Behavioral Economics

Lecture 2: Belief Updating

Daniel Martin
danielmartin@ucsb.edu

Questions

- Do people update their beliefs correctly?
- Today:
  - What does this even mean?!?
  - What is the evidence that people don't update beliefs correctly?
  - What are some proposed biases?

Belief updating in the wild

- People learn facts about the world (that are relevant to decision-making)
  1. Learning about stocks (invest or not)
  2. Learning about today's weather (carry an umbrella or not)
  3. Learning about the job applicant (hire or reject)
- Decision-makers learn about these facts from information (analyst report, weather forecast, CV, etc.)
- How does this information impact what they should believe?

Belief updating

- Primitives
  - States $\omega \in \Omega$
  - Evidence (signal, information) $e \in E$
  - Conditional probability $\Pr(e|\omega)$
  - Prior $\Pr(\omega)$
- Bayes' Rule

$$\overbrace{\Pr(\omega|e)}^{\text{posterior/updated belief}} = \frac{\overbrace{\Pr(e|\omega)}^{\text{conditional probability}} \overbrace{\Pr(\omega)}^{\text{prior}}}{\underbrace{\sum_{\omega' \in \Omega} \Pr(e|\omega') \Pr(\omega')}_{\text{normalizing factor}}}$$

- Why standard? Tractable and strongly normative

Belief updating

- But do people *actually* update beliefs according to Bayes' Rule?
- Early experimental evidence from Kahneman and Tversky

Kahneman & Tversky (1972)

- Consider a person who is tested for a disease (states sick or fine: $\Omega = \{S, F\}$)
- Disease prevalence is 15% in the general population, and the test has an accuracy of 80 percent
- What is the chance the person is sick conditional on a positive test result ($e = p$)?

$$\Pr(S|p) = \frac{\Pr(p|S) * \Pr(S)}{\Pr(p|S) * \Pr(S) + \Pr(p|F) * \Pr(F)}$$

$$\Pr(S|p) = \frac{0.15 * 0.8}{0.15 * 0.8 + 0.85 * 0.2}$$

$$\Pr(S|p) = 0.41$$

Kahneman & Tversky (1972)

- The most common answer was 80% completely ignoring the prior!
- Most answers were more than 50%
- Doctors are famously terrible at this
- But people got a lot of practice/exposure on these types of calculations during the COVID pandemic
- Begs the question: What is the impact of feedback?

Esponda, Vespa & Yuksel (2024)

- Same problem as Kahneman & Tversky (1972), but different framing
- Project is a success or failure conditional on a signal being positive or negative
- 100 projects in total, 15 of which are successes and the remaining 85 are failures
- Interface produces a signal, positive or negative, with 80 percent reliability

Esponda, Vespa & Yuksel (2024)

Repeated the task 200 times to see if people can learn over time with feedback

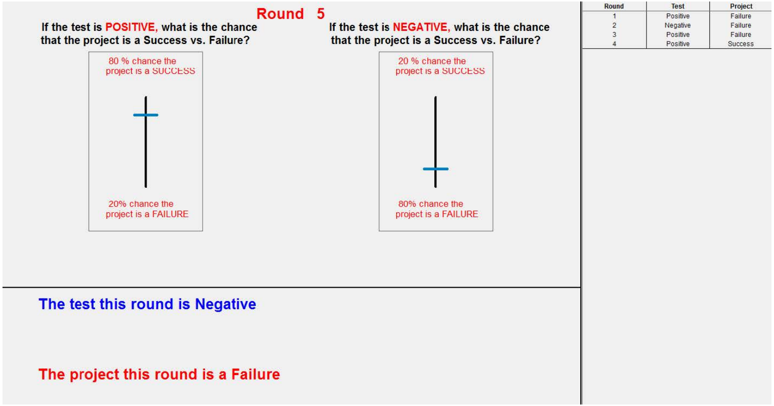

Esponda, Vespa & Yuksel (2024)

