The role of linguistic interpretation in human failures of reasoning

Salvador Mascarenhas
The University of Oxford
Ecole Normale Supérieure

ESSLLI 2016, week 2, lecture #4

1 The conjunction fallacy

Scandinavian peninsula [Slide 4]

> The Scandinavian peninsula is the European area with the greatest percentage of people with blond hair and blue eyes. This is the case even though (as in Italy) every possible combination of hair and eye color occurs. Suppose we choose at random an individual from the Scandinavian population.

Which is more probable?

1. The individual has blond hair.
2. The individual has blond hair and blue eyes.

Linda [Slide 5]

> Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

Which is more probable?

1. Linda is a bank teller.
2. Linda is a bank teller and is active in the feminist movement.

Conjunction fallacy [Slide 6]

- About 85% of subjects rank 2. higher than 1.
- According to probability theory, and taking the strings in 1. and 2. to be interpreted in the most obvious (and naive) way, this is a fallacy. For recall that $P(\varphi \wedge \psi) \leq P(\varphi)$ and $P(\varphi \wedge \psi) \leq P(\psi)$.
- Enormous amounts of ink have been spilled on this fallacy. It is certainly Tversky and Kahneman's most famous result, and it has puzzled psychologists, economists, and statisticians for decades.
- It is also an exception to the trend in psychology I've been criticizing: this fallacy has actually been investigated somewhat systematically from an *interpretation-based* perspective

The conjunction fallacy in the heuristics and biases tradition [Slide 7]

- Three classes of experiments by Tversky and Kahneman (1983):
  1. Direct transparent: as above, direct comparison between A and $A \wedge B$
  2. Direct subtle: A, $A \wedge B$, and a number of (alleged) fillers
  3. Indirect: one group sees $A \wedge B$ together with (alleged) fillers, another group sees A, B, and (alleged) fillers
- (Which type of CF effect is most puzzling to you?)
- T&K's account:
  1. Substitution of "representativeness" question for "probability" question (caveat: this is not an interpretation-based account)
  2. Conjunctive option ranks higher on representativeness than the relevant conjunct

Representativeness heuristic [Slide 8]

- "It is [...] natural and economical for the probability of an event to be evaluated by the degree to which that event is representative of an appropriate mental model [in the sense of prototype]." (Tversky & Kahneman, 1983)

Representativeness
"An assessment of the degree of correspondence between a sample and a population, an instance and a category, an act and an actor." (T&K, 1974)

- When the token and type can be described in the same terms, representativeness reduces to similarity or prototypicality.
- The notion can be made precise in terms of prototype theory (Kahneman and Frederick, 2002) and exemplar theory (Nilson, Juslin, and Olsson, 2008).

2 Alternative approaches

Probabilities vs. frequencies [Slide 10]

- There is evidence that the interpretation of "what is more probable?" is not the expected probability-theoretical one.
- Gigerenzer (1991) shows that, when the problem is stated in terms of explicit *frequencies* rather than probabilities, conjunction effects all but disappear.

  Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

  Which is more probable?

1. Bank tellers: ____ out of 100
2. Bank tellers who are active in the feminist movement: ____ out of 100

A simple secondary implicature? [Slide 11]

- T&K (1983) consider the possibility that "bank teller" is in fact interpreted as "bank teller who is not active in the feminist movement." This would be (to some extent) an absolving interpretation (NB: would still require something like base-rate neglect)
- They controlled for this by using the variants below

1. Linda is a bank teller *whether or not she is active in the feminist movement.*
2. Linda is a bank teller and is active in the feminist movement.

A simple secondary implicature? [Slide 12]

1. Linda is a bank teller *whether or not she is active in the feminist movement.*
2. Linda is a bank teller and is active in the feminist movement.

- Since (1) literally expresses ignorance about whether Linda is a feminist, it is implausible that it be strengthened to "not a feminist"
- Conjunction fallacy effects were still found, *though the rates decreased: 57%*
- (Asking subjects to bet on the original Linda problem has a comparable mitigating effect: 56% commit the conjunction fallacy)

T&K on "whether or not" [Slide 13]

> "the extension of 'Linda is a bank teller whether or not she is active in the feminist movement' *clearly includes* the extension of 'Linda is a bank teller and is active in the feminist movement'" (Tversky & Kahneman, 1983, p. 299)

- Does that sound right to you?

