Rational Inattention and Perceptual Distance

David Walker-Jones

University of Toronto

May 9th, 2020

Motivation:

- In many discrete choice settings learning is costly for agents
  - Learning about options takes time
  - Time is money
- Significant implications for standard economic questions
  - Makes welfare analysis more challenging
    - Direct impact: learning costs reduce surplus
    - Indirect impact: agents may make choice mistakes, so revealed preference analysis is more difficult
- Empirical researchers may wish to fit a costly learning model
- Theorists may wish to use a costly learning model for predictions in strategic settings
  - But what costly learning model?

Standard tool: Shannon Entropy

- Agent behaviour is well understood (Matějka & McKay, 2015)
  - Tractable
  - Strong predictions
- Lacks perceptual distance (Dean & Neligh, 2019)
  - Some outcomes may be more similar than others
  - Similar outcomes should be more difficult to differentiate
  - Not possible with Shannon Entropy
  - Can result in unrealistic predictions for behaviour

Perceptual distance:

- Agent faced with screen with 100 red and blue balls
- Wants to identify colour that constitutes the majority
  - Set of options: {Red, Blue}
  - If correct get prize
- Outcomes with 49 and 51 red balls are similar
  - Hard to differentiate between states
  - Behaviour should be similar in states
- Outcomes with 1 and 99 red balls are dissimilar
  - Easy to differentiate between states
  - Behaviour should be dissimilar in states
- Not possible with Shannon Entropy
  - Predicts behaviour is same when 1 or 49 balls are red
  - Easy to reject experientially (Dean & Neligh, 2019)

We provide an alternative that generalizes Shannon Entropy:

- Axiomatic foundation
    - Clear what structure is being imposed
- Incorporates perceptual distance
    - Explains choice behaviour studied by (Dean & Neligh, 2019)
- Maintains tractability
- New caveats for multinomial logit
- The natural way of generalizing Shannon Entropy to incorporate perceptual distance

Formal Setting:

- Agent learns about measurable space $(\Omega, \mathcal{F})$
  - Finite states: $\omega \in \Omega$
  - Set of events: $\mathcal{F}$
- Prior belief: $\mu \in \Delta(\Omega)$
  - Starting point for learning
  - Representation of knowledge agent already has

How do we think of the agent learning?

- Begin with agent trying to learn state of world
  - Natural starting point
  - Useful benchmark
- Learn through series of yes or no question
  - Tractable
  - Supported by psych literature (Stewart, Chater, & Brown, 2006; Noguchi & Stewart, 2014, 2018)

Formal Setting:

- Yes or no question: binary partition of state space
  - Two mutually exclusive events whose union is Ω
  - $\mathcal{P}^b = \{A_1, A_2\}$ s.t. $A_1 \cup A_2 = \Omega$, $A_1 \cap A_2 = \emptyset$, $A_i \neq \emptyset$
- Realized event $\mathcal{P}^b(\omega)$ is event that contains realized state ω
  - $\mathcal{P}^b(\omega) = A_i \in \{A_1, A_2\}$ iff $\omega \in A_i$
- Cost of learning the realized event of a binary partition given the agents beliefs, $C(\mathcal{P}^b, \mu)$, is primitive of this paper

Formal Setting:

- Learning strategy: $S^b = (\mathcal{P}_1^b, \ldots, \mathcal{P}_n^b)$ s.t.:
  - Always reveals the state of the world
    - $\forall \mu \in \Delta(\Omega)$: learning realized event of each $\mathcal{P}_i$ reveals ω
- Without loss to assume agent does not want to change learning strategy partway through learning
  - Theorem 1: order questions by their ‘difficulty’
  - Efficient no matter the belief
  - True for any belief reached during learning process

Formal setting:

- Expected cost of a learning strategy:

$$C(S^b,\ \mu) \equiv C(\mathcal{P}_1^b,\ \mu) + \mathbb{E}\Big[C\Big(\mathcal{P}_2^b,\ \mu(\cdot|\mathcal{P}_1^b(\omega))\Big) + \cdots + C\Big(\mathcal{P}_n^b,\ \mu(\cdot|\cap_{i=1}^{n-1}\mathcal{P}_i^b(\omega))\Big)\Big]$$

- Realized cost of learning strategy:
  - Pay cost to learn realized event of $\mathcal{P}_1^b$ given belief
  - Update belief
  - Pay cost to learn realized event of $\mathcal{P}_2^b$ given belief
  - Update belief
  - ...
  - Pay cost to learn realized event of $\mathcal{P}_n^b$ given belief

Example:

