Scribe Notes: Privacy and Economics

Hyoungtae Cho
hcho5@cs.umd.edu

Jay Pujara
jay@cs.umd.edu

Naomi Utgoff
utgoff@econ.umd.edu

12/1/2010

Abstract

The connections between Game Theory and the subjects of Privacy and Economics are an active field of research. Some topics considered by game theorists include valuing private information, and modeling the behavior of non-standard utility functions (eg. altruists, malicious players) in classic settings such as auctions and congestion games. We present the dynamics of players in a virus inoculation game as an example of current research in this area. Our independent research extends this work using a mechanism called *imitation dynamics* and analyzes the interplay of altruistic and selfish players. We show analytical bounds and initial conditions that result in the spread of altruism.

1 Background

1.1 Motivations

Privacy in the colloquial sense has many different meanings; the term is applied to everything from heavy window curtains to the defense of abortion rights. The goal of game theorists and economists is to provide a more formal setting to study privacy. One way of formalizing the question of privacy is to assign a value to private information an agent possesses. Specifically, if sharing private information can help achieve a social optimum, how much should each agent be compensated for contributing his or her information?

A more game theoretical approach to the question of privacy is to consider non-standard utility functions, or cases where player's utilities are interdependent. Two such cases are altruism, where agents seek to maximize the utility of other players, and malice, where players try to minimize the utility of other players. A related concept is the *Stackelberg Threshold*, or the fraction of agents that must be controlled by a central authority to achieve some desired social outcome. The impact of these utility functions is evaluated in common game theoretic scenarios such as congestion games, network formation, and auctions.

Popular metrics used in these studies are the "Price of Malice" or the ratio of social cost with malicious players to the Price of Anarchy, as well as the price of altruism, the ratio of the social cost with altruistic players to the Price of Anarchy. In some cases, adding malicious players can improve the social cost, and this situation is called "the Windfall of Malice". In addition to studying these quantities, game theorists try to characterize the existence of equilibria, convergence of games, hardness of computing equilibria, as well as understand how to better design mechanisms to converge to a social optimum.

1.2 Key Findings

We present findings from three papers representing current research in the field of Privacy and Economics.

On the Value of Private Information

Kleinberg et al. seek to characterize the value of information using the Shapely Value [4]. Specifically, the calculate the value of agent i's information x_i by looking at the average contribution of i's information over all $(n!)$ possible arrival orders (π).

$$x_i = \frac{1}{n!} \sum_{\pi \in S_n} v(S(\pi, i)) - v(S(\pi, i) - \{i\})$$

This contribution of private information is evaluated in three scenarios: a marketing survey, recommendation systems, and collaborative filtering. An intuitive result from this analysis is that the agents who benefit from sharing information depends greatly on the setting. In the case of a marketing survey, where a company seeks to offer a new product, the agents who are in the majority benefit from sharing their preferences. Conversely, when attempting to recommend interesting products, those agents with novel preferences are best-compensated.

Congestion Games with Malicious Players

In [1], Babaioff et al. present the scenario of a congestion game where some fraction ϵ of flow is controlled by a malicious player. Rather than

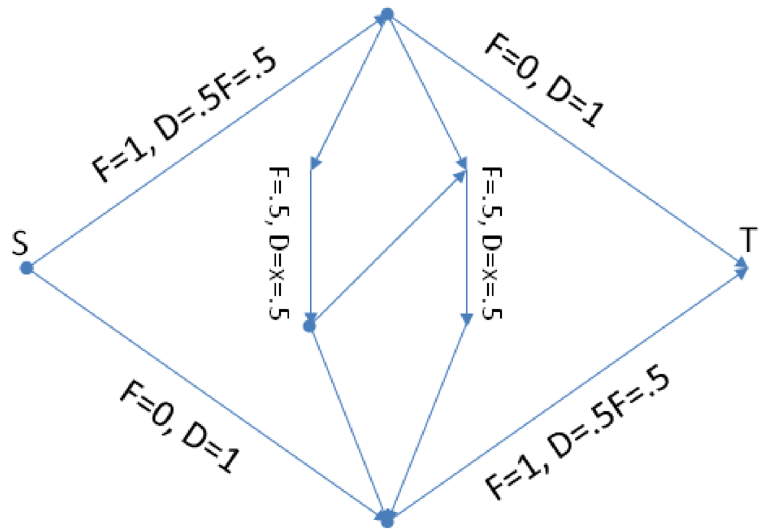

Figure 1: Nash flow with no malicious flow

using the usual definition of Price of Malice, this paper uses a different computation - the change of the Nash delay in graph G with total flow v as the malicious flow approaches 0, or:

$$lim_{\epsilon\to 0}\frac{D(G,\epsilon v)-D(G)}{\epsilon D(G)}=\frac{\Delta^{\epsilon}_{\text{delay}}}{\epsilon\cdot\text{delay}}$$

In some cases they show that malicious players seeking to increase the latency across edges maximally can cause selfish players to make choices that decrease the total delay, which they term the "Windfall of Malice". As seen in Figure 1, players suffer a delay of 1.5 by choosing the upper path, splitting over the two downward connections, and the combining in the last link to the destination. A malicious player can route their flow over all variable cost links (upper path, Z connector, lower path) to maximally increase the delay on each link. However the Nash equilibrium of selfish players, shown in Figure 2, results in a delay of 1.483 which is lower than the Nash delay, even with additional malicious flow added to the network.

