

Sheep, contrarians and saboteurs: disruptors in friendship games

David S. Dixon

Department of Economics
University of New Mexico

On a network where games are played among immediate network neighbors - or friends - a player's optimal strategy in one subnetwork may be non-optimal in another. For the two friendship-based payoffs examined, Nash equilibrium is a function of the number of players playing each strategy, rather than the specific strategies of individual players. With all players facing the same strategies, payoffs, and Nash equilibria, the network reaches a static equilibrium when every subnetwork reaches a subgame Nash equilibrium and no player has incentive to change strategies. Payoffs are introduced that are based on the global state of the network rather than the payoffs in the local subnetwork. Three global strategies are examined: sheep (always choose with the majority), contrarians (always choose against the majority), and saboteurs (intentionally prevent an equilibrium by switching between the sheep and contrarian payoffs). The gradual introduction of players with these global payoffs results in a shift from static equilibrium to dynamic equilibrium and, in some cases, to disequilibrium outcomes. Different relative percentages of those playing each strategy lead to new outcomes, new dynamics, or both. The results show that even small percentages of the global-payoff players can disrupt outcomes significantly. The results also imply that, of the three, the sheep are most disruptive.

1. Network games

Galeotti et al. (2010) present the theoretical basis for, and some examples of, network games. In these games, the players are distributed on a random network and the payoffs are functions of the expressed preferences of the immediate neighbors on the network. For a

model of strategic substitutes, the payoff is such that, if at least one neighbor is paying the cost, none of the other neighbors has an incentive to also pay it. This is a free-rider model, similar to the private provision of a public good. For a model of strategic complements, the payoff is highest for the choice that is supported by a majority of neighbors. This is similar to a network externality, where adopting the most common word processing software, for example, maximizes the ability to share documents with neighbors. Lamberson (2011)¹ shows that, on a friendship network with one of these games, the network reaches a static equilibrium for strategic complements and two equilibria for strategic substitutes. Dixon (2011b) demonstrates that these results can be reproduced using an agent-based model (ABM). That equilibrium outcomes are somewhat sensitive to the topology of the friendship network.(Dixon, 2011a)

Lamberson (2011) adopts the term *friend* for these network neighbors, reflecting the fact that adjacent nodes in a social network can be quite distant geographically. Lamberson also examines the effect of clustering: the extent to which a player’s friends are friends amongst themselves. Lamberson finds that clustering can increase the equilibrium provision of a public good - perhaps inefficiently - in a strategic substitutes model. In a strategic complements model, however, clustering can improve the diffusion of new ideas or technologies. Lamberson explores numerical simulations on a regular random network (all nodes have exactly the same number of friends) and on a Bernoulli random network (each node has a randomly distributed number of friends). Using ABM simulations, Dixon (2011a) examines friendship games with other network topologies: a small world network (a Bernoulli network where some connections are reformed to increase clustering), and a power-law network in which the number of friends has a power-law distribution. Dixon finds that equilibrium outcomes are somewhat sensitive to the topology of the friendship network.

A note on terminology: A player in the game theory sense is the same as a node in the networks discussed, which is the same as an agent in the agent-based modeling sense, which is called a turtle in NetLogo agent-based models. That is, the terms player, node, agent, and turtle are synonymous. Which term is used will depend on context, though there is incomplete consistency in this draft. One intentional inconsistency arises in the context of game theory, where usages like “the strategy an agent plays” simply flows better than “the strategy a player plays”.

1.1. The strategic complements game

Suppose there are two strategies, x and y . If an agent has k friends, then, at any given instance, there are k_x of them playing strategy x , and k_y of them playing strategy y . For the strategic substitutes models, the payoff for playing strategy x is

$$\pi_x(k_x) = f(k_x) - c_x \tag{1}$$

¹For an updated version see (Lamberson, 2015).

and the payoff for playing strategy y is

$$\pi_y(k_x) = f(k - k_x) - c_y \quad (2)$$

where f is a non-decreasing function and c_x and c_y are the costs of play x and y , respectively.

The adoption of a standard is a strategic complement: an agent chooses what most of its friends choose. The decision of friends has a positive affect in that an agent tends to take the same choice as friends that favor a strategy.

The strategic complements model presented in Lamberson (2011) is simply this: play strategy x if four or more neighbors (in a random network of mean degree 10) are playing x . That is

$$\pi_x(k_x) = \begin{cases} 1 & k_x \geq 4 \\ 0 & \text{otherwise} \end{cases}$$

1.2. The strategic substitutes game

For the strategic substitutes models, the payoff for playing strategy x is

$$\pi_x(k_x) = 1 - c_x \quad (3)$$

where $0 < c_x < 1$ and the payoff for playing strategy y is

$$\pi_y(k_x) = \begin{cases} 1 & k_x \geq 1 \\ 0 & \text{otherwise} \end{cases}$$

The provision of a public good is a strategic substitute: an agent needn't provide it unless none of its friends do. The decision of friends has a negative affect in that an agent tends to take the opposite choice of any friend that favors a strategy.

The strategic substitutes model presented in Lamberson (2011) is simply this: play strategy x if fewer than four neighbors (in a random network of degree 10) are playing x . That is, costs are zero and

$$f(k_x) = \begin{cases} 1 & k_x \leq 4 \\ 0 & \text{otherwise} \end{cases}$$

2. Network Topology

The models in Lamberson (2011) feature 1000 players on a random network of degree 10. The number of other players to which a player is connected is that player's *degree*. In some cases it is a regular random network, meaning that connections are made at random until every node has exactly ten neighbors. In other cases it is a Bernoulli random network with an edge probability of 0.01. That is, for a network potentially connecting all players to all players, there is a probability of one in one hundred that a given connection will actually exist. The degrees of the nodes are distributed binomially, with a mean degree of approximately 10. That is, players have, on average, ten friends.

In the ABM developed for this paper, there are five ways in which a random network can be generated. These are referred to as the *regular*, *Bernoulli*, *preferential attachment*, *truncated power law*, and *small world* network models. The following are descriptions of these network models.

2.1. Regular Random Network

One way to form a random network is for each agent to make two friends, but only with other agents that don't already have two friends. This results in a regular random network where the nodes have a uniform degree of two, ensuring that the average degree is two. It is a high connectivity network: no agents will end up completely disconnected from the network.

