Bounded Rationality and Local Interactions in ABMS

Models and Simulations in Philosophy
April 23rd, 2013

Last Week:

- Two Platonic Puzzles: Justice and Meaning
- ABMs vs. Equilibrium Explanations in Classical Economics and Mathematical Biology

ABMs vs. Classical Economic Models

Classic Models

- Rational

ABMs

ABMs vs. Classical Economic Models

Classic Models

- Rational

ABMs

- Boundedly Rational

ABMs vs. Classical Economic Models

Classic Models

- Rational
- Homogeneous agents

ABMs

- Boundedly Rational

ABMs vs. Classical Economic Models

Classic Models

- Rational
- Homogeneous agents

ABMs

- Boundedly Rational
- Heterogenous Agents

ABMs vs. Classical Economic Models

Classic Models

- Rational
- Homogeneous agents
- Global Interaction

ABMs

- Boundedly Rational
- Heterogenous Agents

ABMs vs. Classical Economic Models

Classic Models

- Rational
- Homogeneous agents
- Global Interaction

ABMs

- Boundedly Rational
- Heterogenous Agents
- Local interactions

ABMs vs. Classical Economic Models

Classic Models

- Rational
- Homogeneous agents
- Global Interaction
- Equilibria

ABMs

- Boundedly Rational
- Heterogenous Agents
- Local interactions

ABMs vs. Classical Economic Models

Classic Models

- Rational
- Homogeneous agents
- Global Interaction
- Equilibria

ABMs

- Boundedly Rational
- Heterogenous Agents
- Local interactions
- Dynamics

Getting More Precise

But we did **not**:

- Explain what rational meant, let alone boundedly rational

But we did **not**:

- Explain what rational meant, let alone boundedly rational
- Explain how one might represent local interactions among agents

Getting More Precise

But we did **not**:

- Explain what rational meant, let alone boundedly rational
- Explain how one might represent local interactions among agents
- Investigate dynamic population-level models and contrast them with ABMs

Outline

Outline

1. Review
2. Rationality: Bounded and Not
   - Classic Decision Theory
   - Bounded Rationality

Outline

Outline

Outline

	Sun	Rain
Read	2	3
Biergarten	4	-2
Listen to Nickelback	-10	-10

	Sun	Rain
Read	2	3
Biergarten	4	-2
Listen to Nickelback	-10	-10

Decision Matrices: Like game matrices, except one of the "players" is "Nature", which is responsible for the state of the world (generally, in columns).

	Sun	Rain
Read	2	3
Biergarten	4	-2
Listen to Nickelback	-10	-10

Decision Matrices: Payoffs to the decision-maker depend upon the unknown state of nature and what action she chooses (in the rows).

Dominance

	Sun	Rain
Read	2	3
Biergarten	4	-2
Listen to Nickelback	-10	-10

	Sun	Rain
Read	2	3
Biergarten	4	-2
Listen to Nickelback	-10	-10

Dominance: If the outcome of some action a_1 (e.g., Listen to Nickelback) is worse than that of another a_2 (e.g., Read) regardless of the state of the world, do not choose a_1.

Worst-Case

	Sun	Rain
Read	2	3
Biergarten	4	-3

Worst-Case

	Sun	Rain
Read	2	3
Biergarten	4	-3

Worst-Case: Each action has a worst-case payoff. E.g., For Read, it's 2. For Biergarten, it's -3.

	Sun	Rain
Read	2	3
Biergarten	4	-3

	Sun	Rain
Read	2	3
Biergarten	4	-3

Minimax: Pick the action with the best worst-case payoff. Here, it's Read.

- But suppose you look outside, and it's a beautiful spring day in Munich.

- But suppose you look outside, and it's a beautiful spring day in Munich.
- You read the weather forecast, which claims the chance of rain is .5%.

Decision Matrices

- But suppose you look outside, and it's a beautiful spring day in Munich.
- You read the weather forecast, which claims the chance of rain is .5%.
- Minimax ignores the probability of rain.

Decision Matrices

- But suppose you look outside, and it's a beautiful spring day in Munich.
- You read the weather forecast, which claims the chance of rain is .5%.
- Minimax ignores the probability of rain.
- We'd like some decision rule that simultaneously considers payoffs/losses and probability.

Suppose you fully believe the weather forecast, which claims the chance of rain is .5%.

	Sun	Rain
Read	2	3
Biergarten	4	-3

The expected utility of Biergarten is:

$$\begin{aligned} \text{SEU}(Biergarten) &= p(Sun) \cdot 4 + p(Rain) \cdot -3 \\ &= 995 \cdot 4 + .005 \cdot -3 \\ &= 3.965 \end{aligned}$$

Suppose you fully believe the weather forecast, which claims the chance of rain is .5%.

