a memory constraint, a decision rule is then a mapping from the sequence of signals received so far to an action to be taken if required. As in Wilson (2014), we model memory constraints by finite automata, so every decision rule must be implementable through a finite number of internal memory states (not to be confused with states of the world). If called on exogenously to make a terminal choice, which ends the game, the decision-maker chooses an action based on the memory state she is in. Updating cannot be Bayesian, in general, because each memory state corresponds to a category of posterior beliefs about the state of the world so that only a finite set of such probabilities can be captured. After observing a signal, the decision-maker may choose to make a transition to another memory state and this represents the process of updating. Thus, as explained below, the timing is such that the decision only depends on the current memory state whilst the updating can depend both on this and on the signal received.

By considering changing worlds, it becomes possible to replicate optimal Bayesian inference in a constrained setting. When the expected horizon tends to infinity in the fixed-worlds environment, more and more signals will contribute to the optimal inference and a finite state automaton with its memory bounded will be unable to process the large amount of information on offer before a decision has to be made. With changing worlds, however, if the state of the world is not too persistent, observations from many periods prior to the current one should have diminishing relevance for a decision, if required to be made. This possibility comes from the drastically different dynamics of the Bayesian updating in changing worlds, which consists of two parts, information from the signal received and accounting for the transition of the state of the world. For each distinct signal, these two effects give rise to a fixed point for the posterior such that the posterior remains unchanged when that signal is received, and the fixed-point is the attraction point toward which the posterior moves. In the fixed-worlds environment, the only fixed points are zero and one. In the extreme case of the i.i.d. process, the posterior immediately jumps to its fixed point once a signal is received, regardless of the current posterior.

We exploit this feature and show that even severely restricted memory can replicate optimal inference. In particular, we show that when the two signals are of similar strengths, then the unconstrained optimum can be implemented by a simple rule where the decision-maker only takes the most recent signal into account when taking her action, provided that the persistence is sufficiently weak or the signals are sufficiently informative. This simple rule only requires two memory states to implement, and may be interpreted as the availability heuristic, as it only uses the most recent signals. But this is not biased reasoning as it implements the same decision as Bayesian inferences would.

Now, what if the signals are not informative enough for this simple rule to implement the unconstrained optimal payoff? In this case, an unbounded number of memory states is necessary to achieve that payoff, but we show that, for any given memory constraint (no matter how large), there is a range of strengths of the signals under which the availability heuristic is the constrained optimal rule. In other words, the simple rule that takes the action based only on the most recent signal, which requires only two memory states, performs better than other finite automata (even with randomization) with larger sizes. As a result, if we would impose a small cost for each additional memory state, the optimal number of the states would be two for a range of parameters. Thus, for a range of signal strengths, under zero cost the optimal number of memory states is unbounded, but a small cost brings the optimal number to two. Perhaps surprisingly, this range *expands* as the persistence increases. In this constrained case, the availability heuristic does exhibit a recency bias as a fully rational agent should take previous signals into account, but even for a small cost of memory capacity, the availability heuristic is the most efficient decision rule.

We also consider the other polar case where the constrained optimal rule completely ignores recent informative signals. When the information structure is very asymmetric in the sense that one of the two signals is very informative while the other is mildly so, we show that the constrained optimal rule is overconfident in the strong signal to the extent that the decision-maker is fully convinced of the state of the world forever after one such signal, completely ignoring the possibility of regime change and informative signals of such a change. This occurs even in the Shiryaev (1978) model of regime change, where one of the states of the world is absorbing but the other may change. For a prior that is relatively high on the absorbing state of the world, we show that a two memory-state finite automaton is optimal for any given size constraint, according to which after one strong signal, a "surprise," indicating the other (non-absorbing) state of the world, the decision-maker acts as if the state of the world is the non-absorbing one and ignores the possibility of regime change. In contrast, a perfectly rational decision-maker would eventually be convinced of the absorbing state of the world, if she received sufficiently many signals indicating this and hence would act completely differently. This result holds qualitatively for other priors and is robust to the introduction of cost of memory states.

Related literature

The strand of literature most related to our work pertains to the issue of learning with a finite state automaton, starting with Hellman and Cover (1970) which solves the simple hypothesis problem with finite memory, but using a limit-of-the-means objective function rather than discounted lifetime utility (or, equivalently, a continuation probability $1 - \eta$). In economics, Wilson (2014) adopts a similar approach but studies the case where η is close to zero and the state of nature is fixed. Hu (2022) studies constrained optimal finite automata in an asymmetric signal structure similar to ours, but in the Wilson (2014) setting with fixed worlds. Other related problems are studied by Kocer (2010), Monte (2007), and Monte and Said (2014). The study of decision making under finite state automata is closely related to the study of decision problems under imperfect recall as in Piccione and Rubinstein (1997) and other related papers, as will become clearer during our analysis.

We discuss certain points of comparison with the preceding papers. Extending results in Wilson (2014), we establish a generalised modified multi-self consistency in the changing-worlds environments, which implies an optimal partition of the posterior beliefs, each cell corresponds to a memory state and both transition and action rules are sequentially rational with respect to such beliefs. This result also allows us to derive conditions under which the optimal decision rule with unconstrained memory can be implemented by a deterministic finite state automaton. As in Monte and Said (2014), this demonstrates the lack of need for any sophistication in the decision rule in special environments as they also show that a two-state finite automaton can implement the unconstrained optimum. However, as in Hellman and Cover (1970), Monte and Said use the limit-of-means payoff criterion (that is, the limit case where η converges to zero) and focus on the symmetric case, and for the constrained case they only consider two- or three-state finite automata. Thus their result is considerably generalized in our paper. We show that the availability heuristic can outperform *all* finite automata with size no larger than K for *any* given K, no matter how large, for a range of persistence where the memory constraint is binding. This result does require $\eta > 0$, but η can be arbitrarily small. Moreover, we do not need symmetry for the result to hold, as long as the two signals are not too asymmetric in terms of their informativeness.

A pertinent observation relevant to our ignoring regime change result, in the Shiryaev model of regime change, appeared in a Boston Federal Reserve discussion paper of 2010, on the 2008 financial crisis. Writing on the optimistic forecasts of housing prices prevalent at that time, Gerardi et al. (2010) documented several regional indicators of a decline in housing prices (as well as a warning from Shiller, who is quoted in the paper) but somehow these were never sufficient to trigger a full-scale alarm about a crisis. Ignoring these "small" signals, most forecasters did not revise their forecasts downwards (something incompatible with a fully Bayesian analysis). Nevertheless, ex post many have argued that the increase in housing prices was a "bubble," and it would eventually decline. That is, the current state of nature itself was not fixed but could switch to an absorbing state, as in the Shiryaev model of regime change. Our results show that it can be (constrained) optimal to ignore the possibility of regime changes, despite the continual arrival of informative signals indicating otherwise.[1]

2 Framework

The framework generalizes the one in Wilson (2014) to allow for changing worlds. Time is discrete with infinitely many periods $t = 0, 1, 2, ...$ There are two possible states of the world, $\theta \in \Theta = \{H, L\}$. The prior at period 0 is given by $\mathbf{P}_0 \in \Delta(\Theta)$ and we use the notation $p_0 = \mathbf{P}_0(H)$. The state of the world $\theta_t \in \Theta$ evolves according to a Markov process over time and is unobservable, with the transition matrix given by Table 1, where Δ^θ is the probability that the state of the world $\theta \in \{H, L\}$ persists to the next period and $1 - \Delta^\theta$ is the probability that the state of the world switches from θ to the other state, with $\Delta^H,\ \Delta^L \in [0, 1]$. When $\Delta^H + \Delta^L = 1$, we have the i.i.d. environment. In contrast, when $\Delta^H + \Delta^L = 2$, $\Delta^H = 1 = \Delta^L$ and we have the

[1]Such behaviour has been used as a building block to understand the financial crisis; see, for example, Gennaioli et al. (2013), where investors neglect tail risks in formulating their strategies in a model of shadow banking.

	H	L
H	Δ^H	$1-\Delta^H$
L	$1-\Delta^L$	Δ^L

Table 1: Transition matrix of states of the nature

fixed-worlds environment. We assume

$$\Delta^H + \Delta^L \in [1, 2], \tag{1}$$

and hence both states are somewhat persistent (except for at the limit).

In each period, with probability η the decision-maker (DM) has to take a terminal action $a \in A = \{a^H, a^L\}$ with the utility function

$$u(a^H, H) = u^H > 0,\ u(a^L, L) = u^L > 0,\ u(a^H, L) = 0 = u(a^L, H).$$

The state of the world, though unobservable, can be inferred from signals received from a finite set X. Conditional on the current state of the world θ, the probability of receiving signal x is given by μ_x^θ. We use $\xi_x \equiv \mu_x^H/\mu_x^L$ to denote the likelihood ratio of signal x.

The timing of the model is as follows. In period 0 if the DM is called upon to make the final decision, then it will be done only based on the prior. Otherwise, at the end of the period the state of the world may transit according to the transition probabilities, and, after the transition, the DM receives a signal at the beginning of period 1, x_1. If the DM is called upon to take a decision in period 1, then x_1 is taken into account, as well as the possible transition (though unobservable) that may have happened. Otherwise, the state of the world may transit at the end of period-1 according to the transition probabilities, and so on. A decision rule, denoted by D, then maps a partial history of signal realizations, $\mathbf{x} = (x_1, ..., x_t)$, at period $t = 0, 1, 2, ...$, to a terminal action that the DM would take if called upon, where the history is empty when $t = 0$.

Without any memory constraint, the DM chooses the optimal action depending on the posterior belief at the time to make a decision, and we use p to denote the posterior on H. Given p, the optimal action is to choose a^H if $\xi(p) \equiv p/(1-p) > \xi^* \equiv u^L/u^H$ and it is a^L if $\xi(p) < \xi^*$, and the DM is indifferent between the two actions when $\xi(p) = \xi^*$, where $\xi(p)$ denotes the likelihood ratio for the posterior p. When the current posterior on H is p, the likelihood ratio of the updated posterior p' following a possible transition and then receiving signal x is given by

$$\xi(p') = \gamma(p; \xi_x) \equiv \frac{\xi(p)\Delta^H + (1 - \Delta^L)}{\xi(p)(1 - \Delta^H) + \Delta^L} \times \xi_x. \tag{2}$$

According to (2), the likelihood ratio of the new posterior is a product of two terms: the first term takes into account the transition probabilities of the state of the world and updates the posterior accordingly, and the second term is the likelihood ratio of the signal received and is multiplied by the updated posterior obtained in the first term. The effect of the second term is familiar: when $\xi_x > 1$, it increases the posterior and decreases it otherwise. The effect of the first term depends on the current position of the posterior relative to the invariant distribution, whose likelihood ratio is denoted by $\bar{\xi}$, and the first term always moves toward $\bar{\xi}$. Moreover, when $\Delta^H + \Delta^L = 1$, i.e., the i.i.d. case, the first term is always equal to $\Delta^H/(1-\Delta^H) = \bar{\xi}$ and is independent of $\xi(p)$. That is, in the i.i.d. case the new posterior is independent of the previous posterior and hence memory does not matter. When $\Delta^H = \Delta^L = 1$, i.e., the fixed-worlds environment, the first term is $\xi(p)$, and the memory matters most. As a result, the overall direction of the posterior change depends not only on ξ_x, but also on $\xi(p)$, and the following lemma shows that the direction can be characterized by the fixed-point of γ. Note that γ strictly increases with p as long as $\Delta^H + \Delta^L \in (1, 2)$.

Lemma 2.1. *Suppose that $\Delta^H + \Delta^L \in (1, 2)$. Then, for each $x \in X$, there is a unique $\bar{p}_x$ such that*

$$\xi(\bar{p}_x) = \gamma(\bar{p}_x; \xi_x). \tag{3}$$

$\bar{p}_x < \bar{p}_y$ if and only if $\xi_x < \xi_y$. Moreover, $\gamma(p; \xi_x) < \xi(p)$ if and only if $p > \bar{p}_x$.

A notable feature of Bayesian updating under changing worlds is that $\gamma(p, \xi_x) < \xi(p)$ if $p > \bar{p}_x$ while $\gamma(p, \xi_x) > \xi(p)$ if $p < \bar{p}_x$, that is, the posterior always moves toward the fixed-point $\bar{p}_x$. Moreover, except for the extreme case where $\Delta^H + \Delta^L = 1$ (that is, the i.i.d. case), $\gamma(p, \xi_x) \in (\xi(\bar{p}_x), \xi(p))$ if $\bar{p}_x < p$ and $\gamma(p, \xi_x) \in (\xi(p), \xi(\bar{p}_x))$ if $\bar{p}_x > p$, that is, the posterior will be closer to the fixed-point but would never reach it by receiving signal x (but could pass it or reach it by other signals). In the limit case where $\Delta^H + \Delta^L \to 2$ (that is, the fixed-world environment), $\bar{p}_x \to 1$ if $\xi_x > 1$ and $\bar{p}_x \to 0$ if $\xi_x < 1$. Finally, it is straightforward to verify that if $\xi_x = 1$, then $\xi(\bar{p}_x) = \bar{\xi}$. Hence, $\bar{p}_x > \bar{p}$ if $\xi_x > 1$ and $\bar{p}_x < \bar{p}$ if $\xi_x < 1$.

Finite automata and multiself consistency

We focus on decision rules that can be implemented by a finite automaton with a given number of memory states. A *stochastic finite-state automaton* (SFSA) or simply an *automaton* is a tuple $M = \langle Q, \sigma, d, q^o \rangle$, where Q is a finite set of *memory states*, $\sigma : Q \times X \to \Delta(Q)$ is a *transition function*, $d : Q \to A$ is an *action rule* and $q^o \in Q$ is the *initial memory state*.[2] The automaton makes decisions as follows: at period 0, if called upon to take the terminal action, it is based on $d(q_0)$ with $q_0 = q^o$, the initial memory state. In period 1, after receiving the signal x_1, the DM then moves to the next memory state according to the transition probabilities given by $\sigma(q_0, x_1)$. Similarly, at a given period t, the DM begins with the memory state $q_{t-1} \in Q$, and, after receiving signal x_t, the memory state moves to q_t according to the transition probabilities $\sigma(q_{t-1}, x_t)$. If called upon to take a terminal action, it is then given by $d(q_t)$. When all the transition rules are deterministic, we also call the automaton a *deterministic finite-state automaton* (DFSA).

Define $\mathcal{M}_K$ to be the set of all finite state automata with $|Q| \leq K$. Denote the ex-ante expected payoff from the automaton M to be $V(M)$. Given the memory constraint K, an optimal automaton is one which solves the maximisation problem

$$\max_{M \in \mathcal{M}_K} V(M). \tag{4}$$

Note that the optimization is constrained by the upper bound K on the number of memory states of the automaton, M. This maximization problem consists of two parts: first, a choice of the number of memory states below K and the action rule for each memory state; second, taking the number of memory states as given, optimization over transition probabilities. The first step is essentially a discrete problem while the second optimizes over continuous variables.

[2]Kalai and Solan (2003) show that it is with no loss of generality to assume that the action rule is deterministic.

