Optimal Memory with Sequential Learning: Signals or Posterior Beliefs

Lars Wittrock (joint with Collin Raymond)

December 7, 2023

Maastricht University

Feedback

This project is work in progress.

Any feedback is very welcome. Especially:

- Additional papers we should cite?
- Is the theoretical model clear?
- Many extensions are possible. Which is the most interesting?

Introduction

Motivation

Problems with sequential learning:

- Judging unfair coins (Cover and Hellman, 1970),
- Estimate flood risk based on experiences (Bordalo et al., 2023),
- Financial analyst forecasting a stock's performance,
- Homicide detective investigating suspects,
- ...

Motivation

Problems with sequential learning:

- Judging unfair coins (Cover and Hellman, 1970),
- Estimate flood risk based on experiences (Bordalo et al., 2023),
- Financial analyst forecasting a stock's performance,
- Homicide detective investigating suspects,
- ...

Role of memory:

- Memory of information is essential ...
- ... but memory capacity is limited! (Miller, 1956; Cowan, 2010, ...)

Motivation

Problems with sequential learning:

- Judging unfair coins (Cover and Hellman, 1970),
- Estimate flood risk based on experiences (Bordalo et al., 2023),
- Financial analyst forecasting a stock's performance,
- Homicide detective investigating suspects,
- ...

Role of memory:

- Memory of information is essential ...
- ... but memory capacity is limited! (Miller, 1956; Cowan, 2010, ...)

→ What information do people keep track of?

Motivation

Economic models:

- Traditional theory (e.g. multi-armed bandit problems, Wald sequential testing, etc)
    - People remember *(posterior) beliefs*
    - Update beliefs sequentially

- Recent behavioral models (e.g. belief formation, etc)
    - People remember *individual signals*
    - Form a posterior belief only when prompted

Motivation

Economic models:

- Traditional theory (e.g. multi-armed bandit problems, Wald sequential testing, etc)
  - People remember *(posterior) beliefs*
  - Update beliefs sequentially
- Recent behavioral models (e.g. belief formation, etc)
  - People remember *individual signals*
  - Form a posterior belief only when prompted

Consequence:

- 'Rational' decision maker: equivalent
- Boundedly rational DM: can lead to different final beliefs/choices

This project

Research Question: How do people process sequential information?

- Do people keep track of posterior beliefs or individual signals?

This project

Research Question: How do people process sequential information?

- Do people keep track of posterior beliefs or individual signals?

Preview: Context of the situation matters!

- If signals are easy to keep track of: remember signals
- If posteriors are easy to keep track of: remember posteriors

Context matters

Two examples:

- Financial analyst
  - Task: forecast whether stock increases in value
  - Information: large number of data sets
  - Clear task and many signals $\rightarrow$ remember posterior
- Homicide detective
  - Task: assess probability of guilt for different suspects
  - Information: few clues
  - Less certain task and few signals $\rightarrow$ remember signal

Context matters

Two examples:

- Financial analyst
  - Task: forecast whether stock increases in value
  - Information: large number of data sets
  - Clear task and many signals $\rightarrow$ remember posterior
- Homicide detective
  - Task: assess probability of guilt for different suspects
  - Information: few clues
  - Less certain task and few signals $\rightarrow$ remember signal

Focus on two dimensions:

- Uncertainty about decision-relevant dimension
- Number of signals

Literature

Related literature

- **Memory limitations**
    - Memory capacity is limited
      (Miller, 1956; Cowan, 2010; Ericsson and Kintsch, 1995; Oberauer et al., 2016)
    - People optimize how they use their memory
      (Miller, 1956; Bays and Husain, 2008; Ma et al., 2014; Brady et al., 2009)

