Is Rational Disagreement Possible?

DeGroot, Lehrer, and Aumann's Theorems

Models and Simulations in Philosophy
May 5th, 2014

Review

Question: "Peer" Disagreement: Can rational individuals with the same evidence disagree?

We've discussed two answers.

Review

Answer 1: No.

- "Equal-weight view" = In the face of disagreement,
  - Suspend judgment
  - Adjust probabilistic credences appropriately (often, by averaging)
- Motivated by **uniqueness thesis**, which asserts
  - There is a unique state of belief (or agnotisicism) warranted by one's evidence

Today's Class

Questions:

- What should one believe when there are several people who disagree?
- Does favoring one's own views actually permit disagreement?
- What about the more common cases in which there is non-shared evidence?

Today's Class

Today: Two formal models of disagreement.

- **Aumann** - Is disagreement possible if evidence is not shared?
- **DeGroot/Lehrer** - Is disagreement possible among multiple individuals even when non-equal weight is given to different peers?

Outline

The Scene

The Scene: Mr. Boddy is found dead at the scene of a crime.

The Scene: There are six suspects.

The Detectives

There are two detectives on the scene:

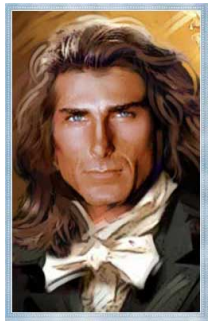

Detective Farbeo

Detective Kleidum

Their Evidence

The two detectives take evidence back to their respective labs:

- Detective Farbeo analyzes the color of hair of killer, and
- Detective Kleidum analyzes the fabric of killer's clothing.

Farbeo's Evidence

Brown

Black

White

Kleidum's Evidence

Cotton

Polyester

Farbeo's Evidence

Brown | Black | White

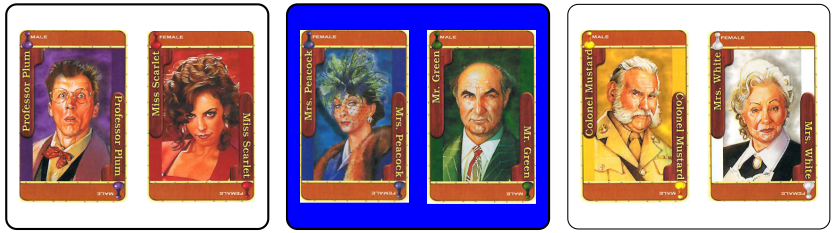

Suppose Farbeo learns the killer's hair is **black**.

Kleidum's Evidence

Cotton | Polyester

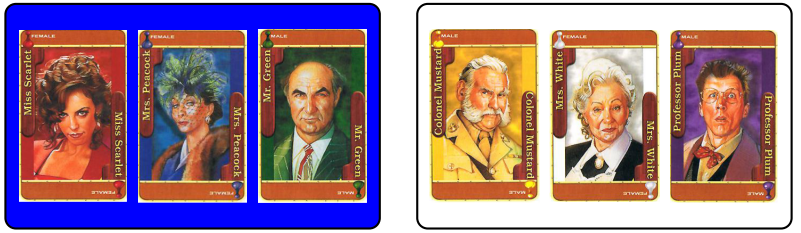

Kleidum learns the killer was wearing **cotton.**

Their Evidence

- To avoid biasing their findings, the detectives are **not** permitted to share their evidence before trial.
- However, they are *repeatedly* asked at a press conference about their beliefs about the proposition
  - φ: " The killer's name begins with the letter 'P.' "

Their Evidence

Question: After the press conference, can the detectives disagree about the probability of φ?

Assumption 1: Our detectives are good Bayesian agents.

Bayesianism

Bayesianism is the conjunction of two theses:

- Beliefs = Probabilities
- Beliefs updated by conditionalization

Assumption 2: Before collecting evidence, the detectives agreed about the likelihood that each suspect was the killer.

For simplicity, let's assume they agreed each suspect was equally likely to be the killer, i.e.,

$$P_F(Scarlet) = P_F(Green) = \ldots = \frac{1}{6}$$

$$P_K(Scarlet) = P_K(Green) = \ldots = \frac{1}{6}$$

Farbeo's Beliefs

Brown Black White

Before Conference: $P_F(\varphi|\text{“}Black\ Hair\text{'}) = \frac{1}{2}$.

Kleidum's Beliefs

Cotton

Polyester

Before Conference: $P_K(\varphi|\text{“}Cotton\text{”}) = \frac{1}{3}$.

Before the press conference

Question: Can the detectives learn anything about φ from the press conference?

On first glance, the answer seems to be “no.”

Farbeo's Beliefs

Brown

Black

White

Before Conference: Farbeo knows the killer has black hair, and is therefore, either Peacock or Green.

