THE EFFECTS OF ALTRUISM AND SPITE ON GAMES

by

Po-An Chen

A Dissertation Presented to the
FACULTY OF THE USC GRADUATE SCHOOL
UNIVERSITY OF SOUTHERN CALIFORNIA
In Partial Fulfillment of the
Requirements for the Degree
DOCTOR OF PHILOSOPHY
(COMPUTER SCIENCE)

August 2011

Acknowledgements

First and foremost, I would like to thank my advisor, David Kempe, for his guidance, support throughout this journey and being a thoughtful and patient mentor, as well as the other members of my committee, Ming-Deh Huang, Shanghua Teng, Milind Tambe and Harrison Cheng, for their helpful comments and suggestions. I would also like to thank my colleagues and collaborators, Mary David, Bart de Keijzer, Guido Schaefer, Xiaoxun Sun, William Yeoh and Xiaoming Zheng, for all the stimulating discussions we have had. Last but not least, I would like to thank my family, especially, my parents, sister and grand parents for their support, understanding and patience with me, as well as my good friends, Yu-Chuan Lin, Ekapong Kachai, Gu Le Le Lu, Ga-Ming Ngo, Jia Tan, Hsuien-Yen Chang, Ta-Lun Huang, and Chih-Wen Lo for their encouragement and companionship.

Table of Contents

List of Tables

List of Figures

Abstract

Standard game theory assumes purely selfish or rational individual behavior, which means that every player will just act to optimize his own payoff function regardless of the effects that their choices may have on the others. However, many phenomena where people do care about others' benefits can be observed in the real world. Experiments also show discrepancy between experimental results and theoretical predictions with the assumption of selfishness. Various explanations with "not entirely selfish" players perceiving "other-regarding" payoffs have been proposed. One of them is altruism and spite among players.

Selfish outcomes have been observed to be drastically downgraded from the optimal one in several natural games. Since players are not totally selfish, these predictions may have been simply too pessimistic. Our goal in this thesis is studying the impact of partially altruistic and spiteful behavior on the outcome of games, and specifically the social welfare, in a social or economic network environment.

We develop and analyze a game-theoretic model with partially altruistic and spiteful players situated in an economic or social network environment. We show the effects of such a model on several problems: traffic routing, congestion games, network vaccination, and auctions. The trend of impact from altruism is different across classes of games.

In particular, improvements on the Price of Anarchy are shown in routing games with non-atomic partially altruistic users. However, this trend is not the case for congestion games with atomic partially altruistic players in which the Price of Anarchy is increasing with altruism, yet some special cases of congestion games still exhibit the trend of improvements. Introducing partial altruism into network vaccination games can result in no stable outcome, which even changes the game dynamics completely. We then draw a roadmap of a few interesting future directions.

Chapter 1

Introduction

Standard game theory assumes purely selfish or rational individual behavior, which means that every player will just act to optimize his or her own payoff function regardless of the effects that their choices may have on the others. However, many phenomena where people do care about others' benefits or about their intentions can be observed in the real world. Directly benefiting others, people do volunteering work; people make donations to NGOs or charities. Indirectly benefiting the society, people choose to be environmentally-friendly by driving hybrid cars, etc. On the other hand, people sometimes do want to harm or reduce the benefits of other people who they dislike. These observations seem to be inconsistent with the assumption of selfishness.

Furthermore, the assumption of selfishness has been repeatedly questioned by economists and sociologists, with experiments showing a discrepancy from the prediction with such an assumption. While the model of selfishness produces outcomes that are quite consistent with some experiments such as competitive experimental markets, it fails in public goods and trust and reciprocity experiments, for instance, public goods contribution games and ultimatum bargaining [57, 58]. (A more detailed discussion can be found in Chapter 3):

- In public goods contribution games, each player with a fixed initial endowment is free to make a donation at a cost to himself to a common pool that will give a social benefit greater than the contribution. Each should contribute all his endowment at the social optimal outcome, but contributing nothing is a dominant strategy, i.e., everyone will be free riding. Yet, with as many as 10 or more players, experiments found that players do contribute to the common pool. Total contributions can be expected to lie between 40% to 60% of the social optimum.

- In ultimatum bargaining, the first player proposes a proportion to divide a fixed amount of money, and the second player gets to accept or reject it. The first player can make only one offer. If the second player accepts the offer, the money is shared accordingly; if the second player rejects the offer, both players receive nothing. According to the theory with selfishness, any division leaving the second player with any non-zero amount should be accepted, so the first player should demand at least the greatest amount less than the entire amount. In experiments, surprisingly the first players do not propose a demand near to this amount, only 50 to 60 percent of the total being demanded in general and ungenerous demands (but still significantly less than the total) even being rejected frequently.

All these imply that even for simple games in controlled environments (in the absence of personal interaction or repeated experiments), participants do not really act selfishly;

their behavior can be either friendly or the opposite, malicious, to other players. Many explanations have been considered for this phenomenon. Most of them adopted simple or more complicated revised models to explain experimental results, which are either consistent or inconsistent with selfishness, for various kinds of games. Different types of models with "other-regarding payoffs" have been developed: Levine brought up a model of relative spite and altruism [58], which can be regarded as incorporating fairness, not in the sense that players have a particular target to be considered "fair" by them, but in the sense they are willing to be more altruistic to another player who is more altruistic towards them. Fehr and Schmidt considered a model with an innate sense of fairness [33], where in addition to purely selfish players, there are players who dislike inequitable outcomes. They incur disutility when experiencing inequity if they are worse off or better off in their original payoff terms than the other players. "Positive reciprocity" is a motivation to repay generous or helpful behavior of another by behaving generously or helpfully to the other person; thus, positively reciprocal behavior is conditional kindness different from the unconditional kindness motivated by altruism [28]. Gintis et al. explained cooperation among agents by reciprocity [42], where they showed that a high level of cooperation can be attained when social groups have a sufficient fraction of strong reciprocators.

In another line of work, people observe that in many games, the resulting outcomes from selfish actions can be drastically different from the optimal one under a central authority, in terms of certain global measures. In particular, the *social welfare (or cost)*, which is defined as the sum of all players' utilities (or costs), is of interest in this context. The ratio between the *worst* outcome of selfish choices and the socially optimal outcome has been termed the "Price of Anarchy" (PoA) by Koutsoupias and Papadimitriou [55] to measure how much the overall performance of any outcome of selfish choices gets downgraded compared to the optimum. Similarly, the "Price of Stability" (PoS) [2] is then defined as the ratio between the *best* outcome of selfish choices and the socially optimal outcome to capture how bad the overall performance will have to be if we want that no player has an incentive to deviate. The formal definitions of different equilibrium concepts of outcomes and the PoA and PoS are introduced in Section 2.1 and 2.2, respectively. The PoA or PoS of a lot of games have been shown to be quite high. Time and again, the assumption of selfishness is questionable here. Is the PoA or PoS really that bad or good in those games if players actually are *not entirely selfish*, i.e., partially altruistic or spiteful?

Our goal in this thesis thus is

> studying the impact of partially altruistic/spiteful behavior on the outcome of games, and specifically the social welfare.

Specifically, we investigate the question whether and how Nash equilibria and the Price of Anarchy (PoA), the Price of Stability (PoS), or the revenue (for auctions), will change if players are assumed to be partially altruistic/spiteful, embedded in a social or economic network environment.

To this end, we consider a natural model of altruism and spite, extending one proposed by Ledyard [57, p. 154]. The model of altruism and spite is given formally in Section 2.3. Intuitively, we want to model that players will trade off the benefit to themselves against the benefit to others. This can be modeled by assuming that the *perceived* utility of each player is a linear combination of his own a priori payoff and the payoffs of other players. Note that game theory is simultaneously a theory of utility and a theory of play (equilibrium concepts) for strategic games [93]. However, a lot of subsequent development of it has concentrated on the analysis of how to play games and equilibrium concepts without looking into utility theory. This kind of interpretation of game theory assuming players do not care about others' payoffs or about their intentions is only a special-case interpretation of game theory. What we try to do here is studying the PoA/PoS via an interpretation of game theory with a more general utility theory, since the previous analysis of the PoA/PoS is also subject to the limited interpretation of game theory.

The second emphasis of this thesis is modeling the interaction relationship between any pair of players in an economic or social network, which will definitely affect how players perceive the outcomes. The interaction relationship can be nicely modeled as a weighted (directed) edge between a pair of two nodes in a graph, where the weights represent different altruism/spite levels between any pair of players.

For instance, imagine that a player who does not only care about his own benefit but also the other players' benefits, may become happier if his action does not only bring a fair benefit to himself but also to the other players, or a player may become less happy if another player who he competes with more gets allocated something valuable instead of another one who he does not compete with or care about that much. Often, such relationships between players can even be modeled just as that between a player himself and all the others, whose benefits can be treated as a whole (namely, the social welfare or social cost), as a special case so as to obtain interesting results. It roughly corresponds to people's caring or awareness for the overall welfare or cost in the society, and a parameter can be used to capture how strong this feeling is compared to selfish incentives. Nevertheless, in the most general case, each player may care about each other player quite differently. After all, it is natural for any player to prefer/dislike (and therefore to be more altruistic/spiteful towards) any other player in a social or economic setting.

In summary, we develop and analyze a game-theoretic model with partially altruistic/spiteful players defined through their perceived utilities, situated in an economic or social network environment. We show the impact of such a model on several problems, in particular, traffic routing, congestion games, network vaccination, and auctions. The trend of impact from altruism is different across classes of games. In particular, improvements on the PoA are shown in routing games with non-atomic partially altruistic users. However, this trend is not the case for congestion games with atomic partially altruistic players in which worsening of the PoA is proved, yet some special cases of congestion games still exhibit the trend of improvements. Introducing partial altruism into network vaccination games can result in no stable outcome, which even changes the game dynamics completely. We then move on to drawing a roadmap of a few interesting future directions.

We briefly summarize our results here.

1.1 Traffic Routing

We consider traffic routing under the impact of altruism/spite as the first class of games to examine. Roughgarden and Tardos [86] pioneered the study of the PoA for traffic routing networks. They analyze a model proposed by Wardrop and Beckmann et al. [94, 8], in which edges possess traffic-dependent latency functions. When users choose a certain path, they increase the traffic on all edges of the path, and thus also the latency experienced by all other users sharing the path. This model of selfishness assumes, in accordance with much of the game theory literature, that users choose their routes completely without regard to the delay that their choice may cause for other users in the system. The PoA can be bad under this model of selfishness. Nevertheless, as mentioned above, if users are partially altruistic (i.e., when choosing a route, they care about the latency experienced by all other users sharing the path besides minimizing their own latency), we are interested in how the PoA will change. We can instantiate our general model in the context of traffic routing.

The perceived "cost" of a user is a linear combination of his own latency and the increase in latency the users causes others (precise definitions are given in Section 4.1). By varying the *altruism coefficient* β (so each user does not differentiate other users yet β can still vary among users), we can smoothly tune the altruism from spiteful ($\beta = -1$), through selfish ($\beta = 0$), to entirely altruistic ($\beta = 1$).

Our first result is that if all users are (at least) β-altruistic, and $\beta > 0$, then the PoA is always bounded by $1/\beta$, for all networks, arbitrarily many commodities, and arbitrary semi-convex latency functions on the edges. Thus, if a constant amount of altruism is introduced into the system, then the PoA is bounded by a constant. A more general version of our result characterizes precisely the worst-case PoA for any class of latency functions; from this general result, better bounds can be obtained for more restricted classes of functions. Among others, our result implies a bound of $\frac{4}{3+2\beta-\beta^2}$ on the PoA if all latency functions are linear. The general bound also lets us analyze the *spite resistance* of a class of latency functions: the most spite under which the PoA would still be finite.

We next extend our results beyond uniform altruism, and consider arbitrary distributions of altruism among the players. In that scenario, even the existence of Nash Equilibria is not obvious; we use a theorem of Mas-Collel [64] to prove that such games with infinitely many agents indeed have Nash Equilibria. Even for single-commodity flows in arbitrary graphs, prohibitive lower bounds on the PoA are known [12], so we focus here on parallel link networks, studied for instance by Roughgarden [80].

For parallel link networks, we show that for any non-negative distribution of altruism over the users in the network with *average altruism level* $\bar{\beta}$, the PoA with convex edge latency functions is always bounded by $1/\bar{\beta}$. In the specific case where the distribution of altruism has only completely altruistic or completely selfish users, this matches a bound obtained (with a polynomial-time algorithm) by Roughgarden [80].[1] The bound of $1/\bar{\beta}$ follows from a more general result characterizing the PoA for arbitrary classes of convex functions. In fact, that more general result, when applied to the case of a distribution over entirely selfish and entirely altruistic users, implies tighter bounds for Stackelberg routing compared with a result of Swamy [90]. Finally, we show that for the bound we

[1] Roughgarden's bound for Stackelberg routing on parallel link networks applies to arbitrary functions, whereas ours requires convexity.

derive, the worst case is in fact attained by the $\{0, 1\}$ altruism distribution, while the best case is when all users are $\bar{\beta}$-altruistic.

1.2 Congestion Games

We then consider a class of games generalizing routing games (i.e., network congestion games), congestion games with atomic players. Notice that we are dealing with atomic players whose individual size is 1, instead of non-atomic players whose individual size is infinitesimal. Instead of just focusing on the most restricted equilibrium concept to analyze the outcomes, we also relax the equilibrium concepts that we use to more permissive ones. (For the formal definitions of different equilibrium concepts, refer to Section 2.1.) Yet, we are only interested in partial altruism (so partial spite is not discussed) and linear cost functions here.

Now we have seen how relaxing the assumption of selfishness change things. We want to see what would happen if we also look at more general equilibrium concepts. The adoption of Nash equilibria as a prescriptive solution concept implicitly assumes that players are able to reach such equilibria. In particular, in light of several known hardness results for finding Nash equilibria, this assumption is very suspect for computationally bounded players. In response, recent work has begun analyzing the outcomes of natural response dynamics [10, 11, 84], as well as more permissive solution concepts such as correlated or coarse correlated equilibria [5, 44, 85]. This general direction of inquiry has become known as "robust Price of Anarchy."

Thus, our goal is to begin a thorough investigation of the effects of relaxing both of the standard assumptions simultaneously, i.e., considering the combination of weaker solution concepts and notions of partially altruistic behavior by players. In Section 5.1, our formal definition of an n-player *altruistic* congestion game with parameters (β_i) falls under our general model of altruism (however, again with each player i not differentiating other players $j \neq i$ but his parameter β_i can still be different from β_j of another player j for $j \neq i$). Note that our altruism model (despite also considering spite there) in Section 4.1 can be seen as a non-atomic analogy of our model for atomic players here: Informally, player i's cost (or payoff) is a convex combination of $(1-\beta_i)$ times his direct cost (or payoff) and β_i times the social cost (or social welfare) for β_i ranging from pure selfishness ($\beta_i = 0$) to pure altruism ($\beta_i = 1$).

In order to analyze the degradation of system performance in light of partially altruistic behavior, we extend the notion of *robust Price of Anarchy* [84] to games with partially altruistic players, and show that a suitably adapted notion of *smoothness* [84] captures the properties of a system that determine its robust Price of Anarchy. We use these insights to analyze *(linear) congestion games*: Players choose subsets of resources, and as resources are chosen by many players, their cost increases (linearly), to all players using them.

We derive tight bounds on the (robust) PoA for these games for uniform altruism, i.e., $\beta_i = \beta$ for all i. The prevailing trend is rather unexpected: for congestion games, the worst-case robust PoA actually *increases* (as $(5+4\beta)/(2+\beta)$) with increasing altruism β. The intuition behind the increase is the following: there are instances in which all players get stuck choosing the wrong resources. A deviation by one player affects not only him, but also others: for congestion games, the player may increase the cost on the resources he switches to. Thus, partially altruistic players have even stronger disincentive to deviate from the suboptimal strategy, meaning that even worse system states are stable.

The above explanation intuitively corresponds to altruistic players “accepting” more states as “stable”. This suggests that the best stable solution can also be chosen from a larger set, and the PoS should thus decrease. Our results lend partial support to this intuition: for congestion games, we derive an upper bound on the PoS which decreases as $2/(1+\beta)$.

It should be noted that the increase in PoA is not a universal phenomenon. Indeed, for linear symmetric singleton congestion games (in which all players have the same strategy set, consisting of all sets of exactly one resource), we establish a bound of $4/(3+\beta)$ for the PoA with respect to pure Nash equilibria (implying a bound of $4/3$ when $\beta = 0$). This bound is noteworthy not only because it shows improvements resulting from the presence of altruism; it also establishes that pure Nash equilibria can result in strictly lower PoA than weaker solution concepts.[2] In particular, this establishes a natural example of a class of games whose PoA can not be established using the smoothness framework. The results are completed by a more in-depth analysis of the effects on the pure PoA of combining players with different altruism levels in singleton congestion games, giving bounds based on the proportion of players with different altruism levels.

[2]Lücking et al. [62] in the standard selfish setting gave an example of a class of congestion games whose PoA under mixed Nash equilibria can be arbitrarily close to 2.

1.3 Network Vaccination

We turn our attention to another class of network games yet with an entirely different flavor from the previous routing or congestion games: network vaccination games, or equivalently, network inoculation games.

Vaccinations or inoculations protect nodes in social or computer networks so that they will not be affected by outbreaks of epidemics or computer viruses. As was evident from the extensive coverage of vaccinations against virus spreading in networks, there are many factors complicating the allocation of vaccines. Among the most prominent ones are (1) supply shortages, limiting the number of individuals who can be vaccinated, and (2) node autonomy: individuals make their own decisions on whether to get inoculated, which may conflict with the socially optimal strategy. The former naturally leads to optimization problems for allocating the limited amounts of vaccine, while the latter raises natural questions about the inefficiency of outcomes in such settings. We dedicate our efforts solely to address the latter issue of node autonomy.

We study this issue in a model for network vaccinations generalizing one proposed by Aspnes et al. [4]. In this model, vaccinated nodes can never contract the disease, so they are effectively removed from the network. After all vaccination decisions are made, the disease will break out at a node v chosen with some probability p_v and infect all nodes reachable in the network (with the vaccinated nodes removed). There is a cost of C_v associated with node v being vaccinated and a cost of L_v for v being infected. (A formal description of the model is given in Section 6.1.1.) The optimum solution is one finding a set of nodes to vaccinate that minimizes the expected total cost of all nodes.

In reality, the decision of whether to get vaccinated usually lies with the individual nodes, whose interests do not necessarily align with the social goal of minimizing the total cost. Individuals tend to undervaccinate when they are not concerned with the impact of their action on other nodes. It is therefore natural to investigate how inefficient "societally stable" states can become as a result of individual decisions and their externalities. Indeed, Aspnes et al. [4] already showed that each instance of the inoculation game has at least one pure Nash Equilibrium and that the PoA can be $\Theta(n)$ in the worst case (but no worse).

The $\Theta(n)$ lower bound relies on the fact that individuals are entirely selfish in their vaccination decisions, and completely unaware of — or indifferent to — how their decisions may affect others in the network. Meier et al. [66] analyzed the impact of *friendship* on stable outcomes. In their model, a node's utility is the sum of its own cost and a β fraction of the cost of all its neighbors. They showed that for some graphs, this notion of friendship leads to significantly more efficient equilibria, while for others, the improvement is small.

Here, we instead consider a notion of *altruism* using our model, which as been instantiated in the context of routing and congestion games. Thus, our notion models a general feeling of altruism or responsibility for the welfare of society, and the uniform parameter β captures how strong this feeling is compared to selfish incentives. An alternative interpretation of this model is that the cost of the disease is "socialized" to an extent, e.g., that nodes' health insurance rates increase if others catch the disease. The parameter β then captures how steeply the individuals are penalized for others' diseases.

Interestingly, in the inoculation game with altruism, pure Nash Equilibria need not always exist, as we show in Section 6.1.2. Even when they do exist, Nash Equilibria can sometimes be as bad as in the model without altruism. Nonetheless, while the PoS without altruism can also be $\Theta(n)$, a similar notion improves dramatically with the altruistic model.

Since Nash Equilibria may not exist, the notion of "Price of Stability" does not apply directly. Mixed Nash Equilibria are not a natural solution concept here, as vaccination decisions tend to be permanent or very long-term. We therefore instead consider an "Opt-Out" dynamic, and correspondingly define the *Price of Opting Out*. A benevolent authority suggests an initial vaccination assignment S_0. The nodes in S_0 can choose to opt out of being vaccinated, in any order. However, no node $v \notin S_0$ may opt to become vaccinated. Also, nodes, once opted out, cannot opt back in. (Precise definitions are given in Section 6.1.3.) This models a scenario in which individuals may choose to avoid being vaccinated due to various concerns, but the authority will not revise plans vis-à-vis nodes not originally included in the vaccination plan. Our main theorem (stated formally and proved in Section 6.2) is then the following: The Price of Opting Out is at most $1/\beta$.

Thus, in a sense of somewhat limited autonomy among the nodes, our theorem establishes a $1/\beta$ bound on the social inefficiency introduced by individual nodes' decisions. Together with the $\Theta(n)$ bounds for PoA, and PoS without altruism, this result can be interpreted as saying that coordination of vaccination strategies or socialization of healthcare costs alone may not lead to societally desirable outcomes if individual nodes can override suggestions. However, the combination of both, i.e., socializing costs *and* starting with a carefully chosen assignment, may lead to significantly more desirable outcomes, even when individuals get to override the suggested vaccination strategies.

Naturally, it would be desirable to strengthen the results to the outcomes after arbitrary best-response dynamics. Since best-response dynamics may cycle for inoculation games with partial altruism, we cannot focus on stable states alone, but would have to consider all states reachable via best-response dynamics.

1.4 Auctions

As the last example application, we study the impact of social networks, altruism and spite on auctions. Specifically, due to the forming of preference/dislike relations through interactions, we are motivated to study auctions in which the utility of losers is not always 0, but rather depends on the identity of the winner, and the utility the winner derives from the auction. This falls again into our consideration of the partially altruistic/spite model, where a player does not only care about his own utility but also each other player's utility as well. We can capture this setting by instantiating our model in the most general form with a spite/altruism matrix $B = (\beta_{i,j})$, where each $\beta_{i,j} \in (-1, 1)$ for $i \neq j$, and $\beta_{i,i} = 1$ for each i. If bidder j wins the auction and obtains utility $\overline{u}_j$, then bidder i's utility from the auction is $\beta_{i,j} \cdot \overline{u}_j$. Thus, again if $\beta_{i,j} < 0$, then player i is *spiteful* toward player j (or a *foe*); if $\beta_{i,j} > 0$, then player i is *altruistic* toward player j (or a *friend*). Notice that we do not assume B to be symmetric.

Auctions with spite among players have been studied before [14, 15, 63, 69, 91]. However, in all past work, the assumption was that each off-diagonal entry of the spite matrix B was the same (and positive), i.e., all players have the same spite level toward each other. We call this the case of *uniform spite*. While it is interesting as an analysis of the effects of general distrust or future competition between bidders, we are also further interested in taking into account the effects of spite/altruism in social or economic networks on individual behaviors.

Here we study single-item auctions with friends and foes with *non-unform* spite/altruism. We focus on auctions with Bayesian priors. For two subclasses of these auctions, we explicitly describe a Nash Equilibrium. These two subclasses are the following:

1. The valuations are drawn independently from $[0, 1]$ according to an arbitrary (but identical) distribution for all bidders, and the social network of bidders is *regular*. This means that each bidder i has the same number d of non-zero $\beta_{i,j}$ for $j \neq i$ (and the same number $n - d$ of zero $\beta_{i,j}$ for $j \neq i$), and all such non-zero entries have the same value $\beta_{i,j} = \beta$. Notice that we are already dealing with non-uniform spite/altruism since each bidder i has non-zero $\beta_{i,j}$ and zero $\beta_{i,j}$. In this case, we analyze both first-and second-price auctions.

2. The spite/altruism matrix B is a non-negative *triangular* or *block* matrix (defined in Section 7.2), but each bidder's valuation of the item is drawn independently and uniformly from the interval $[0, 1]$, and the auction is first-price.

These characterizations partially generalize recent results of Morgan et al. and Brandt et al. [69, 14], which characterized a Nash Equilibrium for the case of uniform spite. We also point out here that the Equilibrium in the second case is not symmetric: different bidders have different bidding strategies.

Our explicit characterization also allows us to derive two interesting corollaries. For the case of regular social networks, we show that if $\beta < 0$ (i.e., bidders only have foes and neutral other bidders), the expected revenue of the second-price auction dominates the first-price auction. Conversely, if $\beta > 0$ (i.e., bidders only have friends and neutral other bidders), then the expected revenue of the first-price auction dominates the second-price auction. For the case of some social networks under uniform valuation distributions, the explicit characterization allows us to study the effect of changes in the social or economic network on the bidding behavior. Perhaps somewhat surprisingly, an increase in spite does not always lead to an increase in bids. Instead, we show that whether it leads to an increase or decrease in bids depends on whether the recipient of spite is currently overbidding or underbidding.

1.5 Future Directions

Our results on these problems suggest many interesting directions for further research, either extensions or related questions. We will list questions specifically related to traffic routing, congestion games, network vaccination, and auctions in the corresponding chapters. Here, we discuss broader questions beyond our current model of altruism and spite. We are interested in repeated games. In particular, with (non-uniform) altruism/spite that can be known or unknown to players, how can equilibrium strategies be *learned*? What would the PoA and PoS be for such equilibria? More broadly, there are other notions of "other-regarding" (or "not entirely selfish") behavior [34] such as fairness [33] and *reciprocity* [42] in the literature of experimental/behavioral game theory [17]. We are also interested in studying some of these models and analyzing their impacts or effects. Specifically, we discuss a model of fairness by Fehr and Schmidt defined for *inequity averse* and *reciprocal* players whose impact on cooperation in prisoners' dilemma games is shown [34].

1.5.1 Learning in Repeated Games

In *repeated games*, an intriguing general question is whether a similar model with not entirely selfish players is helpful for learning equilibrium strategies if there is already some learning algorithm for selfish players, or whether agents can learn equilibrium strategies using a natural learning algorithm.

For example, in routing there are already results using arbitrary *no-regret* algorithms for the standard selfish model [10], where the *per-time-step* regret of a user is the difference between her average latency and the latency of the best fixed path in hindsight, and an algorithm is no-regret if, for any sequence of flows, the expected regret over internal randomness in the algorithm goes to 0 as the number of time steps, T, goes to infinity. In particular, if each user runs a no-regret algorithm, the average regret over all users also approaches 0. Therefore, it can be assumed that a function $R(T)$, which is an upper bound on the average regret, goes to 0 as T goes to infinity; T_ϵ is then defined as the number of time steps required to get $R(T) = \epsilon$. However, it is possible for a flow f to have regret near 0 and yet still be far from a true Nash flow. We can only expect that *most* users take a nearly-cheapest path given a flow f. Define a flow f to be at ϵ-Nash equilibrium if the average cost under this flow is within ϵ of the minimum cost path under this flow; among other results, it is shown that for no-regret algorithms the *time-average* flow $\hat{f}$ is approaching equilibrium [10], i.e., bounds on the number of time steps before $\hat{f}$ is ϵ-Nash are obtained for a given T_ϵ.

We are interested in how (non-uniform) altruism/spite will have effects on the problem, especially on whether players' plays converge to any equilibria, the time of convergence, the PoA/PoS, etc. Similar questions can be asked for atomic congestion games and other games. The results for uniform altruism/spite may be established by simply extending those for entirely selfish users. However, such extending may not work for non-uniform altruism/spite, since users with different altruism levels would perceive the same network configuration differently, which is similar to the situation at one-shot games with non-uniform altrusm/spite. Existence of equilibria and the PoA/PoS are not obvious.

Another type of questions that one can ask emerges from the setting where each player does *not* know the (non-uniform) altruism/spite levels. We would be interested in finding simple strategies wherein each player adapts his strategy based on the utility or cost derived from earlier runs of games. It is not clear if no-regret algorithms would work here since the (non-uniform) altruism/spite levels are unknown. For example, can we learn bidding strategies when (non-uniform) altruism/spite levels are unknown?

1.5.2 Other-Regarding Payoffs

Broadly, we would like to explore beyond our current model of altruism and spite, under which we try to capture the connection between the "not entirely selfish" behavior and the social welfare. In the behavioral game theory literature, there are many different models of other-regarding payoff functions focusing on different notions of not entirely selfish behavior.

As an example of one of the future directions along this line, in Section 8.2.2 we are going to discuss another model that has been defined for *inequity averse* and *reciprocal* players, and shown its impact on *cooperation* through prisoners' dilemma games [34]. This may further motivate us to formally analyze the effects of inequity averse and reciprocal behavior on certain global measures capturing overall cooperation in the society.

Chapter 2

Preliminaries

We are giving basic definitions and models that we are using throughout this thesis, accompanied by some initial results of interesting examples.

We may need some desirable properties of functions later in this thesis: convexity or semi-convexity. A continuous and differentiable function $f : \mathbb{R} \to \mathbb{R}$ is *convex* if its second derivative is non-negative, i.e., $f''(x) \geq 0$; a discrete function $f : \mathbb{N} \to \mathbb{R}$ is convex if $f(x+1) - f(x) \geq f(x) - f(x-1)$. A (continuous or discrete) function $f : \mathbb{R} \to \mathbb{R}$ is *semi-convex* if $x \cdot f(x)$ is convex. We use $\mathbf{v}$ to denote a vector $(v_1, ..., v_i, ..., v_k)$. So, $\mathbf{0}$ is the all-ones vector and $\mathbf{1}$ is the all-ones vector.

2.1 Games and Equilibrium Concepts

The most essential thing that needs to be defined is what we mean by a game.

Definition 2.1.1 (Game) *A* game *consists of a (finite or infinite) set of players, each of which is assigned a finite set of strategies S_i and a payoff function $p_i : S_1 \times ... \times S_n \to \mathbb{R}^+$ that player i wants to optimize.*

A strategy profile $\mathbf{s} = (s_1, ..., s_n) \in S_1 \times ... \times S_n$ is any combination of strategies for the players, chosen from the Cartesian product of finite sets of strategies $S_1 \times ... \times S_n$.

At pure Nash equilibrium, no player can improve his payoff by unilaterally changing his strategy.

Definition 2.1.2 (Pure Nash Equilibrium) *A strategy profile* $\mathbf{s} = (s_1, ..., s_n)$ *is a* pure Nash equilibrium *if for all players* i*,* $p_i(s_1, ..., s_i, ..., s_n) \geq p_i(s_1, ..., s_i', ..., s_n)$ *for all* $s_i' \in S_i$*.*

We say that a strategy profile $\mathbf{s}$ is an ϵ-Nash equilibrium if no player can improve his payoff by more than ϵ by unilaterally changing his strategy.

Definition 2.1.3 (ϵ-Nash Equilibrium) *For* $\epsilon > 0$*, a strategy profile* $\mathbf{s} = (s_1, ..., s_n)$ *is an* ϵ-Nash equilibrium *if for all players* i*,* $p_i(s_1, ..., s_i, ..., s_n) \geq p_i(s_1, ..., s_i', ..., s_n) - \epsilon$ *for all* $s_i' \in S_i$*.*

We will use more general equilibrium concepts beyond Nash equilibrium in Chapter 5. Therefore, we introduce them here as well.

In a mixed Nash equilibrium, a player's strategy can be any probability distribution over available strategies, and no individual player can improve his expected payoff by choosing another probability distribution.

