

Selective-Memory Equilibrium

Drew Fudenberg

Massachusetts Institute of Technology

Giacomo Lanzani

Harvard University

Philipp Strack

Yale University

We study agents who are more likely to remember some experiences than others but update beliefs as if the experiences they remember are the only ones that occurred. To understand the long-run effects of selective memory, we propose selective-memory equilibrium. We show that if the agent's behavior converges, their limit strategy is a selective-memory equilibrium, and we provide a sufficient condition for behavior to converge. We use this equilibrium concept to explore the consequences of several well-documented biases. We also show that there is a close connection between selective-memory equilibria and the outcomes of mis-specified learning.

I. Introduction

We provide a new conceptual framework for the study of agents who have selective memory in that they are more likely to recall some events than

We thank Chiara Aina, Ian Ball, Abhijit Banerjee, Douglas Bernheim, Pierpaolo Battigalli, Roland Bénabou, Dirk Bergemann, Benjamin Brooks, Jonathan Brownrigg, Simone Cerreia-Vioglio, Gary Charness, Olivier Compte, Roberto Corrao, Martijn de Vries, Glenn Ellison, Sebastian Ebert, Ignacio Esponda, Erik Eyster, Thomas Graeber, Jerry Green, Philippe Jehiel,

Electronically published November 13, 2024

Journal of Political Economy, volume 132, number 12, December 2024.

© 2024 The University of Chicago. All rights reserved. Published by The University of Chicago Press.

<https://doi.org/10.1086/731412>

others. We assume that selective memory is stochastic and exogenous and allow the agent's actions to influence what they observe.¹ In most of the paper, we also assume that agents are unaware of their selective memory, so they update their beliefs as if the experiences they remember are the only ones that occurred.² These assumptions fit evidence from both experimental and real-world settings. Although our work is inspired by the neuroscience and psychology literature on memory, we do not try to develop a model that fully matches the memory formation and retrieval process. Instead, we develop and motivate a solution concept that makes it easy to obtain predictions about long-run actions and beliefs for any given memory distortion.

Our focus is on selective memory's long-run implications. We show that if an agent's behavior converges, their beliefs concentrate on the *memory-weighted likelihood maximizers*, that is, distributions that maximize the likelihood of a distorted version of the true outcome distribution that gives more weight to realizations that are more likely to be remembered. We also provide conditions on the agent's payoff function and the support of their prior that imply that their behavior does converge. Whether or not these conditions are satisfied, when behavior converges, it converges to a *selective-memory equilibrium*, which is a strategy that myopically maximizes their expected payoff against a probability distribution over these maximizers. If all experiences are recalled with the same probability, then memory limitations have no long-run effect. However, if memory is selective and agents are more likely to remember some experiences than others, selective memory can have a persistent effect. For example, an agent who is more likely to recall when they performed well in a task than when they performed poorly will underestimate the task's difficulty and do it too often.

Our framework lets us analyze the long-run consequences of important and widely documented forms of selective memory such as pleasant-memory bias (Mischel, Ebbesen, and Zeiss 1976; Adler and Pansky 2020; Chew, Huang, and Zhao 2020) and the related ego-boosting bias (Zimmermann

Mats Koster, Roger Lagunoff, Shengwu Li, Stephen Morris, Ryan Oprea, Parag Pathak, Jérôme Renault, Frank Schilbach, Giorgio Saponaro, Josh Schwartzstein, Andrei Shleifer, Rani Spiegler, Tomasz Strzalecki, Dmitry Taubinsky, Maximilian Voigt, Alex Wolitzky, Muhamet Yildiz, and Sevgi Yuksel for helpful comments and National Science Foundation grant SES 1951056 for financial support. We thank Neil He for excellent research assistance. This paper was edited by Emir Kamenica.

¹ Memory has been informally described as stochastic since the early stages of the psychology literature, as in James (1890), and recent evidence in neuroscience (e.g., Shadlen and Shohamy 2016) and economics (e.g., Sial, Sydnor, and Taubinsky 2023) supports this interpretation. Schacter (2008) and Kahana (2012) discuss evidence that some experiences are recalled more often than others.

² Reder (2014), Zimmermann (2020), and Gödker, Jiao, and Smeets (forthcoming) provide evidence of partial or complete unawareness of memory biases. The main results extend as stated to the case of an agent who is aware that they sometimes forget but is not aware that their memory is selective, and who does not draw inference from their past actions.

2020), cognitive dissonance (Elkin and Leippe 1986; Chammat et al. 2017; Gödker, Jiao, and Smeets, forthcoming), associativeness (Thomson and Tulving 1970; Tulving and Schacter 1990; Enke, Schwerter, and Zimmermann 2022; Goetzmann, Watanabe, and Watanabe 2022), confirmatory bias (Hastie and Park 1986), and the relative memorability of extreme outcomes (Cruciani et al. 2011). In contrast, earlier papers on selective memory each studied a specific form of memory bias, and most considered only short-run effects.

Under positive-memory bias, the agent is more likely to recall experiences that induce a larger utility. For example, Zimmermann (2020) finds that subjects who received poor scores on an IQ test are more likely to state that they “cannot recall” their test results, even though that answer is pay-off dominated in the experiment, and there were only three things for subjects to recall. Gödker, Jiao, and Smeets (forthcoming) finds that investors are more likely to remember positive returns of stocks they invested in and that their selective memory distorts both their beliefs and their future investment decisions in the direction our model predicts.

We show that positive-memory bias can endogenously generate the same long-run behavior as dogmatic overconfidence in a fixed learning environment. However, we argue that the overconfidence that arises from selective memory is more susceptible to external manipulation through changes in the feedback provided to the agent. For example, coupling negative feedback on one dimension with positive feedback on another will make the negative feedback be recalled more often, which leads to less bias in long-run beliefs.³

Agents with *associative memory* are more likely to recall situations similar to the current decision problem, for example, when they had a similar mood. In general, this can lead the agent to underweight data relative to its true informativeness. However, the simplest version of associativeness, similarity weighting (Kahana 2012), does not alter the possible long-run outcomes for a correctly specified agent because they learn the true consequences of their on-path action.⁴

We also study *extreme-experience bias*, which makes experiences with more extreme payoffs more memorable. We show that moderate risk aversion paired with this bias may explain the extreme risk aversion revealed by the prices of safe and risky assets in financial markets. Moreover, if rarer experiences are more easily recalled, the agent overweights small probabilities as assumed in prospect theory.

³ This is suggested in the management literature by, e.g., Procházka, Ovcari, and Durinik (2020).

⁴ Thus with similarity weighting, all selective-memory equilibria are unitary self-confirming equilibria (Battigalli 1987; Fudenberg and Levine 1993a).

Selective-memory equilibrium resembles Berk-Nash equilibrium (Esponda and Pouzo 2016), which applies to agents with perfect memory but a misspecified prior, as both require that the agent's action is a best reply to a distorted version of the true outcome distributions. Indeed, we show that every uniformly strict Berk-Nash equilibrium (Fudenberg, Lanzani, and Strack 2021) that is not supported by beliefs that assign strictly positive probability to an impossible outcome is equivalent to a uniformly strict selective-memory equilibrium for some memory function and a full-support prior. Moreover, every uniformly strict selective-memory equilibrium is equivalent to a uniformly strict Berk-Nash equilibrium with the appropriate prior support. However, this equivalence fails for equilibria that are not uniformly strict.⁵ In addition, unlike Berk-Nash equilibria, selective-memory equilibria generally do not reduce to self-confirming equilibria when the agent is correctly specified because the agent need not learn the consequence of the equilibrium action. Importantly, the form of misspecification that would lead to the same behavior as a given form of selective memory depends on the environment. That is, particular forms of misspecification and selective memory that coincide under one information structure could lead to very different comparative statics with respect to changes in what the agent observes. To illustrate this, we show that combining positive and negative feedback has qualitatively different effects on agents with ego-boosting memory than on dogmatically overconfident agents.

Related theoretical work.—Mullainathan (2002) studies selective memory where the probability of recalling an observation is the sum of a base rate, an “associativeness” term that measures the experience’s similarity to the current observation, and a “rehearsal” term that indicates whether the experience was recalled in the previous period. As we do, the paper assumes that agents are naïve about their selective memory. It also assumes that signals are Gaussian and are not influenced by the agent’s actions. Afrouzi et al. (2020) also studies an agent forecasting the next realization of a first-degree autoregressive (AR[1]) process. It assumes that the agent knows the data-generating process and chooses which experiences to recall at a cost. Bordalo et al. (2021) shows how memory depends on the relative frequency of various characteristics and can be manipulated by making some observations stand out more. None of these papers addresses our question of determining the agent’s long-run beliefs and actions.

There is also a set of papers that study long-run behavior with selective attention or recording, where whether an experience is recorded determines whether it will be recalled in every future period, as in the model in appendix B.4 (app. B is available online). Compte and Postlewaite

⁵ A selective-memory equilibrium is *uniformly strict* if it is the unique best reply to all the beliefs supported on the memory-weighted likelihood maximizers.

(2004) considers a myopic agent with the choice between a safe action with a known payoff and a risky action whose outcome distribution is unknown. It assumes that taking the risky action is sometimes a dominant strategy so that the agent will eventually take it infinitely often and that periods with good performance are more likely to be recorded. This leads to overconfidence, as in our section IV.A example. Schwartzstein (2014) studies the long-run beliefs of an agent whose attention is based on perceived informational value. The agent recalls all of their observations but naively does not realize that they did not pay attention to some relevant aspects of what they observed. As with selective memory and misspecified beliefs, this can lead the agent to make systematically biased forecasts. Relatedly, Schweizer and de Vries (2022) assumes that for exogenous reasons, the agent weights outcomes differently depending on how extreme they were (compared to other outcomes) at the time they realized. This can lead to probability distortions analogous to those of cumulative prospect theory or selective memory with rare-experience bias (see sec. IV.C).

Wilson (2014) and Jehiel and Steiner (2020) study the optimal use of a finite memory by an agent who receives a stream of exogenous signals until they stop and take a single action. Battigalli and Generoso (2021) proposes a formalism to separate assumptions on the players' objective information and memory in games. Bénabou and Tirole (2002) considers a 2-period model where a time-inconsistent agent receives either a null signal or a bad signal in the first period, and at a cost can change the probability that the second-period self recalls the bad signal. The resulting game need not have a unique equilibrium, but in some cases, it can lead to overconfidence. Jehiel (2021) proposes a multiself-solution concept to model "forgetful liars." Further afield, Malmendier and Nagel (2016), Malmendier, Pouzo, and Vanasco (2020), and Malmendier and Shen (2023) consider models where agents weight events based on their age when the events happened, and Nagel and Xu (2022) analyzes an asset-pricing model where the representative agent has fading memory.

II. Setup

We study a sequence of choices made by a single agent. In every period $t \in \mathbb{N}$, the agent observes a signal s from the finite set S and then chooses an action a from the finite set A . The realized signal s and the chosen action a induce an objective probability distribution $p_{s,a}^* \in \Delta(Y)$ over the finite set of possible outcomes Y .⁶ A *pure strategy* is a map $\sigma: S \rightarrow A$, and the agent's flow payoff is given by the utility function $u: S \times A \times Y \rightarrow \mathbb{R}$.

⁶ We denote objective distributions with a superscript asterisk.

We assume that the agent knows the fixed and i.i.d. (independent and identically distributed) full-support distribution $\zeta \in \Delta(S)$ over signals.⁷ They also know that the map from actions and signals to probability distributions over outcomes is fixed and depends only on their current action and the realized signal. However, they are uncertain about the outcome distributions each signal-action pair induces. To model this uncertainty, we suppose that the agent has a prior μ over data-generating processes $p \in \Delta(Y)^{S \times A}$, where $p_{s,a}(y)$ denotes the probability of outcome $y \in Y$ when signal s is observed and action a is played. The support of μ is Θ ; its elements are the p 's the agent initially thinks are possible. The prior is *correctly specified* if its support contains the true data-generating process $p^* \in \Theta$; if not, the prior is *misspecified*.

To simplify the exposition, we will assume throughout the paper that selective-memory agents are correctly specified, but this is not essential; all results except for proposition 1 are true as stated under the weaker assumption that Bayesian updating is well-defined at every history that is reachable with positive probability. We sometimes consider a prior with *full support*, by which we mean that every possible data-generating process is in the support of the agent's prior, that is, $\Theta = \Delta(Y)^{S \times A}$.

ASSUMPTION 1 (Maintained assumption). The agent is correctly specified.

Objective histories and recalled histories.—We assume that the agent always recalls the signal they just observed. The agent's memory of the outcomes corresponding to past signal-action pairs is distorted by a collection of signal-dependent *memory functions* $m_s : S \times A \times Y \rightarrow [0, 1]$, where $m_s(s, a, y)$ specifies the probability with which the agent remembers a past realization of the signal-action-outcome triplet (s, a, y) when they observe signal s' . We call these triplets *experiences*.

Let $H_t = (S \times A \times Y)^t$ denote the set of all histories of length t , and $H = \cup_t H_t$ the set of all histories. After objective history $h_t = (s_\tau, a_\tau, y_\tau)_{\tau=1}^t$ and signal s_{t+1} , the *recalled periods* R_t are a random subset of $\{1, \dots, t\}$. Period τ is remembered with probability $m_{s_{t+1}}(s_\tau, a_\tau, y_\tau)$, independently of which other periods are remembered, so the probability that $R_t \subseteq \{1, \dots, t\}$ is remembered, given $h_t = (s_\tau, a_\tau, y_\tau)_{\tau=1}^t$ and s_{t+1} , is

$$\mathbb{P}[R_t | (s_\tau, a_\tau, y_\tau)_{\tau=1}^t, s_{t+1}] = \prod_{\tau \in R_t} m_{s_{t+1}}(s_\tau, a_\tau, y_\tau) \prod_{\tau \in \{1, \dots, t\} \setminus R_t} (1 - m_{s_{t+1}}(s_\tau, a_\tau, y_\tau)).$$

For every objective history h_t and set of recalled periods R_t , the *recalled history* $h_t(R_t) \in H_{|R_t|}$ is the subsequence of recalled experiences listed in the order they realized.⁸

⁷ This assumption lets us focus on our key points and can be substantially relaxed.

⁸ Appendix A1 gives a formal definition of recalled histories.

Beliefs.—We assume that the agent recomputes their beliefs each period based on all of their recollections, as opposed to simply updating their period- t beliefs on the basis of their period- t observation, and that the agent is unaware of their selective memory and naïvely updates their beliefs as if the experiences they remember are the only ones that occurred,⁹ so that the posterior probability of every (measurable) $C \subseteq \Theta$ after recalled history $h_t = (s_\tau, a_\tau, y_\tau)_{\tau=1}^t$ is

$$\mu(C|h_t) = \frac{\int_C \prod_{\tau=1}^t p_{s_\tau, a_\tau}(y_\tau) d\mu(p)}{\int_\Theta \prod_{\tau=1}^t p_{s_\tau, a_\tau}(y_\tau) d\mu(p)}. \quad (1)$$

In appendix A3, we show that if agents recognize that their memory is faulty but believe that it is not selective and do not make inferences about unrecalled observations from their recalled past actions, the main results extend as stated.¹⁰

Best responses and optimal policies.—Denote by $\text{BR}(s, \nu)$ the actions that maximize expected utility when signal s is observed and the agent's belief is $\nu \in \Delta(\Theta)$:¹¹

$$\text{BR}(s, \nu) = \operatorname{argmax}_{a \in A} \int_{\Theta} \sum_{y \in Y} u(s, a, y) p_{s, a}(y) d\nu(p).$$

A *policy* $\pi: H \rightarrow A^S$ specifies a pure strategy for every recalled history. We assume that the agent is myopic and uses an *optimal policy*, that is, a map $\pi: H \rightarrow A^S$ such that for every $s \in S$ and recalled history $h_t \in H$, $\pi(h_t)(s) \in \text{BR}(s, \mu(\cdot | h_t))$.¹²

Examples.—We illustrate our model with five commonly studied examples of memory bias. In this subsection, assumptions about the memory function m hold for all $s, s' \in S$, $y, y' \in Y$, and $a \in A$.

⁹ See, e.g., d'Acremont, Schultz, and Bossaerts (2013) for fMRI (functional magnetic resonance imaging) evidence that agents access their accumulated evidence each period when updating beliefs, and Reder (2014) for evidence that agents are often naïve about their selective memory and do not make inferences about their forgotten observations from the actions they remember taking.

¹⁰ Appendix A3 maintains our assumption that the agent remembers an experience either perfectly or not at all. We relax this in app. B.3, where the agent may remember some but not all aspects of a past experience, such as one or two components of a multidimensional outcome. That model assumes that the agent is not fully naïve, because remembering that some experience occurred but not all of its details might lead the agent to question their ability to perfectly recollect the past.

