Is Rational Disagreement Possible?
The DeGroot-Lehrer and Hegselman-Krause Models

Models and Simulations in Philosophy
May 12th, 2014

Review

Question: Can rational individuals with the same evidence disagree?

We've discussed two answers.

Review

Answer 1: Yes.

- It's rational to stick to one's guns. "Psychological" evidence need not outweigh non-psychological evidence [Kelly, 2011]
- Individuals with different evidence might disagree [Feldman, 2011].

Review

Answer 1: No.

- "Equal-weight view" = In the face of disagreement,
  - Suspend judgment [Feldman, 2011].
  - 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

Last Class: We began discussing formal models aimed at answering the same question:

Question: Can rational individuals with the same evidence disagree?

Theorem

[Aumann, 1976] If two individuals have a common prior and their posteriors are common knowledge, then their posteriors are equal.

⇒ Non-shared evidence need not undermine agreement.

Aumann's model is a paradigm of a classical economic model.

Characteristics of Aumann's Model:

- Both agents are ideal, rational Bayesian agents.
- There are only two individuals considered (though the model extends naturally to finite numbers).
- Information (about the posteriors) is completely shared.
  - Even stronger: it's common knowledge.
- The dynamics (e.g., via announcements) leading to consensus are not explored.

ABMs vs. Classical Economic Models

Classic Models	ABMs
Rational/Bayesian agents	Boundedly Rational
Homogeneous agents	Heterogenous Agents
Global Interaction	Local interactions in a network
Equilibria	Dynamics

And many more . . .

Today: Agent-based models (ABMs) for exploring the question of peer-disagreement.

All derive from the DeGroot-Lehrer model of repeated belief averaging.

Review

Characteristics of Models Today:

- Repeated averaging need not be (but can be) Bayesian.
- Agents may be heterogenous with respect to influence. how they assign each other credibility, and how they share information.
  - Averaging in a network [Golub and Jackson, 2010].
  - Averaging those with "close" beliefs [Douven, 2010, Hegselmann and Krause, 2002].
- The dynamics of consensus-formation are explored.

Today's Class

Question: Can rational individuals with the same evidence disagree?

- DeGroot [1974], Lehrer and Wagner [1981] - Disagreement is impossible among multiple individuals even when non-equal weight is given to different peers.
- Golub and Jackson [2010] - "Splitting the difference" can make crowds "wise."
- Douven [2010] - Whether "splitting the difference" is rational (or not) depends upon your goals and context.

Outline

The DeGroot/Lehrer Model

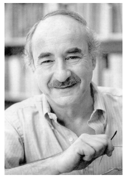

Morris DeGroot

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

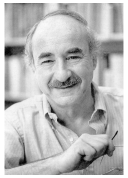

Morris DeGroot

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

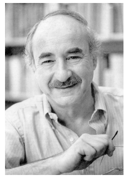

Morris DeGroot

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 particular conjecture about Bayesian networks is true.

- 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

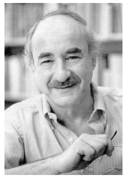

Morris DeGroot

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.

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**?

I will just quote Lehrer.

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 [1976], 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 [1976], 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 [1976], page 330.

The DeGroot/Lehrer Model

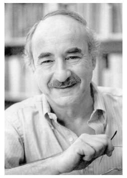

Morris DeGroot

Keith Lehrer

Theorem

[DeGroot, 1974, Lehrer, 1976] In the above model, all individuals beliefs approach a common probability as the number of stages grows larger.

Normally, I won't talk about how proofs proceed in this class.

In this case, the proof is instructive.

Markov Processes

How do students decide what to eat?

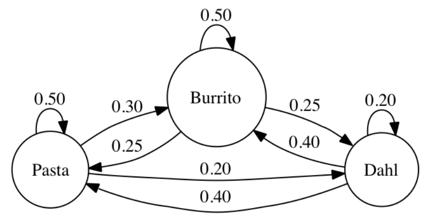

Markov Process = The current state of a system depends only upon its recent past.

