Smoothness Theorem for Incomplete Information Games with Altruistic Players (Final Project in Computational Games Theory Course)

Alon Ardenboim
I.D. 060870953

Abstract

We unify the two notions of smoothness used to prove price of anarchy bounds for games of incomplete information, and for games where players have some concern for the welfare of others (altruistic extensions of games). Using the new notion of smoothness, we prove an extension theorem to bound to price of anarchy for these kind of games. We raise some issues regarding the way we formulated the payoff functions in the altruistic extensions of incomplete information games. Finally, we show a slightly simpler way of defining smoothness in games of incomplete information than the previously introduced notion.

1 Introduction

In many situations, strategic entities may sometime act in a way which benefits themselves, but damages the overall "good" of society (cynically demonstrated in real life by politicians). The conflict between the self interest of strategic players and the desired social goal has received a great focus in the field of game theory, with the prisoner's dilemma being a canonical example of this conflict. In order to measure the inefficiency of such a selfish behavior, a concept named *Price Of Anarchy* (**PoA**) was termed [1]. Loosely speaking, given an objective function f, say, a maximization function such as social welfare, the price of anarchy is the result of dividing the objective function at the worst equilibrium of the game to the maximum value of this objective function. That is:

$$\text{PoA} = \min_{e \in \text{Eq}} \frac{f(e)}{f(\text{OPT})}. \tag{1.1}$$

The price of anarchy can be measured in contexts of different equilibria concepts such as pure \ mixed \ correlated equilibrium, and be called pure \ mixed \ correlated PoA accordingly.

A robust way to measure the PoA of games was introduced by Roughgarden [2] using the *smoothness* notion. To give an intuition on how to use such a notion, let's consider a payoff-dominating maximization goal function $P : S \mapsto \mathbb{R}^+$, $S = S_1 \times \cdots \times S_n$ being the strategy space of the players. By *payoff-dominating* we mean that the sum of players' individual payoff functions is bounded by the objective function. Smoothness in full information games is defined as follows:

Definition 1.1. A game $G = (S, \{P_i\}_{i \in [n]})$ ($P_i : S \mapsto \mathbb{R}^+$ being the payoff function of player i) is (λ, μ)*-smooth* with respect to a strategy profile $\mathbf{s}^* \in S$ and a maximization objective function P if

$$\sum_{i=1}^{n} P_i(s_i^*, \mathbf{s}_{-i}) \geq \lambda \cdot P(\mathbf{s}^*) - \mu \cdot P(\mathbf{s}) \tag{1.2}$$

for every strategy profile $\mathbf{s} \in S$.

As we shall see next, we can show a lower bound on the PoA of a (λ, μ)-*smooth* game. Let's assume we have a payoff dominating maximization objective P that reaches its maximum at $\mathbf{s}^* \in S$. Then if a game G is (λ, μ)-*smooth* with respect to such P and $\mathbf{s}^*$, we have that for every Nash equilibrium strategy profile $\mathbf{s}$,

$$P(\mathbf{s}) \geq \sum_{i=1}^{n} P_i(\mathbf{s}) \geq \sum_{i=1}^{n} P_i(s_i^*, \mathbf{s}_{-i}) \geq \lambda \cdot P(\mathbf{s}^*) - \mu \cdot P(\mathbf{s}) \tag{1.3}$$

where the inequalities are derived from the payoff dominance of P, the fact that $\mathbf{s}$ is a NE, and the smoothness property of G respectively.

Rearranging gives us

$$\text{PoA} = \min_{s \in \text{NE}} \frac{P(\mathbf{s})}{P(\mathbf{s}^*)} \geq \frac{\lambda}{1+\mu}. \tag{1.4}$$

A nice property we can get by using the notion in 1.1 is that we can also show bounds on much broader equilibrium concepts of full information games than pure NE. In fact, the PoA bounds proved using this notion generalize even to a wide equilibrium concept such as coarse correlated equilibrium [2].

