Clear Thinking in an Uncertain World: Human Reasoning and its Foundations

Lecture 9

Eric Pacuit

Department of Philosophy
University of Maryland, College Park
pacuit.org
epacuit@umd.edu

October 28, 2013

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?

- *B* Linda is a bank teller.
- *B & F* Linda is a bank teller and is active in the feminist movement.

Typically a large percentage of people asked say 2 is more probable than 1.

A. Tversky and D. Kahneman. *Extensions versus intuitive reasoning: The conjunction fallacy in probability judgment*. Psychological Review 90 (4): 293 - 315, 1983.

- In the original experiment, 85% of the participants judged B & F to be more likely than B. This contradicts the probability calculus, or so it seems ("conjunction fallacy").
- What shall we conclude? That 85% of us are *irrational*?
- Tversky and Kahneman's studies triggered a tremendous amount of work in Cognitive Psychology, but also in Philosophy as the issue of rationality is at stake.

The motivating assumption behind the corresponding research program is that we are doing quite well in our ordinary reasoning, and so examples such as the Linda case suggest that we should reconsider our theory of rationality and perhaps come up with an alternative that includes non-empirical and empirical considerations.

So how can the experimental findings be explained? Here are four proposals:

1. People implicitly add "and not a feminist" to proposition B.
2. People have problems with the notion of probability. If one uses frequencies instead, the effect will disappear. In fact, the number of people who commit the conjunction fallacy goes down but the effect does not disappear.
3. People do not read '&' in B & F as the logical operator $\wedge$.

K. Tentori, N. Bonini and D. Osherson. *The conjunction fallacy: a misunderstanding about conjunction?*. Cognitive Science, 28, pgs. 467 - 477, 2004.

4. People ask which of the two propositions B and B & F is better confirmed by the background story.

V. Crupi, B. Fitelson and K. Tentori. *Probability, confirmation and the conjunction fallacy*. Thinking and Reasoning 14, pp. 182-199, 2008.

4. People ask which of the two propositions B and $B \,\&\, F$ is better confirmed by the background story.

V. Crupi, B. Fitelson and K. Tentori. *Probability, confirmation and the conjunction fallacy*. Thinking and Reasoning 14, pp. 182-199, 2008.

5. Assume that the participants in the experiments address the following question: *Which of the two options (i.e. B and $B \,\&\, F$) is more probable given that a partially reliable source (i.e. the experimenter) informs you about them?*

L. Bovens and S. Hartmann. *Bayesian Epistemology*. Oxford University Press, 2003.

Probability

Kolmogorov Axioms:

1. For each E, $0 \leq p(E) \leq 1$
2. $p(W) = 1$, $p(\emptyset) = 0$
3. If $E_1, \ldots, E_n, \ldots$ are pairwise disjoint ($E_i \cap E_j = \emptyset$ for $i \neq j$), then $p(\bigcup_i E_i) = \sum_i p(E_i)$

Probability

Kolmogorov Axioms:

1. For each E, $0 \leq p(E) \leq 1$
2. $p(W) = 1$, $p(\emptyset) = 0$
3. If $E_1, \ldots, E_n, \ldots$ are pairwise disjoint ($E_i \cap E_j = \emptyset$ for $i \neq j$), then $p(\bigcup_i E_i) = \sum_i p(E_i)$

- $p(\overline{E}) = 1 - p(E)$ ($\overline{E}$ is the complement of E)
- If $E \subseteq F$ then $p(E) \leq p(F)$
- $p(E \cup F) = p(E) + p(F) + p(E \cap F)$

Conditional Probability

The probability of E *given* F, denoted $p(E|F)$, is defined to be

$$p(E|F) = \frac{p(E \cap F)}{p(F)} = \frac{p(E, F)}{p(F)}$$

Question: Derive the above equation from the following assumptions:

1. $p(\cdot \mid E)$ is a probability measure
2. $p(E \mid E) = 1$
3. If $F_1, F_2 \subseteq E$, then $\frac{p(F_1 \mid E)}{p(F_2 \mid E)} = \frac{p(F_1)}{p(F_2)}$

