

Inefficiencies in Cost and Market Sharing Games

Thomas Kesselheim

Last Update: May 6, 2019

Today, we will wrap up our discussion of the inefficiency due to selfish behavior and show two further results. The first one is about a cost-minimization game, for which we will consider a less pessimistic perspective than the price of anarchy. The second one is about a very similar game but now with a utility-maximization objective.


1 Cost-Sharing Games

We first turn to *fair cost-sharing games*, which are congestion games with delays $d_r(x) = c_r/x$ for constant c_r for every resource $r \in \mathcal{R}$. That is, we have a set \mathcal{N} of n players and a set \mathcal{R} of m resources. Player i allocates some resources, i.e., his strategy set is $\Sigma_i \subseteq 2^{\mathcal{R}}$. Each resource $r \in \mathcal{R}$ has fixed cost $c_r \geq 0$. The cost c_r is assigned in equal shares to the players allocating r (if any).

Social cost turns out to be the sum of costs of resources allocated by at least one player:

$$SC(S) = \sum_{i \in \mathcal{N}} c_i(S) = \sum_{i \in \mathcal{N}} \sum_{r \in S_i} d_r(n_r(S)) = \sum_{\substack{r \in \mathcal{R} \\ n_r(S) \geq 1}} n_r(S) \cdot c_r / n_r(S) = \sum_{\substack{r \in \mathcal{R} \\ n_r(S) \geq 1}} c_r . \quad (1)$$

The price of anarchy for pure Nash equilibria can be as big as the number of players n , even in a symmetric game. For $\epsilon > 0$, consider the example



Edge labels indicate the cost value c_r for this resource. It is a pure Nash equilibrium if all players use the bottom edge, whereas the social optimum would be that all users use the top edge.

Although this is a very stylized example, there are indeed examples of such bad equilibria occurring in reality. A prime example are mediocre technologies, which win against better ones just because they are in the market early and get their share. This way, they are widely supported. Maybe another example are social networks and messaging apps. Many people would prefer not to use, say, Facebook but they cannot switch to an alternative platform unless their friends do.

1.1 Price of Stability

The price-of-anarchy viewpoint is still a pessimistic one because we make statements about the *worst* equilibria. This is different in *price of stability*. For an equilibrium concept Eq , it is defined as

$$PoS_{\text{Eq}} = \frac{\min_{p \in \text{Eq}} SC(p)}{\min_{s \in S} SC(s)} .$$

As the set of equilibria gets larger, the minimum gets smaller and smaller. Therefore, if the respective equilibria exist, we now have

$$1 \leq PoS_{\text{CCE}} \leq PoS_{\text{CE}} \leq PoS_{\text{MNE}} \leq PoS_{\text{PNE}} \leq PoA_{\text{PNE}} \leq \dots .$$

Theorem 9.1. *In a symmetric cost-sharing game, the price of stability for pure Nash equilibria is 1.*


Proof. We explicitly construct a pure Nash equilibrium as follows. It is a symmetric equilibrium, meaning that all players use the same strategy. Consider player 1, and set S_1 to the strategy from Σ_1 that minimizes $\sum_{r \in S_1} c_r$. Set $S_2 = \dots = S_n = S_1$. This certainly minimizes social cost according to Equation (1).

It is also an equilibrium because for each i strategy S_i is a best response against S_{-i} . To see this, consider some alternative $S'_i \in \Sigma_i$. We have

$$c_i(S'_i, S_{-i}) = \sum_{r \in S'_i \cap S_i} \frac{c_r}{n} + \sum_{r \in S'_i \setminus S_i} c_r \geq \frac{1}{n} \sum_{r \in S'_i} c_r \geq \frac{1}{n} \sum_{r \in S_i} c_r = c_i(S) ,$$

where we used that S_i was chosen to minimize $\sum_{r \in S_i} c_r$. □

The core insight of the previous proof is that in a symmetric game every social optimum is a pure Nash equilibrium. For general, asymmetric games, the social optimum is not necessarily a pure Nash equilibrium. Consider the following game with n players. Each player i has source node s_i and destination node t .



