

2009

Good Edge, Bad Edge: How Network Structure Affects a Group's Ability to Coordinate

Daniel P. Enemark

University of California, San Diego, Department of Political Science, denemark@ucsd.edu

Mathew D. McCubbins

University of California, San Diego, Department of Political Science

Ramamohan Paturi

University of California, San Diego, Department of Computer Science

Nicholas Weller

University of Southern California, Department of Political Science and School of International Relations

Follow this and additional works at: http://opensiuc.lib.siu.edu/pn_wp

Recommended Citation

Enemark, Daniel P.; McCubbins, Mathew D.; Paturi, Ramamohan; and Weller, Nicholas, "Good Edge, Bad Edge: How Network Structure Affects a Group's Ability to Coordinate" (2009). *Working Papers*. Paper 27.

http://opensiuc.lib.siu.edu/pn_wp/27

This Article is brought to you for free and open access by the Political Networks Paper Archive at OpenSIUC. It has been accepted for inclusion in Working Papers by an authorized administrator of OpenSIUC. For more information, please contact opensiuc@lib.siu.edu.

Good Edge, Bad Edge: How Network Structure Affects a Group's Ability to Coordinate¹

Daniel Enemark,* Mathew McCubbins,* Ramamohan Paturi,† and Nicholas Weller§

*University of California, San Diego, Department of Political Science

†University of California, San Diego, Department of Computer Science and Engineering

§University of Southern California, Department of Political Science and School of International Relations

Coordination is a core concern in social science. Problems as diverse as trying to decide where to go to dinner, what political candidate to support or which regulatory policy to adopt all contain coordination as a core element. Most coordination problems arise among actors connected in a network, and these connections can both improve and impede a group's ability to achieve coordination. To model how links influence coordination we distinguish between “constraining edges” that make coordination harder by reducing the number of equilibrium outcomes, and “redundant edges” that make coordination easier by merely increasing communication without affecting the number of equilibria. We show experimentally that the addition of constraining edges reduces coordination, while redundant edges improve subjects' ability to solve a coordination problem.

I. Introduction

Coordination is ubiquitous to human experience. Economists and political scientists study vast coordination problems involving millions of humans, from the adoption of technological standards such as computer operating systems to the development of supranational governance such as the European Union. Often, social scientists use simple, two-person stage games to model these large-scale coordination problems. The simple coordination game is difficult because players moving simultaneously are uncertain which of the pure-strategy Nash equilibria they should aim for. This uncertainty, which arises from lack of information about the other player's action, is used as an analogy to explain the difficulty and cost of real-world coordination.

What goes unrecognized in most analyses is that the various individuals attempting to solve a coordination problem rarely have equal amounts of information. They are neither uniformly uninformed, nor are they fully and equally aware of each other's decisions. In reality, the actors in most coordination problems are embedded in a network that connects some pairs and not others. The structure of this network has important consequences for the ability even of perfectly rational actors to coordinate.

¹ Paper for presentation at the annual meeting of the Public Choice Society, Las Vegas, Nevada, March 6, 2009. The authors thank CALIT2 for its generous financial support that made this research possible. McCubbins also thanks the Chancellors Associates Chair VII for financial support.

Despite the importance of these networks, little is known about how their structure affects our ability to solve problems. What features of a network help individuals solve coordination problems? What features hinder coordination? Observational studies have come to differing conclusions about the effect of network structure on coordination and cooperation, finding that network connections can both improve and impede a group's ability to achieve coordination. We explain these conflicting findings by distinguishing between two classes of network connections: those that constrain the number of equilibria to the coordination game, and those that merely increase communication without affecting the number of equilibria. We predict that the former hinders coordination while the latter helps it.

To test our theory, we embed a simple coordination game into networks of varying structure. Our results demonstrate that the ability of human subjects to solve a coordination problem depends crucially on the network, because the network constrains the number of equilibria to the game and defines the amount of communication. In particular, we show that adding "constraining edges" that eliminate equilibria causes groups to be less successful in solving the coordination game (even though these edges also increase communication), whereas adding "redundant edges" that facilitate communication without affecting the number of equilibria causes groups to be more successful. (*Edge* is the term in graph theory for the links or connections in a network.)

