Game Theory
Lecture #15 — Price of Anarchy

Outline:

- Congestion games
- Price of Anarchy
- Influencing Nash equilibria

1 Recap

The last lecture focused on a class of games known as congestion games. The main result that we covered demonstrated that a congestion game is an instance of a potential game, which directly implies that a pure Nash equilibrium is guaranteed to exist in any congestion games. However, we also observed a difference in structure between the global objective and the potential function. This lecture will focus on the implication of these differences on the efficiency of the result Nash equilibria. Further, we will focus on the design of taxation mechanisms, i.e., mechanisms geared at influencing self-interested behavior, with the goal of improving the operational efficiency of the resulting Nash equilibria.

2 Congestion Games and Potential Games

We begin this lecture by reviewing the definition of congestion games and potential games.

Definition 2.1 (Congestion Game) *A congestion game is a game with the following elements:*

- *Resource set:* $\mathcal{R}$
- *Congestion functions: Each resource* $r \in \mathcal{R}$ *is associated with a congestion function* $c_r : \{0, 1, 2, \dots\} \to \mathbb{R}$*, where* $c_r(k)$ *denotes the congestion on resource* r *with* $k \geq 0$ *players*
- *Player set:* N
- *Action sets: Each player* $i \in N$ *is associated with an action set* $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$
- *Players' cost function: Each player* $i \in N$ *is associated with a cost function* $J_i : \mathcal{A} \to \mathbb{R}$ *where for any joint action* $a = (a_i, a_{-i}) \in \mathcal{A}$

$$J_i(a_i, a_{-i}) = \sum_{r \in a_i} c_r(|a|_r)$$

where $|a|_r := |j \in N : r \in a_j|$ captures the number of players selecting resource r in the joint action a

- *Social welfare: The efficiency of a given action profile $a \in \mathcal{A}$ is measured with respect to the social welfare, which takes on the form*

$$C(a) = \sum_{i \in N} J_i(a) = \sum_{r \in \mathcal{R}} |a|_r \cdot c_r(|a|_r)$$

Congestion games can be used to model a wide variety of problems related to engineering systems related to transportation network, task assignment and resource allocation problems. What is a reasonable prediction of behavior in such games? A pure Nash equilibrium would be a viable solution, however at the beginning of last lecture we were unsure whether or not one was guaranteed to exist. To answer this question we focused on drawing an equivalence between congestion games and potential games, defined as follows:

Definition 2.2 (Potential Game) *A game with player set $N = \{1, 2, ..., n\}$, action sets $(\mathcal{A}_i)_{i \in N}$, and utility functions $U_i : \mathcal{A} \to \mathbb{R}$ (or cost functions $J_i : \mathcal{A} \to \mathbb{R}$) is called a* potential game *if there exists a function $\phi : \mathcal{A} \to R$ such that for every action profile $a \in \mathcal{A}$, player $i \in N$, and alternative choice $a_i' \in \mathcal{A}_i$*

$$U_i(a_i', a_{-i}) - U_i(a_i, a_{-i}) = \phi(a_i', a_{-i}) - \phi(a_i, a_{-i}). \tag{1}$$

This function ϕ is called the potential function.

Pure Nash equilibria are guaranteed to exist in any potential game; however, their relevance to engineering systems appears limited when viewed in comparison to congestion games. However, Monderer and Shapley changed this perspective when demonstrating that a congestion game is in fact a potential game as shown in the following theorem.

Theorem 2.1 (Monderer and Shapley, 1996) *Consider any congestion game as defined above. The congestion game is a potential game with potential function*

$$\phi(a) = \sum_{r \in \mathcal{R}} \sum_{k=1}^{|a|_r} c_r(k)$$

This theorem demonstrates that a pure Nash equilibrium is guaranteed to exist in any congestion game. This holds irrespective of the number of players, the number of resources, the resource congestion functions, or the agents' action sets. Further, one such pure Nash equilibrium is the action profile a^* that minimizes the potential function $\phi(\cdot)$, i.e.,

$$a^* \in \arg\min_{a \in \mathcal{A}} \phi(a) = \arg\min_{a \in \mathcal{A}} \sum_{r \in \mathcal{R}} \sum_{k=1}^{|a|_r} c_r(k).$$

However, the differences between the potential function ϕ and the social welfare C can create problems with regards to the efficiency of pure Nash equilibria as the following example demonstrates.

