

Informatics 1 Cognitive Science

Lecture 14: Biases in Human Decision Making

Frank Keller

14 February 2025

School of Informatics
University of Edinburgh
keller@inf.ed.ac.uk

Slide credits: Frank Mollica, Chris Lucas, Mirella Lapata

Representativeness Bias

Availability Bias

Base Rate Neglect

Wason's Card-selection Task

Previous lecture:

- Rational analysis
- Expected utility and utility functions
- Prospect Theory

This lecture:

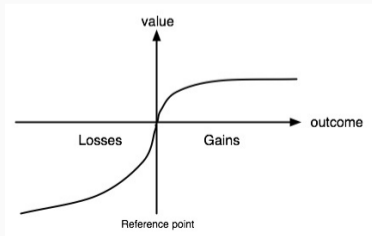
- More on cognitive biases
- A cognitive bias is a deviation from optimal, normatively correct behavior
- Example: choosing a bet even though it doesn't have the highest expected gain

Recap

Framing effects – people choose dominated pairs of bets, i.e., bets that are non-optimal according to probability theory. **How the bets are framed (presented) matters.**

Prospect theory:

- We evaluate bets one at a time.
- $U()$ focuses not on overall wealth, but gains and losses.
- $U()$ flattens out as gains/losses become more extreme, more so for gains. We over-weight extreme probabilities.



Today's Lecture

- More examples of behavior that departs from probability theory (at least under the researchers' original conception of their experiments)
- Are people solving different problems than experimenters thought?
- Are people striking a balance between good decisions and resource constraints?

Representativeness Bias

Representativeness Bias

“Bill is 34 years old. He is intelligent, but unimaginative, compulsive, and generally lifeless. In school, he was strong in mathematics but weak in social studies and humanities.”

Representativeness Bias

“Bill is 34 years old. He is intelligent, but unimaginative, compulsive, and generally lifeless. In school, he was strong in mathematics but weak in social studies and humanities.”

Rank the following statements in terms of how likely they are to be true:

- Bill is a physician who plays poker for a hobby.
- Bill is an architect.
- Bill is an accountant.
- Bill plays jazz for a hobby.
- Bill surfs for a hobby.
- Bill is a reporter.
- Bill is an accountant who plays jazz for a hobby.
- Bill climbs mountains for a hobby.

Representativeness Bias

“Bill is 34 years old. He is intelligent, but unimaginative, compulsive, and generally lifeless. In school, he was strong in mathematics but weak in social studies and humanities.”

Rank the following statements in terms of how likely they are to be true:

- Bill is a physician who plays poker for a hobby.
- Bill is an architect.
- Bill is an accountant.
- Bill plays jazz for a hobby.
- Bill surfs for a hobby.
- Bill is a reporter.
- Bill is an accountant who plays jazz for a hobby.
- Bill climbs mountains for a hobby.

Conjunction Fallacy

- Bill is an accountant. (A)
- Bill plays jazz for a hobby. (J)
- Bill is an accountant who plays jazz for a hobby. ($A \wedge J$)

Most people said $A \wedge J$ is more likely than J. Since J includes both $J \wedge A$ and $J \wedge \neg A$, this appears to be a fallacy.

\wedge : logical and

\neg : logical not

\rightarrow : logical implication (if-then)

(Tversky & Kahneman, 1983)

Conjunction Fallacy

Tversky & Kahneman's explanation:

Bill being an accountant is **representative** of his description.

In terms of probabilities, $P(\text{description}|\text{category})$ is high. This is not the same as $P(\text{category}|\text{description})$.

(Tversky & Kahneman, 1983)

Let's Estimate some Probabilities

Time for a short quiz on Wooclap!



<https://app.wooclap.com/GJCPSB>

Availability Bias

Availability Bias

What's more dangerous?

- Spending an hour on a large aircraft?
- Spending an hour in a typical passenger car?



Availability Bias

Tversky & Kahneman's explanation:

We estimate probabilities by generating or recalling examples

Availability Bias

Tversky & Kahneman's explanation:

We estimate probabilities by generating or recalling examples

Another example:

“Consider the letter R. Is R more likely to appear in

- the first position [in a word]?
- the third position [in a word]?”

152 participants; 105 incorrectly said the first position is more likely.

Tversky & Kahneman (1973). Availability: A heuristic for judging frequency and probability.

Availability can be a useful heuristic; we can't consider all possible events.