For example, a simple form of the strategic substitutes model in Section 1.2 is implemented in NetLogo and simulated as outlined in Section 3. Figure 1a shows the results for a degree 10 regular network of 1000 nodes. This plot overlays the ABM results (black) on an image of the corresponding numerical results in Lamberson (2011) (blue). As with Lamberson's result, about 40 percent of the nodes in the ABM are playing strategy x at equilibrium. This is the characteristic outcome for this game and its payoffs: the network reaches a static equilibrium of 40 percent of the agents playing x and 60 percent playing y . The ABM reaches a slightly higher equilibrium ratio than Lamberson's for all initial values, making this unlikely to be an outcome of stochastic effects.

Similarly, a simple form of the strategic complements model in Section 1.1 is implemented in NetLogo and simulated as outlined in Section 3. Figure 1b shows the results for a degree 10 regular network of 1000 nodes. This plot overlays the ABM results (black) on an image of the corresponding numerical results found by Lamberson (2011) (blue). As with Lamberson's result, initial distributions of more than 35 percent playing x lead to the *all-in* (all players playing x) equilibrium, while initial distributions of 20 percent or fewer playing x result lead to *all-out* (no players playing x) equilibrium. This is the characteristic outcome for this game and its payoffs: the network reaches one of two static equilibria of either a) no players playing x or b) all players playing x . At the boundary of these two outcomes, Lamberson's result shows an initial distribution of 25 percent playing x ending in the *all-out* (no players playing x) equilibrium. The ABM, on the other hand, shows this

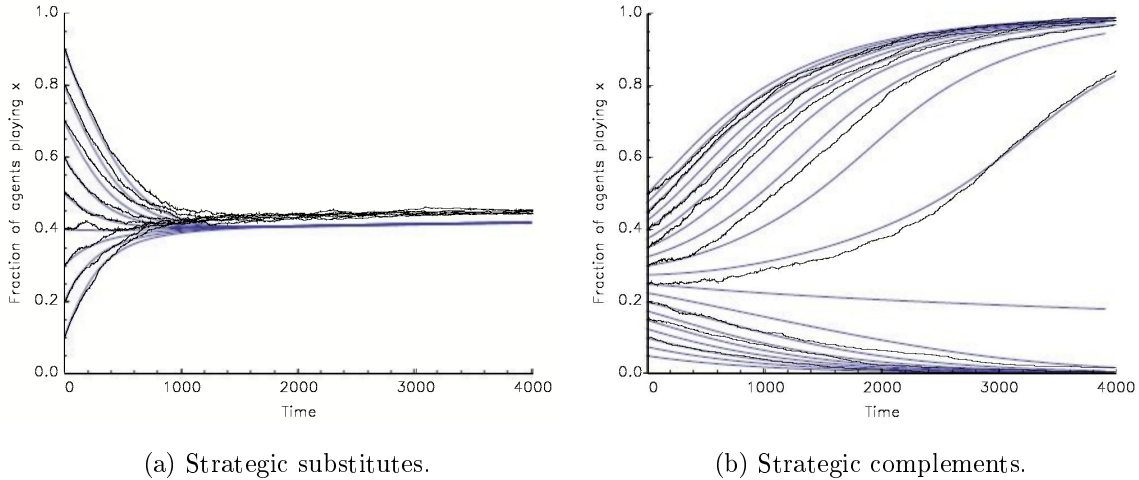


Figure 1: Degree ten regular network overlaid on Lamberson (2011) Figs. 2 and 1, respectively.

initial distribution ending in the *all-in* (all players playing x) equilibrium. This difference, as well as other deviations of the ABM from Lamberson’s, may be due to stochastic effects.

2.2. Bernoulli Random Network

A Bernoulli random network with n nodes and probability p that an edge will exist results in nodes with degrees that are binomially distributed about the mean np . (Erdős and Rényi, 1959) There are two common methods for constructing this type of network. One is the $G(n, M)$ method, which is called the Erdos-Renyi random network in this paper, and the other is the $G(n, p)$ model, which is called the Gilbert random network in this paper.

In the Erdos-Renyi random network, a network is chosen at random, with uniform distribution, from the collection of all possible networks with n nodes and M edges. For the models in this paper, M is not known a priori, so a $G(n, M)$ model is approximated by adding edges between randomly chosen pairs of nodes until the mean degree reaches the desired value.

If all friendship pairs are equally probable with probability p then a Gilbert random network, $G(n, p)$ is formed. (Gilbert, 1959) The mean degree is np , where n is the number of nodes. For the models in this paper, this is created with nested loops: an outer loop over a randomized list of all nodes, and an inner loop over a randomized list of all nodes that haven’t already come up in the outer loop. In the inner loop, a connection is formed if a uniform random draw is less than or equal to p .

No discernible difference is seen between the Erdos-Renyi and Gilbert methods: the degree distributions of both networks very approximate a binomial distribution with R-squared greater than 0.96, and clustering is about 0.018 in both. The Gilbert algorithm runs about one-third faster, however, which is a consideration when generating many thousands of networks

2.3. Preferential Attachment Network


The preferential attachment model (Wilensky, 2005a) is included in the NetLogo (Wilensky, 1999) demo library and is based on an approach by Barabási and Albert (1999). This is an approximation of a *scale-free network*, a network with a power-law distribution of degrees per node. Also called a Pareto distribution, it results in a few nodes having a very large number of connections and many nodes with very few connections. Scale-free networks are seen in academic citations de Solla Price (1965) and in a variety of Internet linkages. Albert and Barabási (2002) find that the probability of a link for a node with degree k is $p(k) = \alpha k^{-\gamma}$ where γ is between 2 and 3. The NetLogo preferential attachment algorithm yields a γ of approximately 2, and a mean degree of approximately 2. The outcome space for a degree two preferential attachment network is compared with a degree two regular network in Figure 2.