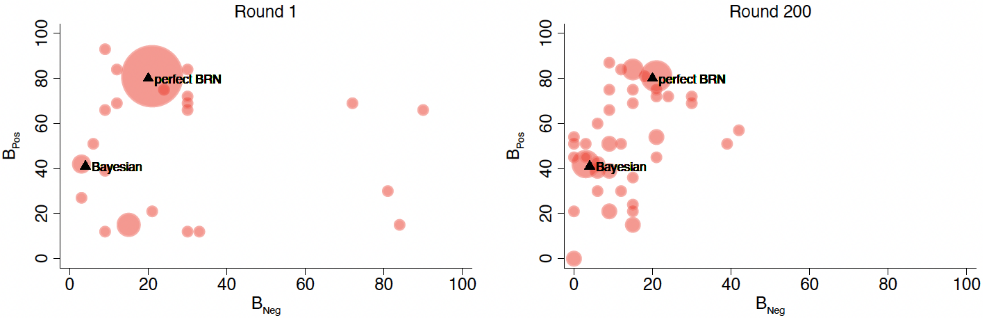

- Persistence of suboptimal behavior in the presence of feedback

Esponda, Vespa & Yuksel (2024)

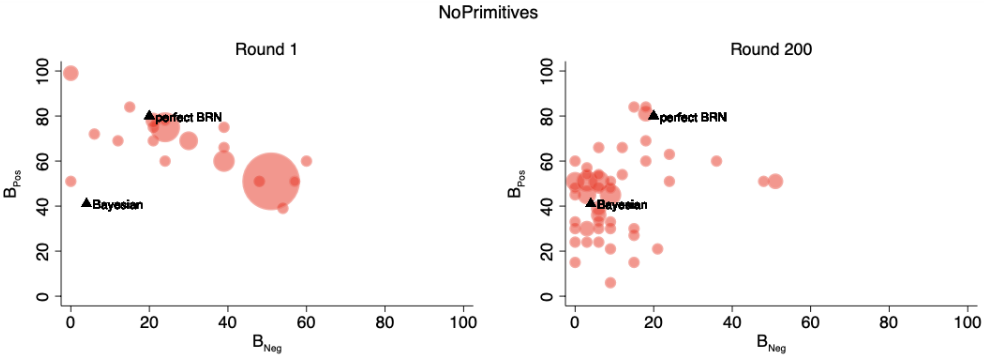

- Do better if told nothing about the primitives

Esponda, Vespa & Yuksel (2024)

- Why does this happen?
  - Primatives might cause people to over-focus on incorrect mental models
  - Could increase confidence in initial answers, limiting engagement with and learning from feedback
- Find beliefs converge closer to Bayesian in new treatment where tell subjects (to whom the message applies) that their initial responses are not correct
- Allow subjects to "lock-in" responses and ignore feedback – use much earlier with accesss to primatives

Tversky & Kahneman (1974)

- A panel of psychologists have interviewed and administered personality tests to 30 engineers and 70 lawyers, all successful in their respective fields. On the basis of this information, thumbnail descriptions of the 30 engineers and 70 lawyers have been written. You will find on your forms five descriptions, chosen at random from the 100 available descriptions. For each description, please indicate your probability that the person described is an engineer, on a scale from 0 to 100.
- Jack is a 45-year-old man. He is married and has four children. He is generally conservative, careful, and ambitious. He shows no interest in political and social issues and spends most of his free time on his many hobbies which include home carpentry, sailing, and mathematical puzzles. The probability that Jack is one of the 30 engineers in the sample of 100 is ___%.
- Then they reverse the probabilities, no change!