To solve this optimization problem, our strategy is to first characterize the optimal transition rules for a given number of memory states. For this part we extend the principle of multiself consistency proposed by Piccione and Rubinstein (1997) for models of imperfect recall to our environment, a methodology borrowed from Wilson (2014). While this principle is essentially based on the first-order conditions for optimal transition probabilities, its main insight is that the optimal transition rules follow a version of "sequential rationality" in the sense that for each memory state, there is a corresponding posterior "belief" so that when the automaton lands at that state and receives a signal, it is optimal to choose the next memory state to go to by comparing the expected continuation payoffs according to that belief. Moreover, as will become clear later, it is better to consider such beliefs as categories than a precise posterior, and their updating rules are in general biased relative to the Bayesian benchmark.

The optimal number of memory states can then be found by comparing the optimal expected payoffs for each number of memory states less than K. In the fixed-world environment, the constraint K is always binding under the full support assumption, i.e., $\xi_x \in (0, \infty)$ for all x (c.f. Wilson (2014) and Hu (2022)). In other words, the unconstrained optimum is not implementable by any DFSA in the fixed-worlds environment. In contrast, K may not be binding under changing worlds. In particular, in the i.i.d. environment, only two memory states are needed to implement the unconstrained optimum. When K is binding, the main finding from the literature is that randomization can be optimal (c.f. Hellman and Cover (1970) and Wilson (2014)) in the fixed-world environment.

Now we move to the characterization of optimal transition rules, taking the number of memory states as given, which we denote by K for now. We show that the optimal transition rule maximizes the expected continuation value with respect to some categories of beliefs, and we first define continuation payoffs and a notion of "beliefs."

First note that once an automaton is fixed, the payoffs are completely determined by the evolution of the pair (θ, q) over time since actions only depend on $q \in Q$. This pair evolves according to a Markov process:

$$\mathbb{P}(\theta', q'|\theta, q) = \sum_{x \in X} \nu^{\theta}_{\theta'} \mu^{\theta'}_x \sigma(q, x)(q'), \tag{5}$$

where $\nu^{\theta}_{\theta} = \Delta^{\theta}$ and $\nu^{\theta}_{\theta'} = 1 - \Delta^{\theta}$ if $\theta' \neq \theta$. The expected payoff from M, $V(M)$, can be decomposed to be the sum of the expected payoff accumulated from the pairs $(q,\theta) \in Q \times \{H, L\}$, that is,

$$V(M) = \sum \left\{ f(q,\theta) u[d(q), \theta],\ (q,\theta) \in Q \times \{H, L\} \right\},$$

where

$$f(q,\theta) = \sum_{t=0}^{\infty} \eta(1-\eta)^t \sum \left\{ \mathbf{P}_0(\theta_0) \prod_{j=0}^{t-1} \nu_{\theta_{j+1}}^{\theta_j} \mu_{x_{j+1}}^{\theta_{j+1}} \sigma(q_j, x_{j+1})(q_{j+1}) : (q_t, \theta_t) = (q, \theta) \right\}, \tag{6}$$

and where the summation is over all $(\theta_j, q_j, x_{j+1}) \in \{H, L\} \times Q \times X$ with $q_0 = q^o$, $j = 0, ..., t-1$. Since the expressions $f(q,\theta)$ does not depend on the decision rule d, for d to be optimal, it must be the case that $d(q)$ solves

$$\max_{a \in A} \sum_{\theta = H, L} f(q, \theta) u(a, \theta),$$

that is, $d(q)$ needs to maximize the expected utility according to the "belief"

$$p(q) = \frac{f(q, H)}{f(q, H) + f(q, L)}.$$

Moreover, $f(q,\theta)$ satisfies the following recursive equations:

$$\begin{aligned} f(q,\theta) &= \eta \mathbf{P}_0(\theta) \mathbf{1}_{q=q^o} + (1-\eta) \sum_{\theta', q' \in Q, x \in X} \sigma(q', x)(q) \nu_{\theta}^{\theta'} \mu_x^{\theta} f(q', \theta') \qquad (7) \\ &= \sum_{\theta', q' \in Q, x \in X} [\eta \mathbf{P}_0(\theta) \mathbf{1}_{q=q^o} + (1-\eta) \nu_{\theta}^{\theta'} \mu_x^{\theta} \sigma(q', x)(q)] f(q', \theta'). \end{aligned}$$

That is, $f(q,\theta)$ is the stationary distribution under the transition probability from (q,θ) to (q',θ') given by

$$T(q', \theta' | q, \theta) = \sum_{x \in X} [\eta \mathbf{P}_0(\theta) \mathbf{1}_{q'=q^o} + (1-\eta) \nu_{\theta'}^{\theta} \mu_x^{\theta'} \sigma(q, x)(q')]. \tag{8}$$

The transition (8) would collapse to (5) if $\eta = 0$. For $\eta > 0$, the difference is that (8) features a return to the initial memory state q^o and initial distribution $\mathbf{P}_0$ with probability η every period. Intuitively, this discrepancy captures the constraints arising from the limited number of memory states. When arriving at the memory state q, the

DM cannot distinguish whether she is there right at the beginning or having been through a few other memory states before arriving there.

The principle of multiself consistency developed by Wilson (2014) extends the above logic not only to the decision rule d but also to the transition rule σ. For this purpose we need to take signals into account as well and we may then define "beliefs" at memory state q as follows:

$$\frac{p(q)}{1-p(q)} = \frac{f(q,H)}{f(q,L)}, \; \frac{p(q,x)}{1-p(q,x)} = \frac{\xi[p(q)]\Delta^H + (1-\Delta^L)}{\xi[p(q)](1-\Delta^H) + \Delta^L} \times \xi(x). \tag{9}$$

According to (9), the belief $p(q)$ is determined by $f(q,H)$ and $f(q,L)$, and $p(q,x)$ is obtained from that belief according to (2) using $p(q)$ as the initial posterior.

To characterize an optimal SFSA, we use $V_q(\theta)$ to denote the continuation value at memory state q conditional on the state of nature being θ. Two memory states are called *equivalent* if they share the same transition rules to any other states or their equivalents, and have the same specified action.

Theorem 2.1. *Consider the changing world environment under assumption (1). Let M be an optimal finite automaton that solves (4) for a given K without equivalent states. We rank the memory states in M according to*

$$p(q_1) \leq p(q_2) \leq \cdots \leq p(q_K),$$

with the convention that if $p(q_i) = p(q_{i+1})$ then $V_{q_i}(H) \leq V_{q_{i+1}}(H)$. Let $\Delta V^\theta_{i,j} = V_{q_i}(\theta) - V_{q_j}(\theta)$.

1. *$\Delta V^H_{i,j} < 0$ and $\Delta V^L_{i,j} > 0$ for all $i < j$, and $\Delta V^H_{j,i}/\Delta V^L_{i,j} \geq \Delta V^H_{k,j}/\Delta V^L_{j,k}$ for all $i < j < k$.*

2. *Define $\bar{\xi}_i = \Delta V^L_{i,i+1}/\Delta V^H_{i+1,i}$, $i = 1, ..., K-1$. Then, in M,*

   (a) for each q_i,

   $$\xi[p(q_i)] \in [\bar{\xi}_{i-1}, \bar{\xi}_i]; \tag{10}$$

   (b) for each q_i, q_j and x, $\sigma(q_i,x)(q_j) > 0$ only if

   $$\xi[p(q_i,x)] \in [\bar{\xi}_{j-1}, \bar{\xi}_j], \tag{11}$$

   where $\bar{\xi}_0 = 0$ and $\bar{\xi}_K = \infty$;

(c) $d(q_i) = a^H$ *only if* $\xi[p(q_i)] \geq \xi^*$ *and* $d(q_i) = a^L$ *only if* $\xi[p(q_i)] \leq \xi^*$.

Theorem 2.1 generalizes the characterization result in Wilson (2014) to changing worlds, based on the principle of modified multi-self consistency proposed by Piccione and Rubinstein (1997). See Online Appendix for formal details. As mentioned earlier, it requires the decision rule be optimal according to the beliefs $p(q_i)$, as $d(q_i) = a^H$ only if $\xi[p(q_i)] \geq \xi^*$ and $d(q_i) = a^L$ only if $\xi[p(q_i)] \leq \xi^*$. Moreover, the transitions are optimal w.r.t. such beliefs and the continuation values as well. Accordingly, Theorem 2.1 identifies a partition of beliefs based on the continuation values, $\{\bar{\xi}_i : i = 0, ..., K\}$, and, from the current memory state q_i and signal x, it is optimal to transit to state q_j only if $p(q_i, x)$, the posterior updated from $p(q_i)$ and signal x according to Bayes rule, lies within $[\bar{\xi}_{j-1}, \bar{\xi}_j]$.

In summary, the optimal SFSA shares the same structure as the unconstrained optimal rule, but uses $p(q)$ as the posterior. In this sense the DM who employs the optimal SFSA satisfies a notion of "sequential rationality." However, there is a key difference between the optimal SFSA and the unconstrained optimal rule—the beliefs $p(q)$ are not derived from Bayes rule. Instead, since $p(q)$'s are derived from the transition matrix (8), which is perturbed by η, these beliefs are biased. Finally, while $p(q_i)$ can be regarded as the "representative belief," it is better to consider categories of beliefs represented by $[\bar{\xi}_{i-1}, \bar{\xi}_i]$ for the memory state q_i, as it is optimal to transit to q_i whenever the "beliefs" land inside that range.

The modified multi-self consistency is based on arguments regarding local optimality, and hence Theorem 2.1 only provides necessary conditions for local optimality. However, it is still powerful for two reasons. First, it significantly reduces the number of possible randomizations, as Theorem 2.1 2 (b) implies that randomization can happen only if the posterior lies on the boundary. Wilson (2014) in fact used this fact to characterize optimal randomization for η close to zero in fixed-worlds environments. Our technique, however, focuses on a second merit. Again, by Theorem 2.1 2 (b), an optimal transition must be deterministic if $\xi[p(q_i), x] \in (\bar{\xi}_{j-1}, \bar{\xi}_j)$, that is, if the posterior lies within the interior of the boundaries. This turns out to be a crucial observation for us to identify some optimal heuristics, to which we turn next.

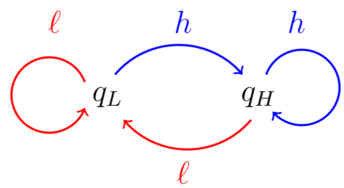

Figure 1: The availability heuristic: M_2^a

3 Recency bias

Here we consider situations that give rise to a simple heuristic rule that is constrained efficient, which exhibits recency bias when the constraint is binding. The recency bias means that the DM only uses the most recent signals to determine her action, disregarding her past observations. In contrast, without frictions the Bayesian inference would take the full history into account and past signals would be relevant for the current action. This bias is related to availability heuristics mentioned in Tversky and Kahneman (1974). These results would also illustrate the key differences between changing and fixed worlds. For this purpose, we focus on the case where $X = \{h, \ell\}$ with $\mu_h^H = \mu = \mu_\ell^L$, that is, h-signal and ℓ-signal are symmetric.

Now, in this environment, this simple heuristic that only utilizes the most recent signals can be represented by the DFSA in Figure 1, which we call M_2^a (the superscript "a" for *availability heuristics* and the subscript for the number of memory states), and in which $d(q_H) = a^H$ and $d(q_L) = a^L$. Thus, the decision-maker takes action a^H when called upon if the most recent signal is h, and she takes action a^L when called upon if the most recent signal is ℓ. We begin with a simple observation.

Lemma 3.1. *Suppose that $\Delta^H + \Delta^L \in [1, 2)$ and that $\mu_h^H = \mu = \mu_\ell^L$. There exists $\mu^* < 1$ such that for all $\mu \in [\mu^*, 1]$, M_2^a implements the unconstrained optimum with $q^o = q_H$ if $\xi(p_0) \geq \xi^*$ and $q^o = q_L$ if $\xi(p_0) < \xi^*$.*

For the extreme case where $\mu = 1$, the logic behind Lemma 3.1 is straightforward: when the two signals fully reveal the true states of the world, past observations are not relevant and only the current signal counts, even for the unconstrained decision-maker. Moreover, M_2^a is the *unique* DFSA that implements the unconstrained optimum (up to equivalent memory states), because the states of the world are changing ($\Delta^H + \Delta^L < 2$).

	H	L
H	$\Delta^H = \nu$	$1 - \Delta^H = 1 - \nu$
L	$1 - \Delta^L = 1 - \nu$	$\Delta^L = \nu$

	h	ℓ
H	$\mu_h^H = \mu$	$\mu_\ell^H = 1 - \mu$
L	$\mu_\ell^H = 1 - \mu$	$\mu_\ell^L = \mu$

Table 2: Left: transition matrix; right: signal distributions

Now, Lemma 3.1 gives a range of signal strengths under which M_2^a is the *unconstrained* optimal rule. When $\Delta^H + \Delta^L < 2$, as we have seen in Lemma 2.1, the fixed points $\bar{p}_h$ and $\bar{p}_\ell$ are within $(0, 1)$ for $\mu < 1$. Assuming that the prior lies within the two fixed-points, $(\bar{p}_\ell, \bar{p}_h)$, the threshold μ^* is the minimum μ for the following two conditions to hold. First, from $\xi(\bar{p}_h)$ as (the likelihood ratio of) the current posterior, one ℓ-signal is sufficient to bring (the likelihood ratio of) the updated posterior below ξ^*, and second, from $\xi(\bar{p}_\ell)$ as (the likelihood ratio of) the current posterior, one h-signal is sufficient to bring (the likelihood ratio of) the updated posterior above ξ^*. As a result, it is always optimal to take a^H (a^L) after a single h (ℓ)-signal. In contrast, in the fixed-worlds environment the fixed points are the extreme points zero and one, and hence for any $\mu < 1$, this condition cannot be fulfilled. Indeed, in the fixed-worlds environment, M_2^a implements the unconstrained optimum only when $\mu = 1$ exactly, and μ^* converges to one as $\Delta^H + \Delta^L$ converges to 2, i.e., the fixed-worlds environment.

Now, for $\mu < \mu^*$, M_2^a does not implement the unconstrained optimum. Assuming that $p_0 \in (\bar{p}_\ell, \bar{p}_h)$ at μ^*, then for slightly lower μ's the constraint becomes binding for any given K. To solve for the constrained optimal rule, we consider only the symmetric case here: $\Delta^H = \nu = \Delta^L$ with $\nu \in (0.5, 1)$, as illustrated in Table 2. This simplification allows us to rule out tedious discussions of various cases that may arise due to asymmetry. In this case, ν measures the persistence of the state of the world, with $\nu = 1$ being the fixed-worlds environment and $\nu = 0.5$ the i.i.d. case. Moreover, the stationary distribution has $\bar{\xi} = 1$.

Given the symmetry, we assume without loss of generality that $\xi^* > 1$. In this case, μ^* is determined by $\gamma(\bar{p}_\ell, \xi_h) = \xi^*$ provided that $\xi(p_0) \in [1, \xi^*]$. The following theorem shows that the unconstrained optimum requires unbounded number of memory states for μ below μ^*, but the constrained optimality is obtained by M_2^a.

Theorem 3.1. *Suppose that* $\Delta^H = \nu = \Delta^L \in (0.5, 1)$ *and that* $\xi(p_0) \in [1, \xi^*]$. *Let*

$K \geq 2$ *be given.*

1. *There exists* $\mu_K < \mu^*$ *such that the unconstrained optimum is not implementable by DFSA with* $|Q| \leq K$ *for all* $\mu \in (\mu_K, \mu^*)$.

2. *There exists* $\tilde{\mu} \in [\mu_K, \mu^*)$ *such that for all* $\mu \geq \tilde{\mu}$, M_2^a *with* $q^o = q_L$ *is the optimal SFSA among* $\mathcal{M}_K$.