Babaioff et. al present some interesting results on the problem of malicious flow in congestion games. They prove that in such games, an equilibrium always exists (given strictly increasing delay functions). For certain classes of delay functions (continuous, weakly concave) they show that a

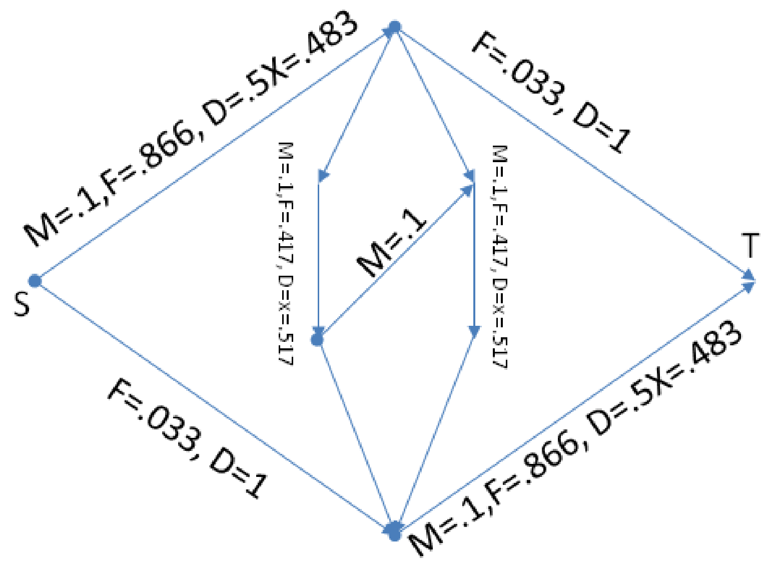

Figure 2: Nash flow with malicious flow of .1

pure equilibrium exists. They also prove bounds for both the price of malice and the windfall of malice. The upper bound on the price of malice for a game with m edges and relative slope d (defined in terms of delay function l, as $d = sup_x \frac{xl'(x)}{l(x)}$), of d(m-1). They also prove that for a game with m edges, the windfall of malice is $-\frac{m^2}{2(m+2)}$. Open problems in this area include proving hardness and convergence results, characterizing games that can result in a windfall of malice, and proving a lower bound on the price of malice.

Spiteful Bidding in Sealed-Bid Auctions

Brandt et al. have explored malicious players in auctions with both complete information and private information [2]. For their analysis, they assume each agent has a utility function of the form $U_i = (1-\alpha)u_i - \alpha \sum_{i \neq j} u_j$, where the utility is a sum of the agent's valuation and the negative utility of the other agent's utilities. The parameter α is called the "spite" parameter. When $\alpha = 0$, the usual results for selfish agents hold, while $\alpha = 1$ models purely malicious players. Considering the case of sealed-bid auctions, they model the valuation of agents using a distribution and prove their results using the Bayes Nash Equilibrium. They find that, surprisingly, a first-price sealed-bid spiteful auction is truthful when players are malicious. They also show that as the spite parameter α increases, the revenue of the auction also increases. Finally, they compare the case of private information (sealed-bid auctions) to a situation where complete information is available. They find that the revenue of a first-price auction increases when private information is revealed, however second-price auctions have a decreased revenue relative to the sealed-bid auction.

2 Inoculation Game

One area of particular interest to our group was the behavior of players when their interactions are governed by the topology of the game. We present a paper by Moscibroda et al. that examines the behavior of a network facing a virus attack [5].

2.1 Model Setup

The inoculation game consists of n players on a grid, where the neighbors of each player are the fellow players in each of the four cardinal directions (eg. up, down, left, right). Players are either inoculated, and secured against virus attacks, or insecure and susceptible to damage if infected by a virus. Each player will choose to inoculate at a cost of 1 or risk suffering a loss of L should they be attacked by a virus. At some point a virus will infect a random node. A player infected by a virus will spread the virus to each of his insecure neighbors, and those neighbors, in turn, will spread to their insecure neighbors. As a result any insecure member of a connected component of insecure nodes will spread the infection to all nodes in the component. Note that the game is only interesting when $L < n$, otherwise everyone will inoculate, and so we will only consider that case in the discussion.