Tversky & Kahneman (1974)

- Phenomenon: base-rate neglect again!
- Explanation: representativeness heuristic
- "Representativeness-based updating": In making an inference, $\Pr(\omega|e)$, about the likelihood that ω is true, we tend to over-use how "representative" or similar e is to what we would expect to see if ω were true
- While $\Pr(e|\omega)$ is of course a central part of proper Bayesian reasoning, we tend to over-use these conditionals and to improperly take into account other (more intuitive) notions of the similarity of e to ω in updating our beliefs

Esponda, Oprea & Yuksel (2023)

- Application of representativeness: Discrimination
- Experimental subjects distort their evaluation of new evidence: interpret such information to be more representative of the group the individual belongs to
- Produces an irrational discriminatory gap in the evaluation of members of the two groups
- The gap disappears when subjects are prevented from contrasting different groups, suggesting it is driven by representativeness
- Also disappears when subjects receive information before learning of the individual's group

Esponda, Oprea & Yuksel (2023)

- Estimate the "type" (a number between 1 and 100) of an unspecified attribute of a fictitious member of one of two groups ("green" or "orange") that differ only in their mean type
- Why completely abstract? Remove other, non-inferential sources of discrimination (such as animus or taste-based discrimination) that might confound
- Information:
  - Told which group the fictitious person is a member of (green or orange)
  - Shown a number of dots on their screen equal to the fictitious person's true type for a split second
- Short exposure to the dots means that subjects receive only a noisy, subjective signal (pro: more natural, con: less control)

Esponda, Oprea & Yuksel (2023)

Round 8

Guess the actual value of the randomly selected person.

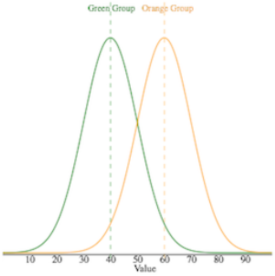

The person is from the Orange Group.

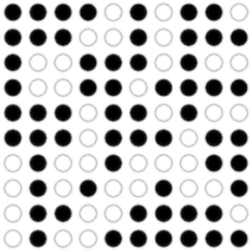

Esponda, Oprea & Yuksel (2023)

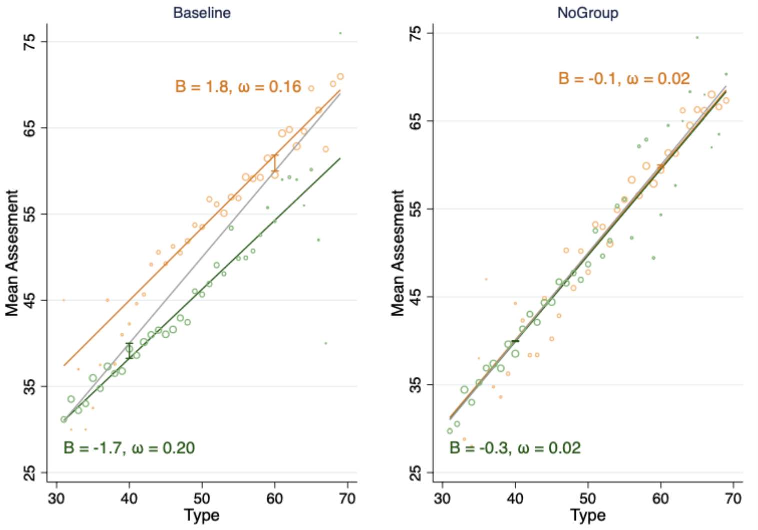

Evolving method

- Following K&T, **MANY** experiments have tested Bayes' Rule
- Standard method now is "marble-cup"/"ball-urn"/"bag-chip" procedure

Example

- Imagine the following experiment similar to Holt & Smith (2009)
- Cup A or cup B is picked to draw marbles from
  - Subjects don't know which was picked
- Cup A is 55% likely to be picked and has 60 green/40 red marbles
- Cup B has 40 green/60 red marbles
- Marbles are drawn one at a time with replacement
- Subjects guess the likelihood that cup A was picked after seeing each marble draw

Example

The **A** cup has a **55%** chance of being picked in each round.
The **B** cup has a **45%** chance of being picked in each round.

The **A** cup contains **60** green marbles and **40** red marbles.
The **B** cup contains **40** green marbles and **60** red marbles.

