The Robust Price of Anarchy of Altruistic Games

Po-An Chen[1], Bart de Keijzer[2], David Kempe[1], and Guido Schäfer[2,3]

[1] Department of Computer Science, University of Southern California, USA, {poanchen, dkempe}@usc.edu
[2] Algorithms, Combinatorics and Optimization, CWI Amsterdam, The Netherlands, {b.de.keijzer, g.schaefer}@cwi.nl
[3] Dept. of Econometrics and Operations Research, VU University Amsterdam, The Netherlands

Abstract. We study the inefficiency of equilibria for several classes of games when players are (partially) altruistic. We model altruistic behavior by assuming that player i's perceived cost is a convex combination of $1-\alpha_i$ times his direct cost and α_i times the social cost. Tuning the parameters α_i allows smooth interpolation between purely selfish and purely altruistic behavior. Within this framework, we study altruistic extensions of cost-sharing games, utility games, and linear congestion games. Our main contribution is an adaptation of Roughgarden's *smoothness* notion to altruistic extensions of games. We show that this extension captures the essential properties to determine the *robust price of anarchy* of these games, and use it to derive mostly tight bounds.

1 Introduction

Many large-scale decentralized systems involve the interactions of large numbers of individuals acting to benefit themselves. Thus, such systems are naturally studied from the viewpoint of game theory, with an eye on the social efficiency of stable outcomes. Traditionally, "stable outcomes" have been associated with pure Nash equilibria of the corresponding game. The notions of *price of anarchy* [9] and *price of stability* [2] provide natural measures of the system degradation, by capturing the degradation of the worst and best Nash equilibria, respectively, compared to the socially optimal outcome. However, the predictive power of such bounds has been questioned on (at least) two grounds: First, the adoption of Nash equilibria as a prescriptive solution concept implicitly assumes that players are able to reach such equilibria, a very suspect assumption for computationally bounded players. In response, recent work has begun analyzing the outcomes of natural response dynamics [3,15], as well as more permissive solution concepts such as mixed, correlated or coarse correlated equilibria. (This general direction of inquiry has become known as "robust price of anarchy".) Second, the assumption that players seek only to maximize their own utility is at odds with altruistic behavior routinely observed in the real world. While modeling human incentives and behavior accurately is a formidable task, several papers have proposed natural models of altruism and analyzed its impact on the outcomes of games [4,5,6,10].

The goal of this paper 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. We formally define the *altruistic extension* of an n-player game in the spirit of past work on altruism (see [10, p. 154] and [4,5,8]): player i has an associated altruism parameter α_i, and his cost (or payoff) is a convex combination of $(1-\alpha_i)$ times his direct cost (or payoff) and α_i times the social cost (or social welfare). By tuning the parameters α_i, this model allows smooth interpolation between pure selfishness ($\alpha_i = 0$) and pure altruism ($\alpha_i = 1$). To analyze the degradation of system performance in light of partially altruistic behavior, we extend the notion of *robust price of anarchy* [15] to games with altruistic players, and show that a suitably adapted notion of *smoothness* [15] captures the properties of a system that determine its robust price of anarchy. We use our framework to analyze the robust price of anarchy of three fundamental classes of games.

1. In a *cost-sharing game* [2], players choose subsets of resources, and all players choosing the same resource share its cost evenly. Using our framework, we derive a bound of $n/(1-\hat{\alpha})$ on the robust price of anarchy of these games, where $\hat{\alpha}$ is the maximum altruism level of a player. This bound is tight for uniformly altruistic players.

2. We apply our framework to *utility games* [16], in which players choose subsets of resources and derive utility of the chosen set. The total welfare is determined by a submodular function of the union of all chosen sets. We derive a bound of 2 on the robust price of anarchy of these games. In particular, the bound remains at 2 regardless of the (possibly different) altruism levels of the players. This bound is tight.

3. We revisit and extend the analysis of *atomic congestion games* [14], in which players choose subsets of resources whose costs increase (linearly) with the number of players using them. Caragiannis et al. [4] recently derived a tight bound of $(5+4\alpha)/(2+\alpha)$ on the pure price of anarchy when all players have the *same* altruism level α.[4] Our framework makes it an easy observation that their proof in fact bounds the robust price of anarchy. We generalize their bound to the case when different players have different altruism levels, obtaining a bound in terms of the maximum and minimum altruism levels. This partially answers an open question from [4]. For the special case of symmetric singleton congestion games (which corresponds to selfish scheduling on machines), we extend our study of non-uniform altruism and obtain an improved bound of $(4-2\alpha)/(3-\alpha)$ on the price of anarchy when an α-fraction of the players are entirely altruistic and the remaining players are entirely selfish.