2.4. Truncated Power Law Network


Although many real-world networks exhibit power-law-distributed degrees with $\gamma = 2$, real-world networks also have practical minimum and maximum degrees. For example, with a power law distribution, degrees of zero and one are most likely, but nodes with degree less than two are not actually on a network. Similarly, there is an effective upper limit on the number of friends an individual can have. Truncating a preferential attachment distribution at upper and lower bounds can yield a distribution with $\gamma = 2$ but a mean degree much greater than two. The truncated power law implemented for this paper is a preferential attachment network where nodes with degree within the truncation range are selected randomly with power-law probability and given an additional edge to another node. This continues until a target mean degree is reached. The outcome space for a degree four truncated power law network is compared with a degree four regular network in Figure 3.

2.5. Small World Network


The small world model (Wilensky, 2005b) is included in the NetLogo (Wilensky, 1999) demo library and is based on an approach by Watts and Strogatz (1998). This is a Bernoulli random network with some of the edges reconnected to create larger hubs and increase clustering.



(a) Strategic substitute, 2-regular network.



(b) Strategic substitute, 2-powerlaw network.



(c) Strategic complement, 2-regular network.







Figure 2: Comparison of degree two regular and power law networks.




(a) Strategic substitute, 4-regular network.



(b) Strategic substitute, 4-truncated network.



(c) Strategic complement, 4-regular network.



(d) Strategic complement, 4-truncated network.

Figure 3: Comparison of degree two regular and power law networks.

3. The ABM Simulation

Following the the numerical models in Lamberson (2011), the NetLogo models update a single, randomly selected agent at each time step. This random sampling means that, for a network with 500 nodes, in the first 500 time steps, some agents may not be updated at all, and others may be updated more than once.

These are the steps in a simulation:

1. Create 500 agents, then create a network by connecting the agents as nodes in a network of the selected topology.
2. Randomly assign agents an initial strategy based on the selected game: strategic complement or strategic substitute. Each run involves ten simulations, each with a different initial strategy distribution:
 - a) For strategic substitutes, starting initial distributions of the fraction of agents playing strategy x are 10 percent through 100 percent in steps of 10 percent.
 - b) For strategic complements, starting initial distributions of the fraction of agents playing strategy x are 10 percent through 55 percent in steps of 5 percent.
3. Each time step, a node is selected at random and that node selects a strategy based on current payoffs based on the node's friends. This may be the same as the strategy already being played.
4. Each simulation proceeds for 4000 time steps, except as noted.

Each game has an average of ten players: the current player and its ten friends. Each player, in its turn, chooses between the two strategies, x or y , based on a payoff computed from the strategy currently played by the friends. If the player had played x in the past, but now chooses y , some of the friends may face a different payoff when their turns come around, and some may switch strategies. Eventually, however, for the games described in Section 1, the entire network - every player and all its friends - reaches a point where there's no longer a reason to switch strategies, ending in a static equilibrium ratio of those playing x and those playing y .

3.0.1. Degree 10 Regular Random Network.

Plots of the ABM results with a degree 10 regular random network are shown in Fig. 4. The equilibrium for strategic substitutes is between 46.6 percent and 47.9 percent playing x . For strategic complements, the all-in or all-out division is between 25 percent and 30 percent playing x .

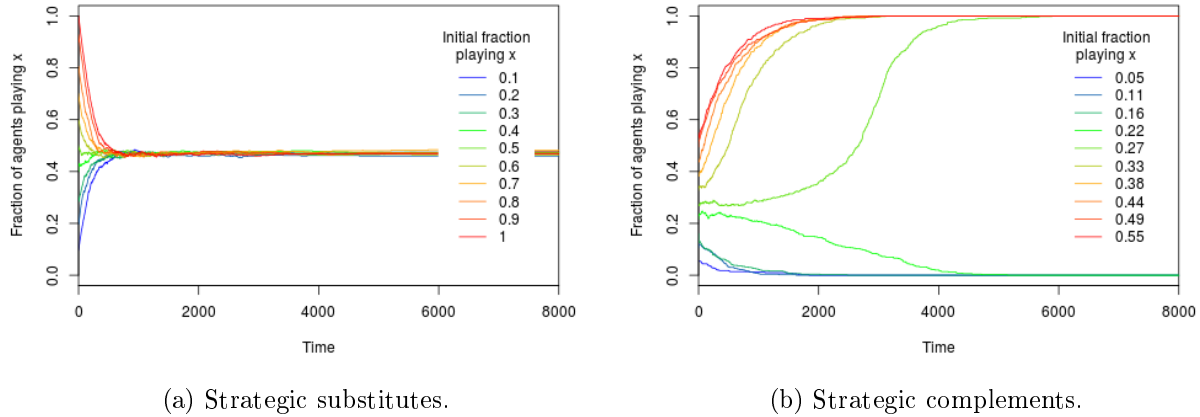


Figure 4: Degree 10 Regular random network.

3.0.2. Degree 10 Bernoulli Random Network.

Plots of the ABM results with a degree 10 Bernoulli random network are shown in Fig. 5. This is an Erdős-Rényi random network, but the results for a Gilbert random network are effectively identical. The equilibrium for strategic substitutes is between 53.0 percent and 54.3 percent playing x . For strategic complements, the all-in or all-out division is between 20 percent and 25 percent playing x .

4. The Global Strategies: Sheep, Contrarians, and Saboteurs

In the games described in Section 1, agents play a strategy based solely on payoffs based on the strategies played by their subnetwork of friends. Eventually, all subnetworks reach a Nash equilibrium and the entire network comes to a static equilibrium. To this game now are added players with payoffs based, not on the subnetwork of friends, but on the entire network. The contrarian payoff is positive for playing the opposite strategy as the majority on the global network. Galam (2004) identifies contrarians as a crucial factor in close elections. The sheep payoff is positive for playing the same strategy as the majority of the players on the global network. Sheep are associated with the bandwagon effect, which Callander (2007) associates with a specific type of voter. The saboteur payoff is positive for playing against the majority when the majority is large, and the saboteur payoff is positive for playing with the majority when the majority is small. The saboteur payoff is a conditional sum of the contrarian and sheep payoff, and is included here merely to add a simple adaptive behavior.