Bayes Theorem

Bayes Theorem: $p(E \mid F) = \frac{p(F \mid E)p(E)}{p(F)}$

Bayes Theorem

Bayes Theorem: $p(E \mid F) = \frac{p(F \mid E)p(E)}{p(F)}$

$$= \frac{p(F \mid E)p(E)}{p(F \mid E)p(E) + P(F \mid \neg E)P(\neg E)}$$

Bayes Theorem

Bayes Theorem: $p(E \mid F) = \frac{p(F \mid E)p(E)}{p(F)}$

$$= \frac{p(F \mid E)p(E)}{p(F \mid E)p(E)+P(F \mid \neg E)P(\neg E)}$$

$$= \frac{p(E)}{p(E)+p(\neg E)x} \text{ where } x = \frac{p(F \mid \neg E)}{p(F \mid E)} \text{ (the likelihood ratio)}$$

Independence

E and F are **independent** iff $p(E, F) = p(E)p(F)$

Independence

E and F are **independent** iff $p(E, F) = p(E)p(F)$ iff $p(E \mid F) = p(E)$ iff $p(F \mid E) = p(F)$

Independence

Suppose Alice has a coin that she knows is either fair or double-headed. Either possibility seems equally likely, so she assigns each a probability of $1/2$. She then tosses the coin twice.

Independence

Suppose Alice has a coin that she knows is either fair or double-headed. Either possibility seems equally likely, so she assigns each a probability of $1/2$. She then tosses the coin twice. Is the event that the first coin landed heads independent of the event that the second coin toss lands heads?

Conditional Independence

E is **conditionally independent** of F given C iff
$p(E \mid F, C) = p(E \mid C)$.

Example: $A =$ yellow fingers, $B =$ lung cancer, $C =$ smoking. A and B are positively correlated, i.e. learning that a person has A raises the probability of B. Yet, if we know C, A leaves the probability of B unchanged.

C is called the *common cause* of A and B

A Few Observations

The Chain Rule: $p(E, F) = p(E \mid F)p(F)$

A Few Observations

The Chain Rule: $p(E, F) = p(E \mid F)p(F)$

Properties of Independence: Suppose that $p(A \mid B) = p(A \mid B, C)$. Then,

A Few Observations

The Chain Rule: $p(E, F) = p(E \mid F)p(F)$

Properties of Independence: Suppose that $p(A \mid B) = p(A \mid B, C)$. Then,

- $p(\neg A \mid B) = p(\neg A \mid B, C)$

A Few Observations

The Chain Rule: $p(E, F) = p(E \mid F)p(F)$

Properties of Independence: Suppose that $p(A \mid B) = p(A \mid B, C)$. Then,

- $p(\neg A \mid B) = p(\neg A \mid B, C)$
- $p(C \mid B) = p(C \mid B, A)$

Joint and Marginal Probability

The **joint probability** of two binary propositional variables A and B can be calculated from three values:

Example. $p(A, B) = 0.4$, $p(A, \neg B) = 0.2$, $p(\neg A, B) = 0.3$. Since $\sum_{A,B} p(A, B) = 1$, we have $p(\neg A, \neg B) = 0.1$.

In general $2^n - 1$ values have to be specified to give the joint probability over n propositional variables.

marginal probability: $p(A) = \sum_B p(A, B)$.

Bayes Networks

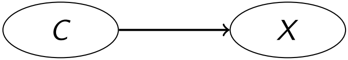

C “causes” X

C “directly influences X

A Few Details

Bayes networks are *directed acyclic graphs* (nodes with directed edges that do not form a cycle) with a probability distribution respecting the Parental Markov Condition (a variable is conditionally independent of its non-descendents given its parents)

Judea Pearl's Example

I'm at work, neighbor John calls to say my alarm is ringing, but my neighbor Mary doesn't call. Sometimes it's set off by minor earthquakes. Is there a burglar?

Variables: Burglar (B), Earthquake (E), Alarm (A), JohnCalls (J), MaryCalls (M).

Network topology reflects "causal" knowledge:

- A burglar can set the alarm off
- An earthquake can set the alarm off
- The alarm can cause Mary to call
- The alarm can cause John to call