Related literature

- **Memory limitations**
  - Memory capacity is limited
    (Miller, 1956; Cowan, 2010; Ericsson and Kintsch, 1995; Oberauer et al., 2016)
  - People optimize how they use their memory
    (Miller, 1956; Bays and Husain, 2008; Ma et al., 2014; Brady et al., 2009)
- **Different assumptions** with sequential learning
  - People remember *posterior beliefs*
    (Cover and Hellman, 1970; Wilson, 2014; Monte, 2005; Monte and Said, 2010; Monte, 2013, 2014; Chatterjee et al., 2022; Chatterjee and Hu, 2023; Hu, 2023; Dow, 1991; Benjamin, 2019)
  - People remember *individual signals*
    (Mullainathan, 2002; Bordalo et al., 2023; Neligh, 2022; Wachter and Kahana, 2023; Enke and Zimmermann, 2019; Graeber et al., 2022; Bénabou and Tirole, 2002; Bénabou and Tirole, 2004; Chew et al., 2020; Fudenberg et al., 2014; Leung, 2023)

Related literature

- **Memory limitations**
  - Memory capacity is limited
    (Miller, 1956; Cowan, 2010; Ericsson and Kintsch, 1995; Oberauer et al., 2016)
  - People optimize how they use their memory
    (Miller, 1956; Bays and Husain, 2008; Ma et al., 2014; Brady et al., 2009)
- **Different assumptions** with sequential learning
  - People remember *posterior beliefs*
    (Cover and Hellman, 1970; Wilson, 2014; Monte, 2005; Monte and Said, 2010; Monte, 2013, 2014; Chatterjee et al., 2022; Chatterjee and Hu, 2023; Hu, 2023; Dow, 1991; Benjamin, 2019)
  - People remember *individual signals*
    (Mullainathan, 2002; Bordalo et al., 2023; Neligh, 2022; Wachter and Kahana, 2023; Enke and Zimmermann, 2019; Graeber et al., 2022; Bénabou and Tirole, 2002; Bénabou and Tirole, 2004; Chew et al., 2020; Fudenberg et al., 2014; Leung, 2023)
- **Information memory**
  - Econ Theory: Posteriors are optimal (Da Silveira et al., 2020)
  - Neuroscience: Signals (and posteriors) are remembered
    (d'Acremont et al., 2013; Shadlen and Shohamy, 2016)

Theory

Model Overview

Environment

- Uncertain state of the world
- Multiple dimensions
- Sequential noisy signals

Model Overview

Environment

- Uncertain state of the world
- Multiple dimensions
- Sequential noisy signals

Decision Problem

- Single decision at the end in one dimension
- Potential uncertainty which dimension is selected

Model Overview

Environment

- Uncertain state of the world
- Multiple dimensions
- Sequential noisy signals

Decision Problem

- Single decision at the end in one dimension
- Potential uncertainty which dimension is selected

Information Processing

- Remember posteriors?
- Remember signals?

Model - Environment

States of the world:

- D dimensions, each with realization 0 or 1
- State is given by $\omega \in \Omega = \times_D d$

Decision:

- Bet $b_{d,k}$: choice between two acts for every dimension d
  - $a_{d,0}$: if dimension d has realization $k = 0$ pays off L, 0 otherwise
  - $a_{d,1}$: if dimension d has realization $k = 1$ pays off L, 0 otherwise
- $\Delta(b)$ is a lottery over bets, with generic element q

Information:

- Prior belief $\psi \in [0,1]^D$
- T periods with signals $s^t \in S$
- Blackwell matrix $G_{\omega,s}$

Model - Decision problem

Summary

- Decision problem: q, L, T
    - T signals
    - At time $T+1$ face bet b drawn from q
    - Choice between $a_{d,0}$ and $a_{d,1}$
    - Payoffs are governed by L

Model - Information Processing (i)

Beginning of each period:

- Initial prior: $\psi \in [0,1]^D$,
- Remembered beliefs: Set $R^t \subseteq D$ with (posterior) beliefs ρ_d^t,
- Remembered signals: Set M^t of pairs (m,τ) for $\tau \leq t-1$ and $m \subset S \cup \emptyset$.

Model - Information Processing (i)

Beginning of each period:

- Initial prior: $\psi \in [0,1]^D$,
- Remembered beliefs: Set $R^t \subseteq D$ with (posterior) beliefs ρ_d^t,
- Remembered signals: Set M^t of pairs (m, τ) for $\tau \leq t-1$ and $m \subset S \cup \emptyset$.

