Good Edge, Bad Edge:
How Network Structure Affects a Group's Ability to Coordinate[1]

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 issue in social interaction. Problems as diverse as trying to decide where to go to dinner, what political candidate to support, or which regulatory policy to adopt all involve coordination. Coordination problems are often solved via decentralized action and the structure of information among the participants can play a decisive role in whether or not groups can coordinate. The structure of information can be represented by a network in which actors are nodes, and the connections between them represent the ability of the two actors to observe each others' action. To model how network links influence coordination we distinguish between "constraining edges" that make coordination harder by reducing the number of coordination solutions, and "redundant edges" that make coordination easier by merely increasing communication without affecting the number of solutions to the coordination problem. We utilize a general networked coordination problem to demonstrate experimentally that the constraints imposed by connections to others can impede coordination, but connections that do not constrain can facilitate coordination. The results help us to understand the importance of network connections, and how different types of connections can influence real-world coordination.

---

[1] Paper for presentation at the annual meeting of the Midwest Political Science Association, April 2-4, 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.

1. Introduction

Coordination is ubiquitous to social experience. Examples of coordination range from the adoption of a domestic currency, to citizens' attempts take collective political action (i.e. voting, protest, rebellion), to the choice of regulatory regimes for domestic and international policies, to methods for shipping goods around the world. Although each of these problems differs in their specifics they are all examples of a general coordination problem in which the defining feature is the presence of multiple solutions, and actors must determine which solutions represent successful coordination. The difficulty of resolving coordination problems depends upon the number of solutions to the problem and the information about others' decisions.

The key to resolving most coordination problems is information about others actions. In most settings individuals attempting to solve a coordination problem have different amounts of information. They are neither uniformly uninformed, nor are they fully and equally aware of each other's decisions. Knowledge of others' decisions and action can be modeled as a network where each actor is a node and an edge between two nodes represents that both connected actors can observe the others' actions. In most real-world settings the network in which actors are embedded does not connect every actor to every other actor, and the structure of the network can have important consequences for the ability of actors to coordinate.

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 solutions to the coordination game, and those that merely increase communication without affecting the number of solutions. 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 solutions to the game and defines the amount of communication. In particular, we show that adding constraining edges that eliminate solutions causes groups to be less successful in solving the coordination game to the extent that some problems become impossible for groups to solve. However, subjects are able to solve these problems when we add new edges that do not reduce the number of solutions to these previously impossible coordination problems. These results demonstrate that additional network connections can make coordination more difficult, but edges can also facilitate coordination and in fact some edges make coordination possible when it was previously impossible. Therefore, when we study and build real-world networks we need to think carefully about the types of edges in the network and how those edges affect collective outcomes.

2. The ubiquity of coordination and networks

Coordination problems are central to social interaction. The ability to solve a coordination problem depends largely on the difficulty of the problem (are solutions available and known?) and the presence of information about others' actions. We briefly discuss three coordination problems that are significant in political situations.

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 Duverger's law. Electoral rules create an incentive for actors to 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/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.[2]

Often, social scientists use simple, two-person stage games to model these large-scale, complex coordination problems. The simple coordination game is difficult because players moving simultaneously are uncertain of which action to take. 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. However, in all of the settings discussed above not only does information affect coordination but so does the number of actors and their possible actions. With multiple actors and even just two a few possible actions per individual, there are thousands, if not millions, of possible outcomes. Only some subset of these many possible outcomes represents successful coordination, and the actors' must discover one of the coordinated outcomes.

---

[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; Heclo 1978; Huckfedlt 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.

3. The experimental task: a general coordination game

The prior examples of coordination all differ how difficult it is to resolve the coordination problem and the type of information, but at their core they are situations in which decentralized decision makers attempt to find a solution to a coordination problem. Therefore, we need a general coordination game for which we can vary the difficulty of the problem and the information structure. To study experimentally how problem difficulty and information affect coordination we use a problem called 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 (Barnier 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 also make it an ideal experimental task.

Traditionally, the GCP is solved by a centralized algorithm that looks at 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

---

[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.

are only paid if they successfully do so. Following Rasmussen (2006), we view this as a coordination game, because there are multiple equilibria. 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 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.

3.1. Previous findings on networked coordination

Previous studies of coordination in a network have yielded two major conclusions: (1) greater numbers of network connections makes the game easier to solve, and (2) asymmetric incentives make the game harder to solve. Kearns and his colleagues were the first to observe the relationship between network connections 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 network with more connections.