Notice that many of these bounds on the robust price of anarchy reveal a counter-intuitive trend: at best, for utility games, the bound is independent

[4] The altruism model of [4] differs from ours in a slight technicality discussed in Section 2 (Remark 1). Therefore, various bounds we cite here are stated differently in [4].

of the level of altruism, and for congestion games and cost-sharing games, it actually *increases* in the altruism level, unboundedly so for cost-sharing games. Intuitively, this phenomenon is explained by the fact that a change of strategy by player i may affect many players. An altruistic player will care more about these other players than a selfish player; hence, an altruistic player accepts more states as "stable". This suggests that the best stable solution can also be chosen from a larger set, and the price of stability should thus decrease. Our results on the price of stability lend support to this intuition: for congestion games, we derive an upper bound on the price of stability which decreases as $2/(1+\alpha)$; similarly, for cost-sharing games, we establish an upper bound which decreases as $(1-\alpha)H_n + \alpha$.

The increase in the price of anarchy is not a universal phenomenon, demonstrated by *symmetric singleton* congestion games. Caragiannis et al. [4] showed a bound of $4/(3+\alpha)$ for pure Nash equilibria with uniformly altruistic players, which decreases with the altruism level α. Our bound of $(4-2\alpha)/(3-\alpha)$ for mixtures of entirely altruistic and selfish players is also decreasing in the fraction of entirely altruistic players. We also extend an example of Lücking et al. [11] to show that symmetric singleton congestion games may have a mixed price of anarchy arbitrarily close to 2 for arbitrary altruism levels. In light of the above bounds, this establishes that pure Nash equilibria can result in strictly lower price of anarchy than weaker solution concepts.

Most proofs are omitted from this short paper; they are available in the full version.

2 Altruistic Games and the Robust Price of Anarchy

Let $G = (N, \{\Sigma_i\}_{i\in N}, \{C_i\}_{i\in N})$ be a finite strategic game, where $N = [n]$ is the set of players, Σ_i the strategy space of player i, and $C_i : \Sigma \to \mathbb{R}$ the cost function of player i, mapping every joint strategy $s \in \Sigma = \Sigma_1 \times \cdots \times \Sigma_n$ to the player's direct cost. Unless stated otherwise, we assume that every player i wants to minimize his individual cost function C_i. We also call such games *cost-minimization games*. A *social cost* function $C : \Sigma \to \mathbb{R}$ maps strategies to social costs. We require that C is *sum-bounded*, that is, $C(s) \leq \sum_{i=1}^{n} C_i(s)$ for all $s \in \Sigma$. We study *altruistic extensions* of strategic games equipped with sum-bounded social cost functions. Our definition is based on one used (among others) in [5], and similar to ones given in [4,6,10].

Definition 1. *Let $\alpha \in [0,1]^n$. The* α-altruistic extension *of G (or simply α-*altruistic game*) is defined as the strategic game $G^\alpha = (N, \{\Sigma_i\}_{i\in N}, \{C_i^\alpha\}_{i\in N})$, where for every $i \in N$ and $s \in \Sigma$, $C_i^\alpha(s) = (1-\alpha_i)C_i(s) + \alpha_i C(s)$.*

Thus, the perceived cost that player i experiences is a convex combination of his direct (selfish) cost and the social cost; we call such a player α_i-*altruistic*. When $\alpha_i = 0$, player i is entirely selfish; thus, $\alpha = \mathbf{0}$ recovers the original game. A player with $\alpha_i = 1$ is entirely altruistic. Given an altruism vector $\alpha \in [0,1]^n$, we let $\hat{\alpha} = \max_{i\in N} \alpha_i$ and $\check{\alpha} = \min_{i\in N} \alpha_i$ denote the maximum and minimum altruism levels, respectively. When $\alpha_i = \alpha$ (a scalar) for all i, we call such games *uniformly α-altruistic games.*

Remark 1. In a recent paper, Caragiannis et al. [4] model uniformly altruistic players by defining the perceived cost of player i as $(1-\xi)C_i(s)+\xi(C(s)-C_i(s))$, where $\xi \in [0,1]$. It is not hard to see that in the range $\xi \in [0, \frac{1}{2}]$ this definition is equivalent to ours by setting $\alpha = \xi/(1-\xi)$ or $\xi = \alpha/(1+\alpha)$.

The most general equilibrium concept we consider is coarse (correlated) equilibria.

