

Behavioral studies of networked human problem- solving

12-1-2012

Background & Motivation

- Social media has many different goals
 - Communication: friendship, product ratings, political news, ...
 - which often leads to decision-making and actions
 - **Coordination**: task delegation & negotiation, orchestrating joint problem-solving, ...
- Analysis of social media has many different goals
 - Identifying communities & modeling homophily
 - Predicting future behavior
 - Improved marketing
 - Summarization of a group
 - ...
- Important line of work: correlating observational data on social structure with **performance** on tasks

“You can observe a lot by just watching” – Y. Berra

Example

NETWORK EFFECTS


THE INFLUENCE OF STRUCTURAL SOCIAL CAPITAL ON OPEN SOURCE PROJECT SUCCESS

PARAM VIR SINGH • YONG TAN • VIJAY S. MOOKERJEE[†]



Data: collaboration networks from thousands of SourceForge projects.

Sample hypothesis: *The number of indirect ties of a project will be positively correlated with success.*



$$\text{Indirect Ties } FD_i = \sum_{j=2}^N d_{ij} w_{ij}$$

$$\begin{aligned} \ln DV_{it} = & \beta_0 + \beta_1 RT_{it} + \beta_2 TP_{it} + \beta_3 SEJD_{it} + \beta_4 SECOR_{it} + \beta_5 mcEDT_{it} + \beta_6 mcEIT_{it} \\ & + \beta_7 mcEDT_{it} \times mcEIT_{it} + \beta_8 mcEC_{it} + \beta_9 (mcEC_{it})^2 + \beta_{10} mcTD_{it} + \beta_{11} (mcTD_{it})^2 \\ & + \beta_{12} EC_{it} + \beta_{13} BC_{it} + \beta_{14} \ln TSIZE_{it} + \beta_{15} \ln PDV_{it} + \beta_{16} \ln BUGS_{it} + \beta_{17} \ln SUP_{it} \\ & + \beta_{17} \ln Page_{it} + \beta_{18} mcAGE_{it} + \beta_{19} (mcAGE_{it})^2 + \kappa TM + \psi FR_{it} + \chi EN_i + \sum_{m=1}^{ui-1} \delta_m UI_{im} \\ & + \sum_{k=1}^{ia-1} \lambda_k IA_{ik} + \sum_{h=1}^{ty-1} \eta_h TY_{ih} + \sum_{g=1}^{os-1} \gamma_g OS_{ig} + u_i + \varepsilon_{it}, \end{aligned}$$

Productivity as function of network-based measurements

Background & Motivation

- Social media has many different goals
 - Communication: friendship, product ratings, political news, ...
 - which often leads to decision-making and actions
 - Coordination: task delegation & negotiation, orchestrating joint problem-solving, ...
- Analysis of social media has many different goals
 - Identifying communities & modeling homophily
 - Predicting future behavior
 - Improved marketing
 - Summarization of a group
 - ...
- **Analysis of observational data is *limited***
 - Example: for viral marketing, it's hard to predict the effect of an *intervention* in the marketing scheme from available data

“You can observe a lot by just watching” – Y. Berra

Overview & Motivation

- Social media has many different goals
 - Communication: friendship, product ratings, political news, ...
 - which often leads to **decision-making** and **actions**
 - **Coordination**: task delegation & negotiation, orchestrating joint problem-solving, ...
- Analyze What do we really know about the underlying processes that are being coordinated by the communication?
 - Identify
 - Predict
 - Improve
 - How do people solve problems collectively?
 - Structure
 - What are the effects of networks on *problem-solving*?
 - ...
 - Does real human behavior approximate game-theoretic models?
- Analysis of **observational data is limited**
 - Example: for viral marketing, it's hard to predict the effect of an *intervention* in the marketing scheme from available data

Behavioral Experiments in Network Science

slides pilfered from:

**Networked Life
CIS 112
Spring 2010
Prof. Michael Kearns**



Background and Motivation

- Network Science: Structure, Dynamics and Behavior
 - sociology, economics, computer science, biology...
 - network universals and generative models
 - empirical studies: network is **given**, hard to explore alternatives
- Navigation and the Six Degrees
 - Travers & Milgram → Watts, Kleinberg
 - distributed all-pairs shortest paths
 - what about **other** problems?
- Behavioral Economics and Game Theory
 - human rationality in the lab
 - typically subjects in pairs

