Selective Memory

Based on work with Giacomo Lanzani and Philipp Strack

Drew Fudenberg

University of Tokyo

February 2024

Motivation

- People have *limited* and *selective* memory: some events are more likely to be recalled.

- Several forms of selective memory have been observed/defined, e.g. positive memory bias, cognitive dissonance reduction, associativeness, and confirmatory bias.

- We study the implications of these biases when an agent chooses actions that maximize their expected utility and learns about their consequences.

Example

- An agent repeatedly observes the outcome of a test, which is either pass, $y = 1$, or fail, $y = 0$.
- Probability of passing is p^*.
- In an incentivized experiment, Zimmermann [2020] finds that
  - people are more likely to recall favorable tests.
  - they overestimate their probability of passing.
- Now suppose an agent with this sort of memory takes an action $a \in \{0, 1\}$, with $u(a, y) = a(y - z)$, so $a = 1$ is optimal if $p^* > z$.
- Then, in the selective memory equilibrium, the agent behaves as if they were misspecified and overconfident.

Related Work

- **Selective memory**: Mullainathan [2002], Bénabou and Tirole [2002], Jehiel [2018], and Heidhues, Koszegi, and Strack [2023]: Short-run effects of specific biases.

- **Endogenous misspecified learning**: We adapt some techniques from Esponda and Pouzo [2016], Esponda, Pouzo, and Yamamoto [2021], and Fudenberg, Lanzani, and Strack [2021]. Other related papers include Bohren and Hauser [2023], and Frick, Iijima, and Ishii [2023].

- **Selective Attention** Compte and Postlewaite [2004], Schwartzstein [2014], and Schweizer and De Vries [2022] study the long-run effects of a related error: instead of not recalling some outcomes, the agent might never observe them.

Setup

- In the general model, each period $t = 1, 2, ...$ the agent observes a signal from the finite set S and chooses an action a from the finite set A.

- This talk will supress the signals.

- Action a induces an objective probability distribution $p_a^* \in \Delta(Y)$ over the finite set of possible outcomes Y.

- Agent's payoff is $u(a, y)$.

- The agent has prior μ_0 over $p \in \Delta(Y)^A$, where $p_a(y)$ denotes the probability of outcome $y \in Y$ when action a is played:
  - Agent knows that the map from actions and signals to probability distributions depends only on their current action and the realized signal.
  - But they are uncertain about the outcome distributions each signal-action pair induces.
- Let Θ denote the support of the agent's prior.
- The prior is correctly specified if $p^* \in \Theta$.
- If $p^* \notin \Theta$, the prior is misspecified.

Maintained Assumption

- We assume the agent is correctly specified.
- This helps clarify the difference between selective memory and misspecified learning.
- Also assume that the true p^* and the models in the support of the agent's prior are mutually absolutely continuous. (Some results hold under weaker assumptions.)

Selective Memory

- In our work, memory is both
  - Probabilistic: Agents may remember different things in the same situations and
  - Selective: Some experiences are more likely to be recalled.
- These features are well documented, see e.g. Kahana [2012].
- We assume that the agent's memory of past periods is distorted by a memory function m that specifies the probability the agent remembers a past realization of (a, y).
- We also assume (except in an extension at the end of the talk) that each observation is remembered or not independent of the others.

Overview and General Framework

- Will present results from two papers.
- Start with concepts and notation used in both.
- In "Selective Memory Equilibrium," $m : A \times Y \to (0, 1]$ is independent of time, so the number of things the agent remembers goes to ∞ with t.
- In "Learning, Limited Memory, and Stochastic Choice," the probability of remembering eeach specific event decreases over time quickly enough that only finite many events are remembered.
- After history $h_t = (a_i, y_i)_{i=1}^t$, the recalled periods R_t are a random subset of $\{1, ..., t\}$.
- $h_t(R_t)$ is the subsequence of h_t that only contains the periods in R_t.