Although this result is very robust in dealing with many interesting game settings, recently, the focus of the computer science community has shifted towards more complex game concepts. One of those concepts deals with Bayesian settings where players don't have any information about the payoff functions of other players, but rather on some prior distribution from which such a function is sampled (see [3, 4] for more information). Another interesting game concept deals with players that don't care entirely about their own payoff function, but also the payoffs of other players. These players can be regarded as "Altruistic" or "Spiteful" players, regarding on the way they encompass the other players' utilities into their own payoffs (for more information, check [5, 6, 7]).

While bounding the PoA using a similar notion to 1.1 was shown to both incomplete information games and altruistic games [4, 7], a unifying smoothness theorem has yet to be achieved.

2 Preliminaries

2.1 Altruism in Games

Let $G = (S = S_1 \times \cdots \times S_n, \{P_i\}_{i \in [n]})$ be a strategic game with n players, where S_i is the strategy space of player i, and $P_i : S \mapsto \mathbb{R}^+$ the payoff function of player i mapping the combined strategy profile $\mathbf{s} \in S$ into her payoff, with every player wanting to maximize her payoff. A payoff dominating *social welfare* function $P : S \mapsto \mathbb{R}^+$ maps the combined strategy profile into the combined welfare of the players. By payoff dominating, we mean that $\sum_{i=1}^n P_i(\mathbf{s}) \leq P(\mathbf{s})$ for every strategy profile $\mathbf{s}$ (you can think of this in a mechanism design context where there are no negative transfers from the mechanism to the players). The altruistic extension we adopted is taken from [7] and defined in the following manner:

Definition 2.1. Given a game G and a vector $\alpha \in [0,1]^n$, the α-*altruistic* extension of G parameterized with a payoff-dominating social welfare function P is defined as a game $G = (S = S_1 \times \cdots \times S_n, \{P_i^\alpha\}_{i\in[n]})$, where P_i^α is the player's modified payoff function:

$$P_i^\alpha(s) = (1-\alpha_i)\cdot P_i(s) + \alpha_i \cdot P(s).$$

We denote this extension as $G^\alpha(P)$.

Therefore, in the altruistic extension of the game, the players take into account the social welfare of the games' outcome.

Chen et al. [7] defined a notion of smoothness in a similar manner to the one proposed in 1.1 in order to achieve bounds on the PoA in games with altruistic settings.

2.2 Games of Incomplete Information

In incomplete information games, each player has a type space T_i, an action space A_i, and a payoff function P_i. Let $T = T_1 \times \cdots \times T_n$ and $A = A_1 \times \cdots \times A_n$ be the combined action and type space. Player i's payoff function $P_i : T_i \times A \mapsto \mathbb{R}^+$ maps both her type and the combined action of the players into her payoff. We assume that the players' type vector is drawn from a commonly known prior distribution F. This distribution may be a product distribution or a joint distribution (for our results, we assume F is a product distribution).

When a player makes her move she only knows of her type, and makes a move that maximizes her payoff in expectation, over the possible types of the other players. When we talk about equilibrium in incomplete information games, we assume that each player doesn't know the other players types, but knows their bidding strategies ("Assuming my types is t_i, I will player mixed strategy $s_i(t_i)$") and the prior distribution on their types. Therefore, in incomplete game equilibrium, given a combined strategy, each player doesn't improve by changing her strategy. More formally:

Definition 2.2. Let $G = (T, A, \{P_i\}_{i\in[n]})$ be an incomplete information game. A strategy profile $\mathbf{s} = (s_1, \ldots, s_n)$ (s_i maps player i's type space T_i into a probability distribution over her action space A_i) is a ***Bayes-Nash equilibrium*** if for every player i, every type $t_i \in T_i$ and action $a'_i \in A_i$,

$$\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}} [P_i(t_i; \mathbf{a})] \geq \mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}_{-i}\sim\mathbf{s}_{-i}(\mathbf{t}_{-i})}} [P_i(t_i; (a'_i, \mathbf{a}_{-i}))]$$

where $F_{-i}^{t_i}$ denotes the distribution on the other players types conditioned on player i being of type t_i.