Definition 2 (Coarse equilibrium). *A* coarse (correlated) equilibrium *of a game G is a probability distribution σ over $\Sigma = \Sigma_1 \times \cdots \times \Sigma_n$ with the following property: if s is a random variable with distribution σ, then for each player i, and all $s_i^* \in \Sigma_i$:*

$$\mathbf{E}_{s\sim\sigma}\left[C_i(s)\right] \leq \mathbf{E}_{s_{-i}\sim\sigma_{-i}}\left[C_i(s_i^*, s_{-i})\right], \tag{1}$$

where σ_{-i} is the projection of σ on $\Sigma_{-i} = \Sigma_1 \times \cdots \times \Sigma_{i-1} \times \Sigma_{i+1} \times \cdots \times \Sigma_n$.

It includes several other solution concepts, such as correlated equilibria, mixed Nash equilibria and pure Nash equilibria.

The *price of anarchy (PoA)* [9] and *price of stability (PoS)* [2] quantify the inefficiency of equilibria for classes of games: Let $S \subseteq \Sigma$ be a set of strategy profiles for a cost-minimization game G with social cost function C, and let s^* be a strategy profile that minimizes C. We define $\text{PoA}(S,G) = \sup\{C(s)/C(s^*) : s \in S\}$ and $\text{PoS}(S,G) = \inf\{C(s)/C(s^*) : s \in S\}$. The *coarse* (or *correlated, mixed, pure*) *PoA* (or *PoS*) of a class of games $\mathcal{G}$ is the supremum over all games in $\mathcal{G}$ and all strategy profiles in the respective set of equilibrium outcomes. Notice that the PoA and PoS are defined with respect to the *original* social cost function C, not accounting for the altruistic components. This reflects our desire to understand the overall performance of the system (or strategic game), which is not affected by different *perceptions* of costs by individuals.[5]

Roughgarden [15] introduced the notion of (λ, μ)-*smoothness* of strategic games with sum-bounded social cost functions and showed that it provides a generic template for proving bounds on the PoA as well as the outcomes of no-regret sequences [3].

The smoothness approach cannot be applied directly to our altruistic games because the social cost function C that we consider here is in general not sum-bounded in terms of C_i^α (which is a crucial prerequisite in [15]). However, we are able to generalize the (λ, μ)-smoothness notion to altruistic games, thereby preserving many of its applications. For notational convenience, we define $C_{-i}(s) = C(s) - C_i(s)$.

[5] If all players have a uniform altruism level $\alpha_i = \alpha \in [0,1]$ and the social cost function C is equal to the sum of all players' direct costs, then for every strategy profile $s \in \Sigma$, the sum of the perceived costs of all players is equal to $(1-\alpha+\alpha n)C(s)$. In particular, bounding the PoA (or PoS) with respect to C is equivalent to bounding the PoA (or PoS) with respect to total perceived cost in this case.

Definition 3. *G^α is* (λ, μ, α)-smooth *iff for any two strategy profiles $s, s^* \in \Sigma$,*

$$\sum_{i=1}^{n} C_i(s_i^*, s_{-i}) + \alpha_i(C_{-i}(s_i^*, s_{-i}) - C_{-i}(s)) \leq \lambda C(s^*) + \mu C(s).$$

Most of the results in [15] following from (λ, μ)-smoothness carry over to our altruistic setting using the generalized (λ, μ, α)-smoothness notion. The following result allows a calculation of the PoA.[6]

Proposition 1. *Let G^α be an α-altruistic game that is (λ, μ, α)-smooth with $\mu < 1$. Then, the coarse (and thus also correlated, mixed, and pure) price of anarchy of G^α is at most $\frac{\lambda}{1-\mu}$.*

For many important classes of games, the bounds obtained by (λ, μ, α)-smoothness arguments are actually tight, even for pure Nash equilibria. This motivates defining the *robust PoA* as the best bound that can be proved using the smoothness technique.

Definition 4. *Let G^α be an α-altruistic game. Its* robust PoA *is defined as $RPoA_G(\alpha) = \inf\{\frac{\lambda}{1-\mu} : G^\alpha$ is (λ, μ, α)-smooth with $\mu < 1\}$. For a class $\mathcal{G}$ of games, we define $RPoA_{\mathcal{G}}(\alpha) = \sup\{RPoA_G(\alpha) : G \in \mathcal{G}\}$.*

We study the robust PoA of three classes of games: they are all described by a set E of *resources* (or *facilities*), and strategy sets $\Sigma_i \subseteq 2^E$ for each player, from which the player can choose a subset $s_i \in \Sigma_i$ of resources. Given a joint strategy s, we define $x_e(s) = |\{i \in N : e \in s_i\}|$ as the number of players that use resource $e \in E$ under s. We also use $U(s)$ to refer to the union of all resources used under s, i.e., $U(s) = \bigcup_{i \in N} s_i$.