Beliefs

- The agent computes beliefs at time t by applying Bayes rule to $h_t(R_t)$ as if those were the only periods that were observed:

$$\mu(C|h_t(R_t)) = \frac{\int_C \prod_{\tau \in R_t} p_{a_\tau}(y_\tau) d\mu(p)}{\int_\Theta \prod_{\tau \in R_t} p_{,a_\tau}(y_\tau) d\mu(p)}.$$

- In particular, agents don't try to draw inferences about un-recalled events from recalled past actions.

- See e.g. Reder [2014] for evidence that agents are often näive about their selective memory and do not make inferences about their forgotten observations from the actions they remember taking.

Agent's Behavior

- Belief μ determine the actions that maximize the current-period subjective expected utility:

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

- A policy π specifies an action for every recalled history.

- We assume the agent acts myopically: For all h_t,

$$\pi(h_t) \in BR(\mu(\cdot | h_t)).$$

Memory Weighted Likelihood Maximizers

- If a standard Bayesian agent always plays action a, their beliefs concentrate on the likelihood maximizers given the true distribution p^*:

$$\text{argmax}_{p\in\Theta} \sum_{y\in Y} p_a^*(y) \log p_a(y).$$

- With selective memory, the analog of the true long run distribution p^* is the memory-weighted outcome distribution p^m,

$$p^m(a,y) := \sum_{y\in Y} m(a,y) p_a^*(y).$$

- And the analog of the maximum likelihood estimate is the set of memory-weighted likelihood maximizers

$$\Theta^m(a) := \text{argmax}_{p\in\Theta} \sum_{y} m(a,y) p_a^*(y) \log p_a(y).$$

Mixed Actions

- “Selective Memory Equillibrium” focused on pure equilibria.
- ”Limited Memory...” extends to mixed actions, as mixed actions are needed to capture the limit of long but finite memories.
- The memory-weighted likelihood maximizers with mixed actions just use the avergage memory weighted outcome distribution

$$\sum_{a \in A} \alpha(a) \sum_{y \in Y} m(a, y) p_a^*(y)$$

in place of

$$\sum_{y \in Y} m(a, y) p_a^*(y).$$

Definition

- A pure selective memory equilibrium is an action a such that there is $\nu \in \Delta(\Theta^m(\delta_a))$ such that $a \in BR(\nu)$.
- A unitary-beliefs selective memory equilibrium is a $\alpha \in \Delta(A)$ such that there is $\nu \in \Delta(\Theta^m(\alpha))$ such that $a \in BR(\nu)$ for all $a \in \mathrm{supp}(\alpha)$.
- A heterogeneous-beliefs selective memory equilibrium is a $\alpha \in \Delta(A)$ such that for all $a \in \mathrm{supp}(\alpha)$ there is $\nu^a \in \Delta(\Theta^m(\alpha))$ such that $a \in BR(\nu^a)$.

- These equilibrium concepts depend on the prior's support but not the prior probabilities.

- Unitary selective memory equilibria always exist.

- As with the existence of Berk-Nash equilibrium, this can be shown by considering a zero-sum game between the agent and an adversary who picks beliefs for the agent

- Heterogeneous-belief selective memory equilibrium requires that every action in $\mathrm{supp}(\alpha)$ is justified by a belief over the likelihood maximizers corresponding to the mixed action α.

- This differs from heterogeneous-belief self-confirming equilibrium, where each action a in the support of α can be a best response to belief over the maximizers corresponding to a alone.

"Selective Memory Equilibrium"

- Now specialize to the case where $m : A \times Y \to (0, 1]$ is independent of time.

Definition

- Action a is a limit action if there is an optimal policy π such that
$$\mathbb{P}_\pi \left[\sup\{t\colon a_t \neq a\} < \infty\right] > 0.$$
- a is a global attractor if for every optimal policy π
$$\mathbb{P}_\pi \left[\sup\{t\colon a_t \neq a\} < \infty\right] = 1.$$

Limit Strategies are Selective Memory Equilibria

Theorem

Every limit action is a selective memory equilibrium.

