THE SOCIAL STRUCTURE OF SCIENCE

Conor Mayo-Wilson

Models and Simulations in Philosophy
June 11th, 2013

Course Outline:

- **First Half**: Models and simulations in ethics and political philosophy
  - Central Question: How can norm following emerge if such behavior is irrational or only one of many types of rational behavior?

Course Outline:

- **First Half**: Models and simulations in ethics and political philosophy
  - Central Question: How can norm following emerge if such behavior is irrational or only one of many types of rational behavior?
- **Second Half**: Models and simulations in philosophy of science and epistemology
  - No single overarching question. Questions are loosely about the "social structure of science."

What in the world is the "social structure of science"?

What in the world is the "social structure of science"?

Good question. But let's go back a bit . . .

Outline

1. Course Outline

Outline

1. Course Outline
2. Individual and Group Rationality
   - Conflict in the Social Sciences
   - Statistical Game Theory?

Outline

1. Course Outline
2. Individual and Group Rationality
   - Conflict in the Social Sciences
   - Statistical Game Theory?
3. Kitcher's Model

Outline

1. Course Outline
2. Individual and Group Rationality
   - Conflict in the Social Sciences
   - Statistical Game Theory?
3. Kitcher's Model
4. NetLogo

Outline

Individual vs. Collective Rationality

For some time, philosophers and social scientists have recognized a tension between individual rationality and group welfare.

For example, consider

- Prisoner's Dilemmas and Free Rider Problems
- Social Choice Theory

For some time, philosophers and social scientists have recognized a tension between individual rationality and group welfare.

For example, consider

- Prisoner's Dilemmas and Free Rider Problems
- Social Choice Theory

These models are supposed to show that rational individuals may not maximize group welfare (or their own, for that matter).

Conversely, the models we've studied show that boundedly rational agents may, despite their limitations, maximize group welfare.

Example: Cooperation (in prisoners' dilemmas) emerges in networks in which agents employ the boundedly-rational strategy "Imitate the Best."

Scientific Rationality

Might scientific inquiry be like other social interactions?

- Scientists, and humans in general, conduct experiments and learn in groups.
- Is it possible that the most reliable individual methods for designing experiments and making inferences from data might cause a group to learn slowly or unreliably?
- Is it possible that seemingly naive experimental design methods and statistical procedures are optimal when employed by groups of scientists?

To my knowledge, Kitcher [1990] was the first to ask these questions explicitly.

Kitcher: Individual vs. Collective Rationality in Science

> *Is it possible that there should be a mismatch between the demands of individual rationality and those of collective (or communal) rationality? Could it turn out that high-minded inquirers, following principles of individual rationality, should do a poor job of promoting the epistemic projects of the community they constitute? Might those with baser motives actually do more to advance their community's epistemic endeavors?*

Kitcher [1990], pp. 6.

- Aren't these questions obvious? Why were they first asked explicitly in 1990?
- Should the answers to these questions be any different from those in economics?

Traditionally, rational choice theory has two types of models:

- Decision Theory
- Game Theory

Traditionally, rational choice theory has two types of models:

- Decision Theory
  - One decision-maker.
  - Payoff depends upon action and unknown state of the world.

Traditionally, rational choice theory has two types of models:

- Decision Theory
  - One decision-maker.
  - Payoff depends upon action and unknown state of the world.
- Game Theory
  - Many decision-makers.
  - Payoff depends upon decision-makers' actions.

Statistical Decision Theory

Statistics has only one type of model:

- Statistical Decision Theory
  - One statistician or scientist (i.e., the decision-maker) who receives data over time.
  - Payoff depends predictive method (i.e. action) and the accuracy of prediction or hypotheses
    - Accuracy of prediction and/or hypothesis is quantified in many different ways: by variance, bias, asymptotic consistency, etc.
  - Accuracy is a function of an unknown state of the world.

Statistical Decision Theory

Statistics has only one type of model:

- Statistical Decision Theory
  - One statistician or scientist (i.e., the decision-maker) who receives data over time.
  - Payoff depends predictive method (i.e. action) and the accuracy of prediction or hypotheses
    - Accuracy of prediction and/or hypothesis is quantified in many different ways: by variance, bias, asymptotic consistency, etc.
  - Accuracy is a function of an unknown state of the world.
- Traditionally, there is no statistical game theory.

Scientific Games?

Perhaps this is with good reason.

- Scientists actions consist of experimenting, hypothesizing, predicting, etc.

Perhaps this is with good reason.

- Scientists actions consist of experimenting, hypothesizing, predicting, etc.
- The success or failure of an experiment, hypothesis, or prediction does not depend upon other scientists' actions, but rather, upon whether the experiment is well-designed, whether the hypothesis is true, and how accurate the prediction is.

Perhaps this is with good reason.

- Scientists actions consist of experimenting, hypothesizing, predicting, etc.
- The success or failure of an experiment, hypothesis, or prediction does not depend upon other scientists' actions, but rather, upon whether the experiment is well-designed, whether the hypothesis is true, and how accurate the prediction is.
- So perhaps there is no analog of game theory in science, where scientists rewards depend upon others actions.