	Sun	Rain
Read	2	3
Biergarten	4	-3

In contrast, expected utility of Read is:

$$\begin{aligned} \textsc{seu}(\mathit{Read}) &= p(\mathit{Sun}) \cdot 2 + p(\mathit{Rain}) \cdot 3 \\ &= 995 \cdot 2 + .005 \cdot 3 \\ &= 2.005 \end{aligned}$$

Three Decision Rules

- Maximize (subjective) expected utility (SEU)

Three Decision Rules

- Maximize (subjective) expected utility (SEU)
- Dominance
- Minimax

- The Standard in Economics: An agent is rational if she acts **as if** she were maximizing expected utility.

- The Standard in Economics: An agent is rational if she acts **as if** she were maximizing expected utility.
- That is, the agent may not act **with the intent** of maximizing expected utility. She may happen to do maximize utility accidentally or uncounsciously (due to practice and training, or genetic predisposition).

- The Standard in Economics: An agent is rational if she acts **as if** she were maximizing expected utility.
- That is, the agent may not act **with the intent** of maximizing expected utility. She may happen to do maximize utility accidentally or uncounsciously (due to practice and training, or genetic predisposition).
- There are a number of arguments for the claim that expected utility maximization is the **unique** rational decision rule; we won't discuss them here.

Bounded Rationality

- "Bounded rationality" is a term of art. No one defines it.

Bounded Rationality

- "Bounded rationality" is a term of art. No one defines it.
- Here is **my** best attempt to explain what I think is meant.

Bounded Rationality

- "Bounded rationality" is a term of art. No one defines it.
- Here is **my** best attempt to explain what I think is meant.
- Call an agent deliberately rational if she is rational and acts **with the purpose** of maximizing expected utility.

An agent is boundedly rational if she is

- **Neither** deliberately rational nor rational (simpliciter),
- but, in certain important contexts, she is rational or approximates rationality (simpliciter).

Deliberate Rationality

What does it take to be deliberately rational?

Deliberate EU Maximization

To act with the purpose of maximizing expected utility (i.e. to be deliberately rational), one must

- Consider all available actions,

Deliberate eu Maximization

To act with the purpose of maximizing expected utility (i.e. to be deliberately rational), one must

- Consider all available actions,
- Consider all possible states of the world,

To act with the purpose of maximizing expected utility (i.e. to be deliberately rational), one must

- Consider all available actions,
- Consider all possible states of the world,
- Know the payoff of each action in each possible state of the world,

To act with the purpose of maximizing expected utility (i.e. to be deliberately rational), one must

- Consider all available actions,
- Consider all possible states of the world,
- Know the payoff of each action in each possible state of the world,
- Assign each state of the world some probability of occurring,

Deliberate EU Maximization

To act with the purpose of maximizing expected utility (i.e. to be deliberately rational), one must

- Consider all available actions,
- Consider all possible states of the world,
- Know the payoff of each action in each possible state of the world,
- Assign each state of the world some probability of occurring,
- Calculate the SEU of each action and compare the results.

Each of these tasks may be difficult for fallible agents with limited time, memory, and computational abilities . . .

Here are some cases in which each of the tasks may be difficult:

- Consider all available actions,
    - Ways to get from upper Manhattan to Brooklyn
    - Chess strategies
    - Any decision where many actions are available.

Here are some cases in which each of the tasks may be difficult:

- Consider all available actions,
  - Ways to get from upper Manhattan to Brooklyn
  - Chess strategies
  - Any decision where many actions are available.
- Consider all possible states of the world,
  - Possible traffic patterns in Manhattan.

Here are some cases in which each of the tasks may be difficult:

- Know the payoff of each action in each possible state of the world,
  - How long will a taxi take if the Brooklyn bridge is congested, there is construction in midtown, etc.?
  - Economic and/or Social policy. What will be the effect of austerity measures in Europe?

Difficulty of Deliberation

Here are some cases in which each of the tasks may be difficult:

- Know the payoff of each action in each possible state of the world,
  - How long will a taxi take if the Brooklyn bridge is congested, there is construction in midtown, etc.?
  - Economic and/or Social policy. What will be the effect of austerity measures in Europe?
- Assign each state of the world some probability of occurring,
  - What's the probability there is construction in midtown today?
  - Climate Policy - What is the probability that mean global surface temperature rises by at least two degrees centigrade?