¹¹ For every $X \subseteq \mathbb{R}^k$, $\Delta(X)$ denotes the set of Borel probability distributions on X endowed with the topology of weak convergence.

¹² Note that we restrict attention to deterministic optimal policies and do not allow the agent to randomize over pure strategies.

EXAMPLE 1 (Utility-dependent memory). In some cases, the probability of remembering an experience depends on its associated utility, so that $m_s(s, a, y) = \Phi(u(s, a, y))$ for some $\Phi: \mathbb{R} \rightarrow [0, 1]$. Agents who are more likely to remember pleasant experiences correspond to monotone increasing Φ ; agents who are more likely to remember extremely high- or low-utility realizations have Φ that is single dipped.¹³

EXAMPLE 2 (Positive-memory bias). *Positive-memory bias* is the tendency to overremember experiences that reflect positively on oneself, such as a high test score (see Mischel, Ebbesen, and Zeiss 1976 for early experimental evidence of positive-memory bias and Adler and Pansky 2020 for a survey). To model this, we let one dimension $y_1 \in \mathbb{R}$ of the outcome y reflect the self-image consequences of the experience, and specify that $m_s(s, a, y) = \Phi(y_1)$ for some increasing $\Phi: \mathbb{R} \rightarrow [0, 1]$.

EXAMPLE 3 (Cognitive dissonance and ex post regret). *Cognitive dissonance* is a memory bias where the probability of recalling an experience depends on how well the chosen alternative performed compared to the counterfactual payoff the agent would have received under the ex post optimal choice (Elkin and Leippe 1986). This corresponds to $m_s(s, a, y) = \Phi(\max_{a' \in A} u(s, a', y) - u(s, a, y))$ where $\Phi: \mathbb{R}_+ \rightarrow [0, 1]$ is decreasing. If the outcome includes the payoff that would have been obtained with each action, the probability of remembering an outcome is decreasing in what Loomes and Sugden (1982) called “regret” (see Lanzani 2022 for the version without a state space that formally corresponds to the case we have here).

EXAMPLE 4 (Associative memory and similarity weighting). To model *associative memory* (Thomson and Tulving 1970), assume that

$$m_s(s, a, y) > 0 \quad \text{and} \quad \frac{m_s(s, a, y)}{m_s(s', a, y)} > \frac{m_s(s, a, y)}{m_s(s', a, y)},$$

so that a signal is more likely to trigger memories of experiences where the signal was the same. In general, signals represent the conditions under which the choice is made. For example, when in a particular mood, agents tend to recall situations when they were in that mood before (Matt, Vázquez, and Campbell 1992; Mayer, McCormick, and Strong 1995), and professional economic forecasters overweight periods with a macroeconomic context similar to the current one, but only if they lived through them (Goetzmann, Watanabe, and Watanabe 2022).¹⁴

¹³ Because agents never make choices before the signal realizations, there is no way to pin down the relationship between the utilities of two experiences that differ in their signal component. Therefore, both here and in example 3, the definitions of the biases should be interpreted as saying that there are a u and a Φ such that the conditions are satisfied.

¹⁴ Jehiel (2018) studies investors who base their decisions only on the outcomes of projects that were implemented after the same signal and ignore periods when the signal was different, and Bordalo, Gennaioli, and Shleifer (2020) shows how similarity weighting can lead to the attribution and projection biases.

A leading special case is *similarity-weighted memory*, where the probability of recalling a past experience depends only on the context in which the choice is taken: Here there is a metric $d: S^2 \rightarrow \mathbb{R}_+$, and $m_s(s, a, y) = \Phi(d(s, s'))$ for some strictly decreasing function $\Phi: \mathbb{R}_+ \rightarrow [0, 1]$.

EXAMPLE 5 (Confirmatory-memory bias). The agent has *confirmatory-memory bias* (see Hastie and Park 1986 and Esponda, Vespa, and Yuksel 2023 for evidence of the relevance of memory for confirmation bias) if they are more likely to remember experiences that the prior deems more likely. Suppose that the agent has only two hypotheses, as in Lord, Ross, and Lepper (1979) and Rabin and Schrag (1999), so that $\Theta = \{p^0, p^1\}$, with $\mu(p^0) > \mu(p^1)$. Then, confirmatory-memory bias corresponds to

$$\frac{p_{s,a}^0(y)}{p_{s,a}^1(y)} \geq (>) \frac{p_{s,a}^0(y')}{p_{s,a}^1(y')} \Rightarrow m_s(s, a, y) \geq (>) m_s(s, a, y').$$

III. Long-Run Outcomes

Let \mathbb{P}_π denote the probability measure on the set $(S \times A \times Y)^\mathbb{N}$ of sequences of experiences induced by the objective signal and outcome distributions ζ and p^* , the agent's memory m , and policy π .¹⁵

DEFINITION 1. A strategy σ is

- (i) a *limit strategy* if there is an optimal policy π such that

$$\mathbb{P}_\pi[\sup\{t: a_t \neq \sigma(s_t)\} < \infty] > 0 \quad \text{or}$$

- (ii) a *global attractor* if for every optimal policy π

$$\mathbb{P}_\pi[\sup\{t: a_t \neq \sigma(s_t)\} < \infty] = 1.$$

In words, strategy σ is a limit strategy if there is positive probability that it will be played in every period after some random but finite time, and it is a global attractor if it is a limit with probability 1. This section gives some general results about limit strategies.¹⁶ Section IV then discusses the consequences of some specific memory biases.

A. Selective-Memory Equilibrium

To characterize the strategies that can arise as limit behavior, we define for each strategy σ and signal s' the set of *memory-weighted likelihood maximizers* after s' :

¹⁵ This is the unique extension from the probabilities of the finite histories $(S \times A \times Y)^t$, $t \in \mathbb{N}$.

¹⁶ Example 11 shows that limit strategies may not exist without further assumptions.

$$\Theta_s^m(\sigma) := \operatorname{argmax}_{p \in \Theta} \left(\sum_{s \in S} \hat{s}(s) \sum_{y \in Y} m_{s'}(s, \sigma(s), y) p_{s, \sigma(s)}^*(y) \log p_{s, \sigma(s)}(y) \right). \quad (2)$$

These are the elements of Θ that maximize the likelihood of the memory-weighted outcome distribution induced by σ . Note that only the relative sizes of the weights m matter for determining $\Theta_s^m(\sigma)$: if $\hat{m}(\cdot) = \lambda m(\cdot)$ for some $\lambda > 0$, then \hat{m} and m have the same memory-weighted maximizers.

ASSUMPTION 2 (Maintained assumption). For every $s', s \in S, a \in A, \sigma \in A^S, y \in Y$, and $p \in \Theta_s^m(\sigma), p_{s,a}^*(y) > 0$ implies that $p_{s,a}(y) > 0$.

This assumption requires that no data-generating process that is a best memory-weighted fit to some strategy can be ruled out in finite time.¹⁷

DEFINITION 2. A strategy σ is

- (i) a *selective-memory equilibrium* if for all $s \in S$ there is $\nu_s \in \Delta(\Theta_s^m(\sigma))$ such that $\sigma(s) \in \text{BR}(s, \nu_s)$, or
- (ii) a *uniformly strict selective-memory equilibrium* if for all $s \in S$ and all $\nu \in \Delta(\Theta_s^m(\sigma)), \{\sigma(s)\} = \text{BR}(s, \nu)$.

In a selective-memory equilibrium σ , the action played after each signal s is a best reply to some belief over memory-weighted likelihood maximizers given σ . The uniformly strict version adds the restriction that there is the same unique best reply for each of these maximizers. Both concepts allow the actions played in response to different signals to be justified by different beliefs because which memories are triggered depends on the current realization of the signal.

THEOREM 1. Every limit strategy is a selective-memory equilibrium.

To prove the theorem, we fix a limit strategy σ and suppose by contradiction that is not a selective-memory equilibrium. This means that $\sigma(s')$ is not a best reply to any belief in $\Theta_{s'}^m(\sigma)$ for some $s' \in S$. If $\sigma(s')$ is not a best reply to any belief in Θ , it is never played in response to s' , so it cannot be a limit strategy. If it is a best reply to some belief in Θ but not in $\Theta_{s'}^m(\sigma)$, lemma A3 implies that there exists an experience with objective positive probability under σ that is recalled with positive probability. Lemma A4 then shows that since σ is a limit strategy, for some time t , there is an action sequence a' such that if the agent plays a' and then σ afterward, there is positive probability that the induced sequence of beliefs makes σ optimal at all periods $\tau \geq t + 1$.

Under a policy that converges to σ , when there is an experience with positive probability of being recalled, a variation of the Borel-Cantelli

¹⁷ The misspecified-learning literature usually makes the stronger assumption that the true data-generating process and all the agent's mental models are mutually absolutely continuous. We relax this to allow the natural full-support prior on the action-independent models in our overconfidence application.

lemma (claim 1) implies that a.s. (almost surely) the recalled history is long and that the associated empirical frequency after s' converges to the distribution given by $p_{s,a(s)}^*$. Thus, the strong law of large numbers implies that a large recalled history is representative of the memory-based outcome distribution. With this, we can extend Berk's (1966) concentration result to the beliefs given the recalled experiences to show that distributions that do not maximize the memory-weighted likelihood have vanishing posterior probability on a set of representative recalled histories that has objective probability converging to 1. But then σ must be a selective-memory equilibrium, as otherwise, it could not be a best reply to these beliefs concentrated on the maximizers.

Theorem 1 provides a learning foundation for some equilibrium concepts that have been used in recent work. For example, Kőszegi, Loewenstein, and Murooka (2021) proposes an equilibrium concept where the agent is more likely to remember successes than failures if they are in a good mood, and the agent's mood is determined by their self-esteem, which is a function of the number of past successes they remember. This is a case of our model where the agent's mood is an action chosen to match their perceived probability of succeeding at a task (i.e., their perceived ability). Our equilibrium concept then coincides with Kőszegi, Loewenstein, and Murooka's (2021) "self-esteem personal equilibrium," and theorem 1 shows that any long-run learning outcome must be such an equilibrium.

We also provide a foundation for Berk-Nash equilibrium based on selective memory. For example, section IV.B shows that positive-memory bias can lead to overconfidence. Overconfidence has been modeled as the result of exogenous misspecification; the fact that it can be endogenously derived from a well-documented memory bias provides a microfoundation for Berk-Nash equilibrium in this context. More generally, proposition 6 shows that any Berk-Nash equilibrium can be microfounded through selective memory. Finally, section V.A shows that in our setting, the long-run action of an agent with underinference (Phillips and Edwards 1966) must be a selective-memory equilibrium.

B. Global Convergence to Equilibrium

We now give a sufficient condition for the agent's strategy to globally converge to a uniformly strict selective-memory equilibrium, which a fortiori implies that such equilibria exist.¹⁸

ASSUMPTION 3 (Identification).

- (i) For all $(s, s', a) \in S^2 \times A$, $\sum_{y \in Y} p_{s,a}^*(y) m_{s'}(s, a, y) > 0$.

¹⁸ Fudenberg, Lanzani, and Strack (2024) show that the heterogeneous-belief version of selective-memory equilibrium does exist.

(ii) There is a $\hat{p} \in \Theta$ such that for every $(s, s', a) \in S^2 \times A$,

$$\operatorname{argmax}_{p \in \Theta} \left(\sum_{y \in Y} m_i(s, a, y) p_{s,a}^*(y) \log p_{s,a}(y) \right) = \{\hat{p}\}.$$

The substantial assumption here is assumption 3ii, which requires that the best fit for the remembered distribution is independent of the agent’s action. This assumption is always satisfied if the agent correctly believes that their actions have no influence on the distribution of outcomes and has the same memory function for each action, as in the examples in section IV. Note that without an assumption along these lines, the data-generating process that best explains the agent’s observations given one action a_1 could lead to a belief that makes another action a_2 optimal, which then, in turn, could lead to a belief that makes a_1 optimal. Clearly, such cycles would preclude global convergence; see example 11. Beyond that, the assumption requires that for each (s, s', a) , there is an outcome with a positive probability of being remembered. Our next result considers closed balls $B_\varepsilon(\hat{p})$ around the data-generating process \hat{p} of assumption 3ii, where the distance used to define the ε balls is the maximum of the total variation distance between their signal-action contingent distributions.¹⁹

THEOREM 2.

(i) Under assumption 3, for every optimal strategy π and every $\varepsilon > 0$,

$$\mathbb{P}_\pi \left[\lim_{t \rightarrow \infty} \mu(B_\varepsilon(\hat{p}) | h_t(R_t)) = 1 \right] = 1.$$

(ii) If in addition $\operatorname{BR}(s, \delta_{\hat{p}})$ is a singleton for all s , then $\hat{\sigma}$ is a global attractor, where

$$\hat{\sigma}(s) = \operatorname{BR}(s, \delta_{\hat{p}}) \quad \forall s \in S.$$

It is the unique selective-memory equilibrium, and it is uniformly strict.

The proof starts by using a “mixingale law of large numbers” (see, e.g., Hall and Heyde 2014) to conclude that the outcome frequency converges to the one predicted by the true data-generating process and the agent’s actions. We then address the complication posed by the fact that memory is stochastic, so even when the agent has played many times, their beliefs can be very different from one period to the next, unlike in learning models with perfect memory, preventing the use of classical martingale arguments for beliefs (see, e.g., Battigalli et al. 2019).

¹⁹ Formally, $B_\varepsilon(\hat{p}) = \{q \in \Delta(Y)^{A \times S} : \max_{s \in S, a \in A} \|q_{s,a} - \hat{p}_{s,a}\|_{TV} \leq \varepsilon\}$.

We first use the Chernoff inequality to provide an upper bound on the probability that the recalled empirical frequency significantly diverges from the memory-distorted version of the actual empirical frequency. This upper bound is then combined with the Borel-Cantelli lemma to show that for every $\gamma \in (0, 1)$, there exists a random but a.s. finite time after which any signal-action pair with frequency at least γ does not have a large deviation from the memory-distorted empirical distribution of its induced consequences and has recalled frequency bounded away from 0.

Assumption 3 implies that for every signal-action pair, model \hat{p} is the unique model that best fits the memory-adjusted theoretical distribution. Thus, because recalled memories are representative after signal-action pairs that have positive frequency, and pairs with low frequency have negligible impact on beliefs, beliefs concentrate on \hat{p} . When there is a unique best reply to \hat{p} , this implies that the agent's behavior converges as well.

REMARK 1. In our model, the set of recalled histories is not only stochastic but nonmonotonic: the agent might remember a past event one day and not another, which fits the evidence on memory retrieval (see, e.g., Kahana 2012). Appendix B.4 analyzes the limit implications of an alternative model where the memory function determines the probability that an experience is recalled in the period just after it occurs. If it is recalled, it is never forgotten; if not, it is never remembered. Because experiences recalled at later dates include all those recalled earlier, in this alternative model, the agent's past actions do not convey additional information. As with the model we present here, any limit action must be a selective-memory equilibrium.

IV. Specific Forms of Selective Memory

A. Similarity-Weighted Memory and Self-Confirming Equilibrium

DEFINITION 3. Strategy σ is a (unitary) *self-confirming equilibrium* if for all $s \in S$ there is $\nu_s \in \Delta(\Theta)$ such that for all $p \in \text{supp} \nu_s$, $p_{s,\sigma(s)} = \hat{p}_{s,\sigma(s)}^*$ and $\sigma(s) \in \text{BR}(s, \nu_s)$.

Unitary self-confirming equilibrium (Battigalli 1987; Fudenberg and Levine 1993a) requires that the action played is a best response to a belief that is correct on the equilibrium path but possibly incorrect about off-path actions.²⁰

²⁰ Fudenberg and Kreps (1988) shows how such actions can be the long-run limit of myopic learning, and Fudenberg and Kreps (1995) shows that any long-run outcome with purely myopic players and a single agent in each player role must be a unitary self-confirming equilibrium. Fudenberg and Levine (1993b) shows that when there are many agents in each player role, the long-run outcome must be a heterogeneous-beliefs self-confirming equilibrium, whether players are myopic or not.

PROPOSITION 1. For an agent with similarity-weighted memory (example 4), a strategy is a selective-memory equilibrium if and only if it is a self-confirming equilibrium.

More generally, this conclusion holds whenever $m_s(s, a, y)$ does not depend on a or y , as the true distribution is the best fit for every signal, so the weight assigned to each signal does not matter. However, similarity weighting can change the set of selective-memory equilibria when the agent is misspecified.²¹

B. Ego-Boosting Memory Bias and Overconfidence

It is well established that many people are more likely to recall situations that reflect positively on themselves.²² This leads to a particular kind of pleasant-memory bias: they are more likely to remember experiences that boost their self-assessment than those that reduce it.