According to Theorem 3.1 (1), as soon as μ moves below μ^*, the unconstrained optimum requires an unbounded number of memory states. The reason is the following. For μ slightly below μ^*, one h-signal takes the posterior slightly below ξ^* from $\xi(\bar{p}_\ell)$ and another h-signal will take it over. Now, consider the number of ℓ-signals needed from $\bar{p}_h$, denoted by N, so that the resulting posterior, denoted by p^N, has the following property. From $\xi(p^N)$, it takes two h-signals to cross ξ^* but from $\xi(p^{N-1})$ only one h-signal is needed. See Figure 2 for a graphical illustration. Since at $\mu = \mu^*$ we need $p^N = \bar{p}_\ell$, which cannot be satisfied unless N approaches infinity according to Lemma 2.1, for μ slightly below μ^* we need an arbitrarily large N. We can then choose μ_K to ensure that such $N > K$. This implies that the unconstrained optimal decision rule requires at least $K + 1$ memory states to implement, since following any two posteriors $\xi(p^n)$ and $\xi(p^{n'})$ with $n < n'$, we can find a sequence of signals ($N - n'$ ℓ signals followed by an h-signal) that lead to different actions under the unconstrained optimal rule, they must correspond to different memory states.

A formal argument uses the Myhill-Nerode Theorem to explicate this logic to prove the result.[3] Note that this only works for a range, as for μ close to half and hence as the signals become almost uninformative, the fixed-points converge to half as well and in that case the unconstrained optimal rule only sticks to a^L.

Theorem 3.1 implies a discontinuity in the number of memory states necessary to implement the unconstrained optimum. For $\mu \geq \mu^*$, only two memory states are required and M_2^a implements the unconstrained optimum, but for $\mu < \mu^*$ and slightly

[3]This theorem characterizes deterministic finite-state automata, a characterization based on the following observation: if two distinct partial histories of signals, $\mathbf{x}$ and $\mathbf{y}$, lead the automaton to the same memory state, q, then $\mathbf{xz}$ and $\mathbf{yz}$ will also result in the process reaching the same memory state, q', which might not be the same as q. This observation allows for a neat characterization of number of states needed to implement a decision rule, based on the categories of partial histories it induces.

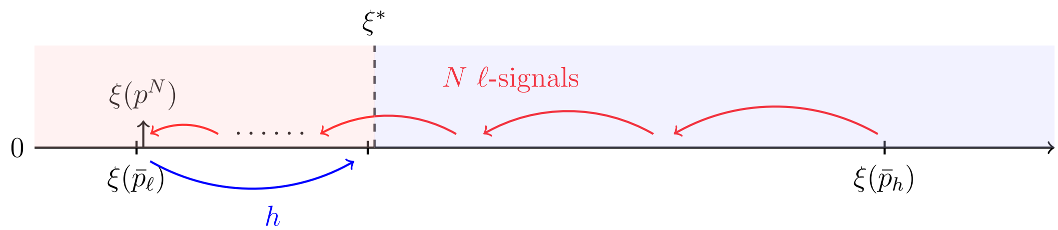

Figure 2: Determination of state-complexity for $\mu < \mu^*$

below, an unbounded number is required. Theorem 3.1 (2), however, shows that M_2^a is constrained optimal for a range of μ's below μ^*, and that M_2^a is the optimal SFSA under an arbitrary constraint K. In other words, for $\mu \geq \tilde{\mu}$, bigger finite automata (deterministic or stochastic) do not perform better than M_2^a.

Note that, nevertheless, the optimal payoff (constrained or unconstrained) is continuous as μ moves below μ^*, but this bounded-rationality approach emphasizes the optimal decision procedure and reveals discontinuity in that procedure. Notice also that the availability heuristic can also be implemented by a DM with bounded recall and with memory capacity of one. What we have shown here is that a very simple rule is indeed optimal, even if other more complicated schemes are available to the DM, including more complicated deterministic or stochastic transitions of memory states.

The proof of Theorem 3.1 (2) has three components. The first component shows that M_2^a uniquely implements the unconstrained optimum for all $\mu \geq \mu^*$, up to equivalent memory states. Note that μ^* is defined by the property that, starting from $\xi(\overline{p}_\ell)$, a single h-signal takes the posterior likelihood ratio to ξ^*. Multiplicity is not possible because with the prior in between the two fixed points, $\xi(\overline{p}_\ell)$ is never reached with a finite number of ℓ-signals, so the posterior can never be at ξ^*, the indifference level for the two actions (if the DM is called on to take one).

The second part of the proof demonstrates local optimality against other two-state automata, with μ slightly lower than μ^*. This uses Theorem 2.1 (b): Since M_2^a specifies that it is strictly optimal to follow the required transition at $\mu = \mu^*$, continuity of optimal value functions in μ ensures strict optimality close enough to μ^*.

Finally, to demonstrate that M_2^a is optimal against larger automata, we consider first those that implement very different transition and decision rules than M_2^a, which are clearly sub-optimal. However, there could be automata with these rules being very close to M_2^a but with replica states. Replica states of q_H and q_L will have the same actions but could transition to other replica states. But such transitions will be optimal if they have the same continuation values as the original state, so such automata cannot do strictly better. The beliefs in these replica states might, however, differ (from Theorem 2.1) because the paths taken to get to them might involve more or fewer transitions.

Intuitively, the different beliefs across different replicas of q_H reflect the different subsets of partial histories of signals that would lead to the original memory state q_H. Any of such partial histories will necessarily end with an h-signal, and we show that such beliefs will lie strictly above ξ^* for a range of μ's below μ^* based on the following observation. At μ^*, any posterior resulting from any partial history that ends with h will lie above ξ^*. Since K is finite, the process will always stay away from the fixed points, and hence at μ^* such belief will lie strictly above and we appeal to continuity to show that it remains so for μ slightly below. The formal argument uses the contraction mapping theorem.

In general, the threshold $\tilde{\mu}$ from Theorem 3.1 (2) depends on K, in addition to other parameters such as η and ν, and is not tractable. However, for $K = 2$, we are able to fully characterize $\tilde{\mu}$, and extend this result to the fixed-worlds environment ($\nu = 1$) and restore continuity.

Proposition 3.1. *Suppose that $\Delta^H = \nu = \Delta^L \in (0.5, 1]$ and that $\xi(p_0) = 1 < \xi^*$. Let $K = 2$. The threshold $\tilde{\mu}(\nu)$ above which M_2^a with $q^o = q_L$ is optimal among $\mathcal{M}_2$ is continuous and strictly increases with ν for all $\nu \in (0.5, 1]$. Moreover,*

$$\mu^*(0.5) = \tilde{\mu}(0.5) < \tilde{\mu}(1) < 1.$$

Note that by Lemma 3.1, M_2^a implements the unconstrained optimal rule for $\mu \in [\mu^*, 1]$ and hence Proposition 3.1 has a bite only for $\mu \in [\tilde{\mu}, \mu^*)$. However, as mentioned earlier, μ^* converges to unity as ν approaches unity, the fixed-worlds environment, a result follows from continuity of the unconstrained payoffs. Proposition 3.1 also implies that the threshold $\tilde{\mu}$ is strictly below one even at $\nu = 1$, and hence, even under the fixed world, there is a range of μ's for which M_2^a is constrained optimal. In the other

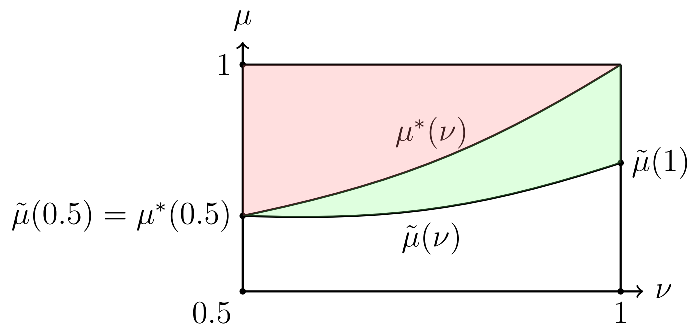

Figure 3: Red area: M_2^a unconstrained optimal; green area: M_2^a constrained optimal under $K = 2$

extreme, $\mu^* = \tilde{\mu}$ when $\nu = 0.5$, the i.i.d. environment. Thus, since both μ^* and $\tilde{\mu}$ strictly increase as ν increases from $\nu = 0.5$ and they begin at the same point at $\nu = 0.5$, but $\tilde{\mu}$ ends at a lower point as ν approaches the unity, the range where M_2^a is constrained optimal under $K = 2$ increases as ν increases, as depicted in Figure 3. In other words, while it is true that the availability heuristic is more likely to be optimal (unconstrained or constrained) when the information structure is closer to i.i.d., it is more likely to be constrained optimal, that is, to be optimal as a heuristic or a *simplified* rule when the information structure is closer to the fixed-worlds environment.

The less-is-more result also shows that our results are robust to introduction of costs of adding memory states. Specifically, consider a model where each additional memory state is costly and hence the optimal K is endogenous. To do so, suppose that it costs the DM cK to have K memory states. We have the following result.

Proposition 3.2. *Suppose that $\Delta^H = \nu = \Delta^L \in (0,5,1)$ and that $\xi(p_0) \in [1, \xi^*]$. Suppose that it costs cK to have K memory states. For each c sufficiently small, there exists $\tilde{\mu} < \mu^*$ such that for all $\mu \geq \tilde{\mu}$, optimal $K = 2$ and M_2^a is the optimal SFSA.*

According to Proposition 3.2, for any cost sufficiently small but positive, there is a range of μ's below μ^* under which the optimal choice is $K = 2$. In contrast, in the same range of parameters, for exactly zero cost the optimal K is unbounded. Again, this reveals the stark discontinuity in optimal procedure, and the constrained optimal rule is cognitively "cheap" but performs relatively well. This result then gives a precise formulation of the arguments for simple but effective heuristics used in Gigerenzer and Todd (1999).

Now we briefly discuss the role of the assumption that the information structure is given by Table 2. First the assumption that $\Delta^H = \Delta^L$ is not important, and all our results will hold qualitatively as long as $\Delta^H + \Delta^L \in (1, 2)$ and $\xi(p_0)$ is close to $\bar{\xi}$. This last assumption that $\xi(p_0)$ is close to $\bar{\xi}$, however, is indispensable for M_2^a to be optimal to avoid the possibility that the posteriors lie outside the fixed-points, $\bar{\xi}_h$ and $\bar{\xi}_\ell$. Regarding the symmetry assumption that $\mu_h^H = \mu = \mu_\ell^L$, both Theorem 3.1 and Proposition 3.1 will hold qualitatively if $\mu_h^H \neq \mu_\ell^L$ but are not too far from one another. This result follows almost immediately from Theorem 2.1. Indeed, since the logic of proof relies on the strict optimality of transition, it holds by continuity for small changes in the parameters, including μ_h^H. For the same reason, if we would introduce a third signal which is small in the sense that its likelihood ratio is close to unity, then it will be optimal to ignore it in the sense that there is no transition after seeing the small signal while the constrained optimal rule would follow M_2^a for the two more informative signals for ν close to one under Table 2. However, these results rely on μ_h^H and μ_ℓ^L being close to one another, and in the next section we discuss the other case where they are rather asymmetric.

4 Ignoring regime change

Up to now we have focused on the symmetric or a nearly symmetric situation where μ_h^H and μ_ℓ^L are close. Now we consider the asymmetric case where the strengths of the two signals are very different. In particular, we consider the benchmark case where $\mu_h^H = 1$ but keep $\mu_\ell^L = \mu \in (0, 1)$, and where $\Delta^H = 1 > \nu = \Delta^L$, the transition matrix according to the Shiryaev (1978) model of regime change. In this case, regardless of the value of ν, the ℓ-signal fully reveals that the current state of the world is L, but the h-signal increases the posterior on H. Moreover, $\Delta^H = 1$ means that the state H is an absorbing state. Hence, when $\nu < 1$, even when the fully revealing signal ℓ appears, there can still be *regime change* and learning is never dispensable. See also Table 3 for a summary of the transition matrix and signal distributions.

Here we focus on the case where $\xi(p_0) \geq \xi^*$, which allows us to demonstrate the

	H	L
H	$\Delta^H = 1$	$1 - \Delta^H = 0$
L	$1 - \Delta^L = 1 - \nu$	$\Delta^L = \nu$

	h	ℓ
H	$\mu_h^H = 1$	$\mu_\ell^H = 0$
L	$\mu_\ell^H = 1 - \mu$	$\mu_\ell^L = \mu$

Table 3: Left: transition matrix; right: signal distributions

main insight that holds for the other case as well. We will comment on the other case where $\xi(p_0) < \xi^*$ toward the end of the section. Back to the case where $\xi(p_0) \geq \xi^*$, there are two polar situations where the unconstrained optimum requires only two memory states. The first is the fixed-worlds environment, where $\nu = 1$, and the unconstrained optimum can be implemented by M_2^b depicted in the left panel in Figure 4 with $q^o = q_H$ and $d(q_H) = a^H$ and $d(q_L) = a^L$. In this case, an h-signal continues to increase the posterior on H and taking a^H is optimal before any ℓ-signal, but one ℓ-signal reveals that the state of nature is L and remains there forever. The other polar situation occurs with ν close to zero and regime change is very likely from state L, and M_2^a implements the unconstrained optimum. Intuitively, for ν small, an h-signal is sufficient to bring the posterior above ξ^*, regardless of the current posterior. The condition on ν for this to happen is given by

$$\frac{\nu}{1-\nu}\xi^* \leq \frac{1}{1-\mu}. \tag{12}$$

If $K = 2$, the constraint is binding for ν below one and any value above the maximum value that makes (12) hold, a threshold denoted by $\bar{\nu}$. That is, for $\nu \in (\bar{\nu}, 1)$, any DFSA with two memory states cannot implement the unconstrained optimum; in fact, the number of memory states needed converges to infinity as ν approaches one. Indeed, for ν close to one, after one ℓ-signal it would require a large number of h-signals to bring the posterior to reach ξ^* again, a number that converges to infinity as ν approaches one. We have the following theorem regarding the constrained optimal rule when ν is in between these two polar situations.

Theorem 4.1. *Suppose that $\Delta^H = 1$ and $\Delta^L = \nu \in [0, 1]$ and that $\mu_h^H = 1 > \mu = \mu_\ell^L$. Suppose that $\xi(p_0) \geq \xi^*$.*

1. *For any $K \geq 2$ given, there exists $\tilde{\nu}_K < 1$ such that for all $\nu \geq \tilde{\nu}_K$, the optimal SFSA is M_2^b with $q^o = q_H$.*

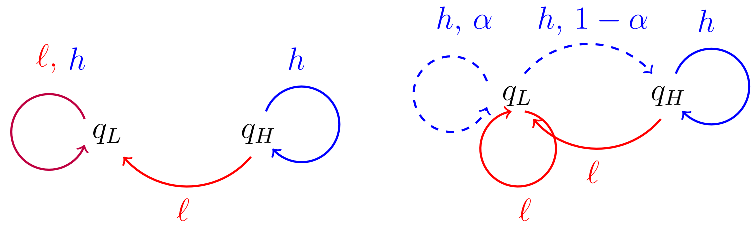

Figure 4: Left: M_2^b; Right: $M_2^a(\alpha)$

2. *Let $K = 2$. For $\nu \in (\bar{\nu}, \tilde{\nu}_2)$, the optimal SFSA is $M_2^a(\alpha)$with $q^o = q_H$ and with $\alpha \in (0, 1)$.*

According to Theorem 4.1 (1), M_2^b is the constrained optimal rule for a range of changing-worlds environments. As a result, even though there is still a positive probability of regime change after receiving an ℓ-signal, the DM behaves as if she is fully convinced that the state of the world remains at L forever and ignores the possibility of regime change, with her belief stuck at L. In contrast, an unconstrained DM will continue to change her posterior after any number of ℓ-signals. In fact, since H is an absorbing state, the unconstrained DM's posterior will eventually move to H with full conviction almost surely, a stark contrast to the constrained optimal rule.