The social cost of the game can be expressed as the sum of the inoculation costs and the infection costs. Since the cost of inoculation is 1, the inoculation costs are simply the number of inoculated nodes. To express the infection costs requires considering each component. The probability of a component of size k will be infected is $\frac{k}{n}$. In equilibrium, we expect all components to have the same size, K. Note that in equilibrium, the inoculation cost must equal the expected cost of infection, so $1 = \frac{K}{n}L$, and the size of each component is $\frac{n}{L}$ (but the optimal conditions favor smaller components to minimize the total expected cost). If γ nodes are secured, the remaining $n-\gamma$ nodes must be insecure. The probability of infection for any component is $\frac{K}{n}$ and the loss suffered for each node is L. Since all nodes are in equal-sized components, we assume they have the same probability of

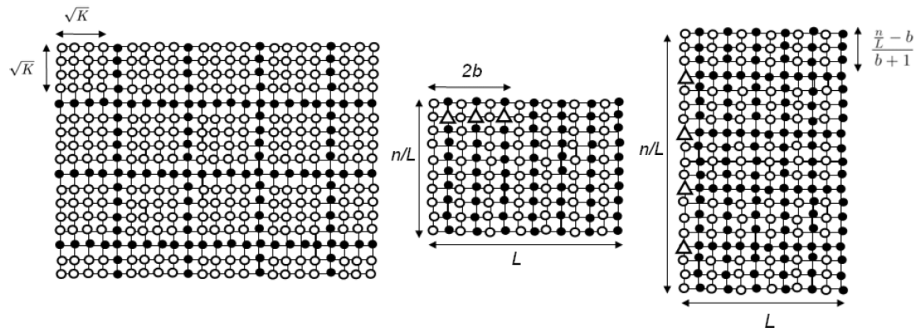

Figure 3: Examples of the socially optimal upper bound, a Nash equilibrium, and the result when players fear a group of malicious players

suffering loss L. Hence:

$$\text{Cost} = \text{Inoculation Cost} + \text{Total Infection Cost}$$

$$\text{Cost} = \text{Inoculated Nodes} + \sum_{\text{components}} P(\text{Infection}) \cdot \text{Size} \cdot \text{Infection Cost}$$

$$\text{Cost} = \text{Inoculated Nodes} + \sum_{i \in \text{components}} \frac{\text{Size}_i}{n} \cdot \text{Size}_i \cdot L$$

$$\text{Cost} = \gamma + (n - \gamma) \cdot \frac{K}{n} \cdot L$$

2.2 Social Optimum

To derive the bounds for the social optimum, we start by finding the bounds for γ. The upper bound for γ is $\frac{n}{2}$, in the case that alternating diagonal stripes of nodes are secure, the remaining nodes will have no insecure neighbors. For the lower bound of γ, we consider that all insecure components are circles of size K. The radius of such a circle would be $r = \sqrt{\frac{K}{\pi}}$. The circumference of each circle would be secure nodes, so that each component would have $2\pi\sqrt{\frac{K}{\pi}} = 2\sqrt{K\pi}$ secure nodes on its boundary. Since each secure node can border at most two components, we must have at least $\sqrt{K\pi}$ secure nodes per component. How many such components are there? We can divide the insecure nodes by the component size, $\frac{n-\gamma}{K}$. Solving $\gamma \geq \frac{n-\gamma}{K} \cdot \sqrt{K\pi} = (n-\gamma)\sqrt{\frac{\pi}{K}}$ results in $\gamma \geq n \cdot \frac{\sqrt{\frac{\pi}{K}}}{1+\sqrt{\frac{\pi}{K}}}$.

Now we can derive the lower bound for the social optimum. Using the lower bound of γ for the inoculation cost and the upper bound of γ in the infection cost, we can derive the lower bound for the optimal social cost:

$$\text{Cost}_{\text{opt}} \geq n \cdot \frac{\sqrt{\frac{\pi}{K}}}{1+\sqrt{\frac{\pi}{K}}} + (n - \frac{n}{2})\frac{LK}{n}$$

We can also solve the equation to find the optimal K under these lower bound assumptions, arriving at $K = \frac{2}{3}\sqrt{\pi} \cdot (\frac{n}{L})^{\frac{2}{3}}$. Plugging this back into the cost equation, the bound for cost is

$$\text{Cost}_{\text{opt}} \geq \frac{1}{3}\sqrt{\pi} \cdot n^{\frac{2}{3}} L^{\frac{1}{3}}$$

Using a similar technique, but assuming that all components are squares, we can find an upper bound of $\text{Cost}_{\text{opt}} \leq 4n^{\frac{2}{3}}L^{\frac{1}{3}}$. This scenario is depicted in the leftmost diagram of Figure 3.

2.3 Nash Equilibrium

The Nash equilibrium is relatively straightforward in this game. As mentioned earlier, we know that in equilibrium the expected cost of infection must equal the cost of inoculation, so components will be sized $\frac{n}{L}$. Consider the middle diagram of Figure 3. Each component is of size $\frac{n}{L}$ and half the nodes are inoculated. The inoculation cost is $\frac{n}{2}$ and the expected infection cost is $\frac{1}{2}\frac{n}{L}L = \frac{n}{2}$, yielding a total cost of n. As an upper bound, we know the inoculation cost can be at most n when all nodes are inoculated. Similarly, the infection cost of a component of size $\frac{n}{L}$ is n, so the infection cost is also bounded by n. Hence the upper bound of the Nash equilibrium is $2n$.