Transition Matrices

Markov processes can be described by transition matrices:

	Pasta	Burrito	Dahl
Pasta	.5	.3	.2
Burrito	.25	.5	.25
Dahl	.4	.4	. 2

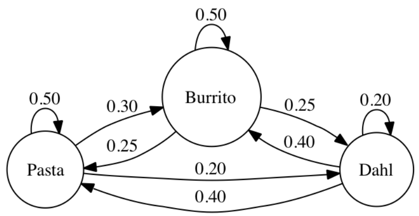

Transition Matrices

Markov processes can be described by transition matrices:

$$T = \begin{pmatrix} .5 & .3 & .2 \\ .25 & .5 & .25 \\ .5 & .25 & .25 \end{pmatrix}$$

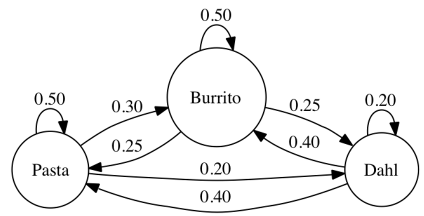

Transition Matrices

The transition matrix for two stages is obtained by squaring the original matrix:

$$T^2 = \begin{pmatrix} .5 & .275 & .225 \\ .5 & .275 & .225 \\ .5 & .275 & .225 \end{pmatrix}$$

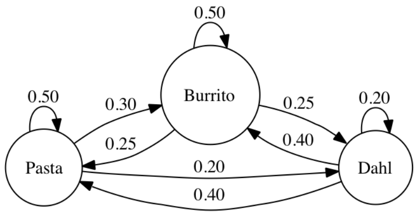

Stationary Limits

The transition matrix for n many stages is obtained by taking the n^{th} power of the original matrix:

$$T^n = \begin{pmatrix} .5 & .275 & .225 \\ .5 & .275 & .225 \\ .5 & .275 & .225 \end{pmatrix}$$

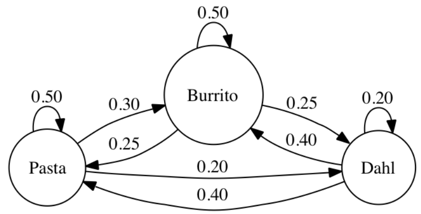

Stationary Limits

Curious: The transition matrix acquired a fixed value, and its rows are identical ...

Stationary Limits

Theorem

Under a wide variety of conditions, the T^n approaches a fixed, limiting matrix T^∞ with one row. The single row represents the probability of where the process will be in the limit, regardless of its starting point.

This is a mathematical fact about matrix multiplication.

So it doesn't matter what the numbers in the matrix represent ...

Weight Matrices

Consider the weight matrix that represents the weights individuals assign to one another.

	Agent 1	Agent 2	Agent 3
Agent 1	$w_{1,1}$	$w_{1,2}$	$w_{1,3}$
Agent 2	$w_{2,1}$	$w_{2,2}$	$w_{2,3}$
Agent 3	$w_{3,1}$	$w_{3,2}$	$w_{3,3}$

"Weight" Matrices

For instance, suppose agents assign each other the following weights:

$$W = \begin{pmatrix} .5 & .3 & .2 \\ .25 & .5 & .25 \\ .5 & .25 & .25 \end{pmatrix}$$

Updating and "Weight" Matrices

Let b_0 be a vector representing all agents' original beliefs.

Then after one stage their new beliefs are taken by multiplying the matrix by the vector.

$$b_1 = W \cdot b_0.$$

Consensus and "Weight" Matrices

$$W^n = \begin{pmatrix} .5 & .275 & .225 \\ .5 & .275 & .225 \\ .5 & .275 & .225 \end{pmatrix}$$

- In general, after n many stages:

$$b_n = W^n \cdot b_0.$$

- Because W^n approaches a limit, so does b_n.

Consensus

$$W^n = \begin{pmatrix} .5 & .275 & .225 \\ .5 & .275 & .225 \\ .5 & .275 & .225 \end{pmatrix}$$

- Because the rows of W^n are all the same, each element of the vector $b_n = W^n \cdot b_0$ is the same.
  - So consensus is reached.