3 Cost-sharing Games

A *cost-sharing game* is given by $G = (N, E, \{\Sigma_i\}_{i \in N}, \{c_e\}_{e \in E})$, where c_e is the non-negative cost of facility $e \in E$. The cost of each facility is shared evenly among all players using it, i.e., the direct cost of player i is defined as $C_i(s) = \sum_{e \in s_i} c_e / x_e(s)$. The social cost function is $C(s) = \sum_{i=1}^{n} C_i(s) = \sum_{e \in U(s)} c_e$.

It is well-known that the pure PoA of cost-sharing games is n [13]. We show that it can get significantly worse when there is altruism. Also we provide an upper bound on the pure PoS when altruism is uniform.

Theorem 1. *For α-altruistic cost-sharing games, the robust PoA is $\frac{n}{1-\hat{\alpha}}$ (where $n/0 = \infty$), and for uniform altruism, the pure PoS is at most $(1-\alpha)H_n + \alpha$.*

[6] All results in this section continue to hold for altruistic extensions of *payoff-maximization games G*: One needs only replace C by Π and μ by $-\mu$ in Definition 3, and replace $\frac{\lambda}{1-\mu}$ by $\frac{1+\mu}{\lambda}$ and $\mu < 1$ by $\mu > -1$ in Definition 4.

4 Utility Games

A *utility game* [16] $G = (N, E, \{\Sigma_i\}_{i \in N}, \{\Pi_i\}_{i \in N}, V)$ is a payoff maximization game, in which Π_i is the payoff function of player i, and V is a submodular[7] and non-negative function on E. Every player i strives to maximize his payoff function Π_i. The social welfare function $\Pi : \Sigma \to \mathbb{R}$ to be maximized is $\Pi(s) = V(U(s))$, and thus depends on the union of the players' chosen resources, evaluated by V. The payoff function of every player i is assumed to satisfy[8] $\Pi_i(s) \geq \Pi(s) - \Pi(\emptyset, s_{-i})$ for every strategy profile $s \in \Sigma$. Intuitively, this means that the payoff of a player is at least his contribution to the social welfare. Moreover, it is assumed that $\Pi(s) \geq \sum_{i=1}^{n} \Pi_i(s)$ for every $s \in \Sigma$; see [16] for a justification of these assumptions. Vetta [16] proved a bound of 2 on the pure PoA for utility games with non-decreasing V; Roughgarden [15] showed that this bound is achieved via a (λ, μ)-smoothness argument. We extend it to altruistic extensions of these games.

Theorem 2. *The robust PoA of α-altruistic utility games is* 2.

5 Congestion Games

In an *atomic congestion game* $G = (N, E, \{\Sigma_i\}_{i \in N}, \{d_e\}_{e \in E})$, every facility $e \in E$ has an associated *delay function* $d_e : \mathbb{N} \to \mathbb{R}$. Player i's cost is $C_i(s) = \sum_{e \in s_i} d_e(x_e(s))$, and the social cost is $C(s) = \sum_{i=1}^{n} C_i(s)$. We focus on *linear* congestion games, i.e., the delay functions are of the form $d_e(x) = a_e x + b_e$, where a_e, b_e are non-negative rational numbers. Pure Nash equilibria of altruistic extensions of linear congestion games always exist [8]; this may not be the case for arbitrary (non-linear) congestion games.[9]

The PoA of linear congestion games is known to be $\frac{5}{2}$ [7]. Recently, Caragiannis et al. [4] extended this result to linear congestion games with uniformly altruistic players. Applying the transformation outlined in Remark 1, their result can be stated as follows:

Theorem 3 ([4]). *The pure PoA of uniformly α-altruistic linear congestion games is at most* $\frac{5+4\alpha}{2+\alpha}$.

The proof in [4] implicitly uses a smoothness argument in the framework we define here for altruistic games. Thus, without any additional work, our framework allows the extension of Theorem 3 to the robust PoA. Caragiannis

[7] A function $f : 2^E \to \mathbb{R}$ is called *submodular* iff $f(A \cup \{x\}) - f(A) \geq f(B \cup \{x\}) - f(B)$ for any $A \subseteq B \subseteq E$, $x \in E$.

[8] We abuse notation and write $\Pi(\emptyset, s_{-i})$ to denote $V(U(s) \setminus s_i)$.