Proof Sketch

- A preliminary lemma shows that if a is a limit strategy, it's sufficient to look at what happens when the agent always plays a.
- By the SLLN the empirical frequency converges to p^*.
- A variation of the Borel-Cantelli lemma implies that almost surely the recalled history is large, and the SSLN implies that a large recalled history is close to the memory-weighted outcome distribution.
- Then extend Berk [1966]'s concentration result: On a set of recalled histories with objective probability $\to 1$, the distributions that don't maximize the memory-weighted likelihood have posterior probability $\to 0$.
- So asymptotically the agent best replies to $\Theta^m(a)$.

- In Zimmermann [2020], the agent observes three noisy reports about their performance relative to other subjects. Each report is either positive, $y = 1$, or negative, $y = 0$.
  - People are more likely to recall favorable tests.
  - They overestimate their probability of passing.
- To model this, we assume that

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

with $\phi < 1$.

Misattribution in a Joint Task

- Besides the IQ test, the agent has a co-worker relationship.
- Outcome space $Y = \{0, 1\}^2$, where y_1 denote the test result and y_2 denotes whether the group project succeeded.
- y_2 independently drawn as a mixture of p_1^* and some q^*

$$p_2^*(y_2) = \alpha p_1^*(y_2) + (1 - \alpha) q^*(y_2)\,,$$

  where α measure the correlation between performance in the two tasks.
- The agent knows this relation and has full support belief on $p_1 \in \Delta(Y)$ and $q \in \Delta(Y)$.

- The distribution that matches the fraction of own successes among the recalled experiences is an element of Θ, so it's the unique memory-weighted likelihood maximizer.
- The perceived fraction of successes is

$$p_1 = \frac{p^*}{\underbrace{p^*}_{\text{Successes}} + \underbrace{(1-p^*) \times \phi}_{\text{Failures}}} = p^* + \frac{p^*(1-p^*)}{\phi/(1-\phi) + p^*}.$$

- This ratio is increasing in the memory's selectiveness ϕ.
- To match their belief about their own ability, the agent's belief about the coworker's ability concentrates on

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

  which is decreasing in the selectiveness and how the two tasks are related.
- The agent thus overestimates their ability and underestimates the co-worker's ability.

- Now suppose the agent starts out with an unbiased belief about their coworkers ability, and each period t chooses the fraction $1 - \alpha$ of work to delegate to them.

- Here the memory-weighted likelihood maximizers do depend on the agents action, so our global convergence theorem doesn't apply, but we can give a direct proof that there is a global attractor.

- As the agent over-remembers their own successes, they become overconfident about their own ability and overly pessimistic about their coworkers, so delegate less work to them, which makes the misinference problem worse.

- And under-delegation in the unique limit strategy

Risk Attitudes and Extreme Experience Bias

- Let $y \in Y \subset \mathbb{R}$ be an amount of money received by the agent, $u(a, y) = v(y)$ for some concave $v : Y \to \mathbb{R}$.
- Safe action $a = 0$ with outcome y_0 and risky action $a = 1$.
- Extreme events are more likely to be recalled: $m(a, y) = h(|y - \mathbb{E}_{p_1^*}(y)|)$, with h increasing.
- Extreme experience bias makes the risky action seem more risky than it is.
- So if p_1^* is symmetric, and choosing the lottery over the safe action isn't a selective memory equilibrium with perfect memory, then it's not a selective memory equilibrium with extreme experience bias.

Selective Memory Equilibrium and Misspecification

- Let $\Theta^{\mathbf{1}}(\sigma)$ be the likelihood maximizers with perfect memory.

Definition

1. Strategy σ is a Berk-Nash equilibrium if it is a best response to some belief ν over elements of $\Theta^{\mathbf{1}}$.
2. A Berk-Nash equilibrium σ is a self-confirming equilibrium if there is $p \in \Theta$ that exactly matches p^* when σ is played.
3. Strategy σ is a uniformly strict Berk-Nash equilibrium if σ is the *unique* best response to every belief over elements of $\Theta^{\mathbf{1}}$.

Definition

A Berk-Nash equilibrium and a selective memory equilibrium are belief equivalent if they prescribe the same behavior, and are supported by the same belief.