??? Linda is a bankteller and is active in the feminist movement, whether or not she is active in the feminist movement.

A primary implicature? [Slide 14]

- Recall: primary implicatures are of the form "it is not the case that the speaker believes that φ." They are *very* difficult to cancel, unlike secondary implicatures
- So we expect the non-conjunctive option to be interpreted along the lines below, *even when we control for the secondary implicature*

1. Linda is a bank teller *and I have no way of deciding whether she is active in the feminist movement or not.*
2. Linda is a bank teller and is active in the feminist movement.

Question
Is it conjunction-fallacyish to resist a statement of ignorance about something one actually doesn't feel ignorant about? Is it irrational? (See Dulany & Hilton, 1991; Mascarenhas, 2014)

An interpretation-based hypothesis for the conjunction fallacy [Slide 15]

For a transparent Linda-type problem whose options are A and $A \wedge B$

1. reasoners calculate primary (ignorance) implicatures for the A option, committing "one" to ignorance about B,
2. reasoners consider that $P(B) > .5$ (base-rate neglect),
3. option A pays a hefty price in the currency of probabilities because it asserts ignorance about a proposition that reasoners do not believe one is ignorant about.

3 The CF as hypothesis testing

The CF within Bayesian confirmation theory [Slide 17]

- T&K imagine that the primary (hard) task subjects are considering is comparing the probabilities that A and that $A \wedge B$, *conditioned* on the description of Linda
- But maybe they are instead asking themselves *which hypothesis about Linda* (A or $A \wedge B$) best explains the description of Linda.
- Crupi et al. (2008) prove that a number of different strategies for measuring increase in firmness from the literature derive the observed responses to conjunction fallacy setups as a kind of hypothesis testing.

Likelihoods [Slide 18]

- A very pedestrian form of hypothesis testing turns out to be enough to illustrate the spirit of the account: just look at *likelihoods* rather than posteriors
- In this view, subjects conclude *reasonably* that the hypothesis that Linda is a bankteller and a feminist increases the probability of the given description more than the hypothesis that she is a bankteller.

$P(B|d) \geq P(B \wedge F|d)$, but $P(d|B) < P(d|B \wedge F)$

- But why should subjects ask themselves a question about *confirmation* as opposed to a simple question about posteriors?

Interpretation to the rescue [Slide 19]

- The task of conditionalizing on the description of Linda renders the description *irrelevant*!
- Any account of pragmatics predicts that hearers will try their best to interpret all of the information they are given in a *relevant* way

So
The target task is one that violates a principle of pragmatics. It is therefore expected (and in fact reasonable) for subjects to come up with a related task that maximizes the relevance of the information given in the setup. Testing hypotheses (NB an *erotetic* process) is an excellent and immediate candidate.

4 Representativeness revisited

Representativeness revisited [Slide 21]

- T&K tell us that subjects substitute a question about *representativeness* for a question about *probabilities*
- But these two functions are fundamentally different.
- Probability is a function from *propositions* to the real line
- Representativeness is a function from *a pair of an individual and a property* to (say, for simplicity's sake) the real line

1. How likely is it that Linda is a bank teller? $P(B(\ell))$
2. How typical an example of a bank teller is Linda? $R(B,\ell)$

The problem [Slide 22]

- For the examples studied by T&K, representativeness works more or less straightforwardly because a topical individual is always available and yields a coherent question about representativeness
- In their examples, the conjunctive option always contains two propositions about the same individual

(1) Linda is a bank teller and Clinton will win the presidential election.

- How likely is (1)? OK
- How typical an example of (1) is Linda? PROBLEM
- With respect to *what* can we ask a representativeness question about (2)?

(2) a. It is raining.
b. The Goldbach conjecture is true.

Properties and individuals or propositions? [Slide 23]

- Representativeness accounts in the market either predict there should be no conjunction effects for such cases or they cannot evaluate such cases.
- (This is not to say that representativeness cannot in principle be rescued: define a notion of naive probability inductively (structure-sensitively) that reduces to representativity of the subject with respect to the VP in the base case.)
- In a pilot study, I tested four scenarios with the following properties.
    1. A description of a topical individual t,
    2. an unlikely (and unrepresentative) statement about t,
    3. a likely statement about something else,
    4. the conjunction of these two statements.