Given the above definition for equilibrium in incomplete information games, we can now define how to measure the inefficiency in such a game:

Definition 2.3. Let $G = (T, A, \{P_i\}_{i\in[n]})$ be an incomplete information game. Let $P : T \times A$ be a payoff dominating maximization objective, and given $\mathbf{t} \in T$, let $OPT(\mathbf{t})$ be the action profile that maximizes P for $\mathbf{t}$. Finally, let S_{BN} be the set of Bayes-Nash equilibrium profiles for G. The *price of anarchy* of the game is:

$$\min_{\mathbf{s}\in S_{BN}} \frac{\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}} [P(\mathbf{t}; \mathbf{a})]}{\mathbf{E}_{\mathbf{t}\sim F} [P(\mathbf{t}; OPT(\mathbf{t}))]}.$$

When F is a product distribution, we talk about *independent PoA* (***iPoA***). Otherwise, we talk about *correlated PoA* (***cPoA***). In our results, we'll focus on bounding the iPoA.

In a recent paper [4], Roughgarden generalized the smoothness notions to games of incomplete information (for a simplification of Roughgarden's smoothness notion, check out appendix A). We'll be interested to capture both the incomplete information settings and the altruistic settings in one unifying notion of smoothness.

3 Incomplete Information Games with Altruistic Extension

The incomplete information settings and the altruistic ones can be easily combined. If we look at a incomplete information game structure $G =$

$(T, A, \{P_i\}_{i\in[n]})$, a payoff dominating social welfare function $P : T \times A \mapsto \mathbb{R}^+$, and a vector $\alpha \in [0,1]^n$, then we can simply modify the players payoffs to:

$$P_i^\alpha(t_i;\mathbf{a}) = (1-\alpha_i)\cdot P_i(t_i;\mathbf{a}) + \alpha_i \cdot \mathbf{E}_{\mathbf{t}_{-i}\sim F_{-i}^{t_i}}\left[P(\mathbf{t};\mathbf{a})\right]. \tag{3.1}$$

Notice that we use the expectation of the social welfare instead of the social welfare itself because each player doesn't know about the other players types. Since the player knows the prior distribution F, she can use the outcome to compute the expectation of the social welfare.

3.1 Smoothness notion

We'll now introduce the smoothness notion we'll use to bound the inefficiency of this type of game. The smoothness notion will be parameterized by a choice function $\mathbf{c}^*$ in a similar manner to the one introduced by Roughgarden [4]. For simplicity of notation, we'll define $P_{-i}^{\mathbf{t}}(\mathbf{a}) = P(\mathbf{t};\mathbf{a}) - P_i(t_i;\mathbf{a})$.

Definition 3.1. Let $G = (T, A, \{P_i\}_{i\in[n]})$ be an incomplete information game, and let $G^\alpha(P)$ be its α-altruistic extension parameterized with a social welfare function $P : T\times A \mapsto \mathbb{R}^+$. $G^\alpha(P)$ is (λ,μ)-*smooth* with respect to the choice function $\mathbf{c}^* : T \mapsto A$ if

$$\sum_{i=1}^{n} P_i(t_i;(c_i^*(\mathbf{t}),\mathbf{a}_{-i})) + \alpha_i\cdot\left(P_{-i}^{\mathbf{t}}(c_i^*(\mathbf{t}),\mathbf{a}_{-i}) - P_{-i}^{\mathbf{t}}(\mathbf{a})\right) \geq \tag{3.2}$$
$$\lambda\cdot P(\mathbf{t};\mathbf{c}^*(\mathbf{t})) - \mu\cdot P(\mathbf{t};\mathbf{a})$$

for every type vector $\mathbf{t}$, and every outcome $\mathbf{a}$ feasible for $\mathbf{t}$.

Using this notion, we'll proof our main theorem:

Theorem 3.2. *Let $G^\alpha(P)$ be the α-altruistic extension of incomplete information G parameterized with a payoff dominating maximization function P. If $G^\alpha(P)$ is (λ,μ)-smooth with respect to an optimal choice function for P, then the iPoA of $G^\alpha(P)$ with respect to P is at least $\frac{\lambda}{1+\mu}$.*

3.2 Proof of the Main Theorem

First, we'll prove lemma 3.3 to help us show our main result, and afterwards we'll use the lemma to prove theorem 3.2.