Definition 2.1.4 (Mixed Nash Equilibrium) *A vector* $(\sigma_1, ..., \sigma_n)$ *of independent probability distributions over strategy sets is a* mixed-strategy Nash equilibrium *if no player can improve his payoff under the product distribution* $\sigma = \sigma_1 \times ... \times \sigma_n$ *via a unilateral deviation:* $E_{\mathbf{s} \sim \sigma}[p_i(\mathbf{s})] \geq E_{\mathbf{s}_{-i} \sim \sigma_{-i}}[p_i(s_i', \mathbf{s}_{-i})]$ *for all* i *and* $s_i' \in S_i$*, where* σ_{-i} *is the product distribution of all* σ_j*'s other than* σ_i*.*

A correlated equilibrium is more general than a mixed Nash equilibrium. The idea can be interpreted as a trusted mediator who draws a strategy profile s from this distribution σ and "recommends" s_i privately to each player i. Assuming that the other players conform to the mediator's recommendation, if no player would want to deviate from the recommended strategy in expectation, the distribution is called a correlated equilibrium. Mixed-strategy Nash equilibria are the correlated equilibria that are also product distributions.

Definition 2.1.5 (Correlated Equilibrium) *A correlated equilibrium is a joint probability distribution $\sigma = (\sigma_1, ..., \sigma_n)$ over the strategy profiles with the property that $E_{\mathbf{s}\sim\sigma}[p_i(\mathbf{s})|s_i] \geq E_{\mathbf{s}\sim\sigma}[p_i(s_i', \mathbf{s}_{-i})|s_i]$ for all i and $s_i, s_i' \in S_i$, i.e., $\sum_{\mathbf{s}_{-i}\in S_{-i}} \sigma_{\mathbf{s}} p_i(s_i, \mathbf{s}_{-i}) \geq \sum_{\mathbf{s}_{-i}\in S_{-i}} \sigma_{\mathbf{s}} p_i(s_i', \mathbf{s}_{-i})$, where σ_s is the probability that the strategy profile is $\mathbf{s}$.*

A coarse correlated equilibrium is more general than a correlated equilibrium. A correlated equilibrium protects against deviations by players knowing their recommended strategy, while a coarse correlated equilibrium is characterized by a weaker property that each player does not want to deviate without seeing his recommendation.

Definition 2.1.6 (Coarse Correlated Equilibrium) *A* coarse correlated equilibrium *or* coarse equilibrium *is a joint probability distribution σ over strategy profiles that satisfy $E_{\mathbf{s}\sim\sigma}[p_i(\mathbf{s})] \geq E_{\mathbf{s}\sim\sigma}[p_i(s_i', \mathbf{s}_{-i})]$ for every i and $s_i, s_i' \in S_i$.*

Alternatively, coarse correlated equilibria can be defined as all the probability distributions which are the limit of the empirical distribution of some *no-regret sequence*, which will be formally described in Section 8.2.1.

Pure Nash equilibria in which all players play pure strategies (no randomization or mixed strategies) may not always exist while there always exist mixed Nash equilibria for any n-player game due to Nash's Theorem [72]. Note that payoffs can be either *utilities* to be maximized or *costs* to be minimized in an optimization sense.

2.2 Social Welfare, Price of Anarchy (PoA), and Price of Stability (PoS)

Given a game, it is natural to consider the social welfare as the global measure, intuitively, which reflects how well everyone as a whole does. There are many possible social welfare functions. We are mostly interested in the *utilitarian* function, i.e., the sum of all players' utilities.

The social optimum maximizes the social welfare over all possible strategy profiles. Usually, the social optimum is not reached with selfish players, because each of the players is only interested in his own utility.

The Price of Anarchy (PoA) is a measure of how well society does when they play by optimizing their own utility functions (reaching equilibrium, for instance, a pure Nash equilibrium) as opposed to choosing the social optimum.

Definition 2.2.1 (Price of Anarchy [55]) *The* Price of Anarchy *or* pure Price of Anarchy *is defined as the ratio of the social optimum welfare to the welfare of the worst pure Nash equilibrium, i.e.,*

$$\sup_{\mathbf{s}\in E(I)} \frac{\sum_{i=1}^{n} p_i(\mathbf{s})}{\sum_{i=1}^{n} p_i(\mathbf{s}^*)},$$

where $E(I)$ is the set of all pure Nash equilibria and $\mathbf{s}^*$ is the social optimum given a game instance I.

In other words, it is the ratio of the largest social welfare achievable to the least social welfare achieved at any Nash equilibrium. It captures how much the overall performance of any outcome of selfish choices gets downgraded compared to the optimum.

There can be a series of notions of generalized Price of Anarchy besides the pure Price of Anarchy corresponding to different general equilibrium concepts.

Definition 2.2.2 (Mixed Price of Anarchy) *The* Price of Anarchy of mixed Nash equilibria *or* mixed Price of Anarchy *is defined as the ratio of the social optimum welfare to the expected welfare of the worst mixed Nash equilibrium, i.e.,*

$$\sup_{\sigma \in E(I)} \frac{E_{\mathbf{s}\sim\sigma}[\sum_{i=1}^{n} p_i(\mathbf{s})]}{\sum_{i=1}^{n} p_i(\mathbf{s}^*)},$$

where $E(I)$ is the set of all mixed Nash equilibria and $\mathbf{s}^*$ is the social optimum given a game instance I.

Definition 2.2.3 (Correlated Price of Anarchy) *The* Price of Anarchy of correlated equilibria *or* correlated Price of Anarchy *is defined as the ratio of the social optimum welfare to the expected welfare of the worst correlated equilibrium, i.e.,*

$$\sup_{\sigma \in E(I)} \frac{E_{\mathbf{s}\sim\sigma}[\sum_{i=1}^{n} p_i(\mathbf{s})]}{\sum_{i=1}^{n} p_i(\mathbf{s}^*)},$$

where $E(I)$ is the set of all correlated equilibria and $\mathbf{s}^*$ is the social optimum given a game instance I.

Definition 2.2.4 (Coarse Price of Anarchy) *The* coarse Price of Anarchy *or* Price of Total Anarchy [1] *is defined as the ratio of the social optimum welfare to the expected welfare of the worst coarse correlated Nash equilibrium, i.e.,*

$$\sup_{\sigma \in E(I)} \frac{E_{\mathbf{s}\sim\sigma}[\sum_{i=1}^{n} p_i(\mathbf{s})]}{\sum_{i=1}^{n} p_i(\mathbf{s}^*)},$$

where $E(I)$ is the set of all coarse correlated equilibria and $\mathbf{s}^*$ is the social optimum given a game instance I.

Similarly, if we consider the best equilibrium instead of the worst equilibrium, then for pure Nash equilibria we have the following definition.

Definition 2.2.5 (Price of Stability [2]) *The* Price of Stability *is defined as the ratio of the social optimum welfare to the welfare of the best pure Nash equilibrium, i.e.,*

$$\inf_{\mathbf{s} \in E(I)} \frac{E_{\mathbf{s}\sim\sigma}[\sum_{i=1}^{n} p_i(\mathbf{s})]}{\sum_{i=1}^{n} p_i(\mathbf{s}^*)},$$

where $E(I)$ is the set of all pure Nash equilibria and $\mathbf{s}^*$ is the social optimum given a game instance I.

In other words, it is the ratio of the largest social welfare achievable to the largest social welfare achieved at any (pure) Nash equilibrium. It captures how bad the world will have to be if we want that no player wants to deviate. Such an outcome can be thought as a stable state stablizing through steps of deviations from some instable system state, and thereby "stability" is used to name this measure.

[1] Blum et al. [11] define the price of total anarchy as the worst-case ratio of the social optimum to the expected average welfare of a "no-regret sequence", which we will discuss more in Chapter 8.

Also, there can be a series of different notions of Price of Stability besides the Price of Stability for pure Nash equilibria corresponding to different general equilibrium concepts. However, we do not analyze them here so we are going to skip these definitions.

2.3 Model of Partial Altruism and Spite

To achieve our goal in this thesis, we need to model altruism and spite by designing a more general model of utility that subsumes the model with only "self-regarding" payoffs as a special case. There are more complicated models in the literature on "other-regarding" payoffs, discussed more in the beginning of the next chapter.

We base our treatment on a simple and elegant suggestion of Ledyard [57]: In a game with n players, the utility of a player i given a strategy profile $\mathbf{s}$ is $p_i(\mathbf{s}) + \beta_i \frac{1}{n} \sum_j p_j(\mathbf{s})$, where the p_i are the individuals' payoff functions. The parameter β_i captures how important the average social welfare is to player i.

We modify and generalize this approach slightly, and assume that the perceived payoff of each player is a linear combination of his own a priori payoff and the payoffs of other players. Players' perceived payoff functions are based on a coefficient matrix $B = (\beta_{i,j})$ representing the concept of interaction relationships (i.e., altruism/spite), where $\beta_{i,j}$ for $j \neq i$ is the value that player i has toward player j, and $\beta_{i,i}$ is the value that player i has toward himself. In a general form, player i's perceived payoff function is the combination

$$\sum_j \beta_{i,j} p_j(\mathbf{s}),$$

where the range of each $\beta_{i,j}$ depends on how we set it to be in different problems, and the interpretation of every $\beta_{i,j}$ can also vary and depend on problems. In different contexts, we can instantiate B with different parameter settings to provide a suitable model for different problems.

For instance, in Chapter 4 on traffic routing and 5 on atomic congestion games with altruism and spite, where users do not differentiate the other users individually ($\beta_{i,j}$ is the same for all $j \neq i$), by setting $\beta_{i,i} = 1 - \beta_i + \beta_i = 1$ and $\beta_{i,j} = \beta_i$ for all $j \neq i$, player i's perceived payoff function is

$$(1 - \beta_i)p_i(\mathbf{s}) + \beta_i \sum_j p_j(\mathbf{s}),$$

where β_i is the user's *altruism level*, representing how much he cares about the social welfare.

In Chapter 4 where $\beta_i \in [-1, 1]$, we adjust this definition a bit to look at the change "rate" of the social welfare in order to make it a meaningful model, since we are dealing with non-atomic users. This then has the advantage of making all utilities comparable on the same scale, and allowing us to model entirely altruistic behavior by setting $\beta_i = 1$. The restriction to values $\beta_i \geq -1$ is justified in Section 4.2. We call $p_i(\mathbf{s})$ the *selfish* part of player i's utility, and $\sum_j p_j(\mathbf{s})$ the *altruistic* part. If $\beta_i < 0$, then player i derives utility from a decrease in social utility; we call such players *spiteful*. In Chapter 5, we only consider *altruistic* players where $\beta_i \in [0, 1]$, since the techniques used there may not directly extend for spiteful players.

In Chapter 6 on network vaccination games with *uniform* altruism, every player not only does not differentiate the other players individually ($\beta_{i,j}$ is the same for all $j \neq i$) but also has the same altruism level by setting $\beta_i = \beta$ for all i.

In Chapter 7 on auctions with spite and altruism, we keep the most general model by setting $\beta_{i,i} = \beta_{i,i}$ and $\beta_{i,j} = \beta_{i,j}$, the perceived utility of bidder i becomes

$$\sum_j \beta_{i,j} p_j(\mathbf{s}),$$

where $\beta_{i,j} \in [-1, 1]$ and $\beta_{i,j}$ for $j \neq i$ is an off-diagonal element in the spite/altruism matrix $B = (\beta_{i,j})$. If $\beta_{i,j} < 0$ for $j \neq i$, then player i loses utility from player j's winning; we call such a player i *spiteful* to player j. If $\beta_{i,j} > 0$ for $j \neq i$, then player i derives utility from another player j's winning; we call such a player i *altruistic* to player j. Note that if every bidder does not differentiate the other bidders, then the model degenerates back to the one with $\beta_{i,j} = \beta_i$ for all $j \neq i$.

We may need to use some other settings for B when facing other problems in the future. This general framework allows us to do that with ease.

2.3.1 Recursive Altruism

Why do we not use a model where a player recursively considers the utility that a player derives from another player's perceived utility function? It should be generally interesting to consider such a model. Such systems of interdependent utility functions have been studied by Bergstrom [9]. He considered the question of when a system of interdependent utility functions induces unique utility functions over allocations by means of the theory of dominant-diagonal matrices.

We restate his main proposition in our context here, where the system of utility functions is linearly interdependent. Note that he proved a more general proposition ([9], Proposition 3).

Proposition 2.3.1 *Consider a system of interdependent utility functions described by the equations* $U_i = \alpha_{i,i}u_i + \sum_{j\neq i}\alpha_{i,j}U_j$ *for all individuals* i*, where* u_i *is the a priori payoff function of individual* i*,* U_j *is the (recursive) utility function of individual* j*, and* $\alpha_{i,j}$ *is a constant coefficient. Let* $A = (\alpha_{i,j})$ *be the altruism/spite matrix consisting of* $\alpha_{i,j}$ *representing altruism level from* i *to* j*. If* $I - A$ *is dominant diagonal, then there are coefficients* $\beta_{i,j}$ *such that the interdependent utilities can be written equivalently as just linear combinations of a priori payoff functions.*

Therefore, using his results, we can directly define our perceived utilities or costs on a priori payoff functions, without bothering starting indirectly from the recursive definitions.

He also identified the class of transformations on interdependent utility functions that are equivalent in the sense of inducing the same preferences over allocations ([9], Proposition 4).

2.3.2 Tolls, Taxes, and Socialization of Costs

By interpreting the altruistic term in our model as a *monetary* cost (instead of perceived cost due to altruism), we can also consider our model as one of *socialized costs* for a cost minimization problem. Player i has to pay a β_i fraction of the cost incurred by other nodes. For example, it can be in the form of health care premiums in a network vaccination game; this payment should provide additional incentives for node players in a network to be vaccinated, as spreading a disease to others will eventually lead to higher costs for them as well. In a routing or congestion game, the altruistic term can be interpreted as a toll or tax; see Section 4.1.2 for more discussion.

2.3.3 Counter-Intuitive Examples

With our model of altruism and spite defined, our first step is to test the intuition that altruism *always* helps in terms of the social welfare. Here we present two counter-intuitive examples of two-player games: There exists a game instance in which introducing altruism strictly decreases the social welfare for both the *worst* pure and *worst* mixed Nash equilibria; there also exists a game instance in which introducing altruism strictly decreases the social welfare for both the *best* pure and *best* mixed Nash equilibria.

Worst Pure/Mixed Nash Equilibrium.

Consider the Game 1 in Table 2.1. Player 1 playing strategy 1 and player 2 playing strategy 1 is the only pure and mixed Nash equilibrium, the social welfare of which is $1 + 5 = 6$. After defining the perceived payoff of a β-altruistic player to be $(1 - \beta)$ times the original payoff plus β times the social welfare (i.e., the convex combination of the original payoff and the social welfare), the payoff matrix with $\beta = 0.8$ becomes the one in Table 2.2. Then, there are two pure Nash equilibria: player 1 playing strategy 1 and player 2 playing strategy 1 (whose social welfare is 6) and Player 1 playing strategy 2 and player 2 playing strategy 2 (whose social welfare is $2 + 2 = 4$, which is now the

Table 2.1: Payoff matrix for game 1 without altruism.

	Player 2's strategy 1	Player 2's strategy 2
Player 1's strategy 1	(1,5)	(0,0)
Player 1's strategy 2	(0.5,3)	(2,2)

Table 2.2: Perceived payoff matrix for game 1 with uniform altruism $\beta = 0.8$.

	Player 2's strategy 1	Player 2's strategy 2
Player 1's strategy 1	(5,5.8)	(0,0)
Player 1's strategy 2	(2.9,3.4)	(3.6,3.6)

worst pure Nash equilibrium since $6 > 4$). So, the social welfare of the worst pure Nash equilibrium gets worse with altruism in Game 1.

For mixed Nash equilibria, besides these two equilibria with $\beta = 0.8$ there is one more: Player 1 plays strategy 1 with probability 0.03 and strategy 2 with probability 0.97 while player 2 plays strategy 1 with probability 0.6 and strategy 2 with probability 0.4, which results in an expected social welfare of $0.03 \cdot 0.6 \cdot 6 + 0.03 \cdot 0.4 \cdot 0 + 0.97 \cdot 0.6 \cdot 3.5 + 0.97 \cdot 0.4 \cdot 4 = 3.697$. Thus, the expected social welfare drops from 6 (player 1 playing strategy 1 and player 2 playing strategy 1 is the only mixed Nash equilibrium without altruism) to $\min\{6, 4, 3.697\} = 3.697$ for the worst mixed Nash equilibrium.

Best Pure/Mixed Nash Equilibrium.

Now we consider the Game 2 in Table 2.3 for the best pure/mixed Nash equilibrium. Player 1 playing strategy 1 and player 2 playing strategy 1 is the only pure Nash equilibrium, the social welfare of which is $3 + 1.2 = 4.2$. This is the best pure Nash equilibrium. With $\beta = 0.5$, the perceived payoff matrix becomes the one in Table 2.4. Then, there

Table 2.3: Payoff matrix for game 2 without altruism.

	Player 2's strategy 1	Player 2's strategy 2
Player 1's strategy 1	(3,1.2)	(0,0)
Player 1's strategy 2	(0,0)	(2.5,0.8)
Player 1's strategy 3	(2.9,1.5)	(2,2)

Table 2.4: Perceived payoff matrix for game 2 with uniform altruism $\beta = 0.5$.

	Player 2's strategy 1	Player 2's strategy 2
Player 1's strategy 1	(3.6,2.7)	(0,0)
Player 1's strategy 2	(0,0)	(2.9,2.05)
Player 1's strategy 3	(3.65,2.95)	(3,3)

is only one pure and mixed Nash equilibrium: player 1 plays strategy 3 and player 2 plays strategy 2, which has a social welfare of $2 + 2 = 4$. This is the best pure Nash equilibrium with $\beta = 0.5$. So, the social welfare of the best pure Nash equilibrium does get worse with altruism in Game 2.

For mixed Nash equilibria, without altruism, besides the only pure Nash equilibrium where player 1 plays strategy 1 and player 2 plays strategy 1, there is also another mixed Nash equilibrium: player 1 plays strategy 1 with probability 0.3 and strategy 3 with probability 0.7 while player 2 plays strategy 1 with probability 0.05 and strategy 2 with probability 0.95, which results in an expected social welfare of $0.3 \cdot 0.05 \cdot 4.2 + 0.3 \cdot 0.95 \cdot 0 + 0.7 \cdot 0.05 \cdot 4.4 + 0.7 \cdot 0.95 \cdot 4 = 2.877$. So, the best mixed Nash equilibrium has an expected social welfare of $\max\{2.877, 4.2\} = 4.2$. Thus, the expected social welfare still drops from 4.2 to 4 (player 1 playing strategy 3 and player 2 playing strategy 2 which is the only mixed Nash equilibrium with $\beta = 0.5$) for the best mixed Nash equilibrium.

We have seen that at least for some game instances, altruism does strictly worsen the (expected) social welfare for both pure and mixed Nash equilibria. We are therefore more interested in how the PoA/PoS for a class of games (i.e., the worst-case PoA/PoS value in a class of games) changes with altruism/spite. We will see two opposite trends in different classes of games in the following chapters.

Chapter 3

Related Work

We divide the related work into sections on different topics.

3.1 Models with Other-Regarding Payoff Functions

Questions about the accuracy of the assumption that users are selfish and rational have been as old as the field of game theory (see, e.g., [57]). Different models have been proposed to model user preferences more accurately.

Ledyard proposed a simple model of altruism [57] as an explanation for the results of public good contribution games, where a player's utility is a linear function of both the player's own monetary payoff and the other players' payoffs. However, this simple model is inadequate to explain some games still, such as ultimatum bargaining. Levine brought up a revised model to remedy this using *relative* spite and altruism [58], where the adjusted utility of a player reflects the player's own utility and his regard for other players. There are two coefficients: the coefficient of altruism tells spiteful players from altruistic players, and the other coefficient reflects the fact that players may have a higher regard for other altruistic players than spiteful ones. The model can be regarded as incorporating fairness, not in the sense that players have a particular target to be considered fair by them, but in the sense they are willing to be more altruistic to another player who is more altruistic towards them.

Fehr and Schmidt considered a model with an innate sense of fairness [33], where besides purely selfish players, there are players who dislike inequitable outcomes. First, they experience inequity if they are worse off in their material terms than the other players, and they also feel inequity if they are better off. Second, nevertheless, it is assumed that players suffer more from inequity that is to their material disadvantage than from inequity that is to their material advantage.

Geanakoplos et al. [41] use the approach of psychological game to model that players care not just about other players' utility, but also their attitudes towards other players depending on how they are treated. However, these models are complicated, and depart radically from decision theory.

There is work aiming at untangling the concepts of altruism, fairness, and reciprocity. Cox [28] designed experiments to identify the actions resulting from trust or reciprocity away from the actions resulting from altruistic or inequity-averse other-regarding preferences that are unconditional on the others' behavior. This is important in obtaining empirical data and information that can help constructing models that can increase the empirical validity of game theory. Charness and Rabin [18] designed a range of simple experiments to show that players are more concerned with increasing social welfare (sacrificing to increase all players' payoffs, especially for low-payoff players) than with reducing differences in payoffs as proposed in inequity-averse models, while they are also driven by reciprocity.

A model somewhat similar to ours was recently studied in the context of contributions to P2P systems by Feldman et al. [35], who posited an intrinsic *generosity* parameter of users, their willingness to contribute to the system. They then study contribution dynamics and their equilibria, akin to many collective behavior scenarios studied by Schelling [87].

While standard game theory addresses the way completely rational players operate, the field of behavioral game theory [17] uses psychological principles and numerous experiments involving human subjects to develop models and theories of reciprocity, limited strategizing, learning, etc. in order to come closer to the real-world human strategic behavior.

3.2 Externalities

The notion of *spite* and *altruism* as defined here broadly falls into the class of *allocation externalities* in auctions: the utility of a bidder depends not exclusively on her own allocation, but also on the allocations of other bidders. There is a large amount of literature on various types of allocation externalities (see, e.g., [49, 50, 51, 16]).

In particular, Jehiel et al. [50] construct revenue-maximizing auctions for the case where each potential buyer has a given constant externality depending on the identity of the winner. Thus, the difference to our model is that in the model of [50], a loser's utility does not depend on the *price* at which the winner obtained the object, only the winner's identity.

3.3 Selfish Routing

The study of the ineffectiveness of Selfish Routing was pioneered within the theory community by the groundbreaking work of Roughgarden and Tardos [86]. It was preceded by work in the economics and traffic engineering communities on congestion models, traffic routing, and the impact of tolls [75, 94, 8]. Since the original paper by Roughgarden and Tardos, a lot of progress has been made on different aspects of the problem, including different objectives [79], Stackelberg strategies in which an altruistic central authority controls a fraction of all traffic [54, 80, 90], the impact of tolls or taxes on the inefficiency [25, 26, 36, 37, 53], atomic games wherein users control non-infinitesimal amounts of traffic [27, 45, 82], and the effects of network structure on the inefficiency [60, 78, 83]. For an excellent overview of many of these results, see the book by Roughgarden [81].

Tradeoffs between individual optimization and social optimum in the context of traffic routing are also considered by Jahn et al. [48]. They posit that users will be willing to incur latency somewhat exceeding a "lowest possible" baseline if advised by a traffic routing system. They experimentally evaluate how centralized routing of users under this restriction compares with unrestricted centralized routing (which may place very heavy burdens on some users).

3.4 Stackelberg Routing

Among other things, our results draw a connection between Stackelberg strategies and tolls on users, in that the altruistic component of a user's utility can be considered as a (traffic-dependent) toll, and entirely altruistic users act as though they submitted to the control of a benevolent authority. Such Stackelberg routing strategies have been studied extensively. In general, the price of anarchy can still be unbounded, even for single-commodity flows where a central authority controls a large constant fraction of the traffic [12]. For linear latency functions, Karakostas and Kolliopoulos [54] recently showed an upper bound of $(4 - X)/3$ on the Price of Anarchy (where $X = \frac{(1-\sqrt{1-\lambda})(3\sqrt{1-\lambda}+1)}{2\sqrt{1-\lambda}+1}$) for arbitrary networks and commodities in which a central authority controls a λ fraction of traffic. For arbitrary latency functions in series-parallel networks, Swamy [90] bounds the price of anarchy by $1 + 1/\lambda$. For parallel link networks with latency functions from a class $\mathcal{C}$ with an upper bound $\rho(\mathcal{C})$ on the price of anarchy in Pigou examples, he shows an upper bound of $\lambda + (1 - \lambda)\rho(\mathcal{C})$.

In the context of Stackelberg routing, a converse direction has been studied by Sharma and Williamson [88] and Kaporis and Spirakis [52]. They ask how much traffic needs to be controlled by a central authority to guarantee any improvement in average latency [88] (called *Stackelberg threshold*) or to guarantee optimality of the resulting Nash Equilibrium [52] (called *Price of Optimum*).

3.5 General Equilibrium Concepts for Congestion Games

The notion of smoothness was proposed by Roughgarden [84]. The basic idea is to bound the sum of cost increases of individual players switching strategies by a combination of the costs of two states. Because these types of bounds capture local improvement dynamics, they bound the PoA not only for Nash equilibria, but also more general classes, including coarse correlated equilibria. The smoothness notion was recently refined in the local smoothness framework by Roughgarden and Schoppmann [85]. They require the types of bounds described above only for nearby states, thus obtaining tighter bounds, albeit only for more restrictive solution concepts and convex strategy sets. Using the local smoothness framework, they obtained optimal upper bounds for atomic splittable congestion games. Nadav and Roughgarden [71] showed that smoothness bounds apply all the way to a solution concept called average coarse correlated equilibrium, but not beyond.

A comparison between the costs in worst-case outcomes under solution concepts of different generality was recently undertaken by Bradonjic et al. [13] under the name "price of mediation": specifically for the case of symmetric singleton congestion games with convex latency functions, they showed that the ratio between the most expensive correlated equilibrium and the most expensive Nash equilibrium can grow exponentially in the number of players.

3.6 Player-Specific Congestion Games

Hoefer and Skopalik established the existence of pure Nash equilibria for several subclasses of atomic congestion games with non-uniform altruism [46]; for the generalization of arbitrary player-specific cost functions, Milchtaich [67] showed existence for singleton congestion games, and Ackermann et al. [1] for matroid congestion games, in which the strategy space of each player is the basis of a matroid on the set of resources.

3.7 Network Vaccination

A number of recent studies have analyzed the spread of worms or viruses on Internet-like topologies by focusing on characterizing the epidemic threshold (the transmission rate at which the disease goes from dying out quickly to infecting a large share of the network) for models such as small-world graphs [95] and preferential attachment models [7, 56]. The epidemic threshold is related to graph properties such as degree distribution, spectral radius and isoperimetric constants [24]. Based on these observations, Dezső and Barabási [31] suggest the vaccination of high-degree nodes in power-law random graphs as a way of increasing the epidemic threshold and thereby reducing the spread of diseases. Similar heuristics with analysis under random graph models with given degree distributions are also presented in [47].

In the context of network inoculation, the model of Aspnes et al. [4] has been extended in several ways. Meier et al. [66] consider the addition of friendship and show that friendship with neighbors can sometimes lead to significantly more efficient network inoculations. Moscibroda et al. [70] instead consider malicious Byzantine players who may misrepresent their actions with an intent to harm other players. (Naturally, this model is more suited to computer networks than social networks.) Perhaps surprisingly, such *malice* can sometimes lead to societally more desirable outcomes, due to the fear of other players. Recently, Diaz et al. [32] showed that the same "windfall of malice" can be achieved with a *mediator*. A mediator is a trusted third party that suggests actions to each player; the players retain free will and can ignore the mediator's suggestions. The concept of a mediator is closely related to that of a correlated equilibrium [5]. If a mediator recommends actions to the players so that it is in the best choice of each player to follow the mediator's recommendation, then the mediator is implementing a correlated equilibrium (defined in Section 2.1).

3.8 Auctions

The notion of individual spiteful behavior and models similar to the one we are proposing have recently been studied in the context of auctions. For single-item auctions, Brandt and Weiss [15] study the behavior of "antisocial" agents, whose utility decreases in their competitors' profit. They analyze full-information equilibria between two players, both of whom have spite level $\beta = -\frac{1}{2}$.

Morgan et al. [69] and Brandt et al. [14] focus on Bayesian Nash Equilibria of first-price and second-price auctions with uniform spite. The results in these two papers are very similar to each other, and differ mostly in the precise model of the utility of the winner, as discussed briefly in Section 7.1. Brandt et al. [14] study the Bayesian setting, and derive symmetric Bayesian-Nash equilibria for spiteful agents in first-price and second-price sealed bid auctions. They show that the expected revenue in secondnd-price auctions is higher than the expected revenue in first-price auctions when all agents are neither completely selfish nor completely spiteful. They also prove that in the presence of spite, complete information reduces the revenue in second-price auctions, while it increases the revenue in first-price auctions.

Vetsikas and Jennings [91] generalize some of these results for multi-unit auctions, still assuming uniform spite among the players, and deriving symmetric Bayes-Nash equilibria for spiteful agents in both m^{th} and $(m+1)^{\text{th}}$ price sealed bid auctions.

Similarly, Liang and Qi [59] study the effects of cooperative or vindictive bidding strategies on the revenue of sponsored search auctions and the existence of truthful strategies and equilibria.

A similar model is also studied in a recent paper by Deng and Qi [30] on auction design for pricing priority rights. Losers in this model also incur a negative utility, albeit one that depends on their own utility for the item, rather than the winner's. The goal in [30] is to design a truthful, egalitarian and budget-balanced auction.

3.9 Altruism and Spite in Game-Theoretic Settings

Several recent papers have analyzed the impact of spiteful or altruistic behavior in several game-theoretic settings with different models. Babaioff et al. [6] studied the impact of spiteful behavior on the outcome of routing games. In their model, there are two types of players: selfish rational players, and malicious players, who seek to maximize the average delay experienced by the rational players (while not caring about their own delay). They quantify the impact of malicious players on the equilibrium, and show that the price of anarchy can sometimes be increased, and in fact decreased at other times, which is similar to the "windfall of malice" shown in network vaccination [70].

Roth [77] considered the effect of malicious or Byzantine players on the PoA when each player has no regret in linear congestion games. Since his assumptions are strictly weaker than in previous work, the bounds proved on two measures of the price of malice hold also for the quantities studied by Babaioff et al. [6] and Moscibroda et al. [70].

Finally, mechanism design for spiteful players in scheduling is considered by Garg et al. [40]. They developed a strategy used by a spiteful agent to create losses to the other players, and analyze the effect of different levels of agents' spite on the losses caused to the other agents.

3.10 Game Theory and Economic/Social Networks

The impact of social network structure on games has recently been studied by Ashlagi et al. [3], under the name *social context games*. They posit that the utility of an agent can be computed from the subutility functions in her neighborhood, according to various competitive or collaborative aggregation functions. The specific games studied in [3] differ from all games that we considered here, and mostly belong to the class of resource selection games.

Chapter 4

Traffic Routing

Traffic routing is a prevalent problem in real-world transportation and computer networks. Traffic routing games or network congestion games are a special case of (general) congestion games, the atomic version of which will be discussed in the next chapter. The results in this chapter can be compared with those in the next chapter. From a technical point of view, it makes sense to study the non-atomic case before the atomic case since the former can be conceptually thought of as a case of the latter when the number of players goes to infinity. The results of this chapter are based on the paper [21].