Your draws:

1	2	3	4
red			

Marble returned to the cup after each draw.

$$\Pr(A|red) = \frac{\Pr(red|A)\Pr(A)}{\Pr(red|A)\Pr(A) + \Pr(red|B)\Pr(B)} = \frac{.4 * .55}{.4 * .55 + .6 * .45} = .44898$$

Example

- Can show representativeness heuristic in this setting

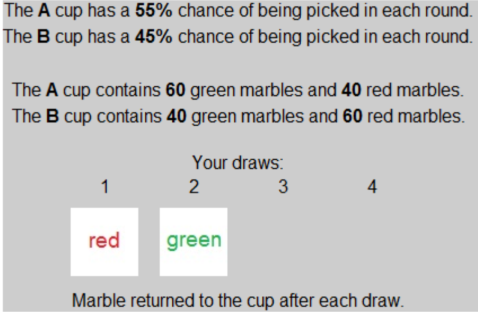

- If cup A contained 50 green and 50 red marbles, then would think A more likely than actually true

Bigger picture

- These types of experiments have generated many observed departures from Bayesian updating!
- Benjamin (2019) wrote an influential survey of papers on belief updating and found a general pattern of underreaction to new information ("prior-biased updating")
  - "The central lesson is that people underweight both the information from signals and their priors—errors that I refer to as underinference and base-rate neglect"
  - "Two biases—prior-biased updating and base-rate neglect—push in opposite directions. A plausible conjecture is that prior-biased updating dominates when the priors are close to 50–50 whereas base-rate neglect dominates when the priors are extreme, but I am not aware of any work that has directly examined how these two biases interact"

Bigger picture

- As mentioned before, Benjamin (2019) found a general pattern of underreaction to new information
- But previous studies showing underreaction (prior-biased updating) use signals with more than 60% accuracy
- People have a tendency to over infer from weak signals, and under infer from strong signals (Augenblick, Lazarus & Thaler 2021)
- See clearly in an experiment with symmetric signal where signal precision = Bayesian posterior

Augenblick, Lazarus & Thaler (2021)

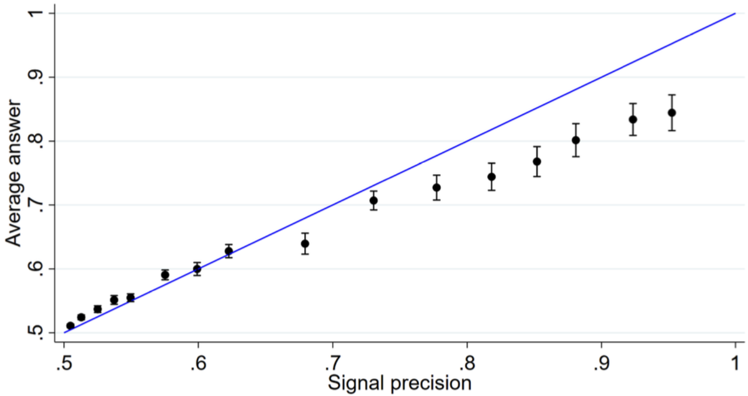

Grether (1980)

- To accommodate these biases, we make the following adjustments

$$\overbrace{p(\omega|e)}^{\text{posterior (not true probability Pr)}} = \frac{\overbrace{\Pr(e|\omega)^{\beta_1}}^{\text{distorted probability}} \overbrace{\Pr(\omega)^{\beta_2}}^{\text{distorted prior}}}{\underbrace{\sum_{\omega' \in \Omega} \Pr(e|\omega')^{\beta_1} \Pr(\omega')^{\beta_2}}_{\text{normalizing factor}}}$$

- β_1 is the weight we place on signals, and β_2 is the weight we place on prior
- $\beta_1 = \beta_2 = 1$ we get the Bayes' rule
- $\beta_1 < 1$ we get underinference, $\beta_1 > 1$ we get overinference
- $\beta_2 < 1$ we get base rate neglect, $\beta_2 > 1$ we get overweighting of priors

Grether (1980)