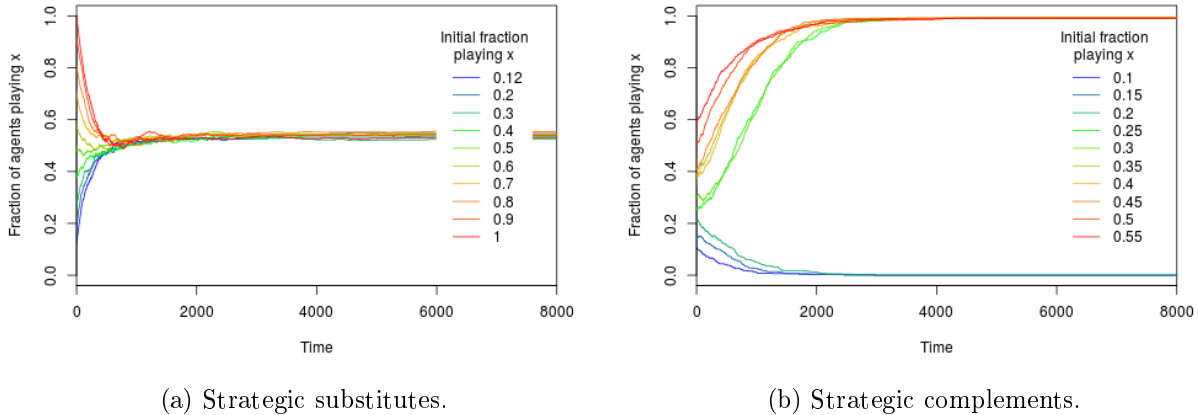


Figure 5: Degree 10 Bernoulli random network.

The introduction of global strategies has a profound effect on the nature of equilibrium outcomes described in Section 3. The single static equilibrium outcome for the strategic substitute payoff becomes two or more equilibrium outcomes. The bimodal static equilibria seen with the strategic complement payoff converge to one, similar to the outcome with power law topologies, or sometimes never reach equilibrium. As the number of saboteurs is increased, previously static equilibria become dynamic, with agents switching strategies frequently.

The following models were simulated using stochastic excursions along the parameters for the percent of the agents facing each global payoff, ranging from 0 to 33 percent. That is, the three-dimensional outcome space is explored with an ensemble of $34 \times 34 \times 34 = 39,304$ simulations for each subnetwork payoff on each network topology. Each simulation is 8000 time steps (ticks) for each of ten initial values of the percent of agents playing strategy x . The completion of an ensemble for each of the two subnetwork payoffs - strategic complement and strategic substitute - for five different network topologies - regular, Bernoulli, power law, truncated power law, and small world - means that there were 393,040 networks constructed and 3,930,400 simulations for a total of more than 31 billion time steps.

4.1. The Strategic Substitute Payoff with Global Payoffs

The consistent characteristic of the strategic substitute payoff - the single equilibrium outcome - persists over various network topologies as shown in Section 2. This feature disappears with only a small increase in the percent of players subject to one of the global payoffs. The single equilibrium outcome is shown in blue in Figure 6. This is the three-dimensional space of percent contrarians, percent saboteurs, and percent sheep. The red




Figure 6: Modes of the strategic substitute payoff on a degree ten Bernoulli random network. Blue indicates regions with single-mode outcomes. Red indicates regions with bi-modal outcomes.

regions illustrate the regions of parameter space which lead to bimodal outcomes. The boundary between the unimodal and bimodal phases is shown in Figure 7. The curvature of this surface shows systematic variation along isoquants of percent saboteurs and percent sheep, for example. The general roughness of the surface is evidence of a large degree of stochastic variation.

Figure 8 plots the outcomes of all 39,304 simulations of the strategic substitute payoff on a degree ten Bernoulli random network. Systematic variation with increasing percent contrarians, percent saboteurs, and percent sheep are evident. The characteristic unimodal outcome persists for all levels of sheep as long as percent contrarians is high and percent saboteurs is low, as shown in the upper-left of the close-ups in Figure 9. For low levels of both contrarians and saboteurs (lower left close-up), the single equilibrium at low levels of sheep appears to switch to a lower equilibrium before switching to bimodal outcomes at higher levels of sheep. A similar but different pattern emerges for high levels of both contrarians and saboteurs (upper-right close-up), while bimodal outcomes are all but universal for low levels of contrarians and high levels of saboteurs (lower-right close-up).

The lower left plot in the lower left close-up in Figure 9 corresponds to the non-perturbed strategic substitute payoff with increasing levels of sheep. The outcome shifts from a single equilibrium outcome to two at 19 percent sheep. This is shown in the simulation plots in Figure 10, with the outcomes shown in the top plots and the activity - the number of agents changing strategy as a function of time - below. The activity plots show that these are both static equilibria.




Figure 7: Phase boundary along the sheep dimension of the strategic substitute payoff on a degree ten Bernoulli random network. Side view and top view. For each value of percent contrarians and percent saboteurs, the surface indicates the value of sheep at which the outcome transitions from one mode to two. The roughness of the surface is a result of stochastic variation among samples. The plot on the right is a top view. Uncolored regions represent outcomes which are unimodal for all values of sheep. Colors represent sensitivity to the sheep parameter, from blue for low sensitivity, to red for high sensitivity.



Figure 8: All simulation outcomes for strategic substitute payoff on a degree ten Bernoulli random network. Each plot has 34 points on the horizontal axis, corresponding to percent sheep from 0 to 33, left to right. Black diamonds represent unimodal outcomes. For multi-modal outcomes, the minimum and maximum mode values are connected with a yellow line, with a blue square at the lower mode, a red circle at the upper mode, and green diamonds at any intermediate modes.