4.1 Preliminaries

We start with the definition of a general congestion game. In a general congestion game, each player's strategy consists of a set of resources that he uses, and the cost of the strategy depends simply on the number of players using each resource. Formally, in a general congestion game $G = (N, E, \{c_e\}_{e \in E}, \{S_i\}_{i \in N})$, we are given a set of players $N = \{1, \ldots, n\}$, a set of *resources* E with *cost functions* $c_e : \mathbb{R} \rightarrow \mathbb{R}$ for every resource $e \in E$, and a strategy set $S_i \subseteq 2^E$ for every player $i \in N$. When S_i is the same for every player i, the game is *symmetric*; when $|S_i| = 1$ for all i, the game is *singleton*. For a joint strategy $\mathbf{s} \in S = S_1 \times ... \times S_n$, define $f_e(\mathbf{s}) = |\{i : e \in s_i\}|$ as the number of players using resource $e \in E$. The social cost function is $C(\mathbf{s}) = \sum_{i=1}^{n} c_i(\mathbf{s})$, where $c_i(\mathbf{s}) = \sum_{e \in s_i} c_e(f_e(\mathbf{s}))$ is the cost of player $i \in N$.

A (traffic) routing game or network congestion game is a special case of a general congestion game, where we are given a (directed) graph $G = (V, E)$. The set of resources is the set of edges E, and a strategy set consists of only paths between a source-sink pair. Each edge is equipped with a flow-dependent *latency function* $c_e(x)$. The meaning is that if the total flow on the edge e is x, then each user experiences a delay $c_e(x)$ on that edge. Thus, we can think of the total traffic as a multi-commodity flow with rates r_i of users between source-sink pairs (s_i, t_i), where the total flow from s_i to t_i is r_i. If f_e denotes the total flow on edge e, then the total latency experienced by a user on a path P is $c_P(f) := \sum_{e \in P} c_e(f_e)$. The total latency experienced by all users is thus $C(f) := \sum_e f_e \cdot c_e(f_e)$. An instance of the routing problem is thus a triple $(G, \mathbf{r}, \mathbf{c})$ (where $\mathbf{r}$ and $\mathbf{c}$ are the vectors of flow rates and edge cost functions). A symmetric singleton routing game or network congestion game therefore has a graph of single-commodity parallel links.

In this chapter, our model is based on the model of Wardrop [94], as described by Roughgarden and Tardos [81, 86], where users that are routed through G are *non-atomic*, i.e., infinitesimally small. We assume that each c_e is a continuously differentiable and monotone nondecreasing function. In addition, for some of our results, we will assume that each c_e is convex, and for others that each c_e is *semi-convex*, i.e., that $x \cdot c_e(x)$ is convex. The *socially optimum solution* for $(G, \mathbf{r}, \mathbf{c})$ is the flow f minimizing $C(f)$, and thus the solution to the convex program

$$\begin{array}{ll} \text{Minimize} & \sum_e f_e \cdot c_e(f_e) \\ \text{subject to} & f \text{ is a feasible multi-commodity flow for } (G, \mathbf{r}, \mathbf{c}). \end{array}$$

The constraints are the standard linear multi-commodity flow constraints; the objective function is convex so long as each c_e is semi-convex. Thus, the optimum can be computed in polynomial time using convex programming [73].

Selfish users do not care about the cost $C(f)$. Their sole goal is to select a path P minimizing their own latency $c_P(f)$. As the goals of different selfish users in minimizing their latency are conflicting with each other, the traffic routing problem can be considered a game, and the "outcome" of this game will be a (pure) *Nash Equilibrium* (Definition 2.1.2): a multi-commodity flow f such that, given f, no user has an incentive to choose a different path. Thus, a flow f is at Nash Equilibrium if and only if for each commodity i, all s_i-t_i paths P with $f_P > 0$ have the same latency $c_P(f)$, and all other s_i-t_i paths have at least the same latency. Nash Equilibria, too, can be computed as solutions to a convex program:

Proposition 4.1.1 ([81], Proposition 2.6.1) *The Nash flows of an instance* $(G, \mathbf{r}, \mathbf{c})$ *are exactly the solutions to the following convex program, and can thus be computed in polynomial time.*

$$\begin{array}{ll} \textit{Minimize} & \sum_e \int_0^{f_e} c_e(t)dt \\ \textit{subject to} & f \textit{ is a feasible multi-commodity flow for } (G, \mathbf{r}, \mathbf{c}). \end{array}$$

If f is a flow at Nash Equilibrium, and f^* the socially optimum flow, then the ratio $\rho(G, \mathbf{r}, \mathbf{c}) := C(f)/C(f^*)$ is the *Price of Anarchy* of the instance $(G, \mathbf{r}, \mathbf{c})$, capturing how much larger $C(f)$ can be than $C(f^*)$ (Definition 2.2.1). Roughgarden and Tardos [86] gave a generalization of Pigou's example [75] (see our different generalization of Pigou's example in Definition 4.2.2 for uniform altruism), showing that if the cost functions can be arbitrary, then the PoA is unbounded, even for networks consisting of two nodes and two parallel links. On the other hand, they proved that if all functions are linear $c_e(x) = a_e x + b_e$, then the Price of Anarchy is at most 4/3.

4.1.1 Altruism and Spite

As defined in Section 2.3, we posit that user i's utility is the combination $(1-\beta_i)p_i(\mathbf{a}) + \beta_i \sum_j p_j(\mathbf{a})$, where $\beta_i \in [-1, 1]$ is the user's altruism level. In order to apply this model to our scenario of traffic routing, we define the payoff of user i on path P as $p_i = -c_P(f)$, where f is the total flow, determined by the actions of all other players. Then, maximizing utility is equivalent to minimizing latency. The traffic routing model assumes that there are infinitely many users, each of whom is infinitesimally small. We can still define the utility function analogously, using the (well-defined) average latency of all users as the altruistic part. However, because users are infinitesimally small and latency functions continuous, the average latency of other users will not depend on an individual user's action. Thus, as long as $\beta \neq 1$, each partially altruistic user will act exactly like a selfish user. A natural model considering the effect the user has on others should instead be based on the *rate* at which the user's action will affect other users. We thus use the following definition of a β-altruistic user:

Definition 4.1.2 *Each β-altruistic user (for $\beta \in [-1,1]$) chooses a path P so as to minimize the* cost function

$$c_P^{(\beta)}(f) \quad := \quad (1-\beta)\sum_{e\in P} c_e(f_e) + \beta \sum_{e\in P}(f_e c_e(f_e))'.$$

The term $\sum_{e\in P} c_e(f_e)$ is the selfish *part of the cost, while $\sum_{e\in P}(f_e c_e(f_e))'$ is the* altruistic *part. $(f_e c_e(f_e))'$ denotes the derivative with respect to f_e. Notice that we can rewrite $c_P^{(\beta)}(f) = \sum_{e\in P} c_e(f_e) + \beta \sum_{e\in P} f_e c_e'(f_e)$.*

While our definition is motivated mathematically, there is a "psychological" interpretation of the underlying choice: in order to behave (partially) altruistically, infinitesimally small users must give infinitesimally small weight to their own payoff, which is achieved implicitly by making the altruistic component the derivative of the social welfare. For an infinitesimal user when entering or leaving an edge, it is not important and is negligible how much the absolute change in the total cost on that edge (for every other user on that edge) is. However, when comparing the change with the size of such an infinitesimal user, it is important how much the *rate* of the absolute change in the total cost on that edge is, meaning how much effect his arrival or leaving has on the other users sharing that edge. Intuitively, one interpretation of the derivative is that the infinitesimal user is projecting what would happen "if everyone acted the same way."

Definition 4.1.2 is similar to the definition of the valuation of a user with a time/money tradeoff of β in the case of network routing with tolls [25]. However, notice that unlike the standard model for tolls, the "edge toll" τ_e a user incurs in our model is *traffic-dependent*, namely $\tau_e := f_e c_e'(f_e)$.

We say that the users are *uniformly β-altruistic* if all users are β-altruistic. More generally, we allow for the case of arbitrary distributions of altruism among the users. In the general case, for each commodity i, we are given an arbitrary *altruism density function* ψ_i on the interval $[-1, 1]$. We only require that all these functions ψ_i be indeed distributions, i.e., forming a Borel measure of total measure 1. If the rate for commodity i is r_i, then the overall altruism density function is $\psi = \frac{1}{\sum_i r_i} \sum_i r_i \psi_i$. The *average altruism* of a distribution ψ is then $\int_{-1}^{1} t\psi(t)dt$. An instance of the partially altruistic traffic routing problem is thus the quadruple $(G, \mathbf{r}, \mathbf{c}, (\psi_i))$. If there is a single commodity with distribution ψ, we write $(G, \mathbf{r}, \mathbf{c}, \psi)$, and if the altruism is uniform, we simplify further to $(G, \mathbf{r}, \mathbf{c}, \beta)$.

Proposition 4.1.3 *Let $(G, \mathbf{r}, \mathbf{c}, \beta)$ be an instance with uniform altruism $\beta \geq 0$. Then, the Nash flows are the optima of the convex program*

$$\begin{array}{ll} \textit{Minimize} & \sum_e \int_0^{f_e} c_e^{(\beta)}(t)dt \\ \textit{subject to} & f \textit{ is a feasible multi-commodity flow for } (G, \mathbf{r}) \end{array}$$

In particular, the instance $(G, \mathbf{r}, \mathbf{c}, \beta)$ always possesses a Nash Equilibrium for $\beta \geq 0$.

The proof of this proposition is virtually identical to that of Proposition 2.6.1 from [81]. The proof there only used the fact that each agent was minimizing a sum of monotone increasing functions $\sum_e g_e(f_e)$ to conclude that the Nash Equilibrium was the flow minimizing the (convex) objective $\sum_e \int_0^{f_e} g_e(t)dt$. Thus, it applies equally to $g_e(t) := c_e^{(\beta)}(t)$.

We also have the following variational inequality to characterize the Nash flows.

Proposition 4.1.4 *Let* $(G, \mathbf{r}, \mathbf{c}, \beta)$ *be an instance with uniform altruism* $\beta \geq 0$*. Then,* f *is a Nash flow for* β*-altruistic users if and only if it minimizes* $\sum_P c_P^{(\beta)}(f)\tilde{f}_P$ *over all feasible flows* $\tilde{f}$*.*

Proof. By fixing Nash flow f, $\sum_P c_P^{(\beta)}(f)\tilde{f}_P$ is the social cost of a feasible flow $\tilde{f}$, where the latency of each path P is the congestion-independent constant $c_P^{(\beta)}(f)$. If a flow f is at Nash equilibrium for β-altruistic users, then a user of commodity i routes to minimize $c_P^{(\beta)}(f)$ over $s_i - t_i$ paths P so $\sum_P c_P^{(\beta)}(f)\tilde{f}_P$ is minimized over all feasible flows $\tilde{f}$. Conversely, if a flow $\tilde{f}$ is not at Nash equilibrium for β-altruistic users, a user can decrease his (perceived) path cost at $\tilde{f}$ by deviation. Thus, $\sum_P c_P^{(\beta)}(f)\tilde{f}_P$ (or, equivalently, $\sum_e c_e^{(\beta)}(f_e)\tilde{f}_e$) is not minimized. ■

The situation is not quite as straightforward for the case of non-uniform altruism distributions ψ, or for negative β. Even for two different values of altruism, there appears to be no natural convex programming formulation for Nash Equilibria. However, using a theorem of Mas-Collel [64], we can still prove the existence of Nash Equilibria.

Theorem 4.1.5 *Each instance* $(G, \mathbf{r}, \mathbf{c}, (\psi_i))$ *has a Nash Equilibrium.*

Proof. Theorem 1 of Mas-Collel [64] proves that each game of infinitely many players has a Nash Equilibrium. A game is characterized by a distribution (Borel measure) over utility functions which are continuous in the action of the player, and the distribution of actions by the remaining players. It is easy to see that each player in the routing game has a utility function $-c_P^{(\beta)}(f)$ continuous in the choice of path P (trivially, since the space of all simple s_i-t_i paths is finite) and in the distribution of other players' strategies f (by continuity of each c_e). The utility for paths not connecting s_i to t_i is $-\infty$ (or an appropriately negative constant). The distribution of altruism values β implies a corresponding distribution over utility functions. Thus, the theorem of Mas-Collel implies the existence of Nash Equilibria for routing games. ■

The proof by Mas-Collel is inherently non-constructive; accordingly, Theorem 4.1.5 does not imply any algorithm for finding such equilibria. Since there always exists a (pure) Nash Equilibrium of instances $(G, \mathbf{r}, \mathbf{c}, (\psi_i))$, we can define the (pure) Price of Anarchy (PoA) with altruism distributions (ψ_i), as $\rho(G, \mathbf{r}, \mathbf{c}, (\psi_i)) = C(f)/C(f^*)$, where f is a Nash flow for $(G, \mathbf{r}, \mathbf{c}, (\psi_i))$, and f^* a socially optimal flow for $(G, \mathbf{r}, \mathbf{c})$. (General definitions can be found in Sections 2.1 and 2.2.)

4.1.2 Taxes and Stackelberg Strategies

Our definition of partial altruism naturally relates to two strategies that have been proposed in the literature for dealing with the selfishness of users: Pigou taxes and Stackelberg strategies.

The idea of taxes or tolls on edges is to charge users a fee for using an edge. The assumption is that money and latency can be measured on the same scale, and users will minimize the (weighted) sum of the two. It is well-known [75] that if the toll charged on each edge e equals the *marginal cost* to others ($f_e^* c_e'(f_e^*)$) at the optimum solution, then the Nash Equilibrium will minimize $C(f)$, i.e., be optimal. Our model of partial altruism can thus be interpreted as charging users a *traffic-dependent* constant fraction of the marginal tax, i.e., with respect to the *current* flow. When the altruism is not uniform, different users will be charged different taxes $\beta_i f_e c_e'(f_e)$ on edges. Our model can thus be considered as investigating the Price of Anarchy when different users have different tradeoffs between taxes and latency, but their tradeoff stays constant across different edges. Similar models were considered, e.g., in [29, 89]. Cole et al. [26, 25] also study optimization problems arising from non-uniform taxation in networks. However, their goal is to minimize the total tolls, subject to forcing the flow to be optimal, whereas we study the Price of Anarchy given the taxation scheme of charging a (user-dependent) fraction of the marginal tax on each edge.

A different strategy for lowering the Price of Anarchy is available when a benevolent central authority controls a λ fraction of the total traffic. The central authority's goal is to route this fraction so as to minimize the total cost $C(f)$, subject to the fact that the remaining users will subsequently route their traffic selfishly. Algorithms for routing flows with this objective are called *Stackelberg strategies*, and the corresponding asymmetric games *Stackelberg games* (see, e.g., [80]).

When the altruism distribution has support $\{0, 1\}$, and the cumulative distribution function of ψ is the step function whose value at 0 is $1 - \lambda$, and whose value at 1 is 1, the altruistic users can be interpreted as a central authority, and their flow as a Stackelberg strategy with the corresponding Price of Anarchy. When the central authority controls a λ fraction of the traffic, the average altruism is exactly λ, and thus, any bound on the Price of Anarchy for average altruism λ gives rise to the same bound for Stackelberg routing. Notice that the converse is not necessarily true: at the moment, it is not known if every optimal Stackelberg strategy gives rise to a Nash Equilibrium of the routing game with altruism support $\{0, 1\}$.

4.2 Uniform Altruism

In this section, we focus on the model of uniformly altruistic users: each user is β-altruistic for $-1 \leq \beta \leq 1$. Thus, the perceived cost of an edge e to the user is $c_e^{(\beta)}(x) = (1-\beta)c_e(x) + \beta \frac{d}{dx}(xc_e(x)) = c_e(x) + \beta x c_e'(x)$. (Notice that for $\beta = 0$, this coincides with selfishness; $\beta = 1$ corresponds to complete altruism, and $c_e^{(1)}(x)$ is exactly the marginal cost of e. For $\beta = -1$, the users are completely spiteful.) Our first result follows directly from the definitions of flows at Nash equilibrium and optimum, and gives a (tight) upper bound on the Price of Anarchy for arbitrary networks, commodities, and arbitrary semi-convex cost functions.

Proposition 4.2.1 *If all cost functions c_e are nondecreasing and semi-convex, then for all networks G and flow rates r, and any altruism level $\beta \in (0, 1]$,*

$$\rho(G, \mathbf{r}, \mathbf{c}, \beta) \quad \leq \quad 1/\beta.$$

Proof. Let $\hat{f}$ be a Nash Equilibrium flow, minimizing the potential function $\Phi(f) = \sum_e \int_0^{f_e} c_e^{(\beta)}(t)dt$, the objective function of the convex program in Proposition 4.1.3. Also, let f^* the optimum flow, minimizing the total cost $C(f) = \sum_e \int_0^{f_e} (tc_e(t))'dt$. Simply from the definition of $c_e^{(\beta)}(t)$, it follows that for any flow f, we have $\Phi(f) \leq C(f) \leq \frac{1}{\beta}\Phi(f)$. Applying the first inequality to f^* and the second to $\hat{f}$, and using the optimality of $\hat{f}$ for Φ, we obtain $C(\hat{f}) \leq \frac{1}{\beta}\Phi(\hat{f}) \leq \frac{1}{\beta}\Phi(f^*) \leq \frac{1}{\beta}C(f^*)$. ■

More generally, we derive a result bounding the Price of Anarchy when all cost functions c_e are drawn from a given class of cost functions. Our characterization will be in terms of the *anarchy value* $\alpha^{(\beta)}(\mathcal{C})$ of a set $\mathcal{C}$ of functions for β-altruistic users, which is defined as a generalization of the anarchy value of functions in [81].

Definition 4.2.2 *1. For any cost function c, the anarchy value $\alpha^{(\beta)}(c)$ of c for β-altruistic users is defined as*

$$\alpha^{(\beta)}(c) \quad = \quad \sup_{r,x\geq 0} \frac{r\cdot c(r)}{x\cdot c(x)+(r-x)\cdot c^{(\beta)}(r)},$$

where $0/0$ is defined to 1.

2. For any class $\mathcal{C}$ of cost functions, the anarchy value for β-altruistic users $\alpha^{(\beta)}(\mathcal{C})$ is

$\sup_{c\in\mathcal{C},c\neq 0}\alpha^{(\beta)}(c)$.

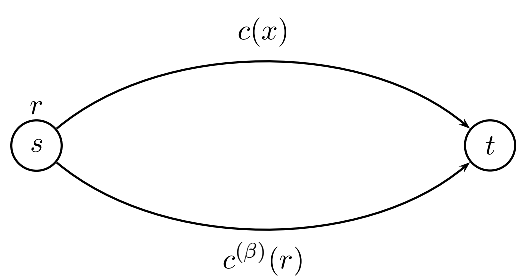

Figure 4.1: Worst-case graphs for uniformly β-altruistic users with rate r

The motivation for this definition of $\alpha^{(\beta)}(c)$ is that it captures the Price of Anarchy for uniformly β-altruistic users in a two-node two-link network, where one link has latency function c and the other has a worst-case constant. See Figure 4.1. Indeed, we will prove this to be the case in Lemma 4.2.7 below. Notice that Lemma 4.2.7 immediately implies that $\alpha(\mathcal{C})$ is a lower bound on the Price of Anarchy in the worst case when all edge latency functions are chosen from $\mathcal{C}$. Our main theorem in this section shows that it is also an upper bound for all networks and arbitrary commodities.

We are mostly interested in $\alpha^{(\beta)}(\mathcal{C})$ when it is finite. In particular, this suggests defining the *spite resistance of $\mathcal{C}$ as the least altruistic behavior that $\mathcal{C}$ could support.* Formally, $b_c = \inf\{\beta \mid \alpha^{(\beta)}(c) < \infty\}$, and $b_{\mathcal{C}} = \inf_{c\in\mathcal{C}} b_c$. It is not difficult to show that $b_c = -\inf_r \frac{c(r)}{rc'(r)}$, and that $\alpha^{(\beta)}(c) = \infty$ for $\beta \leq b_c$. Using L'Hôpital's rule, one sees that the monotonicity and convexity of c imply that $b_c \geq -\lim_{r\to\infty} \frac{c(r)}{rc'(r)} \geq -1$ for all c, which also motivates our earlier restriction to altruism values $\beta \geq -1$.

Theorem 4.2.3 *Let $\mathcal{C}$ be a set of cost functions, and $(G, \mathbf{r}, \mathbf{c})$ an instance with cost functions $c_e \in \mathcal{C}$. Then,*

$$\rho(G, \mathbf{r}, \mathbf{c}, \beta) \quad \leq \quad \alpha^{(\beta)}(\mathcal{C}).$$

Proof. Fix an instance $(G, \mathbf{r}, \mathbf{c})$ with cost functions $c_e \in \mathcal{C}$. Let f^* be an optimal flow and f a Nash flow for β-altruistic users. By rearranging Definition 4.2.2, we obtain the bound $x \cdot c_e(x) \geq \frac{r\cdot c_e(r)}{\alpha^{(\beta)}(\mathcal{C})} + (x-r)\cdot c_e^{(\beta)}(r)$ for any $x, r \geq 0$. Applying this bound to each edge e, with $x = f_e^*$ and $r = f_e$, we bound

$$\begin{aligned} C(f^*) &= \sum_{e\in E} f_e^* c_e(f_e^*) \\ &\geq \frac{1}{\alpha^{(\beta)}(\mathcal{C})} \cdot \sum_{e\in E} f_e c_e(f_e) + \sum_{e\in E} (f_e^* - f_e)\cdot c_e^{(\beta)}(f_e) \\ &= \frac{C(f)}{\alpha^{(\beta)}(\mathcal{C})} + \sum_{e\in E}(f_e^* - f_e)\cdot c_e^{(\beta)}(f_e). \end{aligned}$$

It remains to show that $\sum_e f_e^* \cdot c_e^{(\beta)}(f_e) \geq \sum_e f_e \cdot c_e^{(\beta)}(f_e)$. To this end, applying the variational inequality in Proposition 4.1.4 to f and f^* proves the desired inequality. ■

As a corollary of Theorem 4.2.3, we can obtain a tight bound in the case where the cost functions are polynomials of degree at most p with non-negative coefficients. We denote this class by $\mathcal{C}_p$.

Theorem 4.2.4 *If* $(G, \mathbf{r}, \mathbf{c})$ *has cost functions in* $\mathcal{C}_p$*, then for any altruism value* $\beta \in (-1/p, 1]$,

$$\rho(G, \mathbf{r}, \mathbf{c}, \beta) \quad \leq \quad \left((\tfrac{1+\beta p}{1+p})^{1/p} (\tfrac{1+\beta p}{1+p} - 1 - \beta p) + 1 + \beta p \right)^{-1}.$$

Proof. First, notice that $b_{\mathcal{C}_p} = -1/p$. It can be easily verified that all subsequent calculations stay valid for $\beta > -1/p$, while for $\beta \leq -1/p$, the Price of Anarchy is unbounded.

As observed in [81], it suffices to focus only on polynomials $c(x) = ax^i$ with $x \leq p$. For any instance $(G, \mathbf{r}, \mathbf{c})$ with arbitrary polynomials can be equivalently transformed into one with only such monomials, by replacing each edge with cost function $c_e(x) = \sum_{i=0}^{p} a_i x^i$ by a directed path of $p+1$ edges, the i^{th} edge of which has cost function $\tilde{c}_{e,i}(x) = a_i x^i$. In order to compute the anarchy value $\alpha(c)$ of a nonzero polynomial function $c(x) = ax^i$, we use the equivalent characterization that

$$\alpha^{(\beta)}(c) \quad = \quad \sup_{r \geq 0} \left(\tfrac{\lambda c(\lambda r)}{c(r)} + (1-\lambda)(1 + \tfrac{\beta r c'(r)}{c(r)}) \right)^{-1},$$

where $\lambda \in [0,1]$ solves $c^{(1)}(\lambda r) = c^{(\beta)}(r)$, and 0/0 is defined to 1. To prove this equivalent characterization, we first observe that

$$\tfrac{d}{d\lambda}(c(\lambda r)\lambda r + c^{(\beta)}(r)(r-\lambda r)) \;=\; c^{(1)}(\lambda r)r - c(r)r \;=\; 0,$$

so there is indeed a value of $\lambda \in [0,1]$ solving $c^{(1)}(\lambda r) = c^{(\beta)}(r)$. By Lemma 4.2.7 below, $\alpha^{(\beta)}(c)$ is the price of anarchy in a two-node two-link network, one of whose links has the cost function $c(x)$, the other link having constant cost $c^{(\beta)}(r)$. Routing λr units of flow on the link with cost $c(x)$, and the rest on the link with cost $c^{(\beta)}(r)$, provides an optimal flow, while the Nash Equilibrium has all of its flow on the link with cost $c(x)$. Thus, the ratio of the cost of a Nash flow to that of an optimal flow is

$$\tfrac{r\cdot c(r)}{c(\lambda r)\lambda r + (c(r)+\beta r c'(r))(r-\lambda r)} \;=\; \left(\tfrac{\lambda c(\lambda r)}{c(r)} + (1-\lambda)(1+\tfrac{\beta r c'(r)}{c(r)})\right)^{-1}.$$

Solving for λ in the special case $c(x) = ax^i$, we obtain $\lambda = (\frac{1+\beta i}{1+i})^{1/i}$, and thus $\frac{c(\lambda r)}{c(r)} = \frac{1+\beta i}{1+i}$ and $\frac{c'(r)}{c(r)} = \frac{i}{r}$. Then,

$$\alpha^{(\beta)}(c) \;=\; \left((\tfrac{1+\beta i}{1+i})^{1/i}(\tfrac{1+\beta i}{1+i} - 1 - \beta i) + 1 + \beta i\right)^{-1},$$

which is independent of a and increasing in i (by a derivative test). Hence, the largest $\alpha(c)$ is attained for $c = x^p$, giving

$$\alpha^{(\beta)}(\mathcal{C}_p) \;=\; \left((\tfrac{1+\beta p}{1+p})^{1/p}(\tfrac{1+\beta p}{1+p} - 1 - \beta p) + 1 + \beta p\right)^{-1},$$

as claimed. ∎

It is not difficult to verify that the previous bound converges to $\frac{1}{\beta}$ as $p \to \infty$; the worst case behavior is in fact attained with polynomials of high degree. However, for $p = 1$, Theorem 4.2.4 also allows us to obtain a tighter bound in the special case that all latency functions are linear.

Corollary 4.2.5 *If* $(G, \mathbf{r}, \mathbf{c})$ *has linear cost functions, then for any* $\beta \in (-1, 1]$,

$$\rho(G, \mathbf{r}, \mathbf{c}, \beta) \quad \leq \quad \frac{4}{3+2\beta-\beta^2}.$$

Notice that for any $\beta > 0$, this bound improves on the bound by Roughgarden and Tardos [86] of 4/3 when all users are completely selfish. As the bound can also be shown to be tight, it thus characterizes exactly the gain by partial positive altruism with linear cost functions, and the spite resistance of linear cost functions. In particular, it shows that linear costs have the highest spite resistance among all classes of cost functions.

Remark 4.2.6 *Our results in this section extend straightforwardly to general non-atomic congestion games (not necessarily network congestion games), so long as all cost functions are nondecreasing. The perceived cost of player i's strategy* s_i *with altruism* β *is* $c_i^{(\beta)}(\mathbf{s}) = \sum_{e \in s_i} c_e^{(\beta)}(f_e) = \sum_{e \in s_i} c_e(f_e) + \beta f_e c_e'(f_e)$, *where* $\mathbf{s} = (s_1, ..., s_n)$ *is a strategy profile,* f_e *is the total number of players using resource* e, *and* c_e *is a nondecreasing function. With the same definitions of* $\alpha^{(\beta)}(\mathcal{C})$, *the proofs of the above proposition and theorems naturally carry over to this more general setting.*

Finally, we show that the bounds derived in Theorem 4.2.3 are indeed tight, even for two-node two-link networks:

Lemma 4.2.7 *Consider a two-node two-link network with flow rate* $r = 1$*, and cost functions* $c_1(x) = c(x)$ *on the first link, and constant cost function* $c_2(x) = c^{(\beta)}(r) = c(r) + \beta r c'(r)$ *for the second link . Then, the Price of Anarchy of this instance is* $\alpha^{(\beta)}(c)$*.*

Proof. It is easy to observe from the definition of c_2 that all β-altruistic users will end up using link 1, so that the total cost of the Nash Equilibrium is $c(1) = rc(r)$, while the socially optimum solution has total cost $\inf_{x \leq 1}(x \cdot c(x) + (r - x) \cdot c(r) + \beta(r - x)c'(r))$. Hence, the price of anarchy is exactly $\alpha(c)$. ■

By applying this characterization together with Theorem 4.2.4 and letting the degree of the polynomial go to ∞, we obtain instances $(G, \mathbf{r}, \mathbf{c})$ whose Price of Anarchy approaches $1/\beta$ arbitrarily closely. Similarly, by choosing $p = 1$, we obtain that the bound in Corollary 4.2.5 is tight.

4.3 Non-Uniform Altruism

In this section, we extend our results to the more general and realistic case where different users can have different altruism levels. In the most general case, we are given a distribution ψ of altruism. The existence of Nash Equilibria in this model was shown non-constructively as Theorem 4.1.5. Even for a single commodity and an altruism distribution with support $\{0, 1\}$, and arbitrarily large constant $\bar{\beta}$, a recent result on Stackelberg routing due to Bonifaci et al. [12] shows that the Price of Anarchy can become unbounded.

Thus, we focus here on the case of single-commodity traffic in *parallel link networks*. Parallel link networks have been studied by Roughgarden [80]; among others, they naturally model the assignment of infinitesimally small jobs to machines with load-dependent latencies. Formally, a parallel link network has two nodes s, t, and m parallel edges $e_1, \ldots, e_m$ from s to t. Our main theorem in this section gives a (tight) upper bound on the Price of Anarchy in the presence of partial altruism for single commodity parallel link networks and arbitrary (convex) cost functions. To make the argument work, we need our set of cost functions $\mathcal{C}$ closed under addition of constants.

Theorem 4.3.1 *If all cost functions c_e are convex and nondecreasing, then for parallel link networks G and flow rates r, and any overall altruism density function ψ with nonnegative support and average altruism $\bar{\beta}$,*

$$\rho(G, \mathbf{r}, \mathbf{c}, \psi) \quad \leq \quad 1/\bar{\beta}.$$

We will prove Theorem 4.3.1 as a corollary of the following more general result, bounding the price of anarchy in terms of the set of functions permissible as edge latencies.

Theorem 4.3.2 *If all cost functions $c_e \in \mathcal{C}$ are convex and nondecreasing, then for parallel link networks G and flow rates r, and any overall altruism density function ψ with non-negative support,*

$$\rho(G, \mathbf{r}, \mathbf{c}, \psi) \quad \leq \quad \left(\int_0^1 \psi(t) \frac{1}{\alpha^{(t)}(\mathcal{C})} dt\right)^{-1}.$$

Proof. Let f denote the flow at Nash Equilibrium. We first show that without loss of generality, we can assume that each link e contains only one type of users (i.e., if users have different altruism values β, β', then they do not share a link) and that the support of ψ is finite. To see this, assume that f has users of altruism values $\beta < \beta'$ sharing an edge e. Now replace all users on e with altruism β by users with altruism β'. f must still be a flow at Nash Equilibrium for the new instance (because β'-altruistic users are on link e in Nash Equilibrium). By repeating this process, we eventually obtain an instance with altruism density ψ' which stochastically dominates ψ and has finite support. For this new ψ', the bound on the Price of Anarchy for f provided by the right-hand side of Theorem 4.3.2 can only be smaller, giving us an even better bound than required. Thus, we can from now on focus on the case described above.

Let $0 \leq \beta_1 < \beta_2 < \ldots < \beta_k \leq 1$ be the (finite) support of ψ, where the rate of β_i-altruistic users is r_i (so $\sum_{i=1}^{k} r_i = r$). We need to show that for all flows g of rate r (in particular the optimum flow), we have

$$C(g) \quad \geq \quad (\sum_{i=1}^{k} \frac{r_i}{r} \frac{1}{\alpha^{(\beta_i)}(\mathcal{C})}) \cdot C(f), \tag{4.1}$$

which we will do by induction on k. The base case $k = 0$ is of course trivial.