Experimental Agenda



- Human-subject experiments at the intersection of CS, economics, sociology, “network science”
- Subjects simultaneously participate in groups of ~ 36 people
- Subjects sit at networked workstations
- Each subject controls some simple property of a single vertex in some underlying network
- Subjects have only *local* views of the activity: state of their own and neighboring vertices
- Subjects have (real) financial incentive to solve their “piece” of a collective (global) problem
- Simple example: graph coloring
 - choose a color for your vertex from fixed set
 - paid iff your color differs from all neighbors when time expires
 - max welfare solutions = proper colorings
- Across many experiments, have deliberately varied *network structure* and *problem/game*
 - networks: inspired by models from network science (small worlds, preferential attachment, etc.)
 - problems: chosen for diversity (cooperative vs. competitive) and (centralized) computational difficulty
- Goals:
 - structure/problem → performance/behavior
 - individual & collective modeling → prediction
 - computational and equilibrium theories
 - exploring counterfactuals

Experiments to Date

- Graph Coloring (Jan 2006; Feb 2007)
 - player controls: color of vertex; number of choices = chromatic number
 - payoffs: \$2 if different color from all neighbors, else 0
 - max welfare states: optimal colorings
 - centralized computation: hard even if approximations are allowed
- Consensus (Feb 2007)
 - player controls: color of vertex from 9 choices
 - payoffs: \$2 if same color as all neighbors, else 0
 - max welfare states: global consensus of color
 - centralized computation: trivial
- Independent Set (Mar 2007)
 - player controls: decision to be a “King” or a “Pawn”; variant with King side payments allowed
 - payoffs: \$1/minute for Solo King; \$0.50/minute for Pawn; 0 for Conflicted King; continuous accumulation
 - max welfare states: maximum independent sets
 - centralized computation: hard even if approximations are allowed
- Exchange Economy (Apr 2007)
 - player controls: limit orders offering to exchange goods
 - payoffs: proportional to the amount of the other good obtained
 - max welfare states: market clearing equilibrium
 - centralized computation: at the limit of tractability (LP used as a subroutine)
- Biased Voting (“Democratic Primary Game”; May 2008)
 - player controls: choice of one of two colors
 - payoffs: only under global agreement; different players prefer different colors
 - max welfare states: all red and all blue
 - centralized computation: trivial
- Networked Bargaining (May 2009)
 - player controls: offers on each edge to split a cash amount; may have hidden deal limits and “transaction costs”
 - payoffs: on each edge, a bargaining game --- payoffs only if agreement
 - max welfare states: all deals/edges closed
 - centralized computation: nontrivial, possibly difficult

Graph Coloring

(Behavioral) Graph Coloring



- Undirected network; imagine a person “playing” each vertex
 - Finite number of colors; each person picks a color
 - Goal: no pair connected by an edge have the same color
 - Computationally well-understood
 - no efficient centralized algorithm
 - strong evidence for computational hardness
 - even extremely weak approximation algorithms
 - ...Yet simple and locally verifiable
- Arguably approximates some real-world social problems that are solved in a distributed way:*
- Picking a room for an LTI course
 - Picking a unique ringtone
 - ...


Experimental Design Variables

- Network Structure
 - six different topologies
 - inspired by network formation models
- Information View
 - three different views
- Incentive Scheme
 - two different mechanisms
- Design space: $6 \times 3 \times 2 = 36$ combinations
- Ran all 36 of them (+2)


Research Questions

- Can large groups of people solve these problems at all?
- What role does network structure play?
 - information view, incentives?
- What behavioral heuristics do individuals adopt?
- Can we do collective modeling and prediction?
 - some interesting machine learning challenges


Choices of Network Structure




Simple Cycle




5-Chord Cycle



20-Chord Cycle




Leader Cycle



Preferential Attachment,

$v = 2$



Preferential Attachment,

$v = 3$

Kearns' idea?

Small
Worlds
Family

Choices of Information Views



Your goal: Choose a color that none of your neighbors have. When everyone in the system achieves this goal, the problem will be solved.

Your current color is **green**.


Click below if you wish to choose a different color:

blue

green

red

Low-information view



1 conflict in your immediate neighborhood.

A thick line indicates a conflict that must be resolved.

A thin line is shown when color choices do not conflict.

