Communication and Diversity:
Bandit Problems and Network Models of Science

Models and Simulations in Philosophy
June 30th, 2014

Review

Last Month: Diversity in science

- What is diversity?
- Why is it good?
- How might it be modeled?

Kitcher, Longino, and Wylie



Different benefits of diversity are extolled by these authors.

- Kitcher: When at least two theories are in competition, diversity may raise community's collective chance of success.
- Longino and Wylie: Diversity of approaches might yield results of different value.

Kitcher, Longino, and Wylie



Two types of models:

- Kitcher: "Classical" game-theoretic model
- Weisberg, Muldoon, Hong, and Page: Agent-based models

Today

Today: Zollman [2010]'s ABM of communication in science.

Review

Zollman's Conclusions:

- More communication is not necessarily better
  - Quick access to information can lead good theories to be abandoned too quickly.
- Dogmatism is not necessarily bad
  - Dogmatic scientists make sure theories get fair hearings.
- Dogmatism and lack of communication together are likely bad
  - The prevent a community from pursuing alternative research methodologies.

Transient Diversity

In general, transient diversity is the goal:

- Scientists ought to explore different research methodologies until the best of a group is found, and then explore no longer.
- Dogmatic scientists (with different beliefs) and lack of communication are just two ways of promoting transient diversity

Outline

Nobel Prize

History: In 2005, Barry Marshall and Robin Warren were awarded the Nobel Prize for the discovery of the bacteria that causes ulcers.

Their landmark paper first appeared around 1983.

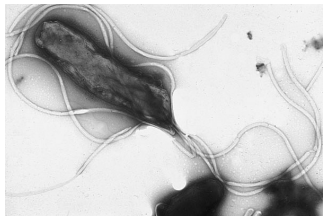

A Brief History of Peptic Ulcers

Here's the funny thing:

- The hypothesis that bacteria causes ulcers has been around since the 19th century.
- Some doctors have successfully treated ulcers with antibiotics since the 1950s.

A Brief History of Peptic Ulcers

Question: Why did Marshall and Warren get so much credit? Why was their finding so important?

Answer: The accepted theory from 1954 to 1985 or so was that

- Bacteria cannot live in the stomach.
- Ergo, bacteria cannot cause ulcers.

A Brief History of Peptic Ulcers

Question: How did this theory become dominant?

The Rise of a Hypothesis

In the middle of the 20th centurty, there were several competing hypotheses about the cause of ulcers:

- Acid
- Bacteria
- Stress (post-Palmer's study)
- And variation on these.

Non-Dogmatism and Quick Dissemination

According to Zollman, two features characterized the medical community in 1954:

1. Sufficient open-mindedness (non-dogmatism) about the causes of ulcers
2. Quick dissemination of research results

Rejection of the Bacterial Hypothesis

How did the bacterial hypothesis become unpopular?

- In 1954, Palmer tests 1000 patients stomachs for bacteria and finds nothing!
- Little did he know his method did not detect the type of bacteria that produce ulcers.

Rejection of the Bacterial Hypothesis

After Palmer: The bacterial hypothesis was nearly universally rejected.

- Lykoudis, a Greek physician, is shunned and fined for treating ulcers with antiobiotics.
- Warren and Marshall cannot get their initial paper accepted into a conference with a 90% acceptance rate!

Transient Diversity

Zollman claims that if the medical community had either

1. Consisted of a few dogmatic defenders of each hypothesis in the middle of the 20th century, or
2. Disseminated Palmer's findings less quickly

Then: Some scientists would have continued to pursue the bacterial hypothesis, and we might have discovered the bacterial cause of peptic ulcers earlier.

Transient Diversity

Can we produce a model that explains the history of ulcers and also bolsters Zollman's (counterfactual) predictions?

Outline

Bandit Problems

Goal: Find the arm (or machine) with the highest payoff.

Tradeoff: To ensure that you find the best machine, you must experiment with inferior ones.

Applications

Typical applications of Bandit Problems:

- Medical Treatment [Berry and Fristedt, 1985]
- Crop choices in Africa [Goyal, 2003]
- Drilling sites [Keller et al., 2005]

Bandit Problems in Philosophy of Science

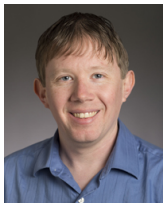

Zollman [2010] claims that bandit problems can be used to represent methodological choices in the sciences.

Here are some examples from Mayo-Wilson et al. [2011].

Bandit Problems in Philosophy of Science

To study animal behavior: Biologists might use

- Population genetic models,
- Game theoretic models,
- Etc.

These techniques corresponds to different "arms" of a slot machine.

Bandit Problems in Philosophy of Science

To explain some human behavior: A psychologist might use any number of theories of concepts

- Exemplar-based
- Prototype-based
- Causal-model theory
- Theory-theory

Again, these modeling techniques or theories corresponds to different "arms" of a slot machine.

Bandit Problems in Philosophy of Science

Understanding the metaphor: Stochasticity

- Models, theories, and techniques are not always successful at explaining observed phenomena. Why?
  - Collected data may be atypical, noisy, and/or subject to error.
  - Experimenters are not perfect in applying models,
  - It may be unclear/indeterminate how to apply the model.
- So models have probabilistic returns, just like slot machines.

