Discussion of “Limits of Learning”

Memory and Attention Conference, Chicago, June 2025

Pedro Bordalo

"Limits of Learning", Rich & Gureckis 2018

- Setup: learn about options through accumulation of experience

$$\text{Choice} \rightarrow \text{Payoff} \rightarrow \text{Learning}$$

- Learning trap: stable wrong beliefs despite extensive opportunity to learn
  - Learn about options that are chosen → information acquisition is endogenous
  - Negative beliefs about an option → not chosen → learning stops
  - Learning traps if RL generalizes poorly from initial experiences (e.g. bad initial draw)
- This paper: multiple features, deterministic payoff (unlike typical bandit problems)

“Limits of Learning”, Rich & Gureckis 2018

- Experiment: packages have features (color, shape, …)
  - Payoff is $\begin{cases} -1 & if\ blue, square \\ 1 & otherwise \end{cases}$
  - Avoiding package gives payoff 0
- Participants decide repeatedly whether to open a package
  - feedback on chosen option → attend to one feature
  - feedback on all options → attend to both features
  - as inferred from majority of choices

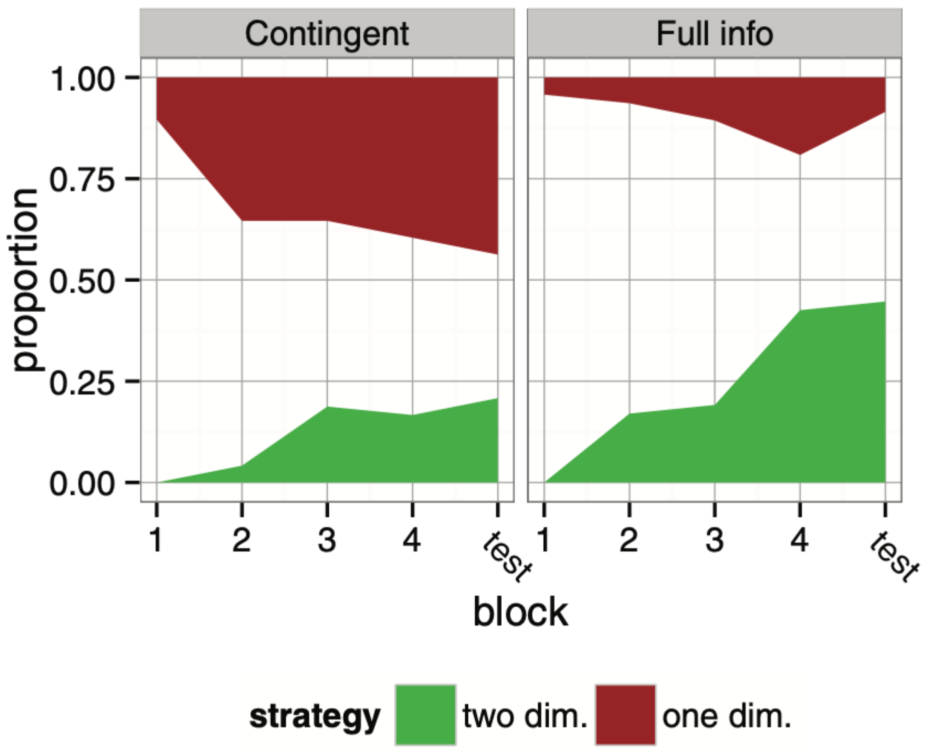

“Limits of Learning”, Rich & Gureckis 2018

- Self-reports of attention allocation at the end of the experiment

“Limits of Learning”, Rich & Gureckis 2018

- Self-reports of attention allocation at the end of the experiment

- In sum: use deterministic setting to link low exploration to lack of attention / overgeneralization
- wrong beliefs (neglect of relevant features) drive suboptimal choices (too little exploration)
- emphasize lack of data due to costly experimentation (like most econ lit), consistent with feedback

Learning traps and cognitive mechanisms

- Information is often freely available
  - availability of info may not be only constraint on low exploration (Schwartzstein 2014)
  - errors in beliefs arise with full information, and can persist with feedback (Hannah et al 2014, Esponda Vespa Yuksel) (why?)
  - can (potentially suboptimal) stable choices emerge through other mechanisms?

Learning traps and cognitive mechanisms

- Information is often freely available
  - availability of info may not be only constraint on low exploration (Schwartzstein 2014)
  - errors in beliefs arise with full information, and can persist with feedback (Hannah et al 2014, Esponda Vespa Yuksel) (why?)
  - can (potentially suboptimal) stable choices emerge through other mechanisms?
- Learning trap relies on developing strong wrong beliefs "early on"
  - "bias of simplification" penalizing full attention; wrong priors, …
  - how do we decide what features to attend to?