For the inductive step, let f be a Nash Equilibrium flow, and g any flow of rate r. For each i, let E_i be the set of edges with positive flow of β_i-altruistic users under f. Notice that by our assumption, the sets E_i are disjoint. For any set E' of edges, let $f(E') = \sum_{e \in E'} f_e$ (similarly, $g(E')$) denote the total flow on E'. Let $E' := E \setminus E_1$ denote the set of all edges not used by β_1-altruistic users.

Intuitively, because the more altruistic users prefer the edge set E' over E_1, we would expect a "good" flow g to do the same. Indeed, we first show that the latency under f on all edges in E_1 is no larger than in E', while the marginal cost, i.e., $(x \cdot c(x))'$, is no larger in E' than in E_1. Let $e \in E_1, e' \in E_j, j > 1$ be arbitrary links with positive flow f. Thus, all users on e have altruism β_1, while all users on e' have altruism $\beta_j > \beta_1$. Because f is at Nash Equilibrium,

$$c_e(f_e) + \beta_1 f_e c'_e(f_e) \leq c_{e'}(f_{e'}) + \beta_1 f_{e'} c'_{e'}(f_{e'}), \tag{4.2}$$

$$c_e(f_e) + \beta_j f_e c'_e(f_e) \geq c_{e'}(f_{e'}) + \beta_j f_{e'} c'_{e'}(f_{e'}). \tag{4.3}$$

Combining appropriately scaled versions of Inequality (4.2) and Inequality (4.3) gives us that

$$c_e(f_e) \leq c_{e'}(f_{e'}), \tag{4.4}$$

$$(1-\xi)c_e(f_e) + (\beta_j - \xi\beta_1) f_e c'_e(f_e) \geq (1-\xi)c_{e'}(f_{e'}) + (\beta_j - \xi\beta_1) f_{e'} c'_{e'}(f_{e'}), \tag{4.5}$$

where ξ for $0 \leq \xi \leq 1$ is a scalar, which we can set later.

Our high-level strategy will be to bound the Nash Equilibrium flow on E' against a restriction g' of g of rate $r - r_1$ on E' by induction, and use a comparison argument for the flow on E_1. We will construct a flow h of rate r_1 whose cost is cheaper than a component of g of the same rate, and which is optimal for modified "residual" edge costs. We can thus compare it against the flow f on E_1 using Theorem 4.2.3.

Define f' to be the restriction of f to the edge set E', i.e., $f'_e = f_e$ for $e \in E'$, and $f'_e = 0$ for $e \in E_1$. Thus, f' is a flow of rate $r' := r - r_1$. Define the modified cost function $\tilde{c}_e(x) := c_e(f'_e + x) + \beta_1 f'_e c'_e(f'_e)$ for all edges e. Thus, $\tilde{c}_e(x)$ is the cost incurred by flow on e if f'_e is unalterable, but not considered part of the actual flow, plus a suitable constant term to "mimic" the altruistic component. This definition of $\tilde{c}_e(x)$ implies that the perceived cost of edge e to β_1-altruistic users is $\tilde{c}_e^{(\beta_1)}(x) = c_e(f'_e+x)+\beta_1 x c'_e(f'_e+x)+\beta_1 f'_e c'_e(f'_e)$. Thus, for $e \in E'$, we have that $\tilde{c}_e^{(\beta_1)}(x) \geq c^{(\beta_1)}(f'_e)$ for all $x \geq 0$, while for $e \in E_1$, because $f'_e = 0$, $\tilde{c}_e^{(\beta_1)}(x) = c^{(\beta_1)}(x+f'_e)$. In particular, this implies that the β_1-altruistic users are at Nash Equilibrium with respect to the modified cost functions $\tilde{c}_e(x)$. Hence, by Theorem 4.2.3, and because $\tilde{c}_e(x) = c_e(x)$ for all $e \in E_1$, we get $C(f - f') = \tilde{C}(f - f') \leq \alpha^{(\beta_1)}(\mathcal{C}) \cdot \tilde{C}(\tilde{f})$ where $\tilde{f}$ is an optimum flow of rate r_1 with respect to the modified edge cost functions $\tilde{c}_e$.

In order to compare f' against the part of g on the edge set E', it will be useful to assume that $g(E') \geq f(E')$. We will show next that we can make this assumption w.l.o.g. For assume that it did not hold. Then, let $e \in E_1, e' \in E'$ be edges with $g_e > f_e > 0$ and $g_{e'} < f_{e'}$. (The existence of e, e' follows from the assumption $g(E') < f(E')$). By the bound on the derivatives in Inequality (4.5), and using the convexity of the edge latency functions, we show that $(g_e c_e(g_e))' \geq (f_e c_e(f_e))' \geq (f_{e'} c_{e'}(f_{e'}))' \geq (g_{e'} c_{e'}(g_{e'}))'$ in the following. The first and last inequalities hold simply by the semi-convexity of c_e and $c_{e'}$, and the second equality is obtained by setting $\xi = \frac{1-\beta_j}{1-\beta_1}$ to get $1 - \xi = \beta_j - \xi\beta_1$ so $(f_e c_e(f_e))' = c_e(f_e) + f_e c'_e(f_e) \geq c_{e'}(f_{e'}) + f_{e'} c'_{e'}(f_{e'}) = (f_{e'} c_{e'}(f_{e'}))'$. Thus, g can be made cheaper by moving some of its flow from e to e'. By repeating this process, we can thus assume that $g(E') \geq f(E')$.

Let γ be such that $C(f - f') = \gamma C(f)$. Because f' and $f - f'$ use disjoint edge sets, we get $C(f') = (1-\gamma)C(f)$. (Notice that the assumption of disjoint edge sets is indeed crucial here. Due to the non-constant cost of edges, in general, it does not hold that $C(f) + C(f') = C(f + f')$.)

By Lemma 4.3.3 below, we can decompose $g = h + g'$, where g' is a flow of rate r' entirely on E', and h is a flow of rate r_1 satisfying the property (4.7), namely $\tilde{C}(\tilde{f}) \leq \sum_e h_e c_e(g_e) + \sum_e g'_e(c_e(g_e) - c_e(g'_e))$. We can thus apply induction on the flows f' and g' of rate r' on the modified graph with edge set E'. Notice that while f' may not be an Equilibrium flow on E, it is indeed an Equilibrium flow on E'. Thus, we obtain that

$$\begin{aligned} C(g) &= C(g') + \sum_e h(e)c_e(g_e) + \sum_e g'_e(c_e(g_e) - c_e(g'_e)) \\ &\geq (\sum_{i=2}^{k} \frac{r_i}{r'} \frac{1}{\alpha^{(\beta_i)}(\mathcal{C})}) \cdot C(f') + \frac{1}{\alpha^{(\beta_1)}(\mathcal{C})} C(f - f') \qquad (4.6) \\ &= \left((\sum_{i=2}^{k} \frac{r_i}{r'} \frac{1}{\alpha^{(\beta_i)}(\mathcal{C})}) \cdot (1-\gamma) + \frac{1}{\alpha^{(\beta_1)}(\mathcal{C})} \cdot \gamma \right) \cdot C(f). \end{aligned}$$

We next show that $\gamma \leq \frac{r_1}{r}$. By Inequality (4.4), every user on E_1 incurs lower delay than every user on E_j, and consequently on E'. Thus, the average delay $\frac{1}{r_1}C(f - f')$ of users on E_1 is at most the average delay $\frac{1}{r}C(f)$ of all users, so $C(f - f') \leq \frac{r_1}{r}C(f)$.

The lower bound (4.6) is a convex combination of the non-negative terms $\sum_{i=2}^{k} \frac{r_i}{r'} \frac{1}{\alpha^{(\beta_i)}(\mathcal{C})}$ and $\frac{1}{\alpha^{(\beta_1)}(\mathcal{C})}$, with coefficients $(1-\gamma)$ and γ. The anarchy value $\alpha^{(\beta)}(\mathcal{C})$ is a monotone non-increasing function of β, so the weighted average reciprocal anarchy value for altruism levels $\beta_2, \ldots, \beta_k$ is at least the reciprocal for β_1. Thus, the convex combination is minimized when the coefficient γ of the smaller term $\frac{1}{\alpha^{(\beta_1)}(\mathcal{C})}$ is as large as possible, i.e., when $\gamma = r_1/r$. Substituting this bound,

$$\begin{aligned} C(g) &\geq ((\sum_{i=2}^{k} \frac{r_i}{r'} \frac{1}{\alpha^{(\beta_i)}(\mathcal{C})}) \cdot \frac{r'}{r} + \frac{1}{\alpha^{(\beta_1)}(\mathcal{C})} \cdot \frac{r_1}{r}) \cdot C(f) \\ &= (\sum_{i=1}^{k} \frac{r_i}{r'} \frac{1}{\alpha^{(\beta_i)}(\mathcal{C})}) \cdot C(f), \end{aligned}$$

completing the inductive step, and thus the proof. ■

Lemma 4.3.3 *Let f' be a flow of rate r' using only edges from E', and define $\tilde{c}_e(x) := c_e(f'_e + x) + \beta_1 f'_e c'_e(f'_e)$. Let g be any flow of rate $r = r' + r_1$, with $g(E') \geq r'$. Let $\tilde{f}$ be the optimum flow of rate r_1 with respect to edge costs $\tilde{c}_e$. Then, g can be decomposed as $g = h + g'$, where g' is a flow of rate r' on E', satisfying*

$$\tilde{C}(\tilde{f}) \leq \sum_e h_e c_e(g_e) + \sum_e g'_e(c_e(g_e) - c_e(g'_e)). \tag{4.7}$$

Proof. Let $\Delta := g(E') - r' \geq 0$ be the amount of "excess flow" that g sends on E', compared to f. We begin by setting $h_e = g_e$ for all edges $e \in E_1$, giving us a flow of rate $r_1 - \Delta$. So we need to add Δ more units of flow to h. Let $E'' := \{e \in E' \mid g_e \geq f'_e\}$ be the set of edges in E' on which g sends more flow than f'. Thus, we have that $g(E'') - f'(E'') \geq g(E') - f'(E') = \Delta$. In particular, we can define a flow h of total rate

Δ on E'', such that $h_e \leq g_e - f'_e$ for all $e \in E''$. For all other edges e, we set $h_e = 0$, and thus obtain a flow h of rate r_1, such that $h_e \leq g_e$ for all edges e. We then have that

$$\begin{aligned}
\sum_e h_e c_e(g_e) &= \sum_{e \in E_1} h_e c_e(h_e) + \sum_{e \in E''} h_e c_e(g_e) \\
&\geq \sum_{e \in E_1} h_e c_e(h_e) + \sum_{e \in E''} h_e c_e(f'_e + h_e),
\end{aligned}$$

where the inequality follows from the monotonicity of the latencies c_e. Next, because $g'_e \geq f'_e$ for all $e \in E'$, and the latency functions are convex, $\frac{c_e(g_e) - c_e(g'_e)}{h_e} \geq c'_e(f'_e)$ for all $e \in E''$ with $h_e > 0$. Combining this bound with the fact that $\beta_1 \leq 1$, we obtain that

$$\begin{aligned}
\sum_e g'_e(c_e(g_e) - c_e(g'_e)) &\geq \sum_{e \in E''} g'_e(c_e(g_e) - c_e(g'_e)) \\
&\geq \sum_{e \in E''} f'_e \beta_1 h_e c'_e(f'_e).
\end{aligned}$$

Summing the previous two inequalities now gives us

$$\begin{aligned}
& \sum_e h_e c_e(g_e) + \sum_e g'_e(c_e(g_e) - c_e(g'_e)) \\
\geq & \quad \sum_{e \in E_1} h_e c_e(h_e) + \sum_{e \in E''} h_e c_e(f'_e + h_e) + \sum_{e \in E''} h_e \beta_1 f'_e c'_e(f'_e) \\
= & \sum_e h_e \tilde{c}_e(h_e) \\
\geq & \ \tilde{C}(\tilde{f})
\end{aligned}$$

where the final inequality follows from the optimality of $\tilde{f}$ with respect to the cost functions $\tilde{c}_e$. ■

Proof of Theorem 4.3.1. If $\mathcal{C}$ is specifically the set of all increasing semi-convex functions, Proposition 4.2.1 implies that $\frac{1}{\alpha^{(t)}(\mathcal{C})} \geq t$. Substituting this bound into the integral gives us that

$$\rho(G, \mathbf{r}, \mathbf{c}, \psi) \quad \leq \quad \left(\int_0^1 \psi(t) t dt\right)^{-1} \quad = \quad 1/\bar{\beta}. \qquad \blacksquare$$

It would of course be desirable to extend Theorems 4.3.1 and 4.3.2 to distributions including negative support. However, such an extension is in general not possible. One can construct scenarios in which almost all of the latency is incurred by a small fraction of spiteful users who together congest a link with very steep increase. At the same time, all altruistic users use links with very small constant latency. Then, the PoA is much larger than 1, while the bounds of both theorems would require it to be close to 1.

An immediate corollary of Theorem 4.3.1 can be obtained by choosing the distribution with a rate of λ users being completely altruistic, and $1-\lambda$ users being completely selfish. Since $\bar{\beta} = \lambda$ for this distribution, Theorem 4.3.1 immediately implies

Corollary 4.3.4 *In parallel link networks, the Price of Anarchy under Stackelberg routing with a λ-fraction of traffic being controlled by a central authority is at most $1/\lambda$.*

This result was of course already proved constructively (and giving efficient algorithms) by Roughgarden [80]; nevertheless, it is interesting that it follows directly from our general result. More generally, by using the same distribution with support $\{0, 1\}$ in Theorem 4.3.2, we obtain the following corollary:

Corollary 4.3.5 *In parallel link networks, the Price of Anarchy under Stackelberg routing with a λ-fraction of traffic controlled by a central authority is at most $(\frac{1-\lambda}{\alpha(\mathcal{C})} + \lambda)^{-1}$.*

Notice that Corollary 4.3.5 improves (albeit in a non-constructive way) a result of Swamy [90] for Stackelberg routing: we bound the PoA under Stackelberg routing by the weighted harmonic mean of the PoA for selfish and altruistic users, whereas Swamy's bounds give the arithmetic mean. It is known that the harmonic mean is always bounded above by the arithmetic mean. We can also show that the case of Stackelberg routing is in fact the worst case for the bound of Theorem 4.3.2, in the sense that the right-hand side is maximized. While the bound of Theorem 4.3.2 will in general not be tight, this nevertheless gives rise to the philosophical interpretation that, conditioned on a given average altruism level $\bar{\beta}$, the scenario in which completely altruistic users or a central authority compensate for completely selfish users is the worst case, while uniform altruism through the population is the best case.

Proposition 4.3.6 *Conditioned on the mean of ψ being any given $\bar{\beta}$, the quantity $\left(\int_0^1 \psi(t)\frac{1}{\alpha^{(t)}(\mathcal{C})}dt\right)^{-1}$ is maximized when ψ has point mass of $\bar{\beta}$ on 1 and $1-\bar{\beta}$ on 0. It is minimized when ψ has a point mass of 1 on $\bar{\beta}$.*

Proof. We will show that $\frac{1}{\alpha^{(\beta)}(\mathcal{C})}$ is concave as a function of β. Both results then follow readily from Jensen's Inequality. To prove concavity, let $p_1, p_2 \geq 0$ satisfy $p_1 + p_2 = 1$. For any cost function $c \in \mathcal{C}$, Definition 4.2.2 thus gives us

$$
\begin{aligned}
& \frac{1}{\alpha^{(p_1\beta_1+p_2\beta_2)}(c)} \\
= \; & \inf_\lambda \frac{\lambda c(\lambda r) + (1-\lambda)c(r) + (1-\lambda)(p_1\beta_1 + p_2\beta_2)c'(r)}{c(r)} \\
= \; & \inf_\lambda \Big(\frac{p_1(\lambda c(\lambda r) + (1-\lambda)c(r) + (1-\lambda)\beta_1 c'(r))}{c(r)} \\
& + \frac{p_2(\lambda c(\lambda r) + (1-\lambda)c(r) + (1-\lambda)\beta_2 c'(r))}{c(r)}\Big) \\
\geq \; & \inf_\lambda \frac{p_1(\lambda c(\lambda r) + (1-\lambda)c(r) + (1-\lambda)\beta_1 c'(r))}{c(r)} \\
& + \inf_\lambda \frac{p_2(\lambda c(\lambda r) + (1-\lambda)c(r) + (1-\lambda)\beta_2 c'(r))}{c(r)} \\
= \; & p_1 \frac{1}{\alpha^{(\beta_1)}(c)} + p_2 \frac{1}{\alpha^{(\beta_2)}(c)}.
\end{aligned}
$$

Finally, we take an infimum over all $c \in \mathcal{C}$ on both sides to complete the proof of concavity. ■

4.4 Conclusions and Future Work

We proved a $1/\beta$ bound on the Price of Anarchy even for worst-case networks, latency functions, and commodities, under the assumption that all users are (at least) β-altruistic, and $\beta > 0$. We extended this result to non-uniform altruism distributions for single-commodity flows in parallel link networks. Among others, this result recovers and improves recent bounds on Stackelberg routing by Roughgarden and by Swamy.

Our work suggests many interesting directions for further research. First, the results should be generalized to (more general) network topologies instead of parallel links. Notice that any such result would immediately imply corresponding bounds on Stackelberg routing, so the lower bound of Bonifaci et al. [12] precludes an extension to arbitrary single-commodity flows. However, an extension to *series-parallel graphs* seems plausible at this point.

While we proved the existence of Nash Equilibria for all routing games with non-atomic users, regardless of the distributions of altruism, the proof is non-constructive. The work of Roughgarden [80] implies that finding the *best* Stackelberg strategy is NP-complete. However, it would be interesting whether Stackelberg strategies meeting our bound can always be found efficiently. Alternatively, in light of recent results proving that finding Nash Equilibria is PPAD-complete [22], it may be possible that finding Nash Equilibria for traffic routing games with two (or more) altruism values is also PPAD-complete.

Another interesting question is whether users can learn equilibrium routing strategies using a natural learning algorithm in repeated games. With or without the information about non-uniform altruism levels, we are aiming at finding simple strategies (in the style of [10]) wherein each bidder adapts her routing strategy based on the cost derived from earlier routing results.

Chapter 5

Congestion Games

In this chapter, we are considering altruism in congestion games where players are atomic. Congestion games include, but are more general than, network congestion games, the non-atomic version of which is (traffic) routing games. It takes a different set of techniques to deal with atomic players. Note that the results in Section 5.2 differ from those in Section 4.2 in terms of the trend of impact of altruism on the PoA (decreasing or increasing the PoA) while the results in Section 5.3 for non-uniform altruism bear some similarity to those in Section 4.3, but lose a constant factor exactly due to atomicity. The results of this chapter are based on the paper [20].

5.1 Preliminaries

We have already given the definition of a general congestion game in the beginning of Section 4.1. In this chapter, we are back to general congestion games, not just network congestion games, which were discussed in the previous chapter. Congestion games are a class of games first proposed by Rosenthal [76]. Rosenthal showed that any congestion game is a *potential game*. A potential game is a game, where existence of pure Nash equilibria can be guaranteed by a *potential function*, which is defined to be a function whose difference in value when any player deviates is equal to the difference in value of his individual utility/cost function. Monderer and Shapley [68] proved that, for any potential game, there is a congestion game with the same potential function. Thus, congestion games are equivalent to potential games. Unlike in the previous chapter where players are *non-atomic*, we deal with *atomic* players, each of which is of size 1, in congestion games here. Congestion games generalize (traffic) routing games or network congestion games in the previous chapter, where strategy spaces are just paths in a graph. We mainly focus on congestion games with linear cost functions. In a *linear* congestion game, the cost function of every resource $e \in E$ is of the form $c_e(x) = a_e x + b_e$, where $a_e, b_e \in \mathbb{R}^+$ are non-negative real numbers.

For congestion games with atomic entirely selfish players, Christodoulou and Koutsoupias proved that the PoA of pure Nash equilibria is 5/2 for linear cost functions while it increases with the highest degree for polynomial functions [23]. With partial altruism introduced, our bounds reveal an unexpected trend: for congestion games, the worst-case robust PoA increases with increasing altruism in contrast to the results in the previous chapter for non-atomic (network) congestion games where the PoA for pure Nash equilibria decreases with increasing altruism. We also show that the increase in the PoA is not a universal phenomenon: For symmetric singleton linear congestion games, we derive a bound on the PoA for pure Nash equilibria that decreases as the level of altruism increases, which is similar to the results in the previous chapter for non-atomic (network) congestion games. There are other results showing the worst-case robust PoA increases with increasing altruism in cost-sharing games as well, and for valid utility games, it remains constant and is not affected by altruism [20].

5.1.1 Altruism

We study β-altruistic congestion games defined as follows, which falls into our general model of altruism in Section 2.3:

Definition 5.1.1 *The β-altruistic congestion game G is defined as the congestion game $G^\beta = (N, \{S_i\}_{i\in N}, \{c_i^{(\beta_i)}\}_{i\in N})$, where for every $i \in N$ and $\mathbf{s} \in S$,*

$$c_i^{(\beta_i)}(\mathbf{s}) \quad = \quad (1-\beta_i)c_i(\mathbf{s}) + \beta_i C(\mathbf{s}).$$

A β_i-altruistic player has his perceived cost as above, in comparison to Definition 4.1.2 using derivatives in the altruistic term.

5.1.2 Smoothness

Many proofs bounding the Price of Anarchy for specific games (e.g., [81, 92]) use the fact that deviating from an equilibrium to the strategy at optimum is not beneficial for any player. The addition of these inequalities, combined with suitable properties of the social cost function, then gives a bound on the equilibrium's cost. Roughgarden [84] recently captured the essence of this type of argument with his definition of (λ, μ)-*smoothness* of a game, thus providing a generic template for proving bounds on the Price of Anarchy. Indeed, because such arguments only reason about local moves by players, they immediately imply bounds not only for Nash equilibria but for all the classes of general equilibria defined in Chapter 2 all the way to the outcomes of no-regret sequences of play [11, 10] (coarse correlated equilibria). Recent work has explored both the limits of this concept [71] and a refinement requiring smoothness only in local neighborhoods [85]. The latter permits more fine-grained analysis of games, but applies only to correlated equilibria and their subclasses.

We extend the concept of (λ, μ)-smoothness to altruistic strategic games. This allows us to quantify the Price of Anarchy of these games with respect to the very broad class of coarse correlated equilibria. For notational convenience, we define $C_{-i}(\mathbf{s}) = C(\mathbf{s}) - c_i(\mathbf{s}) = \sum_{j \neq i} c_j(\mathbf{s})$ for strategy profile $\mathbf{s}$.

Definition 5.1.2 **((λ, μ, β)-smoothness)** *Let G^β be a β-altruistic congestion game with social cost function C. G^β is (λ, μ, β)-smooth if for any two strategy profiles $\mathbf{s}, \mathbf{s}^* \in S$,*

$$\sum_{i=1}^{n} c_i(s_i^*, \mathbf{s}_{-i}) + \beta_i(C_{-i}(s_i^*, \mathbf{s}_{-i}) - C_{-i}(\mathbf{s})) \quad \leq \quad \lambda C(\mathbf{s}^*) + \mu C(\mathbf{s}), \tag{5.1}$$

for player i changing strategy from s_i to s_i^* while the others stay at $\mathbf{s}_{-i}$.

For $\beta = (0, ..., 0) = \mathbf{0}$, this definition coincides with Roughgarden's notion of (λ, μ)-smoothness. To gain some intuition, consider two strategy profiles $\mathbf{s}, \mathbf{s}^* \in S$, and a player $i \in N$ who switches from his strategy s_i under $\mathbf{s}$ to s_i^*, while the strategies of the other players remain fixed at $\mathbf{s}_{-i}$. The contribution of player i to the left-hand side of 5.1 then accounts for the individual cost that player i perceives after the switch plus β_i times the difference in social cost caused by this switch. The sum of these contributions needs to be bounded by $\lambda C(\mathbf{s}^*) + \mu C(\mathbf{s})$.

Therefore, if G^β is (λ, μ, β)-smooth with $\mu < 1$, then the coarse price of anarchy of G^β is at most $\frac{\lambda}{1-\mu}$: for player i changing strategy from s_i to s_i^* while the others stay at $\mathbf{s}_{-i}$, $(1-\beta_i)c_i(\mathbf{s}) + \beta_i C(\mathbf{s}) \leq (1-\beta_i)c_i(s_i^*, s_{-i}) + \beta_i C(s_i^*, s_{-i})$, which implies that $\sum_{i=1}^n (1-\beta_i)c_i(\mathbf{s}) + \beta_i C(\mathbf{s}) \leq \sum_{i=1}^n (1-\beta_i)c_i(s_i^*, s_{-i}) + \beta_i C(s_i^*, s_{-i})$, equivalent to $\sum_{i=1}^n c_i(\mathbf{s}) + \beta_i C_{-i}(\mathbf{s}) \leq \sum_{i=1}^n c_i(s_i^*, s_{-i}) + \beta_i C_{-i}(s_i^*, s_{-i})$; after rearrangement, we derive that $C(\mathbf{s}) = \sum_{i=1}^n c_i(\mathbf{s}) \leq \sum_{i=1}^n c_i(s_i^*, s_{-i}) + \beta_i(C_{-i}(s_i^*, s_{-i}) - C_{-i}(\mathbf{s}))$ so with (λ, μ, β)-smoothness, the bound $\frac{\lambda}{1-\mu}$ is obtained. Note that this also holds for more restricted equilibria including correlated equilibria, mixed Nash equilibria, and pure Nash equilibria since the argument works for arbitrary pairs of strategy profiles $\mathbf{s}, \mathbf{s}^*$.

Definition 5.1.3 *The* robust Price of Anarchy *is defined as the best possible bound on the coarse price of anarchy obtainable by a* (λ, μ, β)*-smoothness argument.*

Later, with uniform altruism we will see that the (λ, μ, β)-smoothness arguments can give us some tight bounds even for pure Nash equilibria.

5.2 Linear Congestion Games

Linear congestion games have the advantage that pure Nash equilibria of their altruistic games always exist [46], which may not be the case for arbitrary congestion games. We consider uniform altruism where $\beta_i = \beta$.

Proposition 5.2.1 *Let* G^β *be a uniformly* β*-altruistic linear congestion game. Then,* G^β *is an exact potential game with potential function* $\Phi^{(\beta)}(\mathbf{s}) = (1-\beta)\Phi(\mathbf{s}) + \beta C(\mathbf{s})$*, where* $\Phi(\mathbf{s}) = \sum_{e \in E} \sum_{x=1}^{f_e(\mathbf{s})} c_e(x)$ *is Rosenthal's potential function.*

Proof. Assume that player i changes strategy from s_i to s'_i while the others stay at $\mathbf{s}_{-i}$. Then, his change in perceived cost is

$$\begin{aligned}
& c_i^{(\beta)}(\mathbf{s}) - c_i^{(\beta)}(s'_i, \mathbf{s}_{-i}) \\
= \; & (1-\beta)\big(\sum_{e \in S_i \setminus S'_i} (c_e(f_e(\mathbf{s})) - c_e(f_e(\mathbf{s}) - 1)) + \sum_{e \in S'_i \setminus S_i} (c_e(f_e(\mathbf{s})) - c_e(f_e(\mathbf{s}) + 1))\big) \\
& + \beta\big(\sum_{e \in S_i \setminus S'_i} (f_e c_e(f_e(\mathbf{s})) - (f_e - 1) c_e(f_e(\mathbf{s}) - 1)) \\
& + \sum_{e \in S'_i \setminus S_i} (f_e c_e(f_e(\mathbf{s})) - (f_e + 1) c_e(f_e(\mathbf{s}) + 1))\big) \\
= \; & \Phi^{(\beta)}(\mathbf{s}) - \Phi^{(\beta)}(s'_i, \mathbf{s}_{-i}).
\end{aligned}$$

■

We first prove a lemma for every i for arbitrary *semi-convex* cost functions c_e.

Lemma 5.2.2 *For player i changing strategy from s_i to s_i^* while the others stay at s_{-i},*

$$\begin{aligned}
C(s_i^*, s_{-i}) - C(\mathbf{s}) \;\leq\; & (\sum_{e \in s_i^*} (f_e + 1) c_e(f_e + 1) - \sum_{e \in s_i^*} f_e c_e(f_e)) \\
& - (\sum_{e \in s_i} f_e c_e(f_e) - \sum_{e \in s_i} (f_e - 1) c_e(f_e - 1)).
\end{aligned}$$

Proof. Simply using the definition of the social cost, we have that

$$
\begin{aligned}
& C(s_i^*, \mathbf{s}_{-i}) - C(\mathbf{s}) = \sum_{j=1}^{n} c_j(s_i^*, \mathbf{s}_{-i}) - \sum_{j=1}^{n} c_j(\mathbf{s}) \\
= & \sum_{e \in s_i^* \setminus s_i} (f_e+1)c_e(f_e+1) + \sum_{e \in s_i \setminus s_i^*} (f_e-1)c_e(f_e-1) + \sum_{e \in s_i^* \cap s_i} f_e c_e(f_e) + \sum_{e \notin s_i^* \cup s_i} f_e c_e(f_e) \\
& - \sum_{e \in s_i^* \setminus s_i} f_e c_e(f_e) - \sum_{e \in s_i \setminus s_i^*} f_e c_e(f_e) - \sum_{e \in s_i^* \cap s_i} f_e c_e(f_e) - \sum_{e \notin s_i^* \cup s_i} f_e c_e(f_e) \\
= & \sum_{e \in s_i^* \setminus s_i} (f_e+1)c_e(f_e+1) - \sum_{e \in s_i^* \setminus s_i} f_e c_e(f_e) - \sum_{e \in s_i^* \cap s_i} f_e c_e(f_e) + \sum_{e \in s_i^* \cap s_i} f_e c_e(f_e) \\
& -(\sum_{e \in s_i} f_e c_e(f_e) - \sum_{e \in s_i \setminus s_i^*} (f_e-1)c_e(f_e-1) - \sum_{e \in s_i^* \cap s_i} f_e c_e(f_e)) \\
= & \sum_{e \in s_i^* \setminus s_i} (f_e+1)c_e(f_e+1) - \sum_{e \in s_i^*} f_e c_e(f_e) \\
& + \sum_{e \in s_i^* \cap s_i} f_e c_e(f_e) + \sum_{e \in s_i^* \cap s_i} (f_e+1)c_e(f_e+1) - \sum_{e \in s_i^* \cap s_i} f_e c_e(f_e) \\
& -(\sum_{e \in s_i} f_e c_e(f_e) - \sum_{e \in s_i \setminus s_i^*} (f_e-1)c_e(f_e-1) \\
& - \sum_{e \in s_i^* \cap s_i} f_e c_e(f_e) + \sum_{e \in s_i^* \cap s_i} (f_e+1)c_e(f_e+1) - \sum_{e \in s_i^* \cap s_i} f_e c_e(f_e))
\end{aligned}
$$

$$
\begin{aligned}
\leq & \sum_{e \in s_i^*} (f_e+1)c_e(f_e+1) - \sum_{e \in s_i^*} f_e c_e(f_e) \\
& -(\sum_{e \in s_i} f_e c_e(f_e) - \sum_{e \in s_i \setminus s_i^*} (f_e-1)c_e(f_e-1) - \sum_{e \in s_i^* \cap s_i} (f_e-1)c_e(f_e-1)) \\
= & \sum_{e \in s_i^*} (f_e+1)c_e(f_e+1) - \sum_{e \in s_i^*} f_e c_e(f_e) - (\sum_{e \in s_i} f_e c_e(f_e) - \sum_{e \in s_i} (f_e-1)c_e(f_e-1)),
\end{aligned}
$$

where the last inequality follows from the semi-convexity of c_e, i.e.,

$$
(f_e+1)c_e(f_e+1) - f_e c_e(f_e) \geq f_e c_e(f_e) - (f_e-1)c_e(f_e-1)
$$

for every e, so

$$\sum_{e\in s_i^*\cap s_i}(f_e+1)c_e(f_e+1)-\sum_{e\in s_i^*\cap s_i}f_ec_e(f_e)\geq\sum_{e\in s_i^*\cap s_i}f_ec_e(f_e)-\sum_{e\in s_i^*\cap s_i}(f_e-1)c_e(f_e-1).$$

■

We now can use the above lemma to bound the $C(s_i^*, s_{-i}) - C(\mathbf{s})$ term and then obtain the Price of Anarchy bound for linear cost functions.