Proposition

1. *Every uniformly strict Berk-Nash equilibrium with support Θ is belief equivalent to a selective memory equilibrium with full support for some memory function.*
2. *Every uniformly strict selective memory equilibrium with support Θ is belief equivalent to a uniformly strict Berk-Nash equilibrium for some Θ'.*

The uniform strictness conditions are **needed**:

- There are Berk-Nash equilibria that are not belief equivalent to any selective memory equilibrium with full support.
- There are selective memory equilibria that are not belief equivalent to any Berk-Nash equilibrium.

Behavioral Implications

- Misspecificied learning has been proposed as an explanation of persistent suboptimal behavior. Examples:
  - Dogmatic overconfidence as an explanation for entrepreneurs' investment choices (Heidhues, Koszegi, Strack 2018)
  - Incorrect beliefs and lemon markets (Esponda 2008).
  - Failure to understand regression to the mean and overly active teaching attitudes (Esponda and Pouzo 2016)
- But some analysts question why the agent doesn't realize that they are misspecified once they have lots of data.
- Our results show that most behavior that can be rationalized with misspecification can be explained by selective memory.

(Strong) Sufficient Condition for Global Convergence

- Suppose there is a $\hat{p} = \{\hat{p}_a\}_{a \in A}$ that is the unique memory weighted likelihood maximizer regardless of how the agent plays.

- Satisfied when the agent correctly believes their actions have no influence on the distribution of outcomes, and has the same memory function for all actions, as in some of our examples.

- We show this implies global convergence, and use it to prove existence of pure selective memory equilibrium.

Partial Naïveté

- Suppose actions have no effect on the outcome distribution, so memory depends only on y. (The paper treats the general case).

- And suppose the agent either doesn't remember their own actions or believes their actions convey no information, so they don't draw inferences about forgotten experiences based on their actions.

- We assume the agent knows how many observations they have made.

- The agent is partially naïve: they believe they remember each occurrence of y outcome with probability $\hat{m}(y)$, instead of the true probability $m(y)$.

Selective Memory Equilibrium with Partial Naïveté

- To adapt selective memory equilibrium to partial naïveté we modify the definition of the memory weighted likelihood maximizers.

- Let

$$\Theta^{m,\hat{m}}(\sigma) = \underset{p\in\Theta}{\operatorname{argmin}} \sum_{y\in Y} m(y)p^*(y)\log\left(\hat{m}(y)p(y)\right)$$
$$+\left(1-\sum_{y\in Y} m(y)p^*(y)\right)\log\left(1-\sum_{y\in Y} p(y)\hat{m}(y)\right).$$

Definition

A *partially naïve selective memory equilibrium* is a strategy σ such that there exists a belief $\nu \in \Theta^{m,\hat{m}}(\sigma)$ with $\sigma \in BR(\nu)$.

- Every limit strategy for $(\Theta, m, \hat{m})$ is a partially naive selective memory equilibrium for that $\Theta, m, \hat{m}$.

- If the agent is aware of their own forgetfulness but believes that their memory function is constant, the equilibria are the same as when the agent is fully naïve.

- If the agent is fully aware of their memory function, then selective memory equilibria coincide with self-confirming equilibria.

- If the agent partially realizes that their memory is selective, the relevant measure of distortion is $m(y)/\hat{m}(y)$, i.e., the ratio between how forgetful the agent is and how forgetful they think they are.

Limited and Selective Memory

- We modify the model to keep the expected number of instances recalled bounded.
- For some fixed integer k,

$$m_t(a, y) = km(a, y)/t.$$

- Here k is the asymptotic expected number of recalled histories.
- We don't expect actions to converge as often the agent will rely on a finite number of observations.
- But we look at the convergence of the empirical action frequency

$$\alpha_t(a) = \frac{1}{t} \sum_{\tau=1}^{t} \mathbb{I}_a(a_\tau)$$

- Let η_α be a distribution over histories defined as follows.