- “inference by sampling” is an active research area.
- Availability should increase with actual probability.
- If extreme events tend to be those with extreme utilities → overweighting gives better expected utility estimates.

Lieder et al. (2018). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources.

Sanborn, A. N., & Chater, N. (2016). Bayesian brains without probabilities. *Trends in cognitive sciences*, 20(12), 883-893.

Availability Bias

Availability can be a useful heuristic; we can't consider all possible events.

- “inference by sampling” is an active research area.
- Availability should increase with actual probability.
- If extreme events tend to be those with extreme utilities → overweighting gives better expected utility estimates.

Lieder et al. (2018). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources.

Sanborn, A. N., & Chater, N. (2016). Bayesian brains without probabilities. *Trends in cognitive sciences*, 20(12), 883-893.

You can find the probabilities for causes of death here:

<https://www.floridamuseum.ufl.edu/shark-attacks/odds/compare-risk/death/>

Base Rate Neglect

Base Rate Neglect

Blue and the green cabs operate in a city:

- **85%** of cabs are **blue**.
- **15%** of cabs are **green**.

A cab was involved in a hit-and-run accident at night.

A witness later identified the cab as a **green** cab.

The witness was tested for their ability to judge the color of cabs under similar conditions. They were correct **80%** of the time but confused it with the other color **20%** of the time.

Was the hit-and-run cab **green or **blue**?**

Base Rate Neglect

Typical human judgment: **green**.

Bayes' theorem:

$$P(\text{green}|\text{witness} = g) = \frac{P(\text{green})P(\text{witness} = g|\text{green})}{P(\text{green})P(\text{witness} = g|\text{green}) + P(\text{blue})P(\text{witness} = g|\text{blue})}$$

That is: $0.15 \times 0.80 / (0.15 \times 0.80 + 0.85 \times 0.20) = 0.12 / (0.12 + 0.17) = 0.41$
less than 0.5 \rightarrow Should say **blue**.

Kahneman & Tversky (1972) via Maya Bar-Hillel (1980).

Base Rate Neglect

It seems that participants **ignore the base rate (the prior probability)**!

What is the explanation for this?

Maya Bar-Hillel: We have heuristics for determining the relevance of information

Kahneman: The green cab is more representative of the witness's report

Base Rate Neglect

It seems that participants **ignore the base rate (the prior probability)**!

What is the explanation for this?

Maya Bar-Hillel: We have heuristics for determining the relevance of information

Kahneman: The green cab is more representative of the witness's report

This phenomenon is important:

- Physicians are subject to base rate neglect in evaluating diagnostic tests: prevalence of the disease (prior) vs. accuracy of the test (likelihood).
- Relevant in legal settings too. See Bar-Hillel (1980) for more.

Examples of Heuristics and Biases

Time for a short quiz on Wooclap!



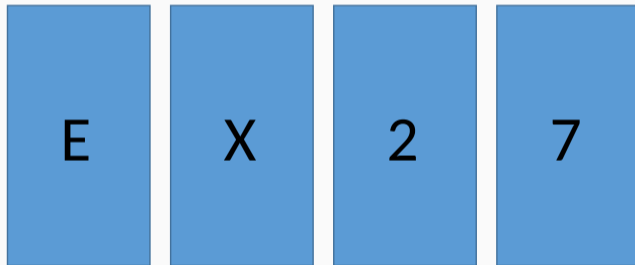
<https://app.wooclap.com/GJCPSB>

Wason's Card-selection Task

Wason's Card-selection Task

The rule:

If there is a vowel on one side of a card, then there is an even number on the other side.

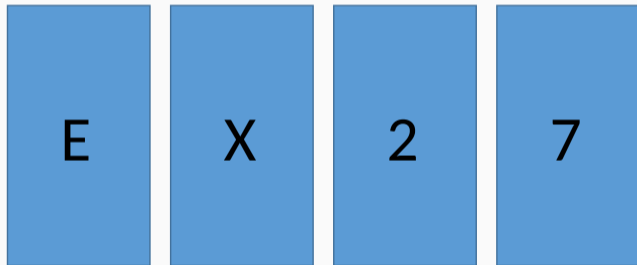


Which cards should we reverse to determine whether the rule is true, assuming cards have letters on one side and numbers on the other?

Wason's Card-selection Task

The rule:

If there is a vowel on one side of a card, then there is an even number on the other side.



According to logic: E and 7.