A player two possible strategies: Either take the direct edge or take the detour via v . The social optimum lets all players choose the indirect path, ending up with social cost $1 + \epsilon$. This, however, is no Nash equilibrium. Player n would opt out and take the direct edge. Therefore, the only pure Nash equilibrium lets all players choose their direct edge, yielding social cost of \mathcal{H}_n . Here, $\mathcal{H}_n = \sum_{i=1}^n \frac{1}{i} = 1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n}$ denotes the n -th harmonic number. We have $\mathcal{H}_n = \Theta(\log n)$.

Theorem 9.2. *The Price of Stability for pure Nash equilibria in fair cost sharing games is at most \mathcal{H}_n .*

Proof. Rosenthal’s potential function for cost-sharing delays is

$$\begin{aligned} \Phi(S) &= \sum_{r \in \mathcal{R}} \sum_{i=1}^{n_r} c_r / i = \sum_{\substack{r \in \mathcal{R} \\ n_r \geq 1}} c_r \cdot \left(1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n_r} \right) \\ &\leq \sum_{\substack{r \in \mathcal{R} \\ n_r \geq 1}} c_r \cdot \mathcal{H}_n \\ &= SC(S) \cdot \mathcal{H}_n . \end{aligned}$$

In $\Phi(S)$ we account for each player allocating resource r a contribution of c_r/i for some $i = 1, \dots, n_r$, whereas in his cost $c_i(S)$ we account only c_r/n_r . Hence, for every state S of a cost-sharing game we have

$$SC(S) \leq \Phi(S) \leq SC(S) \cdot \mathcal{H}_n .$$

Now suppose we start at the optimum state S^* and iteratively perform improvement steps for single players. This eventually leads to a pure Nash equilibrium. Every such move decreases the potential function. For the resulting Nash equilibrium S we thus have $\Phi(S) \leq \Phi(S^*)$ and

$$SC(S) \leq \Phi(S) \leq \Phi(S^*) \leq SC(S^*) \cdot \mathcal{H}_n .$$

This proves that there is a pure Nash equilibrium that is only a factor of \mathcal{H}_n more costly than S^* . \square

2 Market-Sharing Games

Most of our examples so far in this course were cost-minimization games. For the basic definitions there is no real difference when one turns to utility-maximization games instead. However, for the price of anarchy, the story is different, as we will see in the following example.

Let us consider the following *market sharing game*. There are n firms, which are our players N , and m markets M . Each firm can decide to invest in one of these markets. Therefore, for player $i \in N$, the strategy set S_i is a subset of M .

Each market $j \in M$ has a total demand v_j . If k firms invest in the same market, then the market's demand is shared equally. So every firm gets a utility of $\frac{v_j}{k}$.

This way, the utility of player $i \in N$ in state $s \in S$ is

$$u_i(s) = \frac{v_{s_i}}{n_{s_i}(s)} , \quad \text{where } n_j(s) = |\{i \in N \mid s_i = j\}| .$$

The *social welfare* of a state s is defined as the sum of player utilities, or equivalently, as the sum of demands that are fulfilled

$$SW(s) = \sum_{i \in N} u_i(s) = \sum_{\substack{j \in M \\ n_j(s) \geq 1}} v_j = \sum_{j \in \{s_1, \dots, s_n\}} v_j .$$

Example 9.3. *There are n markets $1, \dots, n$; each player can invest in every market. For some $\epsilon > 0$, the demands are $v_1 = n + \epsilon$, $v_2 = \dots = v_n = 1$.*

The social welfare is maximized by each player investing in a different market. In this case, $SW(s) = 2n - 1 + \epsilon$. However, the only pure Nash equilibrium is that all players invest in market 1. Here, $SW(s) = n + \epsilon$.

So far, these games look a lot like the cost-sharing games. And, indeed, they are. We can even interpret them as congestion games: Set $\mathcal{R} = M$, so the markets become the resources, and set $d_j(k) = -\frac{v_j}{k}$ for all $j \in M$ and all k . Now the players' cost functions in the congestion game are exactly the negative utility functions of the market sharing game: $c_i(s) = -u_i(s)$.