McCubbins et al. (2008) confirm that greater connections facilitate faster solutions to symmetric coordination games. This results fits with a broader literature showing that experimental subjects easily solve coordination games with pre-play communication, even when the game requires simultaneous coordination (see for example Blume and Ortmann, 2007). They also introduce asymmetric incentives, an essential aspect of many 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, McCubbins et al. find that asymmetric games are solved more often when connectivity is higher.

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.

4. A theory of coordination and network structure

4.1. The difficulty of finding solutions

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 coordinated solutions is trivially simple. The challenge (in the absence of communication) 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.

Now consider a 16-player graph-coloring game on a network that can be solved with two colors (i.e., a simple circle). 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 solution. Even with symmetric incentives and pre-play communication, the distributed search for solutions hidden among millions of outcomes makes coordination a difficult problem.

4.2 Building a network one link at a time: *Constraining and redundant edges*

To build a network requires that edges are added between nodes. In this paper we focus on how adding edges to a network can change the number of possible solutions to a coordination game as represented by the graph coloring problem. The network structure in the graph coloring problem defines the constraints that a proper coloring must satisfy (that is the network determines which nodes must not be the same color), and the network also determines how information moves throughout the network. We classify edges that can be added to a graph into two types – redundant and constraining edges. A *redundant edge* does not change the number of solutions, but does change how information can move among the nodes. A *constraining edge* changes the solution space by decreasing the number of solutions and also affects information in 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 solutions. This basic insight suggests that changes in network structure can be classified based upon how the change affects the number of solutions to a problem. The players' distributed search for solutions becomes more difficult when the number of solutions decreases, holding the number of outcomes constant.[4] We discuss a measure of the difficulty of coordination in the next section.

To see the effect of network structure on the number of solutions to the coordination game, consider the simple three colorable-network in Figure 3a, 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, 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. The general formula for the number of solutions is described in the next section.

If we add a constraining edge between the second and fourth nodes of the graph in figure 3a, the number of choices available to the fourth node decreases from 2 to 1, reducing the

---

[4] This is true because solutions in graph-coloring, as in the driving game, are symmetric across any player's actions. For example, the driving game has 2 solutions, 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 solutions for each color that Player 1 could pick. If we were able to remove solutions 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.

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} (where *i* indexes nodes from 1 to 16 going from left to right), reducing the number of solutions by half with each new edge, until we arrive at the graph in figure 3b, 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.)

Not all edges, however, decrease the number of solutions; 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 (i.e. connecting edges 1 and 4 in figure 3a). We call this type of edge a *redundant edge*, and it does not make coordination more difficult. In our experiments, where coordination is achieved via decentralized decisions and 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.[6]

---

[6] The categories *redundant edge* and *constraining edge* are exhaustive and mutually exclusive, but this typology only applies in the context of a marginal addition to or subtraction from an existing network. That is, *redundant* or *constraining* is a description of how adding or subtracting a particular edge affects the network, not an inherent property of the edge between some pair of nodes. For example, imagine a "ring network," where the first node is connected to the second, the second to the third, etc, and the sixteenth node is connected to the first. This network requires two colors, and has only two solutions. If we removed only the edge between nodes 3 and 4, we would be left with a line network, which still has only two possible solutions (i.e., it has just as much constraint as the ring network). Thus, we would be removing a redundant edge. Similarly, if we removed only the edge between nodes 7 and 8, we would be removing a redundant edge. But if we removed one edge and then the other, the second action would remove a constraining edge, regardless of which edge is removed second.

For instance, starting with a maximally-constrained network like the one in figure 3b, we can add redundant edges between any two nodes of a different color.[7] 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.

4.3 Measuring the Difficulty of Finding a Solution

To measure the difficulty of solving a particular graph coloring problem we determine the number of possible proper colorings for a graph as well as the number of possible outcomes if each node randomly selects from the graph's chromatic number. We measure the difficulty of solving any a given graph coloring problem by taking the log of the ratio of the number of proper colorings for a network to the number of possible outcomes for a network. The number of possible outcomes for a given network is $k^{(\text{\# of nodes})}$ where *k* is a graph's chromatic number. The number of possible proper colorings of a network is determined by the number of triangles (for 3 colorable graphs) and tetrahedrons (for 4 colorable graphs) and then the number of combinations the remaining, unconstrained nodes can take on. Where *k* represents the chromatic number of a graph, the formula to compute the number of solutions is: $k!\,(k-1)^{(\text{\# nodes not needing maximum chromatic option})}$. The number of proper colorings is affected by the number of constraining links in a network, so when we change the network structure by adding a constraining link our measure of Solution Likelihood also changes. In Section 7 we utilize this measure in our regressions that test the effect of network links on the probability of successful coordination.