II. The ubiquity of coordination and networks in politics

Coordination problems are central to politics. This section provides a brief discussion of three coordination problems prominent in the literature.²

The importance of coordination and information appears in Weingast's (1997) model of citizens' decision to support or revolt against their government. The basic problem is that an individual citizen is unsure about whether other citizens will also revolt if the government transgresses individual rights, and it only makes sense to oppose the government if enough others do it as well. Both of these aspects of the coordination problem are fundamentally about information and who possesses it. Chwe (2000) studies a very similar model of citizen protest, and focuses explicitly on the necessary conditions for an information network to lead citizens to coordinate on protest or not. In Chwe's model the nodes in the network consists of individuals and the edges between two nodes signifies that the two individuals know of the others' decisions.

In electoral politics Cox (1997) discusses how coordination among voters and donors can lead to outcomes consistent with Duverger's law. Electoral rules create an incentive for actors to

² Coordination problems occur across a broad range of social phenomena ranging from solving common pool resource problem (Ostrom 1990; Ostrom and Keohane 1994; Ahn, Ostrom and Walker 2003); to evolution of behavioral strategies (Axelrod 1984, 1997); to the development of social leadership (Calvert 1992); to political participation (McClurg 2003); to which side of the road we should drive on (Lewis 1969); to ending footbinding (Mackie 1996). At least as common as coordination problems in social settings are the presence of network effects in which the decisions of one actor affect the behavior or environment of other actors (Grannovetter 1973, 1974; Hecla 1978; Huckfeldt and Sprague 1987, 2006; Coleman 1988; Ostrom 1990; Wasserman and Faust 1994; Putnam 2000; Scholz, Berardo, and Kile 2008; Fowler 2006; Christakis and Fowler 2008 Watts 2003; see the various chapters in Kahler 2009). Again we simply provide a snippet of the literature that uses networks to understand behavior across a wide range of topics.

coordinate on a candidate or candidates to ensure they are able to maximize the probability of representation in the legislature. Cox's focus is not explicitly on the network that governs information spread between individuals, but one can easily imagine that different actors will be more or less aware of others' voting intentions or financial contributions and that the overall structure of this information will influence coordination.

In the policy realm there are a great many issues in which coordination is important and information networks affect coordination. Policy makers share information among themselves, which can play a role in the ability to solve coordination problems in a decentralized fashion. Scholz, Berardo, and Kile (2008) find that cooperation among estuary management organizations depends on the network in which the organizations operate. Policy actors that are highly connected to others are more likely to work together on policy implementation, suggesting that networks influence coordination. Likewise, Carpenter et al. (2004) show that network structure is both affected by and affects information flow between interest groups attempting to influence policy making. In particular, they conclude that when demand for information increases lobbying firms invest more resources in creating strong ties in their network, but that the resulting network actually impedes the distribution of information. Mintrom and Vergari (1998) find that network connections facilitate the spread of policy ideas between states. Policy entrepreneurs learn about policies and then propose similar policies (a form of coordination) based on information from the networks.

III. An experimental study: distributed graph coloring as a coordination game

We use an experimental adaptation of the Graph Coloring Problem (GCP) to build a better understanding of how network structure influences the ability of groups to achieve coordination. The GCP takes a given network (or "graph") and asks how to color the nodes of the network so that no two connected nodes share the same color. A coloring that satisfies this condition is called a *proper coloring* of the graph, and the minimum number of colors required for a proper coloring is called the *chromatic number* of the graph. The GCP is a workhorse problem in computer science and applied mathematics with applications to difficult allocation problems like air traffic control (Barrier and Brisset 2004). The GCP is used to model complex phenomena because it is simple to understand but difficult to solve,³ and these qualities makes it an ideal experimental task.