Scientific Games?

A Slogan: Science is a game against Nature, not a game among scientists.

Scientific Games?

A number of philosophers, including Kitcher, have recently argued this slogan is highly misleading.

Competition in Science

First, scientific inquiry, as it is currently organized, is competitive:

- Priority Rule - First discoverer gets the credit [Merton, 1957].

First, scientific inquiry, as it is currently organized, is competitive:

- Priority Rule - First discoverer gets the credit [Merton, 1957].
- Matthew Effect - In collaborations, the most famous contributors get the most credit [Merton, 1968].

Second, even in cases in which inquiry is cooperative, scientists' gains and losses depend upon others' actions:

- Scientists depend upon others' data.
- Scientists depend upon their collaborators' skills and efforts.
- Scientists depend upon others to check and replicate their results.
- And so on.

Statistics for Scientific Communities?

- Just as statistics is a set of tools for individual scientists,
- Perhaps there is a set of tools explaining how scientific communities ought to be organized, and in particular, whether individual and collective rationality conflict.

> *How do we best design social institutions for the advancement of learning? The philosophers have ignored the social structure of science. The point, however, is to change it.*

Kitcher [1990], pp. 22.

Philosophers of science have been interested in evaluating and (sometimes) justifying the following features of scientific practice.

1. Reward Structure: How ought scientists be rewarded for discoveries?
    - Is the priority rule a good way of assigning credit to scientists?
        - [Kitcher, 1990] and [Strevens, 2003]: Yes.
    - Should we try to mitigate the Matthew effect?
        - [Strevens, 2006]: Not necessarily.
    - But there are tons of interesting unexplored questions:
        - In general, how should credit be distributed among collaborators?
        - How ought unreliable research be discouraged?

2. Communication: How and when should scientists share data, intermediate results, findings, etc?
    - Is more communication among scientists necessarily good?
        - Zollman [2011]: No.
    - Is publication in highly-specialized journals optimal for communication of scientific results?
        - Mayo-Wilson [2012]: May depend upon whether your goal is quick discovery or accurate dissemination of research findings.

3. Diversity: Is it good for a community to maintain diverse research methodologies and beliefs, even if it's generally agreed that one methodology is inferior?
    - Yes: Kitcher [1990], Strevens [2003], Weisberg and Muldoon [2009], Zollman [2010], Hong and Page [2004].

4. Collaboration: How ought collaboration among scientists be encouraged (if at all)? What are the benefits and costs of collaboration?
   - Ought scientists collaborate as frequently as they do?
     - Weisberg and Muldoon [2009]: Perhaps not.
     - Fittingly, Muldoon has single-authored paper in the works that argues his above earlier co- authored paper has limited conclusions.

4. Collaboration: How ought collaboration among scientists be encouraged (if at all)? What are the benefits and costs of collaboration?
   - Ought scientists collaborate as frequently as they do?
     - Weisberg and Muldoon [2009]: Perhaps not.
     - Fittingly, Muldoon has single-authored paper in the works that argues his above earlier co- authored paper has limited conclusions.
   - Anderson [2012] argues that if scientists collaborate with researchers with complementary skills, then under particular assumptions, the resulting collaboration networks resemble co-authorship networks.
     - Upshot: Current collaborative practices (as evidence by coauthorship) might be rationally justified by their ability to solve problems.

Few features of scientific practice have been investigated rigorously.

Even fewer have been formally modeled. There are a lot of projects here.

Importantly, most of the above models only allow one to draw how possible stories

Organizing Science

- Reward Structure: How is it possible that the priority rule and Matthew effect encourage successful scientific research?
- Communication: How is it possible that increased access to information could cause scientists' methods to become more unreliable?
- Diversity: How is it possible that diverse, but bad, problem solvers outperform homogeneous groups of very good problem-solvers?
- Etc.

Upshot: Philosophers are still relatively far from making good scientific policy with these models.

Case Studies

Case studies in diversity:

- Phlogiston
- Continental Drift
- Structure of DNA
- Zollman discusses a similar example about ulcers.

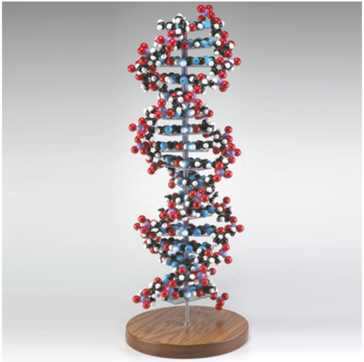

Kitcher's Model Simplified

Kitcher's Goal: To explain why scientists might rationally pursue seemingly inferior methods, as doing so may be good for the community as a whole.

For pedagogical reasons, I will present a simplified version of Kitcher's model.