Overall progress toward a solution:



Your goal: Choose a color that none of your neighbors have. When everyone in the system achieves this goal, the problem will be solved.


Your current color is **green**.

Click below if you wish to choose a different color:

blue

green

red



1 conflict in your immediate neighborhood.

A thick line indicates a conflict that must be resolved.

A thin line is shown when color choices do not conflict.

Overall progress toward a solution:



Your goal: Choose a color that none of your neighbors have. When everyone in the system achieves this goal, the problem will be solved.


Your current color is **green**.

Click below if you wish to choose a different color:

blue

green

red



1 conflict in your immediate neighborhood.

A thick line indicates a conflict that must be resolved.

A thin line is shown when color choices do not conflict.

Overall progress toward a solution:



Choices of Incentive Schemes

Didn't make much difference so I won't discuss

....


- **Collective incentives:**
 - all 38 participants paid if and only if **entire graph** is properly colored
 - payment: \$5 per person for each properly colored graph
 - a “team” mechanism
- **Individual incentives**
 - each participant paid if **they** have no conflicts at the end of an experiment
 - payment: \$5 per person per graph
 - a “selfish” mechanism
- Minimum payout per subject per session: \$0
- Maximum: $19 \times 5 = \$95$

The Experiments: Some Details

- **5 minute (300 second) time limit** for each experiment
- Population demographics: Penn CSE 112 students
- Handout and intro lecture to establish understanding
- Intro and exit **surveys**
- **No communication allowed except through system**
- Experiments performed Jan 24 & 25, 2006
 - Spring 2005: CSE 112 paper & pencil face-to-face experiments
 - Sep 2005: system launch, first controlled experiments
- Jan 24 session: collective incentives; Jan 25 session: individual incentives
- Randomized order of 18 experiments within each session
- First experiment repeated as last to give 19 total per session



The Results: Overview




- 31 of 38 experiments solved
- mean completion time of solved = 82s
 - median = 44s

	Collective Incentive			Individual Incentive		
	Low Info.	Med. Info.	High Info.	Low Info.	Med. Info.	High Info.
Simple Cycle	70	U	26	276	168	25
5-Chord Cycle	134	276, 61	24	221	108	24
20-Chord Cycle	159	39	24	22	39	111
Leader Cycle	24	50	30	74	28	44, 36
Pref. Att., $\nu = 2$	U	U	U	35	83	U
Pref. Att., $\nu = 3$	41	33	U	64	191	U

- 31 of 38 experiments solved
- mean completion time of solved = 82s
- median = 44s

Avg. Experiment Duration for the 3 Info. Views




Effects of Network Structure


	Colors required	Min. degree	Max. degree	Avg. degree	S.D.	Avg. distance	Avg. duration & fraction solved		Distributed heuristic
Simple cycle	2	2	2	2	0	9.76	144.17	5/6	378
5-chord cycle	2	2	4	2.26	0.60	5.63	121.14	7/7	687
20-chord cycle	2	2	7	3.05	1.01	3.34	65.67	6/6	8265
Leader cycle	2	3	19	3.84	3.62	2.31	40.86	7/7	8797
Pref. att., newlinks=2	3	2	13	3.84	2.44	2.63	219.67	2/6	1744
Pref. att., newlinks=3	4	3	22	5.68	4.22	2.08	154.83	4/6	4703

- smaller diameter → better performance
- preferential attachment much harder than cycle-based
- distributed heuristic gives reverse ordering


Heuristic: pick a node with a conflict, and either (a) solve it if there is an unused color or (b) swap the color otherwise. Repeat until solved.




Simple Cycle




5-Chord Cycle




20-Chord Cycle



Leader Cycle



Preferential Attachment,
 $v = 2$




Preferential Attachment,
 $v = 3$


Small
Worlds
Family

Effects of Information View


More information makes cycles easier, preferential attachment graphs harder.




Cycle Convergence Plot: All High Info. Cycles



Cycle Convergence Plot: All Low Info. Cycles



Color Changes for each Node



Effects of Incentive Scheme

	Avg. Experiment Duration (secs)	Avg. Conflict Duration (secs)	Number of Perturbations
Collective Incentive	130.11	5.39	130
Individual Incentive	113.11	4.87	51

Table 1: Changing incentive schemes had a mild effect on average experiment duration and average conflict duration.