Influence

$$W^n = \begin{pmatrix} .5 & .275 & .225 \\ .5 & .275 & .225 \\ .5 & .275 & .225 \end{pmatrix}$$

- The elements of the rows represent how large a role each agent's initial opinion plays in the consensus.
  - Above, agent one's opinion is weighted most heavily in the consensus.

Truth-Seeking

Questions:

- Reliability: Is the consensus closer to the truth than individuals' initial opinions?
  - Variants of this question are addressed by Golub and Jackson [2010] and Douven [2010].
- Speed: How quickly is consensus reached?
  - Very quickly. I will show you simulations in a moment.

Generalization to Lehrer-DeGroot Model

Both the Golub-Jackson and Hegselmann-Krause models generalize the Lehrer-DeGroot model in one important way.

Generalization to Lehrer/DeGroot Model

Common Generalization:

- Agent i's belief is represented by a real number r_i.
  - E.g. A subjective probability of a proposition
  - E.g. A numerical estimate of some quantity (e.g. charge of an electron).
- The truth is likewise represented by a real-number T.
  - E.g. 0 or 1 might represent the truth-value of some proposition.
  - E.g., The "real" value of some quantity (e.g. charge of an electron).

Outline

Golub and Jackson [2010, 2012] are interested in exploring how network structure influences the consensus reached and the reliability of learning.

High School Social Networks

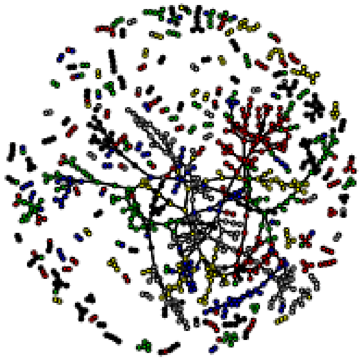

Goodreau et al. [2008]

Academic Social Networks

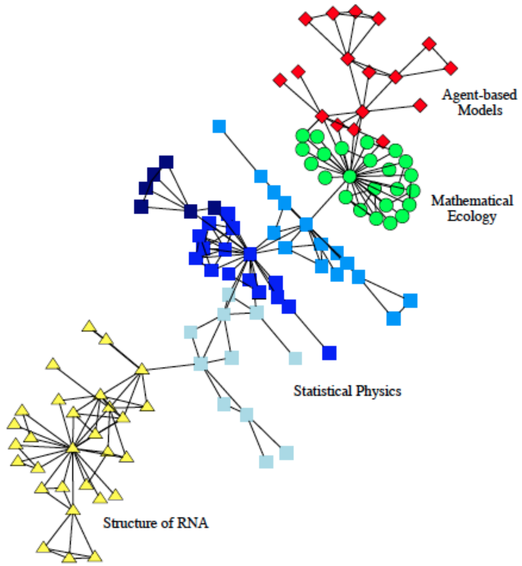

Structure of Real Scientific Networks

What types of structural properties do academic co-authorship networks and other social networks share?

Here are four.

Small Diameters

Finding the Diameter

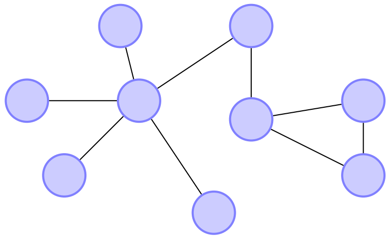

Finding the Diameter

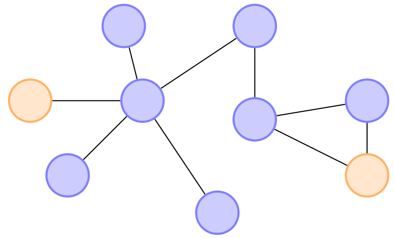

Diameter

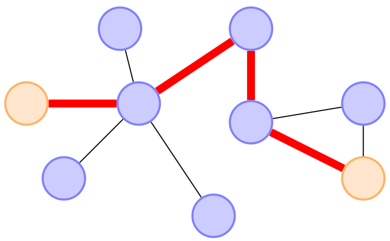

Diameter: The longest-shortest path between any two nodes in the network.