---

[7] 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.

5. 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.

As we increase the number of constraining edges we decrease the number of possible solutions to a graph coloring problem making it more difficult to find a solution to the coordination problem.

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

When we add edges we also increase the information subjects have about the graph. If we add edges that increase information but do not make the problem more difficulty (i.e. redundant edges), then subjects should be able to solve the problem more easily.

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.

We also believe 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 discussed in the next section.

6. 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 network structures. We use 29 different three- and four-color networks, and the unit of analysis is the group of 16 subjects. Each group received every treatment at least once, and the order of treatments was randomized.

To develop the networks, 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 3a). 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 3b), a tessellated line of triangles. The same process is repeated for four-colorable graphs using tetrahedrons.

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 8a. 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 8b.

7. Results[8]

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-

---

[8] Our data are available online at http://dss.ucsd.edu/~denemark/EnemarkMcCubbinsPaturiWeller.csv.

axis, each additional edge reduces the number of solutions. 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 solutions to 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 complement 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. This is a profound result because it shows that a problem that is literally impossible for our subjects to solve can be made possible by adding redundant links that do nothing to directly affect the difficulty of the problem.

We also test the effect of edges on the probability of coordination via statistical analysis. To study our prediction about the probability of coordination we utilize a random-effects logit model to estimate the following regression:

$\Pr(\text{Coordination}) = \alpha + \beta_1 \text{SolutionLikelihood} + \beta_2 \text{NumEdges} + \varepsilon$

In Table 1 we present the results of this regression separately for three and four colorable graphs. We do not find that the random effects are significant, which suggests both that there is not a significant difference between groups in their ability to solve the coordination problem and that we have good balance in the treatments to which we expose the groups. The addition of an edge can affect the solution likelihood if the edge is constraining or not affect the solution likelihood measure if the edge is redundant. Therefore, the effect of the number of redundant edges is estimated from graphs where we add redundant edges to a maximally connected graph.

As we predicted, increases in the number of edges, holding constant the difficulty of the problem, are associated with a greater probability of coordination. This result implies that redundant edges improve coordination, because until we reach the fully constrained networks each addition of an edge also affects the Solution Likelihood, and therefore, the independent effect of edges we identify is estimated by the effect of adding redundant edges to the fully constrained networks. These results confirm the visual results presented in Figures 4 and 5.

The results of these experiments provide strong evidence that how we build a network can have dramatic effects on group coordination. If we build a network in a way that makes the problem more difficult we can quickly get to the point where a group has great difficulty solving a coordination problem. However, if we build a network by adding edges that facilitate communication, but do not affect the difficulty of the problem, then these edges improve group coordination. This is an important result for the design of real-world networks.

8. 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 solutions 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. These findings make it crystal clear that as the number of constraints increases coordination can become impossibly difficult. We also show, however, that the addition of redundant edges can make possible a previously impossible coordination problem.

The effect that networks have on achieving coordination is critical in our experimental setting, and it is also important in real-world coordination. 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 that political actors understand how links in a network affect coordination, and therefore aim to build redundant rather than constraining links.

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 solutions.

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. The efforts on the part of the EU can be seen as their attempt to ensure that Turkey creates a redundant link with the existing EU countries.

References (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. “Communication and Coordination in Social Networks,” 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 Quarter*ly. 56 (December):448-64.

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

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.