Towards Behavioral Modeling

Algorithmic Introspection

Prioritize color matches to high degree nodes. That is, I tried to arrange it so that the high degree nodes had to change colors the least often. So if I was connected to a very high degree node I would always change to avoid a conflict with it. (Sep 2005 comments) If I had a higher degree than the others I was connected to I would usually stay put and avoid changing colors. [many similar comments]

Strategies in the local view: I would wait a little before changing my color to be sure that the nodes in my neighborhood were certain to stay with their color. I would sometimes toggle my colors impatiently (to get the attention of other nodes) if we were stuck in an unresolved graph and no one was changing their color.

Strategies in the global view: I would look outside my local area to find spots of conflict that were affecting the choices around me. I would be more patient in choices because I could see what was going on beyond the neighborhood. I tried to solve my color before my neighbors did.


I tried to turn myself the color that would have the least conflict with my neighbors (if the choices were green, blue, red and my neighbors were 2 red, 3 green, 1 blue I would turn blue). I also tried to get people to change colors by "signaling" that I was in conflict by changing back and forth.

If we seemed to have reached a period of stasis in our progress, I would change color and create conflicts in my area in an attempt to find new solutions to the problem.

When I had two or three neighbors all of whom had the same color, I would go back and forth between the two unused colors in order to inform my neighbors that they could use either one if they had to.

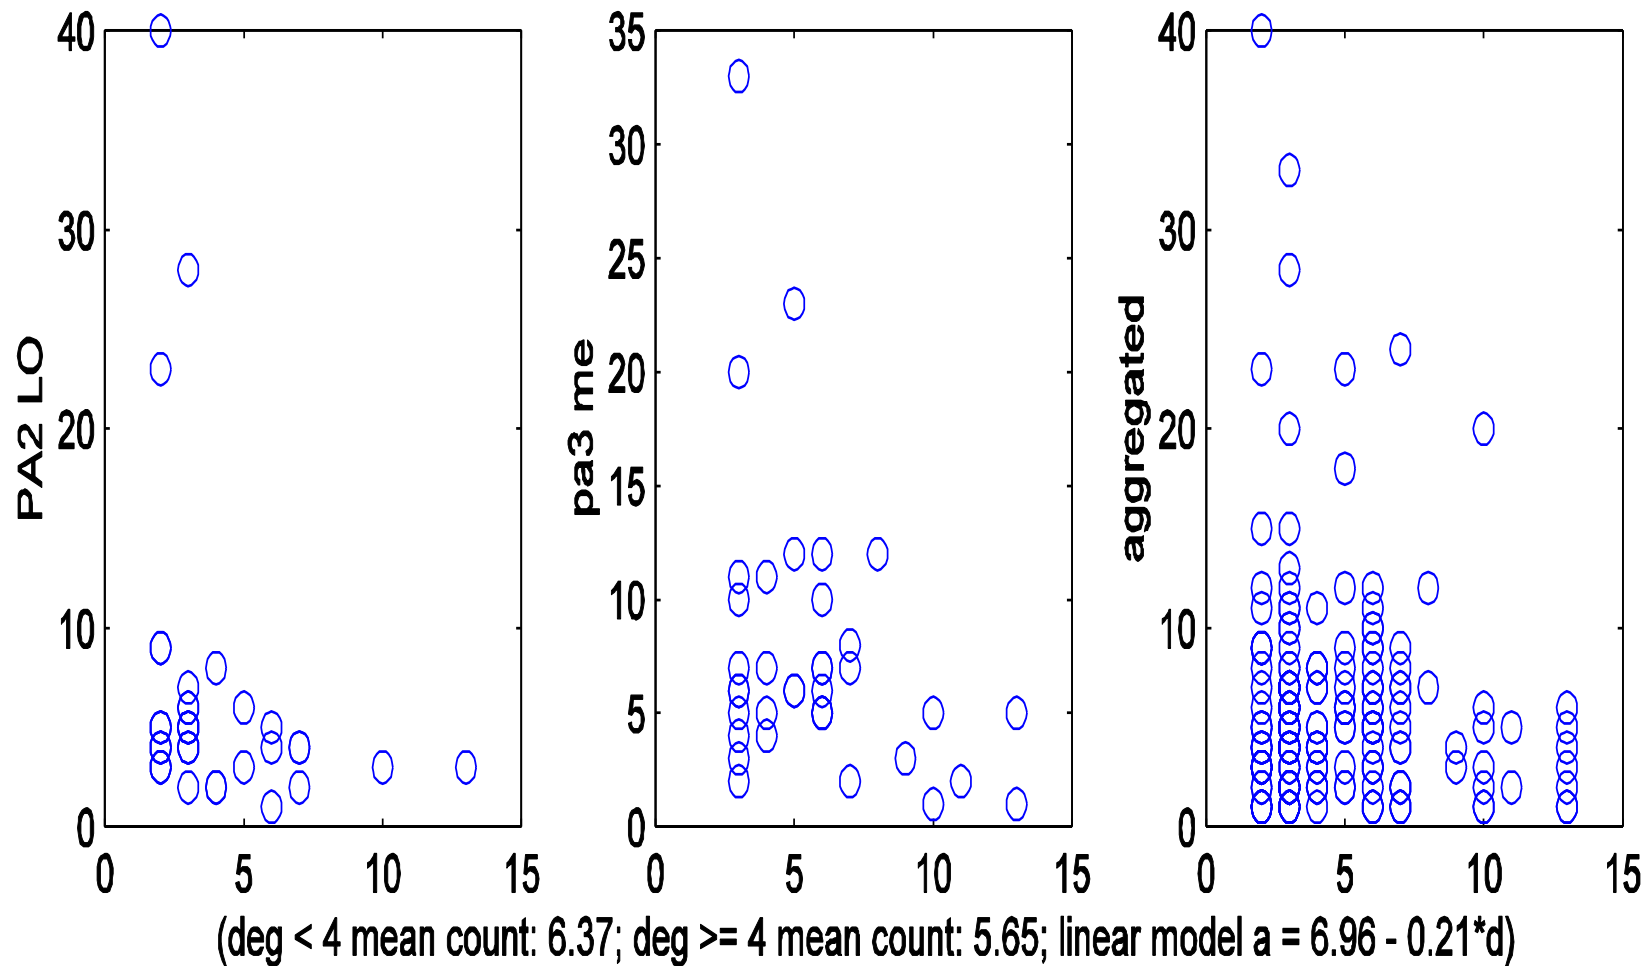
Low-degree nodes spend more time swapping labels.