Empirical Size of Connected Components

	biology	physics	mathematics
number of authors	1,520,251	52,909	253,339
diameter	24	20	27

Newman [2001]

Highly Clustered

Clustering Coefficient

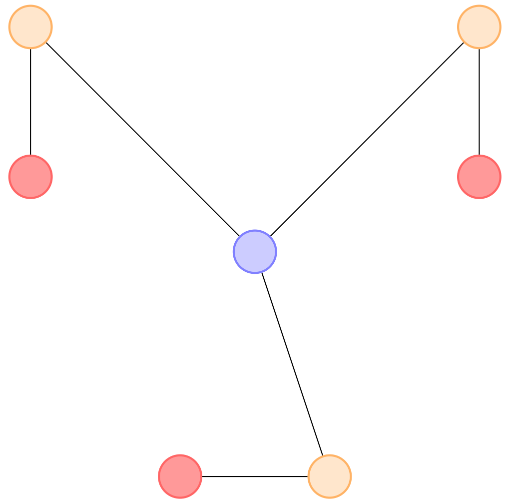

Not clustered

Clustering Coefficient

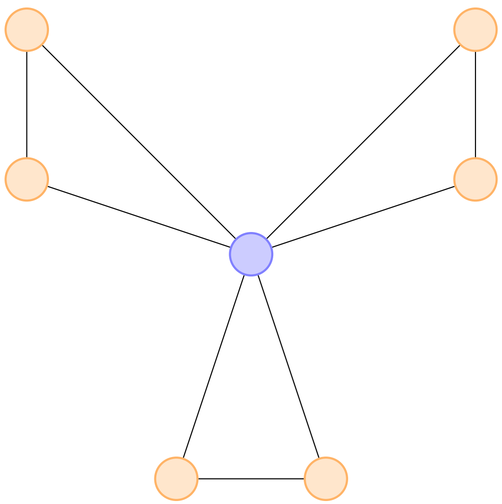

More Clustered

Clustering Coefficient

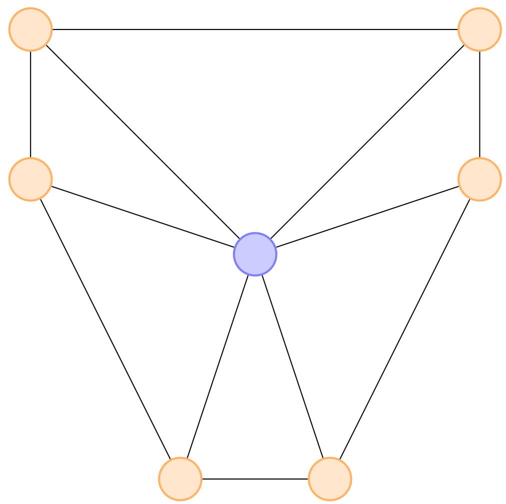

Highly Clustered

High Homophily

Homophily

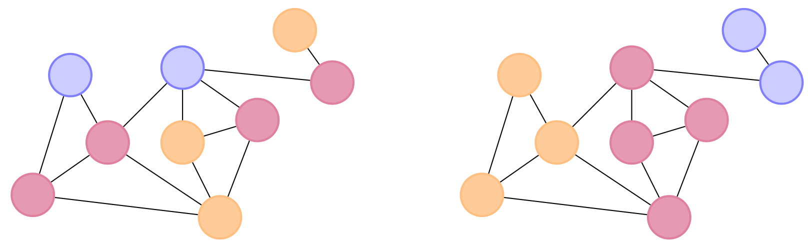

Left: A Non-Homophilous Network
Right: A Homophilous Network

High School Social Networks

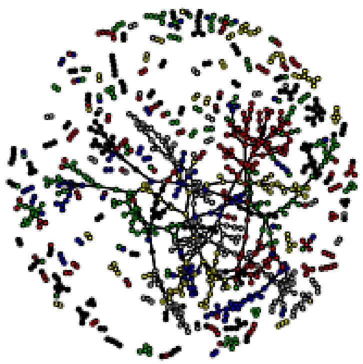

Goodreau et al. [2008]