Consider a situation where the agent observes i.i.d. outcomes $y_t \in Y \subset \mathbb{R}$ that reveal information about an ego-relevant characteristic such as IQ. There are no signals, A is endowed with a linear order, and the agent (correctly) believes that their action does not affect the realized outcome. The next proposition shows that a larger bias leads to a more positive limit belief and higher limit action. This provides a selective-memory foundation for the positive correlation between an agent's happiness and the inaccuracy of their beliefs documented in Alloy and Abramson (1979).

PROPOSITION 2. Suppose that m , m' , and p^* are constant in a , that $m'(a, y) = f(y)m(a, y)$ for some increasing function f , that u is supermodular, and that $\Theta = \Delta(Y)$. The agent's long-run belief with memory m' concentrates on a distribution of outcomes weakly higher in first-order stochastic dominance than the distribution under the long-run belief with memory m , and the limit action with memory m' will be weakly higher than the limit action with memory m .

Intuitively, because the prior assigns positive probability to all action-independent outcome distributions, the memory-weighted likelihood maximizer will be the outcome distribution that exactly matches what the agent remembers. The agent's selective memory makes this recalled history more favorable than the true one, and because the agent's utility function is supermodular, their limit action is weakly higher than the objective optimum.

²¹ Also, even when there is a unique selective-memory equilibrium, and it is objectively optimal, the speed of convergence to the equilibrium can be influenced by similarity weighting. This is similar to what happens with case-based decision theory (Gilboa and Schmeidler 2001) and kernel density estimation, where the optimal bandwidth trades off having enough observations with relying too much on distant values.

²² See, e.g. Mischel, Ebbesen, and Zeiss (1976), Adler and Pansky (2020), Chew, Huang, and Zhao (2020), and Zimmermann (2020).

EXAMPLE 6. Suppose that each period the agent takes an action $a \in \{0, 1\}$, with $u(a, y) = a(y - z)$, $z \in (0, 1)$. Here y is the outcome of an IQ test, which is either pass, $y = 1$, or fail, $y = 0$, so $a = 1$ is optimal if and only if the probability of passing the test exceeds z . The agent passes the test with probability p^* . They always recall passed tests, and they recall failed tests with probability ϕ :

$$m(a, y) = \begin{cases} 1 & \text{if } y = 1, \\ \phi & \text{if } y = 0. \end{cases}$$

In the long run, the agent believes that the probability of passing an IQ test is

$$p = \frac{p^*}{\underbrace{p^*}_{\text{Successes}} + \underbrace{(1 - p^*) \times \phi}_{\text{Failures}}} = p^* + \frac{p^*(1 - p^*)(1 - \phi)}{\phi + (1 - \phi)p^*}.$$

For example, if the true probability p^* is .5, and the agent remembers failing an IQ test with probability .8, in the long run, they believe that they pass the test with probability 5/9. Consequently, they will behave like an exogenously misspecified agent who dogmatically believes their ability to pass is 5/9. Moreover, the difference between p and p^* is monotonic in the agent's selectivity bias ϕ .

This example relates to an experiment by Zimmermann (2020) in which subjects took an IQ test and received three noisy observations of how well they performed relative to other subjects. Zimmermann (2020) finds that all subjects can recall the signals immediately after observing them, but subjects who received negative feedback were less likely to recall the feedback a month later than subjects who received positive feedback: subjects are roughly 20% more likely to state that they "cannot recall" the result of the IQ test if the feedback was negative, even though that answer is payoff-dominated in the experiment and there were only three things for subjects to try to recall.²³ Thus, at least in this experiment, selective memory is a better explanation than selective attention for long-run overconfidence.

Example 6 and proposition 2 also relate to the literature on overconfidence and financial decision-making. Walters and Fernbach (2021) finds that investors are 10% less likely to recall an investment that led to a loss compared to an investment that led to a gain. Moreover, selective memory predicts overconfidence, and overconfidence is reduced when investors

²³ Zimmermann (2020, 339) finds that "negative feedback is indeed recalled with significantly lower accuracy, compared to positive feedback." Here lower accuracy means both that the agents are more likely to report that they do not recall the experience and that they misreport the experience.

rely less on memory. In an incentivized experiment, Gödker, Jiao, and Smeets (forthcoming) finds that subjects overremember good investment outcomes and underremember bad investment outcomes. In line with the prediction of proposition 2, this leads subjects to have overly optimistic beliefs about their investments and reinvest in bad investments more often. Gervais and Odean (2001) studies a different bias: where traders overweight successful trades when learning about their ability, this can lead to overconfidence in a similar way as selective memory.

Ego-boosting bias and misattribution.—We next show how an agent with ego-boosting bias can misinterpret data about other aspects of the world.

EXAMPLE 7. Suppose that, besides taking an IQ test, the agent works on a project with a coworker. The outcome distributions $(p, q) \in [0, 1]^2$ and outcome $(y_1, y_2) \in \{0, 1\}^2$ are two-dimensional, where the first component denotes whether or not the agent passed an IQ test and the second component denotes whether a group project succeeded. The agent passes the IQ test with probability p^* , and the project succeeds with probability $\alpha p^* + (1 - \alpha)q^*$, where $1 - \alpha$ is the share of the work done by the coworker. The agent always remembers experiences with positive IQ test results and remembers experiences with negative test results with probability $\phi \in (0, 1)$. Thus, beliefs concentrate on

$$p = p^* + \frac{p^*(1 - p^*)}{\phi/(1 - \phi) + p^*} \quad \text{and} \quad q = q^* - \frac{\alpha}{1 - \alpha} \frac{p^*(1 - p^*)}{\phi/(1 - \phi) + p^*}.$$

The agent underestimates the coworker’s ability, and the underestimation grows as memory becomes more selective.

To generalize this example, we consider a two-dimensional outcome space $Y = Z \times Z \subset \mathbb{R}^2$, where y_1 corresponds to an ego-relevant characteristic and is distributed according to p^* . The second component y_2 is drawn independently with probability $\alpha p^*(y_2) + (1 - \alpha)q^*(y_2)$ for some $\alpha \in (0, 1)$. The agent knows that the outcomes are independently drawn each period according to these conditions, but does not know p^* or q^* , and their prior belief assigns positive probability to each of these distributions.²⁴

PROPOSITION 3. If m is constant in a and y_2 , increasing in y_1 , and there is $y \in Y$ with $p^*(y) m(y) > 0$, then the agent’s long-run belief about p concentrates on a distribution \hat{p} that is weakly higher in first-order stochastic dominance than p^* , and the agent’s long-run belief about q concentrates on a distribution that is weakly lower than q^* .

“Perhaps the most robust finding in the psychology of judgment is that people are overconfident” (DeBondt and Thaler 1995, 389). The proposition provides an explanation for two commonly found forms of

²⁴ Formally, $\Theta = \{r \in \Delta(Z \times Z) : r(y_1, y_2) = p(y_1)[\alpha p(y_2) + (1 - \alpha)q(y_2)] \text{ for some } p, q \in \Delta(Z)\}$.

overconfidence: (i) overestimation of one's own absolute level of performance and (ii) overestimation of one's performance relative to others (see, e.g., Svenson 1981; Merkle and Weber 2011). For example, Gilovich (2008) finds that 94% of college professors thought they were better than their average colleague.²⁵

Reinforcement through actions.—Actions can play an important role in amplifying the misconceptions caused by selective memory. For example, suppose that in example 7, the agent starts out with an unbiased belief about their coworker's ability and each period t chooses the fraction $1 - \alpha_t$ of work to delegate to them. Because here the memory-weighted likelihood maximizers do depend on the agent's action, theorem 2 does not apply, but as in Heidhues, Kőszegi, and Strack's (2018) analysis of exogenously overconfident agents, there is a global attractor: As the agent overremembers their own successes, they become overconfident about their own ability, and to explain the disappointingly low frequency of successes in the group project, they became overly pessimistic about their coworker's. The agent thus delegates less work to their coworker, whose ability then has a smaller effect on output. To explain the disappointingly low output, the agent becomes even more pessimistic about the coworker's ability, leading to even less delegation in the unique limit strategy.

Changes to the informational environment.—More generally, section V.B shows that the long-run belief induced by selective-memory equilibria can be replicated by exogenous misspecification in any fixed environment, and vice versa. However, selective memory and exogenous misspecification can lead to very different predictions about the effect of changes in information. Suppose, for example, that negative feedback is delivered along with positive feedback about an unrelated trait of the agent. Combining positive and negative information in this way makes a "feedback sandwich," which the management and psychology literatures suggest strengthens the effect of feedback.²⁶ If the positive feedback makes the experiences with failed IQ tests less unpleasant, an agent with positive-memory bias would be more likely to remember them, so their long-run belief would move closer to their actual ability, and they would be less biased about their coworker's ability. In contrast, with exogenous misspecification, positive feedback about an unrelated state would not affect the agent's beliefs about their own or their coworker's ability.

²⁵ Benoit and Dubra (2011) shows how this "I'm-better-than-average effect" can be explained within a purely Bayesian framework; Benoit, Dubra, and Moore (2015) provides more direct evidence for relative overconfidence that rules out the purely rational explanation.

²⁶ Procházka, Ovcari, and Durinik (2020) describes an experiment where bundling negative feedback with positive feedback about an unrelated domain helps agents perform better.

C. *Extreme-Experience Bias and Risk Attitudes*

This section shows that for choices over lotteries, selective memory can generate the same behavior as distorted risk attitudes. We again simplify by supposing that there are no signals, and let the outcome $y \in \mathbb{R}$ be the amount of money received by the agent, with $u(s, a, y) = v(y)$ for some increasing and concave v .

Extreme-experience bias.—Suppose that the agent chooses between a safe action $a = 0$ that induces outcome y_0 and a risky lottery $a = 1$ with expected value \bar{y} . We say that the agent has an *extreme-experience bias* if the probability of remembering an experience m is an increasing function of the distance of the outcome y from its expected value and does not depend on s or a :

$$m(s, a, y) = h(|y - \bar{y}|) \tag{3}$$

for some increasing $h : \mathbb{R}_+ \rightarrow [0, 1]$. Our next result shows that the risky lottery is a selective-memory equilibrium with extreme-experience bias only if it is a selective-memory equilibrium with perfect memory. Moreover, example 10 in appendix B shows that extreme-experience bias can shift the long-run outcome from the lottery to the safe action. To state a result that holds for all concave utility functions, we assume that the true distribution of outcomes is symmetric.²⁷

PROPOSITION 4. Suppose that the distribution p_1^* is symmetric and that the agent thinks all outcome distributions are possible under the risky action.²⁸ If choosing the lottery is not a self-confirming equilibrium, it is not a selective-memory equilibrium with extreme-experience bias.

Because the agent overremembers extreme experiences, the environment seems riskier than it truly is, so in the long run, they do not take the risky action if it would not be optimal for an agent without extreme-experience bias. By making the tail realizations relatively more memorable, extreme-experience bias makes a risk-averse agent act as if they were even more risk averse. This may help explain why the risk aversion needed to match the real-world investment choices is unrealistically high: the agents can be attracted by safe alternatives because they are moderately risk averse, and their memory exaggerates the riskiness of the uncertain alternatives. For example, a single day when the stock market crashed might be more easily remembered than many days of average returns, leading to a biased perception of its riskiness. Indeed, the plausibility of this channel is supported by several studies that show that higher working

²⁷ Extreme-experience bias can have the opposite effect of encouraging risk-taking behavior when the true distribution is very asymmetric with a very low probability of a large payoff.

²⁸ That is, $p_1^*(\bar{y} + c) = p_1^*(\bar{y} - c)$ for all $c \in \mathbb{R}$, and $\Theta = \{p \in \Delta(Y)^A : p_0(y_0) = 1\}$.

memory is associated, either directly or through a proxy measure of cognitive ability, with lower risk aversion at both the intra- and interpersonal levels (see, e.g., Cokely and Kelley 2009; Boyle et al. 2012; Benjamin, Brown, and Shapiro 2013).

Rare-experience bias.—Similarly, some forms of selective memory are equivalent to preferences that arise from distorting outcome probabilities. Suppose that the agent is more likely to remember experiences that happen more rarely, that is, there is a decreasing function $h : [0, 1] \rightarrow [0, 1]$ such that $m(s, a, y) = h(p_1^*(y))$. In this case, in the long run the agent believes that the outcome distribution for the risky action is

$$\frac{h(p_1^*(y))}{\sum_{z \in Y} h(p_1^*(z))}.$$

They will thus act as if they distort probabilities, as in prospect theory (Kahneman and Tversky 1979).²⁹

V. Alternative Models

This section compares our selective memory model with underinference and misspecification, which are two other ways to model similar effects.

A. Underinference

The phenomenon of underinference (Phillips and Edwards 1966) is distinct from selective memory but has similar long-run implications, as we establish in proposition 5. Here agents remember (or are presented with) a record of past observations, so memory is not an issue, and the agent's beliefs are a deterministic function of the sequence of observations. However, they underweight a given observation (s, a, y) when applying Bayes's rule. In particular, they use the deterministic updating rule

$$\mu^U(C | (s_i, a_i, y_i)_{i=1}^t) = \frac{\int_C \prod_{i=1}^t (p_{s_i, a_i}(y_i))^{m(s_i, a_i, y_i)} d\mu(p)}{\int_{\Theta} \prod_{i=1}^t (p'_{s_i, a_i}(y_i))^{m(s_i, a_i, y_i)} d\mu(p')}, \quad (4)$$

for every measurable $C \subseteq \Theta$, where $m(s, a, y) \in [0, 1]$ is the underinference distortion applied to experience (s, a, y) .

²⁹ We view this specification, where m depends on the theoretical frequency p^* , as a convenient modeling shortcut for long-run outcome when instead m depends on the empirical outcome frequency.

As with selective memory, this memory distortion leads beliefs to concentrate on the memory-weighted likelihood maximizers, and as the next result shows, the underinference distortion maps directly to a selective-memory function.³⁰

PROPOSITION 5. If σ is a limit strategy with underinference distortion m , it is a selective-memory equilibrium with memory function m .

A leading special case is *uniform underinference*, where $m(s, a, y) = c < 1$ and the agent discounts all observations by the same amount. In this case, propositions 1 and 5 imply that the limit strategy for a correctly specified agent must be a (unitary) self-confirming equilibrium.³¹ It seems difficult to distinguish selective memory from underinference bias using data about beliefs alone, and none of the data we have found includes information on which histories the subjects recall (see Benjamin 2019 for a survey).

If signals are absent and actions are real-valued, the way actions respond to outcomes can be used to distinguish underinference and selective memory. Under overconfidence, the realization of y_t is sufficient to predict whether a_{t+1} is more or less than a_t . Under selective memory, the set of past experiences retrieved at time $t + 1$ may differ from those at time t , so in general the previous period's outcome and action are not sufficient to predict how actions change. Moreover, the action sequence features a sort of regression to the mean: after a particularly high action, the next action will likely be lower.

In general, with an exogenous data-generating process, the agent's beliefs will converge to the same limit with selective memory as with underinference, so their limit action will be the same. If the data-generating process is endogenous, random memory realizations can induce switches in actions, reducing the set of actions that can be long-run limits for a given memory function. The following example illustrates this possibility.

EXAMPLE 8. There are no signals, $A = \{d', d''\}$, $Y = \{0, 1\}$, $u(\cdot, y) = y$, and the agent knows that the probability of $y = 1$ given action d' is some $c \in (0, 1)$, that is, $p_{d'}(1) = p_{d'}^*(1) = c$ for all $p \in \Theta$. The agent does not know the probability of outcome 1 under action d'' . Their initial belief is that it is larger than that of action d' , so $\text{BR}(\mu) = d''$, although there is $p' \in \Theta$ with $p_{d'}(1) < c$. The truth is that $1 > p_{d''}^*(1) > c$, so action d'' is optimal, but if m is constant and strictly positive, both d' and d'' are selective-memory equilibria. In the underinference model, d' is a limit action for any such m , and if $a_t = d'$, then $a_\tau = d'$ for all $\tau > t$. Instead, with the selective memory model, d' is not a limit action because if d' is played only a finite number of times, there is a positive probability

³⁰ We identify the underinference distortion with the vector of memory functions that do not depend on the current signal.

³¹ Frick, Iijima, and Ishii (2021) shows that uniform underinference leads to the same speed of belief convergence as correct updating in a setting with a fixed outcome distribution.

of forgetting all such experiences and only using the prior to choose the action, which favors action d' .

More generally, selective memory does not generate as much long-run inefficiency as underinference: Whenever the agent believes that the consequences of different actions are independent, if the expected utility of a selective-memory equilibrium a under the memory-weighted likelihood maximizer is lower than the ex ante value of an alternative b , then a is not a limit strategy.

B. Selective Memory and Misspecification

We now relate the long-run implications of selective memory to those of *misspecification* in the sense of the statistics literature, where the true model is not in the support Θ of the agent's prior, and the agent remembers all past observations. The case studied in the misspecification literature has perfect memory, so there $m = 1$ and $\Theta_s^1(\sigma)$ does not depend on s' , so we simply write $\Theta^1(\sigma)$.