After observing signal s^t:

- Update beliefs. Costless (for now)
    - Assume Bayesian updating: $\beta_d^t(\psi, R^t, M^t, s_t) \in [0,1]$
- Remember posterior beliefs?
- Remember signals?

Model - Information Processing (ii)

Strategy:

- Choose strategy $\zeta = (\zeta^1, ..., \zeta^T)$
  - ζ^t: mapping from ψ, R^t, M^t, s^t to R^{t+1} and M^{t+1}

Model - Information Processing (ii)

Strategy:

- Choose strategy $\zeta = (\zeta^1, ..., \zeta^T)$
  - ζ^t: mapping from ψ, R^t, M^t, s^t to R^{t+1} and M^{t+1}

Remember posterior:

- Choose for any d to remember β_d^t at cost c
- If remembered $(y_d^t = 1)$: $d \in R^{t+1}$ and $\rho_d^{t+1} = \beta_d^t$

Model - Information Processing (ii)

Strategy:

- Choose strategy $\zeta = (\zeta^1, ..., \zeta^T)$
  - ζ^t: mapping from ψ, R^t, M^t, s^t to R^{t+1} and M^{t+1}

Remember posterior:

- Choose for any d to remember β_d^t at cost c
- If remembered $(y_d^t = 1)$: $d \in R^{t+1}$ and $\rho_d^{t+1} = \beta_d^t$

Remember signals:

- Choose to remember any previous signal at cost c

$$m_\tau^{t+1} = \begin{cases} \emptyset & \text{if } z_\tau^t = 0 \\ \begin{cases} m_\tau^t & \text{for } \tau \le t-1 \\ s^t & \text{for } \tau = t \end{cases} & \text{if } z_\tau^t = 1 \end{cases}$$

Optimization Problem

Cost from information processing:

$$E[C(\zeta)] = E[\sum_{t=1}^{T}\sum_{d=1}^{D} c \cdot y_d^t(\zeta^t)] - E[\sum_{t=1}^{T}\sum_{\tau=1}^{T} c \cdot z_\tau^t(\zeta^t)]$$

Expected payoff:

$$E[B(\zeta)] = E[\max\{a_{d,0}, a_{d,1}\}|\beta_d(\zeta)]$$

Optimization Problem:

$$\max_{\zeta} E[B(\zeta)] - E[C(\zeta)]$$

Results

Proposition 1: Posteriors

There exist a $\bar{q}$, $\bar{T}$ and $\bar{L}$ so that for all problems q, L, T where the entropy of q is smaller than $\bar{q}$, $T \geq \bar{T}$, $L \geq \bar{L}$, for every $t \leq T$:

- $\tau \leq t$, $(m, \tau) = (\emptyset, \tau)$,
- for at least one d, $\rho_d^{t+1} = \beta_d^t(\psi, R^t, s_t)$ and $d \in R^{t+1}$ (unless $\rho_d^{t+1} = \psi_d$).

Results

Proposition 1: Posteriors

There exist a $\bar{q}$, $\bar{T}$ and $\bar{L}$ so that for all problems q, L, T where the entropy of q is smaller than $\bar{q}$, $T \geq \bar{T}$, $L \geq \bar{L}$, for every $t \leq T$:

- $\tau \leq t$, $(m, \tau) = (\emptyset, \tau)$,
- for at least one d, $\rho_d^{t+1} = \beta_d^t(\psi, R^t, s_t)$ and $d \in R^{t+1}$ (unless $\rho_d^{t+1} = \psi_d$).