Academic Social Networks

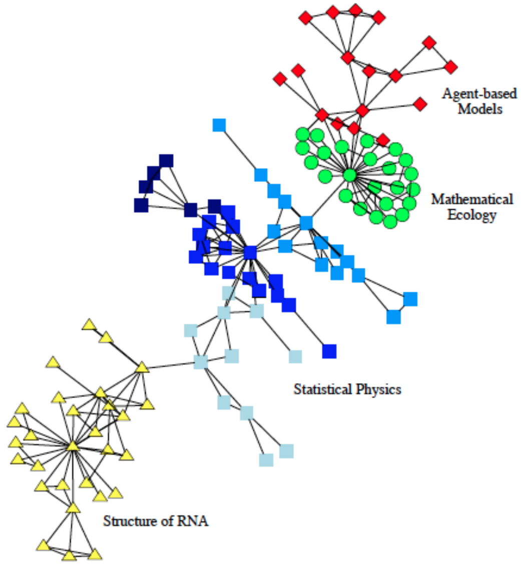

Power Law Degree Distribution

Neighborhoods

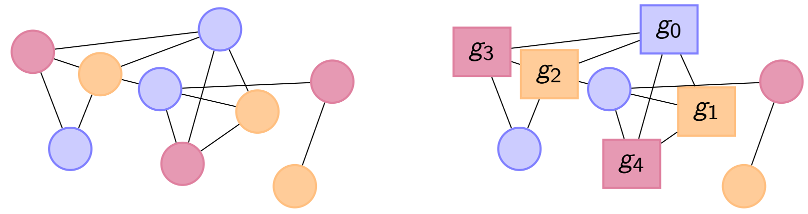

g_0's neighborhood

Degree

The degree of an agent is the number of her neighbors.

Power Law Degree Distribution

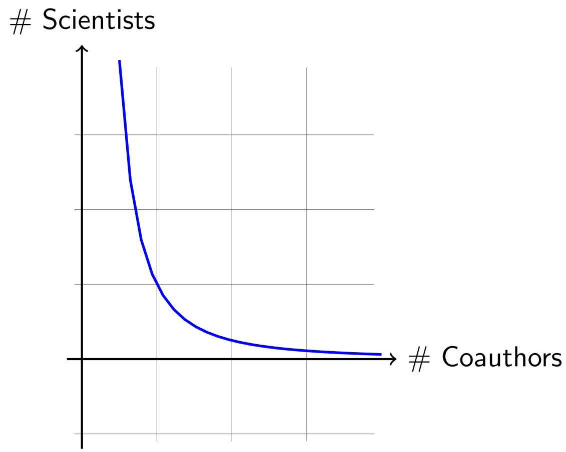

Question: What is the relationship between networks and "weight matrices" above?

Answer: Each undirected (or directed) network can be associated with a unique weight matrix, under the assumption each neighbor is treated equally.

Networks

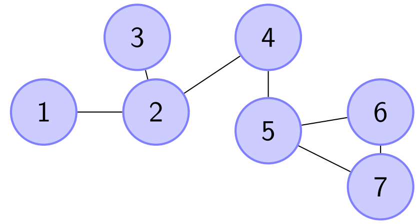

Agent	1	2	3	4	5	6	7
1	$\frac{1}{2}$	$\frac{1}{2}$	0	0	0	0	0
2	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	0	0	0
3	0	$\frac{1}{2}$	$\frac{1}{2}$	0	0	0	0
4	0	$\frac{1}{3}$	0	$\frac{1}{3}$	$\frac{1}{3}$	0	0
5	0	0	0	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$
6	0	0	0	0	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
7	0	0	0	0	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$

Networks as Matrices

Moral: The "weight matrices" are a generalization of networks.

- Matrices more easily represent networks in which influence is
  - Asymmetric
  - Weighted
- So Golub and Jackson talk about properties of the matrix instead of the network.

Random Networks and Convergence

Question: How quickly do random networks reach consensus?

Answer: Quickly. Let me show you.

Influence

$$W^n = \begin{pmatrix} .5 & .275 & .225 \\ .5 & .275 & .225 \\ .5 & .275 & .225 \end{pmatrix}$$

$$b_n = W^n \cdot b_0$$

- **Recall:** The rows of the limiting matrix $W^\infty = \lim_{n\to\infty} W^n$ are identical.
- Elements of the row indicate how much influence each agent has in the consensus.