Table 1: Communication encourages coordination

	Three Colorable Networks	Four Colorable Networks
	DV: Probability of Coordination	DV: Probability of Coordination
Solution Likelihood	1.47 (0.39)**	1.60 (0.29)**
Number of Edges	0.08 (0.028)**	0.08 (0.02)**
Constant	6.58 (1.91)**	6.88 (1.36)**
N	75	104

** = significant at 0.01 level. Regressions estimated using a random effects logistic regression in STATA 10.0.

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

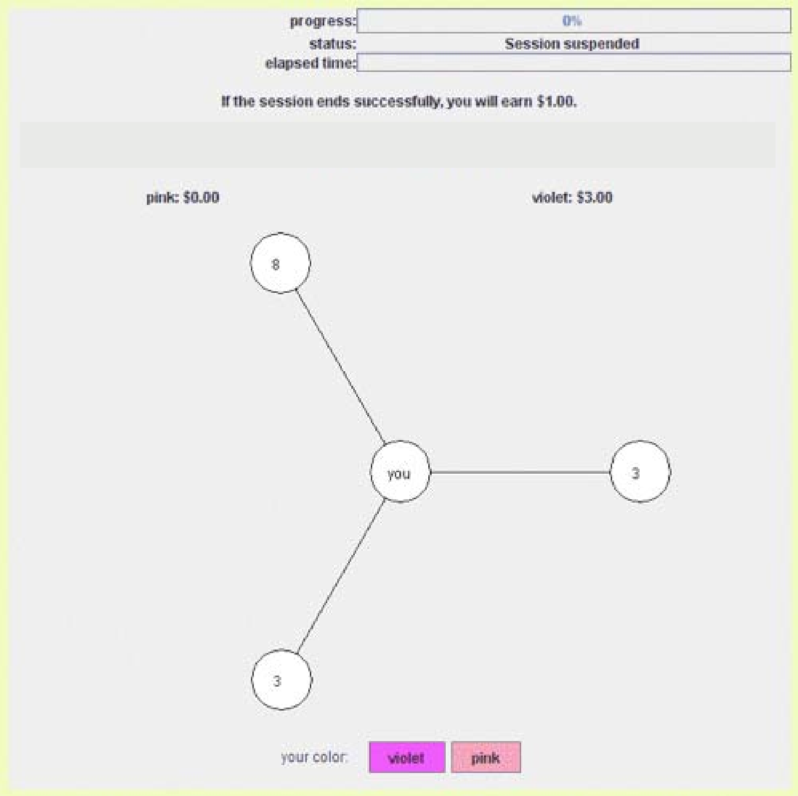

progress: 0%

status: Session suspended

elapsed time:

If the session ends successfully, you will earn $1.00.

pink: $0.00

violet: $3.00

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

P1\P2	Right	Left
Right	1, 1	0, 0
Left	0, 0	1, 1

Figure 3a-c: Building a network

Figure 3a. A minimally-constrained connected three-color network

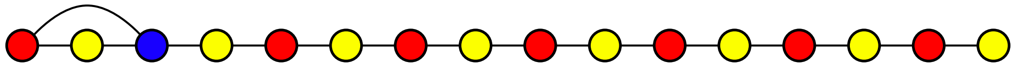

Figure 3b. A maximally-constrained three-color network

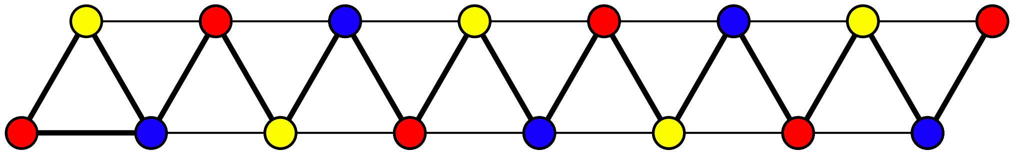

Figure 3c. The maximally-connected three-color network

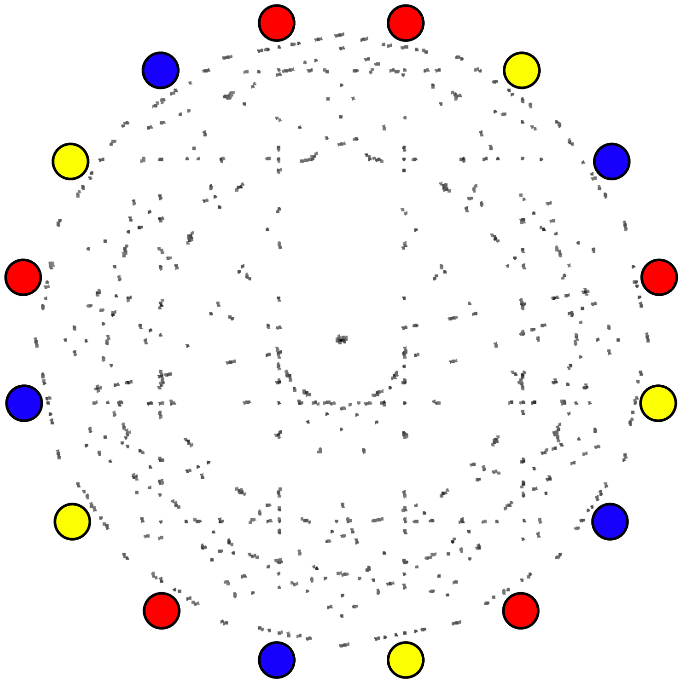

The edges in graph 3.2 are a superset of those in graph 3.1;
the edges in graph 3.3 are a superset of those in 3.2.
The bold edges are those existing in the previous graph.