2.4 Price of Anarchy

Deriving the Price of Anarchy using the results from the previous two sections is easy. Choosing the lower bound of the Nash equilibrium and the upper bound of the social optimum yields $\frac{n}{4n^{\frac{2}{3}}L^{\frac{1}{3}}}$. Using the upper bound of the Nash cost and the lower bound of the social optimum yields $\frac{2n}{\frac{1}{3}\sqrt{\pi}\cdot n^{\frac{2}{3}}L^{\frac{1}{3}}}$. Simplifying, we have:

$$\frac{1}{4}(\frac{n}{L})^{\frac{1}{3}} \leq \text{PoA} \leq \frac{6}{\sqrt{\pi}}(\frac{n}{L})^{\frac{1}{3}}$$

2.5 Bounds with Malicious Players

Moscibroda et al. extended their analysis to an inoculation game with malicious players. These players could lie about their inoculation state to degrade the outcome of players. Two scenarios are considered: the first where the non-malicious players are unaware of the existence of malicious players (oblivious) and the other where the non-malicious players are aware of the number of malicious players but not their identities (non-oblivious). Using a similar means of analysis, they derive bounds for both scenarios.

Intuitively, in the oblivious case, each malicious player can, at best, connect two connected components of size $\frac{n}{L}$, increasing both the likelihood that the nodes will get infected and the infection cost for the component. At worst, they can saturate the grid causing each insecure player to pay a loss of L. Following this intuition, they find that if the number of malicious players b is less than $\frac{L}{2}-1$ and there are s non-malicious players, the price of malice is $s+\frac{nb^2}{L}+\frac{nb}{L}+b^2$, and sL otherwise. Hence the upper bound of the social cost with is $O(min\{sL, s+\frac{nb^2}{L}\})$. The price of malice for $b<\frac{L}{2}-1$ is $\Theta(1+\frac{b^2}{L}+\frac{b^3}{sL})$ (or $\Theta(L)$ otherwise).

In the non-oblivious case, the analysis is more complex. Using the rationale above, if players know there are more than $\frac{n}{L}$ malicious players, all non-malicious players will inoculate, resulting in a cost of s which is no worse than the Nash cost. Considering the other cases, where $b<\frac{n}{L}$ is more complex. Moscibroda et al. prove that the social cost in such situations is at least $\frac{2s+bL}{4}$ and the Price of Malice is at least $\frac{\sqrt{\pi}}{48}(1+\frac{bL}{2s})$. Note that this implies that in some situations the Price of Malice is below 1, corresponding to the same type of result shown for the Windfall of Malice earlier. Since the cost of secure nodes is fixed, as more insecure nodes choose to inoculate for fear of malicious players, both the probability of an infection striking an insecure node and the loss should such an infection occur diminish.

2.6 Stability with Malicious Players

Another topic considered by the paper was the stability of an equilibrium in the face of malicious players. If an equilibrium can survive with b malicious players, the game is defined to be b-stable. In this case, malicious players can repeatedly change their reported state (secure or insecure) to destabilize the equilibrium. Intuitively, the equilibrium will not be stable with malicious players, since by changing their reported state a malicious player can join two insecure components. Indeed, the inoculation game is not 1-stable in general. However in non-grid scenarios (such as an n-clique), a single malicious node cannot destabilize the equilibrium. However, if there are 2 malicious nodes, they can always destabilize the equilibrium, and hence the inoculation game is always 2-instable.

3 Local Interaction Models

3.1 Motivation

We will expand on the inoculation game by considering a local public good provision game to be repeated for multiple rounds, where players' future decisions determined by what they learned from the previous round. Such a game models several interesting real world phenomena:

- Public goods are often provided on a local scale
- Proximity may determine individual benefit
- Typically, people interact again and again, learning from past interactions

Eshel, Samuelson, and Shaked [3] consider a local interaction model on a circle, where each player has only two neighbors. We will consider a different topology: players located in a grid, choosing to provide (or not provide) a public good, and learning by repeated interaction whether to provide the public good in future rounds.

3.2 Model

The game is played by $m \times n$ players, where each player p_{jk} is the jk node on an $m \times n$ grid, where $j = 1, \cdots, m$ and $k = 1, \cdots, n$. In the initial round, each player chooses a strategy from the set $\{A, E\}$. A player who selects strategy A ("altruist") provides 1 unit of public good, shared equally among his vertical and horizontal neighbors, at a cost $c \in [0, 1]$. If a player selects strategy E ("egoist"), he provides no units of public good and incurs no cost. An altruist does not receive any of the public good he provides, and an egoist still stands to benefit from public good, if any, provided by his neighbors, even though the egoist has not made any contribution to his neighbors.

After each round, each player receives payoff equal to the total public good received from their neighbors (if any) minus the cost of providing public good (if that player provided public good). Also after each round, *all* strategies and payoffs become common knowledge to all players. In subsequent rounds, each player looks at himself and his neighbors, and see whether among that group altruists or egoists did better, on average. If the altruist neighbors of an egoist player had higher utility than the egoist neighbors (self included), the egoist will become an altruist, and vice versa. This process is known as *imitation dynamics*. We will now give several examples to clarify the development of this game.

Note that if the game is played only once, there is a unique Nash equilibrium in dominant strategies, namely all players choose "E", so the system dynamics are crucial to developing a long term population of altruists.