Example 2.1 *Consider a two link parallel network congestion game with a High and Low road as shown in*

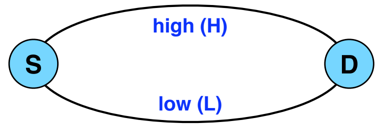

Suppose there are two players, i.e., $N = \{1, 2\}$*, and each player* $i \in N$ *can select either resource, i.e.,* $\mathcal{A}_i = \{\{H\}, \{L\}\}$ *Further, suppose the resource specific congestion functions satisfy* $c_H(k) = k$ *and* $c_L(k) = 2 + \epsilon$ *for any* $k \in \{1, 2\}$*. Accordingly, for this game the players' cost functions can be expressed in matrix form as*

	H	L
H	$2, 2$	$1, 2+\epsilon$
L	$2+\epsilon, 1$	$2+\epsilon, 2+\epsilon$

Player Costs

Observe that (H, H) *is the unique pure Nash equilibrium for this example. Turning our attention towards the potential function and social welfare (i.e., system cost), we have*

	H	L
H	4	$3+\epsilon$
L	$3+\epsilon$	$4+2\epsilon$

System Cost C

	H	L
H	3	$3+\epsilon$
L	$3+\epsilon$	$4+2\epsilon$

Potential Function ϕ

Note that the pure Nash equilibrium of this congestion game is precisely aligned with the pure Nash equilibrium of the identical interest game that would result from assigning each agent a common cost function that is equal to the potential function. This always holds true, as a potential game is an identical interest game in disguise where the common utility / cost function is the potential function. Of particular interest is the action profile that minimizes the potential function, i.e., (H, H)*, which is guaranteed to be a pure Nash equilibrium of the congestion game. Note however, that* (H, H) *does not minimize the social welfare, which is minimized by either* (L, H) *or* (H, L)*.*

3 Price of Anarchy

The example of the previous section demonstrates that a Nash equilibrium can be inefficient with regards to our desired system-level performance metric. In this section, we seek to identify how one should measure this inefficiency. To that, we turn our attention to the well-defined measure of the *price of anarchy* defined as follows:

Definition 3.1 (Price of Anarchy) *Consider a given game G with player set N where each player $i \in N$ is associated with an action set $\mathcal{A}_i$ and cost function $J_i : \mathcal{A} \to \mathbb{R}$. Further, suppose there is a system-level objective defined by $C : \mathcal{A} \to \mathbb{R}$. If a pure Nash equilibrium exists in the game G, then we define the price of anarchy as the ratio of performance between the "worst" Nash equilibrium and optimal action profile, i.e.,*

$$\mathrm{PoA}(G) = \frac{\max_{a^{\mathrm{ne}} \in NE(G)} \; C(a^{\mathrm{ne}})}{\min_{a^{\mathrm{opt}} \in \mathcal{A}} \; C(a^{\mathrm{opt}})} \geq 1$$

where $NE(G) \subseteq \mathcal{A}$ is the set of pure Nash equilibrium in the game G.

The price of anarchy provides an upper bound on the system-level cost associated with any pure Nash equilibrium relative to the optimal system-level cost, i.e., the system-level cost associated with any pure Nash equilibrium a^{ne} is the game G must satisfy

$$C(a^{\mathrm{ne}}) \leq \mathrm{PoA}(G) \cdot \min_{a^{\mathrm{opt}} \in \mathcal{A}} \; C(a^{\mathrm{opt}}).$$

With regards to the previous example is the last section (which has a unique pure Nash equilibrium), the price of anarchy was $\mathrm{PoA}(G) = 4/(3+\epsilon)$. The following example covers the case of non-unique Nash equilibria.

Example 3.1 *Consider a two player game with a cost matrix and system-level cost given by*

	H	L
H	$3,3$	$2,2$
L	$2,2$	$3,3$

Player Costs

	H	L
H	4	3
L	3	2

System Cost

Note that there are two pure Nash equilibria, (H, H) and (L, L), with a system-level cost of 4 and 2, respectively. The optimal system-level cost is 2, and hence we say that the price of anarchy is $4/2 = 2$ which bounds the cost of the worst pure Nash equilibrium. The fact that the best pure Nash equilibrium is the action profile that minimizes the system-level cost does not impact the computation of the price of anarchy.

Price of anarchy guarantees are more impactful if they apply to a large class of games $\mathcal{G}$. For example, the set of games $\mathcal{G}$ could include all congestion games with linear cost functions, all congestion games with polynomial cost functions, or all parallel congestion games. Accordingly, we define the price of anarchy of a set of games $\mathcal{G}$ as the worst price of anarchy of any game $G \in \mathcal{G}$.