- Would insurance policy A benefit agent?
- $\Omega_A = \{A_1, A_2\}$
  - Policy A is beneficial (A_1), or it is not (A_2)
  - Yes or no question: $\mathcal{P}_A^b = \{A_1, A_2\}$
- Belief $\mu_A \in \Delta(\Omega_A)$
- Cost of learning fully determined by $\mu_A(A_1)$
  - Cost of learning if policy A is beneficial: $C(\mathcal{P}_A^b, \mu_A(A_1))$

Example:
Perhaps there are other insurance policies that may benefit them

- Would insurance policy B or D benefit them?
- $\Omega_B = \{B_1, B_2\}, \ \Omega_D = \{D_1, D_2\}$
  - $\Omega = \Omega_A \times \Omega_B \times \Omega_D$
- Yes or no questions: $\mathcal{P}_B^b = \{B_1, B_2\}$, $\mathcal{P}_D^b = \{D_1, D_2\}$
- Beliefs: $\mu_B \in \Delta(\Omega_B)$, $\mu_D \in \Delta(\Omega_D)$
  - $\mu \in \Delta(\Omega)$
- Assume policies are 'similar' enough to have same cost function: $C(\mathcal{P}_A^b, \cdot) = C(\mathcal{P}_B^b, \cdot) = C(\mathcal{P}_D^b, \cdot) = C(\mathcal{P}^b, \cdot)$

Example:
Assume agent learns the state of world by successively learning about different policies

- If the realization of the policies is independent, then the cost of learning the state is:

$$C(\mathcal{P}^b_A, \mu(A_1)) + C(\mathcal{P}^b_B, \mu(B_1)) + C(\mathcal{P}^b_D, \mu(D_1))$$

$$= C(\mathcal{P}^b, \mu(A_1)) + C(\mathcal{P}^b, \mu(B_1)) + C(\mathcal{P}^b, \mu(D_1))$$

- Answers to the questions do not have any 'shared' information
- What if policy realizations are not independent?

Example:

- Suppose the probabilities each policy is beneficial sum to one:
  - Suppose $\mu(A_1) + \mu(B_1) + \mu(D_1) = 1$
  - e.g. $\mu(A_1) = 0.5$, $\mu(B_1) = \mu(D_1) = 0.25$
- Suppose agent further knows only one policy is beneficial
  - If they learn a policy is beneficial they know the state of world
  - If they learn two policies are not they know state of world
- Directly learning about one policy results in indirectly learning about other policies
- Order of inquiry may matter
  - What should be directly learned and what indirectly learned?

Example:
If C is close to constant, then learning about policy A first then B if needed may minimize learning costs

- $$C(\mathcal{P}^b, \mu(A_1)) + (1-\mu(A_1))C\left(\mathcal{P}^b, \frac{\mu(B_1)}{\mu(B_1)+\mu(D_1)}\right)$$
  $$< C(\mathcal{P}^b, \mu(B_1)) + (1-\mu(B_1))C\left(\mathcal{P}^b, \frac{\mu(D_1)}{\mu(D_1)+\mu(A_1)}\right)$$

Example:
If C varies a lot depending on belief, then agent may want to instead avoid policy A and learn about B and then D if needed

- $$C(\mathcal{P}^b, \mu(A_1)) + (1 - \mu(A_1))C\left(\mathcal{P}^b, \frac{\mu(B_1)}{\mu(B_1) + \mu(D_1)}\right)$$
  $$> C(\mathcal{P}^b, \mu(B_1)) + (1 - \mu(B_1))C\left(\mathcal{P}^b, \frac{\mu(D_1)}{\mu(D_1) + \mu(A_1)}\right)$$

Example:

It may also be that order does not impact cost

- $C(\mathcal{P}^b, \mu(A_1)) + (1 - \mu(A_1))C\Big(\mathcal{P}^b, \dfrac{\mu(B_1)}{\mu(B_1) + \mu(D_1)}\Big)$
  $= C(\mathcal{P}^b, \mu(B_1)) + (1 - \mu(B_1))C\Big(\mathcal{P}^b, \dfrac{\mu(D_1)}{\mu(D_1) + \mu(A_1)}\Big)$
- This is what we assume: order does not impact the expected learning cost when questions are 'similar' enough
  - When questions 'similar' enough: which answers are directly learned and which are indirectly learned does not matter
- Relaxes Shannon's assumption that order does not impact the expected learning cost for all questions
  - Shannon assumes all questions are 'similar'
  - Result: all states are the same amount of similar
  - No perceptual distance

Axioms:

- **Axiom 1 (Measurement):** Given a binary partition $\mathcal{P}^b = \{A_1, A_2\}$, $C(\mathcal{P}^b, \mu)$ is determined by $\mu(A_1)$ and $\mu(A_2)$, and we can thus write $C(\mathcal{P}^b, \mu) = C(\mathcal{P}^b, \mu(A_1))$
- Cost of yes/no question fully determined by chance of yes/no
  - Agent is only concerned with differentiating A_1 and A_2
  - Not concerned with which state in A_1 or A_2 is realized
  - Not concerned with conditional distribution given A_1 or A_2