Degree (x) Vs Action Count (y), Unified Preferential Attachment Experiments



(Sep 2005 data)


Low-degree nodes spend more time swapping labels.



(Sep 2005 data)

signaling behaviors

Color Changes for each Node in exppa38new2lowglo



signaling behaviors

Experiments to Date (2010)


- Graph Coloring (Jan 2006; Feb 2007)
 - player controls: color of vertex; number of choices = chromatic number
 - payoffs: \$2 if different color from all neighbors, else 0
 - max welfare states: optimal colorings
 - centralized computation: hard even if approximations are allowed
- Consensus (Feb 2007)
 - player controls: color of vertex from 9 choices
 - payoffs: \$2 if same color as all neighbors, else 0
 - max welfare states: global consensus of color
 - centralized computation: trivial
- Independent Set (Mar 2007)
 - player controls: decision to be a “King” or a “Pawn”; variant with King side payments allowed
 - payoffs: \$1/minute for Solo King; \$0.50/minute for Pawn; 0 for Conflicted King; continuous accumulation
 - max welfare states: maximum independent sets
 - centralized computation: hard even if approximations are allowed
- Exchange Economy (Apr 2007)
 - player controls: limit orders offering to exchange goods
 - payoffs: proportional to the amount of the other good obtained
 - max welfare states: market clearing equilibrium
 - centralized computation: at the limit of tractability (LP used as a subroutine)
- Biased Voting (“Democratic Primary Game”; May 2008)
 - player controls: choice of one of two colors
 - payoffs: only under global agreement; different players prefer different colors
 - max welfare states: all red and all blue
 - centralized computation: trivial
- Networked Bargaining (May 2009)
 - player controls: offers on each edge to split a cash amount; may have hidden deal limits and “transaction costs”
 - payoffs: on each edge, a bargaining game --- payoffs only if agreement
 - max welfare states: all deals/edges closed
 - centralized computation: nontrivial, possibly difficult



Summary of experiments to date. ER stands for Erdős-Renyi, PA for preferential attachment.