Observation 9.4. *Every Market Sharing Game has a pure Nash equilibrium.*

2.1 Price of Anarchy

Interestingly, despite the similarity to cost-sharing games, the price of anarchy is a lot different. Let us first define the price of anarchy for utility-maximization games. The definition is analogous to the one for cost-minimization game.

Definition 9.5. Given a utility-maximization game, let Eq be a set of probability distributions over the set of states S . For some probability distribution p , let $SW(p) = \mathbf{E}_{s \sim p}[SW(s)] = \sum_{s \in S} p(s)SW(s)$ be the expected social welfare. The price of anarchy for Eq is defined as

$$PoA_{\text{Eq}} = \frac{\max_{s \in S} SW(s)}{\min_{p \in \text{Eq}} SW(p)} .$$

So, we swap minima and maxima and the social optimum is now in the numerator and the equilibrium in the denominator. This way, the PoA is still greater than 1. You will also find it defined as the reciprocal. In any case, values closer to 1 are better.

We can also adapt the smoothness definition as follows.

Definition 9.6. A utility-maximization game is called (λ, μ) -smooth for $\lambda > 0$ and $\mu \geq 0$ if, for every pair of states $s, s^* \in S$, we have

$$\sum_{i \in N} u_i(s_i^*, s_{-i}) \geq \lambda \cdot SW(s^*) - \mu \cdot SW(s) .$$

There is again an analogous theorem that smoothness implies a bound on the price of anarchy.

Theorem 9.7. In a (λ, μ) -smooth utility-maximization game, the PoA for coarse correlated equilibria is at most $\frac{1+\mu}{\lambda}$.

Proof. For simplicity, we prove the theorem only for pure Nash equilibria. The generalization to coarse correlated equilibria works exactly as in the case of cost-minimization games.

Let s be a pure Nash equilibrium, s^* be a social optimum. Then we have

$$\begin{aligned} SW(s) &= \sum_{i \in N} u_i(s) \\ &\geq \sum_{i \in N} u_i(s_i^*, s_{-i}) \\ &\geq \lambda \cdot SW(s^*) - \mu \cdot SW(s) . \end{aligned}$$

So $(1 + \mu)SW(s) \geq \lambda SW(s^*)$. □

So, it only remains to give a smoothness proof.

Theorem 9.8. The market-sharing game is $(1, 1)$ -smooth. So $PoA_{\text{CCE}} \leq 2$.

Proof. Observe that

$$\begin{aligned} u_i(s_i^*, s_{-i}) &= \frac{v_{s_i^*}}{n_{s_i^*}(s_i^*, s_{-i})} \\ &\geq \begin{cases} v_{s_i^*} & \text{if } s_i^* \notin \{s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n\} \\ 0 & \text{otherwise} \end{cases} \\ &\geq \begin{cases} v_{s_i^*} & \text{if } s_i^* \notin \{s_1, \dots, s_n, s_1^*, \dots, s_{i-1}^*\} \\ 0 & \text{otherwise} \end{cases} . \end{aligned} \tag{2}$$

Denote by $T = \{s_1, \dots, s_n\}$ all markets invested in in s and by $T^* = \{s_1^*, \dots, s_n^*\}$ all markets invested in in s^* . We can now write

$$\sum_{i \in N} u_i(s_i^*, s_{-i}) \geq \sum_{j \in T^* \setminus T} v_j = \sum_{j \in T^* \cup T} v_j - \sum_{j \in T} v_j .$$

This is because if we take the sum over the terms (2) then every element of T^* exactly appears once unless it is in T .

This gives us

$$\sum_{i \in N} u_i(s_i^*, s_{-i}) \geq \sum_{j \in T^* \cup T} v_j - \sum_{j \in T} v_j \geq \sum_{j \in T^*} v_j - \sum_{j \in T} v_j = SW(s^*) - SW(s) .$$

This is exactly the requirement for smoothness. □

Further Reading

- Chapter 19.3 in the AGT book. (PoS bound)
- Tim Roughgarden's lecture notes <http://theory.stanford.edu/~tim/f13/1/115.pdf> and lecture video <https://youtu.be/VjCKN1-GENI>
- A. Vetta, Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions, FOCS 2002 (Generalization of the result for market sharing game)