$p(B) = .001$

$p(E) = .002$

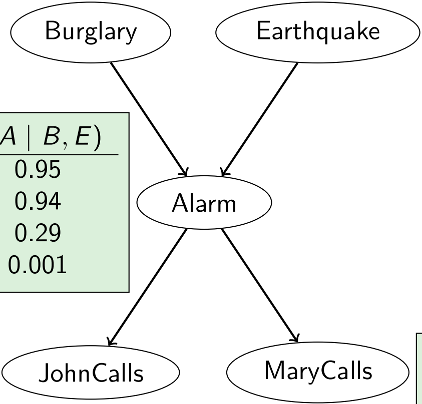

B	E	$p(A \mid B, E)$
T	T	0.95
T	F	0.94
F	T	0.29
F	F	0.001

A	$p(J \mid A)$
T	0.9
F	0.05

A	$p(M \mid A)$
T	0.7
F	0.01

$p(B) = .001$

$p(E) = .002$

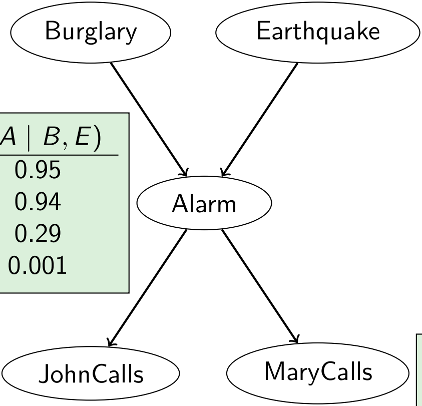

B	E	$p(A \mid B, E)$
T	T	0.95
T	F	0.94
F	T	0.29
F	F	0.001

A	$p(J \mid A)$
T	0.9
F	0.05

A	$p(M \mid A)$
T	0.7
F	0.01

$$p(J, M, A, \neg B, \neg E) = p(J \mid A)p(M \mid A)p(A \mid \neg B \neg E)p(\neg B)p(\neg E)$$

Three Typical Examples

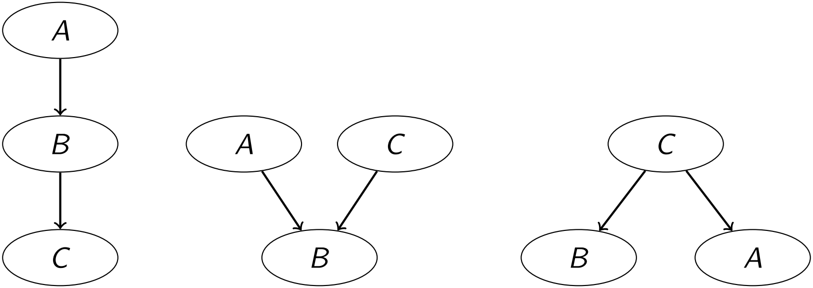

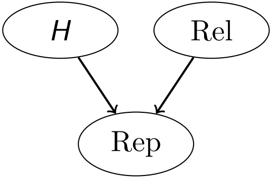

$$
\begin{aligned}
p(H) &= h \\
p(\mathrm{Rel}) &= r \\
p(\mathrm{Rep} \mid H, \mathrm{Rel}) &= 1 \\
p(\mathrm{Rep} \mid \neg H, \mathrm{Rel}) &= 0 \\
p(\mathrm{Rep} \mid H, \neg \mathrm{Rel}) &= a \\
p(\mathrm{Rep} \mid \neg H, \neg \mathrm{Rel}) &= a
\end{aligned}
$$

$$p(H \mid \mathrm{Rep}) =$$

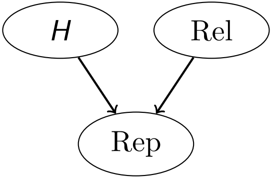

$$
\begin{aligned}
p(H) &= h \\
p(\mathrm{Rel}) &= r \\
p(\mathrm{Rep} \mid H, \mathrm{Rel}) &= 1 \\
p(\mathrm{Rep} \mid \neg H, \mathrm{Rel}) &= 0 \\
p(\mathrm{Rep} \mid H, \neg \mathrm{Rel}) &= a \\
p(\mathrm{Rep} \mid \neg H, \neg \mathrm{Rel}) &= a
\end{aligned}
$$

$$p(H \mid \mathrm{Rep}) = \frac{p(H, \mathrm{Rep})}{p(\mathrm{Rep})} =$$

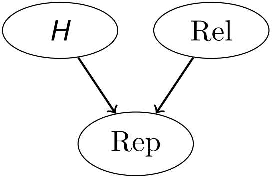

$$
\begin{aligned}
p(H) &= h \\
p(\mathrm{Rel}) &= r \\
p(\mathrm{Rep} \mid H, \mathrm{Rel}) &= 1 \\
p(\mathrm{Rep} \mid \neg H, \mathrm{Rel}) &= 0 \\
p(\mathrm{Rep} \mid H, \neg \mathrm{Rel}) &= a \\
p(\mathrm{Rep} \mid \neg H, \neg \mathrm{Rel}) &= a
\end{aligned}
$$

$$
p(H \mid \mathrm{Rep}) = \frac{p(H, \mathrm{Rep})}{p(\mathrm{Rep})} = \frac{\sum_{\mathrm{Rel}} p(H, \mathrm{Rel}, \mathrm{Rep})}{\sum_{H, \mathrm{Rel}} p(H, \mathrm{Rel}, \mathrm{Rep})}
$$