Task Description	Networks	Incentives/Mechanism	Sample Findings
graph coloring ¹⁷	cycle+chords; PA	differ with neighbors	chords help; importance of information view
coloring and consensus ¹⁰	clique chain w/rewiring	differ/agree with neighbors	opposite structure/task effects
networked trade ¹³	ER; PA; structured; all bipartite	limit orders for trades for opposing good	comparison to equilibrium theory; networked inequality aversion
networked bargaining ³ **	assorted	Nash bargain on each edge	behavioral price of obstinacy
independent set ¹⁵	assorted	kings and pawns with side payments	side payments help; conflict and fairness
biased voting ¹⁴	ER and PA between types; minority power	consensus with competing individual preferences	well-connected minority rules
network formation ¹⁶ **	endogenous to the game	biased voting minus edge expenditures	poor collective performance

as of 2012



- Threshold: switch if 40% of neighbors switched

Biased Voting in Networks

“Democratic Primary Game”

- Cosmetically similar to consensus, with a crucial strategic difference
- Deliberately introduce a tension between:
 - individual preferences
 - desire for collective unity
- Only two color choices; challenge comes from competing incentives
- If everyone converges to same color, everyone gets some payoff
- But different players receive different amounts
 - each player has payoffs for their preferred and non-preferred color
 - e.g. \$1.50 red/\$0.50 blue vs. \$0.50 red/\$1.50 blue
 - can have symmetric and asymmetric payoffs
- High-level experimental design:
 - choice of network structures
 - arrangement of types (red/blue prefs) & strengths of incentives
 - most interesting to coordinate network structure and types

Zak Xavier

game status: Voter Game in progress

If unanimity is reached, your payoff will be

The diagram shows a central node labeled 'you' (blue) connected to eight peripheral nodes. The peripheral nodes are arranged in a circle and contain numerical values. The connections are as follows:



- Blue nodes (+0): Top, Top-Left, Left, Bottom-Left, Bottom, Bottom-Right, Right, Top-Right.
- Red nodes (+6, +4, +5, +1): Bottom-Left, Bottom, Bottom-Right, Top-Right.

Connections (Red lines = Positive, Black lines = Negative):



- Central 'you' node is connected to all peripheral nodes.
- Peripheral nodes are connected to their immediate neighbors in the circle.
- Additional connections: Top-Left (+0) to Top-Right (+1); Left (+0) to Bottom-Left (+6); Bottom-Left (+6) to Bottom (+4); Bottom (+4) to Bottom-Right (+5); Bottom-Right (+5) to Top-Right (+1); Top-Right (+1) to Top (+0).

your color: blue red



coER_0.5



coER_2




coER_1





•Cohesion: what is the effect of local networks that contain *aligned* incentives vs *competing* incentives?




•Cohesion: Erdos-Renyi



coPA_0.5




coPA_1




•Cohesion: Preferential Attachment


- Minority power:
what happens
when a small
number of highly-
connected nodes
have incentives
different from the
majority?