Traditionally, the GCP is solved by a centralized algorithm that considers the entire graph and chooses a color for every node. In our experiments, however, the problem is distributed among 16 subjects, each controlling one node. Each subject can pick among a set of available colors, and each subject can see only those nodes to which he is connected. Subjects can change colors as often as they want, but they have only three minutes to find a proper coloring, and they are only paid if they successfully do so. Following Myerson (1997) and Rasmussen (2006), we view this as a coordination game, because there are multiple pure-strategy Nash equilibria (namely, all

³ The GCP is in the NP-Complete complexity class, meaning that as the size of the input increases the time required to compute the output increases at a rate faster than any polynomial function of the input.

proper colorings). This experimental adaptation of the GCP was first developed by Kearns et al. (2006).

Figure 1 shows the interface subjects use to control their nodes. Note that this interface provides subjects with a few additional pieces of information. Inside each of his neighbors' node is a number representing the number of nodes connected to that node, or in graph theory terms, the node's *degree*. At the top of the screen there is a progress bar showing the portion of the network already solved, and a time bar showing the amount of time remaining. For more information on the experimental protocol, see Appendix B.

Figure 1. The computer interface subjects use to control the color of their nodes



IV. Previous experimental findings on networked coordination

Previous studies of coordination in a network have yielded two major conclusions: (1) greater network connectivity seems to make the game easier to solve, and (2) asymmetric incentives make the game much harder to solve. Kearns and his colleagues were the first to observe the relationship between network connectivity and subjects' ability to solve the problem. They explain that while adding edges "makes the problem more difficult from the isolated viewpoint of any individual subject ... it apparently makes the collective problem easier by reducing the

number of edges coloring conflicts must travel to be resolved.” That is, information flows faster through a more connected network.

McCubbins et al. (2008) confirm the connectivity finding, and introduce asymmetric incentives, an essential aspect of most real-world political and economic coordination problems (Calvert 1992). In their experiments, they identify one of the two colors available to subjects as a “bonus color,” so that if the problem is solved, subjects with that color will receive an additional payment on top of the standard pay. This makes the GCP a coordination game with asymmetric incentives similar to those in the battle of the sexes. Asymmetric incentives drastically reduce coordination in the lab. (In fact, large enough bonuses prevent coordination entirely.) However, just as Kearns et al. find that additional edges yield faster solutions, McCubbins et al. find that asymmetric games are solved more often when connectivity is higher.

McCubbins et al. also find that “when subjects have common, symmetric incentives to coordinate they can successfully achieve coordination regardless of the network structure.” This accords with a broader literature showing that experimental subjects easily solve coordination games with pre-play communication, even when the game requires simultaneous coordination by many players (see for example Blume and Ortmann, 2007).

As described above, Kearns et al. and McCubbins et al. both argue that increases in the number of network edges improve coordination. However, neither of these papers studied systematically the conditions under which additional edges help or hinder network coordination. If we aim to understand how changes in network structure affect collective outcomes, we need to develop a more systematic theory of how adding new edges to a network influences the graph-coloring coordination game. This theory is described below.

V. A theory of constraining and redundant edges

A. The difficulty of finding equilibria

We begin with the simple observation that coordination games—even those with symmetric incentives—are not always easy. Social scientists tend to focus on relatively simple stage games, such as the stag hunt, but computer scientists have shown that coordination problems like the GCP can be an extremely difficult computational problem even from the point of view of a centralized decision-maker (Khanna, Linial, and Safra, 2000).

In the classic two-player coordination stage game, it is immediately obvious to both players which outcomes represent successful coordination. Of the four cells in the driving game, depicted in Figure 2, two represent success (*Right, Right* and *Left, Left*). In this game, the search for equilibria is trivially simple. The challenge is for each player to guess which action the other will take, and the players can easily solve this problem with a moment of pre-play communication.

Figure 2. The driving game, a two-player coordination stage game

$P_1 \backslash P_2$	Right	Left
Right	1, 1	0, 0
Left	0, 0	1, 1

Now consider a 16-player graph-coloring game on a network that can be solved with two colors. Sixteen subjects each have two choices, so there are 2^{16} cells—65,536 possible outcomes. And just like the driving game, only two of those outcomes represent success. In this case the search process represents a serious obstacle to coordination. And that’s only a two-color game; a three-color game on 16 nodes has 43 million cells, a four-color game 4.3 billion. In these games, the difficulty is for all subjects to *find* the same equilibrium. Even with symmetric incentives and pre-play communication, the distributed search for equilibria hidden among millions of outcomes makes graph-coloring a challenging coordination problem.