Theorem 5.2.3 *For linear congestion games with $0 \leq \beta \leq 1$, the robust Price of Anarchy of uniformly β-altruistic linear congestion games is at most $\frac{5+4\beta}{2+\beta}$.*

Proof. Let $c_e(f_e) = a_e f_e + b_e$ for $a_e, b_e \geq 0$ and $f_e \in \mathbb{N}_0$ be the cost function for resource e. By Lemma 5.2.2 and linearity of c_e, for player i changing from s_i to s_i^* while the others stay at s_{-i},

$$\begin{aligned} C(s_i^*, s_{-i}) - C(\mathbf{s}) &= \sum_{e\in s_i^*}(a_e(f_e+1)^2+b_e(f_e+1))-\sum_{e\in s_i^*}(a_ef_e^2+b_ef_e) \\ &\quad -(\sum_{e\in s_i}(a_ef_e^2+b_ef_e)-\sum_{e\in s_i}(a_e(f_e-1)^2+b_e(f_e-1))) \\ &= \sum_{e\in s_i^*}(2a_ef_e+a_e+b_e)-\sum_{e\in s_i}(2a_ef_e-a_e+b_e). \end{aligned}$$

Summing over all players, we get that for $0 \leq \beta < 1$ (the case $\beta = 1$ will be addressed later),

$$\begin{aligned} &\sum_{i=1}^n c_i(\mathbf{s}) \leq \sum_{i=1}^n c_i(s_i^*, \mathbf{s}_{-i}) + \frac{\beta}{1-\beta}\sum_{i=1}^n(C(s_i^*, \mathbf{s}_{-i}) - C(\mathbf{s})) \\ \leq\; &\sum_e(a_ef_e^*(f_e+1)+b_ef_e^*)+\frac{\beta}{1-\beta}(\sum_e f_e^*(2a_ef_e+a_e+b_e)-\sum_e f_e(2a_ef_e-a_e+b_e)). \end{aligned}$$

We can bound $a_e f_e^*(f_e+1) + b_e f_e^* + \frac{\beta}{1-\beta}(f_e^*(2a_e f_e + a_e + b_e) - \sum_e f_e(2a_e f_e - a_e + b_e))$ for every e. It is enough to find c_1, c_2 such that

$$\begin{aligned} & a_e f_e^*(f_e+1) + b_e f_e^* + \frac{\beta}{1-\beta}(f_e^*(2a_e f_e + a_e + b_e) - \sum_e f_e(2a_e f_e - a_e + b_e)) \\ \leq \quad & c_1(a_e f_e^{*2} + b_e f_e^*) + c_2(a_e f_e^2 + b_e f_e), \end{aligned}$$

for every e.

Let $y = f_e^*, x = f_e$ and $a = a_e, b = b_e$ so

$$\begin{aligned} & ay(x+1) + by + \frac{\beta}{1-\beta}(y(2x+a+b) - \sum_e x(2ax - a + b)) \\ = \quad & a(y(x+1) + \frac{\beta}{1-\beta}(y(2x+1) - x(2x-1))) + b(\frac{1}{1-\beta}y - \frac{\beta}{1-\beta}x). \end{aligned}$$

We first find c_1, c_2 such that

$$\begin{aligned} & y(x+1) + \frac{\beta}{1-\beta}(y(2x+1) - x(2x-1)) \\ = \quad & y(x+1) + \frac{\beta}{1-\beta}(y(2x+1) + x)) - \frac{2\beta}{1-\beta}x^2 \\ \leq \quad & (c_1' - \frac{2\beta}{1-\beta})x^2 + c_2 y^2 \\ = \quad & c_1 x^2 + c_2 y^2, \end{aligned}$$

i.e., c_1', c_2 such that

$$\frac{x+1}{y} + \frac{\beta}{1-\beta} \cdot \frac{2x+1}{y} + \frac{\beta}{1-\beta} \cdot \frac{x}{y^2} - c_1'(\frac{x}{y})^2 \leq c_2.$$

Then, we show that such a choice of c_1, c_2 also works for

$$\begin{aligned} & a(y(x+1)+\frac{\beta}{1-\beta}(y(2x+1)-x(2x-1)))+b(\frac{1}{1-\beta}y-\frac{\beta}{1-\beta}x) \\ \leq \quad & c_1(ax^2+bx)+c_2(ay^2+by). \end{aligned}$$

Feasible c_1', c_2 satisfy

$$\begin{aligned} c_2 &= \max_{y\in\mathbb{N},x\in\mathbb{N}_0}(\frac{x+1}{y}+\frac{\beta}{1-\beta}\cdot\frac{2x+1}{y}+\frac{\beta}{1-\beta}\cdot\frac{x}{y^2}-c_1'(\frac{x}{y})^2) \\ &= \max_{x\in\mathbb{N}_0}(x+1+\frac{\beta}{1-\beta}(3x+1)-c_1'x^2). \end{aligned}$$

$(x+1+\frac{\beta}{1-\beta}(3x+1)-c_1'x^2)$ is maximized at $x=1$ for $\frac{1}{3}+\frac{\beta}{1-\beta}\leq c_1' < 1+\frac{3\beta}{1-\beta}$ so $c_2 = 2+\frac{4\beta}{1-\beta}-c_1'$. Also, $\frac{c_2}{1-c_1} = \frac{2+\frac{4\beta}{1-\beta}-c_1'}{1-c_1'+\frac{2\beta}{1-\beta}}$ is minimized at $c_1' = \frac{1}{3}+\frac{\beta}{1-\beta}$. Thus, $c_1 = c_1'-\frac{2\beta}{1-\beta} = \frac{1}{3}-\frac{\beta}{1-\beta}$, $c_2 = \frac{5}{3}+\frac{3\beta}{1-\beta}$, and $\frac{c_2}{1-c_1} = \frac{5+4\beta}{2+\beta}$. Since $-\frac{\beta}{1-\beta}x \leq c_1x$ and $\frac{1}{1-\beta}y \leq c_2y$, $a(y(x+1)+\frac{\beta}{1-\beta}(y(2x+1)-x(2x-1)))+b(\frac{1}{1-\beta}y-\frac{\beta}{1-\beta}x) \leq c_1(ax^2+bx)+c_2(ay^2+by)$ as well.

This bound of $\frac{5+4\beta}{2+\beta}$ approaches 3 when β approaches 1. Thus, $\frac{5+4\beta}{2+\beta}$ also works for the case $\beta = 1$. ■

The following example shows that the upper bound on the robust Price of Anarchy given above is tight for uniformly β-altruistic games, even for pure Nash equilibria.

Lemma 5.2.4 *The Price of Anarchy bound of* $\frac{5+4\beta}{2+\beta}$ *for* $0 \leq \beta \leq 1$ *is tight.*

Proof. Let $n \geq 3$ be the number of player, and E be the set of $2n$ resources. We can divide E into two subsets $E_1 = \{h_1, ..., h_n\}$ and $E_2 = \{g_1, ..., g_n\}$, each of n resources.

$c_e(f_e) = (1+\beta)f_e$ for $e \in E_1$ and $c_e(f_e) = f_e$ for $e \in E_2$. Each player i has two pure strategies: $\{h_i, g_i\}$ and $\{h_{i-1}, h_{i+1}, g_{i+1}\}$.

The social optimum arises when each player selects the first strategy, and has cost $(1+\beta)n + n = (2+\beta)n$. A Nash equilibrium exists when each player selects the second strategy, shown as follows. When each player selects the second strategy, the social cost is $4(1+\beta)n + n = (5+4\beta)n$. If a player deviates to select the first strategy while the other players still select the second strategy, the social cost would become $4(1+\beta)n+n+(3^2-2^2)\cdot(1+\beta)+(2^2-1^2)-(2^2-1^2)\cdot(1+\beta)\cdot 2-(1^2-0^2) = (5+4\beta)n+1-\beta$. Each player selecting the second strategy gives each player a perceived individual cost of $(1-\beta)(4(1+\beta)+1)+\beta(5+4\beta)n = C$. The deviation would give the player a perceived individual cost of $(1-\beta)(3(1+\beta)+2)+\beta((5+4\beta)n+1-\beta) = C'$. We can see that

$$\begin{aligned} C - C' &= (1-\beta)(4(1+\beta)+1)+\beta(5+4\beta)n \\ &\quad -(1-\beta)(3(1+\beta)+2)-\beta((5+4\beta)n+1-\beta) \\ &= (1-\beta)(5+4\beta-5-3\beta)+\beta(-1+\beta) = 0. \end{aligned}$$

So, there is a Nash equilibrium when each player selects the second strategy. We get the Price of Anarchy value of at least $\frac{(5+4\beta)n}{(2+\beta)n} = \frac{5+4\beta}{2+\beta}$ for $0 \le \beta \le 1$. ■

We have shown that the PoA for several general equilibrium concepts can actually increase with β. This is somewhat surprising, and one would have perhaps expected the opposite, the PoA decreases with β. Also notice that this differs from the results in the previous chapter where the PoA improves with altruism. As mentioned in Section 1.2, the increase is due to partially altruistic players having stronger disincentive to move from the suboptimal strategy, making worse system states stable and stay in the set of equilibrium states. Therefore, it is implied that the best stable solution can also be chosen from a larger set so the PoS should decrease.

We now analyze the pure Price of Stability of β-altruistic congestion games. Clearly, an upper bound on the pure Price of Stability extends to the mixed, correlated and coarse Price of Stability. The proof of the following proposition exploits a standard technique to bound the pure PoS of exact potential games (see, e.g., [74]).

Proposition 5.2.5 *The pure Price of Stability of uniformly β-altruistic linear congestion games is at most $\frac{2}{1+\beta}$.*

Proof. Let G^{β} be a uniformly β-altruistic linear congestion game. By Proposition 5.2.1, G^{β} is an exact potential game with potential function $\Phi^{(\beta)}(\mathbf{s}) = (1-\beta)\Phi(\mathbf{s}) + \beta C(\mathbf{s})$, where $\Phi(\mathbf{s}) = \sum_{e \in E} \sum_{i=1}^{f_e(\mathbf{s})} i$ is Rosenthal's potential function. Observe that

$$\begin{aligned} \Phi^{\beta}(\mathbf{s}) &= (1-\beta) \sum_{e \in E} \sum_{i=1}^{f_e(\mathbf{s})} i + \beta C(\mathbf{s}) \\ &= \frac{1-\beta}{2} \sum_{e \in E} (f_e^2(\mathbf{s}) + f_e(\mathbf{s})) + \beta \sum_{e \in E} f_e^2(\mathbf{s}) \\ &= \frac{1+\beta}{2} C(\mathbf{s}) + \frac{1-\beta}{2} \sum_{e \in E} f_e(\mathbf{s}). \end{aligned}$$

We therefore have

$$\frac{1+\beta}{2} C(\mathbf{s}) \leq \Phi^{\beta}(\mathbf{s}) \leq C(\mathbf{s}).$$

Let $\mathbf{s}$ be a strategy profile that minimizes Φ^β, and let $\mathbf{s}^*$ be an optimal strategy profile that minimizes the social cost function C. Note that $\mathbf{s}$ is a pure Nash equilibrium of G^β. We have

$$C(\mathbf{s}) \leq \frac{2}{1+\beta}\Phi^\beta(\mathbf{s}) \leq \frac{2}{1+\beta}\Phi^\beta(\mathbf{s}^*) \leq \frac{2}{1+\beta}C(\mathbf{s}^*),$$

which proves the claim. ■

5.3 Linear Singleton Congestion Games

While the Price of Anarchy for general congestion games, somewhat counter-intuitively, can increase with β, the situation is markedly different for the PoA of *symmetric singleton* congestion games. We have already defined symmetric singleton games in Section 4.1. Recall that in a symmetric singleton congestion game $G = (N, E, \{S_i\}_{i\in N}, \{c_e\}_{e\in E})$, every player chooses one resource (also called *edge*) from $E = \{1, \dots, m\}$, and all strategy sets are identical, i.e., $S_i = E$ for every i. Again, we assume that the cost functions are of the form $c_e(x) = a_e x + b_e$.

In this section, we analyze perhaps the two most fundamental cases with respect to altruistic singleton congestion games: the uniform case, and the case when all altruism levels are in $\{0, 1\}$, i.e., each player is either completely altruistic or completely selfish. For both settings, we establish bounds on the pure PoA which *improve* with the total altruism level in the system, i.e., decrease in β or the fraction of selfish players. This stands in marked contrast to the bounds in the previous section, but is similar to the result of Theorem 4.3.1 in the previous chapter.

Theorem 5.3.1 *The pure Price of Anarchy of uniformly β-altruistic linear symmetric singleton congestion games is $\frac{4}{3+\beta}$.*

While in the previous section, we were able to derive tight bounds via smoothness arguments, this is not possible for altruistic symmetric singleton congestion games. For example, by the above theorem, the Price of Anarchy in the purely selfish setting is 4/3, whereas Lücking et al. [62, Theorem 5.4] showed that the mixed price of anarchy for symmetric singleton congestion games with cost functions $c_e(x) = x$ is $1 + \min\{\frac{m-1}{n}, \frac{n-1}{m}\}$. That is, for $n = m$, the mixed Price of Anarchy approaches 2 as n increases. The bound given in Theorem 5.3.1 can therefore not be derived via a smoothness argument.

Theorem 5.3.1 implies that the pure PoA is 1 if all players are completely altruistic. We remark that this continues to hold true for the more general class of semi-convex cost functions (see Corollary 5.3.4 later).

Proof. [Theorem 5.3.1] Let $\mathbf{s}$ be a pure Nash equilibrium of G^β and $\mathbf{s}^*$ an optimal strategy profile. We write $f_e = f_e(\mathbf{s})$ and $f_e^* = f_e(\mathbf{s}^*)$. For every edge $e \in E$, define $\Delta_e = f_e - f_e^*$. Let E^+ and E^- be the set of edges with $\Delta_e > 0$ and $\Delta_e < 0$, respectively. Define $\Delta = \sum_{e \in E^+} \Delta_e > 0$. Because $\mathbf{s}$ and $\mathbf{s}^*$ assign the same number of players to edges, $\Delta = \sum_{e \in E^+} \Delta_e = -\sum_{e \in E^-} \Delta_e$. If $\Delta = 0$, then the PoA is 1. Hence, we assume that $\Delta > 0$, in which case both E^+ and E^- are non-empty.

By definition, $f_e > f_e^* \geq 0$ for every edge $e \in E^+$. Because s is a Nash equilibrium of G^β, we have for every edge $e \in E^+$ and $\bar{e} \in E$

$$\begin{aligned} & (1-\beta)(a_e f_e + b_e) + \beta((a_e x_e^2 + b_e f_e) + (a_{\bar{e}} x_{\bar{e}}^2 + b_{\bar{e}} f_{\bar{e}})) \\ \leq \quad & (1-\beta)(a_{\bar{e}}(f_{\bar{e}}+1) + b_{\bar{e}}) + \beta\left((a_e(f_e-1)^2 + b_e(f_e-1)) + (a_{\bar{e}}(f_{\bar{e}}+1)^2 + b_{\bar{e}}(f_{\bar{e}}+1))\right), \end{aligned}$$

which is equivalent to

$$(1+\beta)a_e f_e + b_e - \beta a_e \leq (1+\beta)a_{\bar{e}} f_{\bar{e}} + b_{\bar{e}} + a_{\bar{e}}. \tag{5.2}$$

We can use this relation in order to show that

$$\begin{aligned} & \sum_{e \in E^+} \Delta_e((1+\beta)a_e f_e^* + b_e + a_e\Delta_e) + \sum_{e \in E^-} \Delta_e((1+\beta)a_e f_e^* + b_e + \beta a_e \Delta_e) \\ = \quad & \sum_{e \in E^+} \Delta_e((1+\beta)a_e f_e + b_e - \beta a_e \Delta_e) + \sum_{e \in E^-} \Delta_e((1+\beta)a_e f_e + b_e - a_e\Delta_e) \\ \leq \quad & \sum_{e \in E^+} \Delta_e((1+\beta)a_e f_e + b_e - \beta a_e) + \sum_{e \in E^-} \Delta_e((1+\beta)a_e f_e + b_e + a_e) \\ \leq \quad & \Delta(\max_{e \in E^+}((1+\beta)a_e f_e + b_e - \beta a_e) - \min_{e \in E^-}((1+\beta)a_e f_e + b_e + a_e)) \leq 0. \end{aligned} \tag{5.3}$$

The first inequality follows from the definition of Δ_e and because $\Delta_e \geq 1$ for every $e \in E^+$ and $\Delta_e \leq -1$ for every $e \in E^-$; the last inequality follows from (5.2). Thus,

$$
\begin{aligned}
C(\mathbf{s}) &= \sum_{e\in E}(f_e^* + \Delta_e)(a_e(f_e^* + \Delta_e) + b_e) \\
&= \sum_{e\in E}(a_e x_e^{*2} + b_e f_e^*) + \sum_{e\in E^+} \Delta_e(2a_e f_e^* + b_e + a_e\Delta_e) + \sum_{e\in E^-} \Delta_e(2a_e f_e^* + b_e + a_e\Delta_e) \\
&\leq C(\mathbf{s}^*) + (1-\beta)\sum_{e\in E^+} \Delta_e a_e f_e^* + (1-\beta)\sum_{e\in E^-} \Delta_e a_e(f_e^* + \Delta_e) \\
&\leq C(\mathbf{s}^*) + \tfrac{1}{4}(1-\beta)\sum_{e\in E^+} a_e(f_e^* + \Delta_e)^2 \\
&\leq C(\mathbf{s}^*) + \tfrac{1}{4}(1-\beta)C(\mathbf{s}).
\end{aligned}
$$

The first inequality holds because of (5.3). The second inequality uses that $xy \leq \frac{1}{4}(x + y)^2$ for arbitrary real numbers x, y and that $\Delta_e a_e f_e \leq 0$ for every $e \in E^-$. Hence, $(1 - \frac{1-\beta}{4})C(\mathbf{s}) \leq C(\mathbf{s}^*)$ so the pure Price of Anarchy is at most $4/(3 + \beta)$.

To see that this bound is tight, consider a β-altruistic congestion game with two players and two edges $E = \{1, 2\}$ with cost functions $c_1(x) = x$ and $c_2(x) = 2 + \beta$. If the players use different edges, we obtain an optimal strategy profile of cost $3 + \beta$. If both players use edge 1, we obtain a Nash equilibrium of cost 4. ■

Next, we focus on a second very natural special case: when all altruism levels are either 0 or 1. This kind of scenario, in which each player is either completely selfish or completely altruistic, has some natural relationship with *Stackelberg routing games* [80], and constitutes another class of examples where system performance *improves* with the total amount of altruism present. The results of Theorem 5.3.2 remind us of Theorem 4.3.1. However, the bound loses a factor compared with the bound in Theorem 4.3.1. We will see where the loss comes from in several places in the proof due to atomicity.

Theorem 5.3.2 *The pure Price of Anarchy of β-altruistic linear symmetric singleton congestion games and $\beta \in \{0,1\}^n$ is at most $1 + \frac{n_0}{2n+n_0}$, where n_0 is the number of selfish players.*

Let s be a pure Nash equilibrium of G^β and $\mathbf{s}^*$ an optimal strategy profile. Again, let $f_e = f_e(\mathbf{s})$ and $f_e^* = f_e(s^*)$. Based on the strategy profile $\mathbf{s}$, we partition the edges in E into sets E_0, E_1: $E_1 = \{e \in E : \exists i \in N \text{ with } \beta_i = 1 \text{ and } s_i = \{e\}\}$ is the set of edges having at least one altruistic player, while $E_0 = E \setminus E_1$ is the set of edges that are used exclusively by selfish players or not used at all. Let N_1 and N_0 refer to the respective player sets that are assigned to E_1 and E_0. N_1 may contain both altruistic and selfish players, while N_0 consists of selfish players only. Let $k_1 = \sum_{e \in E_1} f_e$ and $k_0 = n - k_1$ denote the number of players in N_1 and N_0, respectively.

The high-level approach of our proof is as follows: we split the total cost $C(\mathbf{s})$ of the pure Nash equilibrium into $C(\mathbf{s}) = \gamma C(\mathbf{s}) + (1-\gamma)C(\mathbf{s})$ for some $\gamma \in [0,1]$ such that $\gamma C(\mathbf{s}) = \sum_{e \in E_0} f_e c_e(f_e)$ and $(1-\gamma)C(\mathbf{s}) = \sum_{e \in E_1} f_e c_e(f_e)$. We bound these two contributions separately to show that

$$\tfrac{3}{4}\gamma C(\mathbf{s}) + (1-\gamma)C(\mathbf{s}) \leq C(\mathbf{s}^*). \tag{5.4}$$

The pure Price of Anarchy is therefore at most $(\frac{3}{4}\gamma + (1-\gamma))^{-1} = \frac{4}{4-\gamma}$. The bound of $1 + \frac{n_0}{2n+n_0}$ then follows by deriving an upper bound on γ in Lemma 5.3.6.

Lemma 5.3.3 *Assume that the cost functions* $(c_e)_{e\in E}$ *are semi-convex. Then there is an optimal strategy profile* $\mathbf{s}^*$ *such that* $f_e \le f_e^*$ *for every edge* $e \in E_1$.

Proof. Let $\mathbf{s}^*$ be an optimal strategy profile, and assume that $f_e^* < f_e$ for some $e \in E_1$. Then there is some edge $\bar{e} \in E$ with $f_{\bar{e}}^* > f_{\bar{e}}$. Consider an altruistic player $i \in N_1$ with $s_i = \{e\}$. (Note that i must exist by the definition of E_1.) Because $\mathbf{s}$ is a pure Nash equilibrium, player i has no incentive to deviate from e to $\bar{e}$, i.e., $C(\{\bar{e}\}, \mathbf{s}_{-i}) \ge C(\mathbf{s})$, or, equivalently,

$$(f_{\bar{e}}+1)c_{\bar{e}}(f_{\bar{e}}+1) - f_{\bar{e}}c_{\bar{e}}(f_{\bar{e}}) \ge f_e c_e(f_e) - (f_e - 1)c_e(f_e - 1). \tag{5.5}$$

Since $f_e^* < f_e$ and $f_{\bar{e}} < f_{\bar{e}}^*$, the semi-convexity of the cost functions implies

$$(f_e^*+1)c_e(f_e^*+1) - f_e^* c_e(f_e^*) \quad \le f_e c_e(f_e) - (f_e - 1)c_e(f_e - 1) \tag{5.6}$$

$$(f_{\bar{e}}+1)c_{\bar{e}}(f_{\bar{e}}+1) - f_{\bar{e}}c_{\bar{e}}(f_{\bar{e}}) \quad \le f_{\bar{e}}^* c_{\bar{e}}(f_{\bar{e}}^*) - (f_{\bar{e}}^* - 1)c_{\bar{e}}(f_{\bar{e}}^* - 1). \tag{5.7}$$

By combining 5.5, 5.6 and 5.7 and re-arranging terms, we obtain

$$(f_e^*+1)c_e(f_e^*+1) + (f_{\bar{e}}^* - 1)c_{\bar{e}}(f_{\bar{e}}^* - 1) \le f_e^* c_e(f_e^*) + f_{\bar{e}}^* c_{\bar{e}}(f_{\bar{e}}^*).$$

The above inequality implies that by moving a player j with $\mathbf{s}_j^* = \{\bar{e}\}$ from $\bar{e}$ to e, we obtain a new strategy profile $\mathbf{s}' = (\{e\}, \mathbf{s}_{-j}^*)$ of cost $C(\mathbf{s}') \le C(\mathbf{s}^*)$. (Note that j must exist because $f_{\bar{e}}^* > f_{\bar{e}} \ge 0$.) Moreover, the number of players on e under the new strategy profile $\mathbf{s}'$ increased by one. We can therefore repeat the above argument (with $\mathbf{s}'$ in place of $\mathbf{s}^*$) until we obtain an optimal strategy profile that satisfies the claim. ■

Note that Lemma 5.3.3 implies that at least for singleton congestion games, entirely altruistic players will ensure that Nash equilibria are optimal.

Corollary 5.3.4 *The pure Price of Anarchy of* 1*-altruistic symmetric singleton congestion games with semi-convex cost functions is* 1.

Henceforth, we assume that $\mathbf{s}^*$ is an optimal strategy profile that satisfies the statement of Lemma 5.3.3.

Lemma 5.3.5 *Define* y^* *as* $y^*_e = f^*_e - f_e \geq 0$ *for every* $e \in E_1$, *and* $y^*_e = f^*_e$ *for all edges* $e \in E_0$. *Then,* $\sum_{e \in E_0} f_e c_e(f_e) \leq \frac{4}{3} \sum_{e \in E} y^*_e c_e(f^*_e)$.

Proof. Consider the game $\bar{G}$ induced by G^β if all k_1 players in N_1 are fixed on the edges in E_1 according to $\mathbf{s}$. Note that all remaining $k_0 = n - k_1$ players in N_0 are selfish. That is, $\bar{G}$ is a symmetric singleton congestion game with player set N_0, edge set E and cost functions $(\bar{c}_e)_{e \in E}$, where $\bar{c}_e(z) = c_e(f_e + z)$ if $e \in E_1$ and $\bar{c}_e(z) = c_e(z)$ for $e \in E_0$. Let $\bar{s}$ be the restriction of s to the players in N_0, and define $\bar{f}$ as $\bar{f}_e = 0$ for $e \in E_1$ and $\bar{f}_e = f_e$ for $e \in E_0$. It is not hard to verify that $\bar{s}$ is a pure Nash equilibrium of the game $\bar{G}$. Let $\bar{\mathbf{s}}^*$ be a socially optimum profile for $\bar{G}$, and for each edge e, let $\bar{f}^*_e$ be the total number of players on e under $\bar{\mathbf{s}}^*$. Then,

$$\begin{aligned} \sum_{e \in E_0} f_e c_e(f_e) &= \sum_{e \in E} \bar{f}_e \bar{c}_e(\bar{f}_e) \\ &\leq \tfrac{4}{3} \sum_{e \in E} \bar{f}^*_e \bar{c}_e(\bar{f}^*_e) \\ &\leq \tfrac{4}{3} \sum_{e \in E} y^*_e \bar{c}_e(y^*_e) = \tfrac{4}{3} \sum_{e \in E} y^*_e c_e(f^*_e), \end{aligned}$$

where the first inequality follows from Theorem 5.3.1 and the second inequality follows from the optimality of $\bar{f}^*$. ■

Notice that the following loses some factor compared with the upper bound of $\frac{n_0}{n}$ on γ in the proof of Theorem 4.3.2.

Lemma 5.3.6 *We have* $\gamma \leq \frac{2n_0}{n+n_0}$.

Proof. The claim follows directly from Theorem 5.3.1 if $N_1 = \emptyset$. Assume that $N_1 \neq \emptyset$, and let $j \in N_1$ with $s_j = \{\bar{e}\}$. Let $\bar{C}(\mathbf{s}) = \sum_{i \in N_0} c_i(\mathbf{s})/k_0$ be the average cost experienced by players in N_0. We first show $c_j(\mathbf{s}) \geq \frac{1}{2}\bar{C}(\mathbf{s})$. If $N_0 = \emptyset$, then $c_j(\mathbf{s}) \geq \frac{1}{2}\bar{C}(\mathbf{s})$ trivially holds. Suppose that $N_0 \neq \emptyset$, and let $i \in N_0$ with $s_i = \{e\}$. Recall that i is selfish. Because $\mathbf{s}$ is a Nash equilibrium, we have $c_i(\mathbf{s}) = a_e f_e + b_e \leq a_{\bar{e}}(f_{\bar{e}}+1) + b_{\bar{e}} \leq 2(a_{\bar{e}} f_{\bar{e}} + b_{\bar{e}}) = 2c_j(\mathbf{s})$, where we lose a factor 2 in the second inequality compared to a corresponding bound in the proof of Theorem 4.3.2. By summing over all k_0 selfish players in N_0, we obtain $c_j(\mathbf{s}) \geq \frac{1}{2}\bar{C}(\mathbf{s})$ and thus $\sum_{j \in N_1} c_j(\mathbf{s}) \geq \frac{1}{2} k_1 \bar{C}(\mathbf{s})$. We have

$$
\begin{aligned}
\gamma &= \frac{\sum_{i \in N_0} c_i(\mathbf{s})}{\sum_{i \in N_0} c_i(\mathbf{s}) + \sum_{j \in N_1} C_j(\mathbf{s})} \\
&\leq \frac{k_0 \bar{C}(\mathbf{s})}{k_0 \bar{C}(\mathbf{s}) + \frac{1}{2} k_1 \bar{C}(\mathbf{s})} \\
&= \frac{2k_0}{n + k_0} \leq \frac{2n_0}{n + n_0},
\end{aligned}
$$

where the last inequality follows because $k_0 \leq n_0$. ■

Proof. [Theorem 5.3.2] Using the above lemmas, we can show that the relation in 5.4 holds:

$$\begin{aligned}
&\tfrac{3}{4}\gamma C(\mathbf{s}) + (1-\gamma)C(\mathbf{s}) = \tfrac{3}{4}\sum_{e\in E_0} f_e c_e(f_e) + \sum_{e\in E_1} f_e c_e(f_e) \\
\leq\ & \sum_{e\in E} y_e^* c_e(f_e^*) + \sum_{e\in E_1} f_e c_e(f_e) = \sum_{e\in E} f_e^* c_e(f_e^*) + \sum_{e\in E_1} (f_e c_e(f_e) - f_e c_e(f_e^*)) \\
\leq\ & \sum_{e\in E} f_e^* c_e(f_e^*) = C(\mathbf{s}^*),
\end{aligned}$$

where the first inequality follows from Lemma 5.3.5, and the last inequality follows from Lemma 5.3.3 and because cost functions are monotone increasing. We conclude that the pure Price of Anarchy is at most $(\frac{3}{4}\gamma + (1-\gamma))^{-1} = \frac{4}{4-\gamma}$. The stated bound now follows from Lemma 5.3.6. ■

In general, this upper bound may not be tight. Nevertheless, we conjecture that the gap can be very small if not tight. Observe that for non-uniform altruism with support $\{0,1\}$ with $n-1$ players being entirely altruistic and 1 player being entirely selfish, Lemma 5.3.6 gives a bound on γ of $\frac{2}{n+1}$, but Theorem 5.3.2 still manages to give a pure PoA bound of $\frac{2n+2}{2n+1}$, which is close to 1 when n is large. (We know that the pure PoA is 1 when every plater is entirely altruistic by Corollary 5.3.4.)

It is not hard to see that the above proof also goes through if every player has altruism level either α or 1. In fact, the only change is in Lemma 5.3.5, where we can replace the $\frac{4}{3}$ by $\frac{4}{3+\alpha}$.