Figure 4. Constraining edges make coordination harder

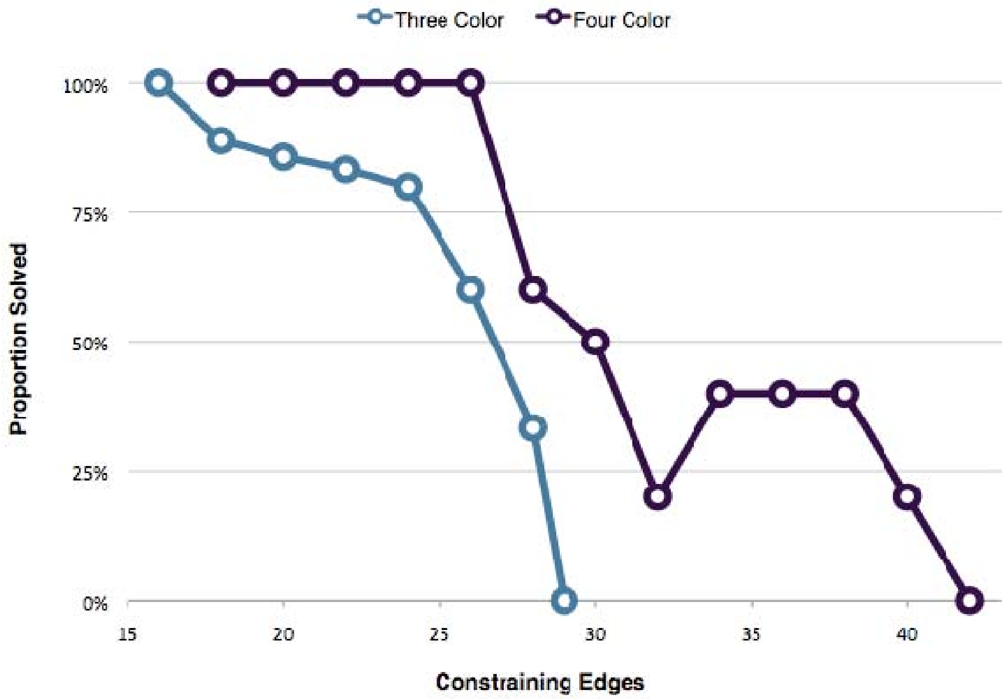

Figure 5. Redundant edges make coordination easier

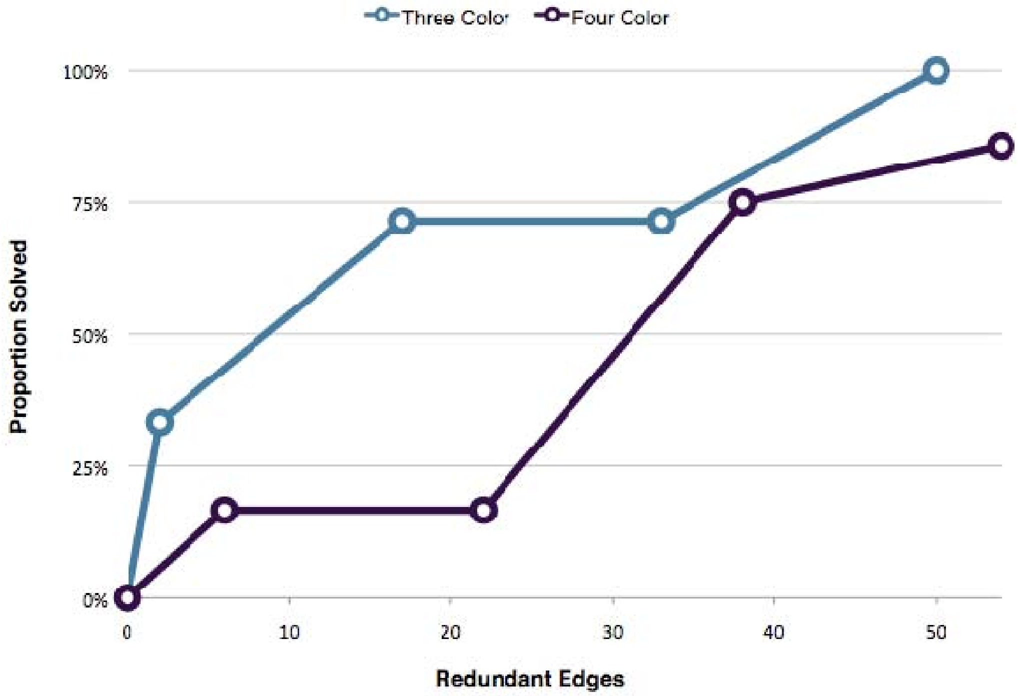

Appendix A: Time-to-Completion Data

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 6 and 7 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.

We also investigate the time it takes subjects to achieve coordination.[9] To do so we estimate via ordinary least squares the following regression:

$$\text{Time to Coordinate} = \alpha + \beta_1 \text{SolutionLikelihood} + \beta_2 \text{NumEdges} + \varepsilon$$

The results for this regression in Table A1 for both three and four colorable networks confirm our predictions. As the coordination problem becomes less difficult (increases in solution likelihood) solutions are more rapid. In addition, holding constant the difficulty of the coordination problem increases in the number of edges (which, by definition are redundant edges) decrease the amount of time to achieve coordination. Both of these results are consistent with our predictions. Although there are many limitations to using R^2 values, it is remarkable that we can explain right around 50% of the variance in time to achieve coordination with only two variables.

---

[9] Each trial was capped at 180 seconds and ended if coordination was not achieved by that point. To ensure that our results were not a function of that time limit we also estimated our regressions with 200 and 300 seconds used as the time an unsuccessful network took. It does not make any difference in our substantive conclusions.