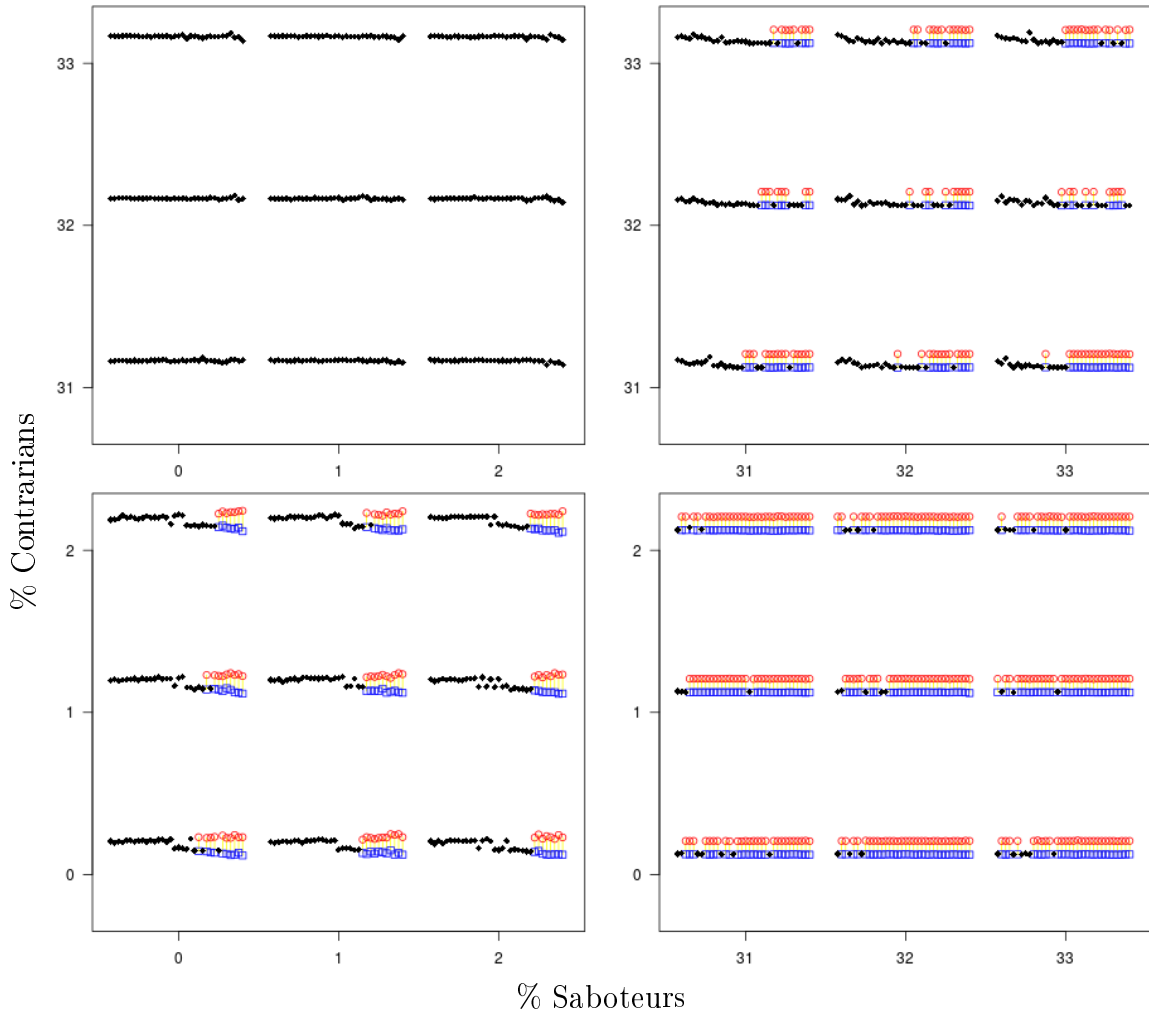






Figure 9: Four corner outcomes for strategic substitute payoff on a degree ten Bernoulli random network. Each plot has 34 points on the horizontal axis, corresponding to percent sheep from 0 to 33, left to right. Black diamonds represent unimodal outcomes. For multi-modal outcomes, the minimum and maximum mode values are connected with a yellow line, with a blue square at the lower mode, a red circle at the upper mode, and green diamonds at any intermediate modes.



(a) Modes - 0 percent contrarians, 0 percent saboteurs, 18 percent sheep.

(b) Modes - 0 percent contrarians, 0 percent saboteurs, 19 percent sheep.



(c) Activity - 0 percent contrarians, 0 percent saboteurs, 18 percent sheep.

(d) Activity - 0 percent contrarians, 0 percent saboteurs, 19 percent sheep.

Figure 10: Bernoulli substitute with mode split at 19 percent sheep.



High levels of saboteurs lead to dynamic equilibria, where the outcome is effectively constant but agents are still changing strategies. Two such cases are shown in Figure 11. The simulation in the right plots, with 33 percent saboteurs, is allowed to run to 20,000 time steps to further illustrate the dynamic nature of these equilibria.

4.2. The Strategic Complement Payoff with Global Payoffs

The outcome space in terms of number of equilibria (modes) is considerably more complicated for the strategic complement payoff compared with the strategic substitute payoff, as shown in Figure 12. The characteristic bimodal outcome seen in Figures 4b and 5b, represented in green in Figure 12, is rare in this outcome space. The large blue volume in this plot shows the dominance of unimodal outcomes similar to the strategic complement outcomes for the degree two topologies in Figure 2. In the region where bimodal outcomes dominate, at low levels of contrarians and saboteurs - there can be three, four, or five outcome modes. Further investigation shows that these are not all equilibria. The complexity of the phase space for this payoff is shown by the phase boundaries in Figure 13. The roughness of these surfaces implies that stochastic effects dominate specific outcomes, but there is clearly some structure evident in the top view plots. Outcomes with more than two modes are restricted to low levels of contrarians and saboteurs, but unimodal outcomes are excluded at the very lowest levels of contrarians and saboteurs (the white area at the far left of Figure 13a). The boundary of this region occurs at a constant sum of percent contrarians and percent saboteurs (the red border of the white area). There is another band of red, indicating that unimodal outcomes are unlikely, where the levels of contrarians and saboteurs sum to just under 33 percent.