Kitcher's Model

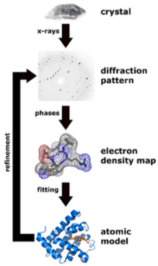

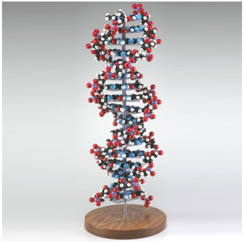

Kitcher's Model

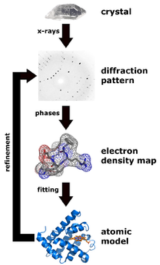

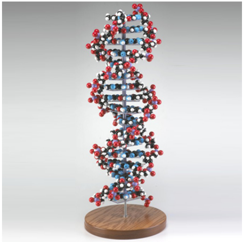

$p(Success) = .1$

$p(Success) = .01$

Kitcher's Model

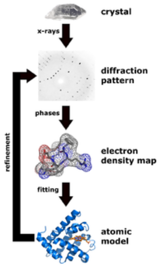

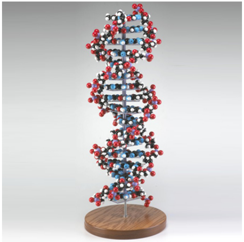

$p(Success) = .1$

$p(Success) = .01$

Kitcher's Model

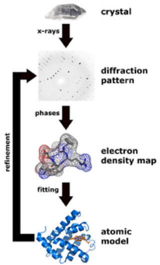

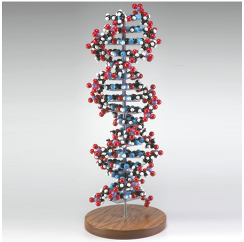

$p(Success) = .1$

$p(Success) = .01$


Kitcher's Model

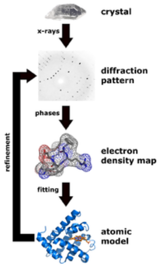

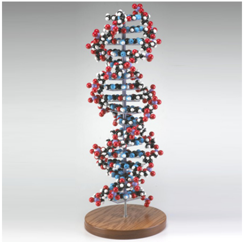

$p(Success) = .1$

$p(Success) = .01$

Kitcher's Model

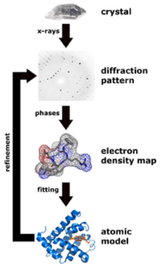

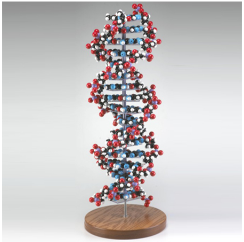

$p(Success) = .1$

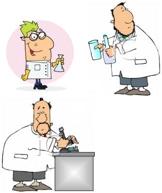

$p(Success) = .01$

Kitcher's Model Simplified

Assumptions of a simplified Kitcher-like model:

- At most one program will succeed.

Kitcher's Model Simplified

Assumptions of a simplified Kitcher-like model:

- At most one program will succeed.
- Scientists successively choose research programs and cannot change their decision once it has been made.

Kitcher's Model Simplified

Assumptions of a simplified Kitcher-like model:

- At most one program will succeed.
- Scientists successively choose research programs and cannot change their decision once it has been made.
- Winning scientists share the rewards.

Kitcher's Model Simplified

Assumptions of a simplified Kitcher-like model:

- At most one program will succeed.
- Scientists successively choose research programs and cannot change their decision once it has been made.
- Winning scientists share the rewards.
- Scientists maximize expected return, where return is the probability of success divided by the number of scientists working on the program.

- Kitcher's Result: Self-interested scientists will eventually choose the less promising program. Why?

- Kitcher's Result: Self-interested scientists will eventually choose the less promising program. Why?
- The utility of the more successful program decreases with each additional scientist: even one is more likely to succeed, one also has to share the rewards with more people.

Upshot: Base motives (e.g. desire for credit) may lead to more optimal division of labor in science than would scientists with "pure motivations" who wish to contribute to the research project with greatest chances of success.

Topics

Topics we'll discuss today:

- Plotting
- Behaviorspace
- Extensions

References I

Anderson, K. (2012). Skill specialization and the formation of collaboration networks. *Unpublished Manuscript*.

Hong, L. and Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences of the United States of America*, 101(46):16385—16389.

Kitcher, P. (1990). The division of cognitive labor. *The Journal of Philosophy*, 87(1):5—22.

Mayo-Wilson, C. (2012). The reliability of testimonial norms in scientific communities. *Forthcoming in Synthese*.

Merton, R. K. (1957). Priorities in scientific discovery: a chapter in the sociology of science. *American sociological review*, 22(6):635659.

Merton, R. K. (1968). The matthew effect in science. *Science*, 159(3810):5663.

Strevens, M. (2003). The role of the priority rule in science. *The Journal of philosophy*, 100(2):55—79.

References II

Strevens, M. (2006). The role of the matthew effect in science. *Studies In History and Philosophy of Science Part A*, 37(2):159—170.

Weisberg, M. and Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, 76(2):225—252.

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

Zollman, K. J. (2011). The communication structure of epistemic communities. In Goldman, A., editor, *Social Epistemology: Essential Readings*, pages 338—350.