Collaborative Learning in Networks And Web-Based Experiments

Duncan Watts

Yahoo! Research

Outline of Talk

- Introduce substantive problem of collaborative learning in networks
- Describe a recent web-based experiment that examines a simple version of the general problem
- Discuss potential for web-based experiments in general, and potential for "experimental macro-sociology"

Complex Problems

- Many problems in science, business, and engineering are “complex” in the sense that they exhibit:
  - Multiplicity of potential solutions
    - In any given competitive situation, many designs/strategies/procedures are conceivably optimal
  - Interdependence among different parts of a solution
    - Changing the value of one parameter can have different effects depending on values of other parameters
- Complexity is often represented abstractly as a “fitness landscape:” a mapping between the parameters of a potential solution and its corresponding “fitness” (Kauffman, 93; Levinthal, 97)
  - “Simple” problems have smooth fitness landscapes, with a single peak
  - “Complex” problems have “rugged” landscapes with many peaks, separated by valleys

Exploration vs. Exploitation

- Successfully navigating a rugged fitness landscape requires some balance between exploitation of known solutions, and exploration for potentially better solutions (March 1991)
  - Too much exploitation leads to suboptimal long-run performance
  - Too much exploration is costly and forgoes short-run advantages of exploitation
- Exploration-Exploitation tradeoff is standard element of all complex optimization algorithms in CS, Stat Mech
  - MCMC algorithms, Markov decision processes, genetic algorithms, etc.
- Here we are less interested in how to solve complex problems optimally (in an algorithmic sense) than how people/organizations actually solve them
  - Unclear what "real" fitness landscapes look like
  - Unclear what the problems solvers know about the landscapes they are navigating

Exploration-Exploitation Tradeoff Arises in Many Forms

- Rational Search (e.g. Radner)
  - N projects with unknown payoffs distributions
  - Must choose between learning current distribution vs. exploring new ones
- Boundedly Rational Search
  - Like above, but with cognitive biases
    - Satisficing (Simon)
    - Prospect Theory (Kahneman and Tversky)
- Organizational Learning (March 1991)
  - Refinement vs. Invention
  - Basic Science vs. Development
- Evolutionary Models (in particular, of organizations)
  - Variation vs. Selection (Hannan and Freeman)

Collaborative Learning In Communication Networks

- In many contexts, the exploration-exploitation tradeoff is complicated by the presence of other problem solvers
  - Potentially helpful because individuals can learn from the experience of others, thus improving collective learning
  - Potentially harmful because learning may also lead the collective may converge on a suboptimal solution
- Information flow within an organization therefore likely to impact its problem solving abilities (Leavitt, 1951; Lazer and Friedman, 2007; Mason et al. 2008,)

Individual vs. Collective Learning

- In a network context, not all individuals are equal
  - Individuals with “central” or “bridging” positions stand to gain from exposure to novel information, complementary ideas, or brokerage opportunities (Granovetter, Burt)
- Again, unclear whether differences in network positions across individuals are good or bad for collective performance
  - Conceivably central or bridging actors can produce efficiencies that are shared by all
  - But opportunity to gain relative advantage may also lead to conflict between individual and collective interests

The Current Project

- Substantive Questions
  - How do individuals collaboratively solve (certain kinds of) “complex” problems?
  - How does the structure of the communication network between them contribute to their collective performance?
  - How does individual position in the network relate to
    - Individual strategy and performance?
    - Collective performance?
- Our approach mostly experimental
  - Seek to exploit recent advances in web-based experimentation, esp. Amazon’s Mechanical Turk (AMT)
- But have also verified experiments with simulations
  - Not discussed today

Screenshot of Experiment

California

To play, click anywhere on the panel on the right to choose a spot to drill for oil. When you are ready, click 'Submit' to confirm your choice.

On the next round and every round after, you will see where you and the other players drilled and how much oil you each got for drilling at that spot.

Hover over the locations or the bars to see the exact amount earned; blacker bars mean more oil.

You always get the same amount of oil per round at any location, no matter how many other people are drilling there.

There will be 15 rounds, and you will be paid based on how much total oil you get across the 15 rounds.

X: Y:

Submit

Time left:
0:45
Current Round: 1

TOTAL: 0

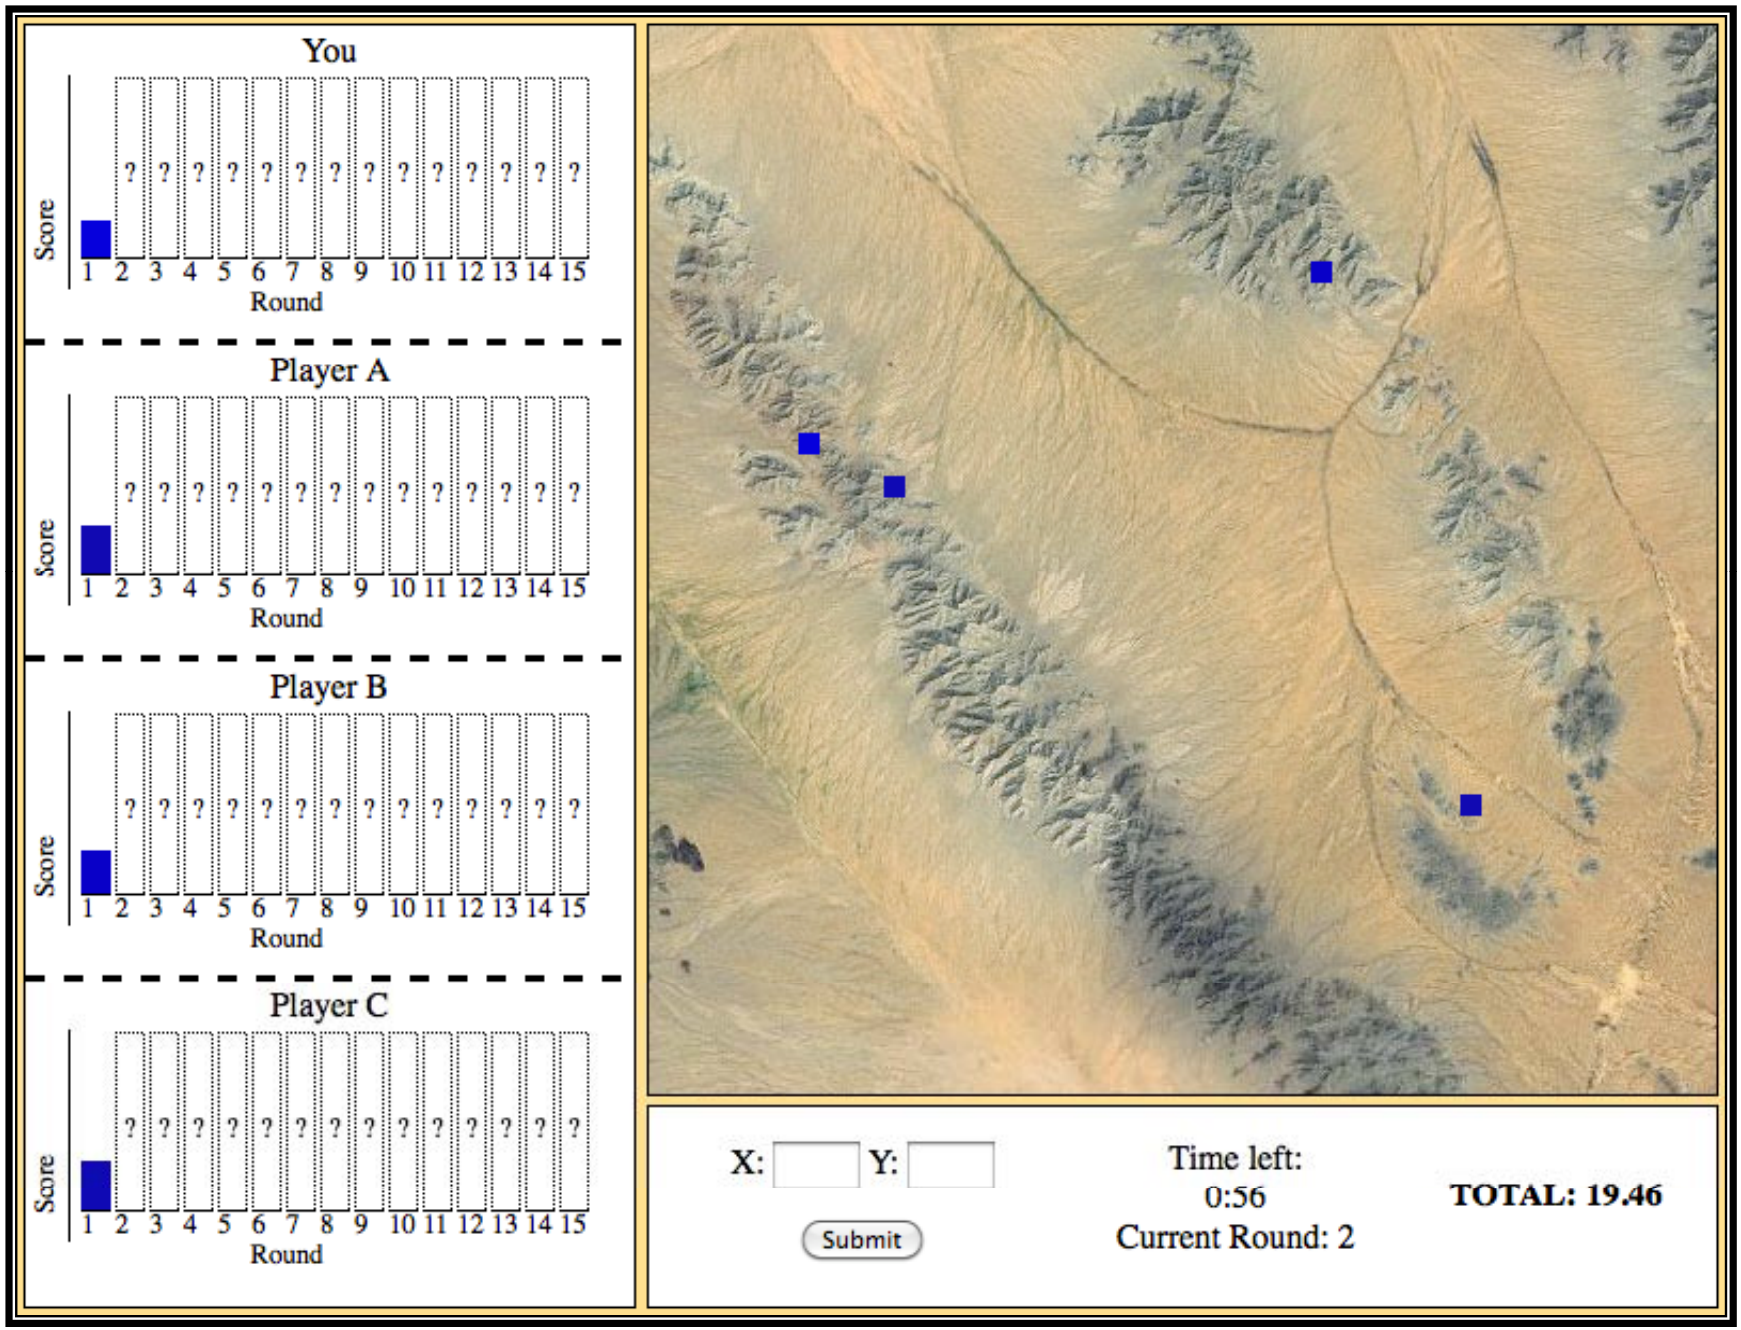
You
Score
? ? ? ? ? ? ? ? ? ? ? ? ? ? ?
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
Round
Player A
Score
? ? ? ? ? ? ? ? ? ? ? ? ? ? ?
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
Round
Player B
Score
? ? ? ? ? ? ? ? ? ? ? ? ? ? ?
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
Round
Player C
Score
? ? ? ? ? ? ? ? ? ? ? ? ? ? ?
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
Round
X:
Y:
Submit
Time left:
0:56
Current Round: 2
TOTAL: 19.46