- The number of times each action-outcome pair (a, y) occurs is Poisson with parameter $k\,[\alpha(a)p_a^*(y)]\,m(a, y)$.

- The number of occurrences for each pair is independent.

- Let F_α be the distribution of beliefs induced by distribution η_α over recalled histories:

$$F_\alpha(B) = \eta_\alpha\{h : \mu(\cdot|h) \in B\}.$$

Definition

A limited-memory equilibrium is an $\alpha \in \Delta(A)$ for which there is a measurable $\rho : \Delta(\Theta) \to \Delta(A)$ such that

1. at every belief ν the action is optimal, i.e., $\rho(\nu) \in \Delta(BR(\nu))$, and
2. α equals the action frequencies induced by ρ, i.e., $\alpha = \mathbb{E}[\rho(\nu)|\nu \sim F_{\alpha}^{m,\mu_0}]$.

- A fixed point condition characterizes the action distribution α in a limited-memory equilibrium: The agent's behavior best replies to the distribution of memory that it induces.

- Note that the set of limited-memory equilibria depends on the prior μ_0 through its effect on posterior beliefs $\mu(\cdot|h)$.

Theorem

A limited-memory equilibrium exists.

- Proof shows that the correspondence that maps $\alpha \in \Delta(A)$ to the set of action distributions induced by the distribution of beliefs F_{α}^{m,μ_0} and some best response selection $\rho \in \Pi_o$ satisfies the conditions of the Kakutani fixed point theorem.

- The next result shows that whenever the agent behavior converges to an action distribution, that distribution is a limited memory equilibrium

Theorem

If α is a limit frequency, then α is a limited-memory equilibrium.

- The first step of the proof uses a law of large numbers for martingale differences to prove that if the empirical action distribution converges to α, then the joint distribution of actions and outcomes converges to a weighted average of the action contingent true DGP.
- Then show this implies that the distribution of recalled experiences converges to η_α^m, which follows from the Poisson limit theorem on the
- Finally use stochastic approximation arguments as in Esponda, Pouzo, and Yamamoto [2021] to show that if the empirical distribution converges to α yet α is not a limited-memory equilibrium, the empirical distribution must move away from α infinitely often.

Relation between Finite and Infinite Memory

Theorem

Suppose $(\alpha^k)_{k\in\mathbb{N}}$ is a sequence of limited-memory equilibria each with memory capacity k and that $\lim_{k\to\infty} \alpha^k = \hat{\alpha}$. Then $\hat{\alpha}$ is a heterogeneous-beliefs selective memory equilibrium.

- First step: show that when $\alpha^k \to \alpha$ the distributions of recalled histories also converge, so the the agent's beliefs concentrate on $\Theta^m(\hat{\alpha})$.
- Then note that each action $\tilde{a}$ for which $\hat{\alpha}(a) > 0$ is a best reply is a best reply to some belief concentrated on $\Theta^m(\hat{\alpha})$.
- There may not be a single belief that makes all of these actions best replies, and we show by example that the limit need not be a unitary-beliefs selective memory equilibrium.

Example

- $Y = \{-1, 1\}$, $A = \{-1, 0, 1\}$, $u(a, y) = ay - 0.1\mathbb{I}_{\{-1,1\}}(a)$:
- The agent's myopic best response is 1 if $\mu(p)>$9/16, -1 if $\mu(p) < 7/16$ and 0 if $7/16 < \mu(p) < 9/16$.
- Suppose the agent believes that p, the probability of $y = 1$, is independent of a and is either equal to 0.9 or 0.1.
- The probability of $y = 1$ is indeed independent of a, and equal to 0.5.
- Here both $p = .1$ and $p = .9$ are likelihood maximizers.
- So $a = 0$ is a unitary belief equilibrium, supported by the belief that both models are equally likely.

- There aren't any unitary beliefs equilibria with positive probability of both $a = -1$ and $a = 1$, because the unique belief that makes actions -1 and 1 indifferent is uniform over 0.9 and 0.1, and at that belief action 0 is preferred to both.