DEFINITION 4. A strategy σ is

- (i) a *Berk-Nash equilibrium* if there exists $\nu \in \Delta(\Theta^1(\sigma))$ such that for all $s \in S$, $\sigma(s) \in \text{BR}(s, \nu)$ or
- (ii) a *uniformly strict Berk-Nash equilibrium* if for all $\nu \in \Delta(\Theta^1(\sigma))$ and all $s \in S$, $\{\sigma(s)\} = \text{BR}(s, \nu)$.

Esponda and Pouzo (2016) shows that only Berk-Nash equilibria can be the long-run outcomes of misspecified learning, and Fudenberg, Lanzani, and Strack (2021) shows that in "rich" environments only uniformly strict Berk-Nash equilibria are stable long-run outcomes.

There is a close relationship between the uniformly strict versions of Berk-Nash equilibrium and selective-memory equilibrium: For a given prior support Θ , every uniformly strict Berk-Nash equilibrium is equivalent to a selective-memory equilibrium with full-support prior for some memory function, and every uniformly strict selective-memory equilibrium is equivalent to a Berk-Nash equilibrium for some support.

DEFINITION 5. A Berk-Nash equilibrium σ with support Θ and a selective-memory equilibrium σ' with support $\tilde{\Theta}$ and memory function m are *belief equivalent* if $\sigma = \sigma'$, and for all $s \in S$ there exists a belief $\nu \in \Delta(\Theta^1(\sigma) \cap \tilde{\Theta}_s^m(\sigma))$ such that $\sigma(s) \in \text{BR}(s, \nu)$.

Two equilibria are belief equivalent if they prescribe the same strategies and if behavior after each signal can be justified by the same belief.

PROPOSITION 6.

- (i) Every uniformly strict Berk-Nash equilibrium σ where $\text{supp} q_{s,a} \subseteq \text{supp} p_{s,a}^*$ for all $q \in \Theta^1(\sigma)$, $s \in S$, and $a \in A$ is belief equivalent to a

- selective-memory equilibrium with full support for some strictly positive memory function.
- (ii) Every uniformly strict selective-memory equilibrium with support Θ is belief equivalent to a uniformly strict Berk-Nash equilibrium for some Θ' .

The idea behind the first part of the proposition is that if we start from a maximizer p with perfect memory but incomplete support, we can choose a memory function that rescales the probability of each (s, a, y) by some constant times $p_{s,a}(y)/p_{s,a}^*(y)$. This makes the recalled frequency equal to p , so p is a weighted-memory likelihood maximizer, and σ is the best reply.³² Here, the absolute continuity requirement is needed because selective memory cannot replicate misspecifications where likelihood-maximizing models assign positive probability to an event that can never be realized. To the best of our knowledge, all of the examples of misspecification studied in the literature satisfy this restriction. The second part of the proposition is trivial: To construct a uniformly strict Berk-Nash equilibrium that leads to the same beliefs and behavior as in the selective-memory equilibrium, we can endow the agent with a degenerate belief that equals the belief in the specified selective-memory equilibrium.

REMARK 2. As we prove in appendix B.2, the uniform-strictness conditions of proposition 6 are needed:

1. There are Berk-Nash equilibria that are not belief equivalent to any selective-memory equilibrium with full support and strictly positive memory function.
2. There are selective-memory equilibria that are not belief equivalent to any Berk-Nash equilibrium.

Moreover, selective-memory equilibria need not be objectively optimal when the agent knows that the distribution of outcomes is independent of their action ($p_{s,a}^* = p_{s,a}'^*$ and $p_{s,a} = p_{s,a}'$ for every $p \in \Theta$, $a, a' \in A$, $s \in S$).

To illustrate the equivalence result, consider a buyer who submits an offer for a good in a double-blind two-sided auction where the price z is set at the buyer's bid, so the seller's dominant strategy is to bid their value. Suppose that the buyer has an exogenously fixed conviction that the price sellers ask is independent of the quality of the good they are selling. If the buyer's value of the good is $x + v + \varepsilon$ where $x \in X \subseteq \mathbb{R}$ is the value for the seller, $v \in V \subseteq \mathbb{R}$ measures the gains from trade, and ε is a noise term, then in the Berk-Nash equilibrium they submit a bid that is too low, as in

³² Every p'' that is outcome equivalent under σ is also a maximizer, and this p'' may not have been an element of Θ . Because σ need not be a best response to some of them, it need not be a uniformly strict selective-memory equilibrium.

Esponda (2008). Proposition 6 shows that memory distortions can, over time, lead the agent to believe that value and bid are independent and thus have the same long-run behavior and beliefs. This is obtained with a memory function that gives more weight to experiences with a larger gap between buyer's values and ask prices.³³

Persistence.—While agents undoubtedly are sometimes misspecified, some recent papers have theoretically questioned how likely these misperceptions are to persist and have proposed mechanisms by which agents might realize that some model not in the support of their initial beliefs better fits the data (Schwartzstein 2014; Fudenberg and Lanzani 2023; He and Libgober 2023; Lanzani 2024). In contrast, an agent with a selective memory and a full-support prior will be able to explain their recollections with one of their conjectured models and so have less reason to learn of their errors.

VI. Discussion

Our equilibrium concept and results make it easy to predict the long-run implications of arbitrary memory biases, which should be of broad use in applied work. We illustrated our framework by showing that it explains how overconfidence can arise from an ego-boosting memory bias and why agents may underestimate their coworkers' abilities even when they are correctly specified. It also lets us explain the excessive levels of risk aversion implied by asset choice as the result of moderate risk aversion paired with an extreme-experience bias that leads agents to overestimate the riskiness of the assets.

Distinguishing between models.—While proposition 6 implies that selective memory and misspecification will have similar long-run implications in a fixed environment, section IV.B shows that the two models have different comparative statics with respect to changes in the environment.³⁴ If we look at the correspondences mapping a true data-generating process into the sets of selective-memory equilibria and uniformly strict Berk-Nash equilibria, our result says that for a fixed p^* we can find an m that makes an element of the image of the Berk-Nash correspondence at p^* an element of the selective-memory equilibrium correspondence at p^* . But this property of the correspondence is lost at a different \tilde{p}^* , which can let us distinguish between the models. For example, the Berk-Nash equilibrium of a degenerate misspecified model has a constant graph, something selective

³³ Specifically $m(a, (x, v)) = k[\sum_{v' \in V} p^*(x, v') \sum_{x' \in X} p^*(x', v)] / p^*(x, v)$ for sufficiently small $k > 0$.

³⁴ Selective memory can arguably be viewed as a form of misspecification, as the agent is not aware of their memory limitations. From that perspective, our results show that the classic misspecification studied in Bayesian statistics is closely related to a psychologically founded form of misspecification.

memory function cannot replicate for nontrivial utility and memory functions. More generally, empirical work might be able to distinguish between the two models based on how agents respond to changes in the true data-generating process (for a concrete example in the context of overconfidence, see the discussion at the end of sec. IV.B).

To distinguish between selective memory and underinference, one can elicit the agent's beliefs.³⁵ Underinference predicts that the likelihood ratio between two data-generating processes θ and θ' always increases between period $t - 1$ and t if the period $t - 1$ outcome was more likely under θ . Selective memory allows for violations of this monotonicity, especially if at the beginning of period t a signal triggering experiences favoring θ' is observed, while this signal is irrelevant with underinference. A more direct way to distinguish selective memory from other sources of mistaken inference, including misspecification, is to elicit both what the agent remembers and what they believe, as in Huffman, Raymond, and Shvets (2022) and Gödker, Jiao, and Smeets (forthcoming), where an important role for memory is observed. This allows one to estimate the memory function and qualitatively distinguish among selective memory, misspecification, and underinference.

Convergence to equilibria.—Theorem 2 gives sufficient conditions for there to be a global attractor. Even when no such strategy exists, one could hope that there is probability 1 of converging to some limit strategy, with which strategy occurs depending both on the agent's prior and on the realized outcomes. We hope to find sufficient conditions for that in future work, along with (presumably weaker) conditions that ensure a positive probability of converging to a limit strategy.

Partial naïveté.—We have assumed that agents treat the experiences they remember as if these were the only ones that happened. Appendix A3 considers agents who are partially aware of their memory limitations. To do this, we assume that agents know calendar time and therefore how many observations they have not remembered.³⁶ For an agent who is aware of their own forgetfulness but unaware that their memory is selective, the selective-memory equilibria under partial and full naïveté coincide. At the other extreme, if agents are fully aware of their memory function, any action that is optimal for the true data-generating process is always a selective-memory equilibrium.

Finite memory.—In our model, the number of recalled experiences converges to infinity, as if the agent had perfect memory. In Fudenberg, Lanzani, and Strack (2024), we modify the model to make the expected

³⁵ However, see Danz, Vesterlund, and Wilson (2022) for practical challenges in belief elicitation.

³⁶ As example 9 in app. A3 shows, less naïve agents can take worse actions and get lower payoffs.

number of recalled periods bounded. Here the agent's beliefs need not converge to a deterministic limit even when the strategy is fixed, which can make the limit behavior stochastic. Thus, instead of characterizing the possible limit strategies, we show that if the frequency with which strategies are used converges, the limit-strategy distribution is generated by a best response to the distribution of memories it generates. We also use this to model the effect of "rehearsal," where an experience recalled in one period is more likely to be recalled again.

Other possible extensions.—It would be relatively easy to extend our analysis to agents who "misremember" and access false memories as opposed to simply forgetting things that happened. A more substantive generalization would be from an agent who believes that outcomes are i.i.d. to an agent who believes that outcomes follow a Markov process. This would let us capture the gambler's fallacy (see Rabin and Vayanos 2010; He 2022) if an outcome is more memorable when it is different from the outcome in the previous period.³⁷ Or it might be much easier for agents to recall whether an experience happened at all than whether it happened five or six times; we could capture this by using a memory function that is concave in the number of times an experience occurred. Another generalization would be to memory functions with recency bias, such as $m_{s,t}(s_\tau, a_\tau, y_\tau) = m_s(s_\tau, a_\tau, y_\tau)f(t - \tau)$ where f is a decreasing function. As with associative memory, when the outcomes are exogenous and f is bounded away from 0, this bias leads only to slower learning, but when outcomes are endogenous, it can prevent the agent from locking on to the optimal action.

Appendix A

Proofs of the Main Results and Partial Naïveté

A1. Preliminaries

For every $t \in \mathbb{N}$, we first explicitly describe the map

$$H_t \times 2^{\{1, \dots, t\}} \rightarrow H,$$

$$(h_t = (s_i, a_i, y_i)_{i=1}^t, R) \mapsto h_t(R_t)$$

that transforms an objective history and a set of recalled periods into the recalled history. Let $n(k, R_t) = \tau$ if τ is the k th-smallest number in R_t , that is, $n(1, R_t) = \tau$ if τ is the first period that is recalled, $n(2, R_t) = \tau$ if τ is the second period that is recalled, and so on. The recalled history is

$$h_t(R_t) = (s_{n(k, R_t)}, a_{n(k, R_t)}, y_{n(k, R_t)})_{k=1}^{|R_t|}. \quad (5)$$

³⁷ This extension could make use of the analysis of belief concentration for misspecified agents with Markov models developed in Fudenberg, Lanzani, and Strack (2023).

Combining equations (1) and (5), we have that the posterior probability of every measurable $C \subseteq \Theta$ after objective history h_t when the recalled periods are $R_t \neq \emptyset$ is

$$\frac{\int_C \prod_{\tau \in R_t} p_{s,a}(y_\tau) d\mu(p)}{\int_{\Theta} \prod_{\tau \in R_t} p_{s,a}(y_\tau) d\mu(p)}. \tag{6}$$

We now state a few lemmas whose proofs are in appendix B. For every $h_t \in H$ let $f(h_t) \in \Delta(S \times A \times Y)$ denote the empirical distribution over signals, actions, and outcomes in $h_t = (s_i, a_i, y_i)_{i=1}^t$, and let

$$\hat{f}(h_t, R_t)(s, a, y) = \frac{1}{|R_t|} \sum_{i \in R_t} \mathbf{1}_{\{(s,a,y)\}}(s, a, y)$$

denote the recalled empirical distribution in objective history h_t when the recalled periods are $\emptyset \neq R_t$. Also, for every $\gamma \in \Delta(S \times A \times Y)$ and $p \in \Delta(Y)^{S \times A}$ let

$$L(\gamma||p) = \sum_{(s,a,y) \in S \times A \times Y} \gamma(s, a, y) \log(p_{s,a}(y))$$

be the log likelihood of the distribution γ with respect to data-generating process p .

The next result shows that the posterior beliefs concentrate on the likelihood maximizers given the recalled empirical distribution.

LEMMA A1. For all Borel-measurable $C, C' \subseteq \Delta(Y)^{S \times A}$, $t \in \mathbb{N}$, $h_t \in H_t$, and $R_t \subseteq \{1, \dots, t\}$,

$$\frac{\mu(C|h_t(R_t))}{1 - \mu(C'|h_t(R_t))} \geq \frac{\mu(C)}{1 - \mu(C')} \exp\left(|R_t| \left[-\sup_{p \in \Theta \setminus C'} L(\hat{f}(h_t, R_t)||p) + \inf_{p \in C} L(\hat{f}(h_t, R_t)||p) \right]\right).$$

Let $\Theta_s^m(\sigma, \varepsilon) = \{p \in \Theta : \exists q \in \Theta_s^m(\sigma), \|p - q\|_\infty \leq \varepsilon\}$ denote an ε ball around the memory-weighted maximizers.

LEMMA A2. If σ is not a selective-memory equilibrium, there are $s' \in S$ and $\varepsilon, C \in \mathbb{R}_{++}$ such that for all $\nu \in \Delta(\Theta)$,

$$\frac{\nu(\Theta_s^m(\sigma, \varepsilon))}{1 - \nu(\Theta_s^m(\sigma, \varepsilon))} > C \Rightarrow \sigma(s') \notin \text{BR}(s', \nu).$$

If σ is a uniformly strict selective-memory equilibrium, there are $\varepsilon, C \in \mathbb{R}_{++}$ such that for all $s \in S$ and $\nu \in \Delta(\Theta)$,

$$\frac{\nu(\Theta_s^m(\sigma, \varepsilon))}{1 - \nu(\Theta_s^m(\sigma, \varepsilon))} > C \Rightarrow \{\sigma(s)\} = \text{BR}(s, \nu).$$

The next lemma says that if an action is an undominated response to some signal s' but cannot be played as a response to s' in any selective-memory equilibrium, then after signal s' the agent must have a nonzero chance of remembering at least one possible experience (s, a, y) .

LEMMA A3. If $\sigma(s') \in \text{BR}(s', \Delta(\Theta)) \setminus \text{BR}(s', \Delta(\Theta_s^m(\sigma)))$, then there is $(\underline{s}, \underline{y}) \in S \times Y$ with $\underline{p}_{s, \sigma(\underline{s})}^*(\underline{y}) > 0$ and

$$m_\ell \left(\underline{s}, \sigma \left(\underline{s} \right), \underline{y} \right) =: \ell > 0.$$

For any $t \in \mathbb{N}$, $\sigma \in A^S$, and $a^t \in A^t$ let $\pi^{(a^t, \sigma)} \in A^H$ be the policy that prescribes action a_τ at period $\tau \leq t$ and action $\sigma(s_\tau)$ at all periods $\tau > t$, and let $\mathbb{P}_{a^t, \sigma}$ be the probability distribution induced by $\pi^{(a^t, \sigma)}$. Throughout the appendix, we let \mathbf{R} , denote the random variable corresponding to the subset of periods recalled after (h_t, s_{t+1}) , while we continue to use the nonbold version, R , for its realizations.

The next lemma shows that if σ is a limit strategy, then for some time t , there is an action sequence a^t such that if the agent plays a^t in the first t periods and then σ afterward, there is positive probability that the induced sequence of beliefs makes σ optimal at all periods $\tau \geq t + 1$.

LEMMA A4. Let $\sigma \in A^S$. If for every $t \in \mathbb{N}$, every $a^t \in A^t$, and every optimal policy $\tilde{\pi}$, $\mathbb{P}_{a^t, \sigma}[\sigma(s_{\tau+1}) = \tilde{\pi}(h_\tau(\mathbf{R}_\tau))(s_{\tau+1}) \text{ for all } \tau \geq t] = 0$, then σ is not a limit strategy.