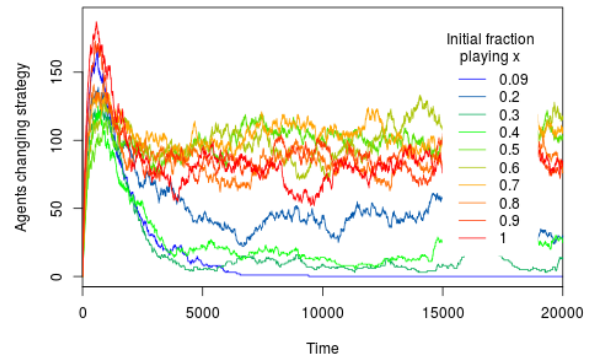
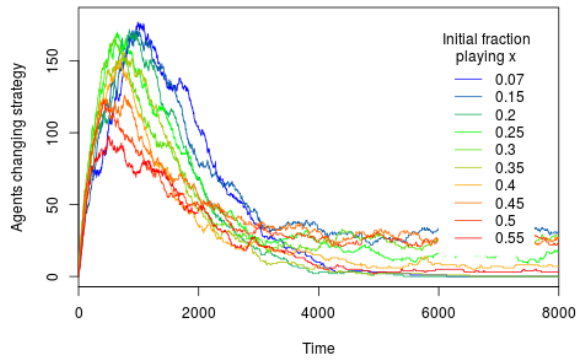
Figure 14 plots the outcomes of all 39,304 simulations of the strategic complement payoff on a degree ten Bernoulli random network. Systematic variation with increasing percent contrarians, percent saboteurs, and percent sheep are evident. The characteristic all-in/all-out outcome dominates at low levels of contrarians and saboteurs, clear from the preponderance of yellow in the lower left corner of Figure 14. The lower-left close-up in Figure 15 reveals that, in addition to the all-in and all-out modes, for non-zero levels of sheep there are also intermediate values, shown as green triangles. It will be shown that some of these are intermediate equilibria, while others represent configurations that never reach equilibrium. The appearance of black diamonds in Figure 13 shows that the red boundary of the white region seen in Figure 13a occurs where levels of contrarians and saboteurs sum to nine percent. Evident also from the overall view in Figure 14, below the major diagonal band of red in Figure 13a, all-in/all-out bimodal outcomes are likely, while above the diagonal bimodal outcomes are closely spaced, with all-in/all-out equilibria never occurring. The dominance of unimodal outcomes at low saboteurs and high contrarians at the top in Figure 13a are evidenced by the black diamonds in the upper right corner of the upper left plot in Figure 14.

Figure 16b shows an outcome for a strategic complement payoff that is remarkably similar to the characteristic outcome for a strategic substitute payoff. The characteristic outcome



(a) Modes - 0 percent contrarians, 10 percent saboteurs, 0 percent sheep.

(b) Modes - 0 percent contrarians, 33 percent saboteurs, 0 percent sheep.



(c) Activity - 0 percent contrarians, 10 percent saboteurs, 0 percent sheep.

(d) Activity - 0 percent contrarians, 33 percent saboteurs, 0 percent sheep.

Figure 11: Bernoulli substitute with dynamic equilibrium.




Figure 12: Modes of the strategic complement payoff on a degree ten Bernoulli random network. Green indicates regions with (expected) bimodal outcomes. Blue indicates regions with a single mode, while yellow, orange, and red indicate outcomes with three, four, and five modes, respectively.

for the strategic complement payoff is shown in Figure 16a, and below it, in Figure 16c, the corresponding activity: the number of agents changing strategy as a function of time. All activity ceases by time step 5000, so these are all clearly static equilibria. Figure 16d shows that, although the outcomes in Figure 16b may be equilibria, they are clearly not static. The large number of contrarians are able to keep the network from ever reaching a non-perturbed equilibrium, while the saboteurs are able to keep the activity churning.

Of particular interest are the outcomes with more than two modes: are there actually multiple equilibrium outcomes or is the network still in flux at the end of the simulation? Two examples are shown in Figure 17. Figure 17a corresponds to the small filled-in region at the left in Figure 13c. The plot below in Figure 17c, shows the corresponding activity. Clearly there are three static equilibrium outcomes. Figure 17b appears in the phase cube as a two-mode outcome because of the closeness of the upper equilibria, but a casual glance notes six equilibria. Figure 16d attests that these are static equilibria.

5. Discussion

The details of establishing correspondence of the ABM to Lamberson (2011) are given in Dixon (2011b). The impacts of network topology are explored in Dixon (2011a). The discussion in Section 4 considers only a Bernoulli random network. The results for other

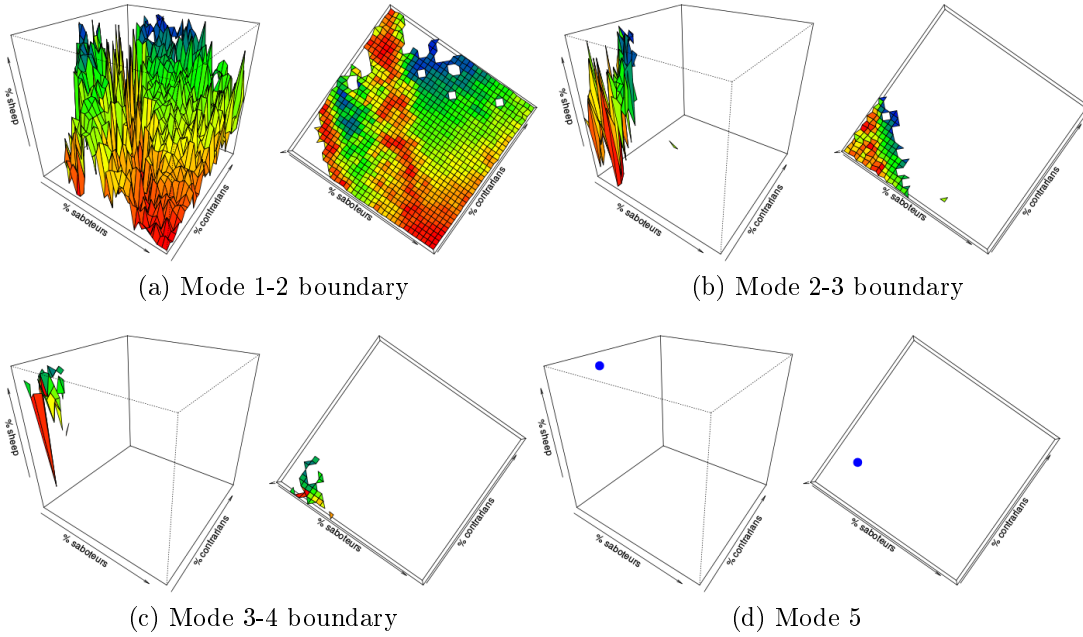


Figure 13: Phase boundaries along the sheep dimension of the strategic complement payoff on a degree ten Bernoulli random network. Side view and top view. For each value of percent contrarians and percent saboteurs, the surface indicates the value of sheep at which the outcome transitions from one mode to the next. Figure (a) is the phase boundary between the first mode and the second. Figure (b) is the boundary between the second mode and the third. Figure (c) is the boundary between the third mode and the fourth. Figure (d) shows the point (7% contrarians and 4% saboteurs) at which there is a single fifth mode outcome.