Definition 3.2 (Price of Anarchy) *Consider a given set of games $\mathcal{G}$ where each game $G \in \mathcal{G}$ has a price of anarchy $\mathrm{PoA}(G)$. Then the price of anarchy of the set of games $\mathcal{G}$ is defined as the worst price of anarchy of any game $G \in \mathcal{G}$, i.e.,*

$$\mathrm{PoA}(\mathcal{G}) = \max_{G \in \mathcal{G}} \mathrm{PoA}(G).$$

Clearly, defining a price of anarchy over a class of games $\mathcal{G}$ is more challenging than defining the price of anarchy of a single game G. However it is also for more valuable as any pure Nash equilibrium in any game $G \in \mathcal{G}$ will have a system-level cost within the price of anarchy bounds. Accordingly, if $\text{PoA}(\mathcal{G})$ is close to 1, then any game in the set $\mathcal{G}$ will have desirable equilibrium behavior.

Example 3.2 *Consider a set of games $\mathcal{G}$ parameterized by $x \in (-\infty, \infty)$*

	H	L
H	$2+x, 2+x$	$1, 1$
L	$1, 1$	$2+x, 2+x$

Player Costs

	H	L
H	2	3
L	4	2

System Cost

Note that the value of x changes the players' cost functions but does not change the resulting system-cost of any action profile. Accordingly, what is the price of anarchy for this family of games? We begin with an equilibrium analysis as a function of x:

- $x < -1$*: There are two Nash equilibrium (H, H) and (L, L), each with a system-level cost of* 2*.*
- $x = -1$*: All action profiles are pure Nash equilibria.*
- $x > -1$*: There are two Nash equilibrium (H, L) and (L, H), with a system-level cost of* 3 *and* 4 *respectively.*

Accordingly, the price of anarchy of the set of games is computed from a worst-case perspective over the three regions defined above. Since the optimal action profile has a system-level cost of 2*, we have that* $\text{PoA}(\mathcal{G}) = 2$ *as the worst Nash equilibrium in any region has a system-cost of* 4*. If we consider the set of games induced by $x < -1$, then the price of anarchy is* $\text{PoA}(\mathcal{G}_{x<-1}) = 1$*.*

3.1 Non-atomic congestion games

How do we characterize the price of anarchy of a set of games $\mathcal{G}$? Further, are there meaningful classes of games for which the price of anarchy is relatively small? These two questions have been researched extensively in the past 20 years in Computer Science with several noteworthy results. Here, we will cover latter question by focusing on a class of congestion games known as non-atomic congestion games. The terminology associated with congestion games is as follows:

- Atomic congestion games: This is the setting that we have commonly studied in lecture involving a discrete set of drivers.

- Non-atomic congestion games: This is the setting that we have studied sporadically in lecture where there are a continuum of drivers that are infinitesimally small

While we have primarily discussed atomic congestion games in this class, it is often easier to characterize price of anarchy guarantees in non-atomic congestion games. Informally, an atomic congestion game approaches a non-atomic game when the number of drivers becomes large.

While we will not delve into the formalities of non-atomic congestion games here, the key differences from atomic congestion games can be summarized here:

- Continuum of player: $[0, 1]$
- Each user $i \in [0, 1]$ with choice a_i experiences a cost $\sum_{r \in a_i} c_r(f_r)$ where $f_r \geq 0$ is the total flow on edge r
- A Nash flow (as opposed to Nash equilibrium) is a feasible flow such that no user $i \in [0, 1]$ will seek to deviate

The following example will shed light on the structure of non-atomic congestion games.

Example 3.3 *Consider a routing problem where 1 unit of mass seeking to traverse from S to D across the following network*

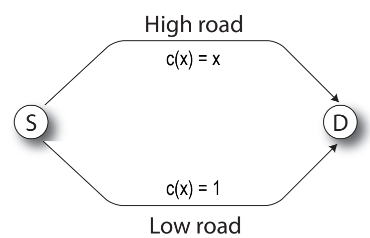

What is the flow that minimizes the total congestion in the network? First note that any feasible flow can be expressed as $(f, 1-f)$ where $f \in [0, 1]$ defines the fraction of traffic to put on the High Road. Accordingly, we have

$$f^{\text{opt}} = \arg\min_{f \in [0,1]} \left(f c_H(f) + (1-f) c_L(1-f) \right) = \arg\min_{f \in [0,1]} \left(f^2 + (1-f) \right) = 1/2,$$

which leads to a system-level cost of $C(f^{\text{opt}}) = 3/4$. Turning our attention to Nash flows, we know that the only flow for which every player $i \in [0, 1]$ does not want to switch choices is when $f^{\text{nf}} = 1$, i.e., all drivers are on the High road. This leads to system-level cost of $C(f^{\text{nf}}) = 1$. Hence, the price of anarchy of this specific game is $\text{PoA}(G) = 4/3$.

