Selfishness Level of Strategic Games*

Krzysztof R. Apt[1,2] and Guido Schäfer[1,3]

[1] Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
[2] University of Amsterdam, The Netherlands
[3] VU University Amsterdam, The Netherlands

Abstract. We introduce a new measure of the discrepancy in strategic games between the social welfare in a Nash equilibrium and in a social optimum, that we call ***selfishness level***. It is the smallest fraction of the social welfare that needs to be offered to each player to achieve that a social optimum is realized in a pure Nash equilibrium. The selfishness level is unrelated to the price of stability and the price of anarchy and in contrast to these notions is invariant under positive linear transformations of the payoff functions. Also, it naturally applies to other solution concepts and other forms of games.

We study the selfishness level of several well-known strategic games. This allows us to quantify the implicit tension within a game between players' individual interests and the impact of their decisions on the society as a whole. Our analysis reveals that the selfishness level often provides more refined insights into the game than other measures of inefficiency, such as the price of stability or the price of anarchy.

1 Introduction

The discrepancy in strategic games between the social welfare in a Nash equilibrium and in a social optimum has been long recognized by the economists. One of the flagship examples is Cournot competition, a strategic game involving firms that simultaneously choose the production levels of a homogeneous product. The payoff functions in this game describe the firms' profit in the presence of some production costs, under the assumption that the price of the product depends negatively on the total output. It is well-known, see, e.g., [1, pages 174–175], that the price in the social optimum is strictly higher than in the Nash equilibrium, which shows that the competition between the producers of a product drives its price down.

In computer science the above discrepancy led to the introduction of the notions of the *price of anarchy* [2] and the *price of stability* [3] that measure the ratio between the social welfare in a worst and, respectively, a best Nash equilibrium and a social optimum. This originated a huge research effort aiming at determining both ratios for specific strategic games that possess (pure) Nash equilibria.

* A full version with all proofs is available at the authors' homepages.

These two notions are *descriptive* in the sense that they refer to an existing situation. In contrast, we propose a notion that measures the discrepancy between the social welfare in a Nash equilibrium and a social optimum, which is *normative*, in the sense that it refers to a modified situation. On an abstract level, the approach that we propose here is discussed in [4], in chapter "How to Promote Cooperation", from where we cite (see page 134): "An excellent way to promote cooperation in a society is to teach people to care about the welfare of others."

Our approach draws on the concept of *altruistic games* (see, e.g., [5] and more recent [6]). In these games each player's payoff is modified by adding a positive fraction α of the social welfare in the considered joint strategy to the original payoff. The ***selfishness level*** of a game is defined as the infimum over all $\alpha \geq 0$ for which such a modification yields that a social optimum is realized in a pure Nash equilibrium.

Intuitively, the selfishness level of a game can be viewed as a measure of the players' willingness to cooperate. A low selfishness level indicates that the players are open to align their interests in the sense that a small share of the social welfare is sufficient to motivate them to choose a social optimum. In contrast, a high selfishness level suggests that the players are reluctant to cooperate and a large share of the social welfare is needed to stimulate cooperation among them. An infinite selfishness level means that cooperation cannot be achieved through such means.

Often the selfishness level of a strategic game provides better insights into the game under consideration than other measures of inefficiency, such as the price of stability or the price of anarchy. To illustrate this point, we elaborate on our findings for the public goods game with n players. In this game, every player i chooses an amount $s_i \in [0, b]$ that he wants to contribute to a public good. A central authority collects all individual contributions, multiplies their sum by $c > 1$ (here we assume for simplicity that $n \geq c$) and distributes the resulting amount evenly among all players. The payoff of player i is thus $p_i(s) := b - s_i + \frac{c}{n} \sum_j s_j$.

In the (unique) Nash equilibrium, every player attempts to "free ride" by contributing 0 to the public good (which is a dominant strategy), while in the social optimum every player contributes the full amount of b. As we will show, the selfishness level of this game is $(1 - \frac{c}{n})/(c - 1)$. This bound suggests that the temptation to free ride (i) increases as the number of players grows and (ii) decreases as the parameter c increases. Both phenomena were observed by experimental economists, see, e.g., [5, Section III.C.2]. In contrast, the price of stability (which coincides with the price of anarchy) for this game is c, which is rather uninformative.