B. Constraining and redundant edges

The players’ distributed search for equilibria becomes more difficult when the number of equilibria decreases, holding the number of outcomes constant.⁴ The number of equilibria in the GCP is determined by the structure of the network. When we add an edge that constrains two nodes so that they no longer can use the same color, we decrease the number of equilibria. We call such an edge a “constraining edge,” and each additional constraining edge makes coordination more difficult.

Not all edges, however, decrease the number of equilibria; if the existing edges in the network constrain two *unconnected* nodes so that they are already forced to choose different colors, adding an edge between those nodes would not decrease the number of solutions to the GCP. We call this type of edge a “redundant edge,” and it does not make coordination more difficult. In fact, in our experiments, where each individual can only see the nodes to which he is connected, adding redundant edges makes coordination easier by increasing the number of nodes the average subject can see.

To see the effect of network structure on the number of equilibria to the coordination game, consider the simple network in Figure 3.1, a line with a single added edge. This is a minimally-constrained connected three-color graph.⁵ If subjects pick colors in order from left to right,

⁴ This is true because equilibria in graph-coloring, as in the driving game, are symmetric across any player’s actions. For example, the driving game has 2 equilibria, one if Player 1 plays *Right* and one if he plays *Left*. Similarly, for a k -color network with x solutions, the GCP has x/k equilibria for each color that Player 1 could pick. If we were able to remove equilibria in a way that reduced symmetry, such as eliminating the *(Left, Left)* equilibrium in the driving game or eliminating any single solution to the GCP, the problem would become easier. This is because each player could eliminate one of the available actions from consideration as it would be less likely to yield a positive outcome.

⁵ Any graph that connects all of n nodes using $n - 1$ edges (the minimum) is called a “tree.” Since the line, like all trees, is two-colorable, one extra edge is required to achieve the minimally-constrained connected three-color graph.

Figure 3.1. A minimally-constrained connected three-color network




Figure 3.2. A maximally-constrained three-color network





Figure 3.3. The maximally-connected three-color network



The edges in graph 3.2 are a superset of those in graph 3.1;
the edges in graph 3.3 are a superset of those in 3.2.
The bold edges are those existing in the previous graph.

and each subject is given three colors from which to choose, then the leftmost node can pick any one of the three colors, the next can pick either of the two remaining colors, and the third node must pick the one color unused by the first two. The rest of the 13 nodes can choose either one of the two colors not chosen by the preceding node. The number of solutions to the GCP on this network is $3 \times 2 \times 1 \times 2^{13}$, or 49,152.

If we add a constraining edge between the second and fourth nodes of the graph in figure 3.1, the number of choices available to the fourth node decreases from 2 to 1, reducing the number of solutions from 49,152 to 24,576. We can continue adding constraining edges between all pairs of nodes v_i and v_{i+2} , reducing the number of solutions by half with each new edge, until we arrive at the graph in figure 3.2, a line of tessellated triangles, for which there are only six solutions. This is a maximally-constrained three-color graph; a three-color network can have no fewer than six solutions because there are six permutations of three colors. (Graph theorists call these graphs uniquely-colorable, because without isometric permutations, there is only one solution.)

Finally, starting with a maximally-constrained network like the one in figure 3.2, we can add redundant edges between any two nodes of a different color.⁶ These edges are “redundant” because they do not affect the conditions required to solve the coordination problem. In our experiments, where the network also defines the information available to each node, adding redundant edges gives subjects more information about each other’s actions.

VI. Hypotheses

We derive two hypotheses from our theory of constraining and redundant edges.

H1: Graphs with more constraining edges will be harder for experimental subjects to solve.

H2: Graphs with more redundant edges, *holding the level of constraint constant*, will be easier for our subjects to solve.