Figure 6. Constraining edges make coordination take longer

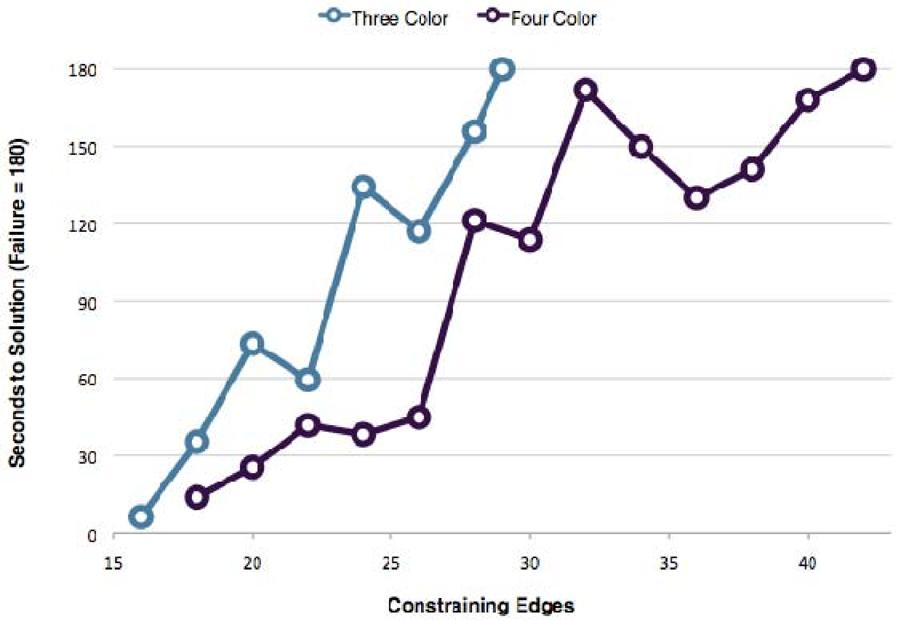

Figure 7. Redundant edges make coordination faster

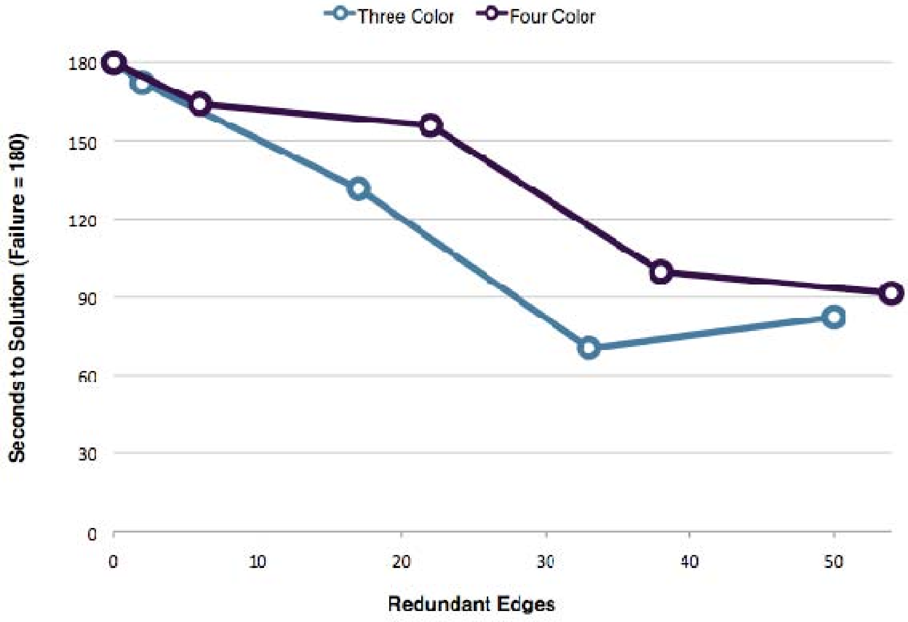

Table A1: Communication encourages faster coordination

	Three Colorable Networks	Four Colorable Networks
	DV: Time to coordinate	DV: Time to coordinate
Solution Likelihood	-46.6 (5.84)**	-38.5 (3.68)**
Number of Edges	-2.02 (0.40)**	-1.94 (0.34)**
Constant	-93.24 (25.6)**	-49.0 (14.1)**
N	75	104
Adj R2	0.46	0.55