In this paper, we define the selfishness level by taking pure Nash equilibrium as the solution concept. This is in line with how the price of anarchy and price of stability were defined originally [2, 3]. However, the definition applies equally well to other solution concepts and other forms of games.

Our Contributions. In this paper, we study the selfishness level of some selected classical and fundamental strategic games. These games are often used to illustrate the consequences of selfish behaviour and the effects of competition. To this aim, we first derive a characterization result that allows us to determine the selfishness level of a strategic game. Our characterization shows that the selfishness level is determined by the maximum *appeal factor* of unilateral profitable deviations from specific social optima, which we call *stable*. Intuitively, the appeal factor of a single player deviation refers to the ratio of the gain in his payoff over the resulting loss in social welfare.

We show that the selfishness level of a finite game can be an arbitrary real number that is unrelated to the price of stability. A nice property of our selfishness level notion is that, unlike the price of stability and the price of anarchy, it is invariant under positive linear transformations of the payoff functions.

We then use the above characterization result to analyze the selfishness level of several strategic games. In particular, we show that the selfishness level of finite ordinal potential games is finite. We also derive explicit bounds on the selfishness level of fair cost sharing games and congestion games with linear delay functions. These bounds depend on the specific parameters of the underlying game, but are independent of the number of players. Moreover, our bounds are tight.

Further, we show that the selfishness level of the Prisoner's Dilemma with n players is $1/(2n-3)$ and that of the public goods game with n players is $\max\{0, (1-\frac{c}{n})(c-1)\}$. Finally, the selfishness level of Cournot competition (an example of an infinite ordinal potential game), Tragedy of the Commons, and Bertrand competition turns out to be infinite.

Related Work. There are only few articles in the algorithmic game theory literature that study the influence of altruism in strategic games [7–11]. In these works, altruistic player behavior is modeled by altering each player's perceived payoff in order to account also for the welfare of others. The models differ in the way they combine the player's individual payoff with the payoffs of the other players. All these studies are descriptive in the sense that they aim at understanding the impact of altruistic behavior on specific strategic games.

Closest to our work are the articles [10] and [8]. Elias et al. [10] study the inefficiency of equilibria in network design games (which constitute a special case of the cost sharing games considered here) with altruistic (or, as they call it, socially-aware) players. As we do here, they define each player's cost function as his individual cost plus α times the social cost. They derive lower and upper bounds on the price of anarchy and the price of stability, respectively, of the modified game. In particular, they show that the price of stability is at most $(H_n + \alpha)/(1+\alpha)$, where n is the number of players.

Chen et al. [8] introduce a framework to study the *robust price of anarchy*, which refers to the worst-case inefficiency of other solution concepts such as coarse correlated equilibria (see [12]) of altruistic extensions of strategic games. In their model, player i's perceived cost is a convex combination of $(1-\bar{\alpha}_i)$ times his individual cost plus $\bar{\alpha}_i$ times the social cost, where $\bar{\alpha}_i \in [0,1]$ is the altruism level of i. If all players have a uniform altruism level $\bar{\alpha}_i = \bar{\alpha}$, this model relates to the one we consider here by setting $\alpha = \bar{\alpha}/(1-\bar{\alpha})$ for $\bar{\alpha} \in [0,1)$. Although not being the main focus of the paper, the authors also provide upper bounds of $2/(1+\bar{\alpha})$ and $(1-\bar{\alpha})H_n + \bar{\alpha}$ on the price of stability for linear congestion games and fair cost sharing games, respectively.

Note that in all three cases the price of stability approaches 1 as α goes to ∞. This seems to suggest that the selfishness level of these games is ∞. However, this is not the case as outlined above.

Other models of altruism were proposed in [7, 9]. Chen and Kempe [9] define the perceived cost of a player as $(1-\beta)$ times his individual cost plus β/n times the social cost, where $\beta \in [0,1]$. Caragiannis et al. [7] define the perceived cost of player i as $(1-\xi)$ times his individual cost plus ξ times the sum of the costs of all other players (i.e., excluding player i), where $\xi \in [0,1]$. Both models are equivalent to the model the we consider here by using the transformations $\alpha = \beta/((1-\beta)n)$ for $\beta \in [0,1)$ and $\alpha = \xi/(1-2\xi)$ for $\xi \in [0, \frac{1}{2})$.