We measure how “hard” or “easy” it is for subjects to solve coordination problems by the frequency with which they find solutions before a three-minute time limit expires. Harder networks are those solved less frequently, while easier networks are those solved more frequently. Because our subjects are only paid for solving the problems they’re given, we can be confident they are motivated to find solutions, and if they fail to solve the problem, it is because the problem is difficult, and not for lack of effort. The results for our frequency measure are shown in Figures 4 and 5.

There is, of course, a probabilistic element to the task. After all, it is possible that if each subject chose a color at random, the outcome would be a solution to the GCP (though the probability of this is less than 1% even for the simplest network we run, as opposed to 50% for the standard 2-person coordination stage game). In addition to the randomness inherent in the task, individual

⁶ A network needn’t be maximally constrained in order to add a redundant edge, but if a network is maximally constrained, all edges that don’t increase chromatic number are redundant.

human behavior is quite unpredictable, and group behavior even more so. Thus we expect to find that some graphs are solved more frequently than others, and not that one half are always solved while the other half are always unsolved.

We also suspect that easier graphs are solved more quickly and harder graphs more slowly, because we believe subjects are motivated to solve problems as quickly as possible. Subjects should attempt to complete problems quickly for two reasons: First, if there is a risk that they will not complete the task within the three-minute deadline, they should work toward a solution as quickly as possible. Second, subjects simply value their own time. We do not, however, create any explicit, controlled incentives for solving problems quickly. In future work we plan to incentivize time-to-completion, and compare the results with the data from this round of experiments, so that we can be confident that our subjects are already attempting to complete problems quickly without explicit incentives. We do, however, present graphs of average time to completion in Appendix A, and these graphs corroborate the results in Figures 4 and 5.

VII. The Experimental Test

Our experiments use a within-subjects design, in which a group of 16 subjects attempts to solve distributed GCPs 30-40 times, with varying networks. The treatments are 29 different three- and four-color networks, and the unit of analysis is the group of 16 subjects. Each group received every treatments at least once, and the order of treatments was randomized.

To develop the networks used as the treatments, we had to use graphs requiring more than two colors, because every connected two-colorable graph is maximally constrained and therefore has no room for additional constraining edges. This was a simple but important improvement over previous work that focused on two-colorable graphs and thus overlooked the effect of constraining edges. Therefore, we began with the least-constrained connected three-color graph (shown in Figure 3.1). We then added constraining edges, two at a time, decreasing the number of solutions so that each successive graph yields $1/4$ the number of solutions of the preceding graph. We added these edges until we arrived at the minimally-connected maximally-constrained graph (shown in Figure 3.2), a tessellated line of triangles.

We then added redundant edges to the maximally-constrained graph. These edges do not change the number of solutions to the GCP, but they do increase the amount of information available to the players. We first connected the ends of the tessellated line to form a ring-lattice of triangles, and then added redundant edges, approximately 16 at a time, until we reached the maximally-connected three-color graph. At this point no more edges could be added without violating three-colorability.

The process of adding constraining and then redundant edges yielded 12 three-color graphs, shown in Appendix B, Figure 9.1. We used the same method to generate four-color graphs, which use tetrahedrons instead of triangles to constrain players to the use of four colors. This yielded 17 four-color graphs, shown in Appendix B, Figure 9.2.

VII. Results⁷

The results of our experiments confirm both of our hypotheses: constraining edges clearly hinder coordination, and redundant edges clearly help it. In fact, successful coordination depends crucially on the number of constraining and redundant edges in the network connecting players. At the minimum number of constraining edges, both three- and four-colorable graphs were solved in every trial. With the addition of more constraining edges, success rates dropped precipitously, and without redundant edges, subjects were completely unable to solve maximally-constrained graphs of either three or four colors.

Figure 4 displays along the x-axis the number of constraining edges in a network and along the y-axis the proportion of networks solved. As we move from left to right along the x-axis, each additional edge reduces the number of equilibria. As we predicted, increases in the number of constraining edges cause groups to be less successful at solving the coordination problem. (For both 3 and 4 colorable graphs once we move beyond the addition of a few non-redundant edges we observe a dramatic decline in the proportion of networks solved. This demonstrates quite dramatically how changing network structure by adding edges can impede coordination.

