Implicatures in Bayesian Dialogues

Christina Pawlowitsch

Université Paris–Panthéon–Assas

Paris Workshop on Games, Decisions, and Language

June 15–16, 2023

Bayesian dialogues

Two individuals report their Bayesian posterior about a certain event back and forth,

- at each step taking into account the additional information contained in the posterior announced by the other at the previous step,
- under the assumption that they have *common knowledge* (Lewis 1969, Aumann 1976) about the prior and the information structure (partition) by which they receive private information.

Converges to a commonly known posterior. Bacharach (1979) shows this for normally distributed random variables.

Geanakoplos and Polemarchakis (1982) show this within a model based on information partitions as used by Aumann (1976).

$\Rightarrow$ dynamic foundation for Aumann's (1976) "agreement" result.

Example of a Bayesian dialogue à la Geanakoplos and Polemarchakis (1982)

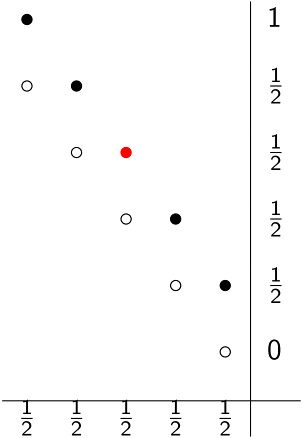

States belonging to A are indicated as filled circles; states not belonging to A as empty circles. The true state belongs to A and is indicated in red.

Example

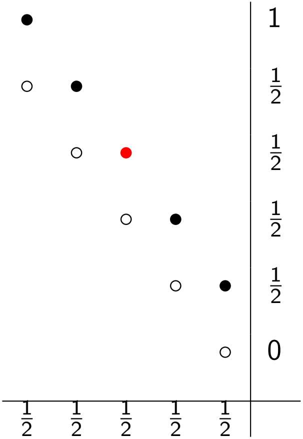

Row says: $1/2$

Example

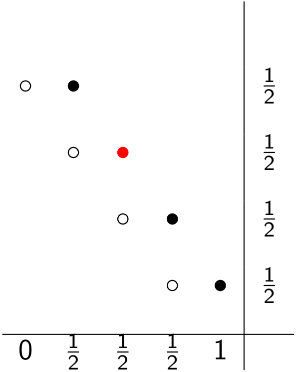

Row says: $1/2$

Column says: $1/2$

Example

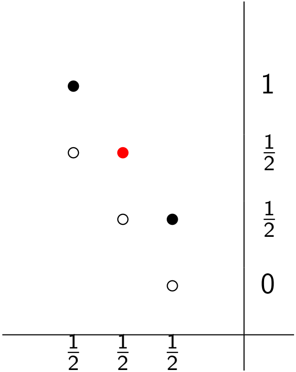

Row says:	$1/2$
Column says:	$1/2$
Row says:	$1/2$

Example

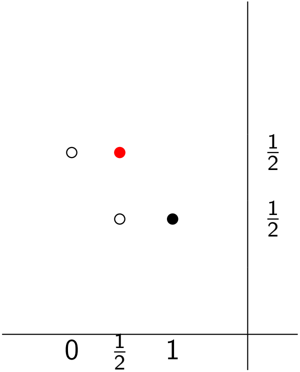

Row says:	$1/2$
Column says:	$1/2$
Row says:	$1/2$
Column says:	$1/2$

Example

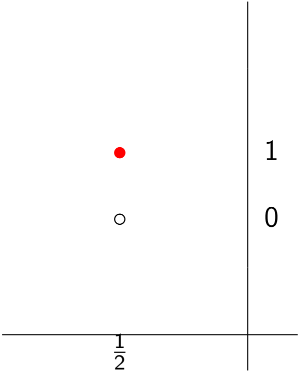

Row says:	$1/2$
Column says:	$1/2$
Row says:	$1/2$
Column says:	$1/2$
Row says:	1

Example

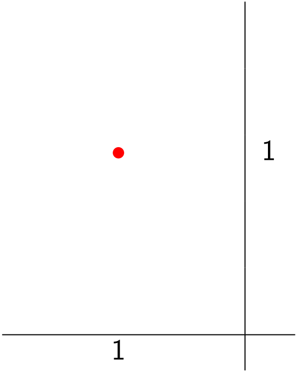

Row says:	$1/2$
Column says:	$1/2$
Row says:	$1/2$
Column says:	$1/2$
Row says:	1
Column says:	1

Will there be another pandemic within the next two years ?