3.3 Initial Conditions and Payoffs

Consider $m \times n = 2 \times 3$, cost $c = 1/10$, and suppose that in the initial round (Round 0), we have the following strategy choices:

$$\begin{array}{ccc} A & A & E \\ E & E & E \end{array}$$

Consider player $p_{1,1}$. Since $p_{1,1}$ is an altruist, he contributes 1 unit of public good (divided evenly between $p_{1,2}$ and $p_{2,1}$) at cost 1/10. He has one altruist neighbor, $p_{1,2}$, who provides 1 unit of public good, divided evenly among $p_{1,1}, p_{2,2}$, and $p_{1,3}$. Therefore, the net payoff to $p_{1,1}$ after round one is $1/3 - 1/10 = 7/30$. Repeating similar calculations for the other players, we see that the payoffs after round 1 are:

$$\begin{array}{rrr} 7/30 & 12/30 & 10/30 \\ 15/30 & 10/30 & 0 \end{array}$$

3.4 Computing the Imitation Dynamics: Learning

Now we want to understand how each player learns from round n to determine what it should do in round $n+1$. Rather than using a Nash analysis, we present the process called "Imitation Dynamics". Consider player $p_{1,1}$.

$$\begin{array}{ccc} A & A & E \\ E & E & E \end{array} \qquad \begin{array}{rrr} 7/30 & 12/30 & 10/30 \\ 15/30 & 10/30 & 0 \end{array}$$

Player $p_{1,1}$ and his neighbor $p_{1,2}$ are altruists, receiving on average $(7/30 + 12/30)/2 = 19/60$. He has one egoist neighbor, $p_{2,1}$, receiving on average $1/2$. Since the egoist did better on average than the altruists (since $1/2 > 19/60$), in the next round $p_{1,1}$ will switch to "E" in the next round.

Repeating similar calculations for the other players, we see that following system evolution:

	Round 0			Round 1
Player	Type	Avg. A Payoff	Avg. E Payoff	Type
$p_{1,1}$	A	0.317	0.5	E
$p_{1,2}$	A	0.317	0.333	E
$p_{1,3}$	E	0.4	0.167	A
$p_{2,1}$	E	0.233	0.417	E
$p_{2,2}$	E	0.4	0.278	A
$p_{2,3}$	E	0	0.222	E
	Round 1			Round 2
Player	Type	Avg. A Payoff	Avg. E Payoff	Type
$p_{1,1}$	E	0	0.389	E
$p_{1,2}$	E	-0.1	0.417	E
$p_{1,3}$	A	-0.1	0.833	E
$p_{2,1}$	E	-0.1	0.167	E
$p_{2,2}$	A	-0.1	0.667	E
$p_{2,3}$	E	-0.1	0.833	E

In this case, the system degenerates to all egoists. Alternatively, we can view the system evolution as follows (with strategies above, and payoffs below):

$$\begin{array}{rrr} A & A & E \\ E & E & E \\ & & \\ 0.23 & 0.4 & 0.33 \\ 0.5 & 0.33 & 0 \end{array} \longrightarrow \begin{array}{rrr} E & E & A \\ E & A & E \\ & & \\ 0 & 0.83 & -0.1 \\ 0.33 & -0.1 & 0.83 \end{array} \longrightarrow \begin{array}{rrr} E & E & E \\ E & E & E \\ & & \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{array}$$

If the cost of providing the public good falls to $c = 1/12 \approx 0.08$, the altruists survive as shown below:

$$\begin{array}{rrr} A & A & E \\ E & E & E \\ & & \\ 0.25 & 0.42 & 0.33 \\ 0.5 & 0.33 & 0 \end{array} \longrightarrow \begin{array}{rrr} E & A & A \\ E & A & E \\ & & \\ 0.33 & 0.75 & 0.25 \\ 0.33 & 0.25 & 0.83 \end{array} \longrightarrow \begin{array}{rrr} A & A & E \\ E & E & E \\ & & \\ 0.25 & 0.42 & 0.33 \\ 0.5 & 0.33 & 0 \end{array}$$

Now it is cheap enough for the altruists to survive. The system will alternate between the first and second states shown above.

An increase in the number of altruists in the initial round may be detrimental to the system in the long run. Suppose that $c = 1/12 \approx 0.08$ and the initial strategy choices are as follows:

$$\begin{array}{ccc} A & A & E \\ E & E & A \end{array}$$

Now, free-riding is more attractive to the egoists at $p_{2,2}$ and $p_{1,3}$, so they will remain egoists in the next round. The system will now degenerate to all egoists instead of the altruists surviving:

$$\begin{array}{ccc} A & A & E \\ E & E & A \\ & & \\ 0.25 & 0.42 & 0.83 \\ 0.5 & 0.83 & -0.08 \end{array} \quad \rightarrow \quad \begin{array}{ccc} E & E & E \\ E & E & E \\ & & \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{array}$$

4 Analysis of Imitation Dynamics

We consider imitation dynamics on the game with general values of the benefit b of a public good and the cost c of producing it. Intuitively, b and c are greater than or equal to zero. Recall that we used b=1 and c=0.1 in the previous chapter. Let us consider larger examples in our local interaction model.