Outline

- Brand loyalty and advertising
  - stark example of (potentially suboptimal) low exploration
- Reasoning and attention to features
  - Predicting value vs problem recognition

Ads as Cues Bordalo, Burro, Gennaioli, Nacamulli, Shleifer (2025)

- Consumers have strong habits among similar goods – e.g. Coke v. Pepsi – despite substantial variation in relative prices over time (low exploration)
- These brands are well known and frequently purchased (information available)
  - standard view: preferences, brand capital (Becker Murphy)

Ads as Cues Bordalo, Burro, Gennaioli, Nacamulli, Shleifer (2025)

- Consumers have strong habits among similar goods – e.g. Coke v. Pepsi – despite substantial variation in relative prices over time (low exploration)
- These brands are well known and frequently purchased (information available)
  - standard view: preferences, brand capital (Becker Murphy)
- These brands are also strongly advertised. Why?
  - George Stigler asked Mr. Wrigley why he kept advertising his chewing gum company despite it being well known. "See that electric train passing by", replied Wrigley, "what would happen if they turned off the electricity?"
  - Ads are cues that make specific wants salient, otherwise may not come to mind

Who does advertising have an impact on?

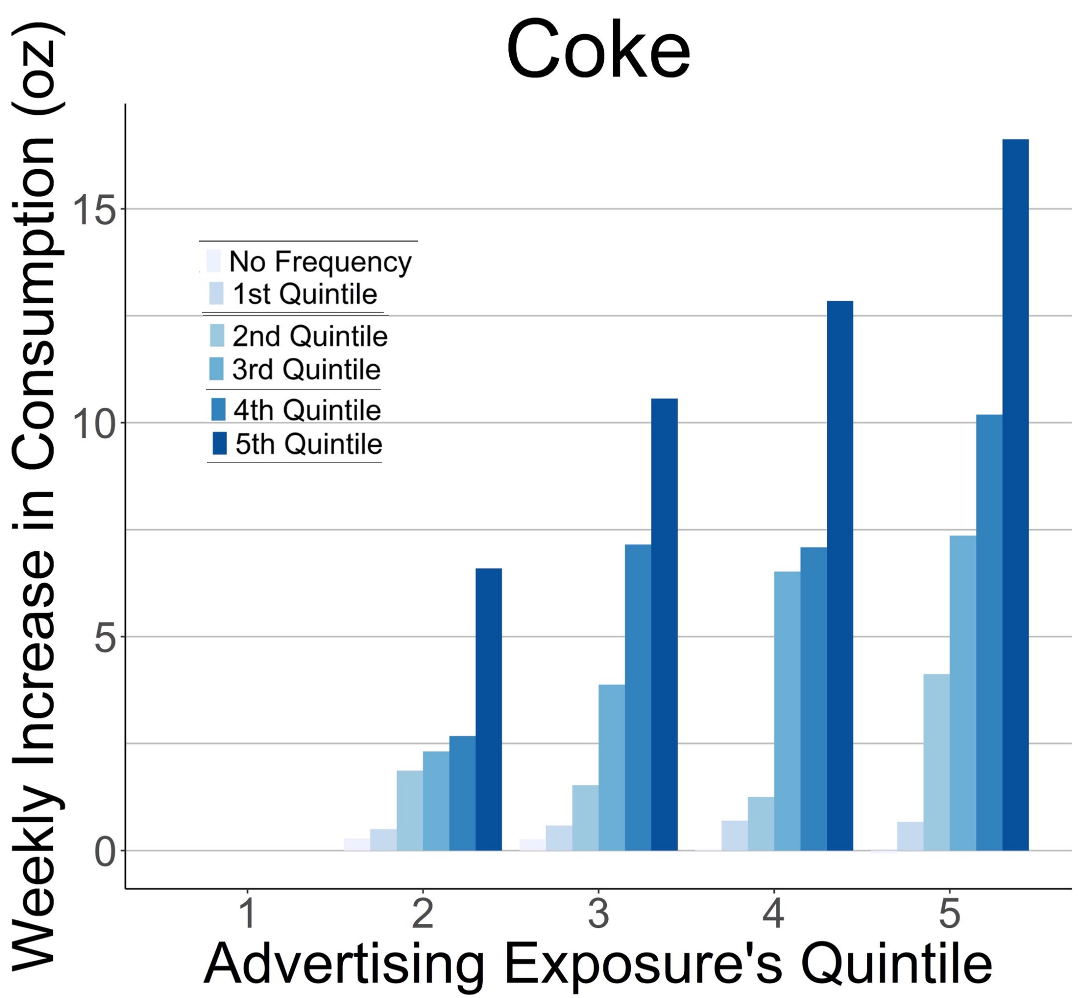

- NielsenIQ Homescan panel on purchases, Nielsen Ad Intel data on ad exposure
- Impact of ads increases with past frequency of purchase. Inconsistent with information.
- Our approach: cognition
  - limited retrieval, shaped by cues and similarity