Difficulty of Deliberation

Here are some cases in which each of the tasks may be difficult:

- Know the payoff of each action in each possible state of the world,
  - How long will a taxi take if the Brooklyn bridge is congested, there is construction in midtown, etc.?
  - Economic and/or Social policy. What will be the effect of austerity measures in Europe?
- Assign each state of the world some probability of occurring,
  - What's the probability there is construction in midtown today?
  - Climate Policy - What is the probability that mean global surface temperature rises by at least two degrees centigrade?
- Calculate each action's SEU and compare the results.
  - Any of the cases above.

By my definition, a boundedly rational agent fails to do one of the following:

- Consider all available actions,
- Consider all possible states of the world,
- Know the (numerical) payoff of each action in each possible state of the world,
- Assign each state of the world some probability of occurring,
- Calculate SEU of all actions and compare the results.

Instead, a boundedly rational agent might decide as follows:

Bounded Rationality

Boundedly rational agents may

- Consider some small subset of available actions
  - E.g., Kasparov considers only three moves a second.

Bounded Rationality

Boundedly rational agents may

- Consider some small subset of available actions
  - E.g., Kasparov considers only three moves a second.
- Consider some small subset of possible states of the world,
  - E.g., Individuals generally ignore low probability states of the world (e.g., airplane crashes) when making particular decisions (e.g., about whether to fly or take a train)

Boundedly rational agents may

- Make qualitative comparisons of outcomes of actions
  - E.g., In the state of the world in which it rains, I might consider taking an umbrella to better than going without one, but I don't know *how much* better.

Boundedly rational agents may

- Make qualitative comparisons of outcomes of actions
  - E.g., In the state of the world in which it rains, I might consider taking an umbrella to better than going without one, but I don't know *how much* better.
- Make qualitative judgments of likelihood; or omit judgments of likelihood at all in other circumstances.
  - E.g., Qualitative Probability - I believe it is likely that it will rain this week, but I do not have an exact percentage (e.g., 95%) that quantifies my belief.
  - E.g. Omission - If you use minimax reasoning, for instance.

Boundedly rational agents may

- Not perform utility calculations and compare all actions.
  - Omit calculations; use rules of thumbs or heuristics to compare actions.
  - We are not all like Darwin when it comes to marriage decisions.

We'll examine some more examples shortly . . .

Outline

- Remember, Fisher argued that if the male/female ratio were not 1:1, then organisms that tended to have more male (resp. female) offspring would have an evolutionary advantage in a mostly female (resp. male) population.

- Remember, Fisher argued that if the male/female ratio were not 1:1, then organisms that tended to have more male (resp. female) offspring would have an evolutionary advantage in a mostly female (resp. male) population.
- Recall, we discussed the fact that Fisher's argument seems to assume large populations in which each organism has a large number of potential mates . . .

The Importance of Local Interactions

- Remember, Fisher argued that if the male/female ratio were not 1:1, then organisms that tended to have more male (resp. female) offspring would have an evolutionary advantage in a mostly female (resp. male) population.
- Recall, we discussed the fact that Fisher's argument seems to assume large populations in which each organism has a large number of potential mates . . .
- It is less easy to see why an organism with few potential mates (e.g., a person born in Wyoming) has any evolutionary advantage, regardless of its genetic tendency to produce offspring of one sex.

Representing Local Interactions

How can we represent local interactions, in which agents typically communicate (or not), cooperate (or not) etc., with a fraction of the population?

Networks

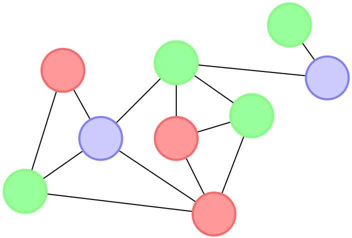

Nodes = Agents
Edges = Indicate which agents "interact"
Colors = "Type" of Agent

Common Social Networks

What are some common social networks that might be represented in this way?

Common Social Networks

What are some common social networks that might be represented in this way?

- Facebook (Edges indicate the "friend" relation)

Common Social Networks

What are some common social networks that might be represented in this way?

- Facebook (Edges indicate the "friend" relation)
- The world-wide web (Edges indicate reciprocal links)

Common Social Networks

What are some common social networks that might be represented in this way?

- Facebook (Edges indicate the “friend” relation)
- The world-wide web (Edges indicate reciprocal links)
- Co-authorship networks (Edges indicate co-authors)

What are some common social networks that might be represented in this way?

- Facebook (Edges indicate the “friend” relation)
- The world-wide web (Edges indicate reciprocal links)
- Co-authorship networks (Edges indicate co-authors)
- Actor Network (Edges indicate the actors have appeared in a movie together)

Common Features of Social Networks

Many social networks share network structure.