4.1 Larger Examples

Observation 1 Let b be the benefit of a public good provided by an Altruist and c be the cost of producing it that the Altruist should pay. As b gets higher or c gets lower, the Egoist has a greater incentive to be an Altruist in the next round.

Proof An agent decides to its type, either Egoist or Altruist, by computing the average payoff of its neighbors for each type and taking the greater payoff. Since b and c are the only factors in computing payoffs, a change in behavior can only be caused by the values of b or c. We assume the values of these parameters to be independent of each other. Since the cost c is subtracted from the Altruist's payoff but not from the Egoist payoff, decreasing c can improve the incentive for Altruists. For a group of Altruists, a higher value of b will result in higher payoffs, and so increasing b can increase the incentive to join a group of Altruists or remain an Altruist.

Let us take an example that we consider the Egoist that has four neighbors in the center.

$$\begin{array}{ccccc} E & E & E & E & E \\ E & A & \underline{A} & A & E \\ E & \underline{A} & \mathbf{E} & \underline{E} & E \\ E & E & \underline{E} & E & E \\ E & E & E & E & E \end{array}$$

We introduce notation for the *average payoff* for Altruists (A) and Egoists (E), denoted Payoff_{avg}.

$$\text{Payoff}_{avg}(E) = \frac{1}{3} \times (\frac{2}{4}b + \frac{1}{4}b + 0) = \frac{1}{4}b$$

$$\text{Payoff}_{avg}(A) = \frac{1}{2} \times [(\frac{2}{4}b - c) + (\frac{1}{4}b - c)] = \frac{3}{8}b - c$$

To detail this calculation, consider the Egoist highlighted in bold text. This Egoist computes the average of three Egoists including itself as well as the average payoff of its two Altruist neighbors. In this case, the benefit of the public goods offered by each of the underlined Altruists is shared by four neighbors, so that it has value $\frac{1}{4}b$. We compute the average Altruist payoff $\text{Payoff}_{avg}(A)$, based on the payoffs of the two underlined Altruists. The upper Altruist neighbor of the Egoist obtains the benefits from two side-neighbors $2 \times \frac{1}{4}b$ and pays the cost of producing a public good c by virtue of being an Altruist, resulting in a total payoff of $\frac{2}{4}b - c$. Similarly, the left Altruist neighbor of the Egoist gains benefits from its upper Altruist neighbor providing one public good, so its payoff is $\frac{1}{4}b - c$. Therefore, $\text{Payoff}_{avg}(A) = \frac{3}{8}b - c$. The Egoist agent will become an Altruist iff $\text{Payoff}_{avg}(A) > \text{Payoff}_{avg}(E)$. In other words, when $(\frac{3}{8}b - c) > \frac{1}{4}b \Leftrightarrow \frac{1}{8}b > c$, the Egoist becomes an Altruist. Otherwise, the agent will remain in an Egoist. To provide a concrete example of this condition, if we have $b = 1$ then the agent will become altruist where $c < \frac{1}{8} = 0.125$.

If the Egoist is surrounded by Altruist neighbors, what will the Egoist become? Can cost influence a transition to Altruism? We consider the following example:

$$\begin{array}{ccccc} E & E & E & E & E \\ E & A & \underline{A} & A & E \\ E & \underline{A} & \mathbf{E} & \underline{A} & E \\ E & E & \underline{A} & E & E \\ E & E & E & E & E \end{array}$$

Interestingly, the answer is *No*. Rather, it never becomes an Altruist, having

$\text{Payoff}_{avg}(E) = b$ and $\text{Payoff}_{avg}(A) = \frac{1}{4}b - c$. We will detail this scenario in the next section.

4.2 Egoist Islands

We discuss a special case where an Egoist is surrounded by Altruist neighbors.

Definition An *Egoist Island* is a set of the Egoist agent(s) which is isolated by surrounding Altruist neighbors. *Altruist Island* is defined as the reverse case, a group of Altruists surrounded by Egoists.

To determine the type of agent, an Egoist agent in the Egoist Island compares its $\text{Payoff}_{avg}(E)$ to $\text{Payoff}_{avg}(A)$. You might think that if an agent is surrounded by neighbors with other types, it will assimilate the type of behaviors of neighbors. However, as we show, the Egoist will not change its type, despite being surrounded by Altruists.

Observation 2.1 An Egoist surrounded by Altruists never converts to Altruist, as long as b and c are positive.

Let us take an Egoist Island example with the stricter condition of Egoist Island than one in the previous section.

$$\begin{array}{ccccc}
A & A & A & A & A \\
A & A & \underline{A} & A & A \\
A & \underline{A} & \mathbf{E} & \underline{A} & A \\
A & A & \underline{A} & A & A \\
A & A & A & A & A
\end{array}$$

In this case, the average payoffs for both types are

$$\text{Payoff}_{avg}(E) = \frac{1}{4}b \times 4 = b$$

$$\text{Payoff}_{avg}(A) = \frac{1}{4}[4 \times (\frac{3}{4}b - c)] = \frac{3}{4}b - c$$

The condition that the Egoist becomes an Altruist is $\frac{3}{4}b - c > b \Leftrightarrow -\frac{1}{4}b > c$, so that it never converts to an Altruist.