Perhaps the most impressive result, however, is that redundant edges can make an otherwise unsolvable problem tractable. Adding these edges does not change the actions, outcomes, incentives, or equilibria of the game; it simply increases the amount of information available to actors. Nonetheless, the frequency with which subjects solved maximally-constrained graphs rose sharply with the addition of redundant edges.

Figure 5 displays the proportion of coordination problems that are successfully solved as we add redundant edges to the fully-constrained three and four colorable networks. With the addition of only a few redundant edges (3 in a three-colorable graph and 6 in a four-colorable graph) networks that were previously unsolvable for subjects become solvable. With the full compliment of redundant edges, the maximally-constrained three-color graph was solved every time. Even the maximally-constrained four-color graph was solved in 85% of trials when all 52 redundant edges were added. This results shows that redundant edges facilitate coordination even when we start with the most difficult coordination problem possible for a given chromatic number.

Figure 6 shows logit coefficients for the effect of constraining and redundant edges, showing that these effects are significant at the .001 level.

⁷ Our data are available online: <http://dss.ucsd.edu/~denemark/GEBEdata.csv>. An annotated Stata do-file with our analysis is also available: <http://dss.ucsd.edu/~denemark/GEBEmodels.do>.

Figure 4. Constraining edges make coordination harder




Figure 5. Redundant edges make coordination easier, holding the number of equilibria constant




Figure 6. Logistic regressions confirm that the effects shown in Figures 4 and 5 are significant

	Likelihood of Solution	
	3-color graphs	4-color graphs
Effect of constraining edges	-.357*** (.112) n = 46	-.296*** (.0634) n = 77
Effect of redundant edges	.0905** (.0322) n = 32	.0943** (.0315) n = 31

Note: Entries are logit coefficients, with standard errors in parentheses.

****** $p < .005$ ******* $p < .001$

VII. Conclusions

To understand large-scale coordination problems, it is essential to consider the network that connects the individuals attempting to coordinate. Our experiments demonstrate that too many constraining connections between players can make coordination infeasible. At the same time, enough redundant edges can make even the hardest problems solvable given a reasonable timeframe.

These findings demonstrate the fundamental point that a network can do two very important things: define the constraints that must be satisfied in order to solve a problem, and facilitate the flow of information. Because a new edge may introduce additional constraints, building more connections may actually impede the ability of a group to coordinate. On the other hand, the free flow of information can help solve an otherwise intractable problem.

Previous work has found that increases in the number of connections helped coordination in both symmetric (Kearns et al. 2006) and asymmetric (McCubbins et al. 2009) games. At first blush these results appear contradictory to our finding that additional edges may impede coordination. This only happens, however, when the new edge constrains the number of equilibria to the coordination game. Because the prior experiments primarily used two-color graphs, and all connected two-color graphs are maximally constrained, these experiments essentially held constraint constant while adding redundant edges. The results from these prior experiments are perfectly consistent with our second hypothesis, that adding redundant edges makes coordination easier. The primary contribution of this paper is to make the theoretical distinction between constraining and redundant edges, and to systematically study how these different types of edges influence coordination.

Expanding the European Union

We end with a brief description of a networked coordination problem in politics. This example shows that networks are an important component of real-world coordination. It also suggests more generally that recognizing the difference between constraining and redundant edges can help us to understand important problems in social science.

The members of the European Union use their shared political institutions to address a multitude of coordination problems, such as the regulation of industry, monetary policy, and social policy. Each time a new state joins the EU it creates links between the new state and existing members. As we demonstrate in this paper, the addition of a constraining edges to a network makes coordination more difficult. Thus real-world political actors who value coordination are likely to avoid the addition of constraining edges. Nonetheless, adding new connections between European states can create immense value by enabling inexpensive communication, which helps solve future coordination problems. This feature is analogous to the beneficial effect of network connections in our experiments; it helps individual nodes share information and find equilibria.