Corollary 5.3.7 *The pure Price of Anarchy of altruistic linear symmetric singleton congestion games with* $\beta \in \{\alpha, 1\}^n$ *is at most* $1 + \frac{(1-\alpha)n_\alpha}{2n+(1+\alpha)n_\alpha}$*, where* n_α *is the number of players with altruism level* α*.*

It would be very interesting to extend this type of results to arbitrary distribution of altruism. We may face more technical difficulties since the allocation of players of different altruism levels is related to the allocation of players at optimum in much more complex way. We would like leave this part as future work.

5.4 Conclusions and Future Work

Intuitively, one would expect the Price of Anarchy of a game to improve when the altruism level β gets closer to $\mathbf{1}$, but we have seen that this is not the case. Indeed, there are important classes of games for which the robust Price of Anarchy turns out to be tight, and actually gets worse as the altruism level of the players increases. The fact that the Price of Anarchy does not *necessarily* get worse in all cases is exemplified by our analysis of symmetric singleton congestion games.

The most immediate future directions include extending the results about non-uniform altruism to arbitrary distribution of altruism and analyzing singleton congestion games with more general functions than linear ones. While the PoA of such functions increases (e.g., the PoA for polynomials increases exponentially in the degree for general congestion games [23]), this also creates room for potentially larger reductions due to altruism. For games where the smoothness argument cannot give tight bounds, would a refined smoothness argument like local smoothness in [85] work? For symmetric singleton congestion games, this seems unlikely, as the PoA bounds are already different between pure and mixed Nash equilibria.

It is also worth trying to apply the smoothness argument or its refinements to analyze the PoA for other dynamics in other classes of altruistic games. Furthermore, while the existence of pure Nash equilibria has been shown for singleton and matroid congestion games with player-specific latency functions [67, 1], the PoA (for pure Nash equilibria or more general equilibrium concepts) has not yet been addressed. Studying the PoA in such a general setting (in which our setting with altruism can be embedded) by either smoothness-based techniques or other methods is undoubtedly intriguing.

Chapter 6

Network Vaccination

We look at yet another class of network games, network vaccination games, in this chapter. However, it has a totally different flavor from congestion (or routing) games: each player actually resides on exactly one node in a social or computer network exclusively and each node has a player, which is not the case for the setting of congestion games; vaccinations or inoculation decisions can be made by node players to protect the social or computer network from being affected by outbreaks of epidemics or computer viruses in the network. However, every node player's selfish decision may not agree with the benefit of the society. For example, vaccination of a node may be a suboptimal decision for this node, but may help protect a large part of the network and therefore reduce the social cost. We will see that here the impact of altruism makes a difference from without altruism even for the existence of pure Nash equilibria. The results of this chapter are based on the paper [19] where, besides the results presented here, there are optimization results due to M. David.

6.1 Preliminaries

Recent epidemics of the Avian Flu and Swine Flu, among others, have reinforced the dramatic vulnerability of our society to outbreaks of epidemic diseases. Similarly, recent worms and viruses (such as Storm or Conficker) have shown us how severe consequences can be for massive infections of nodes in a computer network. Protecting social and computer networks from such outbreaks is a task of paramount social and economic importance.

Strategies for protecting a network fall roughly into two categories: preventive and reactive. Reactive strategies attempt to isolate nodes of the network (individuals or machines) once they have been diagnosed with an infection. Prophylactic vaccinations or inoculations protect nodes so that they will in the future not be affected by an outbreak[1]. Here, we focus on preventive strategies, i.e., the decisions which nodes in a network should be vaccinated.

We first describe the basic model of Aspnes et al. [4] with generalizations, allowing the vaccination cost, infection cost, and probability of initial infection outbreak to vary among nodes, and then extend it to include a notion of altruism. We show that with altruism, there are instances without pure Nash Equilibria. We therefore propose a notion of "opting out" from vaccinations, and define the Price of Opting Out.

[1] In practice, the protection may not be perfect; considering inoculations which succeed only with a certain probability is an interesting direction for future work.

6.1.1 Basic Model

The social or computer network is represented by an undirected graph $G = (V, E)$ of n nodes, each of which can either be vaccinated or unvaccinated. The cost for vaccination is C_v if node v chooses to vaccinate. Once all vaccination decisions are made, exactly one node becomes infected; the infection probability of v is p_v where $\sum_{v \in V} p_v = 1$. From that node, the infection spreads along edges of the graph to all unvaccinated nodes. However, no vaccinated nodes can become infected or pass on the infection. Let S be the set of vaccinated nodes, and $\Gamma_1, \ldots, \Gamma_k$ the connected components of $G \setminus S$, which means removing vaccinated nodes. If node v is unvaccinated and in component Γ_i, its probability of infection is $\sum_{u \in \Gamma_i} p_u$. The cost for becoming infected is L_v, leading to an expected cost of $L_v \cdot \sum_{u \in \Gamma_i} p_u$ if node v chooses to stay unvaccinated.

Since each node v in component Γ_i has expected cost $L_v \cdot \sum_{u \in \Gamma_i} p_u$, the total social cost with set S vaccinating is

$$C(S) = \sum_{v \in S} C_v + \sum_i \sum_{u \in \Gamma_i} p_u \sum_{v \in \Gamma_i} L_v. \tag{6.1}$$

We use S^* to denote the optimum set of nodes to vaccinate, i.e., the set minimizing $C(S)$.

While it is socially optimal to minimize $C(S)$, individual nodes' preferences may not align with this objective. An individual node will choose the strategy (be vaccinated or not be vaccinated) based only on its own tradeoff. That is, a selfish node will vaccinate if $C_v \leq L_v \cdot \sum_{u \in \Gamma_i} p_u$, and not vaccinate otherwise. Let

$$c_v(S) \quad = \quad \begin{cases} C_v & \text{if } v \in S \\ L_v \cdot \sum_{u \in \Gamma_i} p_u & \text{if } v \notin S, v \in \Gamma_i \end{cases} \tag{6.2}$$

denote the cost that node v experiences based on all players' vaccination decisions. Since players will act selfishly, this scenario leads naturally to a game termed the *inoculation game*.

Aspnes et al. [4] already established that there always is a pure Nash Equilibrium in this setting and showed a linear lower bound on the Price of Anarchy and Price of Stability.

Proposition 6.1.1 *Both the Price of Anarchy and the Price of Stability can be* $\Theta(n)$*.*

Proof. A simple example is a star graph with $C_v = C$, $L_v = L$ for every node v, and $C = L + \epsilon$. In the unique Nash Equilibrium, no player vaccinates, while the socially optimal solution vaccinates the center node of the star. The respective social costs are nL and $C + (1 - 1/n) \cdot L$. ■

6.1.2 Altruism

The basic model introduced above assumes that individuals are completely selfish and do not take into account the effects of their actions on other nodes. We therefore study the inoculation game by instantiating our general model of partial altruism, formally defined in Section 2.3. Again, the uniform altruism level of the nodes is denoted by β.

Definition 6.1.2 (Perceived Cost) *The* perceived cost *of a β-altruistic node is the convex combination*

$$c_v^{(\beta)}(S) \quad = \quad (1-\beta) \cdot c_v(S) + \beta \cdot C(S). \tag{6.3}$$

Thus, the perceived cost of a partially altruistic node is the convex combination of the individual cost (selfish part) and the social cost (altruistic part). As usual, a node will choose the strategy (vaccinate or do not vaccinate) which minimizes the perceived cost. The tradeoff is characterized by the following proposition, which can be obtained by simple rearranging.

Proposition 6.1.3 *Let $S \subseteq V \setminus \{v\}$ be the set of other nodes vaccinating, and Γ the component of $G \setminus S$ containing v. Let $\Gamma_1, \ldots, \Gamma_k$ be the subcomponents of Γ resulting from removing v from Γ. Then, v will prefer to be vaccinated if and only if*

$$C_v \quad \leq \quad (1-\beta) \sum_{u \in \Gamma} p_u \cdot L_v + \beta \cdot (\sum_{u \in \Gamma} p_u \sum_{v \in \Gamma} L_v - \sum_{i} \sum_{u \in \Gamma_i} p_u \sum_{v \in \Gamma_i} L_v).$$

Remark 6.1.4 *Our definition of perceived cost is similar to the notion of* friendship *used by Meier et al. [66]. In their case, the altruistic part does not consider the cost of* all *nodes, but just that of the* neighbors *of v in G. Thus, they model more the incentives due to friendship in a social network, while our model captures more a general notion of altruism toward all others.*

Proposition 6.1.5 *There is an instance such that for every β, the Price of Anarchy is* $\Theta(n)$.

Proof. For the complete bipartite graph $K_{2,n-2}$, if $C_v = C$, $L_v = L$ for every node v, and $C = L + \epsilon$, the state in which no node is vaccinated is a Nash Equilibrium regardless of the value of β. The calculation of the ratio is the same as in Proposition 6.1.1. (Notice that for β large enough, the state with both nodes on the left side vaccinated is also a Nash Equilibrium, and the Price of Stability is thus smaller than $\Theta(n)$.) ■

While the inoculation game with selfish players is a potential game and thus possesses pure Nash Equilibria [4], the introduction of partial altruism changes the situation significantly.

Proposition 6.1.6 *There exist instances of the inoculation game with partial altruism in which there is no pure Nash Equilibrium.*

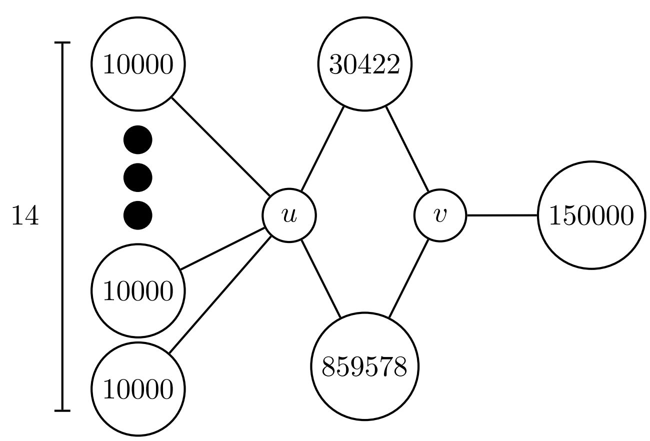

Figure 6.1: A graph without Nash Equilibrium

Proof. Consider the graph in Figure 6.1 with two nodes u, v and cliques of the indicated sizes. Whenever an edge is shown, there is an edge from u (or v) to *all* nodes in the corresponding clique. On the left, there are 14 cliques of size 10000 each. Thus,

$n = 14 \cdot 10000 + 30422 + 859578 + 150000 + 1 + 1 = 1180002$, and for all nodes, we set the infection cost to be $L = 1$ and the vaccination cost to be $C = 4.8015$. The altruism value is $\beta = 0.0000145$. We calculate $C \cdot n = 4.8015 \cdot 1180002 = 5665779.603$ for convenience.

The following calculations will show that

1. no node besides u or v ever wants to be vaccinated,
2. v wants to be vaccinated if and only if u is vaccinated, and
3. u wants to be vaccinated if and only if v is *not* vaccinated.

Thus, this instance encodes a "Matching Pennies" type of game and has no pure Nash Equilibrium.

1. This can be shown by using Proposition 6.1.3 for nodes other than u and v. If both u and v are vaccinated, any node (except u and v) inside a clique wants to stay unvaccinated:

$$\begin{aligned}
5665779.603 &\geq (1 - 0.0000145) \cdot 10000 + 0.0000145 \cdot (10000^2 - 9999^2 - 0^2) \\
&= 100000.14499; \\
5665779.603 &\geq (1 - 0.0000145) \cdot 150000 + 0.0000145 \cdot (150000^2 - 149999^2 - 0^2) \\
&= 150002.175; \\
5665779.603 &\geq (1 - 0.0000145) \cdot 30422 + 0.0000145 \cdot (30422^2 - 30421^2 - 0^2) \\
&= 30422.4411; \\
5665779.603 &\geq (1 - 0.0000145) \cdot 859578 + 0.0000145 \cdot (859578^2 - 859577^2 - 0^2) \\
&= 859590.4638;
\end{aligned}$$

If u is unvaccinated and v is vaccinated, any node (except u and v) inside the connected component merged by u wants to stay unvaccinated:

$$\begin{aligned}5665779.603 &\geq (1-0.0000145)\cdot(14\cdot 10000+30422+859578+1)\\&+0.0000145\cdot((14\cdot 10000+30422+859578+1)^2\\&-(14\cdot 10000+30422+859578)^2-0^2)\\&= 1030015.935;\end{aligned}$$

If both u and v are unvaccinated, any node (except u and v) inside the connected component merged by u and v wants to stay unvaccinated:

$$\begin{aligned}5665779.603 &\geq (1-0.0000145)\cdot(14\cdot 10000+30422+859578+1+150000+1)\\&+0.0000145\cdot((14\cdot 10000+30422+859578+1+150000+1)^2\\&-(14\cdot 10000+30422+859578+1+150000)^2-0^2)\\&= 1180019.11;\end{aligned}$$

If u is vaccinated and v is unvaccinated, any node (except u and v) inside the connected component merged by v wants to stay unvaccinated:

$$\begin{aligned}5665779.603 &\geq (1-0.0000145)\cdot(30422+859578+150000+1)\\&+0.0000145\cdot((30422+859578+150000+1)^2\\&-(30422+859578+150000)^2-0^2)\\&= 1040016.08;\end{aligned}$$

2. By using Proposition 6.1.3, we have the following conditions for v.

   If u is unvaccinated, v wants to get unvaccinated:

$$\begin{aligned}
5665779.603 &\geq (1-0.0000145)\cdot(14\cdot 10000+30422+859578+1+150000+1) \\
&\quad +0.0000145\cdot((14\cdot 10000+30422+859578+1+150000+1)^2 \\
&\quad -(14\cdot 10000+30422+859578+1)^2-150000^2) \\
&= 5660523.46. \qquad (6.4)
\end{aligned}$$

   If u is vaccinated, v wants to get vaccinated:

$$\begin{aligned}
5665779.603 &\leq (1-0.0000145)\cdot(30442+859578+150000+1) \\
&\quad +\beta((30442+859578+150000+1)^2-30422^2-859578^2-150000^2) \\
&= 5669868.455. \qquad (6.5)
\end{aligned}$$

   We thus conclude that 2. can be shown by using (6.4) and (6.5).

3. By using Proposition 6.1.3, we have the following conditions for u.

   If v is vaccinated, then u wants to get unvaccinated:

$$\begin{aligned}
5665779.603 &\geq (1-0.0000145)\cdot(14\cdot 10000+30422+859578+1) \\
&\quad +0.0000145\cdot((14\cdot 10000+30422+859578+1)^2-14\cdot 10000^2 \\
&\quad -30422^2-859578^2) \\
&= 5665668.31. \qquad (6.6)
\end{aligned}$$

If v is unvaccinated, u wants to get vaccinated:

$$\begin{aligned} 5665779.603 &\leq (1 - 0.0000145) \cdot (30442 + 859578 + 150000 + 1 + 14 \cdot 10000 + 1) \\ &\quad +0.0000145 \cdot ((30442 + 859578 + 150000 + 1 + 14 \cdot 10000 + 1)^2 \\ &\quad -(30442 + 859578 + 150000 + 1)^2 - 14 \cdot 10000^2) \\ &= 5666323.17. \end{aligned} \tag{6.7}$$

We thus conclude that 3. can be shown by using 6.6 and 6.7.

■

6.1.3 Opting Out

In light of Proposition 6.1.6, the standard notion of the Price of Stability [2] is not defined for pure Nash Equilibria in our game. Mixed Nash Equilibria are not a natural solution concept in inoculation games, because the decision of whether or not to be vaccinated tends to be permanent or very long-term. Nevertheless, it is of interest to analyze the effect that nodes' autonomy with respect to vaccination decisions has on the social cost.

Since the main concern with individual autonomy is *undervaccination* of the network (see the discussion at the end of Section 6.2), we consider a natural model of *opting out*. A benevolent authority suggests a set of nodes S_0 to vaccinate, such as the optimal solution S^* or an approximation. Subsequently, nodes that were targeted for vaccination have the option to override this decision, e.g., by not showing up for their vaccination. Notice that they cannot opt back in once opting out. However, we do not allow nodes

$v \notin S_0$ to override the decision and become vaccinated instead. Since the resulting dynamic is monotone in the number of vaccinated nodes, it will always converge to a final set of vaccinated nodes. However, this set of nodes may depend on the order in which nodes opt out.

For a starting set S_0, we define $\mathcal{R}(S_0)$ to be the collection of all node sets S such that the opting-out dynamic, starting from S_0, will eventually reach S. Formally, we define $\mathcal{R}(S_0)$ as follows:

Definition 6.1.7 *(1) $S_0 \in \mathcal{R}(S_0)$, and (2) If $S \in \mathcal{R}(S_0)$, and $v \in S$ prefers being unvaccinated, given that all nodes in $S \setminus \{v\}$ are vaccinated, then $S \setminus \{v\} \in \mathcal{R}(S_0)$.*

Note that we do not require each $S \in \mathcal{R}(S_0)$ itself to be stable; we also include in $\mathcal{R}(S_0)$ sets such that in later steps, further nodes will opt out. We then define the *Price of Opting Out* to be the worst-case ratio between the social cost of any set $S \in \mathcal{R}(S^*)$ and the social cost at S^*. Thus, the Price of Opting Out captures the increase in cost due to giving nodes the authority to opt out of vaccinations.

Our notion of the Price of Opting Out bears some similarity with the "Price of Sinking" defined by Goemans et al. [43] in the context of routing games and valid utility games. However, they considered the strongly connected sink component of the best-response graph, and considered bounds on the expected cost under the stationary distribution of a random walk. Our goal is to obtain bounds on *each* reachable state. Naturally, it is a question for future work to consider not only opt-out dynamics, but states reached by arbitrary sequences of best responses.

6.2 The Price of Opting Out

In this section, we bound the price of opting out, by analyzing any state that can be reached from an initially vaccinated set S_0 by a sequence of opt-out moves. Our main theorem is the following:

Theorem 6.2.1 *If S is obtained from S_0 by a sequence of opt-out moves, then $C(S) \leq \frac{1}{\beta} \cdot C(S_0)$.*

Theorem 6.2.1 thus in a sense captures the Price of Limited Autonomy: letting individuals choose not to be vaccinated when an optimal (or near-optimal) solution prescribes that they should be.

Proof. Let $\{v_1, \ldots, v_\ell\} = S_0 \setminus S$ be the set of all nodes who have opted out of being vaccinated, in the order in which they opted out. Let $S_t = S \cup \{v_{t+1}, \ldots, v_\ell\}$ be the set of nodes still vaccinated after t nodes have opted out, and $\Gamma_1^{(t)}, \ldots, \Gamma_{k_t}^{(t)}$ the connected components of $G \setminus S_t$. In particular, $\Gamma_1^{(0)}, \ldots, \Gamma_{k_0}^{(0)}$ are the connected components of the initial vaccinated set S_0. Define

$$\Phi(t) \quad := \quad \textstyle\sum_{v \in S_t} C_v + \beta \cdot \sum_i \sum_{u \in \Gamma_i^{(t)}} p_u \sum_{v \in \Gamma_i^{(t)}} L_v.$$

We will prove by induction that for all t, we have

$$\Phi(t) \quad \leq \quad C(S_0). \tag{6.8}$$

The base case $t = 0$ holds because $\beta \leq 1$, simply substituting the definition of social cost. Consider step t in which node $v = v_t$ decides to opt out of vaccinating. By ceasing to be vaccinated, v merges one or more components $\Gamma_i^{(t)}$ for $i \in M$, forming one new component $\Gamma_j^{(t+1)}$. Then,

$$\begin{aligned}\Phi(t+1) &= \Phi(t) - C_v + \beta \cdot (\textstyle\sum_{u \in \Gamma_j^{(t+1)}} p_u \sum_{v \in \Gamma_j^{(t+1)}} L_v \\ &- \textstyle\sum_{i \in M} \sum_{u \in \Gamma_i^{(t)}} p_u \sum_{v \in \Gamma_i^{(t)}} L_v).\end{aligned}$$

By Proposition 6.1.3, the fact that v chooses to opt out of vaccinating implies that

$$\begin{aligned}C_v &\geq (1-\beta) \textstyle\sum_{u \in \Gamma_j^{(t+1)}} p_u \cdot L_v \\ &\quad + \beta \cdot (\textstyle\sum_{u \in \Gamma_j^{(t+1)}} p_u \sum_{v \in \Gamma_j^{(t+1)}} L_v - \sum_{i \in M} \sum_{u \in \Gamma_i^{(t)}} p_u \sum_{v \in \Gamma_i^{(t)}} L_v) \\ &\geq \beta \cdot (\textstyle\sum_{u \in \Gamma_j^{(t+1)}} p_u \sum_{v \in \Gamma_j^{(t+1)}} L_v - \sum_{i \in M} \sum_{u \in \Gamma_i^{(t)}} p_u \sum_{v \in \Gamma_i^{(t)}} L_v).\end{aligned}$$

Substituting this inequality into $\Phi(t+1)$ shows that $\Phi(t+1) \leq \Phi(t)$, and the claim now follows by induction.

After all v_t have opted out, the total social cost is

$$\begin{aligned}C(S) &= \textstyle\sum_{v \in S} C_v + \sum_i \sum_{u \in \Gamma_i^{(\ell)}} p_u \sum_{v \in \Gamma_i^{(\ell)}} L_v \\ &\leq \tfrac{1}{\beta} \cdot \Phi(\ell) \\ &\leq \tfrac{1}{\beta} \cdot C(S_0),\end{aligned}$$

where the last step followed from the claim we proved by induction. ■

By applying Theorem 6.2.1 to the optimum set S^*, we obtain the following corollary:

Corollary 6.2.2 *The Price of Opting Out is at most* $1/\beta$*.*

Notice, however, that Theorem 6.2.1 is more general. In particular, it applies to approximately optimal starting sets S_0. While computing S^* itself is NP-complete [4], we have shown how to find an $O(\log n)$ approximation to the social cost [19]. Theorem 6.2.1 then guarantees that if we start with the set S_0 vaccinated, after allowing the nodes to opt out, the social cost will be within a factor $O(1/\beta \cdot \log n)$ of optimal, and stable to further opting out of nodes.

We also remark that Theorem 6.2.1 is tight.

Proposition 6.2.3 *There are instances with altruism β where the Price of Opting Out is $1/\beta$.*

Proof. Let $\beta > 0$ be arbitrary, and consider a star graph with n nodes. Let $L_v = 1$ and $C_v = 1 + \beta(n + 1/n - 2)$ for every node v. Then, the inequality $(1 + \beta(n + 1/n - 2)) \cdot n \geq (1 - \beta)n^2 + \beta(n^2 - (n-1) \cdot 1^2)$ for the central node shows that opting out always leads to a solution with no nodes vaccinated (none of the other nodes want to get vaccinated since vaccination by any of them cannot break up the star), for a total cost of n. On the other hand, the optimum solution vaccinates the center node of the star, giving a total cost of $1 + \beta(n + 1/n - 2) + 1 - 1/n = \beta n + O(1)$. As $n \to \infty$, the Price of Opting Out then converges to $1/\beta$. ∎

Instead of the Price of Opting Out, one could study the Price of Opting In. There are two reasons why this is not as natural an approach: (1) It is less realistic that one could force individuals to undergo vaccinations, and (2) The Price of Opting In is always 1. The inequality in Proposition 6.1.3 and some calculation show that if β-altruistic nodes prefer to switch their status to "vaccinated", the social cost always decreases.

Naturally, the most general dynamic one would want to study would combine opting in and opting out, considering any sequence of best-response steps. Ideally, one would want to prove a $1/\beta$ bound for *all* states reachable by such best responses. However, analyzing a full best-response dynamic in this sense appears to be quite challenging, and it is possible that it reaches states much less efficient than $1/\beta \cdot C(S^*)$.

6.3 Conclusions and Future Work

In this chapter, we presented bounds on the Price of Opting Out for a network inoculation game, in the presence of altruism. We believe that the Price of Opting Out result in particular has interesting consequences in terms of policy: it suggests that if individuals have the freedom to opt out of suggested vaccinations, then neither coordination of strategies nor socialization of costs alone will lead to an efficient outcome, yet the combination of both gives outcomes of much lower societal cost.

Naturally, many questions remain for future work. Most directly, our Price of Opting Out result should be generalized (if possible) to a more general notion of autonomy. Since pure Nash Equilibria may not exist in general, a natural (and very strong) result would be to show that *all* states reachable from the optimum by any sequence of individual best responses have the same cost bound. Such sequences could include arbitrary choices to vaccinate or not to vaccinate matching nodes' preferences. Such a result would significantly strengthen the policy implications of our results, since in most scenarios, individuals do have the freedom to decide whether or not they want to be vaccinated. Furthermore, while mixed Nash Equilibria are not an ideal solution concept for the type of game we study here, it would nevertheless be interesting from a theory point of view what they look like, and how efficient or inefficient they are.

Stronger bounds could also be obtained under additional assumptions about the network structure. For instance, most social networks have bounded degrees. Indeed, we can show that even in the basic model of Aspnes et al. without altruism, the Price of Anarchy is bounded by $\sqrt{n\Delta}$ if all degrees are bounded by Δ (whereas the general bound is $\Theta(n)$). The exact impact on the Price of Opting Out or a generalization constitutes an interesting direction for future work. Similarly, it would be interesting to study the impact of other graph parameters.

More generally, the model proposed by Aspnes et al. (which we study here) is somewhat simplistic. It assumes that each edge of the graph will deterministically transmit the infection, and that each vaccination will deterministically protect the node. Assigning (known) probabilities to both types of events would be much more natural, but most likely lead to a significantly more difficult optimization problem. For example, we are given a graph and told who is vaccinated; each edge transmits the infection with probability $p \neq 0, 1$. Then, computing the expected number of infected nodes from a uniformly random outbreak might be $\sharp P$-hard. Since evaluating the objective function might be already hard, things will not be better for the optimization problem.

Finally, our analysis assumes that all nodes know the full topology of the network. This is certainly not true in social networks, and it would be interesting to formulate a natural model of partial knowledge, and analyze its impact on the behavior of individuals in the network. Different nodes may have knowledge of different parts of the network, creating different "views" about connected components whose sizes may affect their decisions. Then, the existence of equilibria is not even clear any more. Let alone the PoA and PoS.

Chapter 7

Auctions

There is no concept of networks in the setting of the standard auction theory. In this chapter, we are going to see how the outcome of an auction is affected by being situated in an economic or social network. We just focus on the most basic auctions, single-item auctions where the auctioneer is selling an item to n bidders. In a standard auction (e.g., first- and second-price auctions), there is no concept of relations among bidders. The purpose here is not to change or design auction mechanisms but to see how equilibrium bidding strategies in a first- or second-price single-item auction change due to bidders having positive or negative relations that form an economic or social network.

7.1 Preliminaries

The traditional view of auctions posits that bidders only care if they win the item(s), and at what price. The utility of bidders not winning the auction is 0, regardless of the actual outcome. If the auction is conducted among perfectly rational strangers, the items are solely for resale, and no future competitive advantage is gained by winning an auction cheaply, this assumption is quite accurate. However, in many realistic scenarios, the bidders are embedded in social and economic networks, which will affect their perception of an auction's outcome. For instance, previous social or economic interactions may have led to positive or negative relationships between bidders. These in turn will cause different perceptions of the auction to a bidder, depending on *which* competitor wins, and at what price. A similar case can be made for potential future economic interactions: if two bidders are likely future collaborators, then one would derive benefit from the other's winning the auction; conversely, if they are likely future competitors, then an outcome advantageous to one is intrinsically threatening to the other.

These observations for auctions motivate our study of auctions in which the utility of losers is not always 0, but rather depends on the identity of the winner, and the utility the winner derives from the auction.

We assume throughout that we have n bidders with valuations v_i. Bids are denoted by b_i. We study auctions in which the auctioneer is trying to sell an item, and consider first- and second-price auctions Both for first- and second-price auctions, the auction mechanism selects as winner a bidder i maximizing b_i (breaking ties arbitrarily, but consistently). Let w be this winner, and $s \in \operatorname{argmax}_{i \neq w} b_i$ be a bidder making the second-highest bid. Then, the *threshold bid* for the winning bidder w is $\tau_w = b_s$.

To distinguish between the utility derived directly from winning in the auction, and the utility derived indirectly from other bidders' utilities due to spite or altruism, we call the former *subutility*, and the latter *perceived utility*. We can see that this also falls under our general model of altruism and spite. The general definition of Chapter 2 specializes to the following (perceived) utilities for bidder i.

1. Second-price auction for selling an item:

$$u_i = \begin{cases} \beta_{i,i} \cdot (v_i - \tau_i) & \text{for } i = w \\ \beta_{i,w} \cdot (v_w - \tau_w) & \text{for } i \neq w \end{cases} \tag{7.1}$$

2. First-price auction for selling an item:

$$u_i = \begin{cases} \beta_{i,i} \cdot (v_i - b_i) & \text{for } i = w \\ \beta_{i,w} \cdot (v_w - b_w) & \text{for } i \neq w \end{cases} \tag{7.2}$$

Remark 7.1.1 *The definition given here coincides with the one used by Brandt et al. [14]. The definition of Morgan et al. [69] differs only in that it omits the factor* $\beta_{i,i}$ *in front of the subutility from himself.*

The larger $|\beta_{i,j}|$, the more important the winning or losing of other bidders becomes to i. (Notice that we do not recursively consider the utility a bidder derives from another bidder's perceived utility. Such *systems of utility functions* studied by Bergstrom [9] have been discussed in Section 2.3.1.)

7.2 Bayesian Auctions

We study auctions in which the auctioneer is selling a single item to spiteful and altruistic bidders. Each bidder's value is private, but drawn from *the same* distribution, F, which is common knowledge among all bidders. Bidders are assumed to maximize expected utility.

We would like to derive general Nash Equilibrium conditions for arbitrary distributions F on $[0, 1]$ and spite levels, for both first- and second-price auctions. Since these conditions are too complicated to solve in general, we focus on two special cases:

1. Networks in which each bidder has the same number d of acquaintances, and feels the same spite/altruism level β toward all of his d neighbors. Thus, we have a social network in which each node has outdegree d, and all bidders have uniform spite/altruism. For this case, we analyze both first- and second-price auctions under arbitrary distributions F of valuations. We show that the revenue of the second-price auction dominates the first-price auction for $\beta < 0$, while the domination is reversed for $\beta > 0$.

2. First-price auctions with *triangular altruism* matrices with non-negative entries or *block spite/altruism* matrices B. Bidders' valuations are drawn *uniformly* from the interval $[0, 1]$. We present a (non-symmetric) Nash equilibrium.

Throughout, we identify the distribution F over $[0, 1]$ with its cumulative distribution function (cdf), and use $f = F'$ to denote its density function. Bidder i's bidding function is b_i: When bidder i has valuation v, she will bid $b_i(v)$. We denote by b_i^{-1} the inverse function of the bidding function, i.e., $b_i^{-1}(b)$ is the valuation v such that bidder i with valuation v would bid b.[1]

We begin by deriving equilibrium conditions for first- and second-price auctions in the fully general setting:

[1] We are thus implicitly assuming that the bidding functions are strictly increasing and continuous.

Lemma 7.2.1 *Assume that all valuation are drawn independently from the same distribution F over* $[0,1]$.