You

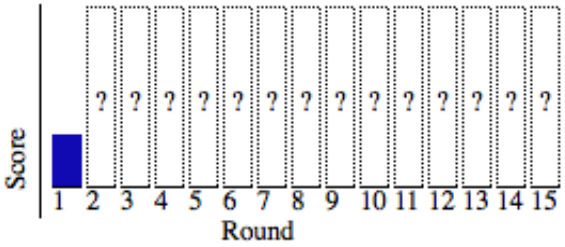

Player A

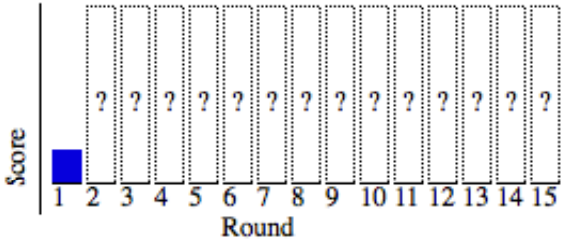

Player B

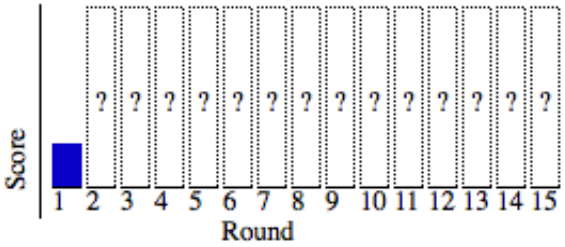

Player C

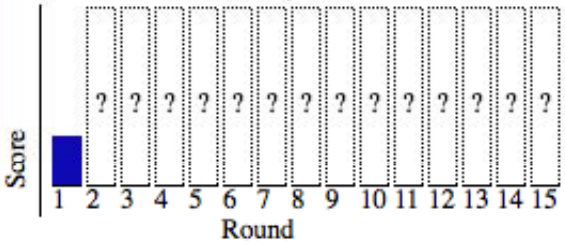

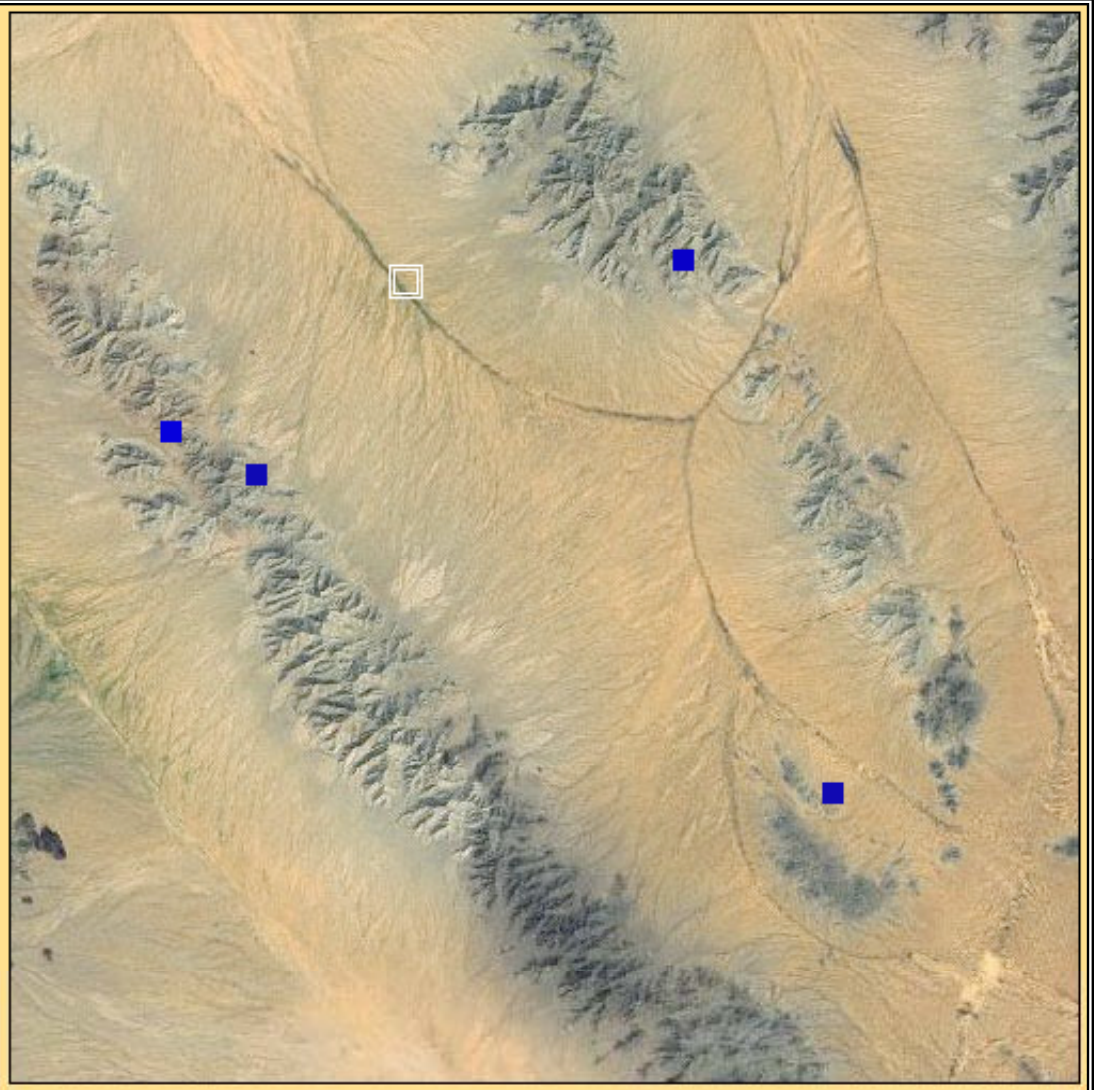

X: 18 Y: 37

Waiting for 2 other players to submit their score.

Time left:
0:38
Current Round: 2

TOTAL: 28.3

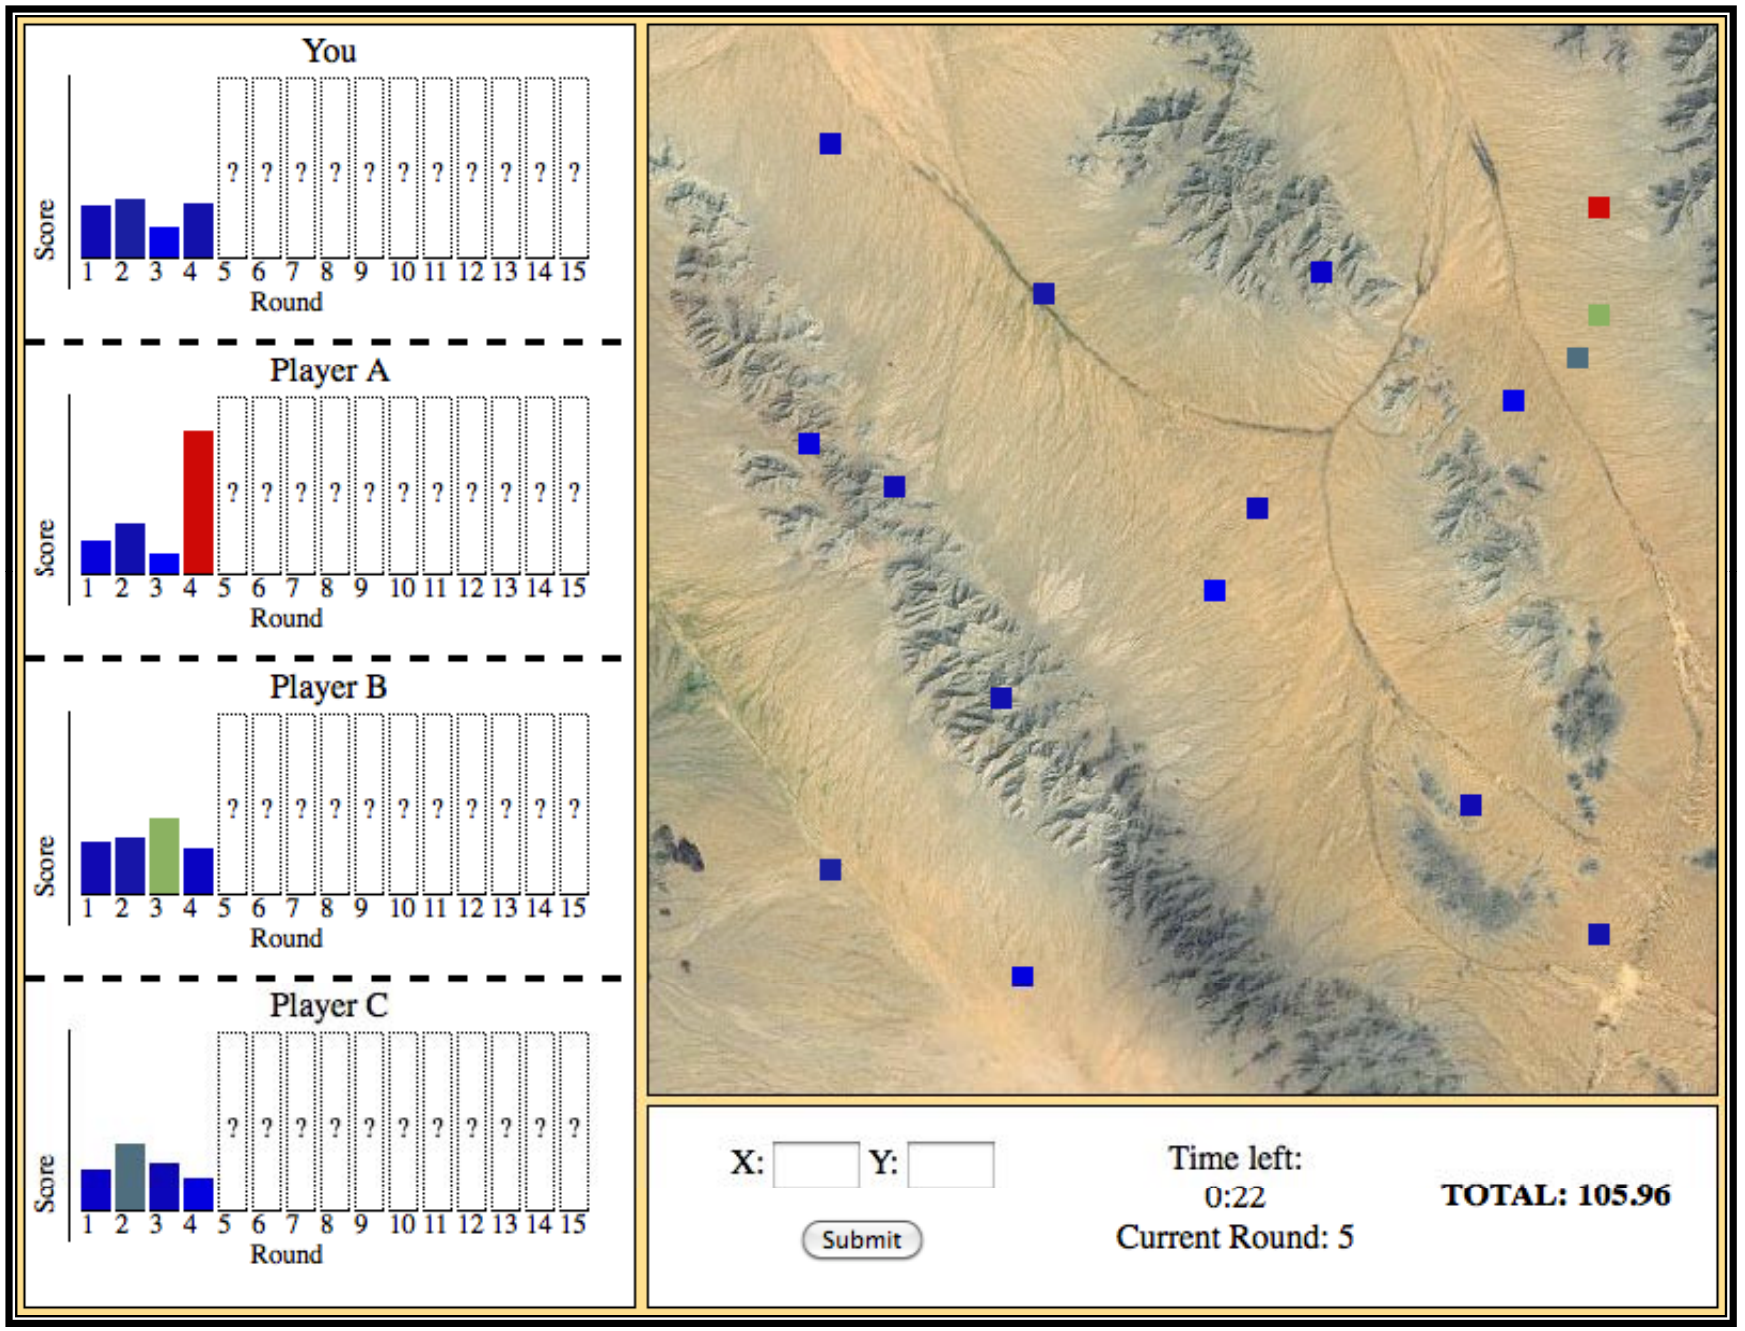
You
Score
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
Round
Player A
Score
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
Round
Player B
Score
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
Round
Player C
Score
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
Round
X:
Y:
Submit
Time left:
0:22
Current Round: 5
TOTAL: 105.96

Generating The Fitness Landscape

- Background generated with 4-octave 2D Perlin noise
  - procedure for generating pseudo-random noise, used to create realistic looking landscapes
- Added to a unimodal 2D Gaussian function with mean chosen uniformly at random and SD = 3
- Normalized so maximum points = 100

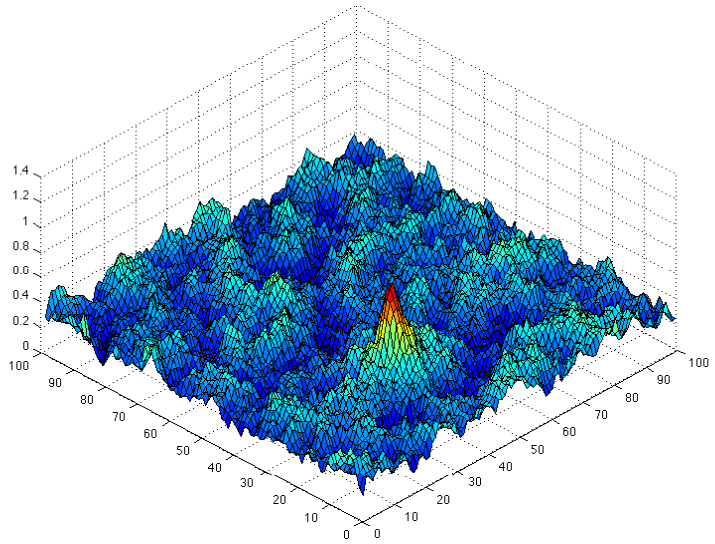

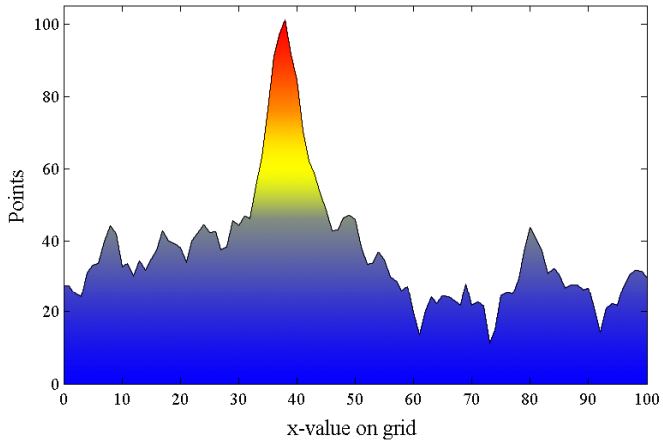

Generating the Networks

- Goal: 16-node fixed-degree graphs with extreme statistics
- Start with fixed-degree random graphs
  - All players have same amount of information
  - Only position in graph can affect success
- Rewire to increase or decrease some graph feature
  - Maximum, Average, Variance
  - Betweenness, Closeness, Clustering, Network Constraint
  - Ensuring connected graph
- Stop when no rewiring improves feature
- Repeat 100 times, keep maximal graph