Farbeo's Beliefs

Cotton

Polyester

Before Conference: Farbeo knows the killer wore cotton.

Farbeo's Beliefs

Cotton

Polyester

Before Conference: Farbeo knows that one in three suspects wearing cotton has a last name beginning with "P."

Farbeo's Perspective

Moral: Farbeo knows that Kleidum will announce $\frac{1}{3}$.

What about Kleidum?

Kleidum's Beliefs

Cotton

Polyester

Before Conference: Kleidum knows the killer wore Cotton, **and** so she can infer the killer is not White or Mustard.

Kleidum's Beliefs

Brown Black White

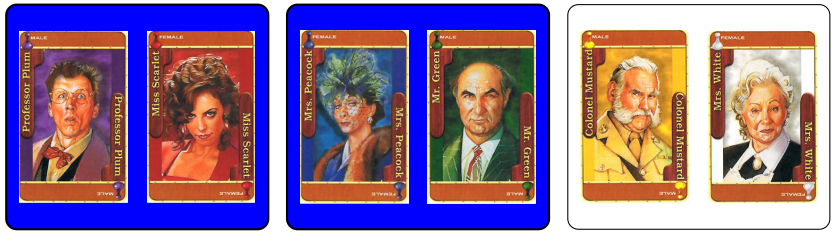

Before Conference: Kleidum knows the killer is not White or Mustard, **and** so she can infer the killer has black or brown hair.

Kleidum's Perspective

Hence, she knows that Farbeo has learned that the killer's hair color is either brown or black. **She cannot infer which color.**

Kleidum's Perspective

- Brown: One of two suspects with brown hair has a last name beginning with "P": Professor Plum.
  - So Farbeo will announce $\frac{1}{2}$.
- **Black:** One of two suspects with black hair has a last name beginning with "P": Mrs. Peacock.
  - So Farbeo will announce $\frac{1}{2}$.

Before the press conference

Moral: Farbeo and Kleidum both know what the other will announce before the press conference.

So it seems that neither can learn anything from hearing what the other has to say.

This appearance is deceiving . . .

Consensus

Suppose Farbeo and Kleidum are repeatedly asked how strongly they believe "The killer's name begins with the letter 'P.' "

I claim the following happens if they announce their views at the same time:

1. **1st Time:** Farbeo announces $\frac{1}{2}$ and Kleidum announces $\frac{1}{3}$.
2. **2nd Time:** Farbeo announces $\frac{1}{2}$ and Kleidum announces $\frac{1}{3}$.
3. **3rd Time:** Farbeo announces $\frac{1}{2}$ and Kleidum announces $\frac{1}{3}$.
4. **4th Time:** Both announce $\frac{1}{2}$.

Question: How is this possible?

Farbeo's Perspective

Suspect	Hair Color	Clothing
Mr. Green	Black	Cotton
Colonel Mustard	White	Polyester
Mrs. Peacock	Black	Cotton
Professor Plum	Brown	Polyester
Ms. Scarlet	Brown	Cotton
Mrs. White	White	Polyester

- Farbeo knows the killer has black hair, and he knows every black-haired suspect wore cotton.
- So he will learn nothing from Kleidum's announcements.
- So he announces $\frac{1}{2}$ come what may.

Learning from Others

Suspect	Hair Color	Clothing
Mr. Green	Black	Cotton
Colonel Mustard	White	Polyester
Mrs. Peacock	Black	Cotton
Professor Plum	Brown	Polyester
Ms. Scarlet	Brown	Cotton
Mrs. White	White	Polyester

Kleidum considers both Scarlet (brown hair) and Peacock (black hair) are still suspects, as they both wore cotton.

Learning from Others

Suspect	Hair Color	Clothing
Mr. Green	Black	Cotton
Mrs. Peacock	Black	Cotton
Mrs. White	White	Polyester
Colonel Mustard	White	Polyester
Professor Plum	Brown	Polyester
Ms. Scarlet	Brown	Cotton

Kleidum knows that **if the killer's hair is brown**, then

Before 1st Announcement:

1. Farbeo is unable to infer anything about the killer's clothing.
2. Farbeo thinks it's possible that the killer wore polyester.
3. Farbeo will announce $\frac{1}{2}$, indicating to Kleidum that the killer has brown or black hair.

Learning from Others

Suspect	Hair Color	Clothing
Mr. Green	Black	Cotton
Mrs. Peacock	Black	Cotton
Mrs. White	White	Polyester
Colonel Mustard	White	Polyester
Professor Plum	Brown	Polyester
Ms. Scarlet	Brown	Cotton

After 1st Announcement:

1. Farbeo has announced $\frac{1}{2}$, indicating that the killer has brown or black hair.
2. Kleidum knows that **if the killer's hair is brown**, then
   - Farbeo knows that if the killer wore polyester, then Kleidum knows the murderer is Plum.
     - Farbeo announces $\frac{1}{2}$ if and only if the killer black or brown hair, and Plum is the only person with brown or black hair who wore polyester.
   - Farbeo knows that Kleidum's 2nd announcement would be 1.