1. *Nash Equilibria of first-price auctions satisfy the following system of differential equations:*

$$\begin{aligned} & \textstyle\sum_{j\neq i} \beta_{i,i}(v-b_i(v)) \cdot \frac{f(b_j^{-1}(b_i(v)))\cdot {b_j^{-1}}'(b_i(v))}{F(b_j^{-1}(b_i(v)))} \\ & + \textstyle\sum_{j\neq i, b_j(1)\geq b_i(v)} \left(\beta_{i,j} b_i(v) - \beta_{i,j} b_j^{-1}(b_i(v))\right) \cdot \frac{f(b_j^{-1}(b_i(v)))\cdot {b_j^{-1}}'(b_i(v))}{F(b_j^{-1}(b_i(v)))} \\ = \; & \beta_{i,i}. \end{aligned} \tag{7.3}$$

2. *Nash Equilibria of second-price auctions satisfy the following system of differential equations:*

$$\begin{aligned} & \beta_{i,i} \cdot \textstyle\sum_{j\neq i} \frac{f(b_j^{-1}(b_i(v))\cdot {b_j^{-1}}'(b_i(v)))}{F(b_j^{-1}(b_i(v)))} \cdot (v-b_i(v)) \\ & + \textstyle\sum_{j\neq i} \frac{\beta_{i,j}}{F(b_j^{-1}(b_i(v)))} \cdot \Big(\\ & (1-F(b_j^{-1}(b_i(v)))) \cdot b_i(v) \cdot \textstyle\sum_{k\neq i,j, b_k(1)\geq b_i(v)} \frac{f(b_k^{-1}(b_i(v)))\cdot {b_k^{-1}}'(b_i(v))}{F(b_k^{-1}(b_i(v)))} \\ & \qquad + b_i(v) \cdot \textstyle\sum_{\ell\neq i} f(b_\ell^{-1}(b_i(v))) \cdot {b_\ell^{-1}}'(b_i(v)) \\ & \qquad - b_i(v) \cdot \textstyle\sum_{k\neq i,j} \frac{f(b_k^{-1}(b_i(v)))\cdot {b_k^{-1}}'(b_i(v))}{F(b_k^{-1}(b_i(v)))} - 1\Big) \\ & + \textstyle\sum_{j\neq i, b_j(1)\geq b_i(v)} \frac{\beta_{i,j}}{F(b_j^{-1}(b_i(v)))} \cdot \Big(- b_j^{-1}(b_i(v)) \cdot f(b_j^{-1}(b_i(v))) \cdot {b_j^{-1}}'(b_i(v))\Big) \\ = \; & -\textstyle\sum_{j\neq i} \beta_{i,j} \end{aligned} \tag{7.4}$$

Proof. We first consider the case of first-price auctions. Denote by $\mathcal{E}_j = [b_j(v_j) > b_k(v_k)$ for all $k \neq j]$ the event that bidder j wins the auction. Using linearity of expectations, the expected utility of bidder i with valuation v is

$$\begin{aligned} & \text{Prob}[\mathcal{E}_i] \cdot \beta_{i,i} \cdot (v - b_i(v)) + \textstyle\sum_{j \neq i} \text{Prob}[\mathcal{E}_j] \cdot \beta_{i,j} \cdot \Big(\mathrm{E}\left[v_j \mid \mathcal{E}_j\right] - \mathrm{E}\left[b_j(v_j) \mid \mathcal{E}_j\right] \Big) \\ = \; & \beta_{i,i} \cdot (v - b_i(v)) \cdot \textstyle\prod_{j \neq i} F(b_j^{-1}(b_i(v))) \\ & + \textstyle\sum_{j \neq i} \beta_{i,j} \cdot \Big(\int_{\min(1, b_j^{-1}(b_i(v)))}^{1} x f(x) \cdot \prod_{k \neq i,j} F(b_k^{-1}(b_j(x))) dx \\ & - \int_{\min(b_j(1), b_i(v))}^{b_j(1)} y f(b_j^{-1}(y)) b_j^{-1\prime}(y) \cdot \prod_{k \neq i,j} F(b_k^{-1}(y)) dy \Big). \end{aligned}$$

At Nash Equilibrium, the bidding strategy b_i must be optimal. Thus, we take a derivative with respect to $b_i(v)$ and set it to 0, obtaining

$$\begin{aligned} & \beta_{i,i} \cdot \Big(\textstyle\sum_{j \neq i} f(b_j^{-1}(b_i(v))) \cdot b_j^{-1\prime}(b_i(v)) \cdot (v - b_i(v)) \cdot \prod_{k \neq i,j} F(b_k^{-1}(b_i(v))) \\ & \qquad - \textstyle\prod_{j \neq i} F(b_j^{-1}(b_i(v))) \Big) \\ & + \textstyle\sum_{j \neq i, b_j(1) \geq b_i(v)} \beta_{i,j} \cdot \Big(\\ & \qquad - b_j^{-1}(b_i(v)) \cdot f(b_j^{-1}(b_i(v))) \cdot b_j^{-1\prime}(b_i(v)) \cdot \textstyle\prod_{k \neq i,j} F(b_k^{-1}(b_i(v))) \\ & \qquad + b_i(v) \cdot f(b_j^{-1}(b_i(v))) \cdot b_j^{-1\prime}(b_i(v)) \cdot \textstyle\prod_{k \neq i,j} F(b_k^{-1}(b_i(v))) \Big) \\ = \; & 0. \end{aligned}$$

Using that

$$\begin{aligned} \textstyle\prod_{k \neq i,j} F(b_k^{-1}(b_i(v))) &= \frac{\prod_k F(b_k^{-1}(b_i(v)))}{F(b_j^{-1}(b_i(v))) \cdot F(v)} \quad \text{and} \\ \textstyle\prod_{j \neq i} F(b_j^{-1}(b_i(v))) &= \frac{\prod_k F(b_k^{-1}(b_i(v)))}{F(v)}, \end{aligned} \tag{7.5}$$

rearranging and canceling yields the condition claimed in the lemma.

For second-price auctions, in addition to $\mathcal{E}_j$, we use $\mathcal{D}_{j,k} = [b_j(v_j) > b_k(v_k) > b_\ell(v_\ell)$ for all $\ell \neq j,k]$ for the event that bidder j wins the auction, and bidder k has the second-highest bid. The expected utility of bidder i with valuation v bidding $b_i(v)$ is

$$\begin{aligned}
& \beta_{i,i} \cdot (\text{Prob}[\mathcal{E}_i] \cdot v - \textstyle\sum_{j\neq i} \text{E}\left[b_j(v_j) \mid \mathcal{D}_{i,j}\right] \cdot \text{Prob}[\mathcal{D}_{i,j}]) \\
& + \textstyle\sum_{j\neq i} \beta_{i,j} \cdot \Big(\text{E}\left[v_j \mid \mathcal{E}_j\right] \cdot \text{Prob}[\mathcal{E}_j] \\
& \qquad - b_i(v) \cdot \text{Prob}[\mathcal{D}_{j,i}] - \textstyle\sum_{k\neq i,j} \text{E}\left[b_k(v_k) \mid \mathcal{D}_{j,k}\right] \cdot \text{Prob}[\mathcal{D}_{j,k}]\Big) \\
= \; & \beta_{i,i} \cdot \textstyle\prod_{j\neq i} F(b_j^{-1}(b_i(v))) \cdot \Big(v - \textstyle\sum_{j\neq i} \int_{b_j(0)}^{b_i(v)} x \cdot f(b_j^{-1}(x)) \cdot b_j^{-1\prime}(x) \cdot \prod_{\ell\neq i,j} F(b_\ell^{-1}(x))dx\Big) \\
& + \textstyle\sum_{j\neq i} \beta_{i,j} \cdot \Big(\int_{\min(1,b_j^{-1}(b_i(v)))}^{1} x \cdot f(x) \cdot \prod_{k\neq i,j} F(b_k^{-1}(b_j(x)))dx \\
& \qquad - b_i(v) \cdot (1 - F(b_j^{-1}(v))) \cdot \textstyle\prod_{k\neq i,j} F(b_k^{-1}(b_i(v))) \\
& - \textstyle\sum_{k\neq i,j} \int_{\min(b_k(1),b_i(v))}^{b_k(1)} x \cdot f(b_k^{-1}(x)) \cdot b_k^{-1\prime}(x) \cdot (1 - F(b_j^{-1}(x))) \cdot \prod_{\ell\neq i,j,k} F(b_\ell^{-1}(x))dx\Big).
\end{aligned}$$

Again, we take a derivative with respect to $b_i(v)$ and set it to 0.

Simplifying the resulting equation using

$$\textstyle\prod_{\ell\neq i,j,k} F(b_\ell^{-1}(b_i(v))) \;=\; \frac{\prod_\ell F(b_\ell^{-1}(b_i(v)))}{F(b_k^{-1}(b_i(v)))F(b_j^{-1}(b_i(v)))F(v)}$$

as well as the two identities (7.5), after rearranging, we get the condition as claimed in the lemma. ■

In general, the system of differential equations (7.3) does not admit a direct solution, due to the interplay between inverses of bidding functions. We therefore next focus on special cases where the particular form of bidding functions allows us to simplify the differential equations further.

7.2.1 Regular Social Networks

As the first special case, we consider regular social networks, i.e., those in which each node has the same out-degree d. Furthermore, we assume that for each pair of bidders (i, j) with a directed edge from i to j, the spite level is the same, $\beta_{i,j} = \beta$ for all i, j with an edge, and let $\beta_{i,i} = \alpha > 0$.

It turns out that under this scenario, both the first-price and second-price auction have a symmetric Bayesian Nash Equilibrium, i.e., a Nash Equilibrium in which all bidding functions are the same, $b_i = b$ for all i.

First-Price Auctions

Theorem 7.2.2 *There exists a Bayesian Nash equilibrium for first-price auctions in which all bidders bid* $b(v) = \mathrm{E}[X \mid X < v]$, *where* X *is a random variable with cdf* $F(x)^{n-1-d\beta/\alpha}$.

Proof. Substituting the symmetric guess $b_i = b$ for all i into the the system of differential equations (7.3), we can simplify the system by using that $b_j^{-1}(b_i(v)) = v$, $b_j^{-1\prime}(b_i(v)) = \frac{1}{b'(v)}$, and $b_j(1) \geq b_i(v)$ for all i, j into

$$\textstyle\sum_{j \neq i} \Big(\beta_{i,i}(v - b(v)) + \beta_{i,j} b(v) - \beta_{i,j} v \Big) \cdot \frac{f(v)}{F(v) b'(v)} \quad = \quad \beta_{i,i},$$

and, using the network structure, simplify further to

$$\Big(((n-1)\alpha - d\beta) \cdot (v - b(v)) \Big) \cdot \frac{f(v)}{F(v) b'(v)} \quad = \quad \alpha.$$

Solving for $b(v)$ gives us $b(v) = v - \frac{1}{n-1-d\beta/\alpha} \cdot \frac{F(v)b'(v)}{f(v)}$. This differential equation has solution

$$b(v) \quad = \quad F(v)^{-(n-1-d\beta/\alpha)} \cdot \int_0^v x \cdot (n-1-d\beta/\alpha) \cdot F(x)^{n-2-d\beta/\alpha} f(x)dx. \qquad (7.6)$$

Thus, we have proved the theorem. ■

Note that the bidding function can be interpreted as the expectation of the highest of $(n-1) - \frac{\beta d}{\alpha}$ private values below v, in spite of the fact that $(n-1) - \frac{\beta d}{\alpha}$ may be a fractional number. Notice that this theorem does not characterize *all* equilibria, and indeed, it seems very likely that this auction also possesses asymmetric Nash Equilibria.

Substituting the uniform distribution over $[0, 1]$ for every bidder's valuation, we obtain the following corollary:

Corollary 7.2.3 *There is a Bayesian Nash equilibrium for first-price auctions with all valuations uniformly distributed in* $[0, 1]$ *in which all bidders bid* $b(v) = (1 - \frac{\alpha}{n \cdot \alpha - \beta d}) \cdot v$.

In particular, when $d = n - 1$, Theorem 7.2.2 and Corollary 7.2.3 recover the results of Brandt et al. [14] who showed that $b(v) = \frac{n-1}{n+\beta} \cdot v$ for uniform spite levels (with $\alpha = 1 + \beta$), and those of Morgan et al. [69] (with $\alpha = 1$).

Second-Price Auctions

We next turn our attention to second-price auctions, and prove the following theorem.

Theorem 7.2.4 *There is a Bayesian Nash equilibrium for the second-price auction with regular friendship graphs in which all bidders bid* $b(v) = \mathrm{E}[X \mid X > v]$*, where* X *is a random variable with cdf*

$$1 - (1 - F(x))^{1+\frac{(n-1)\alpha}{\beta d}}.$$

Proof. We again substitute the symmetric guess $b_i = b$ for all i into the system (7.4) and simplify by using that $b_j^{-1}(b_i(v)) = v$, $b_j^{-1\prime}(b_i(v)) = \frac{1}{b'(v)}$, and $b_j(1) \geq b_i(v)$ for all i, j, canceling $b_i(v) \cdot \sum_{k \neq i,j, b_k(1) \geq b_i(v)} \frac{f(b_k^{-1}(b_i(v))) \cdot b_k^{-1\prime}(b_i(v))}{F(b_k^{-1}(b_i(v)))}$ and $-b_i(v) \cdot \sum_{k \neq i,j} \frac{f(b_k^{-1}(b_i(v))) \cdot b_k^{-1\prime}(b_i(v))}{F(b_k^{-1}(b_i(v)))}$ since $b_k(1) \geq b_i(v)$ for $k \neq i, j$ to obtain

$$\begin{aligned} &\beta_{i,i} \cdot \textstyle\sum_{j \neq i} \frac{f(v)}{F(v)b'(v)} \cdot (v - b(v)) \\ &+ \textstyle\sum_{j \neq i} \frac{\beta_{i,j}}{F(v)} \cdot \Big(- b(v) \sum_{k \neq i,j} \frac{f(v)}{b'(v)} + b(v) \sum_{\ell \neq i} \frac{f(v)}{b'(v)} - v \frac{f(v)}{b'(v)} - 1 \Big) \quad = \quad - \textstyle\sum_{j \neq i} \beta_{i,j}. \end{aligned}$$

Noting that the two sums inside the parentheses almost cancel out, i.e., $- \sum_{k \neq i,j} \frac{f(v)}{b'(v)} + \sum_{\ell \neq i} \frac{f(v)}{b'(v)} = \frac{f(v)}{b'(v)}$, pulling constant terms $\frac{f(v)}{F(v)b'(v)} \cdot (v - b(v))$ out of the sum on the left-hand side and $\frac{1}{F(v)}$ out of the sum on the right-hand side, and using that $\sum_{j \neq i} \beta_{i,j} = d\beta$ and $\beta_{i,i} = \alpha$ for all i, we simplify further to

$$\alpha \cdot (n-1) \cdot \frac{f(v)}{F(v)b'(v)} \cdot (v - b(v)) - \frac{f(v)}{F(v)b'(v)} \cdot d\beta \cdot (v - b(v)) \quad = \quad -(1 - \frac{1}{F(v)}) \cdot d\beta.$$

Rearranging yields the differential equation $b(v) = v + \frac{-\beta d \cdot (1-F(v)) \cdot b'(v)}{((n-1)\alpha - \beta d) \cdot f(v)}$, which has the solution

$$b(v) \quad = \quad \frac{1}{(1-F(v))^{1+\frac{(n-1)\alpha}{\beta d}}} \cdot \int_v^1 x \cdot (1 + \frac{(n-1)\alpha}{\beta d}) \cdot (1 - F(x))^{\frac{(n-1)\alpha}{\beta d}} f(x) dx.$$

Thus, we have proved the theorem. ■

Note similarly that the bidding function can be interpreted as the expectation of the lowest of $1 + \frac{(n-1)\alpha}{\beta d}$ private values above v. Substituting the uniform distribution over $[0, 1]$ for F gives us the following corollary:

Corollary 7.2.5 *There is a symmetric Bayesian Nash equilibrium for the second-price auction with all bids independently and uniformly drawn from* $[0, 1]$ *in which all bidders bid*

$$b(v) \quad = \quad (1 + \frac{\beta d}{(n-1)\alpha - 2\beta d}) \cdot v - \frac{\beta d}{(n-1)\alpha - 2\beta d}.$$

Again, when $d = n - 1$, Theorem 7.2.4 and Corollary 7.2.5 subsume the results for second-price auctions with uniform spite by Brandt et al. [14] who showed that $b(v) = \frac{v-\beta}{1-\beta}$ (with $\alpha = 1 + \beta$), and those by Morgan et al. [69] (with $\alpha = 1$). By combining Corollaries 7.2.3 and 7.2.5, we can compare the expected revenues of the first-price auction and second-price auction.

Theorem 7.2.6 *Assume that the social graph is regular, with uniform spite/friendship values* $\beta < \alpha$, *and that the valuations of all bidders are drawn independently and uniformly from* $[0, 1]$. *Then,*

1. *In the presence of uniform spite ($\beta < 0$), the expected revenue of the second-price auction dominates the expected revenue of the first-price auction.*

2. *In the presence of uniform altruism ($\beta > 0$), the expected revenue of the first-price auction dominates the expected revenue of the second-price auction.*

Proof. Let b_F and b_S denote the bidding functions for first- and second-price auctions, respectively. Also, let $V_{(1)}$ and $V_{(2)}$ be the highest and second-highest valuations among all bidders, respectively. Notice that because all bidders use the same bidding function, the highest valuation always corresponds to the highest bid, and the second-highest valuation to the second-highest bid.

The revenue of the first-price auction is thus $b_F(V_{(1)})$, while the revenue of the second-price auction is $b_S(V_{(2)})$. Notice that both bidding functions are linear, so we can use linearity of expectations. Furthermore, $\mathrm{E}\left[V_{(1)}\right] = \frac{n}{n+1}$, and $\mathrm{E}\left[V_{(2)}\right] = \frac{n-1}{n+1}$. Substituting these in the bidding functions of Corollaries 7.2.3 and 7.2.5 gives us that

$$\begin{aligned}\mathrm{E}\left[b_F(V_{(1)})\right] &= (1 - \tfrac{\alpha}{n\cdot\alpha-\beta d}) \cdot \tfrac{n}{n+1} = \tfrac{(n-1)\alpha-\beta d}{n\cdot\alpha-\beta d} \cdot \tfrac{n}{n+1},\\ \mathrm{E}\left[b_S(V_{(2)})\right] &= (1 + \tfrac{\beta d}{(n-1)\alpha-2\beta d}) \cdot \tfrac{n-1}{n+1} - \tfrac{\beta d}{(n-1)\alpha-2\beta d} = \tfrac{(n-1)\alpha-\beta d}{(n-1)\alpha-2\beta d} \cdot \tfrac{n-1}{n+1} - \tfrac{\beta d}{(n-1)\alpha-2\beta d}.\end{aligned}$$

The difference is

$$\begin{aligned}&\mathrm{E}\left[b_S(V_{(2)})\right] - \mathrm{E}\left[b_F(V_{(1)})\right]\\ =\ & \tfrac{(n-1)\alpha-\beta d}{(n-1)\alpha-2\beta d} \cdot \tfrac{n-1}{n+1} - \tfrac{\beta d}{(n-1)\alpha-2\beta d} - \tfrac{(n-1)\alpha-\beta d}{n\cdot\alpha-\beta d} \cdot \tfrac{n}{n+1}\\ =\ & \tfrac{(n-1)\cdot(n\cdot\alpha-\beta d)((n-1)\alpha-\beta d)-n\cdot((n-1)\alpha-2\beta d)\cdot((n-1)\alpha-\beta d)}{(n+1)\cdot((n-1)\alpha-2\beta d)\cdot(n\cdot\alpha-\beta d)} - \tfrac{\beta d}{(n-1)\alpha-2\beta d}\\ =\ & \tfrac{-\beta d\cdot(n\cdot\alpha-\beta d)}{((n-1)\alpha-2\beta d)\cdot(n\cdot\alpha-\beta d)} - \tfrac{-\alpha(n-1)\beta d+(\beta d)^2}{((n-1)\alpha-2\beta d)\cdot(n\cdot\alpha-\beta d)}\\ =\ & \tfrac{-\beta d\cdot\alpha}{((n-1)\alpha-2\beta d)\cdot(n\cdot\alpha-\beta d)}.\end{aligned}$$

Notice that the denominator is positive for all values of d and all $\beta \in (-1, 1)$. The numerator has the same sign as β. Thus, in the presence of uniform spite, the expected revenue of the second-price auction dominates the first-price auction. In the presence of uniform altruism, the expected revenue of the first-price auction dominates the second-price auction. ■

7.2.2 Uniform Valuations

Next, we consider the case of general matrices B with $\beta_{i,i} > 0$ for all i and $\beta_{i,i} > |\beta_{i,j}|$ for all $j \neq i$, but under the assumption that all valuations are independently and uniformly drawn from $[0, 1]$, i.e., $F(x) = x$ for $x \in [0, 1]$.

In this case, for first-price auctions with arbitrary B, we still do not know how to solve the corresponding differential system. Nonetheless, we can calculate a Bayesian Nash Equilibrium for first-price auctions explicitly if B is a *non-negative triangular* or *block* matrix, because there happens to be a Nash Equilibrium where each bidder bids $b_i(v) = \gamma_i v$ for some constant γ_i. Unfortunately, a guess of $b_i(v) = \gamma_i v$, or even $b_i(v) = \gamma_i v + \xi_i$, does not appear to lead to a solution of the corresponding differential system for second-price auctions, and we are not aware of any explicit characterization of an equilibrium of the second-price auction here.

Definition 7.2.7 *A non-negative upper triangular altruism matrix B is an altruism/spite matrix with $\beta_{i,j} \geq 0$ for $j > i$ and $\beta_{i,j} = 0$ for $j < i$, and $\beta_{i,i} > \beta_{i,j}$ for all $i, j \neq i$.*

Intuitively, this definition captures scenarios where the "friending power" of players is very asymmetric and according to some kind of ranking. Player i has friending power to choose to be altruistic, i.e., $\beta_{i,j} > 0$, or not, i.e, $\beta_{i,j} = 0$, to any player j for $j > i$. However, player i does not have friending power to any player j for $j < i$ so it must be the case that $\beta_{i,j} = 0$ for $j < i$. Notice that a player's friending power may not necessarily be proportional to his outdegree in the social or economic network. We study this model mostly because it is mathematically tractable, while it is doubtful how much it would capture real-world scenarios. Nevertheless, it somehow reflects the asymmetry of knowing and caring about other people in the real world.

Theorem 7.2.8 *Assume that all valuations are drawn independently and uniformly from* $[0, 1]$*. Let* B *be a non-negative upper triangular altruism matrix and* C *a matrix with entries* $c_{i,i} = -(n-1)$ *and* $c_{i,j} = \frac{\beta_{i,j}}{\beta_{i,i}}$ *for* $i \neq j$ *so* $c_{i,j} \geq 0$ *for* $j > i$ *and* $c_{i,j} = 0$ *for* $j < i$*. Then, there is a Bayesian Nash equilibrium for first-price auctions with* B*, where each bidder* i *bids* $b_i(v_i)$ *that satisfies*

1. $b_i(v_i) = \gamma_i v_i$ *with*

$$\gamma_i = \frac{\det(C)}{\det(C)-\det(C_i)},$$

where C_i *is formed by replacing the* i^{th} *column of* C *by all* 1*'s*

2. $b_{i+1}(1) \geq b_i(1)$ *for* $1 \leq i \leq n-1$*.*

Proof. Let $\lambda_i = 1 - \frac{1}{\gamma_i}$ for all i. We first see that $\lambda_i = \frac{\det(C_i)}{\det(C)}$ solves $-(n-1)\lambda_i + \sum_{j>i} \frac{\beta_{i,j}}{\beta_{i,i}} \lambda_j = 1$ for all i by Cramer's rule. (Equivalently, the vector λ of all λ_i entries solves $C \cdot \lambda = \mathbf{1}$.) This is equivalently saying that $\gamma_i = \frac{\det(C)}{\det(C)-\det(C_i)}$ solves for all i,

$$(n-1)\beta_{i,i}(\frac{1}{\gamma_i} - 1) + \sum_{j>i} \beta_{i,j}(1 - \frac{1}{\gamma_j}) = \beta_{i,i}. \qquad (7.7)$$

We then see that $\lambda_1 \leq \lambda_2 \leq ... \leq \lambda_n$: since $-(n-1)\lambda_i + \sum_{j>i} \frac{\beta_{i,j}}{\beta_{i,i}} \lambda_j = 1$ for all i, we can verify these inequalities by simply knowing that $\lambda_n = -1/(n-1)$ and $\frac{\beta_{i,j}}{\beta_{i,i}} \geq 0$ for all $j > i$. Thus, $\gamma_1 \leq \gamma_2 \leq ... \leq \gamma_n$. The inequality on the $b_i(1)$ now follows, i,e, $b_1(1) \leq b_2(1) \leq ... \leq b_n(1)$.

For each bidder i, $(b_j^{-1})'(b_i(v)) = \frac{1}{\gamma_j}$, and $b_j^{-1}(b_i(v)) = \frac{\gamma_i}{\gamma_j} \cdot v$ due to $b_i(v) = \gamma_i v$. The system (7.7) becomes for all i,

$$\sum_{j \neq i} \beta_{i,i}(v - b_i(v)) \cdot \frac{b_j^{-1\prime}(b_i(v))}{b_j^{-1}(b_i(v))} + \sum_{j>i} \beta_{i,j}(b_i(v) - b_j^{-1}(b_i(v))) \cdot \frac{b_j^{-1\prime}(b_i(v))}{b_j^{-1}(b_i(v))} = \beta_{i,i}. \quad (7.8)$$

Since $b_1(1) \leq b_2(1) \leq ... \leq b_n(1)$, we know that $b_j(1) \geq b_i(v)$ for all $j > i$. Therefore, along with $\beta_{i,j} = 0$ for all $j < i$, this system (7.8) becomes for all i,

$$\begin{aligned} & \textstyle\sum_{j \neq i} \beta_{i,i}(v - b_i(v)) \cdot \frac{b_j^{-1\prime}(b_i(v))}{b_j^{-1}(b_i(v))} + \sum_{j \neq i, b_j(1) \geq b_i(v)} \beta_{i,j}(b_i(v) - b_j^{-1}(b_i(v))) \cdot \frac{b_j^{-1\prime}(b_i(v))}{b_j^{-1}(b_i(v))} \\ = \; & \beta_{i,i}. \end{aligned}$$

With $F(x) = x$ and $f(x) = 1$ for all $x \in [0,1]$, this system is equivalent to the general system (7.3). We have shown that there is a Bayesian Nash equilibrium for first-price auctions with B, where each bidder i's bid satisfies (1) and (2). ■

If we solve the system explicitly for the bidding functions, we will see that players with more "friending power" actually bid less aggressively (with $b_i(1) = \gamma_i$ smaller), which roughly agrees with our intuition.

Another natural special case which can be solved easily using our general result (7.3) in Lemma 7.2.1 is that of disjoint cliques of friends in an auction. The bidders form g disjoint groups $S_1, \ldots, S_g$. Within group S_k, all bidders have altruism $\beta^{(k)}$ to each other (and $\beta_{i,i} = 1$). Across groups, bidders are indifferent, i.e., a bidder's altruism or spite level towards any other bidder who is not in his group is 0. Then, C is a block matrix, and the system of linear equations can be solved for each block separately. Due to symmetry, within each group S_k, all bidders will use the same bidding strategy, i.e., $\lambda_i = \lambda_j =: \lambda^{(k)}$ whenever $i, j \in S_k$. The linear equality thus simplifies to $-(n-1)\lambda^{(k)} + (|S_k| - 1) \cdot \beta^{(k)} \cdot \lambda^{(k)} = 1$, with the solution $\lambda^{(k)} = \frac{1}{(|S_k|-1)\cdot\beta^{(k)}-(n-1)}$. Substituting this into the definition of γ_i, we obtain the following corollary:

Corollary 7.2.9 *If the bidders form disjoint cliques S_k with mutual altruism $\beta^{(k)}$, and all valuations are drawn uniformly from $[0, 1]$, then there exists a Bayesian Nash Equilibrium in which each bidder $i \in S_k$ bids $\frac{n-1-\beta^{(k)}(|S_k|-1)}{n-\beta^{(k)}(|S_k|-1)} \cdot v_i$.*

Notice that this corollary reveals several interesting tendencies. First, both λ_i and γ_i are always less than 1, and decreasing in $|S_k|$ and $\beta^{(k)}$. This is not entirely unexpected, as bidders in large or tightly knit cliques feel less of a need to win the auction themselves, since they are more likely to derive utility from a friend's winning. What is perhaps more surprising is that the bidding strategy of a clique S_k does not depend on how large or tightly knit another group $S_{k'}$ is. While this follows readily from our general result, it is not at all apparent a priori, since another tightly knit group might bid lower, allowing group S_k to lower its bids safely as well.

Remark 7.2.10 (Altruism Changes) *We would like to investigate how bidder i's strategy changes if her altruism level $\beta_{i,j}$ toward another member of the network changes in an auction with a triangular altruism matrix or a block altruim/spite matrix. One would intuitively expect that if bidder i's altruism toward bidder j increases, then bidder* *i will always bid lower, i.e., decrease γ_i, because she derives more utility from bidder j's winning. (Indeed, for disjoint cliques, this intuition is borne out.) However, it may turn out that this is not always the case. Our partial results show that in response to the change of one $\beta_{i,j}$, the entire network's strategies adapt, and in some cases, this means that bidder i will lower her bid.*

Remark 7.2.11 (Calculating the Efficiency and Revenue) *It would also be interesting to calculate the efficiency and revenue of an auction with a triangular altruism matrix or a block altruim/spite matrix. Since γ_i can be solved explicitly for these two cases, calculating the expected value of the winning bid seems possible. We leave the analysis of efficiency and revenue (and thus the Price of Anarchy and Price of Stability) for future work.*

7.3 Conclusions and Future Work

We extended the analysis of auctions with spite and altruism among agents to the case of non-uniform spite matrices. We gave explicit characterizations of Nash Equilibria for first-price auctions with valuations drawn uniformly from $[0, 1]$ and triangular altruism matrices or block spite/altruism matrices B, and for first- and second-price auctions with arbitrary valuations and regular social networks.

Many questions remain for future work. For Bayesian auctions, we would first like to move beyond non-negative upper triangular matrices for first-price auctions even only with uniformly distributed valuations. Then, can we find a Nash Equilibrium for first- and second-price auctions in general? It appears that this is significantly more complex:

the fact that first-price auctions had a Nash Equilibrium in which each bidder simply multiplies her bid by a constant was fortuitous. Also, can we extend the analysis of first-price auctions to other distributions, or to priors that are not identical for different bidders?

Having characterized the Nash Equilibrium bidding strategies, we would also like to explicitly compute the revenue and social welfare of the auction. The main obstacle here is to find the expected value of the winning bid, which is now a maximum among n values drawn from different distributions. Calculating the revenue or social welfare would also let us quantify the impact of spite or altruism on the outcome of the auction.

Another intriguing question is whether agents can learn equilibrium bidding strategies using a natural algorithm. Assuming that each agent knows the entire matrix B is certainly unrealistic. Are there simple strategies (in the style of [10]) wherein each bidder adapts her bidding strategy based on the utility derived from earlier auctions?

Another auction model to consider non-uniform spite is full-information auctions where bidders know each other's valuations. For full-information auctions, a large number of questions remain as well. We can start with a model where each bidder just has one spite level as his parameter (different bidders can have different spite levels). For the case of two bidders with full-information, we may want to characterize all ϵ-Nash Equilibria since Nash Equilibria may not exist, and use this characterization to design auctions minimizing the Price of Anarchy or Price of Stability. Then, we would like to extend our results beyond two bidders. This is non-trivial: for instance, even for three bidders, it is possible to construct scenarios with different Nash Equilibria, *in which different bidders win*. However, the main difficulty in designing optimal auctions is that in general, there is not much correlation between the bidder winning the spiteful auction and the bidder winning at the social optimum. Should an extension to larger numbers of bidders succeed, the next step would be to consider arbitrary matrices B again, and to characterize all ϵ-Nash Equilibria in that case. It would also be interesting to prove that the optimal mechanism (in terms of minimizing the PoA or PoS) in all cases must be a simple linear scaling of each bid by a factor depending on spite levels.