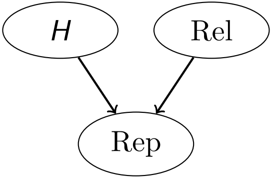

$$\begin{aligned}
p(H) &= h\\
p(\mathrm{Rel}) &= r\\
p(\mathrm{Rep} \mid H, \mathrm{Rel}) &= 1\\
p(\mathrm{Rep} \mid \neg H, \mathrm{Rel}) &= 0\\
p(\mathrm{Rep} \mid H, \neg\mathrm{Rel}) &= a\\
p(\mathrm{Rep} \mid \neg H, \neg\mathrm{Rel}) &= a
\end{aligned}$$

$$p(H \mid \mathrm{Rep}) = \frac{p(H, \mathrm{Rep})}{p(\mathrm{Rep})} = \frac{\sum_{\mathrm{Rel}} p(H, \mathrm{Rel}, \mathrm{Rep})}{\sum_{H,\mathrm{Rel}} p(H, \mathrm{Rel}, \mathrm{Rep})}$$

$$= \frac{p(H) \sum_{\mathrm{Rel}} p(\mathrm{Rel}) p(\mathrm{Rep} \mid H, \mathrm{Rel})}{\sum_{H,\mathrm{Rel}} p(H) p(\mathrm{Rel}) p(\mathrm{Rep} \mid H, \mathrm{Rel})} =$$

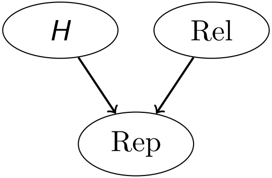

$$
\begin{aligned}
p(H) &= h\\
p(\mathrm{Rel}) &= r\\
p(\mathrm{Rep} \mid H, \mathrm{Rel}) &= 1\\
p(\mathrm{Rep} \mid \neg H, \mathrm{Rel}) &= 0\\
p(\mathrm{Rep} \mid H, \neg \mathrm{Rel}) &= a\\
p(\mathrm{Rep} \mid \neg H, \neg \mathrm{Rel}) &= a
\end{aligned}
$$

$$p(H \mid \mathrm{Rep}) = \frac{p(H, \mathrm{Rep})}{p(\mathrm{Rep})} = \frac{\sum_{\mathrm{Rel}} p(H, \mathrm{Rel}, \mathrm{Rep})}{\sum_{H,\mathrm{Rel}} p(H, \mathrm{Rel}, \mathrm{Rep})}$$

$$= \frac{p(H) \sum_{\mathrm{Rel}} p(\mathrm{Rel}) p(\mathrm{Rep} \mid H, \mathrm{Rel})}{\sum_{H,\mathrm{Rel}} p(H) p(\mathrm{Rel}) p(\mathrm{Rep} \mid H, \mathrm{Rel})} = \frac{h(r + a\bar{r})}{hr + a\bar{r}}$$

L. Bovens and S. Hartmann. *Bayesian Epistemology*. Oxford University Press, 2003.

Which of the two options (i.e. B and B & F) is more probable given that a partially reliable source (i.e. the experimenter) informs you about them?

Assume that the participants compare the conditional probabilities $Pr(B, F \mid \mathrm{Rep}_B, \mathrm{Rep}_F)$ and $Pr(B \mid \mathrm{Rep}_B)$ Note that both condition on different background information, and so $Pr(B, F \mid \mathrm{Rep}_B, \mathrm{Rep}_F)$ can be larger than $Pr(B \mid \mathrm{Rep}_B)$.

Assume that the variables B and F are probabilistically independent and that each proposition is uttered by a partially reliable witness.

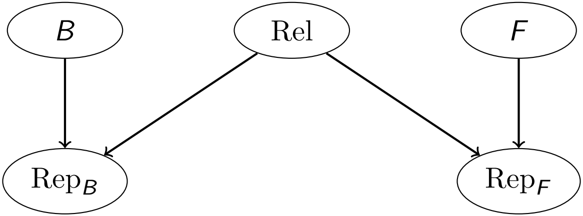
B
Rel
F
Rep_B
Rep_F

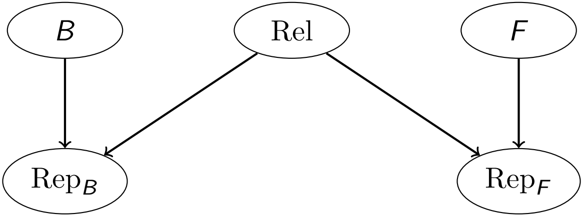

One can then show that $Pr(B, F \mid \mathrm{Rep}_B, \mathrm{Rep}_F) > Pr(B \mid \mathrm{Rep}_B)$ if $Pr(F) > Pr(B)$ and another (arguably plausible) condition holds