The EU avoids adding constraining edges by requiring applicant states to adopt a host of policies designed to align them with the EU's existing equilibrium, turning what would have been a constraining edge into a redundant one. The EU requires that applicants bring their "institutions, management capacity and administrative and judicial systems up to Union standards ... with a view to implementing the *acquis* ... effectively in good time before accession" (European Commission, 2005). (The *acquis* is the full body of existing EU law, an evolving response to the problem of international policy coordination.) The EU can make these strong demands because membership is highly valued, so it can be used as a carrot for inducing applicants to harmonize their policies with existing members.

If a state can manage to adopt all of the required domestic changes, this suggests that the link formed by a new state's membership will be a redundant link that does not complicate coordination. The possibility of a new state making coordination more difficult increases if the new country is significantly different than the countries currently in the EU. The current efforts to make Turkey comply with EU standards can be understood as an attempt to ensure that Turkey's membership does not complicate policy coordination in the Union.

Bibliography (incomplete)

Barnier, N. and P. Brisset. 2004. "Graph Coloring for Air Traffic Flow Management." *Annals of Operations Research* 130, 163 (2004).

Blume, A. and A. Ortmann. 2007. "The effects of costless pre-play communication: Experimental evidence from games with Pareto-ranked equilibria." *Journal of Economic Theory* 132, 274 (2007).

Callander, Steven and Plott, Charles R. 2005. Principles of network development and evolution: an experimental study." *Journal of Public Economics*, 89, 1469–1495

Calvert, Randall. 1992. "Leadership and Its Basis in Problems of Social Coordination." *International Political Science Review* 13, 7 (1992).

Carpenter, Daniel, Kevin Esterling and David Lazer. 2004. "The Strength of Strong Ties: A Model of Contact Making in Policy Networks with Evidence from U.S. Health Politics," *Rationality and Society*. 15(4 November): 411-440.

Chwe, Michael. 2000. "s," *Review of Economic Studies* 67: 1–16.

Cox, Gary W. 1997. *Making Votes Count*. Cambridge University Press

European Commission, "Negotiating Framework [for Turkish Accession]" (Luxembourg, Oct. 3 2005). Accessed 3/1/09 at http://ec.europa.eu/enlargement/pdf/st20002_05_tr_framedoc_en.pdf.

Fowler, James H. 2006. "Connecting the Congress: A Study of Cosponsorship Networks" *Political Analysis* 14 (4): 456-487.

Huckfeldt, Robert and John Sprague. 1987. "Networks in Context: The Social Flow of Political Information." *The American Political Science Review*, Vol. 81, No. 4 (Dec., 1987), pp. 1197-1216

_____. 2006. *Citizens, Politics and Social Communication: Information and Influence in an Election Campaign*

Jensen, Tommy R. and Bjarne Toft. 1994. *Graph Coloring Problems*. Wiley-Interscience

Kearns, M., S. Suri, and N. Montfort. 2006. "An Experimental Study of the Coloring Problem on Human Subject Networks." *Science* 313, 824.

Khanna, S., N. Linial, and S. Safra. 2000. "On the hardness of approximating the chromatic number." *Combinatorica* 20, 393.

Mackie, Gerry. 1996. "Ending Footbinding and Infibulation: A Convention Account."

American Sociological Review, Vol. 61, No. 6. (Dec., 1996), pp. 999-1017

McClurg, Scott D. 2003. "Social Networks and Political Participation: The Role of Social Interaction in Explaining Political Participation." *Political Research Quarterly*. 56 (December): 448-64.

McCubbins, M., R. Paturi, N. Weller. 2009. "Connected Coordination: Network Structure and Group Coordination." *American Politics Research*, forthcoming.

Myerson, M. *Game Theory: Analysis of Conflict* (Harvard University Press, Cambridge, MA, 1997).

Rasmussen, E. 2006. *Games and Information*.

Scholz, John, Ramiro Berardo and Brad Kile. 2008. "Do Networks Enhance Cooperation? Credibility, Search, and Collaboration in Joint Environmental Projects". *Journal of Politics* 70(2): 393-406

Scholz, John T. and Ramiro Berardo. 2008. "Self-Organizing Policy Networks: Risk, Partner Selection and Cooperation in Estuaries." Presented at the 2008 Networks in Political Science Conference, Harvard University.