Fix an arbitrary outcome \underline{y} . Let $n_{s, a, t}$ be the number of times the signal-action pair $(s, a) \in S \times A$ occurred in periods $\{1, \dots, t\}$ and $g_{s, a, t}$ be the frequency of outcomes that realized after signal s and action a until period t , that is,

$$g_{s, a, t}(\underline{y}) = \frac{1}{n_{s, a, t}} \sum_{i=1}^t \mathbf{1}_{\{(s, a, y_i)\}}(s_i, a_i, y_i),$$

with $g_{s, a, t}(\underline{y}) = \mathbf{1}_{\{\underline{y}\}}(\underline{y})$ whenever $n_{s, a, t} = 0$. Similarly, let $\tilde{n}_{s, a, t}$ be the number of times the signal-action pair (s, a) is recalled in period $t + 1$. Also, let $f_{s, a, t}$ be the frequency of outcomes induced by signal s and action a that is recalled at period $t + 1$, with $f_{s, a, t}(\underline{y}) = \mathbf{1}_{\{\underline{y}\}}(\underline{y})$ whenever $\tilde{n}_{s, a, t} = 0$.

For every $(s', s, a) \in S^2 \times A$ and $\varepsilon > 0$ and $t \in \mathbb{N}$, let

$$D_t(s', s, a, \varepsilon) = \mathbf{1}_{(\varepsilon, +\infty)} \left(\left\| \frac{\tilde{n}_{s, a, t}}{n_{s, a, t}} f_{s, a, t}(\cdot) - m_\ell(s, a, \cdot) g_{s, a, t}(\cdot) \right\|_\infty \right)$$

be an indicator function that is 1 if at period t there is a deviation of more than ε between the recalled empirical frequency given s, a and the m_ℓ -memory distorted version of the true empirical frequency. The next lemma shows that it is impossible to have infinitely many periods t where an action-signal pair with realized frequency larger than γ at t has this sort of deviation.

LEMMA A5. For every $\pi \in A^H$, $(s', s, a) \in S^2 \times A$, and $\varepsilon, \gamma > 0$,

$$\mathbb{P}_\pi \left[\sum_{t=1}^{\infty} D_t(s', s, a, \varepsilon) \mathbf{1}_{(\gamma, \infty)} \left(\frac{n_{s, a, t}}{t} \right) \mathbf{1}_{(s')}(s_{t+1}) < \infty \right] = 1. \quad (7)$$

A2. Proofs

A2.1. Proof of Theorem 1

Suppose toward a contradiction that σ is a limit strategy under the optimal policy π , but not a selective-memory equilibrium. By lemma A2 there are $s' \in S$ and $\varepsilon, C \in \mathbb{R}_{++}$ such that for all $\nu \in \Delta(\Theta)$

$$\frac{\nu(\Theta_s^m(\sigma, \varepsilon))}{1 - \nu(\Theta_s^m(\sigma, \varepsilon))} > C \Rightarrow \sigma(s') \notin \text{BR}(s', \nu), \tag{8}$$

and in particular

$$\sigma(s') \notin \text{BR}(s', \Delta(\Theta_s^m(\sigma))). \tag{9}$$

Fix this s' throughout the rest of the proof.

If $\sigma(s') \notin \text{BR}(s', \nu)$ for all $\nu \in \Delta(\Theta)$, we immediately reach a contradiction by definition of optimal policy, since by Kolmogorov 0–1 law (see, e.g., theorem 8.4.4 in Dudley 2018) signal s' will realize infinitely many times \mathbb{P}_π -a.s.

If $\sigma(s') \in \text{BR}(s', \Delta(\Theta))$, equation (9) and lemma A3 imply that there is an experience $(\underline{s}, \sigma(\underline{s}), \underline{y})$ that has objective positive probability under σ and is recalled with positive probability ℓ under signal s' . Now fix an objective history $h_t = (s^t, a^t, y^t) \in H_t$ that has positive probability under an optimal policy π , that is, $\mathbb{P}_\pi[h_t] > 0$. We will show that if the agent plays σ at every period after h_t , $\mathbb{P}_{a', \sigma}$ -a.s. the belief $\mu_\tau(\cdot | h_\tau(\mathbf{R}_\tau))$ reaches a region where no optimal policy prescribes $\sigma(s')$ after signal s' , that is, $\sigma(s') \notin \text{BR}(s', \mu_\tau(\cdot | h_\tau(\mathbf{R}_\tau)))$. By lemma A4, this is enough to obtain the desired conclusion.

By the strong law of large numbers, for every $(s, a, y) \in S \times A \times Y$,

$$\lim_{\tau \rightarrow \infty} f(h_\tau)(s, a, y) = \begin{cases} \zeta(s) p_{s,a}^*(y) & \text{if } a = \sigma(s), \\ 0 & \text{otherwise;} \end{cases} \quad \mathbb{P}_{a', \sigma}\text{-a.s. on the cylinder } h_t.$$

Let $\tilde{p}(\sigma, s') \in \Delta(S \times A \times Y)$ be the induced distribution over remembered experiences

$$\tilde{p}(\sigma, s')(s, a, y) = \begin{cases} \frac{\zeta(s) m_\zeta(s, \sigma(s), y) p_{s, \sigma(s)}^*(y)}{\sum_{\hat{s} \in Y, \hat{s} \in S} \zeta(\hat{s}) m_\zeta(\hat{s}, \sigma(\hat{s}), \hat{y}) p_{\hat{s}, \sigma(\hat{s})}^*(\hat{y})} & \text{if } a = \sigma(s), \\ 0 & \text{otherwise.} \end{cases}$$

For every 2 periods $\tau' > \tau$ and $R_{\tau'} \subseteq \{1, \dots, \tau'\}$, the probability of recalling $R_{\tau'}$ at time τ' conditional on the objective history $h_{\tau'}$ is independent of the recalled periods R_τ at period τ , that is, $\mathbb{P}_{a', \sigma}[\mathbf{R}_{\tau'} = R_{\tau'}, s_{\tau'+1} | h_{\tau'}] = \mathbb{P}_{a', \sigma}[\mathbf{R}_{\tau'} = R_{\tau'}, s_{\tau'+1} | h_{\tau'}, \mathbf{R}_\tau = R_\tau]$. The next claim shows that for every $k \in \mathbb{N}$, $\mathbb{P}_{a', \sigma}$ -a.s. there is a $\tau > t$ such that $s_{\tau+1} = s'$, and the number of periods recalled after h_τ , $s_{\tau+1}$ is larger than k . It is a variation of the Borel-Cantelli lemma based on conditional instead of unconditional probabilities.

To state the claim, for every $t \in \mathbb{N}$, let E_t denote the event that either or both of $|\mathbf{R}_t| \leq k$ and $s_{t+1} \neq s'$ hold.

CLAIM 1. For all $\hat{\tau} \in \mathbb{N}$ and $k \in \mathbb{N}$, $\mathbb{P}_{a', \sigma}[\bigcap_{\tau \geq \hat{\tau}} E_\tau] = 0$.

Proof. For every $\tau \in \mathbb{N}$ and $h = (s_i, a_i, y_i)_{i=1}^\tau$, let $N(h) = \sum_{i=1}^\tau \mathbb{1}_{\{(s, \sigma(s), y)\}}(s_i, a_i, y_i)$ be the number of times $(\underline{s}, \sigma(\underline{s}), \underline{y})$ occurs between periods 1 and τ . For any $j \in \mathbb{N}$, we have

$$\begin{aligned}
\mathbb{P}_{a',\sigma}[\cap_{\tau \in \{\hat{\tau}, \dots, \hat{\tau}+j\}} E_\tau] &= \prod_{\tau=\hat{\tau}}^{\hat{\tau}+j} \mathbb{P}_{a',\sigma}(E_\tau | E_1, \dots, E_{\tau-1}) = \prod_{\tau=\hat{\tau}}^{\hat{\tau}+j} \sum_{h \in H_t} \mathbb{P}_{a',\sigma}(h) \mathbb{P}_{a',\sigma}(E_\tau | E_1, \dots, E_{\tau-1}, h) \\
&= \prod_{\tau=\hat{\tau}}^{\hat{\tau}+j} \sum_{h \in H_t} \mathbb{P}_{a',\sigma}(h) (1 - \mathbb{P}_{a',\sigma}[|\mathbf{R}_\tau| > k, s_{\tau+1} = s' | (|\mathbf{R}_\tau| \leq k, \dots, |\mathbf{R}_{\tau-1}| \leq k, h)]) \\
&= \prod_{\tau=\hat{\tau}}^{\hat{\tau}+j} \sum_{h \in H_t} \mathbb{P}_{a',\sigma}(h) (1 - \mathbb{P}_{a',\sigma}[|\mathbf{R}_\tau| > k, s_{\tau+1} = s' | h]) \\
&\leq \prod_{\tau=\hat{\tau}}^{\hat{\tau}+j} (\mathbb{P}_{a',\sigma}(\{h \in H_\tau : N(h) \leq k\}) + \sum_{h \in H_\tau : N(h) \geq k+1} \mathbb{P}_{a',\sigma}(h) (1 - \mathbb{P}_{a',\sigma}[|\mathbf{R}_\tau| > k, s_{\tau+1} = s' | h])) \\
&\leq \prod_{\tau=\hat{\tau}}^{\hat{\tau}+j} (\mathbb{P}_{a',\sigma}(\{h \in H_\tau : N(h) \leq k\}) + \sum_{h \in H_\tau : N(h) \geq k+1} \mathbb{P}_{a',\sigma}(h) (1 - \ell^{k+1} \zeta(s'))) \\
&= \prod_{\tau=\hat{\tau}}^{\hat{\tau}+j} (\mathbb{P}_{a',\sigma}(\{h \in H_\tau : N(h) \leq k\}) + [1 - \mathbb{P}_{a',\sigma}(\{h \in H_\tau : N(h) \leq k\})] [1 - \ell^{k+1} \zeta(s')]) \\
&= \prod_{\tau=\hat{\tau}}^{\hat{\tau}+j} (1 - \ell^{k+1} \zeta(s') + \mathbb{P}_{a',\sigma}(\{h \in H_\tau : N(h) \leq k\}) \ell^{k+1} \zeta(s')),
\end{aligned}$$

where the second equality follows from the law of iterated expectations; the first inequality follows because for every $\tau \in \{\hat{\tau}, \dots, \hat{\tau} + j\}$,

$$\sum_{h \in H_\tau : N(h) \leq k} \mathbb{P}_{a',\sigma}(h) (1 - \mathbb{P}_{a',\sigma}[|\mathbf{R}_\tau| > k, s_{\tau+1} = s' | h]) \leq \sum_{h \in H_\tau : N(h) \leq k} \mathbb{P}_{a',\sigma}(h);$$

and the second inequality follows from the fact that if signal s' realizes and $(\underline{s}, \sigma(\underline{s}), \underline{y})$ appears at least $k + 1$ times in the objective history, the probability of recalling at least $k + 1$ events is not smaller than ℓ^{k+1} . Since $1 + x \leq e^x$ for all $x \in \mathbb{R}$, the last term is smaller than

$$\exp\left(\sum_{\tau=\hat{\tau}}^{\hat{\tau}+j} -\ell^{k+1} \zeta(s') + \mathbb{P}_{a',\sigma}(\{h \in H_\tau : N(h) \leq k\}) \ell^{k+1} \zeta(s')\right).$$

By definition, $(\underline{s}, \sigma(\underline{s}), \underline{y})$ has objective positive probability under σ , so there is $\hat{\tau} \in \mathbb{N}$ and $\beta \in (0, 1)$ such that for every $\tau \geq \hat{\tau}$, $\mathbb{P}_{a',\sigma}(\{h \in H_\tau : N(h) \leq k\}) < \beta < 1$. Thus,

$$\begin{aligned}
&\lim_{j \rightarrow \infty} \mathbb{P}_{a',\sigma}[\cap_{\tau \in \{\hat{\tau}, \dots, \hat{\tau}+j\}} E_\tau] \\
&\leq \lim_{j \rightarrow \infty} \exp\left(\sum_{\tau=\hat{\tau}}^{\hat{\tau}+j} -\ell^{k+1} \zeta(s') + \mathbb{P}_{a',\sigma}(\{h \in H_\tau : N(h) \leq k\}) \ell^{k+1} \zeta(s')\right) = 0,
\end{aligned}$$

proving the claim: for all $\hat{\tau} \in \mathbb{N}$ and $k \in \mathbb{N}$, $\mathbb{P}_{a',\sigma}[\cap_{\tau \geq \hat{\tau}} E_\tau] = 0$. QED

By the previous claim, for every $k \in \mathbb{N}_{++}$, $\mathbb{P}_{a',\sigma}$ -a.s.

$$\exists \tau > t: \quad s_{\tau+1} = s' \quad \text{and} \quad |\mathbf{R}_\tau| > k \quad (10)$$

CLAIM 2. We have that

$$\mathbb{P}_{a',\sigma}[\{ \tau : s_{\tau+1} = s', \|\hat{f}(h_\tau, \mathbf{R}_\tau) - \check{p}(\sigma, s')\|_\infty > \varepsilon \}] = 0. \quad (11)$$

Proof. Let $\tau \in \mathbb{N}$, $h_\tau \in H_\tau$, and $\varepsilon > 0$. By the Chernoff inequality (see, e.g., Boucheron, Lugosi, and Massart 2013, 23–24),

$$\begin{aligned} & \mathbb{P}_{a',\sigma} \left[\left| \frac{\mathbf{R}_\tau}{\tau} \hat{f}(h_\tau, \mathbf{R}_\tau) - \tilde{p}(\sigma, s') \right| > \varepsilon | (h_\tau, s') \right] \\ & \leq 2|Y| \exp \left(-\varepsilon \tau \left[\log 1/2 - \frac{\log(1/2 + \varepsilon) + \log(1/2 - \varepsilon)}{2} \right] \right). \end{aligned}$$

Since

$$\sum_{k=1}^{\infty} 2 \exp(-\varepsilon k [\log 1/2 - \psi(\varepsilon)]) < \infty,$$

the result follows by the Borel-Cantelli lemma. QED

We show that $\nu(\Theta_s^m(\sigma, \varepsilon)) / [1 - \nu(\Theta_s^m(\sigma, \varepsilon))] > C$ on the histories where conditions (10) and (11) are satisfied. Since they hold $\mathbb{P}_{a',\sigma}$ -a.s., the result follows by (8).

Let $\varepsilon' \in (0, \varepsilon)$ and $\kappa \in \mathbb{R}_{++}$ be such that for all $(s, a) \in S \times A$ and $p \in \Theta_s^m(\sigma, \varepsilon')$, $p_{s,a} \gg p_{s,a}^*$,

$$\frac{\kappa}{2} > \sup_{p' \in \Theta_s^m(\sigma, \varepsilon')} \sum_{s \in S} \zeta(s) \sum_{y \in Y} p_{s,\sigma(s)}^*(y) m_{s'}(s, \sigma(s), y) \log p'_{s,\sigma(s)}(y)$$

and

$$\kappa < \inf_{p \in \Theta_s^m(\sigma, \varepsilon')} \sum_{s \in S} \zeta(s) \sum_{y \in Y} p_{s,\sigma(s)}^*(y) m_{s'}(s, \sigma(s), y) \log p'_{s,\sigma(s)}(y),$$

where their existence is guaranteed by the continuity in p of the memory-weighted log-likelihood and assumption 2. So, by lemma A1

$$\begin{aligned} & \frac{\mu(\Theta_s^m(\sigma, \varepsilon) | h_\tau(R_\tau))}{1 - \mu(\Theta_s^m(\sigma, \varepsilon) | h_\tau(R_\tau))} \geq \frac{\mu(\Theta_s^m(\sigma, \varepsilon') | h_\tau(R_\tau))}{1 - \mu(\Theta_s^m(\sigma, \varepsilon') | h_\tau(R_\tau))} \\ & \geq \frac{\mu(\Theta_s^m(\sigma, \varepsilon'))}{1 - \mu(\Theta_s^m(\sigma, \varepsilon'))} \exp \left(|R_\tau| \left(\inf_{p \in \Theta_s^m(\sigma, \varepsilon')} L(\hat{f}(h_\tau, R_\tau) \| p) - \sup_{p \in \Theta_s^m(\sigma, \varepsilon)} L(\hat{f}(h_\tau, R_\tau) \| p) \right) \right). \end{aligned}$$

The last expression goes to $+\infty$ as $\tau \rightarrow \infty$, since (i) $|R_\tau| \rightarrow \infty$ by equation (10), and (ii) by the definitions of κ and ε' as well as equation (11) we have

$$\begin{aligned} & \liminf_{\tau \rightarrow \infty} \inf_{p \in \Theta_s^m(\sigma, \varepsilon') | (s,a,y)} \sum \hat{f}(h_\tau, R_\tau)(s, a, y) \log(p_{s,a}(y)) \\ & \quad - \sup_{p \in \Theta_s^m(\sigma, \varepsilon) | (s,a,y)} \sum \hat{f}(h_\tau, R_\tau)(s, a, y) \log(p_{s,a}(y)) \\ & = \inf_{p \in \Theta_s^m(\sigma, \varepsilon') | s \in S} \sum_{s \in S} \zeta(s) \sum_{y \in Y} m_{s'}(s, \sigma(s), y) p_{s,\sigma(s)}^*(y) \log(p_{s,\sigma(s)}(y)) \\ & \quad - \sup_{p \in \Theta_s^m(\sigma, \varepsilon) | s \in S} \sum_{s \in S} \zeta(s) \sum_{y \in Y} m_{s'}(s, \sigma(s), y) p_{s,\sigma(s)}^*(y) \log(p_{s,\sigma(s)}(y)) \\ & > \frac{\kappa}{2} > 0. \quad \text{QED} \end{aligned}$$

A2.2. Proof of Theorem 2

First, we introduce some notation. Let

$$\underline{m} := \min_{(s', s, a) \in S^2 \times A} \sum_{y \in Y} p_{s', a}^*(y) m_{s'}(s, a, y) > 0,$$

and for every $(s, s', a, y) \in S^2 \times A \times Y$, let

$$\bar{p}_{s, a}(y|s') := \frac{m_{s'}(s, a, y) p_{s, a}^*(y)}{\sum_{y' \in Y} m_{s'}(s, a, y') p_{s, a}^*(y')}$$

denote the memory-adjusted version of the data-generating process.