Intuitively, for ν close to one, the DM does not expect regime change to happen. Once she receives an ℓ-signal and is convinced of state L, h-signal is a very weak signal for the regime change. Indeed, after an ℓ-signal, it would require a large number of h-signals to convince the unconstrained decision-maker that the state has really changed. A constrained decision-maker, in contrast, faces scarcity in her cognitive resources and might not have enough memory states in her automaton to make such an inference and therefore ignores regime change. Moreover, any stochastic scheme the DM may employ to respond to an h-signal at q_L would in fact make belief at q_L more concentrated at state L and hence make the transition to q_H even less attractive. To see this, under M_2^b, both signals would leave the automaton at q_L with the associated belief based on a mixture of those signals. But a stochastic scheme would have the automaton to leave q_L with some probability at h-signal and hence the corresponding belief would be based on fewer h-signals, leaving it with a higher concentration of ℓ-signals and hence a higher probability on state L. As a result, the constrained optimal rule ignores any

h-signal all together after an ℓ-signal, even though they are informative about regime change. The same logic would hold as long as the ℓ-signal is sufficiently informative about state L, that is, as long as μ_h^H is sufficiently close to one.

Theorem 4.1 (1) also illustrates a less-is-more result, and, for all $\nu \in [\tilde{\nu}_K, 1]$, M_2^b is better than other SFSA with $|Q| \leq K$, either with more memory states or with randomization or both. Thus, a simple decision rule outperforms more complicated ones. This shows that ignoring regime change is a powerful heuristic when the probability of regime change is small, and captures the casual observation that people tend to ignore events with small probability, even though those events imply significant change in stakes. Moreover, one can introduce cost of additional memory states and for an appropriate range of costs the optimal K would be two.

As in Theorem 3.1 (2), we appeal to Theorem 2.1 to prove Theorem 4.1 (1), and, again, the main difficulty is due to the large replica finite automata of M_2^b. As there, it is the beliefs that may differ across the replica memory states. However, for any replica of q_H, to reach it the partial history contains only h-signals, and hence, the resulting belief will float above the threshold. The belief for replicas of q_L is more complicated, as it may involve many h- and ℓ-signals. If $\nu = 1$, one ℓ-signal is sufficient to send the belief to zero. For ν slightly below, however, a slight complication arises as some replicas may have very small probability to be transitioned into. Nevertheless, a simple limit argument shows that the beliefs are still arbitrarily small if ν is close to one.

For lower ν's, according to Theorem 4.1 (2), under $K = 2$, the memory constraint is binding and the optimal SFSA is $M_2^a(\alpha)$ as depicted in the right panel of Figure 4 with $q^o = q_H$ and $d(q_H) = a^H$ and $d(q_L) = a^L$. Note that $M_2^a(\alpha)$ can be regarded as a convex combination between M_2^a and M_2^b, and Theorem 4.1 shows that the optimal SFSA moves from M_2^a gradually to M_2^b as ν increases from zero to one.

We also obtain some comparative statics with respect to ν and η. By (12), the threshold $\bar{\nu}$ only depends on μ and ξ^*. In contrast, the threshold $\tilde{\nu}_K$ depends on K and η as well. Note that, however, neither threshold depends on p_0. When $K = 2$ we can find a closed-form solution for $\tilde{\nu}_K$, and we depict the range of different constrained optimal rules in Figure 5. As depicted there, $\tilde{\nu}_2$ converges to one as η approaches zero, and $\tilde{\nu}_2$ converges to $\bar{\nu}$ as η approaches one. Thus, the red area under which M_2^b is optimal, where $\nu \geq \tilde{\nu}_2$, expands as η increases, while the area under which

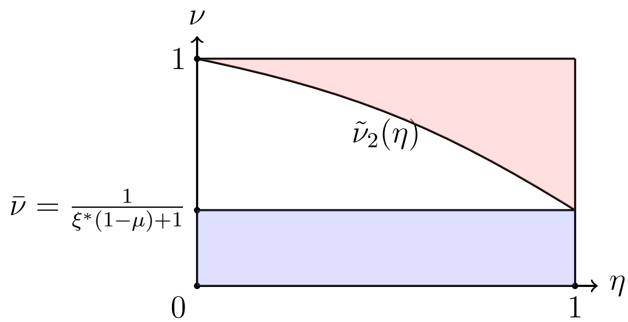

Figure 5: Red area: M_2^b constrained optimal under $K = 2$; blue area: M_2^a unconstrained optimal; white area: $M_2^a(\alpha)$ optimal with $\alpha \in (0, 1)$

randomization is optimal with $\alpha \in (0, 1)$ in $M_2^a(\alpha)$, shrinks. That is, the heuristic of ignoring regime change is more relevant when η is away from zero, and hence such heuristic thinking would be more important for impatient DM's but stochastic schemes would be more important for patient DM's.

Finally, we briefly discuss the case $\xi(p_0) < \xi^*$, while formal details can be found in the Online Appendix. As in the previous case, if ν is sufficiently small and satisfies (12), M_2^a can implement the unconstrained optimum, but with $q^o = q_L$. In contrast to the previous case, at the other extreme where $\nu = 1$, while the unconstrained optimum is still implementable by finitely many memory states, the required number of memory states depends on the informativeness of the signals as well as the distance between $\xi(p_0)$ and ξ^*. More precisely, let

$$N = \lceil \frac{\ln(\xi^*) - \ln[\xi(p_0)]}{\ln(\xi_h)} \rceil. \tag{13}$$

If $\nu = 1$, the unconstrained optimum can be implemented by a DFSA analogous to M_2^b but it needs N additional memory states to account for the fact that one needs N h-signals to cross ξ^* from $\xi(p_0)$. Thus, the unconstrained optimum is implementable if and only if $K \geq N+2$, and Hu (2022) shows that if $K < N+2$, randomization is needed for those transitional memory states under the fixed world. Under changing worlds, we show that there is a memory state analogous to q_L in M_2^b, which is the memory state that the (constrained) automaton goes to whenever an ℓ-signal is received and is self-absorbing. We then obtain a result analogous to Theorem 4.1 for $\xi(p_0) < \xi^*$

as well: for any given K, for a range of ν's below one, the constrained optimal finite automaton ignores regime change in the sense that once an ℓ-signal is received, the automaton enters a memory state that is self-absorbing and takes action a^L.

5 Conclusion

In this paper we studied learning with capacity-constrained information processing, when the unknown state of the world itself evolves over time . By considering two broad classes of information structures, we illustrated new effects due to limited memory. The first is a recency bias according to which the decision-maker only uses the latest signal to determine her action, disregarding useful information from the past. In the changing-world environment with two signals that are similar in strength, a simple two-state finite automaton that employs the availability heuristic outperforms larger ones and is the constrained optimal rule. In a model of regime change with a breakthrough signal, however, the constrained optimal rule simply ignores the possibility of regime change but is fully convinced of the strong signal in the past, disregarding the incoming informative signals indicating otherwise. We also endogenized the optimal memory constraints by introducing costs on memory states, and have demonstrated that these biases are in fact optimal responses even to a small cost of adding additional memory states.

References

[1] Gennaioli, Nicola, Andrei Shleifer, and Robert W. Vishny (2013). "A model of shadow banking," *The Journal of Finance*, 68, 1331-1363.

[2] Gerardi, Kristopher, Christopher L. Foote, and Paul Willen 2010. "Reasonable people did disagree: Optimism and pessimism about the US housing market before the crash," FRB of Boston Public Policy Discussion Paper 10-5.

[3] Gigerenzer, Gerd and Todd, Peter M. (1999). Simple Heuristics That Make Us Smart, Oxford University Press.

[4] Hellman, Martin E., and Cover, Thomas M. (1970). "Learning with Finite Memory," *The Annals of Mathematical Statistics*, 41, 765-782.

[5] Hellman, Martin E., and Cover, Thomas M. (1971). "On memory saved by randomization," *The Annals of Mathematical Statistics*, 1075-1078.

[6] Hu, Tai-Wei (2022). "Forgetful Updating: Ignoring Small Signals and Approximating Big Ones," forthcoming, *Economic Theory.*

[7] Kalai, Ehud, and Eilon Solan (2003). "Randomization and simplification in dynamic decision-making," *Journal of Economic Theory*, 111, 251-264.

[8] Kocer, Yilmaz (2010). "Optimal Plans with Bounded Memory are Modified Multiself Consistent", working paper.

[9] Monte, Daniel (2007). "Learning with bounded memory in games," Ph.D thesis, Yale University.

[10] Monte, Daniel, and Maher Said (2014). "The value of (bounded) memory in a changing world," *Economic Theory*, 56, 59-82.

[11] Nerode, Anil (1958). "Linear Automaton Transformations," *Proceedings of the AMS* , 9, 541-544.

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

[13] Rubinstein, Ariel (1986), "Finite automata play the repeated prisoner's dilemma", *Journal of economic theory*, 39, 83-96.

[14] Shiryaev, Albert N. (1978). *Optimal stopping rules,* Springer.

[15] Stokey, Nancy L. (1989). *Recursive methods in economic dynamics,* Harvard University Press.

[16] Tversky, Amos and Kahneman, Daniel (1974). "Judgment under Uncertainty: Heuristics and Biases", *Science*, 185, 1124-1131.

[17] Wilson, Andrea (2014). "Bounded memory and biases in information processing", *Econometrica*, 82, 2257-2294.

6 Appendix: Proofs

Proof of Lemma 2.1

Let $x \in X$ be given. Then, $\bar{p}_x = p$ solves

$$\frac{p}{1-p} = \xi_x \times \frac{\Delta^H p + (1-\Delta^H)(1-p)}{(1-\Delta^L)p + \Delta^L(1-p)}.$$

This then gives a quadratic equation of p:

$$(1-\Delta^L)p^2 + (\Delta^L - \xi_x \Delta^H)p(1-p) - \xi_x(1-\Delta^H)(1-p)^2 = 0.$$

Assuming that $\Delta^H + \Delta^L \in (1,2)$, when $p = 0$, the left-side is $-\xi_x(1-\Delta^H) < 0$, and when $p = 1$, it is $(1-\Delta^L) > 0$. Since quadratic equations can have at most two zeros, this implies that there exists a unique zero within $(0,1)$. In contrast, when $\Delta^H + \Delta^L = 2$, the solutions are $p = 0$ and $p = 1$. When $\Delta^H + \Delta^L = 1$,

$$\frac{\bar{p}_x}{1-\bar{p}_x} = \xi_x \times \frac{\Delta^H}{1-\Delta^L}. \tag{14}$$

Finally, $\gamma(p;\xi)$ increases with ξ. Hence, $\bar{p}_x$ increases with ξ_x as well. □

Proof of Lemma 3.1

Let $\Delta^H + \Delta^L \in [1,2)$ be given. For each $\mu \in (0.5, 1)$, define

$$g_h(\mu) = \gamma(\bar{p}_\ell, \xi_h) \text{ and } g_\ell(\mu) = \gamma(\bar{p}_h, \xi_\ell).$$

Note that both $\bar{p}_h$ and $\bar{p}_\ell$ are functions of μ, since $\xi_h = \mu/(1-\mu)$ and $\xi_\ell = (1-\mu)/\mu$. It is straightforward to verify that

$$\lim_{\mu\to 1} g_h(\mu) = \infty = \lim_{\mu\to 1} \xi(\bar{p}_h) \text{ and } \lim_{\mu\to 1} g_\ell(\mu) = 0 = \lim_{\mu\to 1} \xi(\bar{p}_\ell).$$

By continuity, then, there exists μ^* such that for all $\mu \in [\mu^*, 1)$,

$$g_h(\mu) \geq \xi^*, \; g_\ell(\mu) \leq \xi^*, \; \xi(\bar{p}_h) \geq \xi(p_0), \text{ and } \xi(\bar{p}_\ell) \leq \xi(p_0). \tag{15}$$

Now, we show that M_2^a implements the unconstrained optimum with $q^o = q_H$ if $\xi(p_0) \geq \xi^*$, a case we consider here; the other case is similar by symmetry. It follows from (15) that (a) the posterior $p \in (\bar{p}_\ell, \bar{p}_h)$ after any sequence of signals: this follows from $p_0 \in (\bar{p}_\ell, \bar{p}_h)$ and Lemma 2.1; (b) for all $p \in (\bar{p}_\ell, \bar{p}_h)$, $\gamma(p, \xi_h) \geq \xi^*$ and $\gamma(p, \xi_\ell) \geq \xi^*$, and hence it is optimal to take action a^H after seeing an h-signal and to take action a^L after seeing an ℓ-signal, regardless of the history.

Proof of Theorem 3.1

Part (1) We first compute

$$
\begin{aligned}
\xi(\bar{p}_h) &= \frac{\nu(\xi_h - 1) + \sqrt{\nu^2(\xi_h - 1)^2 + 4(1-\nu)^2\xi_h}}{2(1-\nu)}, \\
\gamma(\bar{p}_h, \xi_\ell) &= \frac{\nu(\xi_h - 1) + \sqrt{\nu^2(\xi_h - 1)^2 + 4(1-\nu)^2\xi_h}}{2(1-\nu)\xi_h^2},
\end{aligned}
$$

and by symmetry, $\xi(\bar{p}_\ell) = 1/\xi(\bar{p}_h)$ and $\gamma(\bar{p}_\ell, \xi_h) = 1/\gamma(\bar{p}_h, \xi_\ell)$. It is straightforward to verify that both $\bar{p}_h$ and $\gamma(\bar{p}_\ell, \xi_h)$ strictly increase with μ, and $\gamma(\bar{p}_\ell, \xi_h) = 1 = \xi(\bar{p}_h) < \xi^*$ when $\mu = 0.5$. Thus, there exists a unique $\mu^* > 0.5$ such that $\gamma(\bar{p}_\ell, \xi_h) = \xi^* < \xi(\bar{p}_h)$.

Now, for each $\mu \in (0.5, \mu^*)$, define $N(\mu)$ as follows. First, for any sequence of signals $\mathbf{x} \in \{h, \ell\}^*$, extend $\gamma(p, \mathbf{x})$ to denote the posterior after seeing $\mathbf{x}$ from p as the prior. Now, let $N(\mu)$ be the smallest N such that

$$\gamma(\bar{p}_h, \ell^N \circ h) < \xi^*. \tag{16}$$

Thus, by definition, $\gamma(\bar{p}_h, \ell^n \circ h) > \xi^*$ for all $n < N(\mu)$. Now we show that the unconstrained optimal payoff can not be achieved by any DFSA with $|Q| \leq N(\mu)$. To do so, we construct a sequence of partial histories, $\mathbf{x}^i$, $i = 1, ..., N(\mu)$ that each must belong to a distinct cell in the Nerode-Myhill Theorem. Now, let M be sufficiently large so that $\gamma(p_0, h^M)$ is sufficiently close to $\bar{p}_h$ and hence $N(\mu)$ is determined in the same way as (16) if we begin with $\gamma(p_0, h^M)$ instead of $\bar{p}_h$. Now, let

$$\mathbf{x}^i = h^M \circ \ell^i, \ i = 0, 1, 2, ..., N-1. \tag{17}$$

Then, for any $i < j$, let $\mathbf{y}^{i,j} = \ell^{N-j} \circ h$. Then,

$$\gamma(p_0, \mathbf{x}^i \circ \mathbf{y}^{i,j}) > \xi^* \text{ but } \gamma(p_0, \mathbf{x}^j \circ \mathbf{y}^{i,j}) < \xi^*. \tag{18}$$

This implies that $\mathbf{x}^i$ and $\mathbf{x}^j$ cannot be in the same cell and each requires a distinct memory state.