Proposition 2: Signals

There exist a $\bar{T}$, $\bar{D}$, $\bar{q}$ and $\bar{L}$ so that for all problems q, L, T, where $T \leq \bar{T}$, $D \geq \bar{D}$, the entropy of q is larger than $\bar{q}$ and $L \geq \bar{L}$, for every $t \leq T$

- for all t $R^t = \emptyset$.
- For some $\tau \leq t-1$ and s_t $m_\tau^t = m_\tau^{t-1}$ and $m_t^t = s_t$

Extensions

Any q and T:

- Proposition 2: holds for any T
- Proposition 1: may not hold for any q

Extensions

Any q and T:

- Proposition 2: holds for any T
- Proposition 1: may not hold for any q

Different costs for remembering beliefs and signals:

- Proposition 2: unchanged
- Proposition 1: “batch” updating initially

Extensions

Any q and T:

- Proposition 2: holds for any T
- Proposition 1: may not hold for any q

Different costs for remembering beliefs and signals:

- Proposition 2: unchanged
- Proposition 1: "batch" updating initially

Costly updating of posteriors:

- Proposition 2: unchanged
- Proposition 1: repeated "batch" updating

Experiment

Design

Environment: Winning lottery numbers are picked by 3 people:

Team names	**Members**
‘Size’	Hugh & Loa
‘Parity’	Eve & Todd
‘Round’	Iris & Ron

Selected members ensure 75% of numbers are according to their liking.

Task: Guess which member from Team [...] was randomly selected.

Signals: Multiple winning lottery numbers.

Treatments

‘Posterior’ treatment:

- 14 signals
- Certain task: Guess selected member from Team ‘Size’.

‘Signals’ treatment:

- 4 signals
- Uncertain task: Guess selected member from one randomly selected team.

Design

Elicitation

- Surprise question after $\sim 3/4$ of signals
- Question choice:
  - **Posterior:** 'Based on the numbers you have seen, what is the chance that X was randomly selected from Team Y in the beginning?'
  - **Signals:** 'What was the X-th winning number selected by the group of people?'

Design

Elicitation

- Surprise question after $\sim 3/4$ of signals
- Question choice:
  - **Posterior:** 'Based on the numbers you have seen, what is the chance that X was randomly selected from Team Y in the beginning?'
  - **Signals:** 'What was the X-th winning number selected by the group of people?'

Memory load

- Aim: Increase memory cost
- Distraction task after every signal:
  - 8 digit number memory task
  - 4 math questions

Timeline

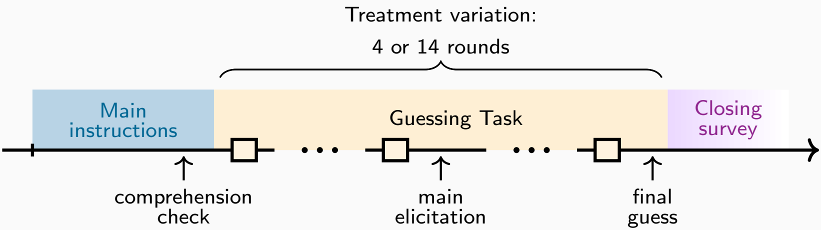

Treatment variation:
4 or 14 rounds
Main instructions
Guessing Task
Closing survey
comprehension check
main elicitation
final guess

Timeline

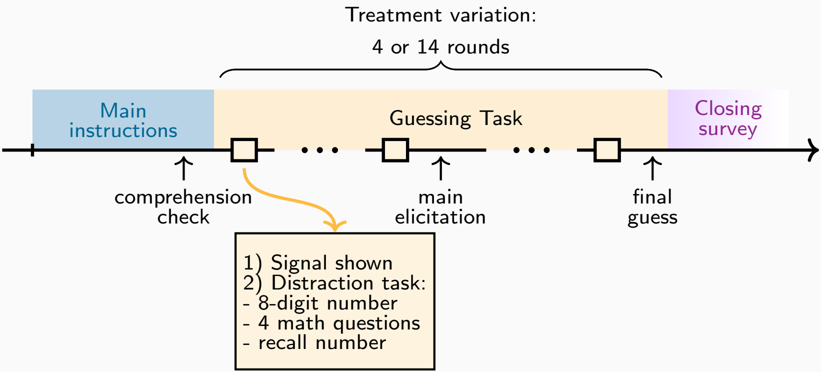
Treatment variation:
4 or 14 rounds
Main instructions
Guessing Task
Closing survey
comprehension check
main elicitation
final guess
1) Signal shown
2) Distraction task:
- 8-digit number
- 4 math questions
- recall number

Hypothesis

Compare: Share of posterior/signal question choices.