In network congestion games, researchers studied the effect of imposing tolls on the edges of the network in order to reduce the inefficiency of Nash equilibria; see, e.g., [13]. From a high-level perspective, these approaches can also be regarded as being normative. Conceptually, our selfishness level notion is related to the *Stackelberg threshold* introduced by Sharma and Williamson [14]. The authors consider network routing games in which a fraction $\beta \in [0,1]$ of the flow is first routed centrally and the remaining flow is then routed selfishly. The Stackelberg threshold refers to the smallest value β that is needed to improve upon the social cost of a Nash equilibrium flow.

2 Selfishness Level

A ***strategic game*** (in short, a game) $G = (N, \{S_i\}_{i\in N}, \{p_i\}_{i\in N})$ is given by a set $N = \{1,\dots,n\}$ of $n > 1$ players, a non-empty set of ***strategies*** S_i for every player $i \in N$, and a ***payoff function*** p_i for every player $i \in N$ with $p_i : S_1 \times \cdots \times S_n \to \mathbb{R}$. The players choose their strategies simultaneously and every player $i \in N$ aims at choosing a strategy $s_i \in S_i$ so as to maximize his individual payoff $p_i(s)$, where $s = (s_1,\dots,s_n)$.

We call $s \in S_1 \times \cdots \times S_n$ a ***joint strategy***, denote its ith element by s_i, denote $(s_1,\dots,s_{i-1},s_{i+1},\dots,s_n)$ by s_{-i} and similarly with S_{-i}. Further, we write (s'_i, s_{-i}) for $(s_1,\dots,s_{i-1},s'_i,s_{i+1},\dots,s_n)$, where we assume that $s'_i \in S_i$. Sometimes, when focussing on player i we write (s_i, s_{-i}) instead of s.

A joint strategy s a ***Nash equilibrium*** if for all $i \in \{1,\dots,n\}$ and $s'_i \in S_i$, $p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i})$. Further, given a joint strategy s we call the sum $SW(s) := \sum_{i=1}^{n} p_i(s)$ the ***social welfare*** of s. When the social welfare of s is maximal we call s a ***social optimum***.

Given a strategic game $G := (N, \{S_i\}_{i\in N}, \{p_i\}_{i\in N})$ and $\alpha \geq 0$ we define the game $G(\alpha) := (N, \{S_i\}_{i\in N}, \{r_i\}_{i\in N})$ by putting $r_i(s) := p_i(s) + \alpha SW(s)$. So when $\alpha > 0$ the payoff of each player in the $G(\alpha)$ game depends on the social welfare of the players. $G(\alpha)$ is then an altruistic version of the game G.

Suppose now that for some $\alpha \geq 0$ a pure Nash equilibrium of $G(\alpha)$ is a social optimum of $G(\alpha)$. Then we say that G is ***α-selfish***. We define the ***selfishness level*** of G as

$$\inf\{\alpha \in \mathbb{R}_+ \mid G \text{ is } \alpha\text{-selfish}\}. \tag{1}$$

Here we adopt the convention that the infimum of an empty set is ∞. Further, we stipulate that the selfishness level of G is denoted by α^+ iff the selfishness level of G is $\alpha \in \mathbb{R}_+$ but G is *not* α-selfish (equivalently, the infimum does not belong to the set). We show below (Theorem 2) that pathological infinite games exist for which the selfishness level is of this kind; none of the other studied games is of this type.

The above definitions refer to strategic games in which each player i maximizes his payoff function p_i and the social welfare of a joint strategy s is given by $SW(s)$. These definitions obviously apply to strategic games in which every player i minimizes his cost function c_i and the social cost of a joint strategy s is defined as $SC(s) := \sum_{i=1}^{n} c_i(s)$. The definition also extends in the obvious way to other solution concepts (e.g., mixed or correlated equilibria) and other forms of games (e.g., subgame perfect equilibria in extensive games).

Note that the social welfare of a joint strategy s in $G(\alpha)$ equals $(1+\alpha n)SW(s)$, so the social optima of G and $G(\alpha)$ coincide. Hence we can replace in the above definition the reference to a social optimum of $G(\alpha)$ by one to a social optimum of G.

Intuitively, a low selfishness level means that the share of the social welfare needed to induce the players to choose a social optimum is small. This share can be viewed as an 'incentive' needed to realize a social optimum. Let us illustrate this definition on three simple examples.