    $r(Coke, cue) \sim f_{Coke} \cdot S(Coke, cue)$
  - ads are such cues (reminders)
  - possibly + differences in tastes

The memory mechanism

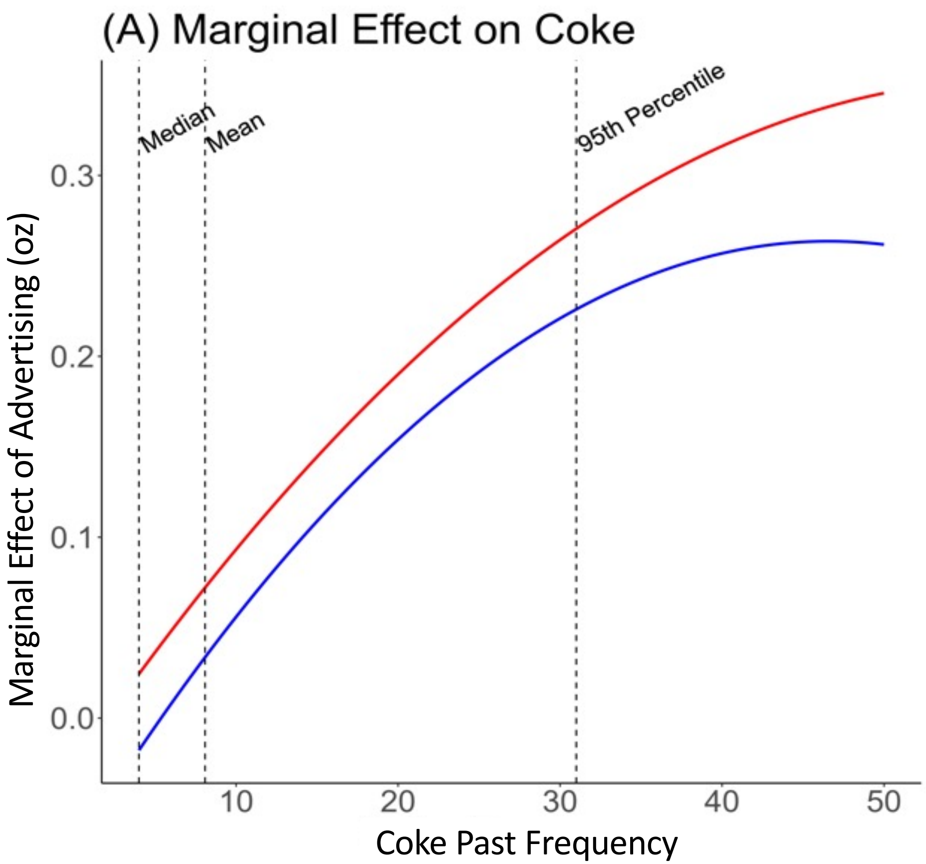

$$r(Coke, cue) \sim f_{Coke} \cdot S(Coke, cue)$$

1. ads cue own brand but also similar brands
   - for a Coke drinker, a Pepsi ad cues Coke
   - larger impact on frequent consumers, who have many memories to draw on

2. ads distract from non-similar brands
   - a Cola ad has: positive spillovers on drinks, but negative spillovers on non-drinks
   - Similar results measuring similarity from ads

Taking stock

- Selective memory breaks the iron-clad link between utility and choice
  - choices reflect not just preferences for a good but what is top of mind, due to habit or cues
  - brand loyalty -- low exploration -- because utility from brand in memory is easier to retrieve / simulate than the alternative
    - suggests information is not enough to change suboptimal behavior

- Identify the role of memory using choice data + non-choice data:
  - i) consumer experiences, ii) exposure to ads/cues, and iii) similarity
  - In RL setting, evidence that recent payoffs or irrelevant context (Niv et al) matters?