power27









power30



- Minority Power: Preferential Attachment




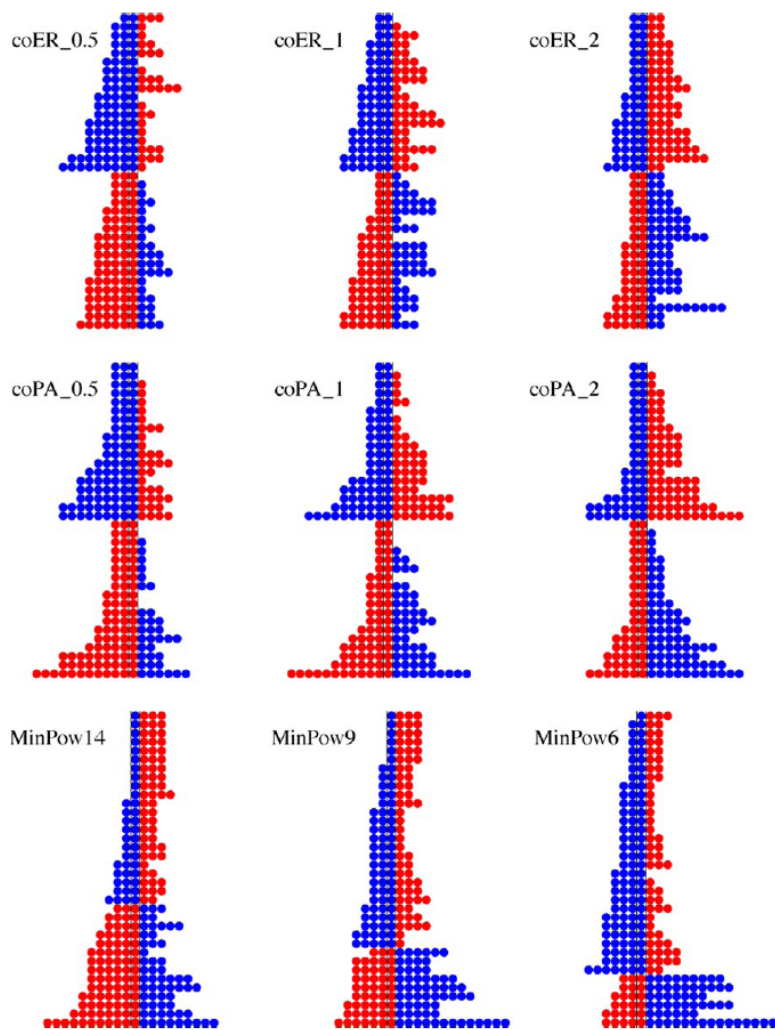
- Row = node x
- Center dot = incentive for x
- Left, right dots = incentive for neighbors of x




Summary of Findings

- 55/81 experiments reached global consensus in 1 minute allowed
 - mean of successful ~ 44s
- Effects of network structure:
 - Cohesion harder than Minority Power: 31/54 Cohesion, 24/27 Minority Power
 - all 24 successful Minority Powers converge to minority preference!
 - Cohesion P.A. (20/27) easier than Cohesion E-R
 - overall, P.A. easier than E-R (contrast w/coloring)
 - within Cohesion, increased inter-group communication helps
 - some notable exceptions...
- Effects of incentives:
 - asymmetric beats weak symmetric beats strong symmetric
 - the value of “extremists”
- Obviousness vs. Minimalism






Behavior correlates well with a simple game-theoretical model (Kearns & Wortman, COLT 2008)



Effects of “Personality”




36 subjects play 81 games each

Experiment:

- look at the 36 average wealth's (averaged over games)
- compare to the average for a random *permutation* of the subjects ids
 - i.e. average $W(i_{p1})$ in game 1, $W(i_{p2})$ in game 2, ... $W(i_{p81})$ in game 81
 - average this over many permutations

So: distribution of wealth in subjects seems pretty much as expected randomly

Effects of “Personality”






- Stubbornness = fraction of time you play preferred color when in minority
- Stubbornness correlates *positively* with wealth
- Switching colors early correlates *negatively* with wealth
- Players were not very stubborn:
 - Only individual three players were stupidly stubborn (defying neighbors as time expired)
 - One did it 3x (but also complied 38/55 times)

Survey results



- Strategies/observations:
 - 27 followed high-degree nodes, or were more stubborn when they had high degree
 - 27 tried signalling (or noticed signals)...many thought it was annoying
 - 24 players: start off with preferred colors
 - 21 noticed or suspected other people being irrationally stubborn
 - 3 mentioned being stubborn themselves
 - 7 varied strategy depending on incentives

Art by Consensus

(small p)




Consensus: inverse of graph-coloring: payoff iff all nodes have a *single* color.




People are different

Figure 3. Average time to global solution for coloring and consensus experiments (solid lines) as a function of edge rewiring in a clique-chain network, and simulation times (dashed lines) on the same networks for distributed heuristics. The parametric structure has the opposite effect on the two problems.






People are
different

Figure 5. Human performance vs. greedy algorithm in networked bargaining, demonstrating the effects of subject obstinacy. Where occlusions occur, blue dots are slightly enlarged for visual clarity. The length of the vertical lines measure the significant effects of subject obstinacy on payoffs.





Summary...

- Artificial small-scale experiments let you examine and *manipulate* things you can't see in large-scale data
- Causality vs correlation



Summary...

- Sometimes good visualizations are more powerful than quantitative measurements



Summary...

- Social data is richer than you think
 - even on highly restricted problems
 - and highly restricted views of real problems
 - e.g. “signals”