Example 1. **Prisoner's Dilemma**

	C	D
C	$2, 2$	$0, 3$
D	$3, 0$	$1, 1$

	C	D
C	$6, 6$	$3, 6$
D	$6, 3$	$3, 3$

Consider the Prisoner's Dilemma game G (on the left) and the resulting game $G(\alpha)$ for $\alpha = 1$ (on the right). In the latter game the social optimum, (C, C), is also a Nash equilibrium. One can easily check that for $\alpha < 1$, (C, C) is also a social optimum of $G(\alpha)$ but not a Nash equilibrium. So the selfishness level of this game is 1.

Example 2. **Battle of the Sexes**

	F	B
F	$2, 1$	$0, 0$
B	$0, 0$	$1, 2$

Here each Nash equilibrium is also a social optimum, so the selfishness level of this game is 0.

Example 3. **Game with a bad Nash equilibrium**
The following game results from equipping each player in the Matching Pennies game with a third strategy E (for edge):

	H	T	E
H	$1,-1$	$-1,\ \ 1$	$-1,-1$
T	$-1,\ \ 1$	$1,-1$	$-1,-1$
E	$-1,-1$	$-1,-1$	$-1,-1$

Its unique Nash equilibrium is (E, E). It is easy to check that the selfishness level of this game is ∞.

Recall that, given a finite game G that has a Nash equilibrium, its ***price of stability*** is the ratio $SW(s)/SW(s')$ where s is a social optimum and s' is a Nash equilibrium with the highest social welfare in G. So the price of stability of G is 1 iff its selfishness level is 0. However, in general there is no relation between these two notions.

Theorem 1. *For every finite $\alpha > 0$ and $\beta > 1$ there is a finite game whose selfishness level is α and whose price of stability is β.*

Further, in contrast to the price of stability (and to the ***price of anarchy***, defined as the ratio $SW(s)/SW(s')$ where s is a social optimum and s' is a Nash equilibrium with the lowest social welfare in G) the notion of the selfishness level is invariant under simple uniform payoff transformations. Given a game G and a value a we denote by $G + a$ (respectively, aG) the game obtained from G by adding to each payoff function the value a (respectively, by multiplying each payoff function by a).

Proposition 1. *Consider a game G and $\alpha \geq 0$.*

1. *For every a, G is α-selfish iff $G + a$ is α-selfish,*
2. *For every $a > 0$, G is α-selfish iff aG is α-selfish.*

This result allows us to better frame the notion of selfishness level. Namely, suppose that the original n-players game G was set up by a designer who has a fixed budget $SW(s)$ for each joint strategy s and that the selfishness level of G is $\alpha < \infty$. Then we should scale $G(\alpha)$ by the factor $a := 1/(1 + \alpha n)$ so that for each joint strategy s its social welfare in the original game G and $aG(\alpha)$ is the same.

By the above proposition, α is the smallest non-negative real such that $aG(\alpha)$ has a Nash equilibrium that is a social optimum. The game $aG(\alpha)$ can then be viewed as the intended transformation of G. That is, each payoff function p_i of the game G is transformed into the payoff function

$$r_i(s) := \frac{1}{1 + \alpha n} p_i(s) + \frac{\alpha}{1 + \alpha n} SW(s).$$

Note that the selfishness level is not invariant under a multiplication of the payoff functions by a value $a \leq 0$. Indeed, for $a = 0$ each game aG has the selfishness level 0. For $a < 0$ take the game G from Example 3 whose selfishness level is ∞. In the game aG the joint strategy (E, E) is both a Nash equilibrium and a social optimum, so the selfishness level of aG is 0.

Theorem 2. *There exists a game whose selfishness level is 0^+, i.e., it is α-selfish for every $\alpha > 0$, but it is not 0-selfish.*

3 A Characterization Result

We now characterize the games with a finite selfishness level. To this end we shall need the following notion. We call a social optimum s ***stable*** if for all $i \in N$ and $s'_i \in S_i$ the following holds: if (s'_i, s_{-i}) is a social optimum, then $p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i})$. In other words, a social optimum is stable if no player is better off by unilaterally deviating to another social optimum.