Outline

- Brand loyalty and advertising
  - stark example of low exploration despite full information
- Reasoning and attention to features
  - Predicting value vs problem recognition

Reasoning about choices

When people face a new problem:

- Approaches such as RL make prediction of value, including based on similar past experiences

- Attend to features that helped predict value

Reasoning about choices

When people face a new problem:

- Approaches such as RL make prediction of value, including based on similar past experiences
- Attend to features that helped predict value

"A cognitive theory of reasoning and choice" (BGLS 2024)

- what kind of problem is this?
- Attend to features that help *recognize the problem*
- output is a method of valuation, not prediction of value

This Model Gives

- Endogenous similarity:

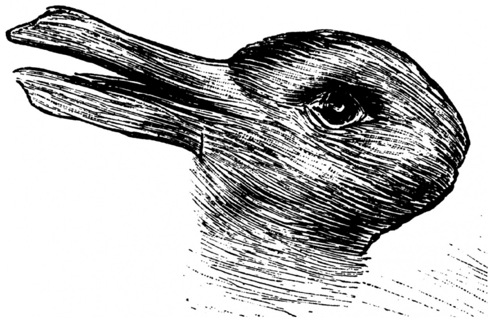

Recognize duck ⇔ focus on beak (neglect mouth)

Recognize rabbit ⇔ focus on mouth (neglect beak)

Cognitive dissonance: neglect relevant, available information

This Model Gives

- Endogenous similarity:

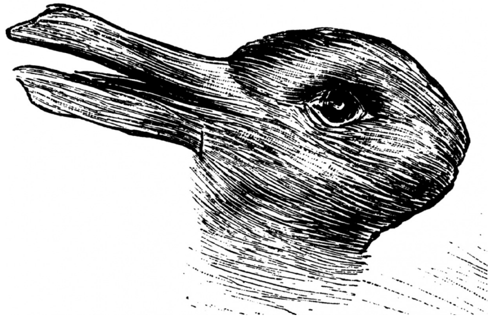

Recognize duck ⇔ focus on beak (neglect mouth)

Recognize rabbit ⇔ focus on mouth (neglect beak)

Cognitive dissonance: neglect relevant, available Information

- Two common patterns of problem solving that shed light on many puzzles:

1. Overgeneralization
    - features relevant in *frequent* past problems are heavily weighted today, even if no longer relevant
    - a person who experienced poverty tends to focus on price, even if no longer poor
    - related to RL result, but due to problem recognition, not learning

This Model Gives

- Endogenous similarity:

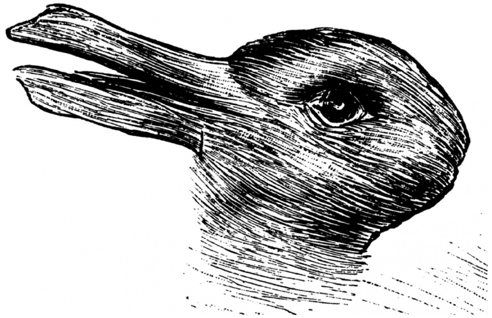

Recognize duck ⇔ focus on beak (neglect mouth)

Recognize rabbit ⇔ focus on mouth (neglect beak)

Cognitive dissonance: neglect relevant, available information

- Two common patterns of problem solving that shed light on many puzzles:

2. Over-differentiation
    - A *change in context* changes categorization though the problem is normatively similar & full information
    - instability: focus on prices in shop, neglect opportunity cost at home
    - related to transfer failure

This Model Gives

- Endogenous similarity:

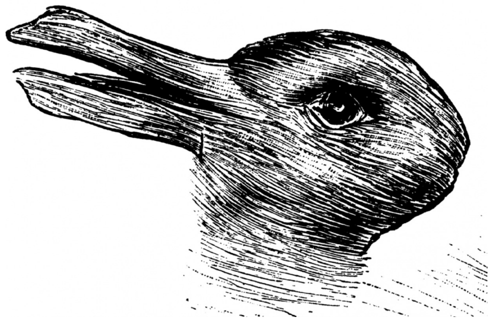

Recognize duck ⇔ focus on beak (neglect mouth)

Recognize rabbit ⇔ focus on mouth (neglect beak)

Cognitive dissonance: neglect relevant, available information

- Two common patterns of problem solving that shed light on many puzzles:

1. Overgeneralization
2. Over-differentiation

- No bias for simplicity: selected category can be more complicated than actual problem

Recognition and learning

Problem recognition yields a valuation method:

- focus on consumption and splurge *vs* focus on opportunity cost and be stingy
- because it does not yield a prediction of value, choice realizations may be less directly tied to learning

Experiment (Esponda, Vespa, Yuksel 2024):

- Inference problem: disease has low base rate, diagnostic test has high accuracy (both given)
- Participants get test result, forecast realization, observe realization (feedback), & repeat
    - consistently neglect base rate for up to 200 rounds, consistent with wrong categorization (BCGKS 2024)
    - If no statistics are given, participants form more accurate beliefs. Suggestive evidence they pay more attention to feedback

Conclusion

This paper shows the promise of bandit type problems to study cognition underlying learning:

- can measure attention to features across problems and link it to retrieval of past experiences, similarity judgments and choice
- assess how information provision (feedback) works and why

Cognitive approach suggests that more than information per se, what matters is how you use it.