Network Features of Interest

- Clustering:
  - Average probability that two neighbors are themselves connected (local density)
- Betweenness
  - Number of shortest paths that pass through node
- Closeness
  - Average shortest path to all other nodes
- Network Constraint
  - Redundancy with neighbors $nc(i) = \frac{1}{k^2} \sum_{j \in N(i)} (1 + \sum_{q \in N(i), q \neq j} w_{qj})^2$

Communication Networks

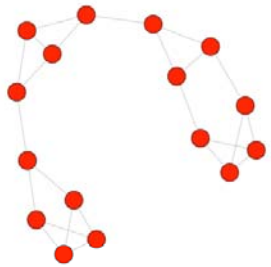

Greatest Average Betweenness

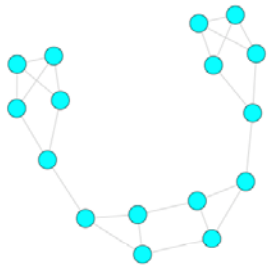

Smallest Maximum Closeness

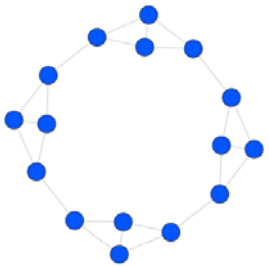

Greatest Average Clustering

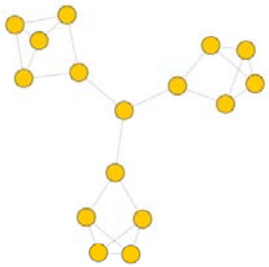

Greatest Maximum Betweenness

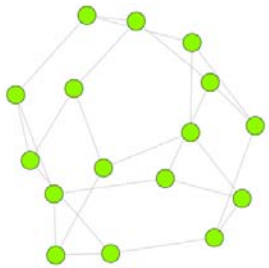

Smallest Average Betweenness

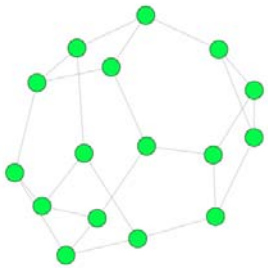

Greatest Maximum Closeness

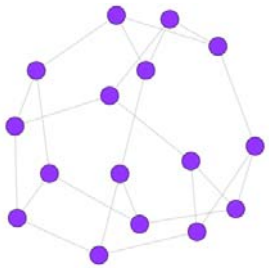

Smallest Average Clustering

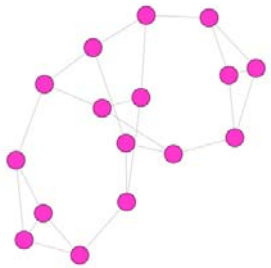

Greatest Variance in Constraint

All Individuals in all networks have 3 neighbors,
All Individuals have the same view of the world

Experiments

- For each session, 16 subjects are recruited from Amazon's Mechanical Turk
  - Standing panel alerted previous day
  - Accept work & read instructions
  - Sit in "waiting room" until enough players have joined
- Each session comprises 8 games
  - One for each network topology
- Each game runs for 15 rounds
  - 100 x100 grid
  - Relative dimensions of peak and landscape adjusted such that peak is found sometimes, but not always
- 171 out of total of 232 games (25 sessions) used
  - 61 games removed because player dropped out

Preview of Results

1. Network structure affects individual search strategy
2. Individual search strategy affects group success
3. Network structure also affects group success directly, via information diffusion
4. Individual and group performance are in tension

Network structure affects individual search strategy

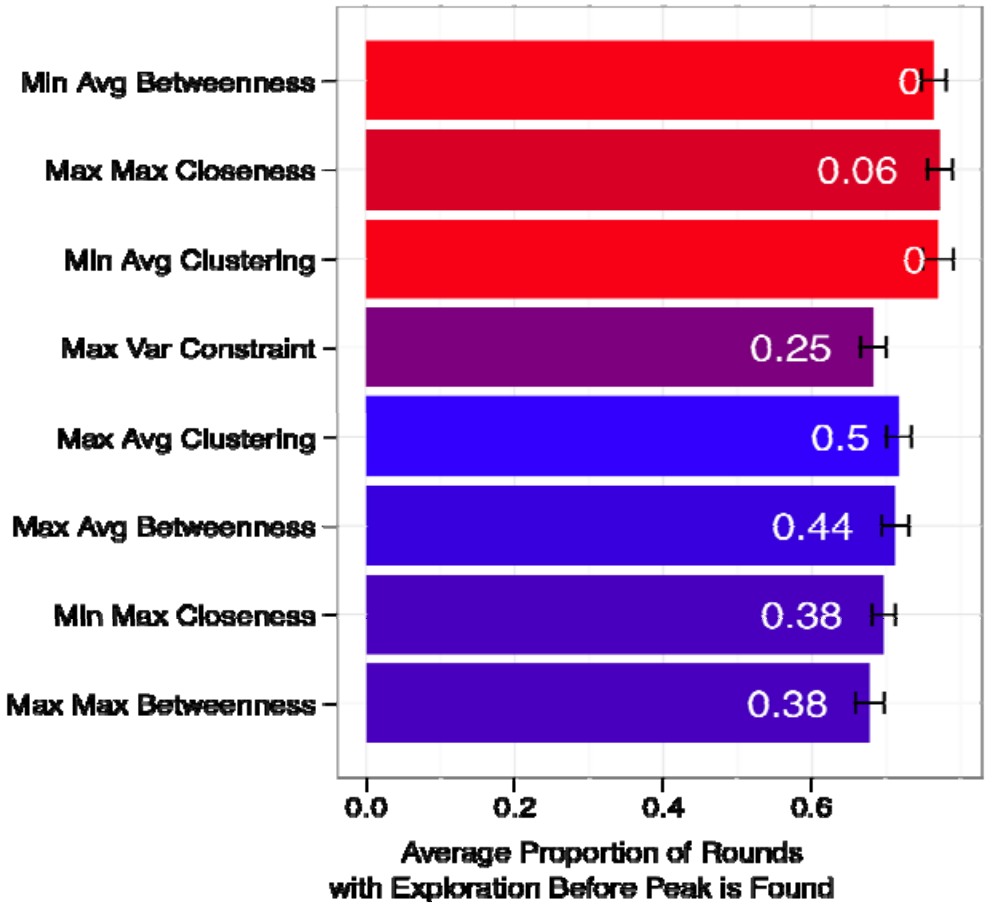

- Networks differ in amount of exploration
- Related to clustering

Network structure affects individual search strategy

- Higher clustering

  ⬇

  Higher probability of neighbors guessing in identical location

- More neighbors guessing in identical location

  ⬇

  Higher probability of copying

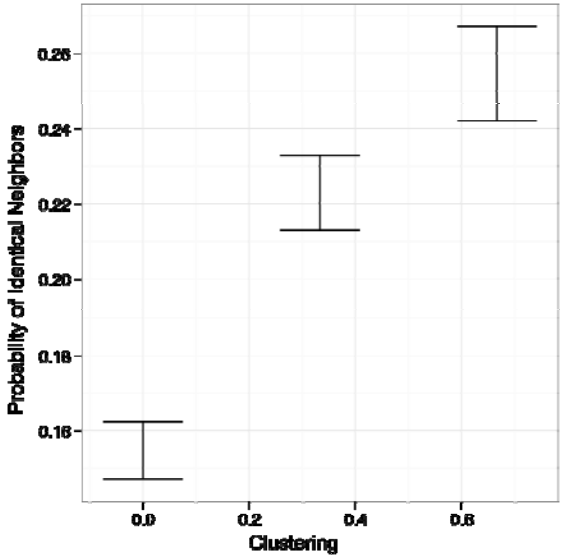

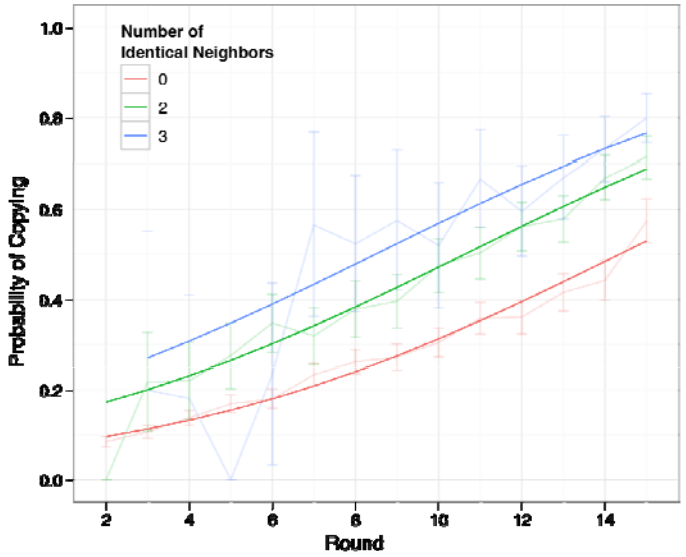

Individual search strategy affects group success

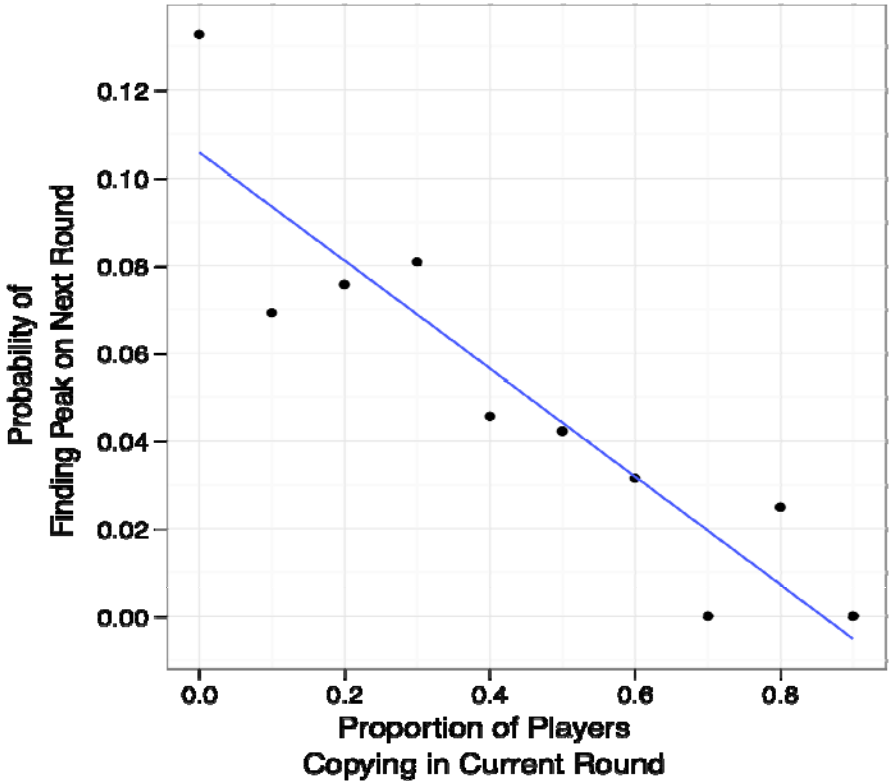

- More players copying each other (i.e., fewer exploring) in current round

↓

Lower probability of finding peak on next round

Individual search strategy affects group success

- No significant differences in % of games in which peak was found
- But pattern similar to differences in exploration