Influence and Network Structure

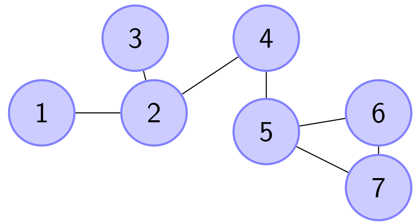

Fact: The influence of an agent is directly proportional to her degree (i.e. number of neighbors).

But there are other ways that agents might **indirectly** acquire influence.

Definition: Say a network is wise if the consensus it reaches is the true value of the unknown quantity.

- E.g., When repeated belief-averaging leads them to conjecture the true value of the charge of an electron.

Question: Under what conditions is the network wise?

Theorem

[Golub and Jackson, 2010] Under one (major) assumption, networks are "wise" precisely when no group has too much influence, and every group takes each other seriously.

Here is their idea in a nutshell:

The Argument

- Main Assumption: Each agent's initial belief is drawn randomly from some probability distribution with a mean equal to the truth T, and some finite variance.
- As the network gets bigger, it follows from the large of law numbers that the average belief of the group is T.
- If agents have equal influence, then the consensus reached will be the average belief of the group.
- So if agents have equal influence, then the consensus will be the true value.

The Main Assumption

Question: How plausible is the main assumption?

The Main Assumption

- The main assumption is essentially impossible to satisfy if the quantity is the truth value of a proposition.
  - Given $r_1, r_2, \ldots r_n < 1$, the average of the r_i's is less than one.
- It is strictly stronger than the assumption that individuals' estimates of a quantity are **unbiased**.
  - E.g., Americans systematically overestimate that proportion of the budget dedicated to foreign aid. Only 4% correctly identify that it is less than 1%.
- Nonetheless, there are cases in which it seems to be reasonable (e.g., Galton's experiment).

Outline

Golub and Jackson [2010] assume that the network is static.

Bounded Confidence Models

- Hegselmann and Krause [2002] propose a model in which one's peers change over time.
- However, in the model, an individual considers only those whose opinions are sufficiently **similar** to his or her own.
  - **Justification:** Anecdotal evidence suggests that we use similarity with our own opinions to gauge the reliability of others.
  - E.g., Political affiliation causes choice of newspaper.
  - Since our opinions change, our peer groups can change.
- Hegselmann and Krause wanted to explain why groups might become polarized and not reach a consensus.

Bounded Confidence Models

Hegselmann and Krause [2002]'s model is the motivation for [Douven, 2010]'s model.

So let me briefly review Hegselmann and Krause [2002]'s model

The Hegselmann-Krause Model

- Agent i's belief is represented by a real number r_i.
  - E.g. A subjective probability of a proposition
  - E.g. A numerical estimate of some quantity (e.g. charge of an electron).
- The truth is likewise represented by a real-number T.
  - E.g. 0 or 1 might represent the truth-value of some proposition.
  - E.g., The "real" value of some quantity (e.g. charge of an electron).

The Hegselmann-Krause Model

- There is some number ρ (for all agents) that represents how "close" others opinions must be to one's own in order for one to take them seriously.
  - If ρ is close to zero, then one only considers the opinions of those who are similar to oneself.
- Agent i is assigned some number τ_i between 0 and 1 that represents how strongly she is "attracted" to the truth.
  - $\tau_i = 0$ will represent an agent who only listens to her peers.
  - $\tau_i = 1$ will represent an agent who has immediate access to the truth.

The Hegselmann-Krause Model

- Time is divided into discrete stages $1, 2, 3 \dots$
- On stage $t + 1$, agent i averages
  - The beliefs of her peers whose opinions are within distance ρ of her own.
  - The truth T.
- The truth is given weight τ_i.
- The remaining weight $1 - \tau_i$ is divided **evenly** among peers.