Figure 14: All simulation outcomes for strategic complement payoff on a degree ten Bernoulli random network. Each plot has 34 points on the horizontal axis, corresponding to percent sheep from 0 to 33, left to right. Black diamonds represent unimodal outcomes. For multi-modal outcomes, the minimum and maximum mode values are connected with a yellow line, with a blue square at the lower mode, a red circle at the upper mode, and green diamonds at any intermediate modes.

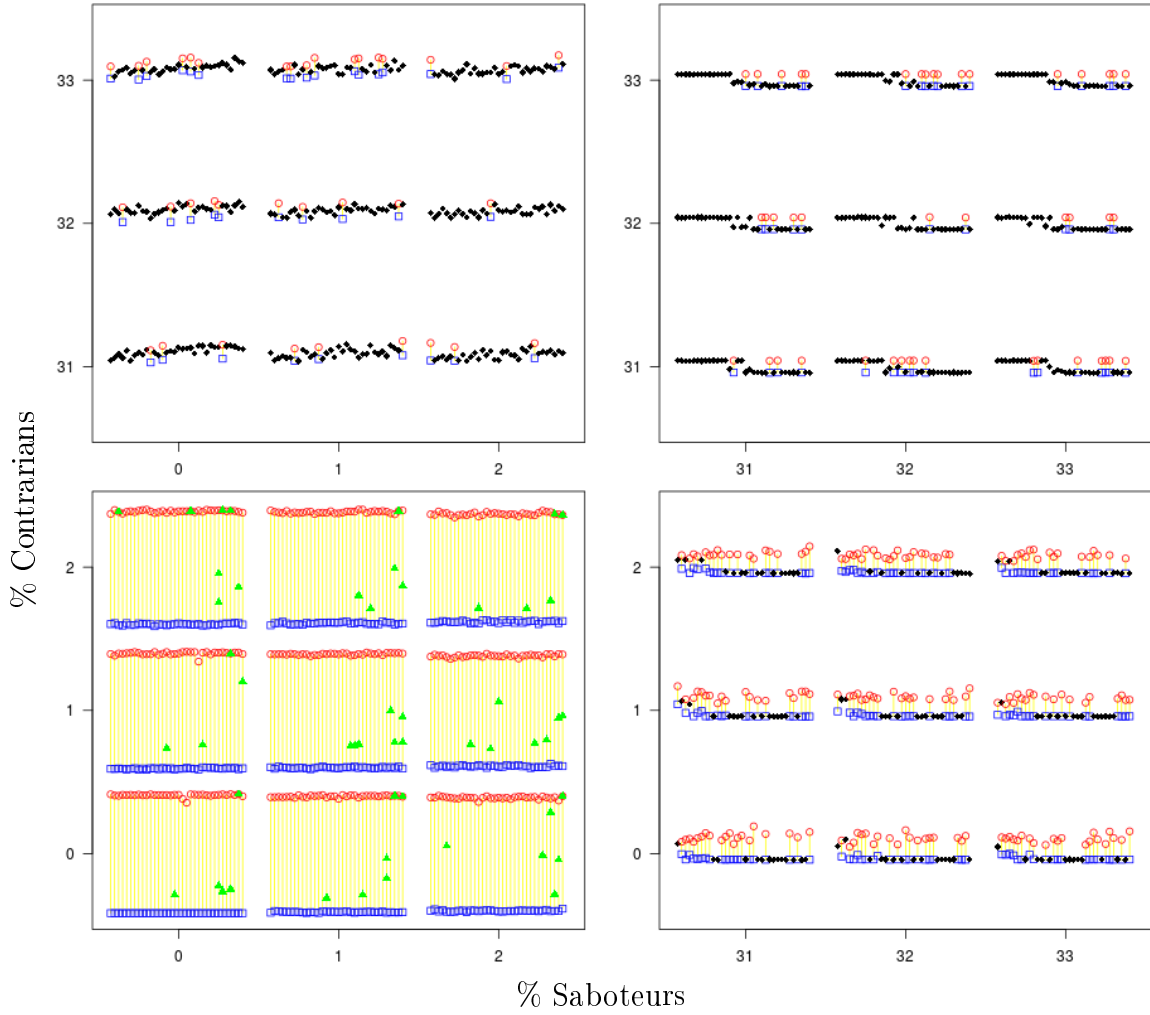




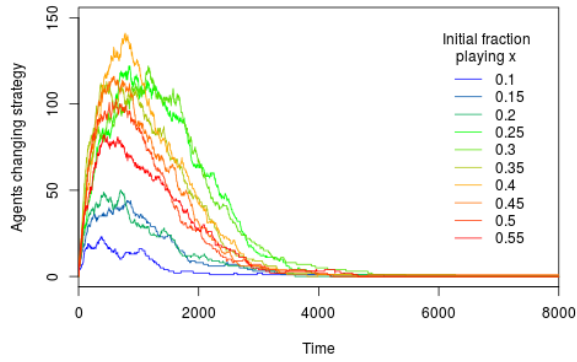
Figure 15: Four corner outcomes for strategic complement payoff on a degree ten Bernoulli random network. Each plot has 34 points on the horizontal axis, corresponding to percent sheep from 0 to 33, left to right. Black diamonds represent unimodal outcomes. For multi-modal outcomes, the minimum and maximum mode values are connected with a yellow line, with a blue square at the lower mode, a red circle at the upper mode, and green diamonds at any intermediate modes.




(a) Modes - 0 percent contrarians, 0 percent saboteurs, 0 percent sheep.



(b) Modes - 33 percent contrarians, 10 percent saboteurs, 0 percent sheep.




(c) Activity - 2 percent contrarians, 2 percent saboteurs, 33 percent sheep.




(d) Activity - 33 percent contrarians, 10 percent saboteurs, 0 percent sheep.


Figure 16: Bernoulli complement with dynamic equilibrium.