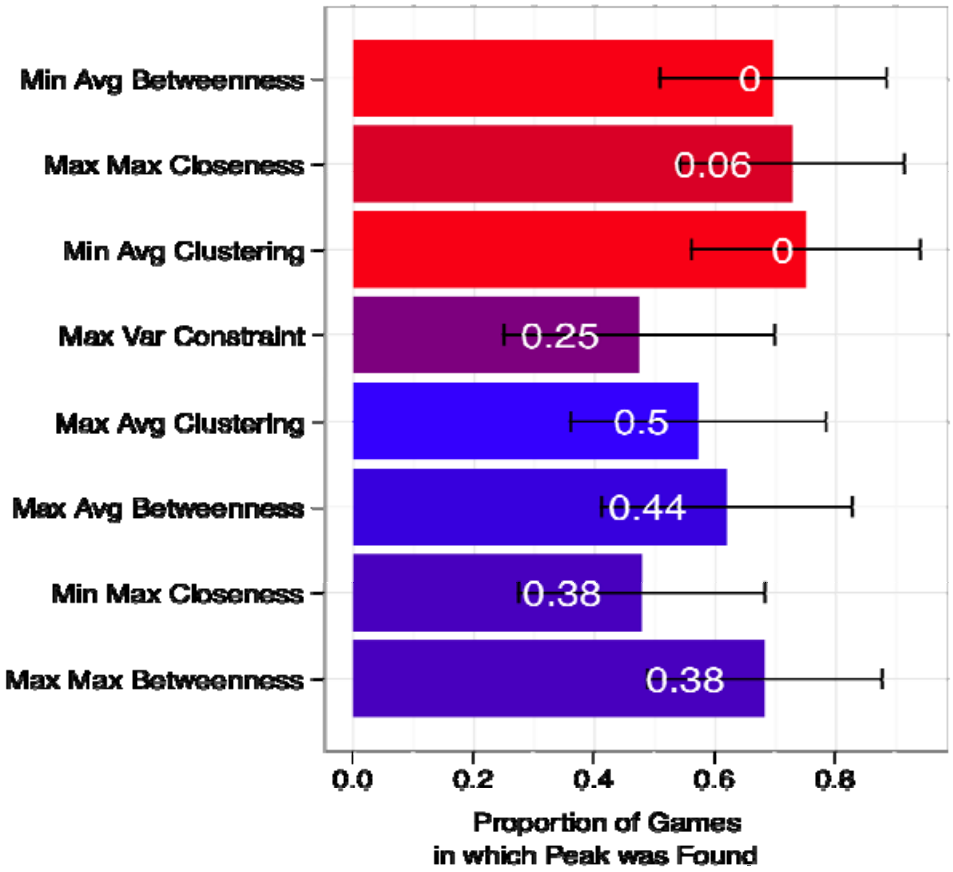

Communication Networks

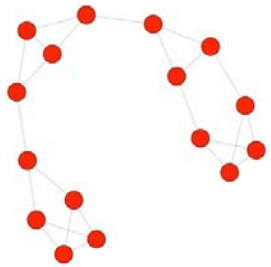

Greatest Average Betweenness

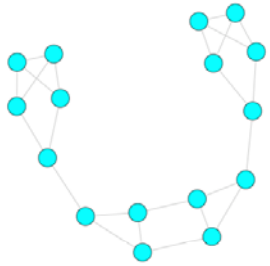

Smallest Maximum Closeness

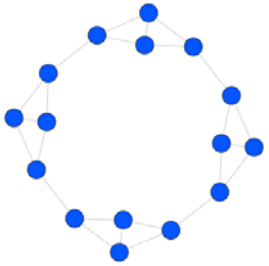

Greatest Average Clustering

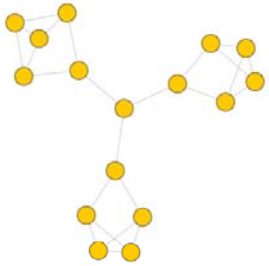

Greatest Maximum Betweenness

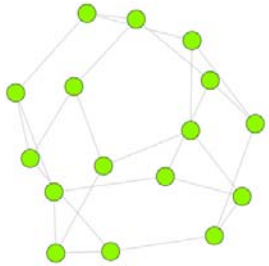

Smallest Average Betweenness

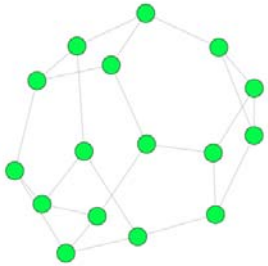

Greatest Maximum Closeness

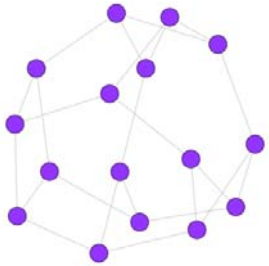

Smallest Average Clustering

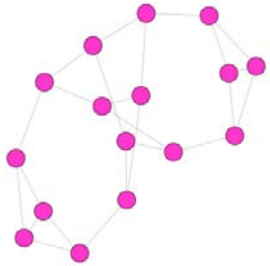

Greatest Variance in Constraint

Network structure affects group success

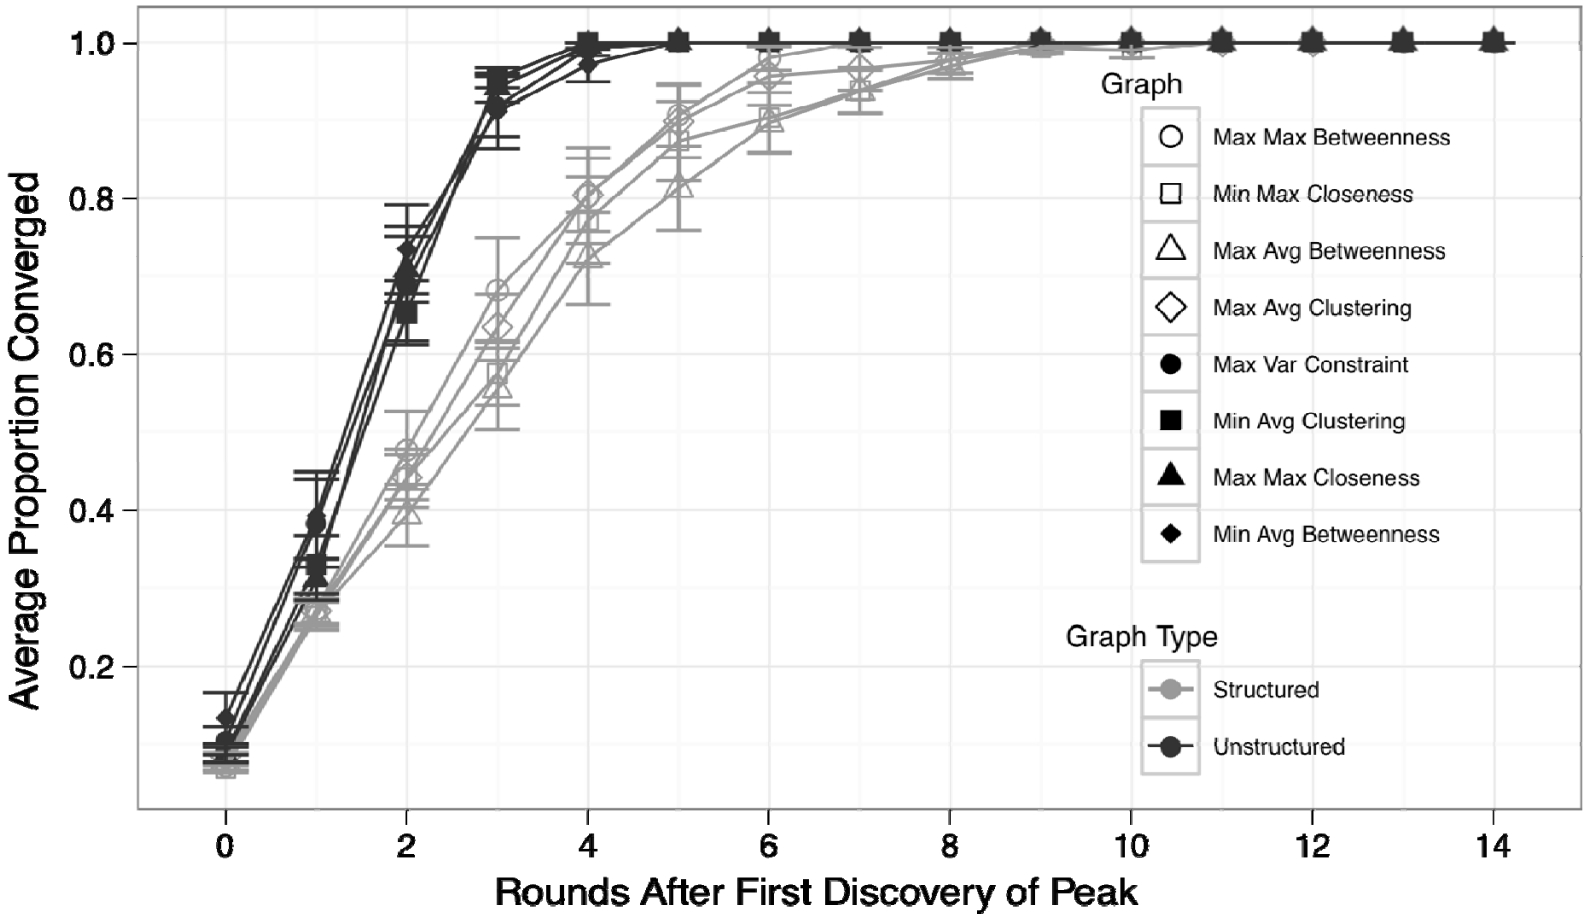

Diffusion of Best Solution

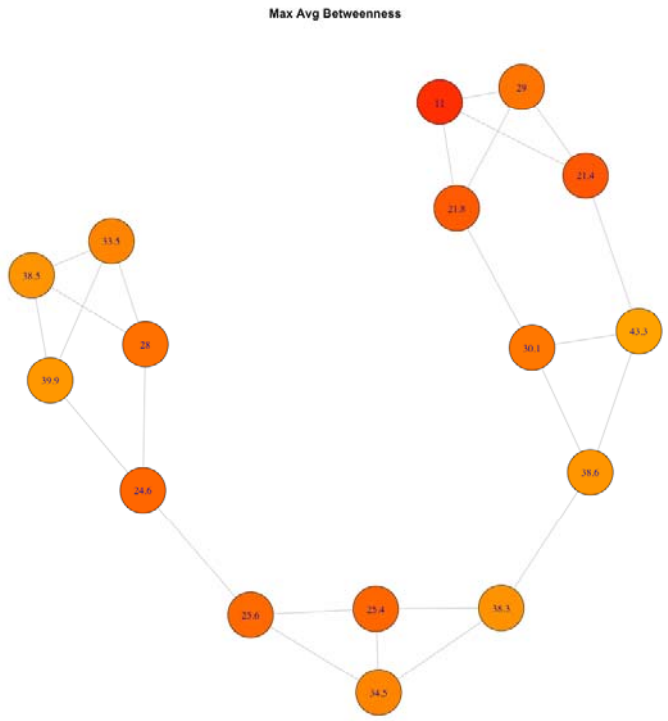

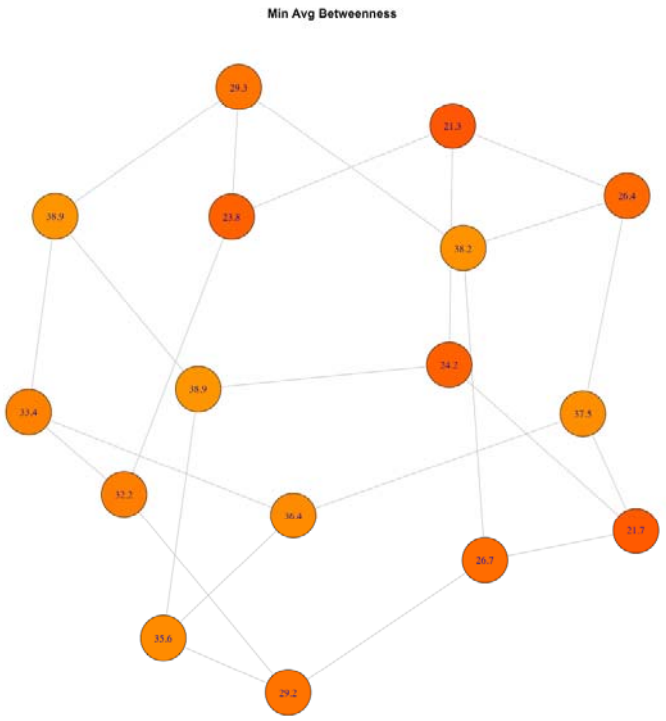

Diffusion of Best Solution

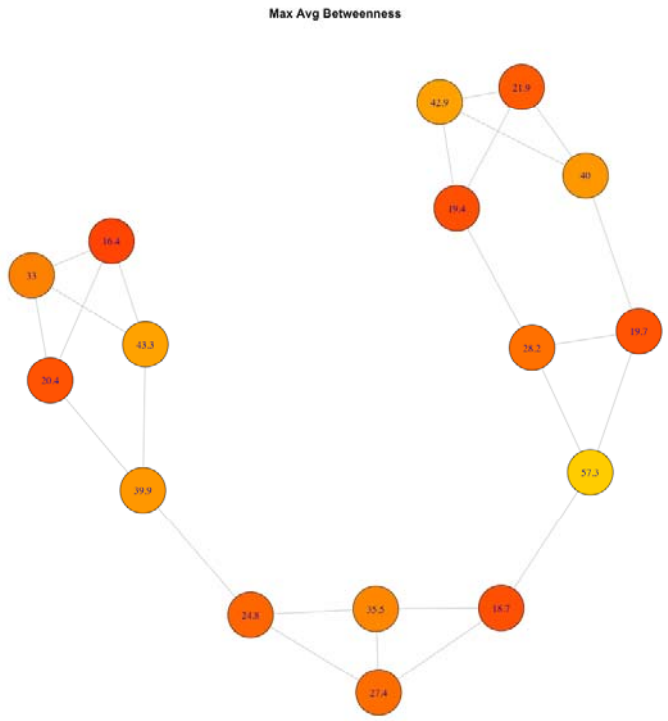

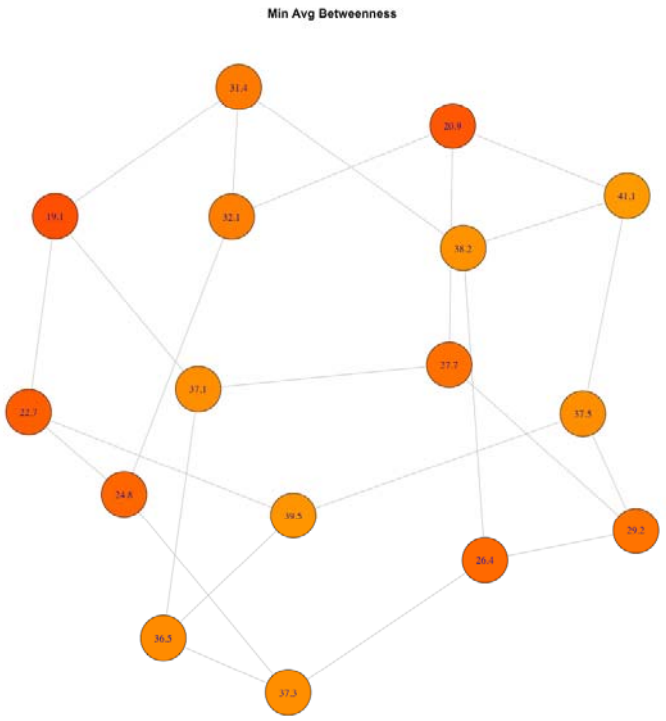

Diffusion of Best Solution

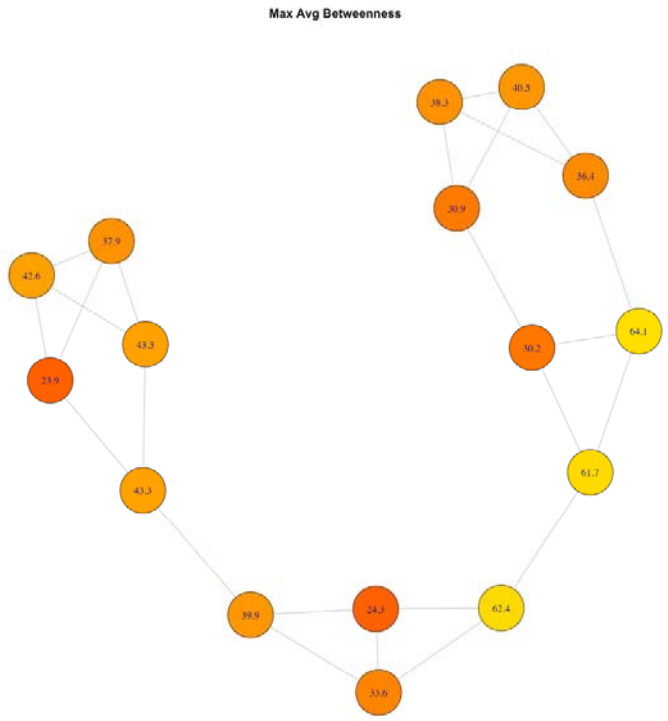

Exp 1039 Trial 8

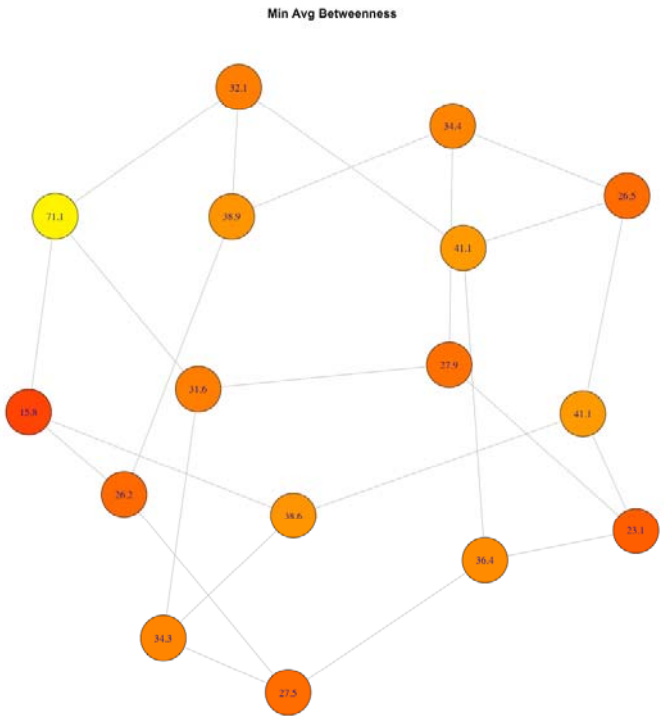

Exp 1037 Trial 6

Diffusion of Best Solution