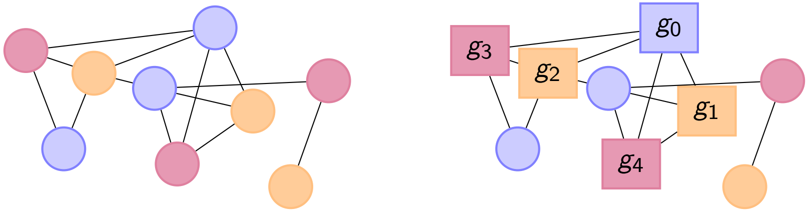

g_0's neighborhood

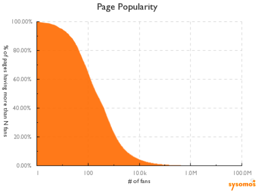

- Power law degree distribution
  - The degree of a node is the number of its neighbors.
  - A power law degree distribution indicates most agents have a few neighbors; few have a modest number; even fewer have many; etc.

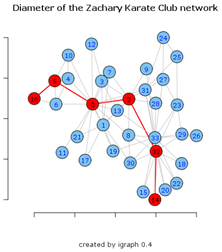

- Small diameter and average-path-length
  - The diameter is the longest path in the network
  - The average-path-length is what it sounds like.
  - E.g., [Milgram, 1967]'s small world experiment
  - E.g., The Kevin Bacon Game
  - E.g., Erdös Numbers

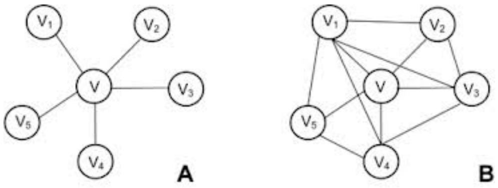

- High clustering
  - Roughly, a graph is clustered if your neighbors' neighbors are your neighbors.
  - E.g., The philosophers with whom I have co-authored papers have also generally co-authored papers with one another.

Average-path-length, diameter, degree distributions, etc. are called network structure.

The next four chapters of [Alexander, 2007] have the same structure:

1. Alexander introduces a two-person game (e.g., the prisoner's dilemma or stag hunt)

The next four chapters of [Alexander, 2007] have the same structure:

1. Alexander introduces a two-person game (e.g., the prisoner's dilemma or stag hunt)
2. He argues that particular actions in the game correspond to particular norms of behavior:
   - Silence in the prisoner's dilemma = Cooperation
   - Hunting Stag = Trust

The next four chapters of [Alexander, 2007] have the same structure:

3. He imagines that agents in a network play the game repeatedly with their neighbors, and use some boundedly rational strategy for learning what action to play next.

The next four chapters of [Alexander, 2007] have the same structure:

3. He imagines that agents in a network play the game repeatedly with their neighbors, and use some boundedly rational strategy for learning what action to play next.
4. He analyzes (i) boundedly rational strategies and (ii) network structures that cause the norm to spread throughout the network.

Imitate the Best and Bounded Rationality

Networks are useful not only for representing local interactions, but also for modeling bounded rationality.

Imitate the Best and Bounded Rationality

Networks are useful not only for representing local interactions, but also for modeling bounded rationality.

In what way is "imitate the best" a bounded rational strategy?

- Does the agent compare all available actions?

- Does the agent compare all available actions?
  - No. She considers only those actions (i) employed by her neighbors and (ii) that were employed on the last stage of inquiry.
  - The agent may not even know which actions are available.

Bounded Rationality

- Does the agent consider all possible states of the world?

Bounded Rationality

- Does the agent consider all possible states of the world?
  - No. In fact, the decision rule doesn't consider states of the world at all.
  - The agent only considers what actions were performed and their payoffs. She may be completely ignorant that payoffs depend upon some unknown "state of the world."

Bounded Rationality

- Does the agent know the payoffs of each action in each state of the world?

- Does the agent know the payoffs of each action in each state of the world?
  - No. As above, she need not even that states of the world exist.

Bounded Rationality

- Does the agent know the payoffs of each action in each state of the world?
  - No. As above, she need not even that states of the world exist.
- Does the agent assign probabilities to each state of the world?

- Does the agent know the payoffs of each action in each state of the world?
  - No. As above, she need not even that states of the world exist.
- Does the agent assign probabilities to each state of the world?
- Does the agent perform calculations and then compare actions?
  - A bit. The agent finds the best action employed by her neighbors. That involves some comparisons, but it requires no calculations.

Outline

Topics

Topics we'll discuss today:

- Declaring vs. Modifying Variables
- Data-Types: Numbers, Strings, Booleans, and Lists
- Operations on variables of all four types
- Printing to the command center
- Global vs. Local Variables

References I

Alexander, J. M. (2007). *The structural evolution of morality*. Cambridge University Press Cambridge.

Milgram, S. (1967). The small world problem. *Psychology today*, 2(1):60–67.