Individual 1: “20 sure” %

Individual 2: “80% sure”

Individual 1: “30% sure”

Individual 2: “60% sure”

Individual 1: “20% sure”

Individual 2: “60% sure”

Individual 1: “20% sure”

Individual 2: “20% sure”

Individual 1: “20% sure”

Individual 2: “20% sure”

Will Trump be the next president of the United States?

Individual 1: “60% sure”

Individual 2: “98% sure”

Individual 1: “90 % sure”

Individual 2: “90% sure”

Individual 1: “80% sure”

Individual 2: “80% sure”

Individual 1: “80% sure”

Individual 2: “80% sure”

Relation to linguistics

A Bayesian dialogue is like asking repeatedly—and in an alternating manner—the same question to two different individuals.

- Hamblin, C. L. (1958). Questions. The Australasian Journal of Philosophy, 36(3): 159–168. [partitions]
- Hamblin, C. L. (1973). Questions in Montague English. Foundations of Language, 10(1): 41–53. [alternative sets as denotations of questions]
- Groenendijk, J. and Stokhof, M. (1984). Studies in the Semantics of Questions and the Pragmatics of Answers. PhD thesis, Universiteit van Amsterdam, Amsterdam.
- Lewis, D. (1988). Relevant implication. Theoria, 54: 161–174.

Properties of the model

The model is not a "game." Individuals are not endowed with a payoff function over the profile of probabilities attributed to the even of interest (the "question under discussion"). The model describes an algorhythm.

Obvious: If diverging preferences over profile of probabilities $\Longrightarrow$ truthfully reporting one's Bayesian posterior might not be optimal.

Intuition (one could have): Always truthfully reporting one's Bayesian posterior, the best the two individuals can do if they have perfectly coinciding interests and these are in line with "wanting to know the truth."

Starting point here: example showing this intuition is wrong.

Example

Let $\Omega = \{a, b, c, d, e, f, g, h, i, j, k\}$, with uniform prior, $\frac{1}{11}$.
Information partitions:

$$\begin{aligned} \mathscr{P}_1 &= \{\{a, b, c, d, e, f\}, \{g, h, i, j, k\}\}, \\ \mathscr{P}_2 &= \{\{a, b, g, h\}, \{c, d, i, j\}, \{e, f, k\}\}. \end{aligned}$$

$A = \{a, b, i, j, k\}$, the event of interest. Suppose $\omega^\star = a$, the true state.

$\{\mathbf{a}^\star, \mathbf{b}\}$	$\{c, d\}$	$\{e, f\}$	$\frac{1}{3}$
$\{g, h\}$	$\{\mathbf{i}, \mathbf{j}\}$	$\{\mathbf{k}\}$	$\frac{3}{5}$
$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{3}$	

Meet, finest common coarsening:
$\mathscr{P}_1 \wedge \mathscr{P}_2 = \{a, b, c, d, e, f, g, h, i, j, k\}$
Join, coarsest common refinement:
$\mathscr{P}_1 \vee \mathscr{P}_2 = \{\{a, b\}, \{c, d\}, \{e, f\}, \{g, h\}, \{i, j\}, \{k\}\}.$

If individual 1 starts:

$$
\begin{array}{ccc|c}
\{\mathbf{a}^\star, \mathbf{b}\} & \{c, d\} & \{e, f\} & \frac{1}{3} \\
\{g, h\} & \{\mathbf{i}, \mathbf{j}\} & \{\mathbf{k}\} & \frac{3}{5} \\
\hline
\frac{1}{2} & \frac{1}{2} & \frac{1}{3} &
\end{array}
$$

$$q_1 = \frac{p(\{a, b\})}{p(\{a, b, c, d, e, f\})} = \frac{1}{3}$$

If individual 1 announces $1/3$, it will become common knowledge that the true state cannot belong to $\{g, h, i, j, k\}$, and therefore should be deleted from the *fund of common knowledge*.