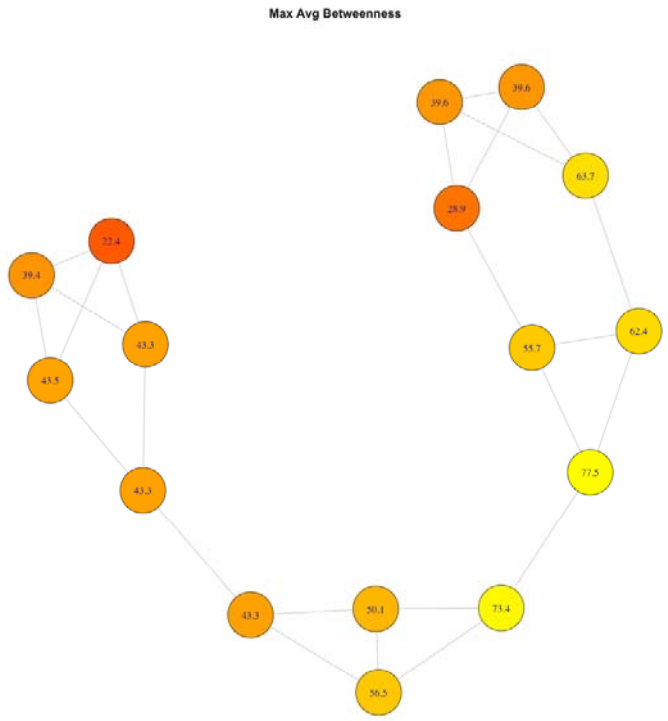

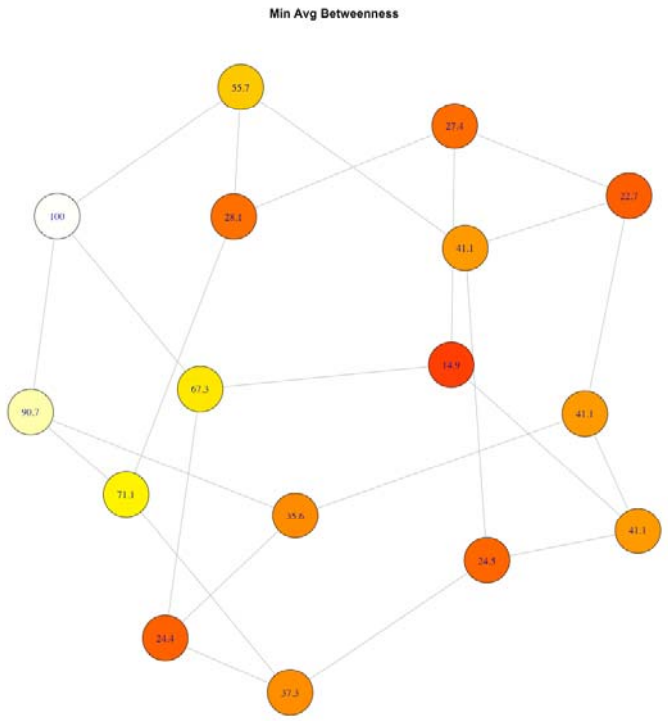

Diffusion of Best Solution

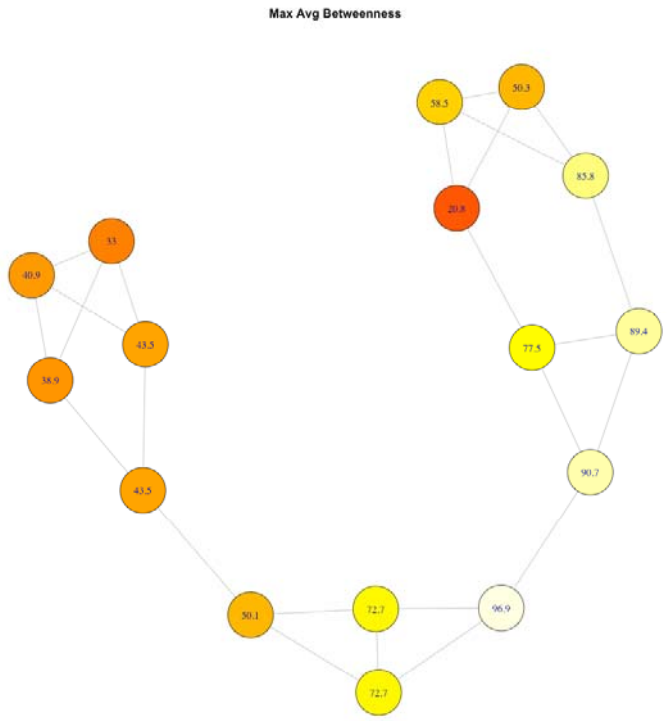

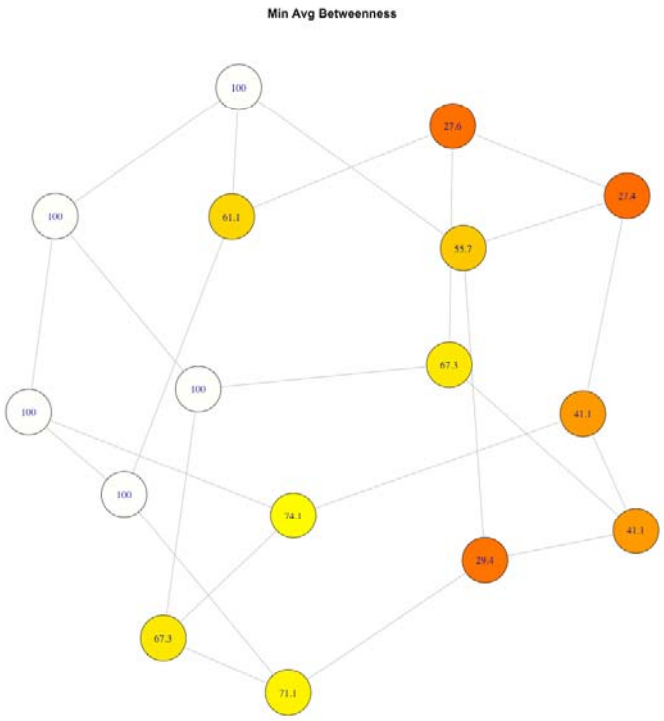

Diffusion of Best Solution

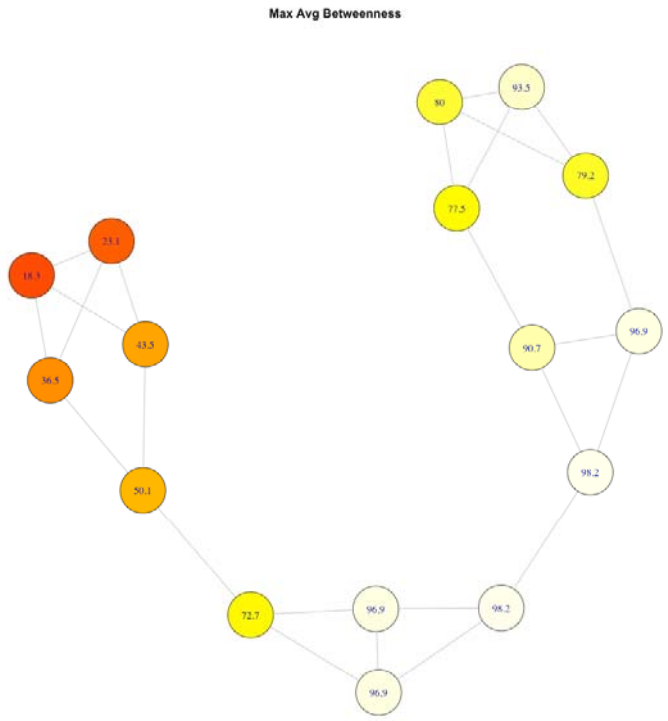

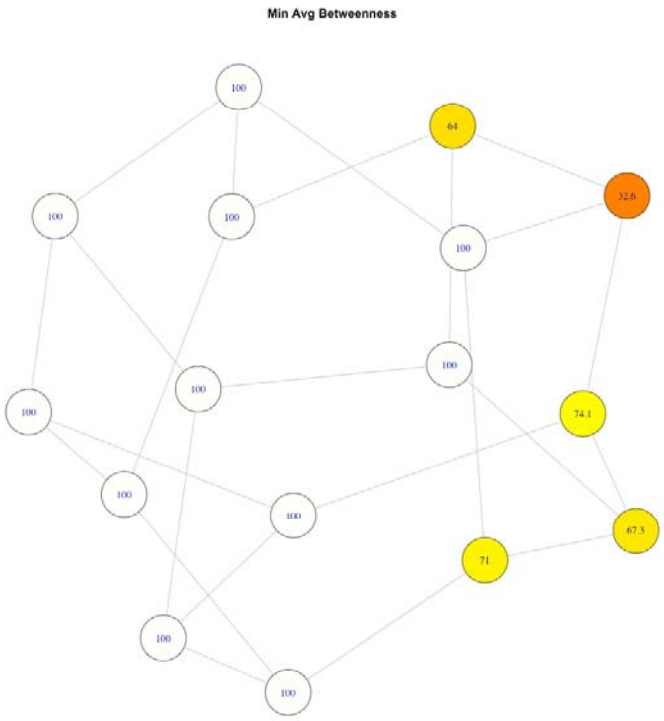

Diffusion of Best Solution

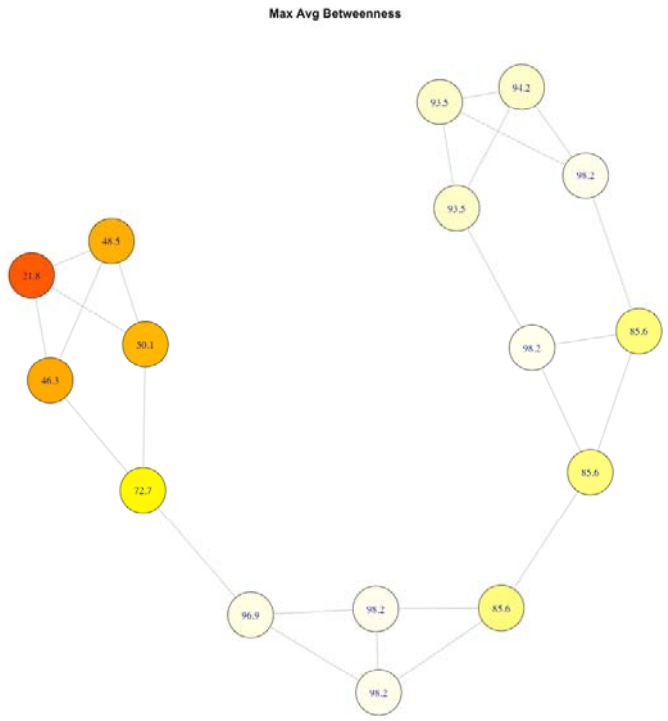

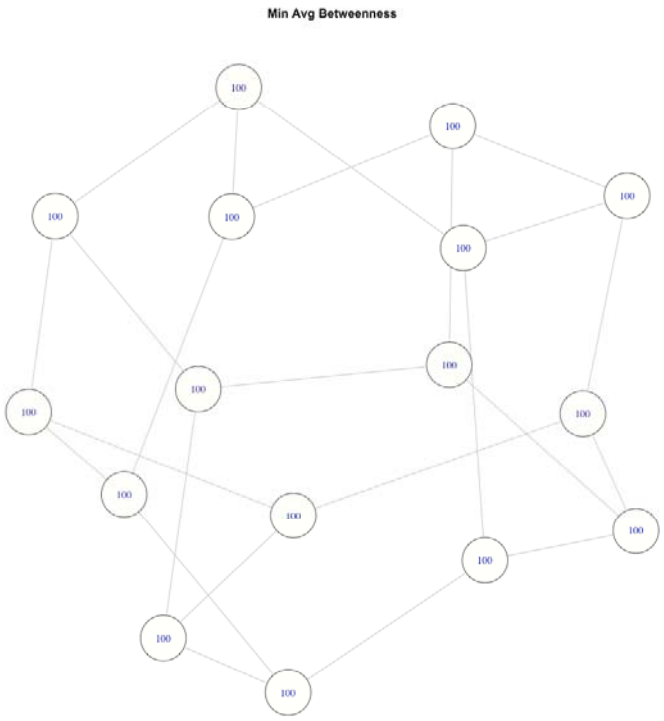

Diffusion of Best Solution

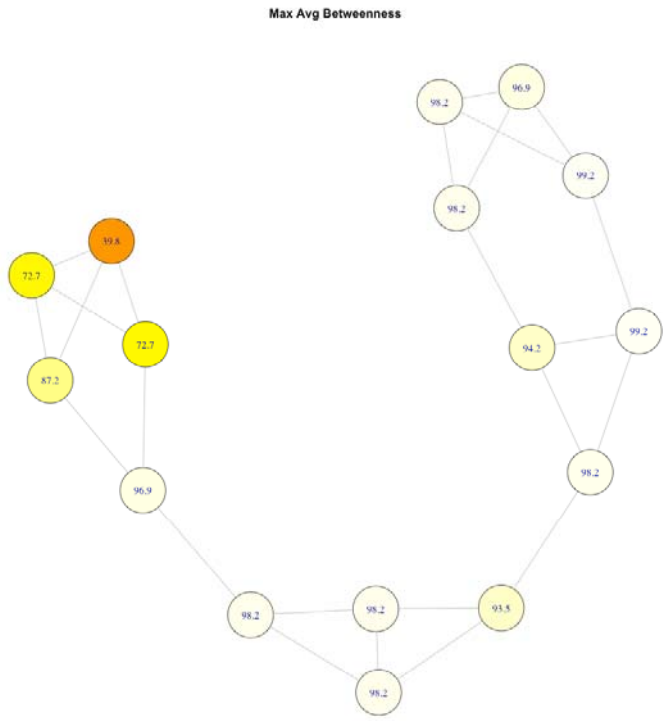

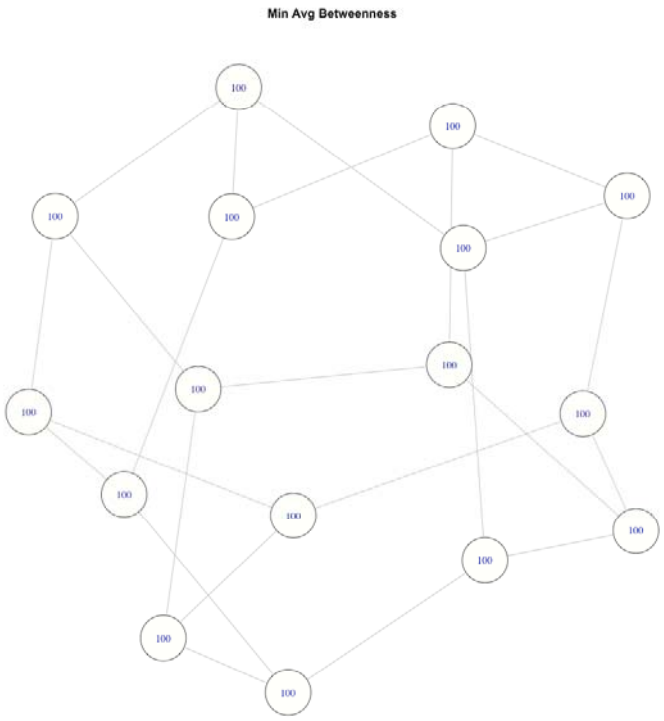

Diffusion of Best Solution

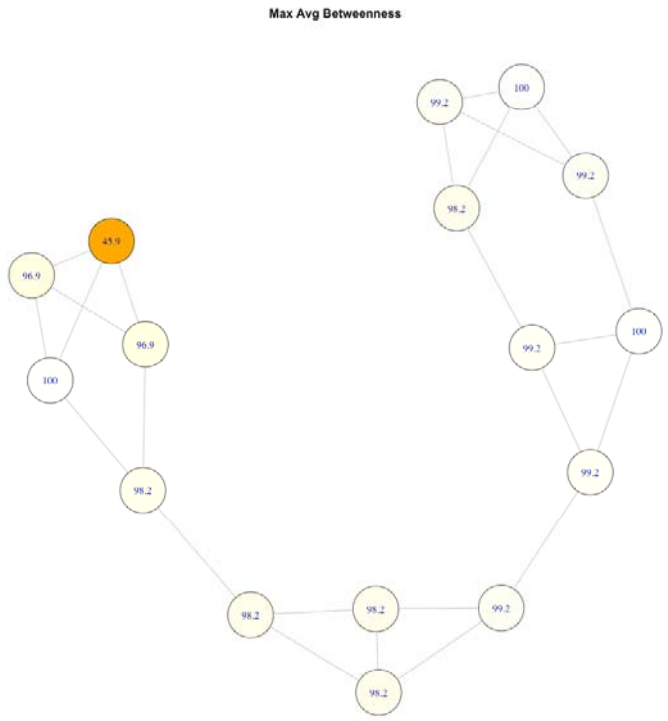

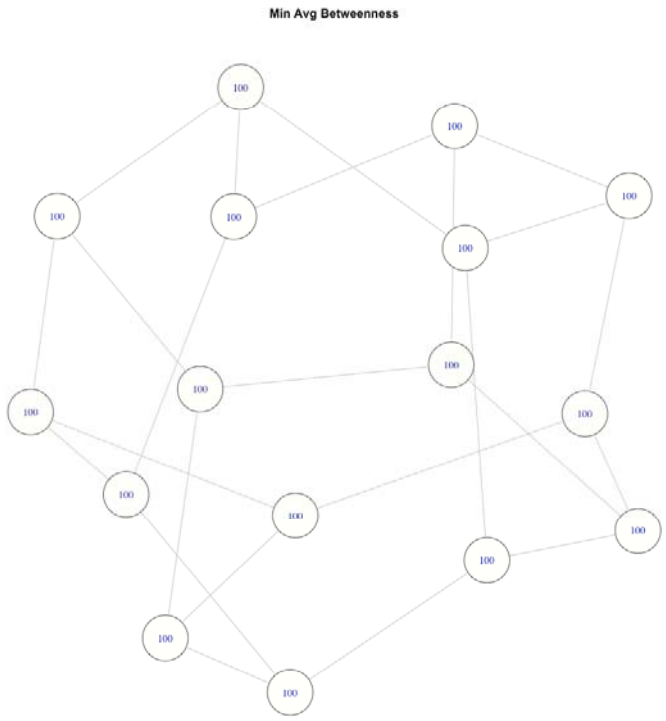

Individual and group performance are in tension

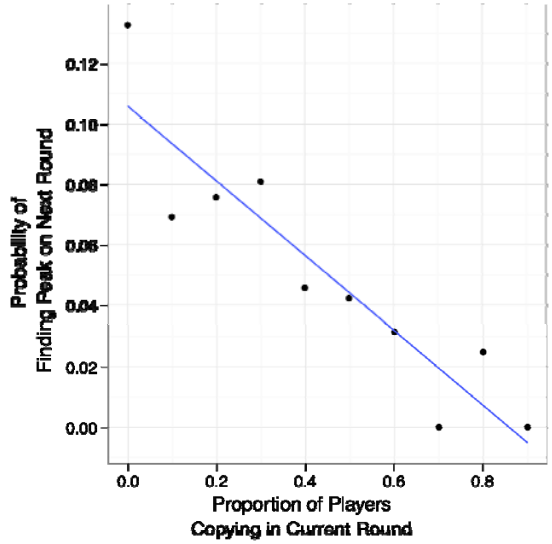

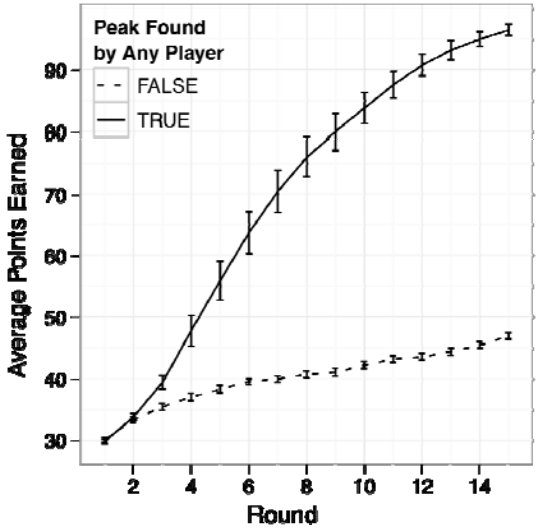

- More likely to find peak with more players exploring (= fewer players copying)
- When peak is found, large difference in points (nearly 2x income)

Individual and group performance are in tension

BUT:

- Per player, higher points on average with “copying” strategy
- → free-riding problem / social dilemma

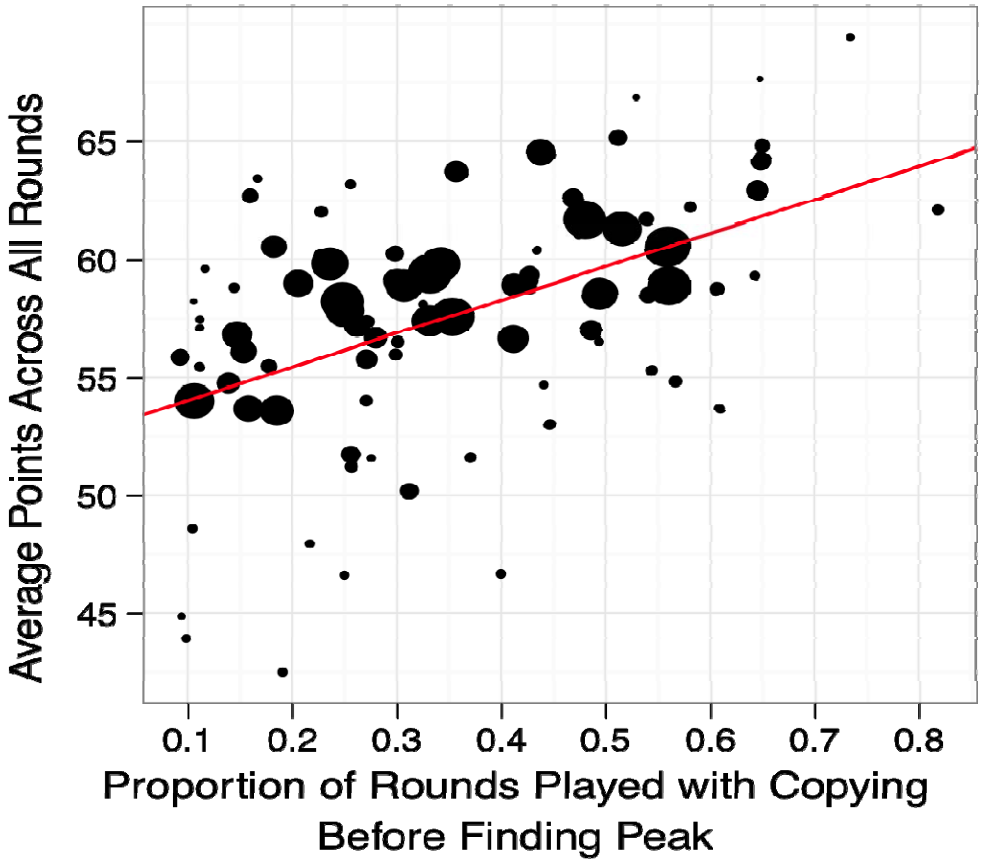