Learning from Others

Suspect	Hair Color	Clothing
Mr. Green	Black	Cotton
Mrs. Peacock	Black	Cotton
Mrs. White	White	Polyester
Colonel Mustard	White	Polyester
Professor Plum	Brown	Polyester
Ms. Scarlet	Brown	Cotton

After 2nd Announcement: Kleidum knows that

1. Farbeo knows that Kleidum's second announcement was $\frac{1}{3}$ and
2. Farbeo knows that, if the killer had worn polyester, then her second announcement would have been 1.
3. So Farbeo knows that the killer did not wear polyester.

Learning from Others

Suspect	Hair Color	Clothing
Mr. Green	Black	Cotton
Mrs. Peacock	Black	Cotton
Mrs. White	White	Polyester
Colonel Mustard	White	Polyester
Professor Plum	Brown	Polyester
Ms. Scarlet	Brown	Cotton

After 2nd Announcement: Hence, Kleidum knows that **if the killer's hair is brown**, then

1. Then the killer wore cotton, not polyester.
2. Farbeo knows the killer is Scarlet, as she is the only person with brown hair wearing cotton.
3. Farbeo's third announcement will be zero, as Scarlet's name does not begin with "P."

Learning from Others

Moral: After the second announcement, Kleidum knows that that **if killer's hair is brown**, then Farbeo's third announcement will be zero.

Learning from Others

After 3nd Announcement: Kleidum knows that

1. **If killer's hair were brown**, then Farbeo's third announcement would have been zero.
2. Farbeo's third announcement was $\frac{1}{2}$, not zero.
3. So the killer's hair is not brown.
4. So the killer's hair is black.
   - She knew it was black or brown before any announcements!
5. So the killer is Peacock or Green.

Post Announcements

Cotton

Polyester

After Conference: $P_K(\varphi) = P_F(\varphi) = \frac{1}{2}$.

An Exercise

Exercise: Determine what Farbeo's and Kleidum's beliefs about φ are before and after the press conference if they learn

- Hair = White and Clothing = Polyester
- Hair = Brown and Clothing = Polyester

After doing the exercise, you'll realize that Farbeo and Kleidum always agree after the press conference.

Black magic?

Question: If Kleidum knew what Farbeo would announce, how did she learn anything from the conference?

Answer: Through the repeated announcements, Kleidum

- Learned what **Farbeo knew** about what **she knew**.
- And then inferred what Farbeo knew.

Common Knowledge

A proposition φ is called common knowledge for Farbeo and Kleidum if

- F knows that φ and K knows that φ,
- F knows that K knows that φ,
- K knows that F knows that φ,
- F knows that K knows that F knows that φ,
- K knows that F knows that K knows that φ,
- And so on.

Knowledge of Priors

Before the conference

- F knows that $P_K(\varphi) = \frac{1}{3}$,
- K knows that $P_F(\varphi) = \frac{1}{2}$

BUT ...

Generating Common Knowledge

Before the conference:

- K does not know that F knows that K knows that $P_F(\varphi) = \frac{1}{2}$.
  - Check this yourself!
- So their posteriors are **not** common knowledge.

Aumann's Agreeing to Disagree

Theorem

[Aumann, 1976] If two Bayesian agents have the same prior probabilities and their posterior probabilities are common knowledge, then their posteriors are equal.

Is Common Knowledge Necessary?

- In practice, the infinite hierarchy in the definition of common knowledge is not necessary.
- The agents only needs to be able to reason about a few "levels" of "higher-order" knowledge.
- Repeated announcements are one way of generating this "higher-level" knowledge.
- This is how the example works.

Aumann's Theorem

Aumann's theorem challenges the common-sense dictum that individuals with different evidence can disagree:

> *The point about the people being peers and sharing their evidence is crucial. No doubt people with different bodies of evidence can reasonably disagree . . . [T]he puzzling case is the one in which each person knows about the other's reasons.*

Feldman [2011], pp. 144.

Question-Begging?

Question: But isn't assuming that individuals share a prior probability question-begging for philosophers who are interested in peer disagreement?

Outline

The DeGroot/Lehrer Model

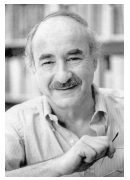

Morris DeGroot

Keith Lehrer

The Model:

- There is a proposition about which several individuals disagree.
- Each individual i initially assigns some probability p_i to the proposition.

