The Model

- 2 states of the world, $\{L, H\}$

- single infinitely-lived DM does not know the state, ex ante belief $\Pr\{H\} = \pi$

- information process:

  - terminates each period w.p. η

  - otherwise generates signal $k \in \mathcal{K} \equiv \{1, 2, ..., K\}$ w.p. μ_k^S, for $S \in \{L, H\}$; $\mu_k^H / \mu_k^L \uparrow$ in k

- at the end, DM guesses the state:

$$\text{payoff} = \begin{cases} 1 & \text{if action} = \text{state} \\ 0 & \text{otherwise} \end{cases}$$

- no discounting

Memory and Strategy

Memory: A set $\mathcal{N} = \{1, 2, ..., N\}$ of available memory states

Memory Rule: When DM receives a signal, he decides which memory state to go to: transition rule described by a stationary signal processing rule

$$\sigma : \mathcal{N} \times \mathcal{K} \to \Delta(\mathcal{N})$$

Initial State: In period 0, DM chooses an initial memory state according to a distribution g^0 over $\mathcal{N}$.

Decision Rule: Conditional on making a decision in memory state i, DM chooses H with probability d_i.

Payoffs and Beliefs

Transition probability: Conditional on $S \in \{L, H\}$, probability of a transition $i \to j$ is

$$\tau^S_{i,j} = (1 - \eta) \sum_{k \in \mathcal{K}} \mu^S_k \sigma^k_{i,j}$$

where $\sigma^k_{i,j}$ is transition prob. $i \to j$ after signal k

Probability distribution over memory states:

$$f^S = \lim_{T \to \infty} \eta g^0 \sum_{t=0}^{T} (T^S)^t$$

Expected payoff:

$$\Pi(g^0, \sigma, d) = \sum_{i \in \mathcal{N}} \left[\pi f^H_i d_i + (1 - \pi) f^L_i (1 - d_i) \right]$$

Beliefs: In memory state i, DM believes

$$\Pr\{H|i\} = \frac{\pi f^H_i}{\pi f^H_i + (1 - \pi) f^L_i}$$

Optimality

Theorem: An optimal memory process exists $\forall\ \eta > 0$:

- If $\max\left\{\frac{\pi}{1-\pi}, \frac{1-\pi}{\pi}\right\} \geq \left(\frac{\mu_1^L \mu_K^H}{\mu_1^H \mu_K^L}\right)^{N-1}$, then the maximum payoff is $\max\{\pi, 1-\pi\}$.

- If $\max\left\{\frac{\pi}{1-\pi}, \frac{1-\pi}{\pi}\right\} < \left(\frac{\mu_1^L \mu_K^H}{\mu_1^H \mu_K^L}\right)^{N-1}$, then the maximizum payoff $\Pi^*(\eta, N)$ is continuous and strictly decreasing in η, and is bounded above by

$$\lim_{\eta \to 0} \Pi^*(\eta, N) \equiv \frac{1 - 2\sqrt{\pi(1-\pi)}r^*}{1 - (r^*)^2}$$

$$\text{where } \frac{1}{r^*} \equiv \left(\frac{\mu_1^L \mu_K^H}{\mu_1^H \mu_K^L}\right)^{\frac{N-1}{2}}$$

Optimal Rule for small η :

No transient states; order such that π_i increasing in i. Then for sufficiently small η :

1. DM ignores all but signals $1, K$

2. No jumps: DM moves up one state after K, down one state after 1

3. Probability of leaving states $1, N$ after opposing info proportional to $\sqrt{\eta}$

4. Interior transition probabilities $(\sigma^K_{i,i+1}, \sigma^1_{i,i-1})$ bounded away from zero. If $\frac{\mu^H_K \mu^L_K}{\mu^H_1 \mu^L_1} \leq 1$, no randomization after K; if reverse, no randomization after 1.

5. initial state, decision rule

Corollary: The DM's beliefs in state i satisfy

$$\lim_{\eta \to 0} \frac{\pi_i}{1-\pi_i} = \sqrt{\frac{\pi}{1-\pi}} \left(\frac{\mu_1^L \mu_K^H}{\mu_1^H \mu_K^L} \right)^{i - \frac{N+1}{2}}$$

Observations and Implications

- $\frac{\Pr\{H|i\}}{\Pr\{L|i\}}$ positive and finite

- if DM moves from state i to $i+1$ after a K-signal, beliefs change by

$$\frac{\mu_K^H}{\mu_K^L} \Big/ \frac{\mu_1^H}{\mu_1^L}$$

  –as if DM is a Bayesian who replaced a 1-signal with the K-signal

Biases:

1. Order Matters: 1st impressions bias in SR, last impressions bias in LR

2. Polarization:

   - SR: two individuals w/ different initial beliefs can move in opposite directions after same info

   - LR: strictly positive probability of eventual disagreement among 2 individuals with same memories who observe same info

3. Overconfidence/Underconfidence:

   - beliefs tend to be too extreme after short information sequences

   - beliefs will eventually be too conservative, with prob near 1; underconfidence especially likely with precise signals

Possible Relation to Preferences

Finite memory $\Rightarrow$ close relation between payoffs and beliefs:

- Generalize payoffs slightly: $\boldsymbol{\Delta}_S \equiv$ payoff difference b/w right decision, wrong decision in state S

- Then optimal rule for small η just depends on $\frac{\mu_1^L \mu_K^H}{\mu_1^H \mu_K^L}$ and $\frac{\pi}{1-\pi}\frac{\boldsymbol{\Delta}_H}{\boldsymbol{\Delta}_L}$

- optimal # memory states depends on importance of problem

Possible Modelling Implications

- modelling individuals as finite-state automatons may be realistic
- need randomization
- significant differences from optimal Bayesians:
  - uninformative signals largely ignored
  - agents behave as if they have "confidence thresh-holds"
  - beliefs may never converge
- maybe useful in reputation literature, private monitoring literature