Axioms:

- **Axiom 2 (Self Similarity):** Given a binary partition $\mathcal{P}^b$, and a vector of probabilities (p_1, p_2, p_3) such that $p_1, p_2, p_3 \in [0, 1)$ and $p_1 + p_2 + p_3 = 1$, C is such that:

$$C(\mathcal{P}^b, p_1) + (1 - p_1)C\left(\mathcal{P}^b, \frac{p_2}{p_2 + p_3}\right)$$
$$= C(\mathcal{P}^b, p_2) + (1 - p_2)C\left(\mathcal{P}^b, \frac{p_1}{p_1 + p_3}\right)$$
$$= C(\mathcal{P}^b, p_3) + (1 - p_3)C\left(\mathcal{P}^b, \frac{p_1}{p_1 + p_2}\right)$$

- When questions are 'similar' enough, what is directly learned and what is indirectly learned does not change expected cost
- A question is assumed to be 'similar' to itself

Axioms:

- **Axiom 3 (Weak Continuity):** Given a binary partition $\mathcal{P}^b$, there is a probability $p \in [0, 1]$ such that $C(\mathcal{P}^b, p)$ is continuous at p
- Very weak assumption
- Required of useful theory

Lemma: Given binary partition $\mathcal{P}^b = \{A_1, A_2\}$, if C satisfies our three axioms, then there exists a multiplier $\lambda(\mathcal{P}^b) \in \mathbb{R}_+$ such that for all probability measures μ:

$$C(\mathcal{P}^b, \mu) = -\lambda(\mathcal{P}^b)\Big(\mu(A_1)\log(\mu(A_1)) + \mu(A_2)\log(\mu(A_2))\Big)$$

The convention used here is to set $0\log(0) = 0$.

Given finite Ω, there are a finite number of binary partitions, and we can order them by their associated multipliers

- Let S^{b*} denote a learning strategy that contains all binary partitions ordered by their multipliers
  - Break ties in any way

Theorem 1: If C satisfies our three axioms, then for any probability measure μ:

$$C(S^{b*}, \mu) = \min_{S^b} C(S^b, \mu)$$

- S^{b*} is an efficient learning strategy no matter the prior
- S^{b*} can be used to compute the cost of learning the state
  - Lemma tells us cost of each yes/no question in S^{b*}

How does Theorem 1 help us study inattentive agents?

- Inattentive agents:
  - Must choose one option from a discrete set
  - Option payoffs can differ in different states of world
  - Agent chooses signal structure before deciding option
  - Pays for signal according to cost function
  - May not perfectly observe the state of the world
- We can measure the cost of inattentive learning as the reduction it causes in the cost of learning the state of the world
- If inattentive learning results in posterior $\tilde{\mu}$, cost of learning is:

$$C(S^{b*}, \mu) - C(S^{b*}, \tilde{\mu})$$

Agent's Problem:

- Options: $\mathcal{N} = \{1, \ldots, N\}$
- Each option, $n \in \mathcal{N}$, in each state of the world, $\omega \in \Omega$, has a value to the agent $\mathbf{v}_n(\omega) \in \mathbb{R}$
- Recommendation Strategy: pick chance of selecting each option in each state of the world
  - Pick $F(n|\omega) \in [0, 1]\ \forall \omega \in \Omega, \forall n \in \mathcal{N}$
- $V(n|F) = \mathbb{E}_{F(\omega|n)}[\mathbf{v}_n(\omega)]$
- Learning cost:

$$\mathbf{C}(F(n, \omega), \mu) = \mathbb{E}\Big[C(S^{b*}, \mu) - C(S^{b*}, \mu(\cdot|n))\Big]$$

- Agent's problem:

$$\max_{F \in \Delta(\mathbb{R} \times \Omega)} \sum_{\omega \in \Omega} \sum_{n \in \mathcal{N}} V(n|F) F(n|\omega) \mu(\omega) - \mathbf{C}(F(n, \omega), \mu),$$

$$\text{such that } \forall \omega \in \Omega : \sum_{n \in \mathcal{N}} F(n, \omega) = \mu(\omega)$$

Resultant behaviour:

- Assume the agent behaves optimally
- Let the probability of them selecting option n in state ω be denoted $\Pr(n|\omega)$
- For each event A and option n, let:

$$\Pr(n|A) = \sum_{\omega \in A} \Pr(n|\omega)\mu(\omega|A)$$

- For each options n, let the unconditional chance of selecting n be denoted:

$$\Pr(n) = \sum_{\omega \in A} \Pr(n|\omega)\mu(\omega)$$

Remember, S^{b*} is a learning strategy that contains all the binary partitions ordered by their multipliers