1. *Part i.* Now we will prove the first part of the theorem, namely that

$$\mathbb{P}_\pi \left[\lim_{t \rightarrow \infty} \mu(B_c(\hat{p}) | h_t(\mathbf{R}_t)) = 1 \right] = 1.$$

The first step is to characterize the distribution of outcomes given the realized signals and actions. Consider the stochastic processes $(\mathbf{X}_t^{(s, \hat{a}, \hat{y})})_{(s, \hat{a}, \hat{y}) \in S \times A \times Y, t \in \mathbb{N}}$ defined by

$$\mathbf{X}_t^{(s, \hat{a}, \hat{y})} = (\mathbf{1}_{\{\hat{y}\}}(y_t) - p_{s, \hat{a}}^*(\hat{y})) \mathbf{1}_{\{(s, \hat{a})\}}(s_t, a_t) \quad \forall t \in \mathbb{N}.$$

These stochastic processes correspond to the deviation of the number of times each y has appeared from their expected frequency given the signal realized and action chosen. They are measurable with respect to the filtration $(\mathcal{F}_t)_{t \in \mathbb{N}}$ generated by the stochastic process of histories $(h_t)_{t \in \mathbb{N}}$. These processes are not i.i.d., as previous outcome realizations affect current period choices, but for each $(s, a, y) \in S \times A \times Y$, $\mathbb{E}[\mathbf{X}_t^{(s, a, y)} | \mathcal{F}_{t-1}] = 0$. Consequently, for each $(s, a, y) \in S \times A \times Y$, $(\mathbf{X}_t^{(s, a, y)})_{t \in \mathbb{N}}$ is a mixingale difference sequence, and from the strong law of large numbers for mixingale sequences (see theorem 2.7 in Hall and Heyde 2014 for the version that applies here), $\lim_{n \rightarrow \infty} (1/n) \sum_{t=1}^n \mathbf{X}_t^{(s, a, y)} = 0$, \mathbb{P}_π -a.s. Recall that $n_{s, a, t}$ is the number of periods before t in which the signal was s and the action a was played, and $g_{s, a, t}$ is the empirical distribution over outcomes in these periods. Moreover,

$$\frac{1}{n} \sum_{t=1}^n \mathbf{X}_t^{(s, a, y)} = \frac{n_{s, a, t}}{n} (g_{s, a, t}(y) - p_{s, a}^*(y)),$$

which implies that for every $\hat{\varepsilon} > 0$, $\gamma > 0$, \mathbb{P}_π -a.s.

$$\limsup_{n \rightarrow \infty} \left(\mathbf{1}_{[\gamma, \infty)} \left(\frac{n_{s, a, t}}{n} \right) \cdot \mathbf{1}_{[\hat{\varepsilon}, \infty)} (\|g_{s, a, t} - p_{s, a}^*\|_\infty) \right) = 0. \quad (12)$$

Recall that $\tilde{n}_{s, a, t}$ is the number of times the signal-action pair (s, a) is recalled at time t . By lemma A5, for every $(s, s', a) \in S^2 \times A$, $\hat{\varepsilon} > 0$, and $\gamma > 0$,

$$\mathbb{P}_\pi \left[\sum_{t=1}^{\infty} D_t(s', s, a, \hat{\varepsilon}) \mathbf{1}_{(\gamma, \infty)} \left(\frac{n_{s, a, t}}{t} \right) \mathbf{1}_{(s')}(s_{t+1}) < \infty \right] = 1. \quad (13)$$

In the set identified by equation (13), which has \mathbb{P}_π probability 1, for all $(s, a) \in S \times A$ and every time subsequence $(t_i)_{i \in \mathbb{N}}$ where $n_{s,a,t_i}/t_i > \gamma$,

$$\begin{aligned} & \left\| f_{s,a,t}(\cdot) - \frac{m_{s'}(s, a, \cdot)g_{s,a,t}(\cdot)}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')} \right\|_\infty \\ & \leq \left\| f_{s,a,t}(\cdot) - \frac{\tilde{n}_{s,a,t}}{n_{s,a,t}} \frac{f_{s,a,t}(\cdot)}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')} \right\|_\infty \\ & \quad + \left\| \frac{\tilde{n}_{s,a,t}}{n_{s,a,t}} \frac{f_{s,a,t}(\cdot)}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')} - \frac{m_{s'}(s, a, \cdot)g_{s,a,t}(\cdot)}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')} \right\|_\infty \\ & \leq \left| \frac{n_{s,a,t}}{\tilde{n}_{s,a,t}} - \frac{1}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')} \right| \left\| \frac{\tilde{n}_{s,a,t}}{n_{s,a,t}} f_{s,a,t}(\cdot) \right\|_\infty \\ & \quad + \frac{1}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')} \left\| \frac{\tilde{n}_{s,a,t}}{n_{s,a,t}} f_{s,a,t}(\cdot) - m_{s'}(s, a, \cdot)g_{s,a,t}(\cdot) \right\|_\infty \\ & \leq \left| \frac{1}{\sum_{y' \in Y} (\tilde{n}_{s,a,t})/(n_{s,a,t})f_{s,a,t}(\cdot)} - \frac{1}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')} \right| \\ & \quad + \frac{\hat{\varepsilon}}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')} \\ & \leq \frac{1}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y') - |Y|\varepsilon} - \frac{1}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')} \\ & \quad + \frac{\hat{\varepsilon}}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')} \\ & = \frac{|Y|\hat{\varepsilon}}{\left(\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y') - |Y|\hat{\varepsilon} \right) \sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')} \\ & \quad + \frac{\hat{\varepsilon}}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')}. \end{aligned}$$

Moreover, $\lim_{i \rightarrow \infty} \sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y') = \sum_{y' \in Y} m_{s'}(s, a, y')p_{s,a}^*(y')$, so that the last term in the display above converges to

$$\begin{aligned} & \frac{|Y|\hat{\varepsilon}}{\left(\sum_{y' \in Y} m_{s'}(s, a, y')p_{s,a}^*(y') - |Y|\hat{\varepsilon} \right) \sum_{y' \in Y} m_{s'}(s, a, y')p_{s,a}^*(y')} \\ & \quad + \frac{\hat{\varepsilon}}{\sum_{y' \in Y} m_{s'}(s, a, y')p_{s,a}^*(y')}; \end{aligned}$$

by assumption 3i, this can be taken arbitrarily small by choosing sufficiently small $\hat{\varepsilon}$.

Therefore, equation (13) implies that for every $(s, s', a) \in S^2 \times A$ and $\gamma > 0$,

$$\mathbb{P}_\pi \left[\left[t : \frac{n_{s,a,t}}{t} > \gamma, s_{t+1} = s' \text{ and } \left\| f_{s,a,t}(\cdot) - \frac{m_{s'}(s, a, \cdot)g_{s,a,t}(\cdot)}{\sum_{y' \in Y} m_{s'}(s, a, y')g_{s,a,t}(y')} \right\|_\infty > \hat{\varepsilon} \right] = \infty \right] = 0. \tag{14}$$

Observe that under data-generating process $q \in \Theta$, the log likelihood of any history $(s_i, a_i, y_i)_{i=1}^t$, $\tau \leq t$ that the agent might recall at time t can be rewritten as

$$\begin{aligned} \sum_{i=1}^{\tau} \log q_{s_i, a_i}(y_i) &= \sum_{(s, a) \in S \times A} \tilde{n}_{s, a, t} \sum_{y \in Y} f_{s, a, t}(y) \log q_{s, a}(y) \\ &= \sum_{(s, a) \in S \times A} \tilde{n}_{s, a, t} \left(-D_{\text{KL}}(f_{s, a, t}, q_{s, a}) + \sum_{y \in Y} f_{s, a, t}(y) \log f_{s, a, t}(y) \right), \end{aligned}$$

where $D_{\text{KL}}(q, q')$ denotes the Kullback-Leibler divergence between $q, q' \in \Delta(Y)$. Thus, for every $\varepsilon > 0$,

$$\frac{\mu(B_{\varepsilon}(\hat{p}) | (s_i, a_i, y_i)_{i=1}^{\tau})}{1 - \mu(B_{\varepsilon}(\hat{p}) | (s_i, a_i, y_i)_{i=1}^{\tau})} = \frac{\int_{B_{\varepsilon}(\hat{p})} \exp\left(-\sum_{(s, a) \in S \times A} \tilde{n}_{s, a, t} D_{\text{KL}}(f_{s, a, t}, p_{s, a})\right) d\mu(p)}{\int_{\Theta \setminus B_{\varepsilon}(\hat{p})} \exp\left(-\sum_{(s, a) \in S \times A} \tilde{n}_{s, a, t} D_{\text{KL}}(f_{s, a, t}, q_{s, a})\right) d\mu(q)}.$$

By assumption 3, \hat{p} maximizes the log likelihood and hence minimizes the divergence from \bar{p} after every signal-action pair. Thus, because D_{KL} is jointly lower semicontinuous (see, e.g., lemma 1.4.3 in Dupuis and Ellis 2011), there is $\bar{\varepsilon} > 0$ such that for all $(s', s, a) \in S^2 \times A$ and $q \in \Theta \setminus B_{\varepsilon}(\hat{p})$, $D_{\text{KL}}(\bar{p}_{s, a}(\cdot | s'), q_{s, a}) > \bar{\varepsilon} + D_{\text{KL}}(\bar{p}_{s, a}(\cdot | s'), \hat{p}_{s, a})$.

By definition, $\bar{p}_{s, a}(\cdot | s') \gg f_{s, a, t}$ at every $t \in \mathbb{N}$ such that $s_{t+1} = s'$. Thus, by equations (12) and (14), for every $\gamma > 0$, \mathbb{P}_{π} -a.s.,

$$\begin{aligned} \liminf_{t \rightarrow \infty} (1/t) \sum_{(s, a) \in S \times A} \tilde{n}_{s, a, t} D_{\text{KL}}(f_{s, a, t}, q_{s, a}) \\ \geq \liminf_{t \rightarrow \infty} (1/t) \sum_{(s, a) \in S \times A: n_{s, a, t} \leq \gamma} \tilde{n}_{s, a, t} D_{\text{KL}}(f_{s, a, t}, q_{s, a}) \\ \geq \liminf_{t \rightarrow \infty} (1/t) \sum_{(s, a) \in S \times A: n_{s, a, t} \leq \gamma} \tilde{n}_{s, a, t} (\bar{\varepsilon} + D_{\text{KL}}(\bar{p}_{s, a}(\cdot | s'), \hat{p}_{s, a})) \end{aligned} \quad (15)$$

for every $q \in \Theta \setminus B_{\varepsilon}(\hat{p})$.

Conversely, by assumption 2, we can choose $\varepsilon' < \varepsilon$ small enough that if $\|f_{s, a, t} - \bar{p}_{s, a}(\cdot | s')\|_{\text{TV}} \leq \varepsilon'$, and $p \in B_{\varepsilon'}(\hat{p})$ then

$$D_{\text{KL}}(f_{s, a, t}, p_{s, a}) \leq \frac{\bar{\varepsilon}}{2} + D_{\text{KL}}(\bar{p}_{s, a}(\cdot | s'), \hat{p}_{s, a})$$

and

$$K := \sup_{(s, s') \in S^2, a \in A, p \in B_{\varepsilon'}(\hat{p}), f \in \Delta(Y) : \bar{p}_{s, a}(\cdot | s') \gg f} D_{\text{KL}}(f, p_{s, a}) < \infty.$$

Therefore, for all $\beta \in (0, 1)$, \mathbb{P}_π -a.s.

$$\begin{aligned} \liminf_{t \rightarrow \infty} \frac{\mu(B_{\varepsilon'}(\hat{p})|h_t(\mathbf{R}_t))}{1 - \mu(B_{\varepsilon'}(\hat{p})|h_t(\mathbf{R}_t))} &= \liminf_{t \rightarrow \infty} \frac{\int_{B_{\varepsilon'}(\hat{p})} \exp\left(-\sum_{s \in S, a \in A} \tilde{n}_{s,a,t} D_{\text{KL}}(f_{s,a,t}, p_{s,a})\right) d\mu(p)}{\int_{\Theta \setminus B_{\varepsilon'}(\hat{p})} \exp\left(-\sum_{s \in S, a \in A} \tilde{n}_{s,a,t} D_{\text{KL}}(f_{s,a,t}, q_{s,a})\right) d\mu(q)} \\ &\geq \liminf_{t \rightarrow \infty} \frac{\int_{B_{\varepsilon'}(\hat{p})} \exp\left(-t\beta K - \sum_{s \in S, a \in A: \tilde{n}_{s,a,t}/t > \frac{\beta}{|S \times A|}} \tilde{n}_{s,a,t} D_{\text{KL}}(f_{s,a,t}, p_{s,a})\right) d\mu(p)}{\int_{\Theta \setminus B_{\varepsilon'}(\hat{p})} \exp\left(-\sum_{s \in S, a \in A: \tilde{n}_{s,a,t}/t > \frac{\beta}{|S \times A|}} \tilde{n}_{s,a,t} D_{\text{KL}}(f_{s,a,t}, q_{s,a})\right) d\mu(q)} \\ &\geq \lim_{t \rightarrow \infty} \frac{\mu(B_{\varepsilon'}(\hat{p}))}{1 - \mu(B_{\varepsilon'}(\hat{p}))} \exp\left(t\left[-\beta K + \frac{m}{2}(1 - \beta)\left(\varepsilon - \frac{\varepsilon'}{2}\right)\right]\right), \end{aligned}$$

where the last inequality follows from equations (13) and (15). For β small enough that $-\beta K + (m/2)(1 - \beta)[\varepsilon - (\varepsilon'/2)] > 0$, the right-hand side goes to infinity as t goes to infinity, so the left-hand side must also diverge, which shows that $\mathbb{P}_\pi[\lim_{t \rightarrow \infty} \mu(B_{\varepsilon'}(\hat{p})|h_t(\mathbf{R}_t)) = 1] = 1$. Since $\varepsilon' < \varepsilon$, this proves the first part of the theorem. In particular, for every $\varepsilon > 0$, the random variable T defined by

$$T := \sup\{t \in \mathbb{N} : \mu(B_{\varepsilon'}(\hat{p})|h_t(\mathbf{R}_t)) < 1 - \varepsilon\} \quad (16)$$

is \mathbb{P}_π -a.s. finite.

2. *Part ii.* To prove the second part of the theorem, note that assumption 3ii implies that $\Theta_s^m(\sigma) = \{\hat{p}\}$ for all $s \in S$ and $\sigma \in A^S$. Therefore, every selective-memory equilibrium must prescribe a best reply to a Dirac belief on \hat{p} after every signal. Since there is a unique best response to \hat{p} for every signal s , $\hat{\sigma}$ is the unique selective-memory equilibrium, and it is uniformly strict. By lemma A2, there is an ε such that $\hat{\sigma}(s)$ is the response to s for any belief ν that assigns probability at least $1 - \varepsilon$ to $B_\varepsilon(\hat{p})$. Since by equation (16) \mathbb{P}_π -a.s. there will be a finite time T (that can depend on the sample path) with $\mu(B_\varepsilon(\hat{p})|h_t(\mathbf{R}_t)) > 1 - \varepsilon$ for all $t > T$, the result follows. QED

A2.3. Proof of Proposition 1

We show that for every signal $s \in S$, only data-generating processes p for which $p_{s,\sigma(s)} = p_{s,\sigma(s)}^*$ are memory-weighted likelihood maximizers.