- However when the agent's memory is finite, beliefs will oscillate between being relatively sure that $p = .1$, relatively sure that $p = .9$, and close enough to $1/2, 1/2$ that 0 is optimal.

- So as the memory length goes to infinity, the limit of the limited-memory equilibria assigns positive probability to both -1 and 1, which requires heterogeneous beliefs.

Extension (in Development): "Rehearsal"

- Rehearsal means that an experience is more likely to be recalled if it was recalled last period (see Kandel et al. [2000]).
- Recalled histories are no longer conditionally independent given the realized history.
- Instead there is a Markov chain, the expected number of times experience (a, y) is recalled is proportional to the frequency of a, the probability of y given a, how memorable that experience is, and whether it occurred or was recalled in the last period.
- We conjecture that this Markov chain admits a unique stationary distribution and that this distribution is the limit time-average distribution over recalled histories.
- And plan to characterize what these distributions can be.

Other Extensions

.

- Generalize i.i.d. signals to a Markov process - would allow for the gambler's fallacy if switches in signals are more memorable.

- Relation between selective memory and PT/CPT.

- Allow the number of recalled instances of an experience to be concave in the number of times an experience has occurred.

- Multiple interacting agents.

Thanks!

Maintained Assumption

Assumption

(i) *There is* $p \in \Theta$ *such that for all* $y \in Y$, $a \in A$, *and* $s \in S$, $p^*_{s,a}(y) > 0$ *only if* $p_a(y) > 0$.

(ii) *The agent is correctly specified.*

Back to the text

Selective Memory Equilibrium and Misspecification

Definition

1. Action a is a Berk-Nash equilibrium if for all $s \in S$, there exists $\nu \in \Delta(\Theta^{\mathbf{1}})$ such that $a(s) \in BR(s, \nu)$.
2. A Berk-Nash equilibrium a is a self-confirming equilibrium if there is $p \in \Theta$ such that for all $s \in S$, $p_a = p_a^*$.
3. action a is a uniformly strict Berk-Nash equilibrium if for all $\nu \in \Delta(\Theta^{\mathbf{1}})$, $a = BR(\nu)$.

Definition

A Berk-Nash equilibrium a with support Θ and a selective memory equilibrium a' with support Θ' are belief equivalent if $a = a'$, and there exists a belief $\nu \in \Delta(\Theta^{\mathbf{1}}(a) \cap \Theta_s^m(a))$ such that $a(s) \in BR(s, \nu)$.

Back to slides

Remembered Histories

- After history $h_t = (s_i, a_i, y_i)_{i=1}^t$ and signal s_{t+1}, the probability distribution over recalled periods is given by

$$\begin{aligned} & \mathbb{P}\left[R_t = B | h_t, s_{t+1}\right] \\ = & \prod_{i=1}^{t} \left(\mathbb{I}_{i \in B} m_{s_{t+1}}(s_i, a_i, y_i) + \mathbb{I}_{i \notin B}(1 - m_{s_{t+1}}(s_i, a_i, y_i))\right) \end{aligned}$$

  the recalled periods determines the recalled sub-history. The *recalled history* is the subsequence of recalled experiences $h_t(R_t) = (s_i, a_i, y_i)_{i \in R_t}$.

Back to slides

Counterexamples without Uniform Strictness

- $Y = \{-1, 1\}$, probability of 1 is 0.5 regardless of a. If the agent does not have selective memory, but is misspecified, with $\Theta = [0, .2] \cup [.8, 1]$ as the probability of 1 under both actions.

- Then both $.2$ and $.8$ are likelihood maximizers, which cannot arise from selective memory with full support prior.

- $Y = \{-1, 0, 1\}$, uniform probability over outcomes regardless of a, with $\Theta = \{(1/3, 1/3, 1/3), (1/3, 1/6, 1/2)\}$ and $m(a, y) = \nVdash_{y=-1}$.

- Then both $(1/3, 1/3, 1/3)$ and $(1/6, 1/6, 1/2)$ are memory-weighted likelihood maximizers, but they can't both be maximizers with perfect memory.

Back to slides