Step 2:

$\{\mathbf{a}^{\star}, \mathbf{b}\}$	$\{c, d\}$	$\{e, f\}$	$\frac{1}{3}$
1	0	0	

$$q_2 \;=\; \frac{p(\{a, b\})}{p(\{a, b\})} = 1$$

If individual 2 announces 1, then it will become common knowledge that the true state cannot belong to $\{c, d, e, f\}$, and therefore should be deleted from the *fund of common knowledge*.

Step 3:

$\{\mathbf{a}^{\star}, \mathbf{b}\}$	1
1	

Individual 1 announces also "1," and the process has reached its absorbing state.

If individual 2 starts:

Step 1:

$$
\begin{array}{ccc|c}
\{\mathbf{a}^\star, \mathbf{b}\} & \{c, d\} & \{e, f\} & \frac{1}{3} \\
\{g, h\} & \{\mathbf{i}, \mathbf{j}\} & \{\mathbf{k}\} & \frac{3}{5} \\
\hline
\frac{1}{2} & \frac{1}{2} & \frac{1}{3} &
\end{array}
$$

$$
q_2 = \frac{p(\{a, b\})}{p(\{a, b, g, h\})} = \frac{1}{2}
$$

Step 2:

$$
\begin{array}{cc|c}
\{\mathbf{a}^\star, \mathbf{b}\} & \{c, d\} & \frac{1}{2} \\
\{g, h\} & \{\mathbf{i}, \mathbf{j}\} & \frac{1}{2} \\
\hline
\frac{1}{2} & \frac{1}{2} &
\end{array}
$$

The process ends here, with each of them announcing $1/2$, from this moment on forever.

Aumann's (1976) result: the "Aumann conditions"

Let $(\Omega, \mathscr{B}, p)$ a probability space, $\mathscr{P}_1$ and $\mathscr{P}_2$ two finite partitions of Ω representing the information accessible to individual 1 respectively 2, and $A \in \mathscr{B}$ an event—all of this being common knowledge. If at state ω the posteriors q_1 and q_2 that the individuals attribute to A are common knowledge, then: $q_1 = q_2$.

Illustrated in the example:

$$
\begin{array}{cc|c}
\{\mathbf{a}^\star, \mathbf{b}\} & \{c, d\} & q_{1,1} = \frac{1}{2} \\
\{g, h\} & \{\mathbf{i}, \mathbf{j}\} & q_{1,2} = \frac{1}{2} \\
\hline
q_{2,1} = \frac{1}{2} & q_{2,2} = \frac{1}{2} &
\end{array}
\qquad
\begin{array}{lcl}
q_1 \text{ CK} & \rightarrow & q_{1,1} = q_{1,2} = q_1 \\
q_2 \text{ CK} & \rightarrow & q_{2,1} = q_{2,2} = q_2 \\
\hline
 & \Rightarrow & q_1 = q_2 \\
 & & \text{“Tower Property”}
\end{array}
$$

Absorbing conditions

- A Bayesian dialogue always ends with the “Aumann” conditions: a subset of Ω on which rules a commonly known posterior.
- However: depending on the order of communication,
  it can end with different subsets of Ω, and hence
  with different commonly known posteriors.

If individual 1 starts:

Step 1:

$\{\mathbf{a}^\star, \mathbf{b}\}$	$\{c, d\}$	$\{e, f\}$	$\frac{1}{3}$
$\{g, h\}$	$\{\mathbf{i}, \mathbf{j}\}$	$\{\mathbf{k}\}$	$\frac{3}{5}$
$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{3}$	

Step 2:

$\{\mathbf{a}^\star, \mathbf{b}\}$	$\{c, d\}$	$\{e, f\}$	$\frac{1}{3}$
1	0	0	

Step 3:

$\{\mathbf{a}^\star, \mathbf{b}\}$	1
1	

If individual 2 starts:

Step 1:

$\{\mathbf{a}^\star, \mathbf{b}\}$	$\{c, d\}$	$\{e, f\}$	$\frac{1}{3}$
$\{g, h\}$	$\{\mathbf{i}, \mathbf{j}\}$	$\{\mathbf{k}\}$	$\frac{3}{5}$
$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{3}$	

Step 2:

$\{\mathbf{a}^\star, \mathbf{b}\}$	$\{c, d\}$	$\frac{1}{2}$
$\{g, h\}$	$\{\mathbf{i}, \mathbf{j}\}$	$\frac{1}{2}$
$\frac{1}{2}$	$\frac{1}{2}$	

Is it always rational to speak Bayesian?

Embedding the example in a story – turning it into a game

Suppose there are two professional chess players (game theorists) thrown in prison. The director of the prison calls on them and announces:

"Here are Ω, your information partitions, the prior p, and the event of interest A:

$$\{\mathbf{a}, \mathbf{b}\} \quad \{c, d\} \quad \{e, f\}$$

$$\{g, h\} \quad \{\mathbf{i}, \mathbf{j}\} \quad \{\mathbf{k}\}$$

A state of the world will materialize and each of you will receive information according to his partition.

Then, I will ask one of you, in front of the other:

> *What is the probability that you attribute to the event A?*

After all of us having heard his answer, I will ask the other:

> *What is the probability that you attribute to the event A?*

After all of us having heard his answer, I will turn to the first again and ask:

> *Did A happen or not?*

If his answer is correct, both of you will get free. If not, both of you will sit for the rest of your lives."

The director first calls on the column individual and asks:

> *"What is the probability that you attribute to the event A?"*

$$
\begin{array}{ccc|c}
\{\mathbf{a}^\star, \mathbf{b}\} & \{c, d\} & \{e, f\} & \frac{1}{3} \\[2ex]
\{g, h\} & \{\mathbf{i}, \mathbf{j}\} & \{\mathbf{k}\} & \frac{3}{5} \\
\hline
\mathbf{\frac{1}{2}} & \frac{1}{2} & \frac{1}{3} &
\end{array}
\quad \Rightarrow \quad
\begin{array}{cc|c}
\{\mathbf{a}^\star, \mathbf{b}\} & \{c, d\} & \frac{1}{2} \\[2ex]
\{g, h\} & \{\mathbf{i}, \mathbf{j}\} & \frac{1}{2} \\
\hline
\frac{1}{2} & \frac{1}{2} &
\end{array}
$$

The column individual truthfully says: "$\frac{1}{2}$"

The director then calls on the row individual and asks what is the probability that she attributes to the event A. Individual 1 says ...

What do you think that the row individual says?

Imagine you are the column individual and, at step 2, hear the row individual say: "1/3."

You know that this *cannot* possibly be the truthful Bayesian posterior of individual 1—and this is common knowledge.

But before your announcement at the first step, 1/3 was a possible truthful Bayesian response of individual 1, implying that he has received the information $\{a, b, c, d, e, f\}$)—and this is also common knowledge.

Hearing "1/3," you understand that this is the message that the row individual wanted to convey. You understand that the true state belongs to $\{a, b\}$. You announce "1," and both of you get free.

If you are the row individual, you understand that this is the way that the column individual would reason. Anticipating that, at the second step, you say:

$$\text{“}1/3.\text{”}$$

Game theorists and philosophers of language might cry out and say:

"But this is

... *forward induction*"!

... a *conversational implicature*"!

A linguistic interpretation

A speech act of which it is commonly known that it cannot be possibly true

In the story above, the profitable deviation from truthfully stating one's Bayesian updated belief thrives on the fact that *by doing the deviation*, it will become *common knowledge* that the announced probability cannot possibly be the speaker's Bayesian updated belief at this step.

That she has deviated from the rule—if such a rule can be assumed to be in place—of truthfully stating her Bayesian updated belief at that step.

One can recognize in this movement a *conversational implicature* (Grice 1975):

the phenomenon that the meaning of a speech act (here the announced probability) will be implied by a deviation from some predefined convention how to talk under normal conditions—what Grice calls the “flouting” of a conversational maxim.