Finally, $N(\mu)$ converges to infinity as μ converges to μ^* by Lemma 2.1 and is weakly increasing in μ. As a result, for any given K, there exists a threshold μ_K so that for all $\mu \in (\mu_K, \mu^*)$, $N(\mu) > K$.

Part (2) By Lemma 3.1, M_2^a implements the unconstrained optimum at $\mu = \mu^*$. It is also the unique one up to equivalent memory states: the only possible indifference occurs at ξ^*, but, at μ^*, the posterior never reaches ξ^*; thus, for any K, $V(M_2^a) > V(M')$ for other SFSA $M' \in \mathcal{M}_K$ whose randomization probabilities are ϵ away from that in M_2^a and this inequality holds for a range of μ's below 1. Now we demonstrate local optimality of M_2^a for a range of μ's below μ^* by appealing to Theorem 2.1. A key feature of M_2^a is that

$$V_{q_H}(H) - V_{q_L}(H) = \eta u^H;\ V_{q_L}(L) - V_{q_H}(L) = \eta u^L. \tag{19}$$

This implies that

$$\bar{\xi} = u^L/u^H = \xi^*$$

is the threshold below which q_L is optimal and above which q_H is optimal to transit to.

We need to handle equivalent states. Consider the case where $\xi(p_0) < \xi^*$; the other case is similar by symmetry. Fix some $K' \leq K$, and consider a SFSA M with $K' = I + J$ memory states that is a replica of M_2^a. We show that there is no local deviations that can do better than M_2^a from M. Note that since it is equivalent to M_2^a, for any memory state equivalent to q_θ, its continuation value is still $V_{q_\theta}(H)$ and $V_{q_\theta}(L)$ as given in M_2^a and hence satisfies (19) above, for both $\theta = H, L$. Let equivalent states of q_H be denoted by $q_H^1, ..., q_H^I$, and q_L be denoted by $q_L^1, ..., q_L^J$, with

$$\sigma(q_\theta^i, h)(q_H^j) = \alpha_{\theta ij},\ \sigma(q_\theta^i, \ell)(q_L^j) = \beta_{\theta ij}. \tag{20}$$

Then, we have

$$\begin{aligned}
f(q_H^j, H) &= \sum_{i=1,..,I} [f(q_H^i, H)\nu_H^H + f(q_H^i, L)\nu_H^L]\alpha_{Hij}(1-\eta)\mu \\
&+ \sum_{i=1,...,J} [f(q_L^i, H)\nu_H^H + f(q_L^i, L)\nu_H^L]\alpha_{Lij}(1-\eta)\mu, \\
f(q_H^j, L) &= \sum_{i=1,..,I} [f(q_H^i, H)\nu_L^H + f(q_H^i, L)\nu_L^L]\alpha_{Hij}(1-\eta)(1-\mu) \\
&+ \sum_{i=1,...,J} [f(q_L^i, H)\nu_L^H + f(q_L^i, L)\nu_L^L]\alpha_{Lij}(1-\eta)(1-\mu),
\end{aligned}$$

$$
\begin{aligned}
f(q_L^j, H) &= \eta p_0 g(q_L^j) + \sum_{i=1,..,I} [f(q_H^i, H)\nu_H^H + f(q_H^i, L)\nu_H^L]\beta_{Hij}(1-\eta)(1-\mu) \\
&+ \sum_{i=1,\dots,J} [f(q_L^i, H)\nu_H^H + f(q_L^i, L)\nu_H^L]\beta_{Lij}(1-\eta)(1-\mu), \\
f(q_L^j, L) &= \eta(1-p_0)g(q_L^j) + \sum_{i=1,..,I} [f(q_H^i, H)\nu_L^H + f(q_H^i, L)\nu_L^L]\beta_{Hij}(1-\eta)\mu \\
&+ \sum_{i=1,\dots,J} [f(q_L^i, H)\nu_L^H + f(q_L^i, L)\nu_L^L]\beta_{Lij}(1-\eta)\mu.
\end{aligned}
$$

We claim that

$$
\bar{\xi}_\ell < \frac{f(q_L^j, H)}{f(q_L^j, L)} < \xi(p_0) \le \xi^* < \frac{f(q_H^i, H)}{f(q_H^i, L)} < \bar{\xi}_h \text{ for all } i = 1, .., I, \text{ and } j = 1, .., J. \quad (21)
$$

To show this, first note that $\{f(q, H), f(q, L)\}_{q \in Q}$ are the stationary distribution of a finite Markov chain and hence it is the fixed point of the transition mapping, which is a contraction (Stokey and Lucas, 1989, Lemma 11.3 and Theorem 11.4). So we only need to show that the inequalities in (21) are preserved by the transition, assuming that the input distribution satisfies it with weak inequalities.

Consider the first set of inequality in (21). We show that for each $\{f(q, H), f(q, L)\}_{q \in Q}$ satisfying (21) with weak inequalities,

$$
\xi^* \le \frac{[f(q, H)\nu_H^H + f(q, L)\nu_H^L]}{[f(q, H)\nu_L^H + f(q, L)\nu_L^L]} \frac{\mu}{1-\mu} \le \bar{\xi}_h.
$$

Moreover, the first inequality is strict for $q = q_H^i$ and the second is strict for $q = q_L^i$. Note that the middle term is equal to $\gamma[p(q), \xi_h]$ with $p(q)/[1-p(q)] = f(q, H)/f(q, L)$. Now, since at μ, $\gamma(\bar{p}_\ell, \xi_h) \ge \xi^*$, the first inequality follows from (21) which implies that $p(q) \in [\bar{p}_\ell, \bar{p}_h]$. When $q = q_H^i$, the inequality is strict as (21) implies that $p(q)/[1 - p(q)] > \xi^*$. The second inequality follow from the fact that $p(q) \le \bar{p}_h$, and, when $q = q_L^i$, the inequality is strict as $p(q_L^i) < \bar{p}_h$. The result then follows for $f(q_H^j, H)/f(q_H^j, L)$ immediately as the new $f(q_H^j, H)/f(q_H^j, L)$ is in between the ratios given by the middle term in (21).

Now we turn to q_L^j. We show that for each $\{f(q, H), f(q, L)\}_{q \in Q}$ satisfying (21) with weak inequalities,

$$
\bar{\xi}_\ell \le \frac{[f(q, H)\nu_H^H + f(q, L)\nu_H^L]}{[f(q, H)\nu_L^H + f(q, L)\nu_L^L]} \frac{1-\mu}{\mu} \le \xi(p_0).
$$

Moreover, the first inequality is strict for $q = q_H^i$ and the second is strict for $q = q_L^i$. Again, note that the middle term is equal to $\gamma[p(q), \xi_\ell]$ with $p(q)/[1-p(q)] = f(q,H)/f(q,L)$, and the argument follows a similar reasoning to above. $\square$

Proof of Proposition 3.1

Since $\xi(p_0) = 1 = \bar{\xi}$, it is straightforward to verify that

$$f^H = \frac{\eta p_0 + (1-\eta)(1-\nu)}{1-(1-\eta)(2\nu-1)} = 0.5 \text{ and } f^L = \frac{\eta(1-p_0) + (1-\eta)(1-\nu)}{1-(1-\eta)(2\nu-1)} = 0.5,$$

where $f^\theta = f(q_H, \theta) + f(q_L, \theta)$ in M_2^a. This then implies that

$$\begin{aligned} f(q_L, H) &= \eta p_0 + 0.5(1-\eta)(1-\mu) = 0.5[\eta + (1-\eta)(1-\mu)]; \\ f(q_L, L) &= \eta(1-p_0) + (1-\eta)0.5\mu = 0.5[\eta + (1-\eta)\mu]; \\ f(q_H, H) &= 0.5\mu(1-\eta); \\ f(q_H, L) &= 0.5(1-\mu)(1-\eta). \end{aligned}$$

To ensure local optimality, we need to check

$$\xi[p(q_H)] = \frac{f(q_H, H)}{f(q_H, L)} > \xi^*, \qquad \xi[p(q_H, \ell)] = \xi[p(q_H)]\xi_\ell < \xi^*; \tag{22}$$

$$\xi[p(q_L)] = \frac{f(q_L, H)}{f(q_L, L)} < \xi^*, \qquad \xi[p(q_L, h)] = \xi[p(q_L)]\xi_h > \xi^*. \tag{23}$$

The first inequality in (22) then becomes

$$\xi[p(q_h)] = \frac{\mu}{1-\mu} = \xi_h > \xi^*. \tag{24}$$

For the second, we have

$$\begin{aligned} \xi[p(q_H, \ell)] &= \frac{\xi[p(q_H)]\nu + (1-\nu)}{\xi[p(q_H)](1-\nu) + \nu} \times \xi_\ell \\ &= \frac{\xi_h \nu + (1-\nu)}{\xi_h(1-\nu) + \nu} \times \xi_\ell = \frac{\nu + (1-\nu)\xi_\ell}{\xi_h(1-\nu) + \nu} < 1 = \xi(p_0) \le \xi^*, \end{aligned}$$

where the first inequality follows from $\xi_\ell < 1 < \xi_h$. This proves the second inequality in (22). Now consider (23). First,

$$\xi[p(q_L)] = \frac{\eta + (1-\eta)(1-\mu)}{\eta + (1-\eta)\mu} < 1 = \xi(p_0),$$

since $\mu > 1/2$. Moreover,

$$\begin{aligned}
\xi[p(q_L,h)] &= \frac{[\eta+(1-\eta)(1-\mu)]\nu+[\eta+(1-\eta)\mu](1-\nu)}{[\eta+(1-\eta)(1-\mu)](1-\nu)+[\eta+(1-\eta)\mu]\nu} \times \xi_h \\
&= \frac{\eta+(1-\eta)[\nu+(1-2\nu)\mu]}{\eta+(1-\eta)[(1-\nu)-(1-2\nu)\mu]}\frac{\mu}{1-\mu},
\end{aligned}$$

which decreases with ν and increases with μ. Thus, if we let $\tilde{\mu}(\nu)$ to be the lowest μ satisfying $\xi[p(q_L,h)] > \xi^*$, M_2^a is locally optimal for all $\mu \geq \tilde{\mu}$. Moreover, $\tilde{\mu}(\nu)$ increases with ν. At $\nu = 1$, $\tilde{\mu}(1) < 1$. At $\nu = 0.5$, using (24), $\tilde{\mu}(0.5) = \mu^*(0.5)$. $\square$

Proof of Proposition 3.2

Let $\hat{c}$ be the difference between the expected payoff from M_2^a against the payoff from taking a^L immediately. Let $c < \hat{c}$, and let K be such that $cK > p_0u^H + (1-p_0)u^L$. Then, by Theorem 3.1, there exists $\tilde{\mu}$ such that M_2^a is optimal among $\mathcal{M}_K$ for all $\mu \geq \tilde{\mu}$. In that range it is then optimal to choose $K = 2$. $\square$

Proof of Theorem 4.1

Proof of (1) Let $K > 2$ be given. We know that M_2^b is the unique optimal SFSA when $\nu = 1$. Thus, for any K, $V(M_2^b) > V(M')$ for other SFSA $M \in \mathcal{M}_K$ whose randomization probabilities are ϵ away from that in M_2^b and this inequality holds for a range of ν's below 1. Now we demonstrate local optimality of M_2^b for a range of ν's below 1 by appealing to Theorem 2.1.

First we compute the continuation values for M_2^b:

$$\begin{aligned}
V_{q_H}(H) &= u^H,\ V_{q_L}(L) = \frac{\eta}{1-(1-\eta)\nu}u^L,\ V_{q_L}(H) = 0, \\
V_{q_H}(L) &= (1-\eta)\left\{\nu[\mu V_{q_L}(L) + (1-\mu)V_{q_H}(L)] + (1-\nu)V_{q_H}(H)\right\}, \\
&= \frac{1}{1-\nu(1-\eta)(1-\mu)}\left\{\frac{\mu(1-\eta)\nu\eta}{1-(1-\eta)\nu}u^L + (1-\eta)(1-\nu)u^H\right\}.
\end{aligned}$$

Thus, the threshold above which transiting to q_H is optimal according to Theorem 2.1 (and below which transiting to q_L is optimal), denoted by $\tilde{\xi}$, is given by

$$\tilde{\xi} = \frac{V_{q_L}(L) - V_{q_H}(L)}{V_{q_H}(H) - V_{q_L}(H)} = \frac{\eta\xi^* - (1-\eta)(1-\nu)}{1-\nu(1-\eta)(1-\mu)}.$$

Note that $\xi[p(q,\ell)] = 0$ for both q. Thus, local optimality only requires

$$\xi[p(q_H)] > \tilde{\xi} \text{ and } \xi[p(q_L,h)] < \tilde{\xi}. \tag{25}$$

As in the proof of Theorem 3.1 (2), we need to handle equivalent states. Fix some $K' \leq K$, and consider a SFSA M with $K' = I + J$ memory states that is a replica of M_2^b. We show that there is no local deviations that can do better than M_2^b from M. Note that since it is equivalent to M_2^b, for any memory state equivalent to q_θ, its continuation value is still $V_{q_\theta}(H)$ and $V_{q_\theta}(L)$ as given in M_2^b above, for both $\theta = H, L$. Let equivalent states of q_H be denoted by $q_H^1, ..., q_H^I$, and q_L be denoted by $q_L^1, ..., q_L^J$, with

$$\sigma(q_H^i,h)(q_H^j) = \alpha_{hij},\ \sigma(q_H^i,\ell)(q_L^j) = \alpha_{\ell ij},\ \sigma(q_L^i,h)(q_L^j) = \beta_{hij},\ \sigma(q_L^i,\ell)(q_L^j) = \beta_{\ell ij}. \tag{26}$$

Then,

$$\begin{aligned}
f(q_H^i,H) =& \eta p_0 g(q_H^i) + (1-\eta)\left\{\sum_{j=1}^{I}[f(q_H^j,L)(1-\nu) + f(q_H^j,H)]\alpha_{hji}\right\},\\
f(q_H^i,L) =& \eta(1-p_0) g(q_H^i) + (1-\eta)\left\{\sum_{j=1}^{I} f(q_H^j,L)\nu(1-\mu)\alpha_{hji}\right\},\\
f(q_L^i,H) =& (1-\eta)\left\{\sum_{j=1}^{J}[f(q_L^j,H) + f(q_L^j,L)(1-\nu)]\beta_{hji}\right\},\\
f(q_L^i,L) =& (1-\eta)\left\{\sum_{j=1}^{J} f(q_L^j,L)\nu[(1-\mu)\beta_{hji} + \mu\beta_{\ell ji}] + \sum_{k=1}^{I} f(q_H^k,L)\nu\mu\alpha_{\ell ki}\right\}.
\end{aligned} \tag{27}$$

As in the proof of Theorem 3.1 (2), we show that $p(q_H^i)$ and $p(q_L^j,h)$ satisfy (25) for all i and j, and we follow the same methodology, i.e., we take the equations in (27) as simultaneous equations and use the contractions mapping theorem to show that the inequalities in (25) hold.