Finally, we would like to extend these results beyond single-item auctions to more complex settings. A particularly promising direction would be the context of keyword auctions [59], as well as various combinatorial settings: for example, path auctions (a special case of auctions on a set systems), where an auctioneer tries to buy a path from bidders residing on edges in a graph; given the structure of bidders inherent from the auction's nature, how the friend or foe relations are formed is not even clear in this context.

Chapter 8

Discussion, Conclusions, and Future Work

8.1 Discussion and Conclusions

Standard game theory has traditionally posited the assumption that individual players are selfish and rational. They only care about the outcome of an economic interaction and their personal gain and loss through the interaction, not other players' gain and loss. However, our experiences and experiments in laboratories sometimes profoundly contradict this. For instance, our caring for the welfare of society or underprivileged people could drive us to some altruistic behavior sometimes, even sacrificing our own benefits. Our concern for the environment could guide us to some environmentally-friendly decisions, sometimes even giving up our own convenience. This could happen even when we know that the effect of a single person's decision is negligible. Furthermore, how can we explain a few players' contribution to the common pool in the experiments for public goods contribution games (see [57] for more discussion)? How can we explain a bidder's dislike for the winning of some competing bidders and indifference for that of other irrelevant bidders, which are not discussed in the standard auction setting where only allocation matters? There may be some psychological account for such regard in a player's interest, which would lead us to a more general interpretation of game theory via a more general utility theory. We think that a model that takes into account altruism and the opposite of it, spite or negative altruism, in real humans would be a better model of how real humans actually behave.

On the other hand, standard selfish theory usually predicts bad outcomes in terms of some global measures. It would be very interesting to see what we can do with our model that incorporates altruism and spite in real humans via defining players' "perceived" utility functions. We want to investigate whether standard game theory is too pessimistic (or maybe it turns out to be too optimistic) about what will happen with real humans in a system by revisiting some global measures on outcomes with our model. We can see that monetary offsets (such as taxes and tolls) added to individual players' payoffs have been designed to incentivize selfish players to deviate towards considering social welfare. Our consideration of altruism, which means empathy for others, can thus be viewed as in effect playing a similar role as such monetary terms.

We consider altruism as well as spite, i.e., negative altruism, in the model, since the opposite of altruistic behavior, spiteful behavior, can also be observed in reality. For instance, consider a scenario of players situated in social networks, where players' previous or future interactions affect them to form dislikes or preferences towards other players. A player would possibly experience a negative utility when some player that he has spite towards gains benefits.

Our choice for the model assumes that each player's (perceived) utility is a linear combination of his own payoff and all the other players' payoffs, given the social welfare as the global objective that we are looking at. We call a player's own payoff the selfish part and the rest,[1] roughly corresponding to the other players' payoffs, the altruistic (spiteful) part. An obvious advantage of this model is simplicity, which at the same time does not prevent us from meaningful and interesting analysis. Besides simplicity, we will first argue that the altruistic part and the global objective that we are concerned with should be considered together as a bundle. We specifically consider a *utilitarian social welfare*, i.e., the sum of all players' utilities, as the objective, so our current choice for perceived utility functions is a match. Then, we will argue that linearly combining the selfish and altruistic parts is a natural choice from the economic literature. On the other hand, the choice for *equilibrium solution concepts* is actually independent from the choice for *players' utilities*, since no matter how players' utility functions change, we still need to analyze the system at some kind of steady state. Thus, Nash equilibria are still the first natural choice for analyzing a one-shot game with full information, and Bayesian-Nash equilibria are a natural choice for a Bayesian setting, though both of them are not necessarily the only choices. Due to several known hardness results for finding Nash equilibria, recent work has begun analyzing the outcomes of natural response dynamics and more permissive solution concepts such as correlated or coarse correlated equilibria. In Chapter 5, we explored the analysis for more permissive equilibrium concepts for congestion games.

We are particularly interested in the effects of altruism on the utilitarian social welfare, which is the global objective that we are concerned with in this proposal. It is quite intuitive to model altruism as a perceived utility function with the altruistic part

[1] If we use the social welfare or the average social welfare here, the altruistic part may still contain at least some fraction of a player's own payoff.

that reflects this utilitarian social welfare, which shows altruistic players' caring for the global objective. Nevertheless, the utilitarian social welfare may not be the only global objective that can represent and measure what happens in a system. In other words, the function for the altruistic part is not a fixed choice but goes with the global objective. For instance, if the objective is not the utilitarian social welfare, but the *minimum utility* among all the players[2], and we want to maximize such a minimum utility, then our current choice for the altruistic part may not be the right choice anymore. Notice that it is also possible that even though the utilitarian social welfare looks at the sum of utilities, an individual's altruism could be about helping the minimum utility, or the other way round. This means that it is not necessary that a player's altruistic part agrees with the global objective. In this thesis, we only consider the models with an agreement between the choices of individual altruism and the global objective.

Actually, one of the future directions is to consider other interesting objective functions as well as their corresponding altruistic part designs. More broadly, what is the relationship between the choice of an objective function and the choice of the altruistic part to induce meaningful effects? For example, maybe considering a weighted utilitarian objective function is interesting. Even for the purely selfish model, not too many different objective functions have been explored.

Although our altruistic part was originally motivated to reflect a psychological account for altruistic/spiteful behavior, it can also be interpreted as a monetary term usually used to suppress selfish behavior in economics, i.e., taxes or tolls. That is, our design can be considered as a "psychological toll", while the other is a "monetary toll".

[2] In selfish routing, besides the average latency objective (the social cost), Roughgarden also considered the maximum latency objective to minimize [79].

If looking at the altruistic part this way, a linear combination of the selfish part and the altruistic part as a perceived utility can be generally viewed as a *quasi-linear* utility function in the microeconomic literature [65]. Quasi-linear utility functions are linear in one argument, generally the monetary term. Formally, such a utility function could be written as $U(x, y) = u(x) + by$, where b is a constant. Quasi-linearity is one of the restrictions on utility functions that enable us to draw inferences about all *indifference curves* between consumption of commodities x and y from a single curve. With quasi-linearity, the indifference curves can be shifted outward as consumption of y increases *without changing their slope*. Quasi-linearity also allows us to trade off smoothly between $u(x)$ and y using the same unit, which a non-quasi-linear utility may not provide.

We would like to point out that the choice for the equilibrium solution concept (theory of play) is independent from the choice for utility functions (theory of utility). Equilibrium solution concepts provide a way to analyze the system at a steady state where no individual player wants to deviate, given that every player is driven by the utility function that he perceives. No matter what is the choice of the utility function, we need to use a reasonable and natural solution concept. After all, game theory consists of two parts: theory of play and theory of utility. Currently, we think that intuitively Nash equilibrium for full-information games and Bayesian-Nash equilibrium for a Bayesian setting would be the first natural choices, when applying these stronger equilibrium solution concepts is possible (when they exist). Nonetheless, we also adopt more permissive equilibrium concepts such as mixed Nash equilibria, correlated equilibria, and coarse correlated equilibria in Chapter 5. This also means that we do not have to stick with some equilibrium concept if there exists no such a equilibrium. Notice that there are other equilibrium solution concepts that we probably do not want to use here. For example, dominant strategies may not exist for traffic routing due to the nature of the problem under study; iterated best response may cause cycles so the convergence can be unclear. For now, our work is more about analyzing a steady state given enough information and less about how strategies are derived and equilibrium is reached. Since we are planning to consider repeated games with learning where a Nash equilibrium may not be reached, equilibrium solution concepts such as ϵ-Nash equilibria or more permissive solution concepts like correlated equilibrium will become more and more attractive.

We demonstrated that our proposed model of altruism with current choices for players' utility functions and equilibrium solution concepts have effects on our current choice of global objective, i.e., social welfare. The trend of impact from altruism is different across classes of games. Improvements on the PoA are shown in some games while this trend is not the case for the other games in which worsening of the PoA is proved. Nevertheless, these choices are not just fixed but context-based, and what we are trying to propose is not only a specific model but also a flexible framework to facilitate analyzing the effects of altruism on the outcomes of strategic games and mechanisms. Furthermore, deciding truly what the "right" model is can be difficult. It may involve designing experiments to support or disprove the theory that is based on the proposed models. Thus, our work probably also suggests several interesting directions and problems for experimental work.

8.2 Future Work

We have summarized many interesting concrete extended questions in their corresponding related chapters. Now, we propose a few future directions in a bit more detail.

8.2.1 Learning in Repeated Games

In the real world, people often face a situation where they get to play a game multiple times or even for a long period of time. Then, how does a player choose his strategy in each run of such a *repeated game*? What kind of equilibria will the game converge to? How good or bad will the social welfare be at such equilibrium? These are all interesting questions to ask about repeated games, and they have been studied, for instance, *fictitious play*, under the theory of *learning* in games [39]. Intuitively, people will adapt their decisions run by run, which means that players would learn to play in repeated games. Among others, *no-regret* algorithms, which we will formally define later, is one of the most popular classes of learning algorithms with certain desirable properties. For example, there are works using no-regret algorithms to ensure convergence to ϵ-Nash equilibrium (their definition of ϵ-Nash equilibrium is a variation of the standard definition) [10] and analyzing the Price of Total Anarchy (at coarse correlated equilibrium) when everyone plays no-regret algorithms [11].

A no-regret sequence $\sigma_1, ..., \sigma_T$ of probability distributions over strategy profiles is defined by the property that the total expected payoff of each player is at most $o(T)$ less than that of the best fixed strategy in hindsight. Formally, for all i and $s'_i \in S_i$,

$$E[\sum_{t=1}^{T} p_i(s^t)] \geq E[\sum_{t=1}^{T} p_i(s'_i, s^t_{-i})] - o(T),$$

where $s^t \sim \sigma^t$ and $s^t_{-i} \sim \sigma^t_{-i}$ for every t. Intuitively, when each player plays the game more and more times, his total expected payoff will get closer and closer to that of the best fixed strategy in hindsight. Various *no-regret algorithms* are guaranteed to generate a no-regret sequence, for example, the multiplicative weights learning algorithm [61, 38] is a popular no-regret algorithm.

With our model of altruism and spite, there are more questions to be asked. In general, when we assume that the (especially, non-uniform) altruism levels are given, we would be interested to know how players with different altruism levels would converge to some kind of equilibrium within a reasonable number of time steps, as well as the PoA/PoS at such an equilibrium. (Recall that we have already analyzed the PoA at coarse correlated equilibria for linear congestion games in Section 5.2.) Furthermore, if the altruism levels are *not* given, a natural question would be how every player adaptively learns his equilibrium strategy, which means implicitly learning these altruism levels. In the following, we try to specify these questions in more concrete problem settings that we have studied.

No-regret algorithms have been shown to give fast convergence to ϵ-Nash equilibria for the standard selfish routing model [10]. The *per-time-step* regret of a user is defined as the difference between her average latency and the latency of the best fixed path in hindsight. That is, formally,

$$\frac{1}{T}\sum_{t=1}^{T} c^t - \min_{P\in\mathcal{P}} \frac{1}{T}\sum_{t=1}^{T}\sum_{e\in P} c_e(f_e^t),$$

where c_e is the edge cost function, T is the number of time steps, $f^1, ..., f^T$ is a series of flows, and a user experienced latencies $c^1, ..., c^T$. An algorithm is no-regret if, for any sequence of flows, the expected regret over internal randomness in the algorithm goes to 0 as T goes to infinity. In particular, if each user runs a no-regret algorithm, the average regret over all users, i.e., $\frac{1}{T}\sum_{t=1}^{T}\sum_{e\in E}\ell_e(f_e^t)f_e^t - \frac{1}{T}\min_{P\in\mathcal{P}}\sum_{t=1}^{T}\sum_{e\in P}\ell_e(f_e^t)$, also approaches 0. So it can be assumed that a function $R(T)$, which is an upper bound on the average regret, goes to 0 as T goes to infinity, i.e.,

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{e\in E}\ell_e(f_e^t)f_e^t \leq R(T) + \frac{1}{T}\min_{P\in\mathcal{P}}\sum_{t=1}^{T}\sum_{e\in P}\ell_e(f_e^t),$$

where $R(T)$ depends on the network size, n, and the maximum possible latency. Then, T_ϵ is defined as the number of time steps that it takes to get $R(T) = \epsilon$. However, it is possible for a flow f to have regret near 0 and yet still be far from a true Nash flow. We cannot expect that all users are taking cheapest paths at any time. We can only expect that *most* users take a nearly-cheapest path given a flow f. That is to say, equivalently, that a flow f is at ϵ-Nash equilibrium if the average cost under this flow is within ϵ of the minimum cost path under this flow,i.e., $C(f) - \min_{P\in\mathcal{P}}\sum_{e\in P}\ell_e(f_e) \leq \epsilon$.

Among other results, it was mainly shown that for no-regret algorithms, the *time-averaged* flow $\hat{f}$ is approaching equilibrium [10]. Specifically, for a given T_ϵ, bounds on the number of time steps before $\hat{f}$ is ϵ-Nash are obtained. This is basically showing that no-regret algorithms can be used to arrive at ϵ-Nash equilibrium for entirely selfish users. Therefore, naturally we would then want to ask how partially altruistic or spiteful users use no-regret algorithms to arrive at some kind of equilibrium as well. What is the impact of our altruistic and spiteful model on the convergence to equilibria? How will the convergence time change? The impact of uniform altruism and spite may be not hard to establish, but the impact of non-uniform altruism or spite is still challenging. Similar questions can be asked for (non-uniform) partially altruistic or spiteful users in atomic congestion games and other games, too.

When the altruism/spite matrix is not given, in the auctions setting, a bidder cannot use it to directly compute his bid so he must learn from the results (utilities) of earlier runs of auctions thereby to adjust his bid accordingly. We want to find learning algorithms that can gradually lead to equilibria, i.e., converging to the correct bids that altruistic bidders should have, within a reasonable number of runs. In the routing or congestion games setting, a player also needs to adaptively learn his choice of a route or resources, which implicitly involves learning the unknown (non-uniform) altruism levels of the other players. It is not clear if we can directly use certain learning algorithms such as no-regret algorithms that we employ when the altruism levels are given.

In another line of work, we have the PoA results where every player is running no-regret algorithms with uniform altruism, i.e., the coarse PoA with uniform altruism, in linear congestion games. Extending these results to non-uniform altruism may be technically challenging. Similar questions can be asked in the non-atomic setting and even for other games as well.

8.2.2 Other-Regarding Payoffs

Broadly, we would like to think beyond our current model of altruism and spite under which we try to capture the relation between not entirely selfish behavior and social welfare. In Chapter 3, we have seen several different models of other-regarding payoff functions focusing on different notions. Here, as an example of one of the future directions along this line, we are going to discuss another model that has been defined for *inequity averse* and *reciprocal* players, its impact on *cooperation* was shown in some simple classic games [34]. This may further motivate us to formally analyze the effects of inequity averse and reciprocal behavior on certain global measures capturing overall cooperation in the society.

Fehr and Schmidt [33] proposed a model capturing some sense of "fairness" in which a player is altruistic towards other players if their original payoffs are below an equitable benchmark but is spiteful when the other players' original payoffs are more than this benchmark. In several applications, for example, resource allocation, it seems natural to assume that an equitable allocation is an equal original payoff for all players. They use the following utility function to capture the notion of inequity aversion.

$$p_i(x_1, ..., x_n) = x_i - \frac{\alpha_i}{n-1} \sum_{j \neq i} \max\{x_j - x_i, 0\} - \frac{\beta_i}{n-1} \sum_{j \neq i} \max\{x_i - x_j, 0\}$$

Table 8.1: Prisoners' dilemma (Table 1 of [34]).

	Player 2 Cooperate	Player 2 Defect
Player 1 Cooperate	(2,2)	(0,3)
Player 1 Defect	(3,0)	(1,1)

Table 8.2: Prisoners' dilemma for inequity averse players (Table 2 of [34]).

	Player 2 Cooperate	Player 2 Defect
Player 1 Cooperate	(2,2)	$(0-3\alpha, 3-3\beta)$
Player 1 Defect	$(3-3\beta, 0-3\alpha)$	(1,1)

with $0 \leq \beta_i \leq \alpha_i$ (note that $\frac{\partial p_i}{\partial x_i} \geq 0$ if and only if $x_i \geq x_j$), where x_i is the original payoff allocated to player i. Intuitively, players experience inequity if they are worse off in their material terms than the other players, and they feel inequity if they are better off; players suffer more from inequity that is to their disadvantage in original payoffs than from inequity that is to their advantage in original payoffs.

Now, using this definition of an inequity averse player's perceived utility function, the original utilities in a two-player prisoners' dilemma in Table 8.1 give the perceived utilities in Table 8.2. For two players, the perceived utility of player i is $x_i - \alpha_i(x_j - x_i)$ if player i is worse off than player j (i.e., $x_j - x_i \geq 0$) or $x_i - \beta_i(x_i - xj)$ if player i is better off than player j (i.e., $x_i - x_j \geq 0$). Here, the two players are assumed to have the same parameters, i.e., $\alpha_1 = \alpha_2 = \alpha$ and $\beta_1 = \beta_2 = \beta$.

If player 2 is expected to cooperate, player 1 would choose between utility allocations (2,2) and (3,0). The perceived utility of player 1 at (2,2) is 2 since there is no inequity; the perceived utility of player 1 at (3,0) is $3-3\beta$ since there is inequity in which player 1 is better off. Thus, player 1 will reciprocate the expected cooperation of player 2 if $\beta > \frac{1}{3}$.

Given player 1's defection, if player 2 cooperates, the perceived utility of player 2 is then $0 - 3\alpha$; yet, if player 2 defects instead, the utility would become 1. This is saying that player 2 will always reciprocate defection because cooperating while the other defects gives less utility and more inequity. In Table 8.2, if $\beta > \frac{1}{3}$, there are two equilibria: Both cooperate and both defect. If the players believe that the other player cooperates, it is rational for each of them to cooperate. Inequity averse and reciprocal players are therefore conditional cooperators. They reciprocate cooperation in response to expected cooperation and defection in response to expected defection.

A prisoners' dilemma game is this basically turned into a cooperation game. Motivated by this, taking a step further, we would like to systematically explore and analyze the impact of inequity averse and reciprocal players in terms of certain global measures reflecting the cooperation level in the system for some suitable class of games. We are not trying to find a right or ultimate model with other-regarding payoffs, but more to explore the connection between the utility theory and the global measure that we are optimizing in general.

Bibliography

[1] H. Ackermann, H. Röglin, and B. Vöcking. Pure nash equilibria in player-specific and weighted congestion games. In *Proc. 4th Workshop on Internet and Network Economics (WINE)*, 2006.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proc. 45th IEEE Symp. on Foundations of Computer Science*, pages 295–304, 2004.

[3] I. Ashlagi, P. Krysta, and M. Tennenholtz. Social context games. In *Proc. 4th Workshop on Internet and Network Economics (WINE)*, pages 675–683, 2008.

[4] J. Aspnes, K. Chang, and A. Yamposlkiy. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. *Journal of Computer and System Sciences*, 72(6):1077–1093, 2006.

[5] R. J. Aumann. Subjectivity and correlation in randomized strategies. *J. Mathematical Economics*, 1(1):67–96, 1974.

[6] M. Babaioff, R. Kleinberg, and C.Papadimitriou. Congestion games with malicious players. In *Proc. 8th ACM Conf. on Electronic Commerce*, pages 103–112, 2007.

[7] A. Barabási, R. Albert, and H. Jeong. Mean-field theory for scale-free random networks. *Physica A*, 272:173–187, 1999.

[8] M. Beckmann, C. McGuire, and C. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.

[9] T. Bergstrom. Systems of benevolent utility functions. *J. of Public Economic Theory*, 1(1):71–100, 1999.

[10] A. Blum, E. Even-Dar, and K. Ligett. Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games. In *Proc. 25th ACM Symp. on Principles of Distributed Computing*, pages 45–52, 2006.

[11] A. Blum, M. T. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *Proc. 40th Annual ACM Symposium on Theory of Computing*, STOC '08, pages 373–382. ACM, 2008.

[12] V. Bonifaci, T. Harks, and G. Schäfer. The impact of stackelberg routing in general networks. Technical Report COGA Preprint 020-2007, TU Berlin, 2007.

[13] M. Bradonjic, G. Ercal-Ozkaya, A. Meyerson, and A. Roytman. On the price of mediation. In *Proc. 11th ACM Conference on Electronic Commerce*. ACM, 2009.

[14] F. Brandt, T. Sandholm, and Y. Shoham. Spiteful bidding in sealed-bid auctions. In *Proc. 22nd Intl. Joint Conf. on Artificial Intelligence*, pages 1207–1214, 2007.

[15] F. Brandt and G. Weiß. Antisocial agents and Vickrey auctions. In *Proc. 8th Workshop on Agent Theories, Architectures and Languages*, pages 335–347, 2001.

[16] I. Brocas. Endogenous entry in auctions with negative externalities. *Theory and Decision*, 54(2):125–149, 2003.

[17] C. F. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, 2003.

[18] G. Charness and M. Rabin. Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, 117:817–869, 2002.

[19] P. Chen, M. David, and D. Kempe. Better vaccination strategies for better people. In *Proc. 12th ACM Conference on Electronic Commerce*, 2010.

[20] P. Chen, B. de Keijzer, D. Kempe, and G. Schäfer. Robust price of anarchy for atomic games with altruistic players. Submitted, 2011.

[21] P. Chen and D. Kempe. Altruism, selfishness, and spite in traffic routing. In *Proc. 9th ACM Conf. on Electronic Commerce*, pages 140–149, 2008.

[22] X. Chen and X. Deng. Settling the complexity of two-player nash-equilibrium. In *Proc. 47th IEEE Symp. on Foundations of Computer Science*, pages 261–270, 2006.

[23] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *In Proc. 37th Annual ACM Symposium on Theory of Computing*. ACM, 2005.

[24] F. Chung. *Spectral Graph Theory*. American Mathematical Society, 1997.

[25] R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proc. 35th ACM Symp. on Theory of Computing*, pages 521–530, 2003.

[26] R. Cole, Y. Dodis, and T. Roughgarden. How much can taxes help selfish routing? *Journal of Computer and System Sciences*, 72:444–467, 2006.

[27] R. Cominetti, J. Correa, and N. Stier-Moses. Network games with atomic players. In *Proc. 33rd Intl. Colloq. on Automata, Languages and Programming*, pages 525–536, 2006.

[28] J. C. Cox. How to identify trust and reciprocity. *Games and Economic Behavior*, 46:260–281, 2004.

[29] S. Dafermos. The traffic assignment problem for multiclass-user transportation networks. *Transportation Science*, 6:73–87, 1972.

[30] X. Deng and Q. Qi. Priority right auction for komi setting, 2009. Available at SSRN: http://ssrn.com/abstract=1405294.

[31] Z. Dezső and A. Barabási. Halting viruses in scale-free networks. *Physical Review E*, 65, 2002.

[32] J. Diaz, D. Mitsche, N. Rustagi, and J. Saia. On the power of mediators. In *Proc. 5th Workshop on Internet and Network Economics (WINE)*, pages 455–462, 2009.

[33] E. Fehr and K. Schmidt. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114:817–868, 1999.

[34] E. Fehr and K. M. Schmidt. The economics of fairness, reciprocity and altruism - experimental evidence and new theories. Technical Report Munich Discussion Paper No. 2005-20, Department of Economics, University of Munich, 2005.

[35] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-riding and white-washing in peer-to-peer systems. In *Proc. ACM SIGCOMM'04 Workshop on Practice and Theory of Incentives in Networked Systems (PINS)*, 2004.

[36] L. Fleischer. Linear tolls suffice: New bounds and algorithms for tolls in single source networks. *Theoretical Computer Science*, 348:217–225, 2005.

[37] L. Fleischer, K. Jain, and M. Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *Proc. 45th IEEE Symp. on Foundations of Computer Science*, pages 277–285, 2004.

[38] Y. Freund and R. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29:79–103, 1999.

[39] D. Fudenberg and D. Levine. *The Theory of Learning in Games*. The MIT Press, 1998.

[40] N. Garg, D. Grosu, and V. Chaudhary. An antisocial strategy for scheduling mechanisms. In *Proc. 19th IEEE Intl. Parallel and Distributed Processing Symp.*, 2005.

[41] J. Geanakoplos, D. Pearce, and E. Stacchetti. Psychological games and sequential rationality. *Games and Economic Behavior*, 1(1):60–79, March 1989.

[42] H. Gintis, S. Bowles, R. Boyd, and E. Fehr. *Moral Sentiments and Material Interests : The Foundations of Cooperation in Economic Life*. MIT Press, 2005.

[43] M. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, pages 142–154, 2005.

[44] J. Hannan. Approximation to bayes risk in repeated plays. *Contributions to the Theory of Games*, 3:97–193, 1957.

[45] A. Hayrapetyan, E. Tardos, and T. Wexler. The effect of collusion in congestion games. In *Proc. 37th ACM Symp. on Theory of Computing*, pages 89–98, 2006.

[46] M. Hoefer and A. Skopalik. Altruism in atomic congestion games. In *Proc. 17th European Symposium on Algorithms*, volume 5757 of *Lecture Notes in Computer Science*. Springer, 2009.

[47] M. Jackson. *Social and Economic Networks*. Princeton University Press, 2008. forthcoming.

[48] O. Jahn, R. Möhring, A. Schulz, and N. Stier-Moses. System-optimal routing of traffic flows with user constraints in networks with congestion. *Operations Research*, 53:600–616, 2006.

[49] P. Jehiel and B. Moldovanu. Allocative and informational externalities in auctions and related mechanisms. In *Proc. 9th World Congress of the Econometric Society*, 2006.

[50] P. Jehiel, B. Moldovanu, and E. Stacchetti. How (not) to sell nuclear weapons. *American Economic Review*, 86(4):814–829, 1996.

[51] P. Jehiel, B. Moldovanu, and E. Stacchetti. Multidimensional mechanism design for auctions with externalities. *Journal of Economic Theory*, 85(2):258–294, 1999.

[52] A. Kaporis and P. Spirakis. The price of optimum in stackelberg games on arbitrary single commodity networks and latency functions. In *Proc. 19th ACM Symp. on Parallel Algorithms and Architectures*, pages 19–28, 2006.

[53] G. Karakostas and S. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous users. In *Proc. 45th IEEE Symp. on Foundations of Computer Science*, pages 268–276, 2004.

[54] G. Karakostas and S. Kolliopoulos. Stackelberg strategies for selfish routing in general multicommodity networks. Technical Report 2006/08, McMaster University, 2006.

[55] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proc. 17th Annual Symp. on Theoretical Aspects of Computer Science*, pages 404–413. Springer, 1999.

[56] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal. Stochastic models for the web graph. In *Proc. 41st IEEE Symp. on Foundations of Computer Science*, pages 57–65, 2000.

[57] John Ledyard. Public goods: A survey of experimental resesearch. In J. Kagel and A. Roth, editors, *Handbook of Experimental Economics*, pages 111–194. Princeton University Press, 1997.

[58] David Levine. Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1:593–622, 1998.

[59] L. Liang and Q. Qi. Cooperative or vindictive: Bidding strategies in sponsored search auctions. In *Proc. 3rd Workshop on Internet and Network Economics (WINE)*, pages 167–178, 2007.

[60] H. Lin, T. Roughgarden, E. Tardos, and A. Walkover. Braess's paradox, fibonacci numbers, and exponential inapproximability. In *Proc. 32nd Intl. Colloq. on Automata, Languages and Programming*, pages 497–512, 2005.

[61] N. Littlestone and M. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994.

[62] T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. *Theoretical Computer Science*, 406:187–206, October 2008.

[63] E. Maasland and S. Onderstal. Auctions with financial externalities. *Economic Theory*, 32(3):551–574, 2007.

[64] A. Mas-Colell. On a theorem of Schmeidler. *J. of Mathematical Economics*, 13:201–206, 1984.

[65] A. Mas-Collel, W. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[66] D. Meier, Y. Oswald, S. Schmid, and Roger Wattenhofer. On the windfall of friendship: Inoculation strategies on social networks. In *Proc. 9th ACM Conf. on Electronic Commerce*, pages 294–301, 2008.

[67] I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996.

[68] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.

[69] J. Morgan, K. Steiglitz, and G. Reis. The spite motive and equilibrium behavior in auctions. *Contributions to Economic Analysis & Policy*, 2(1):1102–1127, 2003.

[70] T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: Byzantine players in a virus inoculation game. In *Proc. 25th ACM Symp. on Principles of Distributed Computing*, pages 35–44, 2006.

[71] U. Nadav and T. Roughgarden. The limits of smoothness: A primal-dual framework for price of anarchy bounds. In *Proc. 6th Workshop on Internet and Network Economics (WINE)*, pages 319–326, 2010.

[72] J. F. Nash. Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. USA*, 36:48–49, 1950.

[73] I. E. Nesterov, A. S. Nemirovskii, and Y. Nesterov. *Interior Point Polynomial Algorithms in Convex Programming*. Soc for Industrial & Applied Math, 1994.

[74] N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[75] A. Pigou. *The Economics of Welfare*. Macmillan, 1920.

[76] R. W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[77] A. Roth. The price of malice in linear congestion games. In *Proc. 4th Workshop on Internet and Network Economics (WINE)*, pages 118–125, 2008.

[78] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67:341–364, 2003.

[79] T. Roughgarden. The maximum latency of selfish routing. In *Proc. 15th ACM Symp. on Discrete Algorithms*, pages 973–974, 2004.

[80] T. Roughgarden. Stackelberg scheduling strategies. *SIAM J. on Computing*, 33:332–350, 2004.

[81] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[82] T. Roughgarden. Selfish routing with atomic players. In *Proc. 16th ACM Symp. on Discrete Algorithms*, pages 1184–1185, 2005.

[83] T. Roughgarden. On the severity of braess's paradox: Designing networks for selfish users is hard. *Journal of Computer and System Sciences*, 72:922–953, 2006.

[84] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proc. 41st Annual ACM Symposium on Theory of Computing*, pages 513–522. ACM, 2009.

[85] T. Roughgarden and F. Schoppmann. Local smoothness and the price of anarchy in atomic splittable congestion games. In *Proc. 22nd ACM-SIAM Symposium on Discrete Algorithms*, 2011.

[86] T. Roughgarden and E. Tardos. How bad is selfish routing? In *Proc. 41st IEEE Symp. on Foundations of Computer Science*, 2000.

[87] T. Schelling. *Micromotives and Macrobehavior*. Norton, 1978.

[88] Y. Sharma and D. Williamson. Stackelberg thresholds in network routing games or the value of altruism. In *Proc. 8th ACM Conf. on Electronic Commerce*, 2007.

[89] M. Smith. The marginal cost taxation of a transportation network. *Transportation Resesarch, Part B*, 13:237–242, 1979.

[90] C. Swamy. The effectiveness of stackelberg strategies and tolls for network congestion games. In *Proc. 18th ACM Symp. on Discrete Algorithms*, 2007.

[91] I. Vetsikas and N. Jennings. Outperforming the competition in multi-unit sealed bid auctions. In *Proc. 6thIntl. Conf. on Autonomous Agents and Multiagent Systems*, pages 702–709, 2007.

[92] A. Vetta. Nash equlibria in competitive societies with applications to facility location, traffic routing and auctions. In *Proc. 43rd IEEE Symp. on Foundations of Computer Science*, pages 416–425, 2002.

[93] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

[94] J. Wardrop. Some theoretical aspects of road traffic research. In *Proc. of the Institute of Civil Engineers, Pt. II*, volume 1, pages 325–378, 1952.

[95] D. Watts and S. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.