Hypothesis: More subjects choose the posterior question in the 'posterior' treatment and vice versa.

Results

Main Hypothesis

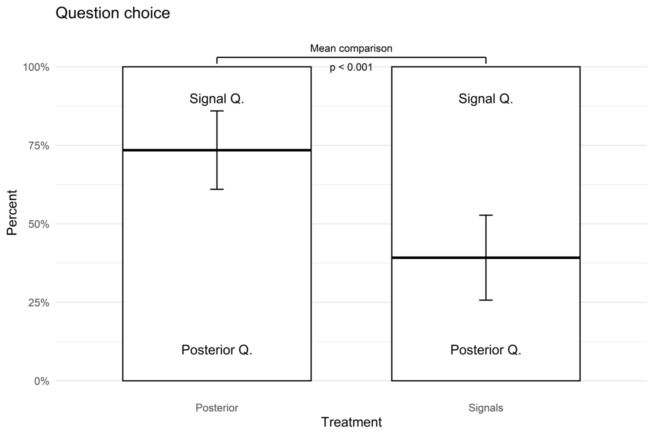

Figure 1: Question choice across the two treatments. Error bars show the 95% confidence interval around the depicted sample mean.

Question choice: Posterior

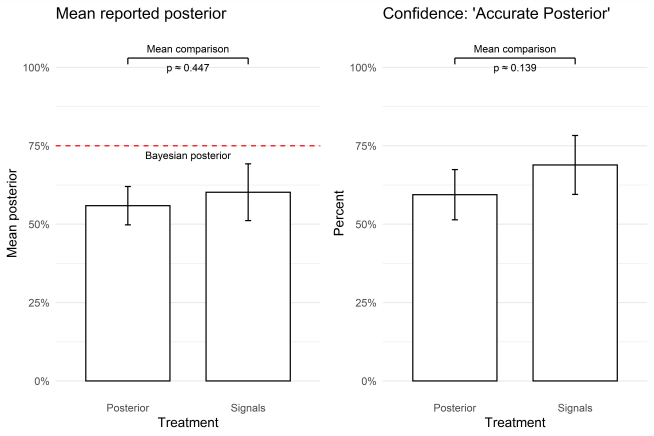

Figure 2: Average reported posterior beliefs and confidence by subjects who chose this question. The red dashed line indicates the Bayesian posterior given the 11 or 3 previous signals subjects have seen in the respective treatment.

Question choice: Signals

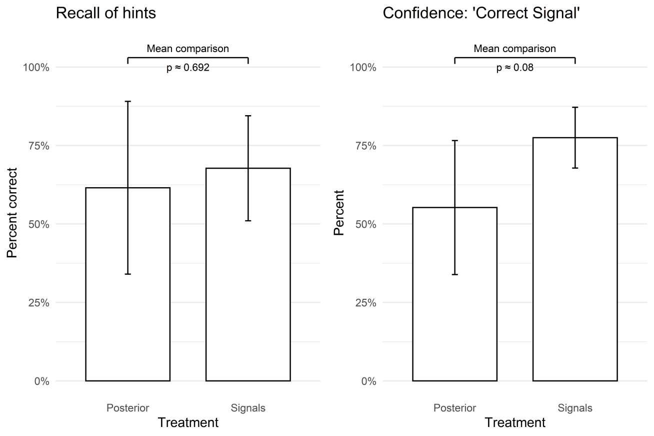

Figure 3: Percentage correct recall of the signal and confidence by subjects who chose this question.

Conclusion

Conclusion

Summary

- People rationally choose what information to remember
- In some environments remember ‘posteriors’ and in others remember ‘signals’.

Conclusion

Summary

- People rationally choose what information to remember
- In some environments remember ‘posteriors’ and in others remember ‘signals’.

Next steps

- Theory: Cost of updating beliefs?
- Theory: Non-Bayesian updating?
- Experiment: Different factors of decision environments?
- Experiment: Show impact on final beliefs/decisions?