Individual and group performance are in tension

- Most successful node in structured networks ~ performance of median node in unstructured networks

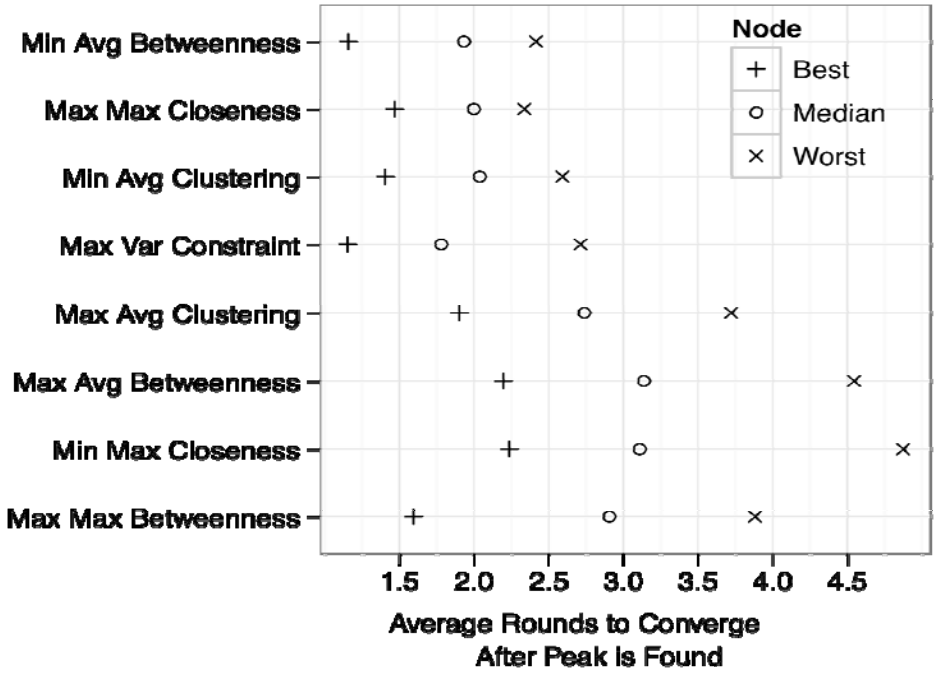

Individual Performance Is Combination of Individual Position and Collective Performance

Greatest Maximum Betweenness

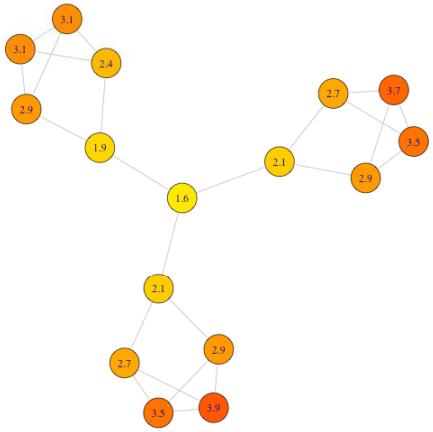

Greatest Average Betweenness

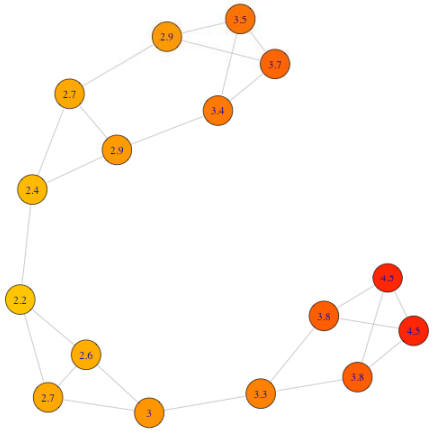

Greatest Maximum Closeness

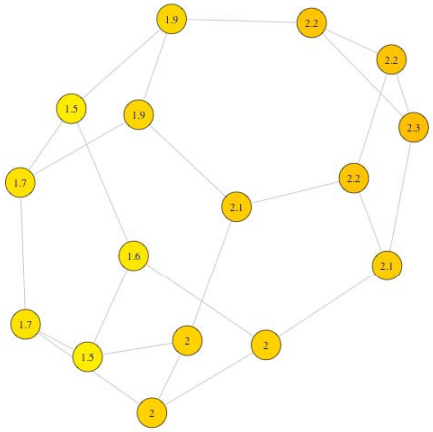

Summary of Results

- Network structure affects individual search strategy
  - Clustering increases copying (exploitation)
- Individual search strategy affects group success
  - More copying means lower probability of finding maximum
- Network structure affects group success directly
  - Networks with lower average path length spread information more quickly
- Individual and group performance are in tension
  - Individuals can improve own relative performance by free-riding on others' exploration
  - Best position in structured networks as good as average position in unstructured networks

What About “Real” Problem Solving?