Bounded Confidence Models

Formally:

- Let $b_{i,t}$ be agent i's belief at time t.
- Let $N_\rho(i,t)$ be those peers whose opinions differ from i's by no more than ρ at stage t.
- Let N be the number of peers in $N_\rho(i,t)$. Then:

$$b_{i,t+1} = \tau_i \cdot T + \sum_{j \in N_\rho(i,t)} \frac{1}{N} \cdot b_{j,t}.$$

How to Model "Splitting the Difference"

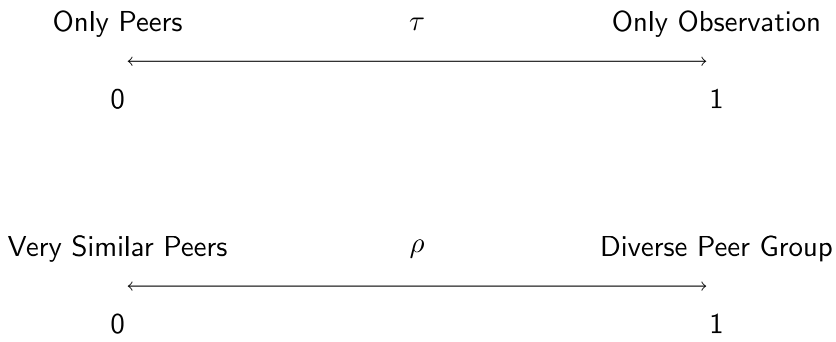

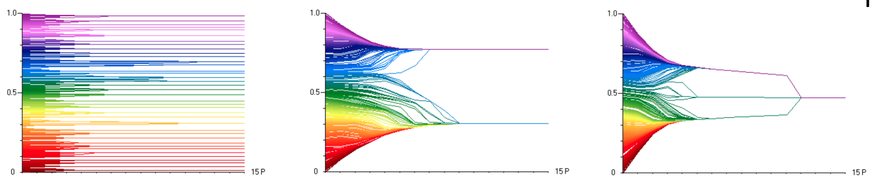

$\rho = .01$ $\rho = .15$ $\rho = .25$

Doven's Model

How does Douven [2010] use the Hegselmann-Krause model?

- Compares two types of communities:
  - **Feldman communities**: Individuals assign high weights to their peers and less to truth.
  - **Kelly communities**: Individuals assign high weight to truth and less to peers.
- He evaluates these communities in two respects:
  - Accuracy: How close their beliefs get to the truth.
  - Speed: How fast they get there.

Douven's Model

You tell me! What does Douven find?

- Accuracy: Feldman $>$ Kelly
- Speed: Kelly $>$ Feldman

Douven's Model

Explanation: These results can be explained in one or two sentences if you know the central limit theorem (or any probabilistic inequality like Chebyshev's), but I'll leave that for you to work out.

A Tradeoff

General Trend: In many future models, we'll see a tradeoff between speed and accuracy. Watch out for it.

Preview

Next Week: We'll start studying a different, but related question: is diversity good for science?

NetLogo

Today we will discuss:

- Agents Commands:
  - Turtles
  - Links
  - Patches
- Agent sets

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.

Douven, I. (2010). Simulating peer disagreements. *Studies in History and Philosophy of Science Part A*, 41(2):148–157.

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

Golub, B. and Jackson, M. O. (2010). Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, pages 112–149.

Golub, B. and Jackson, M. O. (2012). How homophily affects the speed of learning and best-response dynamics. *Forthcoming in Annals of Economics and Statistics.*

References II

Goodreau, S. M., Handcock, M. S., Hunter, D. R., Butts, C. T., and Morris, M. (2008). A statnet tutorial. *Journal of statistical software*, 24(9):1.

Hegselmann, R. and Krause, U. (2002). Opinion dynamics and bounded confidence models, analysis and simulation. *Journal of Artifical Societies and Social Simulation*, 5(3).

Kelly, T. (2011). Peer disagreement and higher order evidence. In Goldman, A. and Whitcomb, D., editors, *Social epistemology: Essential readings*, pages 183–217. Oxford University Press.

Lehrer, K. (1976). When rational disagreement is impossible. *Nous*, 10(3):327–332.

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

Newman, M. E. (2001). The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences*, 98(2):404—409.