Lemma 3.3. *Let $G^{\alpha}(P)$ be the α-altruistic extension of an incomplete information game G parameterized with P, let $\mathbf{s}$ be a Bayes-Nash equilibrium strategy profile, and let $\mathbf{c}^*$ be a choice function. Then for every i and every $t_i \in T_i$ we have:*

$$\begin{aligned} &\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;\mathbf{a})\right] \geq \\ &\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;(c_i^*(\mathbf{t}),\mathbf{a}_{-i})) + \alpha_i\cdot\left(P_{-i}^{\mathbf{t}}(c_i^*(\mathbf{t}),\mathbf{a}_{-i}) - P_{-i}^{\mathbf{t}}(\mathbf{a})\right)\right]. \end{aligned} \tag{3.3}$$

Proof. Let $\mathbf{s}$ be a Bayes-Nash equilibrium for $G^{\alpha}(P)$. Therefore, for every player i, for every type t_i she might have, and for any arbitrary strategy s_i^* she might play,

$$\begin{aligned} &\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[(1-\alpha_i)\cdot P_i(t_i;\mathbf{a}) + \alpha_i\cdot\mathbf{E}_{\mathbf{t}'_{-i}\sim F_{-i}^{t_i}}\left[P((t_i,\mathbf{t}'_{-i});\mathbf{a})\right]\right] \geq \\ &\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}_{-i}\sim\mathbf{s}_{-i}(\mathbf{t}_{-i})\\ a_i^*\sim s_i^*(t_i)}}\left[(1-\alpha_i)\cdot P_i(t_i;(a_i^*,\mathbf{a}_{-i})) + \alpha_i\cdot\mathbf{E}_{\mathbf{t}'_{-i}\sim F_{-i}^{t_i}}\left[P((t_i,\mathbf{t}'_{-i});(a_i^*,\mathbf{a}_{-i}))\right]\right]. \end{aligned} \tag{3.4}$$

Using a linearity of expectation argument, and the fact that $\mathbf{t}_{-i}$ and $\mathbf{t}'_{-i}$ are sampled from the same distribution, we can simplify (3.4) to get:

$$\begin{aligned} &\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;\mathbf{a})\right] \\ \geq\ &\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}_{-i}\sim\mathbf{s}_{-i}(\mathbf{t}_{-i})\\ a_i^*\sim s_i^*(t_i)}} \\ &\left[P_i(t_i;(a_i^*,\mathbf{a}_{-i})) + \alpha_i\cdot\left(P(\mathbf{t};(a_i^*,\mathbf{a}_{-i})) - P_i(t_i;(a_i^*,\mathbf{a}_{-i})) - (P(\mathbf{t};\mathbf{a}) - P_i(t_i;\mathbf{a}))\right)\right] \\ =\ &\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}_{-i}\sim\mathbf{s}_{-i}(\mathbf{t}_{-i})\\ a_i^*\sim s_i^*(t_i)}}\left[P_i(t_i;(a_i^*,\mathbf{a}_{-i})) + \alpha_i\cdot\left(P_{-i}^{\mathbf{t}}(a_i^*,\mathbf{a}_{-i}) - P_{-i}^{\mathbf{t}}(\mathbf{a})\right)\right]. \end{aligned} \tag{3.5}$$

Let s_i^* be the following strategy for player i:

1. Sample the other players types given her own type t_i. Denote this as $\mathbf{t}^*_{-i}$.
2. Play $c_i^*(t_i,\mathbf{t}^*_{-i})$.