It will turn out that to determine the selfishness level of a game we need to consider deviations from its stable social optima. Consider a deviation s'_i of player i from a social stable optimum s. If player i is better off by deviating to s'_i, then by definition the social welfare decreases, i.e., $SW(s_i, s_{-i}) - SW(s'_i, s_{-i}) > 0$. If this decrease is small, while the gain for player i is large, then strategy s'_i is an attractive and socially acceptable option for player i. We define player i's ***appeal factor*** of strategy s'_i given the social optimum s as

$$AF_i(s'_i, s) := \frac{p_i(s'_i, s_{-i}) - p_i(s_i, s_{-i})}{SW(s_i, s_{-i}) - SW(s'_i, s_{-i})}.$$

In what follows we shall characterize the selfishness level in terms of bounds on the appeal factors of profitable deviations from a stable social optimum.

Theorem 3. *Consider a strategic game* $G := (N, \{S_i\}_{i\in N}, \{p_i\}_{i\in N})$.

1. *The selfishness level of* G *is finite iff a stable social optimum* s *exists for which* $\alpha(s) := \max_{i\in N,\, s'_i \in U_i(s)} AF_i(s'_i, s)$ *is finite, where* $U_i(s) := \{s'_i \in S_i \mid p_i(s'_i, s_{-i}) > p_i(s_i, s_{-i})\}$.
2. *If the selfishness level of* G *is finite, then it equals* $\min_{s\in SSO} \alpha(s)$*, where* SSO *is the set of stable social optima.*
3. *If* G *is finite, then its selfishness level is finite iff it has a stable social optimum. In particular, if* G *has a unique social optimum, then its selfishness level is finite.*
4. *If* $\beta > \alpha \geq 0$ *and* G *is* α*-selfish, then* G *is* β*-selfish.*

4 Examples

We now use the above characterization result to determine or compute an upper bound on the selfishness level of some selected games. First, we exhibit a well-known class of games (see [15]) for which the selfishness level is finite.

Potential Games. Given a game $G := (N, \{S_i\}_{i\in N}, \{p_i\}_{i\in N})$, a function $P : S_1 \times \cdots \times S_n \to \mathbb{R}$ is called an ***ordinal potential function*** for G if for all $i \in N$, $s_{-i} \in S_{-i}$ and $s_i, s'_i \in S_i$, $p_i(s_i, s_{-i}) > p_i(s'_i, s_{-i})$ iff $P(s_i, s_{-i}) > P(s'_i, s_{-i})$. A game that possesses an ordinal potential function is called an ***ordinal potential game***.

Theorem 4. *Every finite ordinal potential game has a finite selfishness level.*

In particular, every finite congestion game (see [16]) has a finite selfishness level. We derive below explicit bounds for two special cases of these games.

Fair Cost Sharing Games. In a fair cost sharing game, see, e.g., [17], players allocate facilities and share the cost of the used facilities in a fair manner. Formally, a fair cost sharing game is given by $G = (N, E, \{S_i\}_{i \in N}, \{c_e\}_{e \in E})$, where $N = \{1, \dots, n\}$ is the set of players, E is the set of facilities, $S_i \subseteq 2^E$ is the set of facility subsets available to player i, and $c_e \in \mathbb{R}_+$ is the cost of facility $e \in E$. It is called a *singleton* cost sharing game if for every $i \in N$ and for every $s_i \in S_i$: $|s_i| = 1$. For a joint strategy $s \in S_1 \times \cdots \times S_n$ let $x_e(s)$ be the number of players using facility $e \in E$, i.e., $x_e(s) = |\{i \in N \mid e \in s_i\}|$. The cost of a facility $e \in E$ is evenly shared among the players using it. That is, the 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 given by $SC(s) = \sum_{i \in N} c_i(s)$.

We first consider singleton cost sharing games. Let $c_{\max} = \max_{e \in E} c_e$ and $c_{\min} = \min_{e \in E} c_e$ refer to the maximum and minimum costs of the facilities, respectively.

Proposition 2. *The selfishness level of a singleton cost sharing game is at most* $\max\{0, \frac{1}{2} c_{\max}/c_{\min} - 1\}$. *Moreover, this bound is tight.*

This result should be contrasted with the price of stability of H_n and the price of anarchy of n for cost sharing games [17]. Cost sharing games admit an exact potential function and thus by Theorem 4 their selfishness level is finite. However, one can show that the selfishness level can be arbitrarily large (as $c_{\max}/c_{\min} \to \infty$) even for $n = 2$ and two facilities.

We next derive a bound for arbitrary fair cost sharing games with non-negative integer costs. Let L be the maximum number of facilities that any player can choose, i.e., $L := \max_{i \in N, s_i \in S_i} |s_i|$.