Under the name of the “Cooperative Principle,” Grice isolates four main conversational maxims:

- Quantity: “1. Make your contribution as informative as required (for the current purpose of the exchange). 2. Do not make your contribution more informative than is required.”
- Quality: “Try to make your contribution one that is true.”
  - “1. Do not say what you belief to be false.”
  - “2. Do not say that for which you do lack adequate evidence.”
- Relation: “Be relevant.”
- Manner: “Be perspicuous.”
  - “1. Avoid obscurity of expression.”
  - “2. Avoid ambiguity.”
  - “3. Be brief (avoid unnecessary prolixity).”
  - “4. Be orderly.”

In the example above, the maxim flouted can be said to be that of Quality, which here takes the specific form that one ought to truthfully announce one's Bayesian updated belief at the current state of a conversation.

To "flout" a maxim, says Grice, is to "blatantly fail to fulfill it."

"Blatantly" her comes to mean *in the face of common knowledge*.

In the example above, the implicature can also be said to be triggered by a "clash" of maxims:

In a situation in which the Bayesian updated beliefs of the two individuals are already common knowledge, reporting one's Bayesian updated belief amounts to making a perfectly *irrelevant* speech act: there is a clash between *Quality* and *Relation*.

If the row individual, at step 2, announces her original posterior, she can be said to sacrifice *quality* in order to save *relevance*.

Conversational Implicature (Grice 1975)

The phenomenon that the meaning of a speech act, here the announced probability, will be implied by a deviation from some predefined convention how to talk under normal conditions—what Grice calls the "flouting" of a conversational maxim.

The maxim flouted here: "Always truthfully state your Bayesian posterior."

Under the name of the "Cooperative Principle," Grice isolates four main conversational maxims:

- Quantity: "1. Make your contribution as informative as required (for the current purpose of the exchange). 2. Do not make your contribution more informative than is required."
- Quality: "Try to make your contribution one that is true."
- Relation: "Be relevant."
- Manner: "Be perspicuous."

Maxim flouted here: Quality, which takes the specific form that one ought to truthfully announce one's Bayesian posterior at the current state of a conversation.

Based on:

PAWLOWITSCH (2021),

"Strategic Manipulation in Bayesian Dialogues,"

Synthese 199: 11279–11303.

Literature

- Aumann. 1976. "Agreeing to disagree," *The Annals of Statistics*.
- Bacharach. 1979. "Normal bayesian dialogues." *Journal of the American Statistical Association*.
- Geanakoplos and Polemarchakis. 1982. "We can't disagree forever," *Journal of Economic Theory*.
- Grice. 1975. "Logic and conversation." In Cole, P. & Morgan, J. L. (Eds.), *Syntax and Semantics*, Vol. 3, Speech Acts (pp. 41–58). New York: Academic Press.
- Sebenius and Geanakoplos. 1983. "Don't bet on it : contingent agreements with asymmetric information," *Journal of the American Statistical Association*.
- Fagin, Halpern, Moses, and Vardi. 1995. *Reasoning about Knowledge.* MIT Press (first published as an IBM mimeo, 1988).

Forerunners of Bayesian Dialogues:

- Lacan, 1945. "Le temps logique et l'assertion de certitude anticipée : un nouveau sophisme," *Les Cahiers d'Art, 1940–1944*; reprinted in: *Ecrits* (1966). Le Seuil: Paris.
- Littlewood, 1953. *A Mathematician's Miscellany*, Methuen & Co, London.
- Dalkey, 1969. "The Delphi method: an experimental study of group opinion," *United States Air Force Project Rand*.

- Dalkey, 1969. "The Delphi method: an experimental study of group opinion," *United States Air Force Project Rand*.

More recent work:

- POLEMARCHAKIS, 2016. "Rational dialogs." Working Paper.
- GEANAKOPLOS AND POLEMARCHAKIS, 2023. "Rational dialogs." Forthcoming in Revue Economique.
- DI TILIO, LEHRER, and SAMET. 2022. "Monologues, dialogues, and common priors." *Theoretical Economics*.