Using again the fact that $\mathbf{t}^*_{-i}$ and $\mathbf{t}_{-i}$ are taken from the same distribution and applying to equation (3.5), we get (3.3) which completes the proof. □

Now, we can finally prove our main theorem:

Proof of theorem 3.2. Let $\mathbf{s}$ be a Bayes-Nash equilibrium for game $G^{\alpha}(P)$. Using the payoff dominance property of P, and linearity of expectation, we get:

$$\begin{aligned} \mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P(\mathbf{t};\mathbf{a})\right] &\geq \mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[\sum_{i=1}^{n} P_i(t_i;\mathbf{a})\right] \qquad (3.6)\\ &= \sum_{i=1}^{n}\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;\mathbf{a})\right]. \end{aligned}$$

Since F is a product distribution over the types, we know that

$$\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;\mathbf{a})\right] = \mathbf{E}_{t_i\sim F_{-i}}\left[\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;\mathbf{a})\right]\right] \qquad (3.7)$$

for every i. Therefore, by applying lemma 3.3 we get that for the optimal choice function $\mathbf{c}^*$ and for every i,

$$\begin{aligned} &\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;\mathbf{a})\right] \geq \qquad (3.8)\\ &\mathbf{E}_{t_i\sim F_{-i}}\left[\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;(c_i^*(\mathbf{t}),\mathbf{a}_{-i})) + \alpha_i\cdot\left(P_{-i}^{\mathbf{t}}(c_i^*(\mathbf{t}),\mathbf{a}_{-i}) - P_{-i}^{\mathbf{t}}(\mathbf{a})\right)\right]\right] =\\ &\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;(c_i^*(\mathbf{t}),\mathbf{a}_{-i})) + \alpha_i\cdot\left(P_{-i}^{\mathbf{t}}(c_i^*(\mathbf{t}),\mathbf{a}_{-i}) - P_{-i}^{\mathbf{t}}(\mathbf{a})\right)\right], \end{aligned}$$

where the last transition is again due to the fact that F is a product distribution. Plugging back to (3.6) gives us:

$$\begin{aligned} \mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P(\mathbf{t};\mathbf{a})\right] &\geq \sum_{i=1}^{n}\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;(c_i^*(\mathbf{t}),\mathbf{a}_{-i})) + \alpha_i\cdot\left(P_{-i}^{\mathbf{t}}(c_i^*(\mathbf{t}),\mathbf{a}_{-i}) - P_{-i}^{\mathbf{t}}(\mathbf{a})\right)\right]\\ &= \mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[\sum_{i=1}^{n} P_i(t_i;(c_i^*(\mathbf{t}),\mathbf{a}_{-i})) + \alpha_i\cdot\left(P_{-i}^{\mathbf{t}}(c_i^*(\mathbf{t}),\mathbf{a}_{-i}) - P_{-i}^{\mathbf{t}}(\mathbf{a})\right)\right]\\ &\geq \mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[\lambda\cdot P(\mathbf{t};\mathbf{c}^*(\mathbf{t})) - \mu\cdot P(\mathbf{t};\mathbf{a})\right]\\ &= \lambda\cdot\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P(\mathbf{t};\mathbf{c}^*(\mathbf{t}))\right] - \mu\cdot\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P(\mathbf{t};\mathbf{a})\right], \qquad (3.9) \end{aligned}$$

where the second and last transitions stem from linearity of expectation, and the third transition is due to the fact that $G^\alpha(P)$ is (λ, μ)-smooth with respect to $\mathbf{c}^*$.

Since $\mathbf{s}$ is an arbitrary Bayes-Nash equilibrium for $G^\alpha(P)$, and $\mathbf{c}^*$ is an optimal choice function (returns an action profile which maximizes P for every $\mathbf{t}$), we get from (3.9) that the iPoA of $G^\alpha(P)$ is at least $\frac{\lambda}{1+\mu}$ with respect to P as needed.

□

3.3 Extension to Cost Minimization Games

In cost minimization games, each player i has a cost function $C_i : T_i \times A \mapsto \mathbb{R}^+$ he wishes to minimize, and there exists a global *social cost* function $C : T \times A \mapsto \mathbb{R}^+$ which is cost-dominated. That is,

$$C(\mathbf{t}; \mathbf{a}) \leq \sum_{i=1}^{n} C_i(t_i; \mathbf{a}) \tag{3.10}$$

for every type vector $\mathbf{t}$ and every combined action $\mathbf{a}$.