Weingast, Barry. 1997. "The Political Foundations of Democracy and the Rule of Law." *American Political Science Review*. (June 1997) 91: 245-63.

Appendix A: Time-to-Completion Data

As discussed in the body of the paper, subjects are not explicitly incentivized to solve problems quickly, but we believe they are attempting to solve them quickly to avoid the risk of failure and because they value their own time. Figures 7 and 8 show that adding constraining edges increases mean time-to-completion, and adding redundant edges decreases mean time-to-completion, corroborating the results in Section VII.

Figure 7. Constraining edges make coordination harder





Figure 8. Redundant edges make coordination easier, holding the number of equilibria constant



Appendix B: Experimental Procedure

In all of the experiments reported in this paper the networks consisted of 16 nodes/ subjects. After subjects reported to the experiment they were placed at a computer behind partitions so that they could not see the other participants. Before the experiment began we read aloud the directions to all the subjects in the room to ensure that the procedures, rules, and incentives are all common knowledge. In addition, all subjects take a short quiz (with payment for correct answers) to make sure they understand the experiment.

To ensure that subjects did not repeatedly choose the same color over and over again (thereby possibly facilitating coordination), we used a palette of 10 colors and for each coordination game a subject was randomly given the minimum number of needed colors from a palette of ten possible colors. In addition, subjects in a given game chose from different colors to ensure that if they were able to see another's monitor the subject could not learn anything about how that subject was acting. The central server presented each subject's terminal with the number of colors utilized for a given network (equal to that network's chromatic number). The information displayed on each computer was controlled by our central server that utilized a software program developed and shared by Michael Kearns and Stephen Judd at the University of Pennsylvania.

We conducted experiments with five different groups of 16 subjects. We collected data from 179 different attempts to solve the graph coloring problem. Of those 179 attempts, 75 trials involved three-color graphs and 104 involved four-color graphs. The payment for coordination was \$1 per subject and groups had three minutes to achieve coordination. Each experimental trial ends either when the time limit is reached or the group achieves coordination successfully, and this is known to all subjects in the experiment.

In the experiment each subject controls the color of one node in the network so coordination is the result of distributed actions. However, subjects do have more information than just the color of their own node. The screen that subjects saw during the experiment contained the following information.

- Local View: Subjects are able to see their node and the neighboring nodes to which they are connected. Each node in their local neighborhood contains a number in its center that tells the subject how many total edges a neighboring node has. This allows them to see the color they have chosen for their node as well as the colors chosen by their neighbors
- Color choices: subjects can see the three or four colors from which they can choose

- **Elapsed Time Bar:** This bar kept track of the amount of time since the session began, and allows subjects to determine how much time is remaining before the time limit.
- **Completion Percentage Bar:** This bar provides information about how close the entire network is to completion. The percent completed represents the number of edges without a coloring conflict divided by the total number of edges in the graph.

During the actual experiment the bars for elapsed time and completion percentage are updated in real time. Figure 1 displays a typical screen shot that a subject sees before the experiment begins. By looking at the screen a subject with this picture can determine that he is connected to three nodes and that one of those nodes has eight total edges (and the other two nodes each have three total edges). The subject can also see that he can choose between pink and violet during this session. During the experiment subjects continue to see this screen shot, but the progress bar and elapsed time bars change to reflect the global condition of the network.

The screen shot shows that although subjects have a tremendous amount of information available to them during the experiment, they do not know the structure of the entire network nor do they know who their geographic neighbors are in the experiment. In addition, subjects are assigned to their node randomly at the beginning of each session within a given experiment. Therefore, even if they discover to whom they are connected in a given session that will only last for one session. This procedure ensures subjects do not always occupy the same position in a network when we repeat network structures with different bonus parameters. Because subjects do not know the entire structure of the network they are not able to learn anything about how their choices relate to the group's success or failure for a given network type.

Figures 9.1 and 9.2 show the three- and four-color graphs, respectively, used in our study.

Figure 9.1. Three-color graphs




Figure 9.2 Four-color graphs