- It may be that some binary partitions in S^{b*} are redundant
  - $\mathcal{P}_i^b \in S^{b*}$ is redundant if no matter their prior μ the agent always knows the state of the world before getting to $\mathcal{P}_i^b$
- Lets remove all redundant binary partitions from S^{b*}
  - Denote the resultant learning strategy $\tilde{S}^{b*}$

Suppose the binary partitions in $\tilde{S}^{b*}$ have M different multipliers

- Let λ_1 denote the multiplier associated with all binary partitions with the lowest multiplier
- Let $\mathcal{P}_{\lambda_1}$ denote the unique partition that provides the same information as all the binary partitions with multiplier λ_1
- ...
- Let $\lambda_M > 0$ denote the multiplier associated with all binary partitions with the M^{th} lowest multiplier
- Let $\mathcal{P}_{\lambda_M}$ denote the unique partition that provides the same information as all the binary partitions with multiplier λ_M

Theorem 2: If the agent's behaviour is optimal, then $\forall\, n \in \mathcal{N}$, and $\forall\, \omega \in \Omega$ such that $\mu(\omega) > 0$, the probability that option n is selected in state ω satisfies:

$$\Pr(n|\omega) = \frac{\Pr(n)^{\frac{\lambda_1}{\lambda_M}} \Pr(n|\mathcal{P}_{\lambda_1}(\omega))^{\frac{\lambda_2-\lambda_1}{\lambda_M}} \ldots \Pr(n|\cap_{i=1}^{M-1}\mathcal{P}_{\lambda_i}(\omega))^{\frac{\lambda_M-\lambda_{M-1}}{\lambda_M}} e^{\frac{\mathbf{v}_n(\omega)}{\lambda_M}}}{\sum_{\nu\in\mathcal{N}} \Pr(\nu)^{\frac{\lambda_1}{\lambda_M}} \Pr(\nu|\mathcal{P}_{\lambda_1}(\omega))^{\frac{\lambda_2-\lambda_1}{\lambda_M}} \ldots \Pr(\nu|\cap_{i=1}^{M-1}\mathcal{P}_{\lambda_i}(\omega))^{\frac{\lambda_M-\lambda_{M-1}}{\lambda_M}} e^{\frac{\mathbf{v}_\nu(\omega)}{\lambda_M}}}$$

- Generalization of Matějka and McKay's (2015) useful necessary condition for the optimal behaviour of the agent

Theorem 3: Optimal choice behaviour is identical to the behavior produced by a random utility model where $\forall \omega \in \Omega$ such that $\mu(\omega) > 0$, each option $n \in \mathcal{N}$ has perceived value:

$$u_n = \tilde{v}_n + \alpha_n + \epsilon_n,$$

where $\tilde{v}_n = \frac{\mathbf{v}_n(\omega)}{\lambda_M}$, ϵ_n has an iid Gumbel distribution, and:

$$\alpha_n = \tfrac{\lambda_1}{\lambda_M} \log(N\Pr(n)) + \tfrac{\lambda_2 - \lambda_1}{\lambda_M} \log(N\Pr(n|\mathcal{P}_{\lambda_1}(\omega)))$$

$$+ \cdots + \tfrac{\lambda_M - \lambda_{M-1}}{\lambda_M} \log(N\Pr(n| \cap_{i=1}^{M-1} \mathcal{P}_{\lambda_i}(\omega)))$$

- Fitted values can be biased even if $\Pr(n) = \frac{1}{N}\ \forall n$
- Contradicts rational inattention model with Shannon Entropy (Matějka & McKay, 2015)

We created a generalization of Shannon Entropy that:

- Has an axiomatic foundation
- Incorporates perceptual distance
  - Explains choice behaviour studied by (Dean & Neligh, 2019)
- Tractably implementable
  - Generalizes useful characterizations from Shannon Entropy
- Produces new caveats for multinomial logit
  - Predicts a new kind of informational bias
  - Seems natural in many contexts
  - Cannot be identified with the unconditional choice probabilities

Thanks for your (in?)attention!

Questions or comments?

Dean, M., & Neligh, N. L. (2019). Experimental tests of rational inattention.

Matějka, F., & McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, *105*(1), 272–98.

Noguchi, T., & Stewart, N. (2014). In the attraction, compromise, and similarity effects, alternatives are repeatedly compared in pairs on single dimensions. *Cognition*, *132*(1), 44–56.

Noguchi, T., & Stewart, N. (2018). Multialternative decision by sampling: A model of decision making constrained by process data. *Psychological review*, *125*(4), 512.

Stewart, N., Chater, N., & Brown, G. D. (2006). Decision by sampling. *Cognitive psychology*, *53*(1), 1–26.