We'll define the α*-altruistic* extension of an incomplete information parameterized with a minimization function C in an analogous way to the maximization version (3.1). Let $G = (T, A, \{C_i\}_{i\in[n]})$ be a cost minimization incomplete information game. In the extension game, each player i has a modified cost function,

$$C_i^\alpha(t_i; \mathbf{a}) = (1 - \alpha_i) \cdot C_i(t_i; \mathbf{a}) + \alpha_i \cdot \mathbf{E}_{\mathbf{t}_{-i} \sim F_{-i}^{t_i}} \left[C(\mathbf{t}; \mathbf{a})\right], \tag{3.11}$$

for $\alpha \in [0, 1]^n$.
The smoothness notion we'll use here is the following:

Definition 3.4 (Smoothness - Minimization Version)**.** Let $G = (T, A, \{C_i\}_{i\in[n]})$ be an incomplete information cost-minimization game, and let $G^\alpha(C)$ be its α-altruistic extension parameterized with a social cost function $C : T \times A \mapsto \mathbb{R}^+$. $G^\alpha(C)$ is (λ, μ)*-smooth* with respect to the choice function $\mathbf{c}^* : T \mapsto A$ if

$$\sum_{i=1}^{n} C_i(t_i; (c_i^*(\mathbf{t}), \mathbf{a}_{-i})) + \alpha_i \cdot \left(C_{-i}^{\mathbf{t}}(c_i^*(\mathbf{t}), \mathbf{a}_{-i}) - C_{-i}^{\mathbf{t}}(\mathbf{a})\right) \leq \tag{3.12}$$
$$\lambda \cdot C(\mathbf{t}; \mathbf{c}^*(\mathbf{t})) + \mu \cdot C(\mathbf{t}; \mathbf{a})$$

for every type vector $\mathbf{t}$, and every outcome $\mathbf{a}$ feasible for $\mathbf{t}$.

Like before, $C^{\mathbf{t}}_{-i}(\mathbf{a}) = C(\mathbf{t};\mathbf{a}) - C_i(t_i;\mathbf{a})$. In this case, like many other cases of smoothness theorem, we have a completely analogous proof for the extension theorem (left unproved):

Theorem 3.5 (Main Theorem - Minimization Version)**.** *Let $G^{\alpha}(C)$ be the α-altruistic extension of incomplete information G parameterized with a cost dominated minimization function C. If $G^{\alpha}(C)$ is (λ,μ)-smooth with respect to the optimal choice function for C, then the iPoA of $G^{\alpha}(C)$ with respect to P is at most $\frac{\lambda}{1-\mu}$.*

4 Issues and Future Work

One major flaw in the formulation of the payoff function of an altruistic incomplete information game as defined above (3.1) is that it doesn't take into account the other players' actions when trying to determine the types of the other players. That is, if the games in hand is the first bid auction, then if the action taken by player i is bidding 9 Shekels for the item, when another player tries to determine player i's true type, 8 Shekels or less may still be a feasible prediction.

Future work in this field may try to find a smoothness notion for a an altruistic payoff function that takes into account the other players' actions when trying to determine the players' types. It is desired to find a simple smoothness notion like the one in (3.2) that uses only fixed types (no probabilistic arguments used) in order to derive the bounds shown in theorem 3.2. This is needed because of the intricacy of the definitions of games of incomplete information and their altruistic extension, caused by the use of expectation in the definition of the equilibrium and the payoff function.

References

[1] E. Koutsoupias and C. H. Papadimitriou, "Worst-case equilibria," in *STACS*, pp. 404–413, 1999.

[2] T. Roughgarden, "Intrinsic robustness of the price of anarchy," in *STOC*, pp. 513–522, 2009.

[3] J. C. Harsanyi, "Games with incomplete information played by bayesian players, parts i, ii, and iii," *Management Science*, vol. 14, pp. 159–182, 320–334, and 486–502, 1967.

[4] T. Roughgarden, "The price of anarchy in games of incomplete information," in *ACM Conference on Electronic Commerce*, pp. 862–879, 2012.

[5] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, and E. Papaioannou, "The impact of altruism on the efficiency of atomic congestion games," in *TGC*, pp. 172–188, 2010.