Fix $\hat{s} \in S$ and suppose that p is such that $p_{\sigma(\hat{s}),\hat{s}} \neq p_{\sigma(\hat{s}),\hat{s}}^*$. By the Gibbs inequality,

$$\sum_{y \in Y} p_{s,\sigma(s)}^*(y) \log p_{s,\sigma(s)}^*(y) \geq \sum_{y \in Y} p_{s,\sigma(s)}^*(y) \log p_{s,\sigma(s)}(y)$$

for all $s \in S$, with strict inequality for $s = \hat{s}$. This, together with $d(\hat{s}, \hat{s}) = 0$ and $\Phi(0) > 0$, implies that

$$\begin{aligned}
& \sum_{s \in S} \zeta(s) \sum_{y \in Y} m_\xi(s, \sigma(s), y) p_{s, \sigma(s)}^*(y) \log p_{s, \sigma(s)}(y) \\
&= \sum_{s \in S} \zeta(s) \Phi(d(s, \hat{s})) \sum_{y \in Y} p_{s, \sigma(s)}^*(y) \log p_{s, \sigma(s)}(y) \\
&< \sum_{s \in S} \zeta(s) \Phi(d(s, \hat{s})) \sum_{y \in Y} p_{s, \sigma(s)}^*(y) \log p_{s, \sigma(s)}^*(y) \\
&= \sum_{s \in S} \zeta(s) \sum_{y \in Y} m_\xi(s, \sigma(s), y) p_{s, \sigma(s)}^*(y) \log p_{s, \sigma(s)}^*(y),
\end{aligned}$$

proving that $p \notin \Theta_s^m(\sigma)$. QED

A2.4. Proof of Proposition 2

If assumption 3i is not satisfied, that is, no objectively possible outcome has a strictly positive probability of being remembered, beliefs remain constant over time and thus trivially converge. Otherwise, by theorem 2, we also know that beliefs converge. We first derive the long-run belief for $\tilde{m} \in \{m, m'\}$. Because the memory function \tilde{m} and the probability distribution over outcomes p^* are independent of the agent's action, this long-run belief is unique and independent of a , so we suppress the dependence of p and \tilde{m} on a .

Because $\Theta = \Delta(Y)$, for every σ the unique memory-weighted likelihood maximizers have the distribution

$$p^{\tilde{m}}(y) = \frac{\tilde{m}(y)p^*(y)}{\sum_{z \in Y} \tilde{m}(z)p^*(z)},$$

and by theorem 2 the beliefs concentrate on p^m . Moreover, $p^{m'}(y) = w(y)p^m(y)$, where $w(y) = f(y)(\sum_{z \in Y} m(z)p^*(z)/\sum_{z \in Y} m'(z)p^*(z))$ is nondecreasing, so $z \mapsto \sum_{x \leq z} (p^{m'}(x) - p^m(x)) = \sum_{x \leq z} p^m(x)(w(x) - 1)$ is quasiconvex. It equals 0 for $z < \min_{y \in Y} y$ and for $z \geq \max_{y \in Y} y$, so it is nonpositive for $z \in [\min_{y \in Y} y, \max_{y \in Y} y]$, and $p^{m'}$ dominates p^m in first-order stochastic dominance. Every limit action must be optimal given p^m for $\tilde{m} \in \{m, m'\}$ by theorem 1, so the agent's action must be weakly higher under m' than under m . QED

A2.5. Proof of Proposition 3

From theorem 2, we know that beliefs converge. Because (y_1, y_2) are subjectively independent conditional on the value of p , the learning problem decouples across the two dimensions. By proposition 2, the long-run belief about p is weakly higher than the true distribution p^* . The probability with which an outcome is remembered is independent of the second component, so the agent learns $\alpha p^*(y_2) + (1 - \alpha)q^*(y_2)$. They infer q to be

$$q(y_2) = \frac{\alpha p^*(y_2) + (1 - \alpha)q^*(y_2) - \alpha p(y_2)}{1 - \alpha}.$$

Thus, $q - q^* \equiv [\alpha/(1 - \alpha)](p^* - p)$, and as p is greater than p^* in first-order stochastic dominance, it follows that q is lower than q^* in first-order stochastic dominance. QED

A2.6. Proof of Proposition 4

If $a = 1$ is not a unitary self-confirming equilibrium, then the safe action $a = 0$ is preferred to the risky action $a = 1$, so $\sum_{y \in Y} v(y) \hat{p}_1^*(y) < v(y_0)$. Because the prior assigns positive probability to all distributions induced by action a_1 the unique memory-weighted likelihood maximizer \hat{p} under action 1 is such that

$$\hat{p}_1(y) := \frac{\hat{p}_1^*(y) h(|y - \bar{y}|)}{\sum_{z \in Y} \hat{p}_1^*(z) h(|z - \bar{y}|)}.$$

Therefore, if $a = 1$ is a selective-memory equilibrium when $m(y) = h(|y - \bar{y}|)$, then $v(y_0) \leq \sum_{y \in Y} \hat{p}_1(y) v(y)$. We prove that this cannot be the case by showing that the distribution \hat{p}_1 is second-order stochastically dominated by \hat{p}_1^* . To see this, observe that as \hat{p}_1^* is symmetric around \bar{y} and $h(|y - \bar{y}|)$ is symmetric around \bar{y} it follows that \hat{p}_1 is symmetric around \bar{y} . As h is increasing it follows that $\hat{p}_1 - \hat{p}_1^*$ changes its sign from positive to negative and back to positive so $\sum_{y \leq z} \hat{p}_1^*(y)$ and $\sum_{y \leq z} \hat{p}_1(y)$ cross only once, at $z = \bar{y}$. And since v is concave, theorem 3 and footnote 19 of Machina and Pratt (1997) imply that

$$\sum_{y \in Y} v(y) \hat{p}_1^*(y) \geq \frac{\sum_{y \in Y} \hat{p}_1^*(y) h(|y - \bar{y}|) v(y)}{\sum_{y \in Y} \hat{p}_1^*(y) h(|y - \bar{y}|)}$$

and that the risky action cannot be a selective-memory equilibrium. QED

A2.7. Proof of Proposition 5

Suppose toward a contradiction that σ is a limit strategy under the optimal policy π but not a selective-memory equilibrium. Then by lemma A2 there are $s' \in S$ and c , $C \in \mathbb{R}_{++}$ such that if

$$\frac{v(\Theta_{s'}^m(\sigma, c))}{1 - v(\Theta_{s'}^m(\sigma, c))} > C \quad \text{then} \quad \sigma(s') \notin \text{BR}(s', v). \quad (17)$$

Let $h_t = (s^t, a^t, y^t)$ be a history with positive probability under π . We show that if the agent plays the strategy $\pi^{(a', \sigma)}$, then a.s. the underinference belief $\mu^U(\cdot | (s^t, a^t, y^t))$ is asymptotically in a region where no optimal policy prescribes σ after signal s' . Since a.s. signal s' occurs infinitely many times, the same arguments as in lemma A4 show that this implies the desired conclusion.

By the strong law of large numbers,

$$\lim_{\tau \rightarrow \infty} f(h_\tau)(s, a, y) = \begin{cases} \zeta(s) p_{s,a}^*(y) & \text{if } a = \sigma(s), \\ 0 & \text{otherwise,} \end{cases} \quad \mathbb{P}_{a', \sigma} \text{-a.s.} \quad (18)$$

Next, we express the posterior as a function of the observed frequencies and show that it concentrates on the memory-weighted likelihood maximizers, so the result follows by equation (17). By assumption 2 and the continuity in p of the memory-weighted log-likelihood we can choose $\kappa, c' \in \mathbb{R}_{++}$ so that for all $(s, a) \in S \times A$ and $p \in \Theta_{s'}^m(\sigma, c')$, $p_{s,a} \gg p_{s,a}^*$,

$$\kappa/2 > \sup_{p' \in \Theta^*_t(\sigma, c')} \sum_{s \in S} \zeta(s) \sum_{y \in Y} p'_{s, \sigma(s)}^*(y) m_s(s, \sigma(s), y) \log p'_{s, \sigma(s)}(y)$$

and

$$\kappa < \inf_{p' \in \Theta^*_t(\sigma, c')} \sum_{s \in S} \zeta(s) \sum_{y \in Y} p'_{s, \sigma(s)}^*(y) m_s(s, \sigma(s), y) \log p'_{s, \sigma(s)}(y).$$

By equation (18) and the definition of κ and c' , a.s. on the cylinder h_t we have

$$\begin{aligned} K &:= \lim_{t \rightarrow \infty} \inf_{p' \in \Theta^*_t(\sigma, c')}(s, a, y) \sum f(h_t)(s, a, y) m(s, a, y) \log(p'_{s, a}(y)) \\ &\quad - \lim_{t \rightarrow \infty} \sup_{p' \in \Theta^*_t(\sigma, c')}(s, a, y) \sum f(h_t)(s, a, y) m(s, a, y) \log(p'_{s, a}(y)) \\ &= \inf_{p' \in \Theta^*_t(\sigma, c')} \sum_{s \in S} \zeta(s) \sum_{y \in Y} m_s(s, \sigma(s), y) p'_{s, \sigma(s)}^*(y) \log p'_{s, \sigma(s)}(y) \\ &\quad - \sup_{p' \in \Theta^*_t(\sigma, c')} \sum_{s \in S} \zeta(s) \sum_{y \in Y} m_s(s, \sigma(s), y) p'_{s, \sigma(s)}^*(y) \log p'_{s, \sigma(s)}(y) \\ &> \kappa/2 > 0. \end{aligned}$$

By lemma A1,

$$\begin{aligned} &\frac{\mu(\Theta^m_s(\sigma, c)|h_t)}{1 - \mu(\Theta^m_s(\sigma, c)|h_t)} \\ &\geq \frac{\mu(\Theta^m_s(\sigma, c')) \exp\left(-\sup_{p' \in \Theta^*_t(\sigma, c')} - \sum_{(s, a, y)} t f(h_t)(s, a, y) m(s, a, y) \log(p'_{s, a}(y))\right)}{(1 - \mu(\Theta^m_s(\sigma, c')) \exp\left(-\inf_{p' \in \Theta^*_t(\sigma, c)} - \sum_{(s, a, y)} t f(h_t)(s, a, y) m(s, a, y) \log(p'_{s, a}(y))\right))}, \end{aligned}$$

which goes to ∞ with t since $K > 0$. QED

A2.8. Proof of Proposition 6

To prove part i, let σ be a uniformly strict Berk-Nash equilibrium, and let p' be an arbitrary element of $\Theta^1(\sigma)$. Since σ is a uniformly strict Berk-Nash equilibrium, for all $s \in S$, $\{\sigma(s)\} = \text{BR}(s, \delta_{p'})$. Moreover, by the absolute continuity condition, $p'_{s, \sigma(s)}^*(y) = 0$ implies $p'_{s, \sigma(s)}(y) = 0$, so³⁸

$$K := \max_{(s, a, y) \in S \times A \times Y} \frac{p'_{s, a}(y)}{p'_{s, a}^*(y)} < \infty.$$

Define \tilde{m} as $\tilde{m}_s(s, a, y) = p'_{s, a}(y)/K p'_{s, a}^*(y)$ if $p'_{s, a}^*(y) > 0$ and $\tilde{m}_s(s, a, y) = 1/2$ otherwise. Then, for an agent with a full-support prior and memory function \tilde{m} the memory-weighted likelihood maximizers for strategy σ after signal s' are the elements of

³⁸ We use the convention that $0/0 = 0$.

$$\begin{aligned} & \operatorname{argmax}_{p \in \Delta(Y)^{S \times A}} \sum_{s \in S} \zeta(s) \sum_{y \in Y} \tilde{m}_s(s, \sigma(s), y) p_{s, \sigma(s)}^*(y) \log p_{s, \sigma(s)}(y) \\ &= \operatorname{argmax}_{p \in \Delta(Y)^{S \times A}} \sum_{s \in S} \zeta(s) \sum_{y \in Y} \frac{p'_{s, \sigma(s)}(y)}{K} \log p_{s, \sigma(s)}(y) = \operatorname{argmax}_{p \in \Delta(Y)^{S \times A}} \sum_{s \in S} \zeta(s) \sum_{y \in Y} p'_{s, \sigma(s)}(y) \log p_{s, \sigma(s)}(y). \end{aligned}$$

Thus, p' maximizes the memory-weighted likelihood for all $s' \in S$, so σ is a selective-memory equilibrium with a full-support prior.

Part ii, the converse direction, is trivial: take Θ' to be a singleton p such that for all $a \in A$ and $s \in S$, $p_{s, a}(y) = p'_{s, a}(y)$ for some $p' \in \Theta'_s(\sigma)$. QED

A3. Partial Naïveté

So far, we have assumed that agents treat the experiences they remember as if these were the only ones that happened. This section considers agents who are at least partially aware of their memory limitations. We suppose throughout this section that actions have no effect on the outcome distribution. We also assume that the agent either does not remember their actions or believes that they convey no information. Finally, we suppose that agents know the current period and so know how many observations they have forgotten. If the agent believes that they remember an occurrence of signal $s \in S$ and outcome $y \in Y$ with probability $\hat{m}(s, y) \in (0, 1]$ instead of the true probability $m(s, y)$, the subjective likelihood of recalling the periods R_t after (h_t, s') under data-generating process p is proportional to

$$\left[\sum_{s \in S} \zeta(s) \sum_{z \in Y} p_s(z) (1 - \hat{m}_s(s, z)) \right]^{t - |R_t|} \prod_{i \in R_t} \zeta(s_i) p_{s_i}(y_i) \hat{m}_{s_i}(s_i, y_i),$$

where $|R_t|$ is the number of events the agent remembers. Thus, the subjective log-likelihood equals

$$(t - |R_t|) \log \left[\sum_{s \in S} \zeta(s) \sum_{z \in Y} p_s(z) (1 - \hat{m}_s(s, z)) \right] + \sum_{y \in Y, s \in S, t \in R_t} \mathbf{1}_{\{(s, y)\}}(s_t, y_t) \log(p_s(y) \hat{m}_s(s, y)). \quad (19)$$

Because the expected value of $|R_t|/t$ is $\sum_{y \in Y} \sum_{s \in S} \zeta(s) p_s^*(y) m_s(s, y)$, (19) suggests the following generalization of the definition the memory-weighted likelihood maximizers:

$$\begin{aligned} \Theta_s^{m, \hat{m}}(\sigma) = \operatorname{argmax}_{p \in \Theta} & \left(1 - \sum_{s \in S} \sum_{y \in Y} m_s(s, y) \zeta(s) p_s^*(y) \right) \log \left(1 - \sum_{s \in S} \sum_{y \in Y} \zeta(s) p_s(y) \hat{m}_s(s, y) \right) \\ & + \sum_{y \in Y} \sum_{s \in S} m_s(s, y) \zeta(s) p_s^*(y) \log(\hat{m}_s(s, y) p_s(y)). \end{aligned}$$

DEFINITION 6. A selective-memory equilibrium for a partially naïve agent is a strategy σ such that for every $s \in S$ there exists a belief $v \in \Theta_s^{m, \hat{m}}(\sigma)$ with $(s, v) \in \operatorname{BR}(s, v)$.

For an agent who is aware of their own forgetfulness but not aware that their memory is selective, that is, who believes that their memory function m is constant, $\Theta_s^{m, \hat{m}} = \Theta_s^m$ and the selective-memory equilibria of a partially and fully naïve

agent coincide.³⁹ This shows that what matters for our results is that the agent is unaware that their memory is selective, not that they are unaware of their forgetfulness. At the other extreme, if agents are fully aware of their memory function, selective-memory equilibrium reduces to unitary self-confirming equilibrium because $\delta_{p^*} \in \Theta_s^{m,m}$.

The next result, whose proof is omitted, follows by observing that for a partially naïve agent, the posterior probability of C after an objective history (h_t, s_{t+1}) when the recalled periods are R_t is

$$\frac{\int_C \left(\prod_{\tau \in R_t} \hat{m}_{s_{\tau+1}}(s_\tau, y_\tau) p_{s_\tau}(y_\tau) \right) \left(1 - \sum_{s \in S} \sum_{y \in Y} \zeta(s) p_s(y) \hat{m}_{s_{\tau+1}}(s, y) \right)^{t-|R_t|} d\mu(p)}{\int_\Theta \left(\prod_{\tau \in R_t} \hat{m}_{s_{\tau+1}}(s_\tau, y_\tau) p_{s_\tau}(y_\tau) \right) \left(1 - \sum_{s \in S} \sum_{y \in Y} \zeta(s) p_s(y) \hat{m}_{s_{\tau+1}}(s, y) \right)^{t-|R_t|} d\mu(p)},$$

and then using an argument analogous to the proof of theorem 1.

PROPOSITION 7. When the agent is partially naïve, every limit strategy is a selective-memory equilibrium.