We will now shift attention from defining the price of anarchy of a single non-atomic congestion game to a family of non-atomic congestion games. Specifically, we will focus on a class of non-atomic congestion games with affine congestion functions. This set of games, denoted by $\mathcal{G}$, has the following elements:

- Continuum of player: $[0, x]$, where $x > 0$ is the total mass of traffic
- Set of edges $\mathcal{R}$ where each edge $r \in \mathcal{R}$ has congestion function

$$c_r(f_r) = a_r \cdot f_r + b_r$$

  where $a_r, b_r \geq 0$ are edge-specific coefficients and $f_r \geq 0$ is the total flow on edge r

Accordingly, our goal here is to characterize the price of anarchy of the set of games $\mathcal{G}$? Note that the previous example G is an example of an affine congestion game, i.e., $G \in \mathcal{G}$. Hence, our price of anarchy must satisfy $\text{PoA}(\mathcal{G}) \geq 4/3$ as the price of anarchy is a worst-case analysis over all games in the set $\mathcal{G}$. How much worse is the price of anarchy when we consider the full set of non-atomic congestion games with affine costs? The answer is quite surprising.

Theorem 3.1 (Roughgarden, 2003) *The price of anarchy of the set of non-atomic congestion games with affine congestion functions is* $4/3$.

According, this theorem demonstrates that for any congestion game with non-atomic users and affine latency (or congestion) functions, the system cost associated with a Nash equilibrium will be at most 33% more than the cost associated with the optimal routing profile. This result holds true irrespective of the network structure, the mass of traffic, the set of feasible flows, and the affine congestion functions on each edge in the network. The two link example highlighted above is in fact the worst example in the set of games $\mathcal{G}$!

4 Influencing Nash equilibria

The surprising result highlighted above demonstrates that Nash flows have a system-level cost close to the optimal system-level cost for certain class of games. Nonetheless, here we ask the question as to whether the system operator can influence self-interested behavior to improve the efficiency of the Nash flows. In this section we consider the use of taxes (or monetary fees) for that purpose. Consider the framework of non-atomic congestion games where each edge $r \in \mathcal{R}$ is now associated with a taxation function $t_r : \mathbb{R}^+ \to \mathbb{R}^+$, where $t_r(f_r)$ is the toll on road r given a flow $f_r \geq 0$. Each player now tries to minimize a "new" cost function which incorporates both the congestion experienced and the toll. Now, the cost associated to player i given a choice a_i and a feasible flow f is of the form

$$\sum_{r \in a_i} c_r(f_r) + t_r(f_r).$$

Here, it is important to highlight that while the tolls influence the cost function of the players, it does not impact the system-level objective, which for any feasible flow f still takes on the form

$$\sum_{r \in \mathcal{R}} f_r \cdot c_r(f_r).$$

The question that we want to ask here is whether or not tolls can be employed to improve the efficiency of the resulting Nash flow. The following theorem answers this question in a surprising way.

Theorem 4.1 (Sandholm, 2002) *: Let G represent any non-atomic congestion game with congestion functions that are increasing and differentiable. Consider a special toll known as a Pigouvian tolls where each edge $r \in \mathcal{R}$ is associated with a tolling function*

$$t_r(f_r) = f_r \cdot c_r'(f_r)$$

for any $f_r \geq 0$. Then any Nash flow of the resulting game with tolls optimizes the system-cost, i.e., $\mathrm{PoA}\,(\mathcal{G}^t) = 1$.

This theorem is quite powerful and demonstrates how tolls can be effectively employed to improve the efficacy of self-interested behavior. While Pigouvian tolls might look complicated, effectively what they do is charge users for the dis-utility that they cause other drivers in the network. The following example will shed light into the working of Pigouvian tolls in affine congestion games.