[6] P.-A. Chen and D. Kempe, "Altruism, selfishness, and spite in traffic routing," in *ACM Conference on Electronic Commerce*, pp. 140–149, 2008.

[7] P.-A. Chen, B. de Keijzer, D. Kempe, and G. Schäfer, "The robust price of anarchy of altruistic games," *CoRR*, vol. abs/1112.3680, 2011.

A Simpler Smoothness Notion for Games of Incomplete Information

In [4], Roughgarden introduced the following notion of smoothness for games of incomplete information (adapted for our notions):

Definition A.1 (Roughgarden's notion)**.** Let $G = (T, A, \{P_i\}_{i\in[n]})$ be an incomplete games and P a maximization social welfare function. G is (λ, μ)*-smooth* with respect to the choice function $\mathbf{c}^* : T \mapsto A$ if

$$\sum_{i=1}^{n} P_i(t_i; (c_i^*(\mathbf{t}), \mathbf{a}_{-i})) \geq \lambda \cdot P(\mathbf{t}; \mathbf{c}^*(\mathbf{t})) - \mu \cdot P(\mathbf{t}'; \mathbf{a}) \qquad \text{(A.1)}$$

for every type vectors $\mathbf{t}$, $\mathbf{t}'$, and every action profile $\mathbf{a}$ feasible for $\mathbf{t}'$.

Using this notion, Roughgarden was able to prove price of anarchy bounds in these kind of games. We'll show that we can use the following simpler notion of smoothness instead:

Definition A.2. Let $G = (T, A, \{P_i\}_{i\in[n]})$ be an incomplete games and P a maximization social welfare function. G is (λ, μ)*-smooth* with respect to the choice function $\mathbf{c}^* : T \mapsto A$ if

$$\sum_{i=1}^{n} P_i(t_i; (c_i^*(\mathbf{t}), \mathbf{a}_{-i})) \geq \lambda \cdot P(\mathbf{t}; \mathbf{c}^*(\mathbf{t})) - \mu \cdot P(\mathbf{t}; \mathbf{a}) \tag{A.2}$$

for every type vector $\mathbf{t}$, and every action profile $\mathbf{a}$ feasible for $\mathbf{t}$.

Notice that the main difference is the use of only a single type vector $\mathbf{t}$. This is a weaker requirement then the one Roughgarden used, and should simplify the proofs of the bounds of specific games. We'll prove the following claim analogous to the one proved on [4]:

Theorem A.3. *Let G be an incomplete information game. If G is (λ, μ)-smooth with respect to an optimal choice function for a payoff dominating social welfare function P, then the iPOA of G with respect to P is at least $\frac{\lambda}{1+\mu}$.*

The proof, like the one of theorem 3.2, will be twofold. First we'll prove an helpful lemma similar to lemma 3.3, and afterwards we'll use it to prove theorem A.3.

Lemma A.4. *Let G be an incomplete information game, let $\mathbf{s}$ be a Bayes-Nash equilibrium strategy profile for G, and let $\mathbf{c}^*$ be a choice function. Then for every i, and every $t_i \in T_i$, we have:*

$$\mathbf{E}_{\substack{\mathbf{t}_{-i} \sim F_{-i}^{t_i} \\ \mathbf{a} \sim \mathbf{s}(\mathbf{t})}} \left[P_i(t_i; \mathbf{a}) \right] \geq \mathbf{E}_{\substack{\mathbf{t}_{-i} \sim F_{-i}^{t_i} \\ \mathbf{a} \sim \mathbf{s}(\mathbf{t})}} \left[P_i(t_i; (c_i^*(\mathbf{t}), \mathbf{a}_{-i})) \right]. \tag{A.3}$$

Proof. Let $\mathbf{s}$ be a Bayes-Nash equilibrium for G. Therefore, for every player i, her type t_i, and every strategy s_i^* she might adapt,

$$\mathbf{E}_{\substack{\mathbf{t}_{-i} \sim F_{-i}^{t_i} \\ \mathbf{a} \sim \mathbf{s}(\mathbf{t})}} \left[P_i(t_i; \mathbf{a}) \right] \geq \mathbf{E}_{\substack{\mathbf{t}_{-i} \sim F_{-i}^{t_i} \\ \mathbf{a}_{-i} \sim \mathbf{s}_{-i}(\mathbf{t}_{-i}) \\ a_i^* \sim s_i^*(t_i)}} \left[P_i(t_i; (a_i^*, \mathbf{a}_{-i})) \right]. \tag{A.4}$$

Let s_i^* be the following strategy for player i:

1. Sample the other players types given her own type t_i. Denote this as $\mathbf{t}^*_{-i}$.
2. Play $c_i^*(t_i, \mathbf{t}^*_{-i})$.