Moreover, as with notions of partial naïveté in cursed equilibrium and quasi-hyperbolic discounting, one can define a parametric notion of partial naïveté by assuming that $\hat{m}_{s_t}(s, y) = (1 - \alpha) + \alpha m_{s_t}(s, y)$. For $\alpha = 0$ the agent is fully naïve and unaware of their memory limitations. For $\alpha = 1$ the agent is sophisticated and understands their memory limitations, and so has correct long-run beliefs.

The next example shows that the amount of naïveté can have a nonmonotonic effect when there are more than two actions.

EXAMPLE 9. Suppose that the agent has three alternatives. They can either “do nothing,” $a = n$ with certain payoff of 0, do a quick job $a = k$ with payoff 1 if the job succeeds and -1 otherwise, or do a careful and time-consuming job $a = h$ at cost 0.6 that yields $1 - 0.6 = 0.4$ if the project succeeds and -1.6 otherwise. The probability of success in the quick job is some unknown $p \in [0, 1]$, while the probability of success in the careful job is $2p$ for $p \leq 0.45$ and $2p/11 + 9/11$ otherwise. The agent’s prior assigns positive probability to all $p \in [0, 1]$, where p is a reflection of the agent’s ability.

The true probability is $p^* = 0.2$, so $\mathbb{E}_{p^*}[u(n, \cdot)] > \mathbb{E}_{p^*}[u(k, \cdot)] > \mathbb{E}_{p^*}[u(h, \cdot)]$. Suppose that the agent has ego-boosting bias, in that they always recall successes and recall failures with probability 0.03. Here welfare is nonmonotone in the amount of partial naïveté of the agent. For a fully sophisticated agent, the unique selective-memory equilibrium is the objectively optimal n , while a naïve

³⁹ This is true in particular when the agent is fully naïve and \hat{m} is identically 1, even though the maximand becomes ill-defined. To see why, note that when $\hat{m}_{s_t}(\cdot) = k_{s_t}$ for some constants $k_{s_t} < 1$, the maximand is

$$\begin{aligned} & \left(1 - \sum_{s \in S} \sum_{y \in Y} m_{s_t}(s, y) \zeta(s) p_{s_t}^*(y) \right) \log(1 - k_{s_t}) + \sum_{y \in Y} \sum_{s \in S} m_{s_t}(s, y) \zeta(s) p_{s_t}^*(y) \log(k_{s_t}) \\ & + \sum_{y \in Y} \sum_{s \in S} m_{s_t}(s, y) \zeta(s) p_{s_t}^*(y) \log(p_{s_t}(y)). \end{aligned}$$

The first terms do not depend on p , so $\Theta_s^{m,\hat{m}} = \Theta_s^m$, and in particular complete naïveté is reached in the limit where all $k_{s_t} \rightarrow 1$.

agent has two selective-memory equilibria, n and k , with the latter sustained by the incorrect belief that their ability is so high that k is better than h . However, if the agent believes that they recall the failures with probability 0.12, playing the worst action h is a selective-memory equilibrium because the agent ends up believing that the probability of success is 0.5, which makes h the unique best reply.

References

- Adler, O., and A. Pansky. 2020. "A 'Rosy View' of the Past: Positive Memory Biases." In *Cognitive Biases in Health and Psychiatric Disorders*, edited by T. Aue and H. Okon-Singer, 139–71. London: Academic Press.
- Afrouzi, H., S. Y. Kwon, A. Landier, Y. Ma, and D. Thesmar. 2020. "Overreaction and Working Memory." Working Paper no. 27947 (October), NBER, Cambridge, MA.
- Alloy, L. B., and L. Y. Abramson. 1979. "Judgment of Contingency in Depressed and Nondepressed Students: Sadder but Wiser?" *J. Experimental Psychology: General* 108 (4): 441–85.
- Battigalli, P. 1987. "Comportamento Razionale ed Equilibrio nei Giochi e nelle Situazioni Sociali." PhD diss., Bocconi Univ.
- Battigalli, P., A. Francetich, G. Lanzani, and M. Marinacci. 2019. "Learning and Self-Confirming Long-Run Biases." *J. Econ. Theory* 183:740–85.
- Battigalli, P., and N. Geroso. 2021. "Information Flows and Memory in Games." IGER Working Paper no. 678 (December), Innocenzo Gasparini Inst. Econ. Res., Milan.
- Bénabou, R., and J. Tirole. 2002. "Self-Confidence and Personal Motivation." *Q.J.E.* 117:871–915.
- Benjamin, D. J. 2019. "Errors in Probabilistic Reasoning and Judgment Biases." In *Handbook of Behavioral Economics: Applications and Foundations*, edited by B. D. Bernheim, S. DellaVigna, and D. Laibson, 69–186. Amsterdam: North-Holland.
- Benjamin, D. J., S. A. Brown, and J. M. Shapiro. 2013. "Who Is 'Behavioral'? Cognitive Ability and Anomalous Preferences." *J. European Econ. Assoc.* 11:1231–55.
- Benoit, J.-P., and J. Dubra. 2011. "Apparent Overconfidence." *Econometrica* 79:1591–625.
- Benoit, J.-P., J. Dubra, and D. A. Moore. 2015. "Does the Better-than-Average Effect Show That People Are Overconfident?: Two Experiments." *J. European Econ. Assoc.* 13:293–329.
- Berk, R. H. 1966. "Limiting Behavior of Posterior Distributions When the Model Is Incorrect." *Ann. Math. Statist.* 37:51–58.
- Bordalo, P., K. Coffman, N. Gennaioli, F. Schwerkter, and A. Shleifer. 2021. "Memory and Representativeness." *Psychological Rev.* 128:71–85.
- Bordalo, P., N. Gennaioli, and A. Shleifer. 2020. "Memory, Attention, and Choice." *Q.J.E.* 135:1399–442.
- Boucheron, S., G. Lugosi, and P. Massart. 2013. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford: Clarendon.
- Boyle, P., L. Yu, A. Buchman, and D. Bennett. 2012. "Risk Aversion Is Associated with Decision Making among Community-Based Older Persons." *Frontiers Psychology* 3:205.
- Chammat, M., I. El Karoui, S. Allali, J. Hegège, K. Lehongre, D. Hasboun, M. Baulac, et al. 2017. "Cognitive Dissonance Resolution Depends on Episodic Memory." *Scientific Reports* 7:41320.

- Chew, S. H., W. Huang, and X. Zhao. 2020. "Motivated False Memory." *J.P.E.* 128 (10): 3913–39.
- Cokely, E. T., and C. M. Kelley. 2009. "Cognitive Abilities and Superior Decision Making under Risk: A Protocol Analysis and Process Model Evaluation." *Judgment and Decision Making* 4:20–33.
- Compte, O., and A. Postlewaite. 2004. "Confidence-Enhanced Performance." *A.E.R.* 94:1536–57.
- Cruciani, F., A. Berardi, S. Cabib, and D. Conversi. 2011. "Positive and Negative Emotional Arousal Increases Duration of Memory Traces: Common and Independent Mechanisms." *Frontiers Behavioral Neuroscience* 5:86.
- d'Acremont, M., W. Schultz, and P. Bossaerts. 2013. "The Human Brain Encodes Event Frequencies While Forming Subjective Beliefs." *J. Neuroscience* 33:10887–97.
- Danz, D., L. Vesterlund, and A. J. Wilson. 2022. "Belief Elicitation and Behavioral Incentive Compatibility." *A.E.R.* 122 (9): 2851–83.
- De Bondt, W. F., and R. H. Thaler. 1995. "Financial Decision-Making in Markets and Firms: A Behavioral Perspective." In *Handbooks in Operations Research and Management Science*, vol. 9, edited by R. A. Jarrow, V. Maksimovic and W. T. Ziemba, 385–410. Amsterdam: North-Holland.
- Dudley, R. M. 2018. *Real Analysis and Probability*. London: Chapman & Hall/CRC.
- Dupuis, P., and R. S. Ellis. 2011. *A Weak Convergence Approach to the Theory of Large Deviations*. Hoboken, NJ: Wiley.
- Elkin, R. A., and M. R. Leippe. 1986. "Physiological Arousal, Dissonance, and Attitude Change: Evidence for a Dissonance-Arousal Link and a 'Don't Remind Me' Effect." *J. Personality and Soc. Psychology* 51:55–65.
- Enke, B., F. Schwerter, and F. Zimmermann. 2022. "Associative Memory and Belief Formation." Working Paper no. 26664 (January), NBER, Cambridge, MA.
- Esponda, I. 2008. "Information Feedback in First Price Auctions." *RAND J. Econ.* 39:491–508.
- Esponda, I., and D. Pouzo. 2016. "Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models." *Econometrica* 84:1093–130.
- Esponda, I., E. Vespa, and S. Yuksel. 2023. "Mental Models and Learning: The Case of Base-Rate Neglect." Working paper.
- Frick, M., R. Iijima, and Y. Ishii. 2021. "Welfare Comparisons for Biased Learning." Cowles Found. Discussion Paper no. 2149R (February), Yale Univ., New Haven, CT.
- Fudenberg, D., and D. Kreps. 1988. "Learning, Experimentation and Equilibrium in Games." Working paper.
- . 1995. "Learning in Extensive-Form Games I. Self-Confirming Equilibria." *Games and Econ. Behavior* 8:20–55.
- Fudenberg, D., and G. Lanzani. 2023. "Which Misspecifications Persist?" *Theoretical Econ.* 18:1271–315.
- Fudenberg, D., G. Lanzani, and P. Strack. 2021. "Limit Points of Endogenous Misspecified Learning." *Econometrica* 89:1065–98.
- . 2023. "Pathwise Concentration Bounds for Misspecified Bayesian Beliefs." *Theoretical Econ.* 18:1585–622.
- . 2024. "Limited Memory, Learning, and Stochastic Choice." Working paper.
- Fudenberg, D., and D. K. Levine. 1993a. "Self-Confirming Equilibrium." *Econometrica* 61:523–45.
- . 1993b. "Steady State Learning and Nash Equilibrium." *Econometrica* 61:547–73.
- Gervais, S., and T. Odean. 2001. "Learning to Be Overconfident." *Rev. Financial Studies* 14:1–27.

- Gilboa, I., and D. Schmeidler. 2001. *A Theory of Case-Based Decisions*. Cambridge: Cambridge Univ. Press.
- Gilovich, T. 2008. *How We Know What Isn't So*. New York: Free Press.
- Gödker, K., P. Jiao, and P. Smeets. Forthcoming. "Investor Memory." *Rev. Financial Studies*.
- Goetzmann, W. N., A. Watanabe, and M. Watanabe. 2022. "Evidence on Retrieved Context: How History Matters." Working Paper no. 29849 (March), NBER, Cambridge, MA.
- Hall, P., and C. C. Heyde. 2014. *Martingale Limit Theory and Its Application*. London: Academic Press.
- Hastie, R., and B. Park. 1986. "The Relationship between Memory and Judgment Depends on Whether the Judgment Task Is Memory-Based or On-Line." *Psychological Rev.* 93:258–268.
- He, K. 2022. "Mislearning from Censored Data: The Gambler's Fallacy and Other Correlational Mistakes in Optimal-Stopping Problems." *Theoretical Econ.* 17:1269–312.
- He, K., and J. Libgober. 2023. "Evolutionarily Stable (Mis)specifications: Theory and Application." Working paper.
- Heidhues, P., B. Köszegi, and P. Strack. 2018. "Unrealistic Expectations and Misguided Learning." *Econometrica* 86:1159–214.
- Huffman, D., C. Raymond, and J. Shvets. 2022. "Persistent Overconfidence and Biased Memory: Evidence from Managers." *A.E.R.* 112:3141–75.
- James, W. 1890. *The Principles of Psychology*. New York: Holt.
- Jehiel, P. 2018. "Investment Strategy and Selection Bias: An Equilibrium Perspective on Overoptimism." *A.E.R.* 108 (6): 1582–97.
- . 2021. "Communication with Forgetful Liars." *Theoretical Econ.* 16:605–38.
- Jehiel, P., and J. Steiner. 2020. "Selective Sampling with Information-Storage Constraints." *Econ. J.* 130:1753–81.
- Kahana, M. J. 2012. *Foundations of Human Memory*. Oxford: Oxford Univ. Press.
- Kahneman, D., and A. Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47:263–92.
- Köszegi, B., G. Loewenstein, and T. Murooka. 2021. "Fragile Self-Esteem." *Rev. Econ. Studies* 89 (4): 2026–60.
- Lanzani, G. 2022. "Correlation Made Simple: Applications to Salience and Regret Theory." *Q.J.E.* 137:959–87.
- . 2024. "Dynamic Concern for Misspecification." Working paper.
- Loomes, G., and R. Sugden. 1982. "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty." *Econ. J.* 92:805–24.
- Lord, C. G., L. Ross, and M. R. Lepper. 1979. "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence." *J. Personality and Soc. Psychology* 37:2098–109.
- Machina, M., and J. Pratt. 1997. "Increasing Risk: Some Direct Constructions." *J. Risk and Uncertainty* 14:103–27.
- Malmendier, U., and S. Nagel. 2016. "Learning from Inflation Experiences." *Q.J.E.* 131:53–87.
- Malmendier, U., D. Pouzo, and V. Vanasco. 2020. "Investor Experiences and Financial Market Dynamics." *J. Financial Econ.* 136:597–622.
- Malmendier, U., and L. S. Shen. 2024. "Scarred Consumption." *American Econ. J. Microeconomics* 16 (1): 322–55.
- Matt, G. E., C. Vázquez, and W. K. Campbell. 1992. "Mood-Congruent Recall of Affectively Toned Stimuli: A Meta-Analytic Review." *Clinical Psychology Rev.* 12:227–55.

- Mayer, J. D., L. J. McCormick, and S. E. Strong. 1995. "Mood-Congruent Memory and Natural Mood: New Evidence." *Personality and Soc. Psychology Bull.* 21:736–46.
- Merkle, C., and M. Weber. 2011. "True Overconfidence: The Inability of Rational Information Processing to Account for Apparent Overconfidence." *Organizational Behavior and Human Decision Processes* 116:262–71.
- Mischel, W., E. B. Ebbesen, and A. M. Zeiss. 1976. "Determinants of Selective Memory about the Self." *J. Consulting and Clinical Psychology* 44:92–103.
- Mullainathan, S. 2002. "A Memory-Based Model of Bounded Rationality." *Q.J.E.* 117:735–74.
- Nagel, S., and Z. Xu. 2022. "Asset Pricing with Fading Memory." *Rev. Financial Studies* 35:2190–245.
- Phillips, L. D., and W. Edwards. 1966. "Conservatism in a Simple Probability Inference Task." *J. Experimental Psychology* 72:346–354.
- Procházka, J., M. Ovcari, and M. Durinik. 2020. "Sandwich Feedback: The Empirical Evidence of Its Effectiveness." *Learning and Motivation* 71:101649.
- Rabin, M., and J. L. Schrag. 1999. "First Impressions Matter: A Model of Confirmatory Bias." *Q.J.E.* 114:37–82.
- Rabin, M., and D. Vayanos. 2010. "The Gambler's and Hot-Hand Fallacies: Theory and Applications." *Rev. Econ. Studies* 77:730–78.
- Reder, L. M. 2014. *Implicit Memory and Metacognition*. New York: Psychology.
- Schacter, D. L. 2008. *Searching for Memory: The Brain, the Mind, and the Past*. New York: Basic Books.
- Schwartzstein, J. 2014. "Selective Attention and Learning." *J. European Econ. Assoc.* 12:1423–52.
- Schweizer, N., and M. de Vries. 2022. "Limited Attention and Dynamically Distorted Beliefs." Working paper.
- Shadlen, M. N., and D. Shohamy. 2016. "Decision Making and Sequential Sampling from Memory." *Neuron* 90:927–39.
- Sial, A. Y., J. R. Sydnor, and D. Taubinsky. 2023. "Biased Memory and Perceptions of Self-Control." Working Paper no. 30825 (January), NBER, Cambridge, MA.
- Svenson, O. 1981. "Are We All Less Risky and More Skillful than Our Fellow Drivers?" *Acta Psychologica* 47:143–48.
- Thomson, D. M., and E. Tulving. 1970. "Associative Encoding and Retrieval: Weak and Strong Cues." *J. Experimental Psychology* 86:255–262.
- Tulving, E., and D. L. Schacter. 1990. "Priming and Human Memory Systems." *Science* 247:301–6.
- Walters, D. J., and P. M. Fernbach. 2021. "Investor Memory of Past Performance Is Positively Biased and Predicts Overconfidence." *Proc. Nat. Acad. Sci. USA* 118:e2026680118.
- Wilson, A. 2014. "Bounded Memory and Biases in Information Processing." *Econometrica* 82:2257–94.
- Zimmermann, F. 2020. "The Dynamics of Motivated Beliefs." *A.E.R.* 110:337–61.