** = significant at 0.01 level. Regressions estimated using OLS in STATA 10.0

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 palate of 10 colors and for each coordination game a subject was randomly given the minimum number of needed colors from a palate 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 8a and 8b show the three- and four-color graphs, respectively, used in our study.

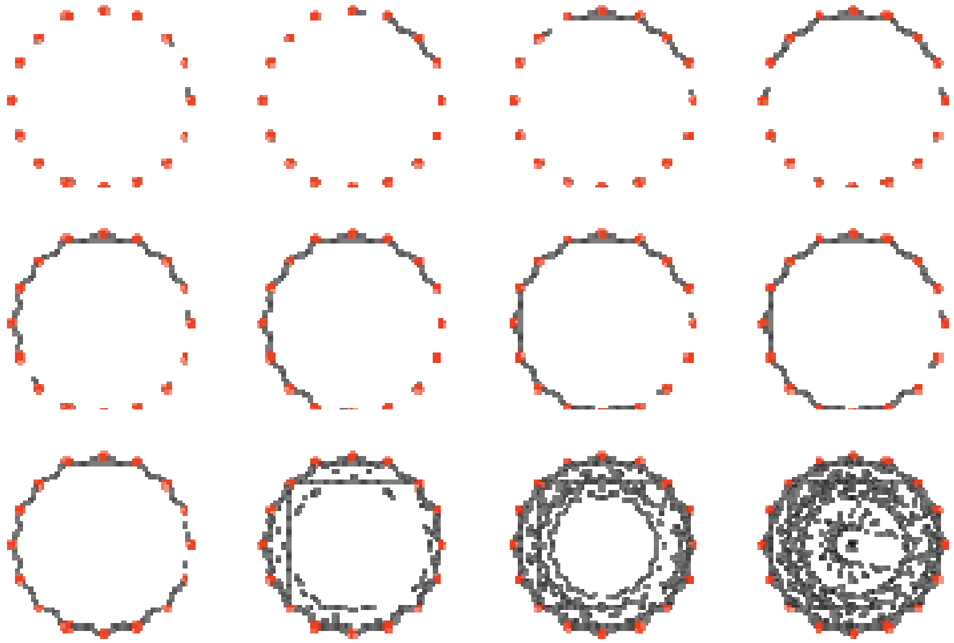

Figure 8a. Three-color graphs

Figure 8b Four-color graphs