An example [Slide 24]

> The Scandinavian peninsula, including Sweden, Norway, and Finland, has a high percentage of people with blond hair and blue eyes. This is the case even though (as in the US) every possible combination of hair and eye color occurs. Suppose we choose at random an individual from the Scandinavian peninsula.

Questions (two conditions)

- To what degree does the individual represent a typical example of the following statement?
- How likely do you think the following statement is?

1. The individual has dark hair.
2. Norway has universal healthcare.
3. The individual has dark hair, and Norway has universal healthcare.

Design [Slide 25]

- Each subject was assigned either to the representativeness condition or the probability condition
- Each subject sees four different stories, and rates only one sentence for each story. They were allowed to refuse to answer by saying the question was ill-formed
- A third group compared the conjunctive option with its unrepresentative/unlikely conjunct

Results [Slide 26]

1. The individual has dark hair.
2. Norway has universal healthcare.
3. The individual has dark hair, and Norway has universal healthcare.

A between-subjects analysis showed that

- Representativeness and likelihood did not differ significantly for (1),
- but differed greatly for (2): likelihood was much greater.
- Representativeness and likelihood differed significantly for (3).
- The likelihood judgments displayed a significant conjunction error ((3) $>$ (1)), but representativeness judgments did not.

Results [Slide 27]

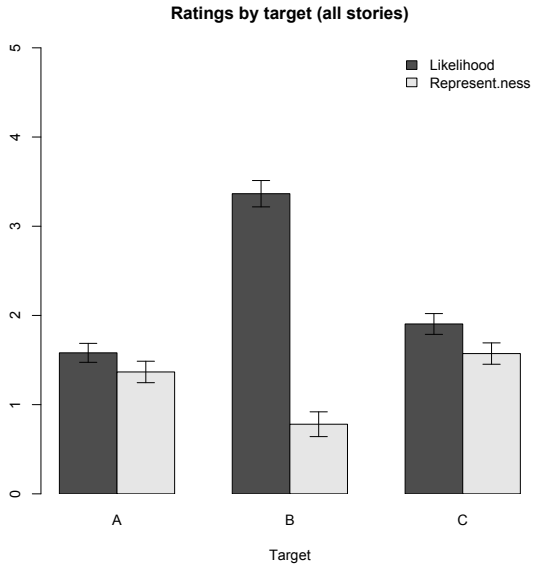

Just a pilot… [Slide 28]

- The design needs serious improvement: the question about representativeness was (to my sensibilities) ill formed
- Though perhaps strangely only a negligible portion of subjects choose the response "Cannot answer."
- But at the very least: we've shown that conjunction errors exist for sentences where the predictions from representativeness are dubious

References

Crupi, Vincenzo, Branden Fitelson and Katya Tentori (2008). Probability, confirmation, and the conjunction fallacy. *Thinking & Reasoning*, 14(2):182–199.

Dulany, Don E. and Denis J. Hilton (1991). Conversational implicature, conscious representation, and the conjunction fallacy. *Social Cognition*, 9(1):85–110.

Gigerenzer, Gerd (1991). How to make cognitive illusions disappear: Beyond "heuristics and biases". In W. Stroebe and M. Hewstone, editors, *European Review of Social Psychology*. Chichester, England: Wiley.

Kahneman, Daniel and Shane Frederick (2002). Representativeness revisited: attribute substitution in intuitive judgment. In Thomas Gilovich, Dale Griffin and Daniel Kahneman, editors, *Heuristics and Biases: The Psychology of Intuitive Judgment*, pages 49–81. Cambridge: Cambridge University Press.

Mascarenhas, Salvador (2014). *Formal Semantics and the Psychology of Reasoning: Building new bridges and investigating interactions*. Ph.D. thesis, New York University.

Nilson, Hakan, Peter Juslin and Henrik Olsson (2008). Exemplars in the mist: the cognitive substrate of the representativeness heuristic. *Scandinavian Journal of Psychology*, 49:201–212.

Tentori, Katya, Nicolao Bonini and Daniel Osherson (2004). The conjunction fallacy: a misunderstanding about conjunction? *Cognitive Science*, 28:467–477.

Tversky, Amos and Daniel Kahneman (1974). Judgment under uncertainty: heuristics and biases. *Science*, 185:1124–1131.

Tversky, Amos and Daniel Kahneman (1983). Extensional versus intuitive reasoning: the conjunction fallacy in probability judgment. *Psychological Review*, 90:293–315.