Proposition 3. *The selfishness level of a fair cost sharing game with non-negative integer costs is at most* $\max\{0, \frac{1}{2} L c_{\max} - 1\}$. *Moreover, this bound is tight.*

Remark 1. We can bound the selfishness level of a fair cost sharing game with non-negative rational costs $c_e \in \mathbb{Q}_+$ for every facility $e \in E$ by using Proposition 3 and the following scaling argument: Simply scale all costs to integers, e.g., by multiplying them with the least common multiplier $q \in \mathbb{N}$ of the denominators. Note that this scaling does not change the selfishness level of the game by Proposition 1. However, it does change the maximum facility cost and thus q enters the bound.

Linear Congestion Games. In a congestion game $G := (N, E, \{S_i\}_{i \in N}, \{d_e\}_{e \in E})$ we are given a set of players $N = \{1, \dots, n\}$, a set of facilities E with a delay function $d_e : \mathbb{N} \to \mathbb{R}_+$ for every facility $e \in E$, and a strategy set $S_i \subseteq 2^E$ for every player $i \in N$. For a joint strategy $s \in S_1 \times \cdots \times S_n$, define $x_e(s)$ as the number of players using facility $e \in E$, i.e., $x_e(s) = |\{i \in N \mid e \in s_i\}|$. The goal of a player is to minimize his individual cost $c_i(s) = \sum_{e \in s_i} d_e(x_e(s))$. The social cost function is given by $SC(s) = \sum_{i \in N} c_i(s)$. Here we call a congestion game *symmetric* if there is some common strategy set $S \subseteq 2^E$ such that $S_i = S$ for all

i. It is *singleton* if every strategy $s_i \in S_i$ is a singleton set, i.e., for every $i \in N$ and for every $s_i \in S_i$, $|s_i| = 1$. In a *linear* congestion game, the delay function of every facility $e \in E$ is of the form $d_e(x) = a_e x + b_e$, where $a_e, b_e \in \mathbb{R}_+$ are non-negative real numbers.

We first derive a bound on the selfishness level for symmetric singleton linear congestion games. As it turns out, a bound similar to the one for singleton cost sharing games does not extend to symmetric singleton linear congestion games. Instead, the crucial insight here is that the selfishness level depends on the *discrepancy* between any two facilities in a stable social optimum. We make this notion more precise.

Let s be a stable social optimum and let x_e refer to $x_e(s)$. Define the *discrepancy* between two facilities e and e' under s as

$$\lambda(x_e, x_{e'}) = \frac{2a_e x_e + b_e}{a_e + a_{e'}} - \frac{2a_{e'} x_{e'} + b_{e'}}{a_e + a_{e'}}. \tag{2}$$

It can be shown that $\lambda(x_e, x_e') \in (-1, 1)$. Let $\lambda_{\max}(s)$ be the maximum discrepancy between any two facilities under s. Further, let $\lambda_{\max}$ be the maximum discrepancy over all stable social optima, i.e., $\lambda_{\max} = \max_{s \in SSO} \lambda_{\max}(s)$.

Let $\Delta_{\max} := \max_{e \in E}(a_e + b_e)$ and $\Delta_{\min} := \min_{e \in E}(a_e + b_e)$. Further, let $a_{\min}$ be the minimum non-zero coefficient of a latency function, i.e., $a_{\min} = \min_{e \in E: a_e > 0} a_e$.

Proposition 4. *The selfishness level of a symmetric singleton linear congestion game is at most* $\max\{0, \frac{1}{2}(\Delta_{\max} - \Delta_{\min})/((1 - \lambda_{\max})a_{\min}) - \frac{1}{2}\}$. *Moreover, this bound is tight.*

Observe that the selfishness level depends on the ratio $(\Delta_{\max} - \Delta_{\min})/a_{\min}$ and $1/(1 - \lambda_{\max})$. In particular, the selfishness level becomes arbitrarily large as $\lambda_{\max}$ approaches 1.

We next state a bound for the selfishness level of arbitrary congestion games with linear delay functions and non-negative integer coefficients, i.e., $d_e(x) = a_e x + b_e$ with $a_e, b_e \in \mathbb{N}$ for every $e \in E$. Let L be the maximum number of facilities that any player can choose, i.e., $L := \max_{i \in N, s_i \in S_i} |s_i|$.

