The Hare and the Tortoise:
The Network Structure of Exploration and Exploitation

David Lazer david_lazer@harvard.edu

Allan Friedman, allan_friedman@ksg.harvard.edu

John F. Kennedy School of Government, Harvard University

NCDG National Center for Digital Government

Program on Networked Governance

Whether as team members brainstorming, or policy experts sharing best-practices, problem solvers communicate and share ideas. This poster examines the how the structure of these communication networks can affect outcomes. We present simulation results demonstrating that more efficient communication networks can actually lower performance in the long run. This effect is manifest in complex (e.g. rugged) problem spaces, where extensive searches are expensive for the individual. When actors can communicate easily or quickly, average performance improves rapidly improves initially, but harder-to-find optimal solutions are less likely to be discovered.

Acknowledgements: This project is being supported by NSF Grant # 0131923.

The Model: Networked Actors in Complex spaces

- The world is represented as a problem space with peaks of good solutions separated by valleys of inferior solutions. A more rugged space has more peaks and valleys. We model a multi-dimensional rugged problem space using Stuart Kauffman's NK model for epistatic evolution (Kauffman 1991).

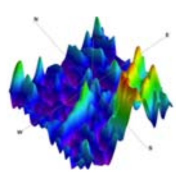

A rugged space with peaks and valleys. If an agent were on a peak, they would be unable to determine if it were the best possible solution.

- Actors are extremely myopic, and cannot see the best solution, but can compare the merits of two points. Each actor wants to find the highest point possible. They follow they same decision rule:

***If** an agent can see another agent at a higher point, copy the best agent visible.*	⇨ Exploitation
***Otherwise**, look at impact of randomly changing one dimension. Move only if it is an improvement*	⇨ Exploration

- The variable of interest is the network structure that determines the number and distribution of agents that a single agent can see and the velocity and synchronicity of information flow.

Network Archetypes

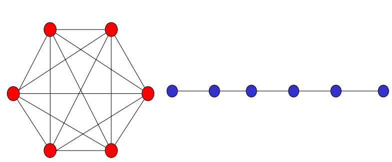

Fully Connected	Linear
Number of neighbors = population	Number of neighbors = 2
Avg. distance between nodes = 1	Avg. distance between nodes = population / 2

Results: Linear vs. full networks

The fully connected network improves quickly, but the linear network outperforms it in the long run.

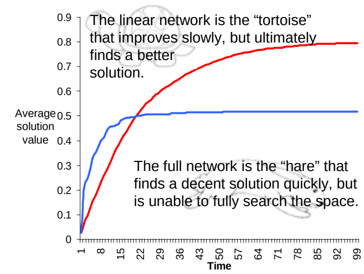

Having more opportunities to share solutions results in a *lower* outcome. This is because less information sharing preserves heterogeneity in the population, allowing actors to search more of the problem space for a better solution.

Heterogeneity of actors in the solution space

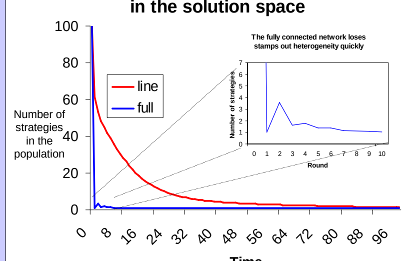

Maximizing the distance between nodes appears to maximize long-term solutions.

Random Networks (Erdos-Renyi)

- We found a curvilinear relationship between density and final score. Moderately low density measures (p=.04-6) outperformed very low density graphs and higher density graphs.
- Graphs with very low density often contain isolated cliques that reduce average score. If all nodes are not connected, then they cannot share the best solution found, lowering the optimal score.
- Eliminating all multi-clique networks confirms a negative (monotonic) relationship between network density and average long-term solution.
- Since average path length is negatively correlated with density, this supports the hypothesis that higher scores are a function of higher distance between nodes.

Velocity and synchronicity

Velocity refers to how fast a message can spread from actor to actor.

•Rather than being able to copy an agent every turn, agents are able to exploit (copy) once every *n* time-steps.

Synchronous communication involves all agents communicating together every *n* periods.

Asynchronous communication involves each agent having a 1/*n* chance of allowing exploitation.

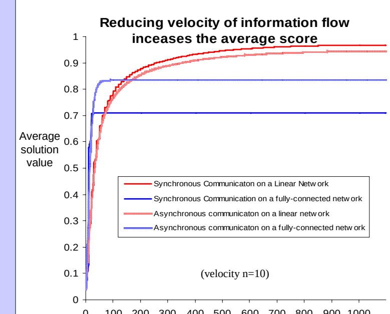

Using asynchronous communication helps the full network behave more like a sparse network, improving results over synchronous communications in the long run.

Small World Networks

- Small world networks have local clustering but short average path lengths (Watts & Strogatz 1998).
- Increasing the average path length increases overall score.

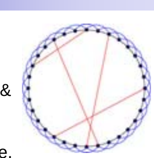

Conclusions

There is a trade-off between networks that explore and exploitation. Network communication patterns can drive heterogeneity, which in turn can affect long-term social outcomes. **More efficient communication networks can actually lower performance in the long run.**

Examples

- Digital government project teams (Binz-Scharf 2003)
- Technological diffusion (Diamond 1997)
- Creative groups (Leenders et al. 2003)