That is,

$$\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;\mathbf{a})\right] \geq \mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}_{-i}\sim\mathbf{s}_{-i}(\mathbf{t}_{-i})\\ \mathbf{t}^*_{-i}\sim F_{-i}^{t_i}}}\left[P_i(t_i;(c_i^*(t_i,\mathbf{t}^*_{-i}),\mathbf{a}_{-i}))\right]. \tag{A.5}$$

Since $\mathbf{t}_{-i}$ and $\mathbf{t}^*_{-i}$ are sampled from the exact same distribution, we get that (A.3) holds, which completes the proof. □

Now we are for the proof of the bound achieved using this simpler notion:

Proof of theorem A.3. Let $\mathbf{s}$ be a Bayes-Nash equilibrium for G. Using the payoff dominance property of P, and linearity of expectation, we get:

$$\begin{aligned}\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P(\mathbf{t};\mathbf{a})\right] &\geq \mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[\sum_{i=1}^{n} P_i(t_i;\mathbf{a})\right] \qquad (A.6)\\ &= \sum_{i=1}^{n}\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;\mathbf{a})\right].\end{aligned}$$

Since F is a product distribution over the types, we know that

$$\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;\mathbf{a})\right] = \mathbf{E}_{t_i\sim F_{-i}}\left[\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;\mathbf{a})\right]\right] \tag{A.7}$$

for every i. Therefore, by applying lemma A.3 we get that for the optimal choice function $\mathbf{c}^*$ and for every i,

$$\begin{aligned}\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;\mathbf{a})\right] &\geq \mathbf{E}_{t_i\sim F_{-i}}\left[\mathbf{E}_{\substack{\mathbf{t}_{-i}\sim F_{-i}^{t_i}\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;(c_i^*(\mathbf{t}),\mathbf{a}_{-i}))\right]\right]\\ &= \mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;(c_i^*(\mathbf{t}),\mathbf{a}_{-i}))\right],\end{aligned}$$

where the last transition is again due to the fact that F is a product distribution. Plugging back to (A.6) gives us:

$$\begin{aligned}\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P(\mathbf{t};\mathbf{a})\right] &\geq \sum_{i=1}^{n}\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P_i(t_i;(c_i^*(\mathbf{t}),\mathbf{a}_{-i}))\right]\\ &= \mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[\sum_{i=1}^{n}P_i(t_i;(c_i^*(\mathbf{t}),\mathbf{a}_{-i}))\right]\\ &\geq \mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[\lambda\cdot P(\mathbf{t};\mathbf{c}^*(\mathbf{t}))-\mu\cdot P(\mathbf{t};\mathbf{a})\right]\\ &= \lambda\cdot\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P(\mathbf{t};\mathbf{c}^*(\mathbf{t}))\right]-\mu\cdot\mathbf{E}_{\substack{\mathbf{t}\sim F\\ \mathbf{a}\sim\mathbf{s}(\mathbf{t})}}\left[P(\mathbf{t};\mathbf{a})\right], \qquad (A.8)\end{aligned}$$

where the second and last transitions stem from linearity of expectation, and the third transition is due to the fact that G is (λ, μ)-smooth with respect to $\mathbf{c}^*$.

Since $\mathbf{s}$ is an arbitrary Bayes-Nash equilibrium for G, and $\mathbf{c}^*$ is the optimal choice function (returns the action profile which maximizes P for every $\mathbf{t}$), we get from (A.8) that the iPoA of G is at least $\frac{\lambda}{1+\mu}$ with respect to P as wanted.

□