(Oberauer et al. 1999; Wason, 1968; participants were Edinburgh first-year psychology undergraduates)

Wason's Card-selection Task

Why **E** \wedge **7** ($p \wedge \neg q$) and not X or 2 ($\neg p \vee q$)?

The rule: p (vowel) implies q (even number).

If we flip over **E** and see an odd number ($p \wedge \neg q$),
we have falsified the rule.

If we flip over **7** and see a vowel ($\neg q \wedge p$),
we have falsified the rule.

Wason's Card-selection Task

Why **E** \wedge **7** ($p \wedge \neg q$) and not **X** or **2** ($\neg p \vee q$)?

The rule: p (vowel) implies q (even number).

If we flip over **X**, it doesn't matter if the number is even or odd – neither falsifies the rule.

If we flip over **2**, it doesn't matter if the letter is a vowel or not – neither falsifies the rule.

Wason's Card-selection Task

The logical answer is $p \wedge \neg q$.

In one representative experiment, participants chose:

- $p \wedge q$ (18 of 34 participants)
- p (7 of 34)
- $p \wedge \neg q$ (1 of 34)
- $p \wedge \neg p$ (1 of 34)
- $p \wedge q \wedge \neg q$ (4 of 34)
- $p \wedge \neg p \wedge q$ (1 of 34)
- Everything (2 of 34)

(Wason (1968). Reasoning about a rule. Quarterly Journal of Experimental Psychology.)

Wason's Card-selection Task

What's going on?

Wason:

- The task requires reasoning in accordance with logical rules.
- People do not reason in accordance with logical rules.

So, people are bad at logic and “formal operations” in general.

Time for a short quiz on Wooclap!



<https://app.wooclap.com/GJCPSB>

An alternate view of the card-selection task:

There are two competing hypotheses:

1. The $p \rightarrow q$ rule is true.
2. The $p \rightarrow q$ rule is false.

We choose examples to get the most information about the truth of the rule – including evidence **for** its truth.

(“A Rational Analysis of the Selection Task as Optimal Data Selection”, Oaksford & Chater, 1994)

A rule:

“If you eat tripe (p), you feel sick (q)”

Selecting:

- [$p/\neg p$]: “Hello tripe-[eater/non-eater], how do you feel?”
- [$q/\neg q$]: “Hello [sick/well] person, have you had tripe recently?”

(“A Rational Analysis of the Selection Task as Optimal Data Selection”, Oaksford & Chater, 1994)

Alice's policy:

- Ask tripe-eaters how they feel.
- Ask sick people if they have eaten tripe lately.

$(p \wedge q)$

Bob's policy:

- Ask tripe-eaters how they feel.
- Ask well people if they have eaten tripe recently.

$(p \wedge \neg q)$

Which policy seems more reasonable?

Alice's policy:

- Ask tripe-eaters how they feel.
- Ask sick people if they have eaten tripe lately.

$(p \wedge q)$

Bob's policy:

- Ask tripe-eaters how they feel.
- Ask well people if they have eaten tripe recently.

$(p \wedge \neg q)$

Which policy seems more reasonable?

Alice's policy takes into account the probability of the evidence: being sick is more informative than being well.

(“A Rational Analysis of the Selection Task as Optimal Data Selection”, Oaksford & Chater, 1994)

Information-seeking

Let's assume people choose the options that are most informative for testing the rule:

- In this view, the probability of p and q should affect q vs. $\neg q$ choices.
- Choices of p , q , and $\neg q$ depend on one another – some combinations are more informative than others.

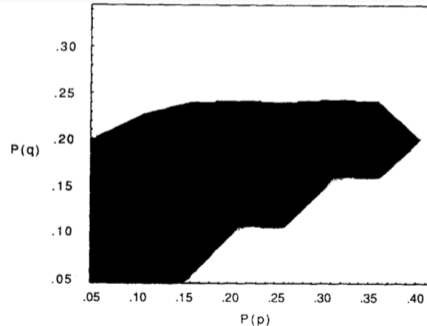


Figure 3. Plot of $P(p)$ against $P(q)$ with $P(M_1) = .5$, showing the region R (in black) where $E[I_R(q)] > E[I_R(\text{not-}q)]$.

- Classic experiments about judgment and decision-making:
 - Paired lotteries (framing effects)
 - Representative bias
 - Availability bias
 - Base rate neglect
 - Wason's card selection task
- Understanding behavior as rational, given limitations and expectations