Now, we show that if the input $\{f(q_H^i,H), f(q_H^i,L)\}$ satisfies the first inequality in (25), i.e., $f(q_H^i,H)/f(q_H^i,L) \geq \tilde{\xi}$ for all $i = 1, ..., I$, the the output $\{f'(q_H^i,H), f'(q_H^i,L)\}$ satisfies $f'(q_H^i,H)/f'(q_H^i,L) > \tilde{\xi}$ for all $i = 1, ..., I$. Now, note that

$$\frac{f'(q_H^i,H)}{f'(q_H^i,L)} \geq \min\left\{\frac{p_0}{1-p_0}, \frac{f(q_H^j,L)(1-\nu) + f(q_H^j,H)}{f(q_H^j,L)\nu(1-\mu)}, j = 1, ..., I\right\},$$

and hence it suffices to show that

$$
\begin{aligned}
\frac{p_0}{1-p_0} &> \tilde{\xi}, && (28)\\
\frac{f(q_H^j, L)(1-\nu) + f(q_H^j, H)}{f(q_H^j, L)\nu(1-\mu)} &= \frac{1}{\nu(1-\mu)}\left[(1-\nu) + \frac{f(q_H^j, H)}{f(q_H^j, L)}\right] > \tilde{\xi}. && (29)
\end{aligned}
$$

The inequality (28) follows immediately from $\xi(p_0) \geq \xi^*$. The inequality (29) follows immediately from $f(q_H^i, H)/f(q_H^i, L) \geq \tilde{\xi} > 0$, and $\nu(1-\mu) < 1$.

Now consider the second inequality in (25). We first show that the inequality holds at $\nu = 1$ uniformly across the transition probabilities given by (26), and then appeals to continuity to show that it holds for ν slightly below. First note that when ν converges to the unity, the third equation in (27) immediately implies that

$$
f(q_L, H) = \sum_{j=1}^{J} f(q_L^j, H) = (1-\eta)\frac{\sum_{j=1}^{J} f(q_L^j, L)(1-\nu)}{\eta}
$$

converges to zero uniformly across the transition probabilities, and hence $f(q_L^j, H)$ also converges uniformly. We claim that $f(q_L^j, H)/f(q_L^j, L)$ converges to zero across the transitions probabilities as well. Note that $f(q_L^i, L) > 0$ except for the limit case where β_{hji}, $\beta_{\ell ji}$ and α_{hki} converge to zero for all j and k. Thus, we need to show that $f(q_L^i, H)/f(q_L^i, L)$ converges to zero first by taking those transition probabilities to zero, and then taking ν to the unity. This ensures continuity at the limit. Now fix some i. Note that when β_{hji}, $\beta_{\ell ji}$ and α_{hki} all converge to zero for all j, $f(q_L^j, L) > 0$ for some $j \neq i$, even at the limit. Now, consider the case that $\lim_{\beta_{hji}\to 0} f(q_L^j, L) > 0$ and take β_{hji} to zero. Now, by applying the L'Hopital's Rule,

$$
\begin{aligned}
\lim_{\beta_{hji}\to 0} \xi(q_L^i) &= \frac{\lim_{\beta_{hji}\to 0}[f(q_L^j, H) + f(q_L^j, L)(1-\nu)]}{\lim_{\beta_{hji}\to 0} f(q_L^j, L)\nu(1-\mu)} \\
&= \frac{\lim_{\beta_{hji}\to 0} f(q_L^j, H)}{\lim_{\beta_{hji}\to 0} f(q_L^j, L)\nu(1-\mu)} + \frac{1-\nu}{\nu(1-\mu)}.
\end{aligned}
$$

Thus,

$$
\lim_{\nu\to 1}\lim_{\beta_{hji}\to 0} \xi(q_L^i) = \frac{\lim_{\beta_{hji}\to 0,\nu\to 1} f(q_L^j, H)}{\lim_{\beta_{hji}\to 0,\nu\to 1} f(q_L^j, L)(1-\mu)} = 0,
$$

where we have used $\lim_{\beta_{hji}\to 0,\nu\to 1} f(q_L^j, H) = 0$. Similarly, this implies that $\xi(q_L^i, h)$ converges to zero as well uniformly as ν converges to one. So there exists $\tilde{\nu}_K < 1$ such that the inequality (25) holds for all $\nu \geq \tilde{\nu}_K$.

Proof of (2) Here we consider $K = 2$. Since after seeing ℓ, $p(q, \ell) = 0$ for any q, Theorem 2.1 implies that the only relevant randomization to consider is $\sigma(q_L, h)(q_L) = \alpha = 1 - \sigma(q_L, h)(q_H)$ for some $\alpha \in [0, 1]$. We first compute the continuation values

$$\begin{aligned}
V_{q_H}(H) &= u^H,\ V_{q_L}(H) = \frac{(1-\eta)(1-\alpha)u^H}{1-(1-\eta)\alpha}, \\
V_{q_L}(L) &= \frac{\eta[1-(1-\eta)\nu(1-\mu)]u^L}{\{[1-(1-\eta)\nu][1-(1-\eta)\nu(1-\mu)\alpha]\}} \\
&+ \frac{(1-\eta)(1-\nu)(1-\alpha)[1-\alpha(1-\eta)^2\nu(1-\mu)]}{[1-(1-\eta)\alpha]\{[1-(1-\eta)\nu][1-(1-\eta)\nu(1-\mu)\alpha]\}}u^H, \\
V_{q_H}(L) &= \frac{(1-\eta)(1-\nu)}{1-(1-\eta)\nu(1-\mu)}\left\{1+\frac{(1-\eta)\nu\mu(1-\alpha)[1-\alpha(1-\eta)^2\nu(1-\mu)]}{[1-(1-\eta)\alpha][1-(1-\eta)\nu][1-(1-\eta)\nu(1-\mu)\alpha]}\right\}u^H \\
&+ \frac{\eta(1-\eta)\nu\mu}{[1-(1-\eta)\nu][1-(1-\eta)\nu(1-\mu)\alpha]}u^L.
\end{aligned}$$

Now, note that the ex ante payoff is given by $p_0 V_{q_H}(H) + (1 - p_0)V_{q_H}(L)$. Since $V_{q_H}(H) = u^H$, optimal α is independent of p_0. Let $F(\alpha; \nu) = V_{q_H}(L)$. Optimal α is determined by maximizing $F(\alpha; \nu)$. We compute the derivative of F w.r.t. α:

$$\begin{aligned}
F'(\alpha;\nu) = &\frac{(1-\eta)^2\eta\nu\mu}{[1-(1-\eta)\nu][1-(1-\eta)\nu(1-\mu)\alpha]^2} \\
&\times\left\{\frac{(1-\nu)[\alpha^2(1-\eta)^2\nu(1-\mu)-1]}{[1-(1-\eta)\alpha]^2}u^H+\nu(1-\mu)u^L\right\}.
\end{aligned} \tag{30}$$

Note that the term

$$\left\{\frac{(1-\nu)[\alpha^2(1-\eta)^2\nu(1-\mu)-1]}{[1-(1-\eta)\alpha]^2}u^H+\nu(1-\mu)u^L\right\}$$

is strictly decreasing in α for all $\nu < 1$ while the first term is always positive. This implies that the optimal α is unique and is determined by the FOC. Moreover, the optimal solution increases with ν. Now,

$$\lim_{\nu\to 1} F'(\alpha;\nu) = \frac{(1-\eta)^2\mu(1-\mu)}{[1-(1-\eta)(1-\mu)\alpha]^2}u^L > 0$$

for all α uniformly, and hence the optimal $\alpha = 1$ for ν close to 1 by continuity, that is for $\nu \geq \tilde{\nu}_2$. Finally, $F'(0, \nu) \leq 0$ if and only if (recall the condition (12))

$$-(1-\nu)u^H + \nu(1-\mu)u^L \leq 0 \ \Leftrightarrow \nu \leq \bar{\nu},$$

that is, optimal $\alpha = 0$ if and only if $\nu \leq \bar{\nu}$. □

A Online Appendix

A.1 Multiself consistency and Proof of Theorem 2.1

We first extend the modified multi-self consistency to the setup here. Recall the expression $f(q, \theta)$ given by (6) and beliefs $p(q)$ and $p(q, x)$ given by (9).

Definition A.1. A SFSA M satisfies *modified multi-self consistency* under $\mathbf{P}_0$ if

1. for each memory state $q \in Q$ with $\sum_\theta f(q, \theta) > 0$, each signal x, and any q' such that $\sigma(q, x)(q') > 0$,

$$p(q,x)V_{q'}(H)+[1-p(q,x)]V_{q'}(L) \geq p(q,x)V_{q''}(H)+[1-p(q,x)]V_{q''}(L) \text{ for all } q'' \in Q; \tag{31}$$

2. for each memory state $q \in Q$ with $\sum_\theta f(q, \theta) > 0$ and $a = d(q)$,

$$p(q)u(a, H) + [1 - p(q)]u(a, L) \geq p(q)u(a', H) + [1 - p(q)]u(a', L) \text{ for all } a' \in A. \tag{32}$$

The following result is crucial for the proof of Theorem 2.1.

Proposition A.1. *Suppose that M is an optimal SFSA under prior $\mathbf{P}_0$ among those with $|Q| \leq K$.*

1. *(Modified Multi-self Consistency) It satisfies modified multi-self consistency under prior $\mathbf{P}_0$.*

2. *(Revelation Principle) For any $q, q' \in Q$,*

$$p(q)V_q(H) + [1 - p(q)]V_q(L) \geq p(q)V_{q'}(H) + [1 - p(q)]V_{q'}(L). \tag{33}$$

Proof. For any pairs of states of nature and memory states, (θ, q) and (θ', q'), define the set

$$W_{(\theta,q),(\theta',q')} = \bigcup_{n=1}^{\infty} W^n_{(\theta,q),(\theta',q')},$$

where for each $n = 1, 2, ...,$

$$W^n_{(\theta,q),(\theta',q')} = \{[(\theta, q), x_1; (\theta_1, q_1), x_2; ....; (\theta_{n-1}, q_{n-1}), x_n; (\theta', q')] : x_i \in X, q_i \in Q, \theta_i \in \Theta\},$$

that is, the set of possible state transitions from q to q'. Given a state of nature θ and $\mathbf{w} \in W^n_{(\theta,q),(\theta',q')}$, define

$$\mathbb{P}(\mathbf{w}) = \eta(1-\eta)^{n-1} \times \prod_{i=1}^{n} \nu_{\theta_i}^{\theta_{i-1}} \mu_{x_i}^{\theta_i} \sigma(q_{i-1}, x_i)(q_i),$$

where $(\theta_0, q_0) = (\theta, q)$ and $(\theta_n, q_n) = (\theta', q')$. The expected payoff from the SFSA is then

$$V = \sum_{\theta,\theta',q} \mathbf{P}_0(\theta) \sum_{\mathbf{w} \in W_{(\theta,q^o),(\theta',q)}} \mathbb{P}(\mathbf{w}) u[d(q), \theta']. \tag{34}$$

We now prove (31) and (32).

First, consider (32). Suppose, by contradiction, that for some memory state $\hat{q}$ with $f(\hat{q}, \theta) > 0$ such that (32) does not hold, and hence there are actions $a = d(q)$ and $a' \in A$ with the inequality in (32) reversed with a strict inequality. By (6), $f(\hat{q}, \theta) = \sum_{\theta'} \sum_{\mathbf{w} \in W_{(\theta',q^o),(\theta,\hat{q})}} \mathbf{P}_0(\theta') \mathbb{P}(\mathbf{w})$, this then implies that

$$\sum_{\theta',\theta} \mathbf{P}_0(\theta') \sum_{\mathbf{w} \in W_{(\theta',q^o),(\theta,\hat{q})}} \mathbb{P}(\mathbf{w}) u(a, \theta) < \sum_{\theta',\theta} \mathbf{P}_0(\theta) \sum_{\mathbf{w} \in W_{(\theta',q^o),(\theta,\hat{q})}} \mathbb{P}(\mathbf{w}) u(a', \theta). \tag{35}$$

Now, consider the alternative SFSA M', which differs from M only in that $d'(\hat{q}) = a'$. From (34) and (35) it follows that M' gives a strictly higher expected payoff than M, a contradiction to the optimality of M.

Now consider (31). Suppose, by contradiction, that $\sigma(q, x)(q') > 0$ and that for some $q'' \neq q'$,

$$p(q, x) V_{q'}(H) + [1 - p(q, x)] V_{q'}(L) < p(q, x) V_{q''}(H) + [1 - p(q, x)] V_{q''}(L). \tag{36}$$

We denote $p' = \sigma(q, x)(q')$ and $p'' = \sigma(q, x)(q'')$. Now, fix all other transition probabilities other than p' and p'', each term $\mathbb{P}(\mathbf{w})$ in V given by (34) is a polynomial of (p', p'') and, since $\eta \in (0, 1)$, V is differentiable w.r.t. (p', p''). Since M is optimal and $p' = \tau(q, x)(q') > 0$, the FOCs require that $\frac{\partial}{\partial p'} V \geq \frac{\partial}{\partial p''} V$. However, we show below that (36) implies that

$$\frac{\partial}{\partial p''} V > \frac{\partial}{\partial p'} V, \tag{37}$$

a contradiction to the optimality of M.

To prove (37), it is straightforward to verify that

$$\frac{\partial}{\partial p'}V = \sum_{\theta,\theta',\hat{q}} \mathbf{P}_0(\theta) \sum_{\mathbf{w}\in W_{(\theta,q^o),(\theta',\hat{q})}(q,x;q')} \varphi_{(q,x;q')}(\mathbf{w})\frac{\mathbb{P}(\mathbf{w})}{p'}u[d(\hat{q}),\theta'], \tag{38}$$

where

$$W_{(\theta,q^o),(\theta',\hat{q})}(q,x;q') = \{\mathbf{w}\in W_{(\theta,q^o),(\theta',\hat{q})} : (q,x,q') \text{ occurs in } \mathbf{w}\}$$

and $\varphi_{(q,x;q')}(\mathbf{w})$ is the number of repetitions of the transition $(q,x;q')$ within $\mathbf{w}$.