Proposition 5. *The selfishness level of a linear congestion game with non-negative integer coefficients is at most* $\max\{0, \frac{1}{2}(L\Delta_{\max} - \Delta_{\min} - 1)\}$. *Moreover, this bound is tight.*

For linear congestion games, the price of anarchy is known to be $\frac{5}{2}$, see [18, 19]. In contrast, our bound shows that the selfishness level depends on the maximum number of facilities in a strategy set and the magnitude of the coefficients of the delay functions.

Remark 2. We can use Proposition 5 and the scaling argument outlined in Remark 1 to derive bounds on the selfishness level of congestion games with linear delay functions and non-negative rational coefficients.

Prisoner's Dilemma for n Players. We assume that each player $i \in N = \{1, \ldots, n\}$ has two strategies, 1 (cooperate) and 0 (defect). We put $p_i(s) := 1 - s_i + 2\sum_{j \neq i} s_j$.

Proposition 6. *The selfishness level of the n-players Prisoner's Dilemma game is $\frac{1}{2n-3}$.*

Intuitively, this means that when the number of players in the Prisoner's Dilemma game increases, a smaller share of the social welfare is needed to resolve the underlying conflict. That is, its 'acuteness' diminishes with the number of players. The formal reason is that the appeal factor of each unilateral deviation from the social optimum is inversely proportional to the number of players.

In particular, for $n = 2$ we get, as already argued in Example 1, that the selfishness level of the original Prisoner's Dilemma game is 1.

Public Goods. We consider the public goods game with n players. Every player $i \in N = \{1, \ldots, n\}$ chooses an amount $s_i \in [0, b]$ that he contributes to a public good, where $b \in \mathbb{R}_+$ is the budget. The game designer collects the individual contributions of all players, multiplies their sum by $c > 1$ and distributes the resulting amount evenly among all players. The payoff of player i is thus $p_i(s) := b - s_i + \frac{c}{n}\sum_{j \in N} s_j$.

Proposition 7. *The selfishness level of the n-players public goods game is $\max\left\{0, \frac{1-\frac{c}{n}}{c-1}\right\}$.*

In this game, every player has an incentive to "free ride" by contributing 0 to the public good (which is a dominant strategy). The above proposition reveals that for fixed c, in contrast to the Prisoner's Dilemma game, this temptation becomes stronger as the number of players increases. Also, for a fixed number of players this temptation becomes weaker as c increases.

Cournot Competition. We consider Cournot competition for n firms with a linear inverse demand function and constant returns to scale, see, e.g., [1, pages 174–175]. So we assume that each player $i \in N = \{1, \ldots, n\}$ has a strategy set $S_i = \mathbb{R}_+$ and payoff function $p_i(s) := s_i(a - b\sum_{j \in N} s_j) - cs_i$ for some given a, b, c, where $a > c \geq 0$ and $b > 0$.

The price of the product is represented by the expression $a - b\sum_{j \in N} s_j$ and the production cost corresponding to the production level s_i by cs_i. In what follows we rewrite the payoff function as $p_i(s) := s_i(d - b\sum_{j \in N} s_j)$, where $d := a - c$.

Proposition 8. *The selfishness level of the n-players Cournot competition game is ∞.*

Intuitively, this result means that in this game no matter how much we 'involve' the players in sharing the social welfare we cannot achieve that they will select a social optimum.

Tragedy of the Commons. Assume that each player $i \in N = \{1, \ldots, n\}$ has the real interval $[0, 1]$ as its set of strategies. Each player's strategy is his chosen fraction of a common resource. Let (see [20, Exercise 63.1]): $p_i(s) := \max\{0, s_i(1 - \sum_{j \in N} s_j)\}$. This payoff function reflects the fact that player's enjoyment of the common resource depends positively from his chosen fraction of the resource and negatively from the total fraction of the common resource used by all players. Additionally, if the total fraction of the common resource by all players exceeds a feasible level, here 1, then player's enjoyment of the resource becomes zero.