An Egoist Island of size two will also persist.

$$\begin{array}{cccccc}
A & A & A & A & A & A \\
A & A & \underline{A} & \underline{A} & A & A \\
A & \underline{A} & \mathbf{E} & \mathbf{E} & \underline{A} & A \\
A & A & \underline{A} & \underline{A} & A & A \\
A & A & A & A & A & A
\end{array}$$

Then,

$$\text{Payoff}_{avg}(E) = \frac{1}{2} \times (\frac{3}{4}b + \frac{3}{4}b) = \frac{3}{4}b$$

$$\text{Payoff}_{avg}(A) = \frac{1}{3}[3 \times (\frac{3}{4}b - c)] = \frac{3}{4}b - c$$

The condition that the Egoist becomes an Altruist is $\frac{3}{4}b - c > \frac{3}{4}b \Leftrightarrow 0 > c$. So it holds.

However, if the size of Egoist Island is greater than or equal to three, and at least one Egoist that is surrounded by less than three Altruist neighbors, the group will not survive..

Observation 2.2 An Egoist island will remain intact as long there are less than three Egoists.

4.3 Stackelberg Threshold for Altruists

Let us consider the case of an Altruist Island with regards to Stackelberg threshold[6]. The *Stackelberg Threshold* is the minimum number of players that must be controlled by a central authority to achieve the social optimum. In this case, the Stackelberg Threshold is the number of Altruists that must be part of an Altruist Island before the altruists can convert the entire grid to altruism.

Observation 3 The minimum size of an Altruist Island that can survive through successive rounds is 3.

Proof First, take a look at the cases where the number of Altruists is less than 3.

1. When the number of Altruists $N(A)$ is 0, it is trivial.

2. When $N(A) = 1$,

$$\begin{matrix} E & E & E & E & E \\ E & E & \underline{E} & E & E \\ E & \underline{E} & \mathbf{A} & \underline{E} & E \\ E & E & \underline{E} & E & E \\ E & E & E & E & E \end{matrix}$$

$$\text{Payoff}_{avg}(E) = \frac{1}{4} \times (4 \times \frac{1}{4}b) = \frac{1}{4}b$$

$$\text{Payoff}_{avg}(A) = -c$$

   To remain in Altruist, $-c > \frac{1}{4}b$. It never happens as long as b and c are positive.

3. When $N(A) = 2$, the payoffs for two Altruists are same since they are symetric.

$$\begin{matrix} E & E & E & E & E & E \\ E & E & \underline{E} & \underline{E} & E & E \\ E & \underline{E} & \mathbf{A} & \mathbf{A} & \underline{E} & E \\ E & E & \underline{E} & \underline{E} & E & E \\ E & E & E & E & E & E \end{matrix}$$

$$\text{Payoff}_{avg}(E) = \frac{1}{3} \times (3 \times \frac{1}{4}b) = \frac{1}{4}b$$

$$\text{Payoff}_{avg}(A) = \frac{1}{2} \times [2 \times (\frac{1}{4}b - c)] = \frac{1}{4}b - c$$

   Likewise $\text{Payoff}_{avg}(A)$ is less than $\text{Payoff}_{avg}(E)$.

Then, let us consider the case $N(A) = 3$.

$$\begin{matrix} E & E & E & E & E & E \\ E & E & \underline{E} & \underline{E} & E & E \\ E & \underline{E} & \mathbf{A}_1 & \mathbf{A} & \underline{E} & E \\ E & \underline{E} & \mathbf{A} & \underline{E} & E & E \\ E & E & \underline{E} & E & E & E \\ E & E & E & E & E & E \end{matrix}$$

1. We compute the left up the Altruist(A_1) in the Altruist island which has two Altruist neighbors.

$$\text{Payoff}_{avg}(E) = \frac{1}{2} \times (2 \times \frac{1}{4}b) = \frac{1}{4}b$$

$$\text{Payoff}_{avg}(A) = \frac{1}{3} \times [(\frac{2}{4}b - c) + 2 \times (\frac{1}{4}b - c)] = \frac{1}{3}b - c$$

The Altruist has a greater incentive to become an Altruist where $\frac{1}{3}b - c > \frac{1}{4}b \Leftrightarrow \frac{1}{12}b > c$.

2. For the other two Altruists,

$$\text{Payoff}_{avg}(E) = \frac{1}{3} \times (\frac{2}{4}b + 2 \times \frac{1}{4}b) = \frac{1}{3}b$$

$$\text{Payoff}_{avg}(A) = \frac{1}{2} \times [(\frac{2}{4}b - c) + (\frac{1}{4}b - c)] = \frac{3}{8}b - c$$

So they convert to an Altruist, where $\frac{3}{8}b - c > \frac{1}{3}b \Leftrightarrow \frac{1}{24}b > c$.