- Our setup is artificial in several respects:
  - Real-world problems unlike to comprise just two dimensions
    - NK model may be better here (Kauffman, Levinthal, Lazer/Friedman)
  - Unclear how to interpret ruggedness of fitness landscape
    - NK model also has this problem
  - Networks do not resemble organizational networks
    - No hierarchy, division of labor etc.
  - Incentives are also unrealistic
    - Little strategic play, no competition, etc.
- One should therefore be cautious inferring much about managerial or strategic questions from our results
- Nevertheless, similar findings have emerged in other studies
  - March (1991) identified similar “dilemma” for individuals
  - Lazer and Friedman found that short path length -→ rapid convergence
- Also platform should generalize to more realistic scenarios

Web-Based Experimentation

- Our project adds to a small but growing body of web-based experiments
  - Salganik, Dodds, and Watts (2006)
  - Mason and Watts (2009)
  - Paolacci et al (2010), Horton et al. (2010)
  - Suri and Watts (2010)
- Major recent innovation has been use of standing panel to run synchronous experiments
- Allows for three major advances over physical labs
  - Possible to scale up to much larger networks
  - Speedup of hypothesis-testing loop
  - Selection of individuals based on past play

Experimental Macro-Sociology?

- "Can you put an army in a lab?" (Zelditch, 1969)
  - At the time, the answer was "No"
  - Led to emphasis on small-group research
- The web is removing this constraint
  - Synchronous play and sampling are also getting resolved
  - Also growing evidence that people "play" similarly on the web as they do in physical labs (Suri and Watts '10, Paolacci '10)
- In the near future, will increasingly see large scale, networked, lab-style experiments in which micro- and macro- variables can be manipulated and observed
- Still unclear how many experiments of potential interest could be conducted on the web
  - Most things that a real army does are still not online
  - Same true for many problems of interest to economic and organizational sociology
  - But this should be viewed as a challenge

Requisite Book Plug…

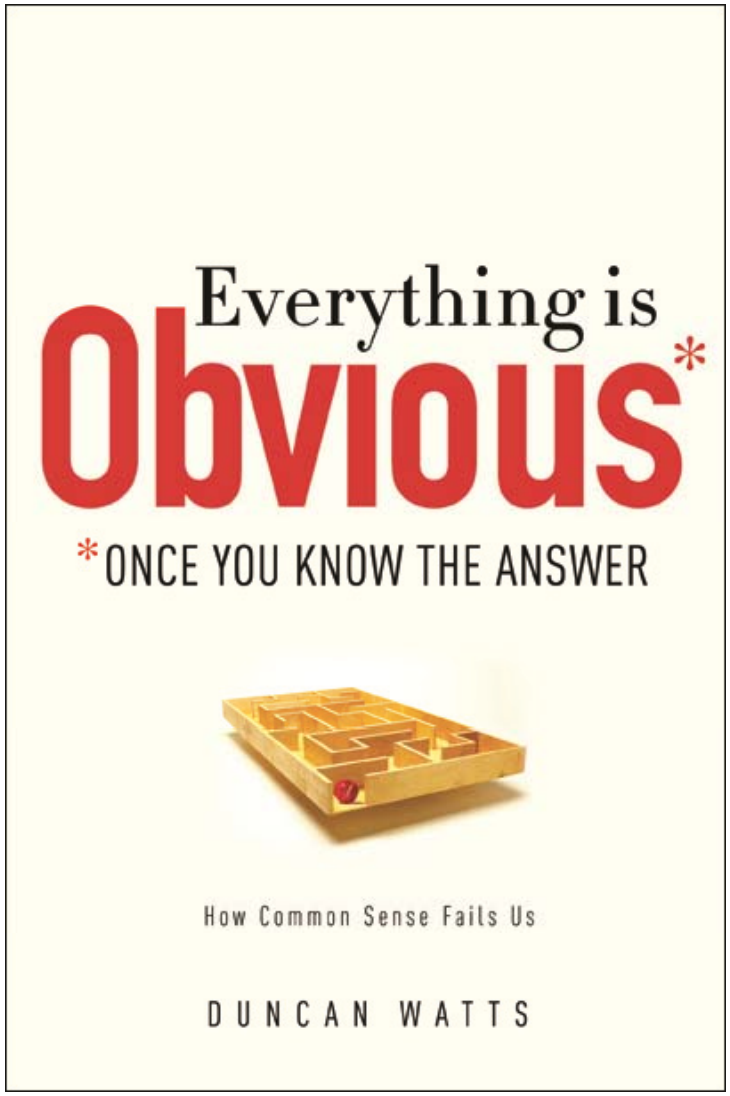

- Why are social phenomena so complex and unpredictable?
- Why do we still feel we can predict and control them?
- What could we do better?
  - In business, government, science

www.everythingisobvious.com

Thank you!

Questions?

Backup Slides

Agent Based Simulation and Comparison with Experiments

Agent Based Model, Based on Real Agents

- Extract individual playing strategies
- Build agent-based simulation where agents play like “real” players
- Explore problem space to discover new hypotheses
  - More complex landscapes
  - Different composition of individual strategies
  - Larger networks
- Return to experiments to test hypotheses

Exploration vs. Exploitation

Probability of exactly copying / guessing within 5 units from neighbor *given maximum has not yet been found*

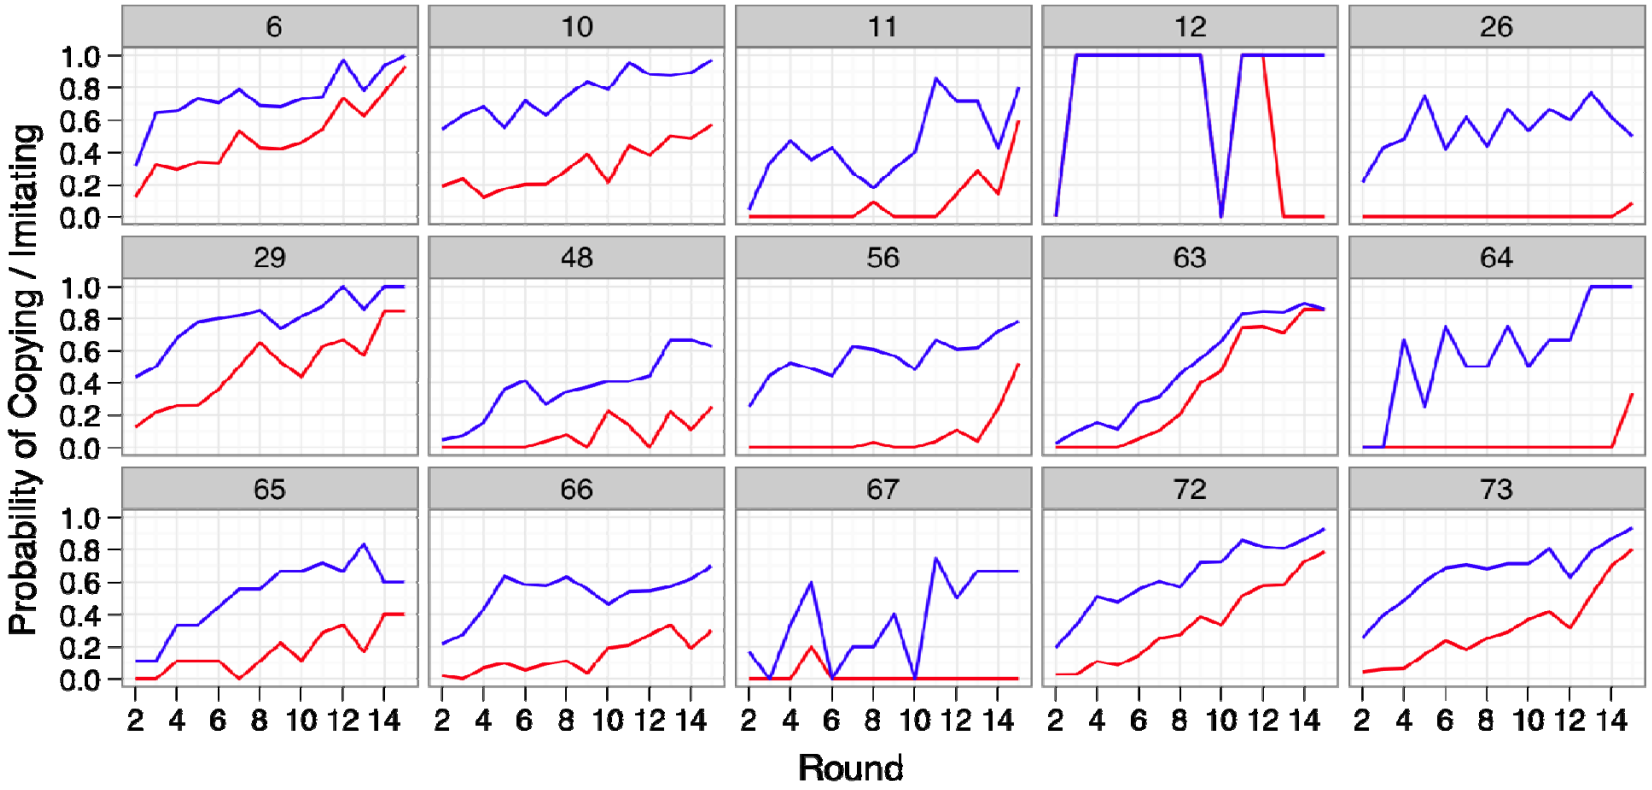

Simulation Details

- Fit linear model to users' probability of copying by round
- Obtain distribution of slopes & intercepts

- On each round:
  - If agent or neighbors have score = 100, copy
  - If agent or neighbors have 60 < score < 100, guess within 3 units of score
  - Else, copy highest score with probability based on intercept, slope & round or explore uniformly at random
- 100 simulated sessions (800 simulated games)

Finding the maximum

- 100 simulated sessions (800 simulated games)
- Maximum is found by at least one agent in 59% of games [63%]
- Maximum is found by all agents in 49% of games [56%]