Proposition 9. *The selfishness level of the n-players Tragedy of the Commons game is ∞.*

Bertrand Competition. Next, we consider Bertrand competition, a game concerned with a simultaneous selection of prices for the same product by two firms, see, e.g., [1, pages 175–177]. The product is then sold by the firm that chose a lower price. In the case of a tie the product is sold by both firms and the profits are split. We assume that each firm has identical marginal costs $c > 0$ and no fixed cost, and that each strategy set S_i equals $[c, \frac{a}{b})$, where $c < \frac{a}{b}$. The payoff function for player $i \in \{1, 2\}$ is given by

$$p_i(s_i, s_{3-i}) := \begin{cases} (s_i - c)(a - bs_i) & \text{if } c < s_i < s_{3-i} \\ \frac{1}{2}(s_i - c)(a - bs_i) & \text{if } c < s_i = s_{3-i} \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 10. *The selfishness level of the Bertrand competition game is ∞.*

5 Concluding Remarks and Extensions

We introduced the selfishness level of a game as a new measure of discrepancy between the social welfare in a Nash equilibrium and in a social optimum. Our studies reveal that the selfishness level often provides more refined insights than other measures of inefficiency.

The definition of the selfishness level naturally extends to other solution concepts and other forms of games. For example, for mixed Nash equilibria we simply adapt our definitions by stipulating that a strategic game G is α-selfish if the social welfare of a mixed Nash equilibrium of $G(\alpha)$ is equal to the optimum social welfare of $G(\alpha)$. The selfishness level of G is then defined as before in (1). For example, with this notion the selfishness level of the Matching Pennies game is 0.

We can also consider subgame perfect equilibria and extensive games. We leave for future work the study of such alternatives.

Acknowledgements: We acknowledge initial discussions with Po-An Chen and thank anonymous reviewers for their valuable comments.

References

1. Jehle, G., Reny, P.: Advanced Microeconomic Theory. Third edn. Addison Wesley, New York, NY (2011)
2. Koutsoupias, E., Papadimitriou, C.H.: Worst case equilibria. In: Annual IEEE Symposium on Theoretical Aspects of Computer Science. (1999) 404–413
3. Schulz, A.S., Moses, N.E.S.: On the performance of user equilibria in traffic networks. In: SODA. (2003) 86–87
4. Axelrod, R.: The Evolution of Cooperation. Basic Books (1984)
5. Ledyard, J.O.: 2. The Handbook of Experimental Economics. In: Public Goods: A Survey of Experimental Research. Princeton University Press (1995) 111–194
6. Marco, G.D., Morgan, J.: Slightly altruistic equilibria in normal form games. Working paper No 185, Center for Studies in Economics and Finance, University of Salerno, Italy (2007) Available from http://www.csef.it/WP/wp185.pdf.
7. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M., Papaioannou, E.: The impact of altruism on the efficiency of atomic congestion games. In: Proc. 5th Symposium on Trustworthy Global Computing. (2010)
8. Chen, P.A., de Keijzer, B., Kempe, D., Schäfer, G.: The robust price of anarchy of altruistic games. In: Proc. 7th Workshop on Internet and Network Economics. (2011) 383–390
9. Chen, P.A., Kempe, D.: Altruism, selfishness, and spite in traffic routing. In: Proc. 10th ACM Conference on Electronic Commerce. (2008) 140–149
10. Elias, J., Martignon, F., Avrachenkov, K., Neglia, G.: Socially-aware network design games. In: Proc. INFOCOM 2010. (2010) 41–45
11. Hoefer, M., Skopalik, A.: Altruism in atomic congestion games. In: Proc. 17th European Symposium on Algorithms. (2009)
12. Roughgarden, T.: Intrinsic robustness of the price of anarchy. In: Proc. 41st Annual ACM Symposium on Theory of Computing. (2009) 513–522
13. Beckmann, M., McGuire, B., Winsten, C.: Studies in the Economics of Transportation. Yale University Press, New Haven (1956)
14. Sharma, Y., Williamson, D.P.: Stackelberg thresholds in network routing games or the value of altruism. Games and Economic Behavior **67**(1) (2009) 174–190
15. Monderer, D., Shapley, L.S.: Potential games. Games and Economic Behaviour **14** (1996) 124–143
16. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. International Journal of Game Theory (2) (1973) 65–67
17. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: Proc. 45th Annual IEEE Symposium on Foundations of Computer Science. (2004) 295–304
18. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proc. 37th Annual ACM Symposium on Theory of Computing. (2005)
19. Awerbuch, B., Azar, Y., Epstein, A.: Large the price of routing unsplittable flow. In: Proc. 37th Annual ACM Symposium on Theory of Computing. (2005) 57–66
20. Osborne, M.J.: An Introduction to Game Theory. Oxford University Press, Oxford (2005)