Now, we show that $\frac{\partial}{\partial p'}V$ is proportional to $p(q,x)V_{q'}(H)+[1-p(q,x)]V_{q'}(L)$:

$$\begin{aligned}
&\left[\sum_{\theta,\theta'} f(q,\theta)\nu^{\theta}_{\theta'}\mu^{\theta'}_x\right][p(q,x)V_{q'}(H)+[1-p(q,x)]V_{q'}(L)]\\
=&\sum_{\theta,\theta'} f(q,\theta)\nu^{\theta}_{\theta'}\mu^{\theta'}_x V_{q'}(\theta')\\
=&\sum_{\theta_0,\theta,\theta',\theta''}\mathbf{P}_0(\theta_0)\sum_{\hat{q}\in Q}\left\{\left[\sum_{\mathbf{w}_q\in W_{(\theta_0,q^o),(\theta,q)}}\mathbb{P}(\mathbf{w}_q)\right]\nu^{\theta}_{\theta'}\mu^{\theta'}_x\left[\sum_{\mathbf{w}_{q'}\in W_{(\theta',q'),(\theta'',\hat{q})}}\mathbb{P}(\mathbf{w}_{q'})\right]\right\}u[d(\hat{q}),\theta'']\\
=&\sum_{\theta_0,\theta,\theta',\theta''}\mathbf{P}_0(\theta_0)\sum_{\hat{q}\in Q}\left\{\sum_{\mathbf{w}_q\in W_{(\theta_0,q^o),(\theta,q)},\mathbf{w}_{q'}\in W_{(\theta',q'),(\theta'',\hat{q})}}\frac{\mathbb{P}[(\mathbf{w}_q,x,\mathbf{w}'_q)]}{\sigma(q,x)(q')}\right\}u[d(\hat{q}),\theta'']\\
=&\sum_{\theta_0,\theta''}\mathbf{P}_0(\theta_0)\sum_{\hat{q}\in Q}\left\{\sum_{\mathbf{w}\in W_{(\theta_0,q^o),(\theta'',\hat{q})}}\varphi_{(q,x;q')}(\mathbf{w})\frac{\mathbb{P}(\mathbf{w})}{p'}\right\}u[d(\hat{q}),\theta''] = \frac{\partial}{\partial p'}V,
\end{aligned}$$

where the last equality follows from (38) and the second last equality follows from $p' = \sigma(q,x)(q')$ and the fact that for any $\mathbf{w}_q \in W_{(\theta,q^o),(\theta,q)}$ and any $\mathbf{w}_{q'} \in W_{(\theta',q'),(\theta'',\hat{q})}$, $(\mathbf{w}_q, x; \mathbf{w}_{q'}) \in W_{(\theta,q^o),(\theta'',\hat{q})}(q,x;q')$ and that each $\mathbf{w} \in W_{(\theta,q^o),(\theta'',\hat{q})}(q,x;q')$ is counted $\varphi_{(q,x;q')}(\mathbf{w})$ times in that list. We have analogous expression for $\frac{\partial}{\partial p''}V$, and hence (36) implies that (37).

Now we prove (33). By modified multi-self consistency, for any $x \in X$ and any q_1, q_2 with $\sigma(q,x)(q_1) > 0$ and $\sigma(q,x)(q_2) > 0$ and any $q_3 \in Q$,

$$\begin{aligned}
[p(q,x)V_{q_1}(H)+[1-p(q,x)]V_{q_1}(L)] &= [p(q,x)V_{q_2}(H)+[1-p(q,x)]V_{q_2}(L)]\\
&\geq [p(q,x)V_{q_3}(H)+[1-p(q,)]V_{q_3}(L)],
\end{aligned}$$

By (9), this implies that for all $x \in X$,

$$\sum_{\theta,\theta'} f(q,\theta)\nu^{\theta}_{\theta'}\mu^{\theta'}_x V_{q_1}(\theta') = \sum_{\theta,\theta'} f(q,\theta)\nu^{\theta}_{\theta'}\mu^{\theta'}_x V_{q_2}(\theta') \geq \sum_{\theta,\theta'} f(q,\theta)\nu^{\theta}_{\theta'}\mu^{\theta'}_x V_{q_3}(\theta'). \tag{39}$$

Thus, (here we assume that the decision rule is deterministic, with no loss of generality because of (32))

$$\begin{aligned}
& p(q)V_q(H) + [1-p(q)]V_q(L) \\
= \;& p(q)\left\{\eta u[d(q),H] + (1-\eta)\left[\sum_{x\in X, q''\in Q, \theta'} \nu_{\theta'}^{H}\mu_x^{\theta'}\sigma(q,x)(q'')V_{q''}(\theta')\right]\right\} \\
+ \;& [1-p(q)]\left\{\eta u[d(q),L] + (1-\eta)\left[\sum_{x\in X, q''\in Q, \theta'} \nu_{\theta'}^{L}\mu_x^{\theta'}\sigma(q,x)(q'')V_{q''}(\theta')\right]\right\} \\
= \;& \eta\{p(q)u[d(q),H] + [1-p(q)]u[d(q),L]\} \\
+ \;& (1-\eta)\sum_{x\in X}\left\{\sum_{q''\in Q}\frac{\sum_{\theta,\theta'} f(q,\theta)\nu_{\theta'}^{\theta}\mu_x^{\theta'}V_{q''}(\theta')}{f(q,H)+f(L,q)}\sigma(q,x)(q'')\right\} \\
\geq \;& \eta\{p(q)u[d(q'),H] + [1-p(q)]u[d(q'),L]\} \\
+ \;& (1-\eta)\sum_{x\in X}\left\{\sum_{q''\in Q}\frac{\sum_{\theta,\theta'} f(q,\theta)\nu_{\theta'}^{\theta}\mu_x^{\theta'}V_{q''}(\theta')}{f(q,H)+f(L,q)}\sigma(q',x)(q'')\right\} \\
= \;& p(q)V_{q'}(H) + [1-p(q)]V_{q'}(L),
\end{aligned}$$

where the first equality follows from the recursive equation for $V_q(\theta)$ for each $\theta = H, L$, the second follows from (9), the inequality follows term by term, first the terms starting with η follow from (32), the terms starting with $(1-\eta)$ follows from (39), again term by term for each x: any term with q'' with $\sigma(q,x)(q'') > 0$ has the same value in the inequality above, and that value is no less than that for the corresponding term with $\sigma(q',x)(q'') > 0$, and the last equality follows from the recursive equation for $V_{q'}(\theta)$. $\square$

Now we are ready to prove Theorem 2.1.

(1) Let $i < j$ be given. By (33),

$$p(q_j)\Delta V_{j,i}^H + [1-p(q_j)]\Delta V_{j,i}^L \geq 0, \text{ and } p(q_i)\Delta V_{i,j}^H + [1-p(q_i)]\Delta V_{i,j}^L \geq 0. \tag{40}$$

Since there are no equivalent states, either $\Delta V_{i,j}^H > 0$ or $\Delta V_{i,j}^H < 0$. By our convention it must be $\Delta V_{j,i}^H > 0$. By the second inequality in (40), $\Delta V_{i,j}^L \geq 0$. Now, if this last inequality is an equality, then we can replace all the transition to q_i to transition to q_j and obtain a higher ex ante payoff, which is a contradiction to the optimality of the SFSA. Now, let $i < j < k$. Again, by (33), we have

$$p(q_j)\Delta V_{j,i}^H + [1-p(q_j)]\Delta V_{j,i}^L \geq 0, \text{ and } p(q_j)\Delta V_{j,k}^H + [1-p(q_j)]\Delta V_{j,k}^L \geq 0, \tag{41}$$

and hence

$$\frac{\Delta V_{i,j}^L}{\Delta V_{j,i}^H} \leq \frac{p(q_j)}{1-p(q_j)} \leq \frac{\Delta V_{j,k}^L}{\Delta V_{k,j}^H}.$$

(2) For part (a), (33) implies that

$$p(q_i)V_{q_i}(H) + [1-p(q_i)]V_{q_i}(L) \geq p(q_i)V_{q_{i+1}}(H) + [1-p(q_i)]V_{q_{i+1}}(L)$$

and hence, by rearranging terms, we have $\xi[p(q_i)] \leq \bar{\xi}_i$. A similar argument holds for $\xi[p(q_i)] \geq \bar{\xi}_{i-1}$.

For (b), let $q \in Q$ be given. By (31), $\sigma(q,x)(q_i) > 0$ only if

$$p(q,x)V_{q_i}(H) + [1-p(q,x)]V_{q_i}(L) \geq p(q,x)V_{q_j}(H) + [1-p(q,x)]V_{q_j}(L)$$

for both $j = i-1$ and $j = i+1$. This then implies (10). Conversely, it is straightforward to verify that if (10) holds, then

$$p(q,x)V_{q_i}(H) + [1-p(q,x)]V_{q_i}(L) \geq p(q,x)V_{q_j}(H) + [1-p(q,x)]V_{q_j}(L)$$

for any $j = 0, ..., K-1$, where $q_0 = q_L$ and $q_{K-1} = q_H$. Note that we need the fact that $\bar{\xi}_i$ increases with i for this, as proved in part (1). Moreover, if $\xi[p(q,x)] \in (\bar{\xi}_{i-1}, \bar{\xi}_i)$, then the above inequality is strict for any $j \neq i$ and hence $\sigma(q,x)(q_i) = 1$.

Finally, (c) follows from (32) and a similar argument.

A.2 Regime change with $\xi(p_0) < \xi^*$

Here we consider the regime change model as in Table 3 with $\xi(p_0) < \xi^*$. As mentioned in the main text, when ν is sufficiently small according to (12), one h-signal can bring the posterior across ξ^* and hence the availability heuristic M_2^a can implement the unconstrained optimum, but with $q^o = q_L$. In contrast, at the other extreme where $\nu = 1$, the DFSA that implements the unconstrained optimum is given by M_{N+2}^b as depicted in Figure 6, where N is given by (13) and where $q^o = q_{L,2}$ and $d(q_{L,i}) = a^L$ for all $i = 1, ..., N+1$ and $d(q_H) = a^H$ (see Hu (2022) for detailed arguments). Intuitively, in the fixed-world environment, $q_{L,1}$ represents the memory state in which the DM has received an ℓ-signal and hence is fully convinced of the state of the world being L and hence it is a self-absorbing memory state. In contrast, at $q_{L,i}$ for $i > 1$, the DM has

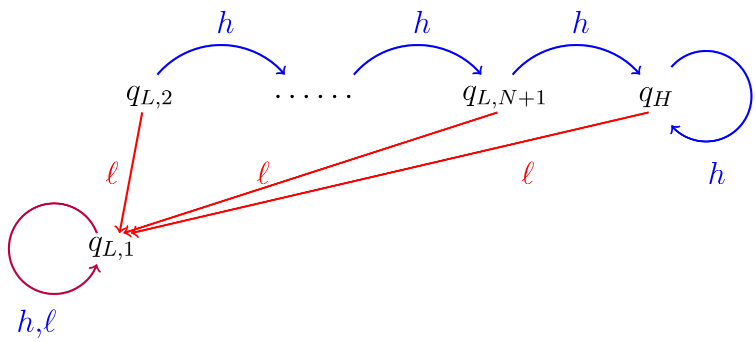

Figure 6: The DFSA, M^b_{N+2}, that implements unconstrained optimum when $p_0 < p^*$

not received any ℓ-signal but have received $i-2$ h-signals, and hence the posterior on H has gone up but has not crossed ξ^* yet, which happens only at q_H.

Now we turn to the case where ν is below one and hence even after an ℓ-signal the state of the world can still change from L to H. In this case, to implement the unconstrained optimum, it requires an unbounded number of memory states. The reason is that, after one ℓ-signal it would require a large number of h-signals to bring the posterior to reach ξ^* again for ν close to one, a number that converges to infinity as ν approaches one. However, for any given constraint K, the following theorem shows that it is optimal to ignore the possibility of regime change when ν is sufficiently high.

Theorem A.1. *Suppose that $\Delta^H = 1$ and $\Delta^L = \nu \in [0,1]$ and that $\mu^H_h = 1 > \mu = \mu^L_\ell$, and that $\xi(p_0) < \xi^*$. If $K \geq N+2$ with N given by (13), then there exists $\tilde{\nu} < 1$ such that for all $\nu \geq \tilde{\nu}$, the optimal SFSA is M^b_{N+2} with $q^o = q_{L,2}$.*

Proof. The proof follows the same steps as in Theorem 4.1. As there, for any given $K \geq N+2$, We know that M^b_{N+2} is the unique optimal SFSA when $\nu = 1$. Thus, for any such K, $V(M^b_{N+2}) > V(M')$ for other SFSA $M \in \mathcal{M}_K$ whose randomization probabilities are ϵ away from that in M^b_{N+2} and this inequality holds for a range of ν's below 1. Now we demonstrate local optimality of M^b_{N+2} for a range of ν's below 1 by appealing to Theorem 2.1. First we compute the value functions under M^b_{N+2}:

$$\begin{aligned}
V_{q_H}(H) &= u^H,\ V_{q_H}(L) = 0,\\
V_{q_{L,1}}(H) &= 0,\ V_{q_{L,i}}(H) = (1-\eta)^{N+2-i}u^H,\ i = 2, ..., N+1\\
V_{q_{L,i}}(L) &= -\frac{\eta\{[1-\nu(1-\mu)]u^L + (1-\nu)u^H\}}{[1-\nu(1-\mu)][1-(1-\eta)\nu(1-\mu)]}[(1-\eta)\nu(1-\mu)]^{N+2-i}\\
&+ \frac{\eta}{1-\nu(1-\eta)}u^L + \frac{(1-\nu)(1-\eta)^{N-i+1}}{1-\nu(1-\mu)}u^H,\ i = 2, ..., N+1,\\
V_{q_{L,1}}(L) &= \frac{\eta}{1-(1-\eta)\nu}u^L.
\end{aligned}$$

Thus,

$$\begin{aligned}
\bar{\xi}_1 &= \frac{V_{q_{L,1}}(L) - V_{q_{L,2}}(L)}{V_{q_{L,2}}(H) - V_{q_{L,1}}(H)}\\
&= \left(\xi^* + \frac{(1-\nu)}{1-\nu(1-\mu)}\right)\frac{\eta[\nu(1-\mu)]^N}{1-(1-\eta)\nu(1-\mu)} - \frac{(1-\nu)}{1-\nu(1-\mu)},\\
\bar{\xi}_i &= \frac{V_{q_{L,i}}(L) - V_{q_{L,i+1}}(L)}{V_{q_{L,i+1}}(H) - V_{q_{L,i}}(H)}\\
&= \left(\xi^* + \frac{(1-\nu)}{1-\nu(1-\mu)}\right)[\nu(1-\mu)]^{N-i+1} - \frac{1-\nu}{1-\nu(1-\mu)},\ i = 2, ..., N,\\
\bar{\xi}_{N+1} &= \frac{V_{q_{L,N+1}}(L) - V_{q_H}(L)}{V_{q_H}(H) - V_{q_{L,N+1}}(H)} = \xi^*.
\end{aligned}$$

Now, the corresponding beliefs are given by

$$\begin{aligned}
\xi(q_{L,1}) &= \frac{(1-\eta)(1-\nu)}{\eta},\\
\xi(q_{L,2}) &= \xi(p_0),\\
\xi(q_{L,i}) &= \frac{1-\nu+\nu\mu p_0}{(1-p_0)[1-\nu(1-\mu)][\nu(1-\mu)]^{i-2}} - \frac{1-\nu}{1-\nu(1-\mu)},\ i = 2, ..., N+1,\\
\xi(q_H) &= \frac{[1-\nu+\nu\mu p_0][1-(1-\eta)\nu(1-\mu)]}{\eta(1-p_0)[1-\nu(1-\mu)][\nu(1-\mu)]^N} - \frac{1-\nu}{1-\nu(1-\mu)}.
\end{aligned}$$

Appealing to Theorem 2.1, local optimality requires

$$\xi(q_H) \geq \xi^*,\ \xi(q_{L,i}, h) \in [\bar{\xi}_i, \bar{\xi}_{i+1}] \text{ for } i = 2, ..., N, \text{ and } \xi(q_{L,1}, h) < \bar{\xi}_1. \tag{42}$$

Now, consider a SFSA M that is a replica of M^b_{N+2} with $K' \leq K$ memory states. We show that M^b_{N+2} is still locally optimal against small deviations from such replica, by showing the corresponding condition for (42). Note that the continuation value does not change with the replica memory states and hence the thresholds $\bar{\xi}_i$'s remain the same. Regarding beliefs, since for the replica state of each $q_{L,i}$ with $i \geq 2$, it can be passed through for at most once, a simple induction argument shows that its belief coincides with the corresponding memory state in M^b_{N+2}. Moreover,

$$\xi(q_{L,i}, h) = \frac{1-\nu+\nu\mu p_0}{(1-p_0)[1-\nu(1-\mu)][\nu(1-\mu)]^{i-1}} - \frac{1-\nu}{1-\nu(1-\mu)} \in [\bar{\xi}_i, \bar{\xi}_{i+1}),$$

which is equivalent to

$$\begin{aligned} &\left(\xi^* + \frac{(1-\nu)}{1-\nu(1-\mu)}\right)[\nu(1-\mu)]^N \\ &\leq \xi(p_0) + \frac{(1-\nu)}{1-\nu(1-\mu)} \leq \left(\xi^* + \frac{(1-\nu)}{1-\nu(1-\mu)}\right)[\nu(1-\mu)]^{N-1}. \end{aligned} \tag{43}$$

When $\nu = 1$, the first inequality is weak and the second is strict, which follow from the definition of N given by (13). For $\nu < 1$, the first becomes strict, and the second is preserved for ν not too small.