(a) Modes - 2 percent contrarians, 2 percent saboteurs, 33 percent sheep.



(b) Modes - 4 percent contrarians, 26 percent saboteurs, 8 percent sheep.



(c) Activity - 2 percent contrarians, 2 percent saboteurs, 33 percent sheep.



(d) Activity - 4 percent contrarians, 26 percent saboteurs, 8 percent sheep.

Figure 17: Bernoulli complement with multiple modes.

network topologies are notable, but not qualitatively different in the space of contrarians, saboteurs, and sheep. The mode spaces for other topologies are shown in the Appendices.

That saboteurs only have a distinct effect on outcomes when majorities are small is not surprising, given their payoff function. The asymmetry of outcomes based on levels of sheep compared to contrarians is, at first, surprising given the apparent symmetry of their payoffs. The asymmetry arises temporally: sheep can vote early and often while the local payoff function has agents swapping strategies.

Sheep have an early influence on pushing the outcome toward the initial configuration. When the initial configuration is close to the non-perturbed equilibrium, local-payoff players are decisive and work with the sheep from the beginning. When the initial configuration is far from the non-perturbed equilibrium, local-payoff players are swapping strategies in the beginning and present weak resistance to sheep from the start. A majority position is established quickly even when it is the opposite of the non-perturbed equilibrium

Contrarians, on the other hand, have an early influence on pushing the outcome away from the initial configuration. When the initial configuration is far from the non-perturbed equilibrium, local-payoff players are swapping strategies but slowing working toward the early preference of contrarians. When the initial configuration is close to the non-perturbed equilibrium outcome, the local-payoff players are decisive and present nearly unified opposition to contrarians from the start. A majority position is established slowly and, when it is, contrarians are faced with changing strategies at a time when local-payoff players are most decisive.

6. Future Work

This paper is the first along an avenue of study examining intentional network disruption. Now that the overall outcome spaces have been scanned with stochastic excursions, the next step will be to look more closely at some regions using Monte Carlo sampling. This will provide a more detailed view of equilibrium and non-equilibrium outcomes in the higher-dimensional space of interplay between local payoffs, global payoffs, initial configurations, network size, and network topology. Further down this road lie examinations of more complex adaptive behaviors.

Other works in progress include investigation of payoffs in a friendship game that are not constrained to closed-form mathematical functions and can incorporate adaptive behaviors such as learning and heuristics. These could enable the construction of models of voters, economic agents, or decision-makers in which the payoff (or utility or fitness) depends on the preferences of multiple groups of friends over multiple conflicting issues. For example, voters may be influenced by a workplace network on issues relating to their livelihood, and by a very different social network on other topics. Another avenue of research is in modeling market models from the production point of view. For example, a hybrid networked market where each producer shares one kind of network with consumers, another kind of network with suppliers, and yet another kind of network with competitors. A third

avenue is exploring evolving network topology, such as emerging social movements, natural disasters, and economic upheaval.

7. Conclusion

The stable equilibria in friendship games established by Lamberson (2011) are intriguing. This result implies that for the strategic complement and strategic substitute payoffs examined, a global static equilibrium results from local Nash equilibria. A game theoretical example of Adam Smith's invisible hand, perhaps. A provision of public goods determined locally by a neighborhood of friends as opposed to voting with their feet as suggested by Charles Tiebout. The establishment of technical standards without the inefficiencies of long marketplace battles and stranded early adopters.

Putting aside the fact that real-world issues are never as simple as our payoff models here, it is instructive that these static equilibria are disrupted by the mere existence of players for whom the payoff is not subgame optimal. The strategic substitute simulations show that the provision of public goods can be shifted away from social optimality simply because many players side with the majority before a true majority has been established, or when many players always oppose the majority preference. The strategic complement simulations show that socially optimal all-in or all-out outcomes are shifted and sometimes never reach equilibrium as a result of non-subgame-optimal payoffs. This could be socially disruptive, but it could also be seen as a good thing: disruptive opposition to extreme social or political agendas can prevent calamitous outcomes even when the majority is hellbent on evil.

8. References

References

- Albert, R. and A. Barabási (2002). Statistical mechanics of complex networks. *Reviews of modern physics* 74(1), 47.
- Barabási, A. and R. Albert (1999). Emergence of scaling in random networks. *Science* 286(5439), 509.
- Callander, S. (2007). Bandwagons and momentum in sequential voting. *The Review of Economic Studies* 74(3), 653–684.
- de Solla Price, D. (1965). Networks of scientific papers. *Science* 149(3683), 510.
- Dixon, D. S. (2011a, October). An Agent-Based Adaptation of Friendship Games: Observations on Network Topologies. Computational Social Sciences Society of the Americas: Second Annual Conference of the Computational Social Sciences Society of the Americas.

- Dixon, D. S. (2011b, June). Preliminary results from an agent-based adaptation of friendship games. Western Economics Association International: 86th Annual Conference of the Western Economics Association International.
- Erdős, P. and A. Rényi (1959). On random graphs, I. *Publicationes Mathematicae (Debrecen)* 6, 290–297.
- Galam, S. (2004). Contrarian deterministic effects on opinion dynamics:"the hung elections scenario". *Physica A: Statistical Mechanics and its Applications* 333, 453–460.
- Galeotti, A., S. Goyal, M. Jackson, F. Vega-Redondo, and L. Yariv (2010). Network games. *Review of Economic Studies* 77(1), 218–244.
- Gilbert, E. (1959). Random graphs. *The Annals of Mathematical Statistics* 30(4), 1141–1144.
- Lamberson, P. (2011). Friendship-based Games. 37th Annual Conference of the Eastern Economics Association.
- Lamberson, P. (2015). Network games with local correlation and clustering. *Social Science Research Network*.
- Watts, D. J. and S. H. Strogatz (1998). Collective dynamics of 'small-world' networks. *nature* 393(6684), 440–442.
- Wilensky, U. (1999). Netlogo. Accessed 31 August 2011.
- Wilensky, U. (2005a). Netlogo preferential attachment model. Accessed 31 August 2011.
- Wilensky, U. (2005b). Netlogo small worlds model. Accessed 19 October 2014.

A. Other Network Topologies

See <http://www.unm.edu/~ddixon/Contrarians> .