Therefore, we conclude that if there exist three Altruists controlled by a central authority, then they all will survive in next round when $\frac{1}{24}b > c$. It is not difficult to analyze that an Altruist island with more than three Altruists also survives, since we can control the position of Altruists so they may have at least one Altruist neighbor(like A_1) which takes benefits from its two Altruist neighbors and that only one Egoist has two Altruist neighbors.

$$\begin{array}{cccccc}
E & E & E & E & E & E \\
E & E & \underline{E} & \underline{E} & E & E \\
E & \underline{E} & \mathbf{A}_1 & \mathbf{A} & \underline{E} & E \\
E & \underline{E} & \mathbf{A} & \underline{E} & E & E \\
E & E & \underline{E} & E & E & E \\
E & E & E & E & E & E
\end{array}$$

In addition, surrounding Egoist neighbors($\underline{E}$) will convert to Altruists at $c > \frac{1}{12}b$. As a result, if $\frac{1}{24}b > c$, after repeated rounds, the Altruists dominate a game and all the positions are taken by them except four corners.

4.4 Corner Effects

Even if the Altruists dominate a game, they cannot completely convert a grid to Altruism, as Egoists at each of the four corners have an incentive to maintain their strategies.

Observation 4 If the Egoist is put on the corner surrounded by Altruists, its strategy will not change.

$$\begin{array}{cccc}
E & A & . & . \\
A & A & . & . \\
. & . & . & . \\
. & . & . & .
\end{array}$$

We did not go over any details of these cases but an Egoist at a corner surrounded by Altruists converts to an Altruist.

$$\text{Payoff}_{avg}(E) = \frac{2}{3}b$$

$$\text{Payoff}_{avg}(A) = \frac{1}{2} \times [2 \times (\frac{1}{3}b + \frac{1}{4}b - c)] = \frac{7}{12}b - c$$

Hence, $\text{Payoff}_{avg}(A) - \text{Payoff}_{avg}(A) > 0 \Leftrightarrow (\frac{7}{12}b - c) - \frac{2}{3}b > 0 \Leftrightarrow -\frac{1}{12}b > c$, and it never changes its type.

However, interestingly, a slight change can result in a different outcome. How can we make a game where egoists do not exist at all? Have another egoist beside it!

$$\begin{array}{cccc} E_1 & E_2 & A & . \\ A & A & A & . \\ . & . & . & . \\ . & . & . & . \end{array}$$

For E_1,

$$\text{Payoff}_{avg}(E) = \frac{1}{2} \times [\frac{1}{2}b + (\frac{1}{3}b + \frac{1}{4}b)] = \frac{13}{24}b$$

$$\text{Payoff}_{avg}(A) = \frac{1}{3}b + \frac{1}{4}b - c = \frac{7}{12}b - c$$

$\text{Payoff}_{avg}(A) - \text{Payoff}_{avg}(E) > 0 \Leftrightarrow (\frac{7}{12}b - c) - \frac{13}{24}b > 0 \Leftrightarrow \frac{1}{24}b > c.$

For E_2,

$$\text{Payoff}_{avg}(E) = \frac{1}{2} \times [\frac{1}{2}b + (\frac{1}{3}b + \frac{1}{4}b)] = \frac{13}{24}b$$

$$\text{Payoff}_{avg}(A) = \frac{1}{2} \times [(\frac{1}{3}b + \frac{2}{4}b - c) + (\frac{1}{3}b + \frac{1}{4}b - c)] = \frac{17}{24}b - c$$

$\text{Payoff}_{avg}(A) - \text{Payoff}_{avg}(E) > 0 \Leftrightarrow (\frac{17}{24}b - c) - \frac{13}{24}b > 0 \Leftrightarrow \frac{1}{6}b > c.$

In sum, if c is less than $\frac{1}{24}b$ for the Egoist island with the size of two at the corner then the Egoists have a greater incentive to convert to an Altruist.

5 Conclusion

We presented current research in the area of Privacy and Economics, including some interesting open questions. Grid-based models of agent interaction remain an interesting area of study. Our novel contribution is to model a game of Altruists and Egoists using the method of *imitation dynamics*, illustrate some interesting behaviors, and prove some simple bounds for the evolution of these games.

References

[1] M. Babaioff, R. Kleinberg, and C.H. Papadimitriou. Congestion games with malicious players. *Games and Economic Behavior*, 67(1):22–35, 2009.

[2] F. Brandt, T. Sandholm, and Y. Shoham. Spiteful bidding in sealed-bid auctions. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1207–1214, 2007.

[3] I. Eshel, L. Samuelson, and A. Shaked. Altruists, egoists, and hooligans in a local interaction model. *American Economic Review*, 88(1):157–179, 1998.

[4] J. Kleinberg, C.H. Papadimitriou, and P. Raghavan. On the value of private information. In *Proceedings of the 8th conference on Theoretical aspects of rationality and knowledge*, pages 249–257. Morgan Kaufmann Publishers Inc., 2001.

[5] T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: Byzantine players in a virus inoculation game. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, pages 35–44. ACM, 2006.

[6] Y. Sharma and D.P. Williamson. Stackelberg thresholds in network routing games or the value of altruism. *Games and Economic Behavior*, 67(1):174–190, 2009.