The DeGroot/Lehrer Model

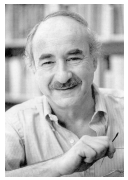

Morris DeGroot

Keith Lehrer

The Model:

- Each individual i assigns every individual j (including himself!) a non-zero weight $w_{i,j}$.
  - The weights represent how reliable i believes j is **relative** to others in the group.
- $0 \leq w_{i,j} \leq 1$.
- For any individual, the weights sum to 1, i.e., $\sum_j w_{i,j} = 1$.

The DeGroot/Lehrer Model

Example: Suppose I am in a meeting with Hannes and Stephan, and we disagree.

- I think Hannes and Stephan are about equally likely to be correct, and both are far more reliable than me.
- Accordingly, I set the following weights:
  - Hannes gets weight .45
  - Stephan gets weight .45
  - I get weight .1.

The DeGroot/Lehrer Model

The Model: In DeGroot and Lehrer's model, the weights dictate how individuals update their beliefs.

The DeGroot/Lehrer Model

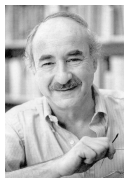

Morris DeGroot

Keith Lehrer

The Model:

- Time is divided into discrete stages.
- Let i's degree of belief on stage t be represented by $p_{i,t}$
- On stage $t+1$, individual i updates his belief to be a weighted-average of everyone's beliefs from stage t.

$$p_{i,t+1} = \sum_j w_{i,j} \cdot p_{j,t}$$

The DeGroot/Lehrer Model

Example: Suppose I am in a meeting with Hannes and Stephan, and we disagree about how likely it is that a job candidate will get hired.

- I assign the following weights:
  - Hannes gets weight .45
  - Stephan gets weight .45
  - I get weight .1.
- Initially, our beliefs are as follows:
  - $p_{Hannes} = .8$
  - $p_{Stephan} = .4$
  - $p_{Conor} = .5$

The DeGroot/Lehrer Model

After 1 stage, my belief is equal to

$$\begin{aligned} p^* &= .45 \cdot p_{Hannes} + .45 \cdot p_{Stephan} + .1 \cdot p_{Conor} \\ &= .45 \cdot .8 + .45 \cdot .4 + .1 \cdot .5 \\ &= .59 \end{aligned}$$

The DeGroot/Lehrer Model

Note that this process of taking a weighted-average is similar the models of "splitting the difference" that you have seen.

Two Additional Features:

- It allows individuals to treat others as reliable to different degrees.
- It works when their are multiple individuals who disagree.

The DeGroot/Lehrer Model

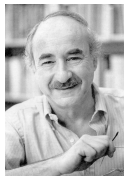

Morris DeGroot

Keith Lehrer

Theorem

[DeGroot, 1974, Lehrer and Wagner, 1981] In the above model, all individuals beliefs approach a common probability as the number of stages grows larger.

Quickness of Convergence

Question: How quickly do opinions of the individuals converge?

Answer: Very quickly. Let me show you a simulation.

The DeGroot/Lehrer Model

The model raises at least three questions:

1. Why should individuals assign non-zero weight to others?
2. Why should individuals repeat the averaging process?
3. Why should the weights remain **constant**?

Answers: I will just quote (but not defend) Lehrer's justifications.

Non-Zero Weights

> *First, the respect assumption, weakened as indicated, may be taken as a condition of a community of experts. If some members of a group respect each other, give positive weight to the probability assignments of each other, but give no weight to the probability assignments of others, then they form a separate and distinct community. Only when each member of a group communicates respect for each other member, either directly or through a chain, does a community of inquiry exist.*

Lehrer and Wagner [1981], page 330.

Repeated Averaging

> *[R]efusing to shift from state 1 to state 2 is equivalent to assigning a weight of 0 to other members of the group at this stage. This amounts to the assumption that there is no chance that one is mistaken and no chance that others in the group with whom one disagrees are correct. In short, the only alternative to the iterated aggregation converging toward a consensual probability assignment is individual dogmatism at some stage.*

Lehrer and Wagner [1981], page 331.

Constant Weights

> *The constancy condition is sustained by the assumption that members of the community . . . acquire no new information . . . The constancy assumption amounts to the requirement that a person who forms an estimate of the reliability of others as indicators of truth apply that estimate consistently until he obtains new information.*

Lehrer and Wagner [1981], page 330.

Preview

Next Week: Agent-based models of peer disagreement.

NetLogo

Today we will discuss:

- Procedures and Reporters
- Global vs. Local Variables

References I

Aumann, R. J. (1976). Agreeing to disagree. *The annals of statistics*, 4(6):12361239.

DeGroot, M. H. (1974). Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121.

Feldman, R. (2011). Reasonable religious disagreements. In Goldman, A. and Whitcomb, D., editors, *Social Epistemology: Essential Readings*, pages 137—158. Oxford University Press.

Lehrer, K. and Wagner, C. (1981). *Rational Consensus in Science and Society: A Philosophical and Mathematical study*, volume 24. D. Reidel.