References i

Bays, P. M. and Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, 321(5890):851–854.

Bénabou, R. and Tirole, J. (2004). Willpower and personal rules. *Journal of Political Economy*, 112(4):848–886.

Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. In *Handbook of Behavioral Economics - Foundations and Applications 2*, pages 69–186. Elsevier.

Bordalo, P., Conlon, J., Gennaioli, N., Kwon, S. Y., and Shleifer, A. (2023). Memory and probability. *The Quarterly Journal of Economics*, 138(1):265–311.

Brady, T. F., Konkle, T., and Alvarez, G. A. (2009). Compression in visual working memory: using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, 138(4):487.

Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics*.

Chatterjee, K., Guryev, K., and Hu, T.-W. (2022). Bounded memory in a changing world: Biases in behaviour and belief. *Journal of Economic Theory*, 206:105556.

Chatterjee, K. and Hu, T.-W. (2023). Learning with limited memory: Bayesianism vs heuristics. *Journal of Economic Theory*, 209:105642.

Chew, S. H., Huang, W., and Zhao, X. (2020). Motivated false memory. *Journal of Political Economy*, 128(10):3913–3939.

Cover, T. and Hellman, M. (1970). The two-armed-bandit problem with time-invariant finite memory. *IEEE Transactions on Information Theory*, 16(2):185–195.

Cowan, N. (2010). The magical mystery four: How is working memory capacity limited, and why? *Current directions in psychological science*, 19(1):51–57.

Da Silveira, R. A., Sung, Y., and Woodford, M. (2020). Optimally imprecise memory and biased forecasts. *NBER Working Paper*.

d'Acremont, M., Schultz, W., and Bossaerts, P. (2013). The human brain encodes event frequencies while forming subjective beliefs. *Journal of Neuroscience*, 33(26):10887–10897.

Dow, J. (1991). Search decisions with limited memory. *The Review of Economic Studies*, 58(1):1–14.

Enke, B. and Zimmermann, F. (2019). Correlation neglect in belief formation. *Review of Economic Studies*.

Ericsson, K. A. and Kintsch, W. (1995). Long-term working memory. *Psychological review*, 102(2):211.

Fudenberg, D., Levine, D. K., et al. (2014). Learning with recency bias. *Proceedings of the National Academy of Sciences*, 111:10826–10829.

Graeber, T., Zimmermann, F., and Roth, C. (2022). Stories, statistics, and memory. *SSRN Electronic Journal*.

Hu, T.-W. (2023). Forgetful updating and stubborn decision-makers. *Economic Theory*, 75(3):781–802.

References ii

Leung, B. T. K. (2023). A simple model of memory-based beliefformation. *Working Paper*.

Ma, W. J., Husain, M., and Bays, P. M. (2014). Changing concepts of working memory. *Nature neuroscience*, 17(3):347–356.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63(2):81.

Monte, D. (2005). Reputation with a bounded memory receiver. *Working Paper*.

Monte, D. (2013). Bounded memory and permanent reputations. *Journal of Mathematical Economics*, 49(5):345–354.

Monte, D. (2014). Learning with bounded memory in games. *Games and Economic Behavior*, 87:204–223.

Monte, D. and Said, M. (2010). Learning in hidden markov models with bounded memory. *MPRA Paper*, (23854).

Mullainathan, S. (2002). A memory-based model of bounded rationality. *The Quarterly Journal of Economics*, 117(3):735–774.

Neligh, N. L. (2022). Rational memory with decay. *SSRN Electronic Journal*.

Oberauer, K., Farrell, S., Jarrold, C., and Lewandowsky, S. (2016). What limits working memory capacity? *Psychological bulletin*, 142(7):758.

Shadlen, M. N. and Shohamy, D. (2016). Decision making and sequential sampling from memory. *Neuron*, 90(5):927–939.

Wachter, J. A. and Kahana, M. J. (2023). Associative learning and representativeness. *SSRN Electronic Journal*.

Wilson, A. (2014). Bounded memory and biases in information processing. *Econometrica*, 82(6):2257–2294.