- Popular because easy to estimate
- Divide by other state to remove normalizing constant

$$\frac{p(\omega|e)}{p(\omega'|e)} = \frac{\Pr(e|\omega)^{\beta_1}}{\Pr(e|\omega')^{\beta_1}} \frac{\Pr(\omega)^{\beta_2}}{\Pr(\omega')^{\beta_2}}$$

- Then take log-linear form

$$\ln \frac{p(\omega|e)}{p(\omega'|e)} = \beta_1 * \ln \frac{\Pr(e|\omega)}{\Pr(e|\omega')} + \beta_2 * \ln \frac{\Pr(\omega)}{\Pr(\omega')}$$

- Then just regress!

Kovach (2020)

- Can consider other distortions to Bayes' Rule

$$\overbrace{p(\omega|e)}^{\text{posterior (not true probability Pr)}} = \underbrace{\frac{\overbrace{\Pr(e|\omega)}^{\text{conditional probability}} \quad \overbrace{\Pr(\omega)}^{\text{prior}} \quad \overbrace{w(\omega)}^{\text{wishful thinking}}}{\sum_{\omega' \in \Omega} \Pr(e|\omega') \Pr(\omega') w(\omega')}}_{\text{normalizing factor}}$$

- $w(\omega)$ is how much the agent wants state ω to be true
- Can also be interpreted as tendency to report that state
- Adds constant to estimation before!

Evolving methods

- What are some other methods for testing for correctly updating beliefs?
  - Martingale property (average of posteriors equals the prior)
  - Studying belief movement (Augenblick & Rabin 2021)
  - Do choices (such as between bets) satisfy NIAS condition (Caplin & Martin 2015) more later
  - Willingness to pay for different bets (de Clippel, Moscariello, Ortoleva & Rozen 2024) more later

Martingale property

- The Martingale property is that the average of posteriors equals the prior
- This is necessary, but not sufficient for Bayes' rule
  - In other words, there are non-Bayesian models that satisfy this property
- Martingale property is satisfied when the posterior is a linear combination of the Bayesian posterior and prior (Epstein, Noor, Sandroni et al. 2010)

$$\overbrace{p(\omega|e)}^{\text{posterior (not true probability Pr)}} = \lambda * \overbrace{\Pr(\omega|e)}^{\text{Bayesian posterior}} + (1-\lambda) * \overbrace{\Pr(\omega)}^{\text{prior}}$$

- $0 < \lambda < 1$ you get under-updating and $\lambda > 1$ you get over-updating
- A generalization of cursed inference (Eyster & Rabin 2005) where λ is between 0 and 1

Martingale property

- Even though it's just a necessary condition, the Martingale property is still useful for testing Bayesianism – if Martingale property fails, then Bayesianism does
- An alternative expression is based on belief movement (Augenblick & Rabin 2021)
- Advantage over directly implementing Martingale test is that you can deduce the bias (under the assumption of the Grether model)
  - Kenneth Chan and Sebastian Brown are leveraging this test to see if beliefs about future wages update correctly

Augenblick & Rabin (2021)

- Introduce a test that only uses the prior belief and updated belief (data required are not too demanding)
- They showed that on average the amount of belief movement has to be equal to the amount of uncertainty reduction if people are Bayesian
- **Idea:** The more you are changing your belief, you should become more confident

2 States (Augenblick & Rabin 2021)

- Model of belief dynamics (multiple time periods)
- 2 states and state 1 occurs with probability π
- Belief movement

$$m_{t_1,t_2} \equiv \sum_{\tau=t_1}^{t_2-1} (\pi_{\tau+1} - \pi_\tau)^2 = (\pi_{t_1+1} - \pi_{t_1})^2 + \cdots + (\pi_{t_2} - \pi_{t_2-1})^2$$

- Uncertainty reduction

$$\begin{aligned} r_{t_1,t_2} &\equiv \sum_{\tau=t_1}^{t_2-1} \pi_\tau(1 - \pi_\tau) - \pi_{\tau+1}(1 - \pi_{\tau+1}) \\ &= \pi_{t_1}(1 - \pi_{t_1}) - \pi_{t_2}(1 - \pi_{t_2}) \end{aligned}$$