[9] If players' altruism levels are not uniform, then the existence of pure Nash equilibria is not obvious. Hoefer and Skopalik [8] established it for several subclasses of atomic congestion games. For the generalization of arbitrary player-specific cost functions, Milchtaich [12] showed existence for (symmetric) singleton congestion games and Ackermann et al. [1] for matroid congestion games.

et al. [4] also showed that the bound of Theorem 3 is asymptotically tight. A simpler example (deferred to the full version of this paper) proves tightness of this bound (not only asymptotically). Thus, the robust PoA is exactly $\frac{5+4\alpha}{2+\alpha}$. We give a refinement of Theorem 3 to non-uniform altruism distributions, obtaining a bound in terms of the maximum and minimum altruism levels.

Theorem 4. *The robust PoA of α-altruistic linear congestion games is at most $\frac{5+2\hat{\alpha}+2\check{\alpha}}{2-\hat{\alpha}+2\check{\alpha}}$.*

We turn to the pure price of stability of α-altruistic congestion games. Clearly, an upper bound on the pure price of stability extends to the mixed, correlated and coarse price of stability.

Proposition 2. *The pure PoS of uniformly α-altruistic linear congestion games is at most $\frac{2}{1+\alpha}$.*

Symmetric Singleton Congestion Games. An important special case of congestion games is that of *symmetric singleton congestion games* $G = (N, E, \{\Sigma_i\}_{i\in N}, \{d_e\}_{e\in E})$, in which every player chooses one facility (also called *edge*) from $E = [m]$, and all strategy sets are identical, i.e., $\Sigma_i = E$ for every i. In *singleton linear congestion games*, the focus here, delay functions are also assumed to be linear, of the form $d_e(x) = a_e x + b_e$.

Caragiannis et al. [4] prove the following theorem (stated using the transformation from Remark 1). It shows that the pure PoA does not always increase with the altruism level; the relationship between α and the PoA is thus rather subtle.

Theorem 5 (Caragiannis et al. [4]). *The pure PoA of uniformly α-altruistic singleton linear congestion games is $\frac{4}{3+\alpha}$.*

We show that even the mixed PoA (and thus also the robust PoA) will be at least 2 regardless of the altruism levels of the players, by generalizing a result of Lücking et al. [11, Theorem 5.4]. This implies that the benefits of higher altruism in singleton congestion games are only reaped in pure Nash equilibria, and the gap between the pure and mixed PoA increases in α.

Proposition 3. *For every $\alpha \in [0,1]^n$, the mixed PoA for α-altruistic singleton linear congestion games is at least 2.*

As a first step to extend the analysis to non-uniform altruism, we analyze the case when all altruism levels are in $\{0,1\}$, i.e., each player is either completely altruistic or completely selfish.[10] Then, the system is entirely characterized by the fraction α of altruistic players (which coincides with the average altruism level). The next theorem shows that in this case, too, the pure PoA *improves* with the overall altruism level.

Theorem 6. *Assume that an α fraction of the players are completely altruistic, and the remaining $(1-\alpha)$ fraction are completely selfish. Then, the pure PoA of the altruistic singleton linear congestion game is at most $\frac{4-2\alpha}{3-\alpha}$.*

[10] This model relates naturally to *Stackelberg scheduling games* (see, e.g., [6]).

Acknowledgements. We thank anonymous reviewers for useful feedback, and for pointing out the similarities between our work on congestion games and that of [4].

References

1. H. Ackermann, H. Röglin, and B. Vöcking. Pure Nash equilibria in player-specific and weighted congestion games. In *Proc. 2th Workshop on Internet and Network Economics)*, 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 Symposium on Foundations of Computer Science*, 2004.
3. 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*, pages 373–382, 2008.
4. I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, and E. Papaioannou. The impact of altruism on the efficiency of atomic congestion games. In *Proc. 5th Symposium on Trustworthy Global Computing*, 2010.
5. P.-A. Chen, M. David, and D. Kempe. Better vaccination strategies for better people. In *Proc. 12th ACM Conference on Electronic Commerce*, 2010.
6. P.-A. Chen and D. Kempe. Altruism, selfishness, and spite in traffic routing. In *Proc. 10th ACM Conference on Electronic Commerce*, pages 140–149, 2008.
7. G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *In Proc. 37th Annual ACM Symposium on Theory of Computing*, 2005.
8. M. Hoefer and A. Skopalik. Altruism in atomic congestion games. In *Proc. 17th European Symposium on Algorithms*, 2009.
9. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proc. 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 4–6 1999.
10. 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.
11. T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. *Theoretical Computer Science*, 406:187–206, October 2008.
12. I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996.
13. N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
14. T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
15. T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proc. 41st Annual ACM Symposium on Theory of Computing*, pages 513–522, 2009.
16. A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proc. 43rd Symposium on Foundations of Computer Science*, 2002.