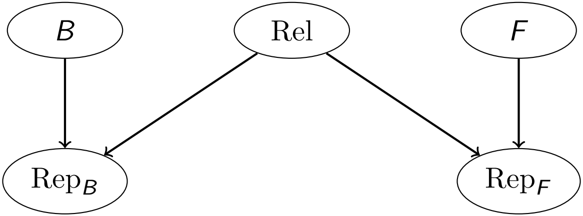

One can then show that $Pr(B, F \mid \mathrm{Rep}_B, \mathrm{Rep}_F) > Pr(B \mid \mathrm{Rep}_B)$ if $Pr(F) > Pr(B)$ and another (arguably plausible) condition holds

Specifically, if $p(B) = b$, $p(F) = f$, the reliability parameter is ρ and the relevant likelihoods are a, then $p(B, F \mid \mathrm{Rep}_B, \mathrm{Rep}_F) - p(B \mid \mathrm{Rep}_B) > 0$ when

$$a^2\overline{f} + \overline{a}\rho(a - f(a + \overline{b})) < 0$$

S. Hartmann and W. Meijs. *Walter the banker: the conjunction fallacy reconsidered*. Synthese, 184, pgs. 73 - 87, 2012.

- It does not seem to be correct that B and F are independent.
- In fact, both variables are negatively relevant given the background story S about Linda (i.e. that she is in her early thirties,...).
- *The Sophisticated Witness Model* takes the background story S into account. S is also uttered by the experimenter, and the participants consider S in their probabilistic judgment.
- The participants in a psychological experiment expect to be fooled. So why should they take as certain what the experimenter says?

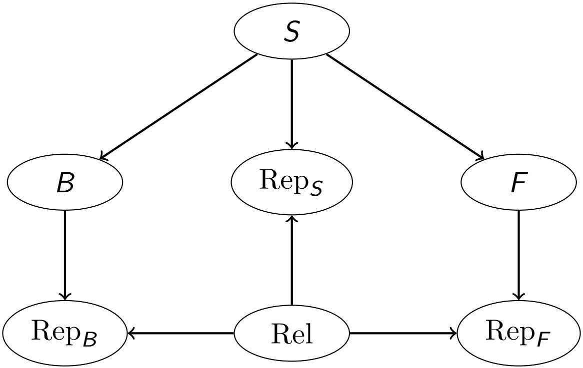
S
B
Rep_S
F
Rep_B
Rel
Rep_F

Walter the Banker

Suppose you are a philosopher who knows a little, but not much about the current status of formal epistemology. Suppose also that you are intimately connected with the university of Leuven, and that through this you know that there will be a conference on probabilistic paradoxes next week. Suppose that on one morning you meet a stranger in the commuter train to his work and you start talking.

Walter the Banker, continued

Now suppose that when the stranger finds out that you are a philosopher, he starts talking about an acquaintance of his, who happens to be a philosopher as well. In summary, you learn that this guy's name is Walter, that he is 31 years old, outspoken, bright, with a background in philosophy and physics and did some work on probabilistic modeling. The stranger continues to make some comments about the merits of formal epistemology in general and then proceeds to inform you that, by the way, Walter is a banker, working as a manager for one of the world's largest financial services providers.

Walter the Banker, continued

Now suppose you arrive at your stop and after saying goodbye, you leave the train. Call this scenario 1. Now consider scenario 2, which is like scenario 1, but before you say your goodbyes, the stranger mentions one more fact: that Walter will be speaking at a conference in Leuven next week.

Walter Formalized

- H_1 Walter is 31 years old, outspoken, bright, and has a background in physics and philosophy.
- H_2 Walter works at a bank.
- H_3 Walter attends a conference on formal epistemology in Leuven.

Claim. It can be the case that

$$p(H_1, H_2, H_3 \mid \mathrm{Rep}_{H_1}, \mathrm{Rep}_{H_2}, \mathrm{Rep}_{H_3}) > p(H_1, H_2 \mid \mathrm{Rep}_{H_1}, \mathrm{Rep}_{H_2})$$