Bandit Problems in Philosophy of Science

Understanding the metaphor: Exploration vs. Exploitation

- Nonetheless, some models are generally more successful than others.
- Scientists do not want to spend their careers pursuing inferior models.
- But they need to explore to figure out which are successful.

Zollman's Model

Zollman [2010] makes the following simplifications to the general bandit problem:

- Arms only give payoffs 0 or 1.
- So agents want to find the arm with the greatest probability of obtaining a payoff of 1.

Zollman's Model

This allows Zollman [2010] to model agents as simple Bayesian learners whose beliefs are represented by beta distributions:

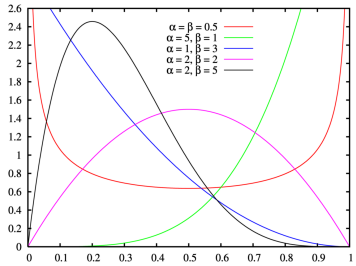

The curve $p(x)$ represent how probable the agent believes it to be that the arm pays off x% of the time.

Zollman's Model

How Beta distributions work:

- For each arm, there are numbers α and β such that the agent initially believes the probability that the arm will give a payoff is $\frac{\alpha}{\alpha+\beta}$.
- After n observations, of which s are 1, the agent then believes the probability of success is:

$$\frac{\alpha + s}{\alpha + \beta + n}.$$

- So the bigger α and β are, the more dogmatic the agent: her beliefs will change more slowly.

Zollman's Model

Zollman [2010] assumes agents are myopic: they always pick the arm which they believe to have the highest expected payoff.

Networks

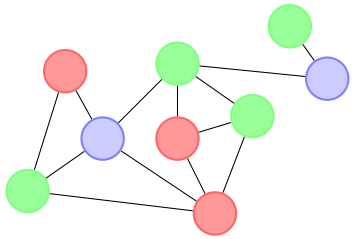

Nodes = Agents
Edges = Indicate which agents can view each other's data.

Bandit Problems in Networks

Again, applying bandit problems in a social setting is not new [Goyal, 2003].

Outline

Criteria for success

Criteria for Success: All agents converge to playing the arm with highest expected payoff.

Zollman's Conclusions

Zollman's Conclusions:

- More communication is not necessarily better
  - Networks with more edges converge on playing the best arm less often
- Dogmatism is not necessarily bad
  - When agents have bigger αs and βs in densely connected networks, they converge more often.
- Dogmatism and lack of communication together are likely bad
  - Sparsely connected networks with agents with big αs and βs don't converge to playing the true arm.

Outline

Individual and Group Rationality

Zollman, like others, seems to find a conflict between individual and group rationality:

- Rational individual agents will employ the seemingly "best" model or theory, but
- If everyone employs the **seemingly** best model, then the **actually** best model or theory may be prematurely abandoned. So the group may suffer.

Two Standards of Rationality

At the outset of class, you read about two types of criteria for rationality:

- Bayesian - Maximize expected utility relative to posterior probability distribution.
- Reliabilist: Converge to true belief or optimal behavior in all possible worlds.

Two Standards of Rationality

Zollman's model employs both standards, but at different levels:

- Individual agents act in a Bayesian way.
- The group is judged by its reliable convergence to the "best" action.

Moral: Hence, the apparent conflict between individual and group rationality might be attributable to two different standards of rationality.

Similar remarks might apply to Kitcher's model.

Independence Thesis

Question: When group and individual rationality are both judged from roughly the same standpoint, can conflicts arise?

Answer: A qualified, "Yes." [Mayo-Wilson et al., 2011].

Under **some** ways of making this question precise using reliabilist standards of rationality, group and individual rationality can diverge.

Bandit Problems in Philosophy of Science

Problem: In bandit problems, payoffs are i.i.d, that is

- For each payoff r and arm a, there is some fixed probability $p(a, r)$ that the payoff will be r each time a is pulled.
- This seems implausible when applied to modeling techniques . . .

Bandit Problems in Philosophy of Science

Question: Why might payoffs of various modeling techniques may change over time?

- At first, the utility of a new model may increase:
  - Scientists become more familiar with and capable of applying the model.
  - Computational and technological improvements may help.
  - The model may offer novel explanations in stagnant areas of science.
- Typically, models then have diminishing returns:
  - The easy applications of a technique are discovered and exhausted.

Model Improvements

Various ways of improving the models:

- Drop the i.i.d assumption (e.g., impose decreasing marginal returns)
- More "realistic" networks
- Dynamic networks

References I

Berry, D. A. and Fristedt, B. (1985). *Bandit Problems: Sequential Allocation of Experiments*. Chapman & Hall.

Goyal, S. (2003). *Learning in Networks: a Survey*. University of Essex.

Keller, G., Rady, S., and Cripps, M. (2005). Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68.

Mayo-Wilson, C., Zollman, K. J., and Danks, D. (2011). The independence thesis: When individual and social epistemology diverge. *Philosophy of Science*, 78(4):653—677.

Zollman, K. J. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 72(1):17—35.