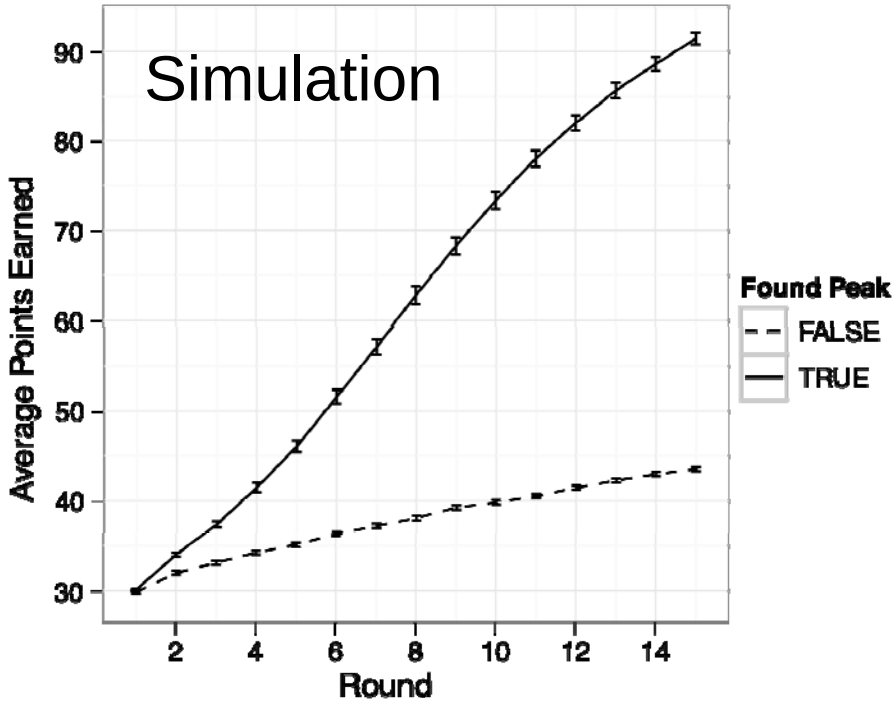

Human Players

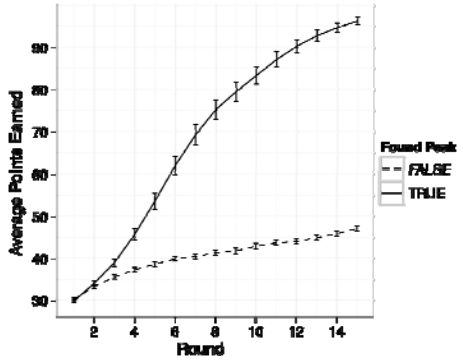

Frequency of Finding Maximum

Human Experiments

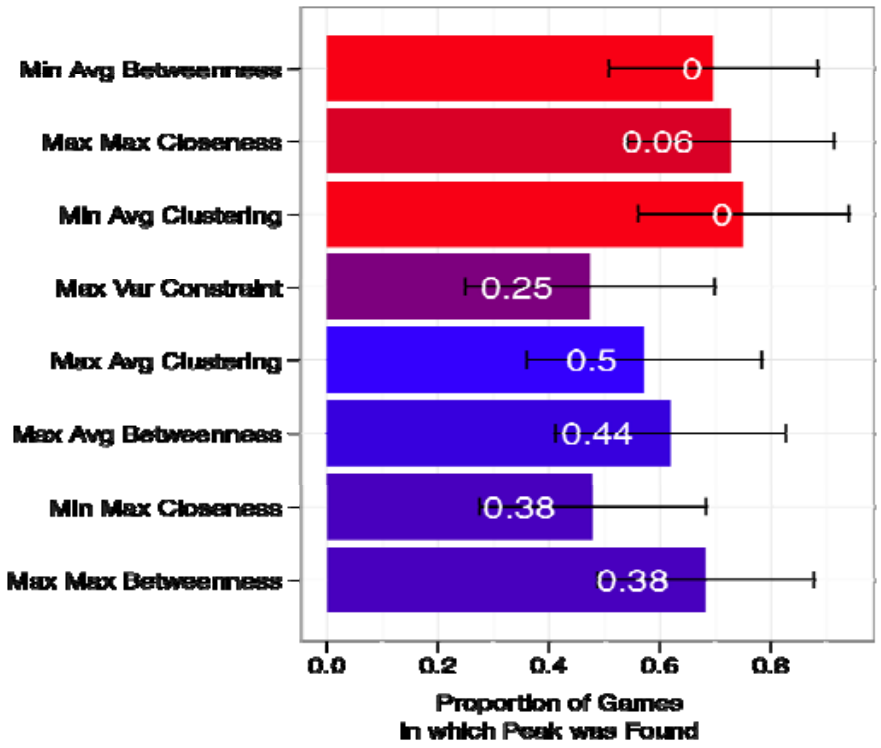

Simulations

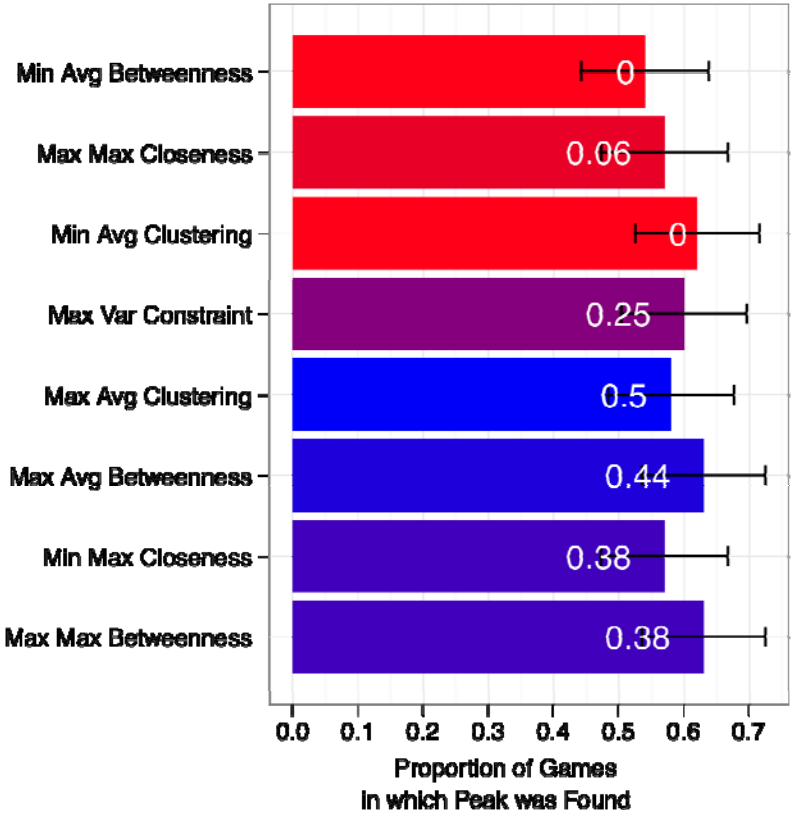

Networks Affect Convergence Time

Simulation

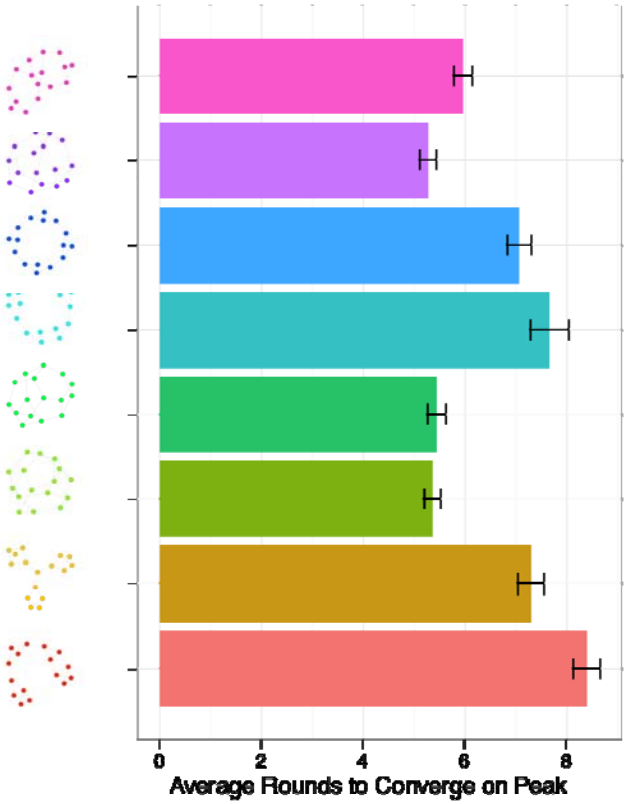

Human Players

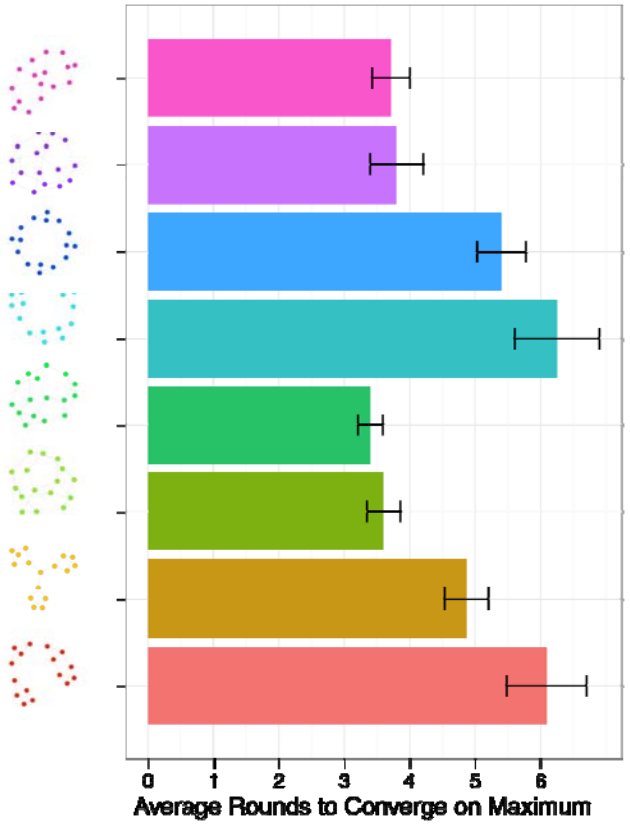

- Replicates findings from experimental work
- Suggests model of player behavior is reasonable

Networks Affect Convergence Time

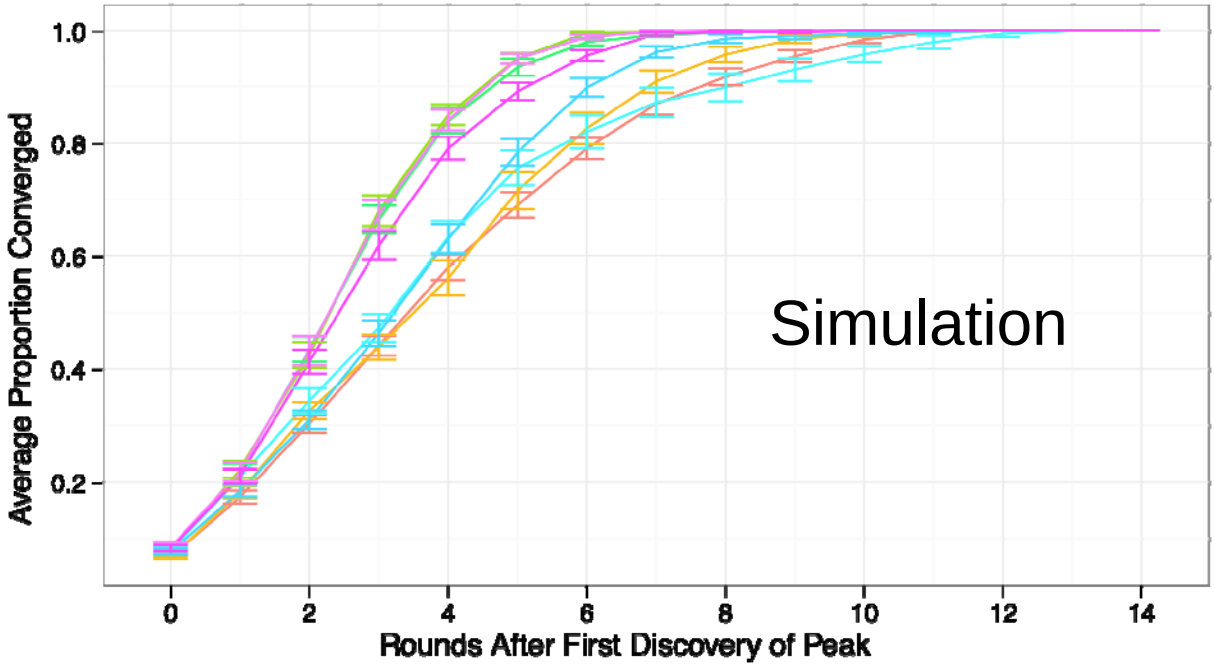

Human Players

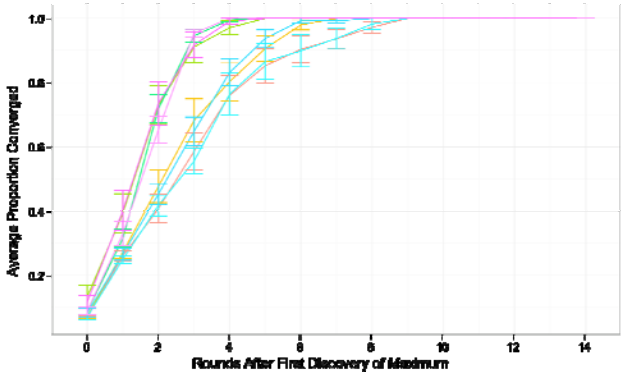

- Replicates findings from experimental work
- Suggests model of player behavior is reasonable

Individual Performance Is Combination of Individual Position and Collective Performance

Greatest Maximum Betweenness

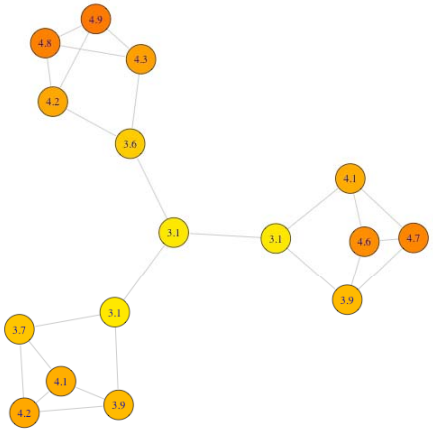

Greatest Average Betweenness

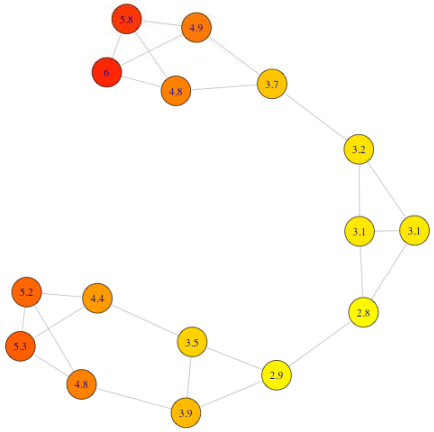

Greatest Maximum Closeness

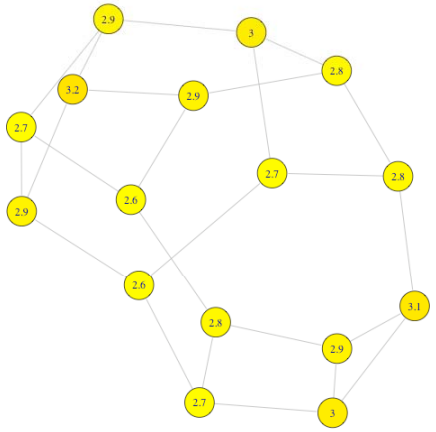

- Individuals in centralized networks perform well, relative to their peers
- All individuals in centralized networks perform poorly relative to individuals in decentralized networks
- Corroborates experimental results

Next Steps

- Explore problem space to discover new hypotheses
  - More complex payoff functions
  - Larger networks
  - Different composition of individual strategies
- Realistic model, but may be over-fit
  - Point threshold & imitation radius learned from known features of payoff functions
  - Copying / round depends on N rounds
- Return to experiments to test hypotheses