Now we consider the replica states of q_H and $q_{L,1}$. Let equivalent states of q_H be denoted by $q_H^1, ..., q_H^I$, and $q_{L,1}$ be denoted by $q_{L,1}^1, ..., q_{L,1}^J$, with

$$\begin{aligned} &\sigma(q_H^i, h)(q_H^j) = \alpha_{hij},\ \sigma(q_H^i, \ell)(q_{L,1}^j) = \alpha_{\ell ij}, \\ &\sigma(q_{L,1}^i, h)(q_{L,1}^j) = \beta_{hij},\ \sigma(q_{L,1}^i, \ell)(q_{L,1}^j) = \beta_{\ell ij}, \\ &\sigma(q_{L,N+1}, h)(q_H^j) = \gamma_{hj},\ \sigma(q_{L,n}, \ell)(q_{L,1}^j) = \gamma_{\ell nj}. \end{aligned} \tag{44}$$

Then, we have the following recursive equations:

$$\begin{aligned} f(q_H^i, H) =& (1-\eta)\left\{[f(q_{L,N+1}, H) + f(q_{L,N+1}, L)(1-\nu)]\gamma_{hi}\right\} \\ &+(1-\eta)\left\{\sum_{j=1}^{I}[f(q_H^j, L)(1-\nu) + f(q_H^j, H)]\alpha_{hji}\right\}, \\ f(q_H^i, L) =& (1-\eta)\left\{\sum_{j=1}^{I} f(q_H^j, L)\nu(1-\mu)\alpha_{hji} + f(q_{L,N+1}, L)\nu(1-\mu)\gamma_{hi}\right\}, \\ f(q_{L,1}^i, H) =& (1-\eta)\left\{\sum_{j=1}^{J}[f(q_{L,1}^j, H) + f(q_{L,1}^j, L)(1-\nu)]\beta_{hji}\right\}, \\ f(q_{L,1}^i, L) =& (1-\eta)\left\{\sum_{j=1}^{J} f(q_{L,1}^j, L)\nu[(1-\mu)\beta_{hji} + \mu\beta_{\ell ji}] + \sum_{k=1}^{I} f(q_H^k, L)\nu\mu\alpha_{\ell ki}\right\} \\ &+(1-\eta)\left\{\sum_{n=2}^{N+1} f(q_{L,n}, L)\nu\mu\gamma_{\ell ni}\right\}, \end{aligned} \tag{45}$$

Now we show that $\xi(q_H^i) \geq \xi^*$ for all i. Using the same methodology as in the proof of Theorem 3.1 (2), i.e., we take the equations in (45) as simultaneous equations and use the contractions mapping theorem, it suffices to show that

$$\frac{f(q_{L,N+1}, H) + f(q_{L,N+1}, L)(1-\nu)]}{f(q_{L,N+1}, L)\nu(1-\mu)} \geq \xi^*,$$

which follows from the earlier result that $\xi(q_{L,N+1}, h) \geq \xi^*$ for ν close to one.

Finally, the result that $\xi(q_{L,1}^i, h) < \bar{\xi}_1$ follows a similar argument to that in the proof of Theorem 4.1 (1) and is omitted. □

Theorem A.1 then extends Theorem 4.1 (1) to the case where $\xi(p_0) < \xi^*$. As Theorem 4.1 (1), this is a less-is-more result, as the optimal SFSA is M_{N+2}^b for any given $K \geq N+2$ for a range of ν's. Moreover, in that case, the DM behaves as if she ignores the possibility of regime change under the constrained optimal rule and is stuck to action a^L after receiving an ℓ-signal, while an unconstrained DM will continue to update her belief and will eventually be fully convinced of state of the world H. However, different from Theorem 4.1 (1), the result in Theorem A.1 requires $K \geq N+2$, with $N \geq 1$ given by (13).

Now we turn to the case where $K < N+2$. In the fixed-worlds environment, Hu (2022) has shown that randomization is optimal, with the optimal SFSA taking the same form with $Q = \{q_{L,1}, q_{L,2}, ..., q_{L,K-1}, q_H\}$, and that optimal randomization occurs at every $q_{L,i}$ for $i = 2, ..., K-1$ with $\sigma(q_i, h)(q_i) = \alpha = 1 - \sigma(q_i, h)(q_{i+1}) \in (0, 1)$ with $q_K = q_H$, and we denote this SFSA by $M_K^b(\alpha)$; see Figure 7 for a graphical representation of $M_K^b(\alpha)$ with $K = 5$. This optimal SFSA features randomization in all intermediate memory states, a feature in contrast to the characterization in Wilson (2014) where randomization occurs at the extreme memory states.

The following result extends this characterization result of randomization to changing worlds.

Theorem A.2. *Suppose that $\Delta^H = 1$ and $\Delta^L = \nu \in [0, 1]$ and that $\mu_h^H = 1 > \mu = \mu_\ell^L$, and that $\xi(p_0) < \xi^*$. If $3 \leq K < N+2$ with N given by (13), there exists $\tilde{\nu} < 1$ such that for all $\nu \geq \tilde{\nu}$, the optimal SFSA takes the form $M_K^b(\alpha)$.*

Proof. By Hu (2022), when $\nu = 1$ the optimal SFSA takes the form $M_K^b(\alpha_1, ..., \alpha_{K-2})$ with $\alpha_1 = \alpha_2 = ...\alpha_{K-2} = \alpha > 0$. Following the same arguments as in the proof of

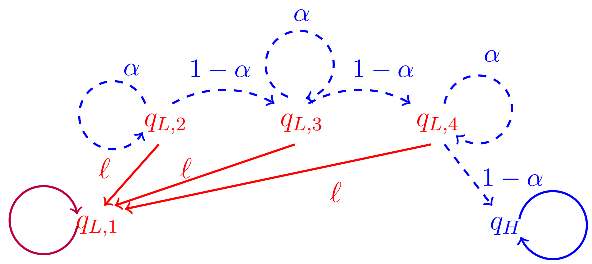

Figure 7: $M_5^b(\alpha)$

Theorem 4.1, for a range of ν's below $\nu = 1$ we only need to consider local optimality. Now we show that for any ν, the locally optimal SFSA of the form $M_K^b(\alpha_1, ..., \alpha_{K-2})$ has the form $\alpha_1 = \alpha_2 = ...\alpha_{K-2} = \alpha > 0$.

First we compute the value functions:

$$\begin{aligned}
V_{q_H}(H) &= u^H,\ V_{q_L}(L) = \frac{\eta}{1-(1-\eta)\nu}u^L,\ V_{q_L}(H) = 0, \\
V_{q_H}(L) &= (1-\eta)\left\{\nu[\mu V_{q_L}(L) + (1-\mu)V_{q_H}(L)] + (1-\nu)V_{q_H}(H)\right\}, \\
&= \frac{1}{1-\nu(1-\eta)(1-\mu)}\left\{\frac{\mu(1-\eta)\nu\eta}{1-(1-\eta)\nu}u^L + (1-\eta)(1-\nu)u^H\right\}, \\
V_{q_i}(H) &= (1-\eta)\{\alpha_i V_{q_i}(H) + (1-\alpha_i)V_{q_{i+1}}(H)\}, i = 1, ..., K-2, \\
V_{q_i}(H) &= \left[\prod_{j=i}^{K-2}\frac{1-\alpha_j}{1-(1-\eta)\alpha_j}\right](1-\eta)^{K-1-i}u^H, \\
V_{q_i}(L) &= \eta u^L + (1-\eta)\{(1-\nu)[\alpha_i V_{q_i}(H) + (1-\alpha_i)V_{q_{i+1}}(H)] \\
&+ \nu[\mu V_{q_L}(L) + (1-\mu)[\alpha_i V_{q_i}(L) + (1-\alpha_i)V_{q_{i+1}}(L)]]\}, i = 1, ..., K-2,
\end{aligned}$$

We get the solution for $V_{q_i}(L)$:

$$\begin{aligned}
V_{q_i}(L) &= \frac{\eta u^L}{1-\nu(1-\eta)} + \frac{1-\nu}{1-\nu(1-\mu)}\left[\prod_{j=i}^{K-2}\frac{1-\alpha_j}{1-(1-\eta)\alpha_j}\right](1-\eta)^{K-1-i}u^H \\
&+ \tilde{C}V_{q_H}(L)[\nu(1-\mu)(1-\eta)]^{K-1-i}\left[\prod_{j=i}^{K-2}\frac{1-\alpha_j}{1-\nu(1-\mu)(1-\eta)\alpha_j}\right],
\end{aligned}$$

where the constant $\tilde{C}$ can be found by meeting the "initial" condition, i.e. by equating the solution $V_{q_i}(L)$ in case $i = K-2$ to the expression $V_{q_{K-2}}(L)$ which can be found explicitly from the formula for $V_{q_i}(L)$, and hence

$$C \equiv \tilde{C}V_{q_H}(L) \quad = \quad \frac{-\eta}{1-\nu(1-\mu)(1-\eta)}\left[u^H + \frac{(1-\nu)u^L}{1-\nu(1-\mu)}\right].$$

Note that the ex ante payoff $p_0V_{q_1}(H)+(1-p_0)V_{q_1}(L)$ is symmetric in $(\alpha_1,...,\alpha_{K-2})$ and supermodular, and hence it is optimal to set $\alpha=\alpha_i$ for all i. By doing so, the ex ante payoff is

$$\begin{aligned} F(\alpha) &= p_0\left[\frac{1-\alpha}{1-(1-\eta)\alpha}\right]^{K-2}(1-\eta)^{K-2}u^H + (1-p_0)\frac{\eta u^L}{1-\nu(1-\eta)} \\ &+ (1-p_0)\frac{1-\nu}{1-\nu(1-\mu)}\left[\frac{1-\alpha}{1-(1-\eta)\alpha}\right]^{K-2}(1-\eta)^{K-2}u^H \\ &+ (1-p_0)C[\nu(1-\mu)(1-\eta)]^{K-2}\left[\frac{1-\alpha}{1-\nu(1-\mu)(1-\eta)\alpha}\right]^{K-2} \end{aligned}$$

Now we show that $F'(\alpha)|_{\alpha=0} > 0 \Leftrightarrow \xi(p_0, h^{K-2}) < \xi^*$, that is if and only if $K-2 < N$ with N given by (13). Now,

$$\begin{aligned} F'(\alpha) &= [p_0 + (1-p_0)\frac{1-\nu}{1-\nu(1-\mu)}][K-2]\left[\frac{1-\alpha}{1-(1-\eta)\alpha}\right]^{K-3}\frac{-\eta}{[1-(1-\eta)\alpha]^2}(1-\eta)^{K-2}u^H \\ &+ [1-p_0]C[\nu(1-\mu)(1-\eta)]^{K-2}[K-2]\left[\frac{1-\alpha}{1-\nu(1-\mu)(1-\eta)\alpha}\right]^{K-3}\frac{\nu(1-\mu)(1-\eta)-1}{[1-\nu(1-\mu)(1-\eta)\alpha]^2} \end{aligned}$$

$$\frac{F'(\alpha)|_{\alpha=0}}{\eta(1-\eta)^{K-2}u^H(K-2)(1-p_0)} = -[\xi(p_0)+\frac{1-\nu}{1-\nu(1-\mu)}] + [\nu(1-\mu)]^{K-2}[\xi^* + \frac{1-\nu}{1-\nu(1-\mu)}]$$

Therefore,

$$\begin{aligned} & F'(\alpha)|_{\alpha=0} > 0 \\ \Longleftrightarrow\ & \frac{F'(\alpha)|_{\alpha=0}}{\eta(1-\eta)^{K-2}u^H(K-2)(1-p_0)} > 0 \\ \Longleftrightarrow\ & -[\xi(p_0)+\frac{1-\nu}{1-\nu(1-\mu)}] + [\nu(1-\mu)]^{K-2}[\xi^* + \frac{1-\nu}{1-\nu(1-\mu)}] > 0 \\ \Longleftrightarrow\ & \xi(p_0)+\frac{1-\nu}{1-\nu(1-\mu)} < [\nu(1-\mu)]^{K-2}[\xi^* + \frac{1-\nu}{1-\nu(1-\mu)}] \\ \Longleftrightarrow\ & \frac{\xi(p_0)+\frac{1-\nu}{1-\nu(1-\mu)}}{[\nu(1-\mu)]^{K-2}} - \frac{1-\nu}{1-\nu(1-\mu)} < \xi^* \\ \Longleftrightarrow\ & \xi(p_0, h^{K-2}) < \xi^*. \end{aligned}$$

The LHS of the preultimate inequality is indeed $\xi(p_0, h^{K-2})$, i.e the result of applying Bayes' rule to $\xi(p_0)$ $K-2$ times in the breakthrough environment for arbitrary ν. $\square$

Theorem A.2 extends the ignoring-regime-change heuristic identified in Theorem A.1 to include randomization when the constraint K is lower than $N+2$. Different from Theorem A.1, however, in this case more memory states can increase the payoff up to $K = N+2$, and hence less-is-more does not hold for $K < N+2$. As a result, there would be three regimes if we would introduce a convex cost function for the memory states. When the cost is exactly zero, then the optimal K would be unbounded. When the cost is sufficiently small, then the optimal $K = N+2$ and the optimal SFSA is deterministic. That is, we can exclude randomization by endogenously determining the memory constraint. When the cost is higher, optimal K can be below $N+2$ and randomization is optimal.

Finally, we remark that all the results in this section would hold if we relax the assumption that $\Delta^H = 1$ but it is close to one, and the assumption that $\mu_h^H = 1$ is not knife-edge either. If we consider $\Delta^H < 1$ but close, the only difference is that when p_0 is at the boundary according to (13) such that $K = N+2$ exactly holds then we need to discuss the optimal SFSA, which may be M_{N+k}^b for $k = 1, 2, 3$, depending on the relative values of Δ^H and Δ^L. Moreover, when $K < N+2$, randomization may take more complicated forms as well. However, the baseline result that the decision-maker behaves as if she ignores the underlying regime change will remain, that is, she is stuck to action a^L once an ℓ-signal is received. Similar situation holds for $\mu_h^H < 1$ but close.