Test: For a Bayesian, these two statistics have to be equal on average

Augenblick & Rabin (2021)

Proposition 6 *Consider the single-signal model with correct priors.*

Given any signal precision θ *and any initial prior* $\pi_t \neq \frac{1}{2}$:

$$\beta > 1 \quad \text{(overreaction)} \quad \Rightarrow \mathbb{E}M_{t,t+1} > \mathbb{E}R_{t,t+1} \quad \text{with } \frac{\partial(\mathbb{E}M_{t,t+1} - \mathbb{E}R_{t,t+1})}{\partial \beta} > 0$$

$$\beta < 1 \quad \text{(underreaction)} \quad \Rightarrow \mathbb{E}M_{t,t+1} < \mathbb{E}R_{t,t+1}$$

$$\alpha > 1 \quad \text{(confirmation bias)} \quad \Rightarrow \mathbb{E}M_{t,t+1} < \mathbb{E}R_{t,t+1} \quad \text{with } \frac{\partial(\mathbb{E}M_{t,t+1} - \mathbb{E}R_{t,t+1})}{\partial \alpha} < 0$$

$$\alpha < 1 \quad \text{(base-rate neglect)} \quad \Rightarrow \mathbb{E}M_{t,t+1} > \mathbb{E}R_{t,t+1} \quad \text{with } \frac{\partial(\mathbb{E}M_{t,t+1} - \mathbb{E}R_{t,t+1})}{\partial \alpha} < 0.$$

Summary: if you have excess belief movement it is either overreaction or base-rate neglect, if you have negative excess belief movement, then it is the other 2 biases

References I

Augenblick, N., Lazarus, E. & Thaler, M. (2021), 'Overinference from weak signals and underinference from strong signals', *arXiv preprint arXiv:2109.09871* .

Augenblick, N. & Rabin, M. (2021), 'Belief movement, uncertainty reduction, and rational updating', *The Quarterly Journal of Economics* **136**(2), 933–985.

Benjamin, D. J. (2019), 'Errors in probabilistic reasoning and judgment biases', *Handbook of Behavioral Economics: Applications and Foundations 1* **2**, 69–186.

Caplin, A. & Martin, D. (2015), 'A testable theory of imperfect perception', *The Economic Journal* **125**(582), 184–202.

de Clippel, G., Moscariello, P., Ortoleva, P. & Rozen, K. (2024), 'Caution in the face of complexity'. Working paper.

Epstein, L. G., Noor, J., Sandroni, A. et al. (2010), 'Non-bayesian learning', *The B.E. Journal of Theoretical Economics* **10**(1), 1–20.

Esponda, I., Oprea, R. & Yuksel, S. (2023), 'Seeing what is representative', *The Quarterly Journal of Economics* **138**(4), 2607–2657.

Esponda, I., Vespa, E. & Yuksel, S. (2024), 'Mental models and learning: The case of base-rate neglect', *American Economic Review* **114**(3), 752–782.

References II

Eyster, E. & Rabin, M. (2005), 'Cursed equilibrium', *Econometrica* **73**(5), 1623–1672.

Grether, D. M. (1980), 'Bayes rule as a descriptive model: The representativeness heuristic', *The Quarterly Journal of Economics* **95**(3), 537–557.

Holt, C. A. & Smith, A. M. (2009), 'An update on bayesian updating', *Journal of Economic Behavior & Organization* **69**(2), 125–134.

Kahneman, D. & Tversky, A. (1972), 'Subjective probability: A judgment of representativeness', *Cognitive psychology* **3**(3), 430–454.

Kovach, M. (2020), 'Twisting the truth: Foundations of wishful thinking', *Theoretical Economics* **15**(3), 989–1022.

Tversky, A. & Kahneman, D. (1974), 'Judgment under uncertainty: Heuristics and biases', *Science* **185**(4157), 1124–1131.