Example 4.1 (Pigouvian Tolls in Affine Congestion Games) *Let G represent any non-atomic congestion game with congestion functions that are affine, i.e., each edge r is associated with a congestion function of the form $c_r(f_r) = a_r \cdot f_r + b_r$ where $a_r, b_r \geq 0$. Given the example above, we know that (i) Optimal flow: $f^{\mathrm{opt}} = (1/2, 1/2)$ and total cost is $C(f^{\mathrm{opt}}) = 3/4$; (ii) Nash flow: $f^{\mathrm{nf}} = (1, 0)$ and total cost is $C(f^{\mathrm{nf}}) = 1$; and (iii) Price of anarchy: $4/3$. Now, suppose each edge r is associated with a Pigouvian toll which takes on the form*

$$t_r(f_r) = a_r \cdot f_r$$

Accordingly, for our example we have

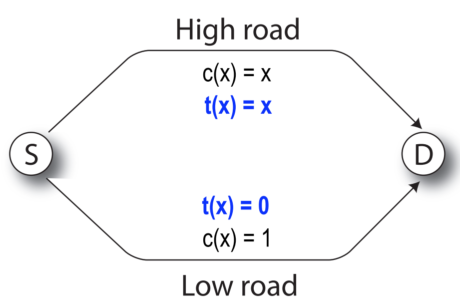

which can be interpreted as the High road now have any pseudo congestion function $\tilde{c}(x) = c(x) + t(x) = 2x$. Given this new pseudo congestion function, the Nash flow is now of of the form $f^{\mathrm{nf}} = (1/2, 1/2)$ as both the High and Low road will have a total pseudo cost of 1. Accordingly, the resulting system-level cost of this Nash flow is $C(f^{\mathrm{nf}}) = 3/4$ which is optimal. Hence, the price of anarchy is 1. This taxation approach ensures resulting Nash flow is optimal irrespective of the network!

5 Conclusion

This lecture focused on characterizing the inefficiency of Nash equilibrium in various of classes of games. Specifically, we introduce the efficiency measures of Price of Anarchy, which provides a worst-case measure on the quality of a Nash equilibrium relative to the optimal action profile. Interestingly, we demonstrated that some interested classes of games, e.g., affine congestion games, actually have a relatively small price of anarchy meaning that the inefficiency associated with self-interested behavior is quite manageable. This need not hold true over all games. Lastly, we illustrated a mechanism for influencing self-interested behavior to improve the efficacy of the resulting Nash flow. Interesting, in congesting games with congestion functions that are both increasing and differentiable, it is always possible to design a taxation mechanism that drives the price of anarchy to 1. These taxation mechanism introduce a way that a system operator can improve the outcome associated with self-interested behavior. This is very much aligned with the objective of the signaling structure and correlated equilibrium covered in the last lecture.

6 Exercises

1. Consider the class of atomic congestion games discussed in class. Here, we focus on the class of parallel congestion games defined as follows:

   - A finite set of resources / roads $\mathcal{R}$.
   - A congestion function for each resource r of the form $c_r : \{1, 2, ...\} \to R$. The cost $c_r(k)$ is the congestion on resource/road r when there are k users.
   - A finite set of players $N = \{1, 2, ..., n\}$.
   - A finite action set of each player $\mathcal{A}_i = \mathcal{R}$. Here, each player can select any single resource $r \in \mathcal{R}$. Let $\mathcal{A} := \mathcal{A}_1 \times ... \times \mathcal{A}_n$ represent the set of joint actions.
   - A cost function for each player i of the form $J_i : \mathcal{A} \to R$ that each player seeks to minimize. The specific form of the cost function is

     $$J_i(a_i = r, a_{-i}) = c_r(|a|_r)$$

     where $|a|_r$ represents the number of player that chose resource r in the action profile a, i.e.,

     $$|a|_r = |\{j \in N : r \in a_j\}|$$

   - System cost of the form

     $$C(a) = \sum_{i \in N} J_i(a) = \sum_{r \in \mathcal{R}} |a|_r \cdot c_r(|a|_r)$$

   (a) Is the above congestion game a potential game with the potential function $C(a)$? Justify your answer.

(b) Consider the following potential function

$$\phi(a) \quad = \quad \sum_{r \in \mathcal{R}} \sum_{k=1}^{|a|_r} c_r(k)$$

Is the above congestion game a potential game with the the potential function ϕ? Justify your answer.

(c) Prove that a pure Nash equilibrium must exist in *any* parallel congestion game.

(d) Derive an anonymous tolling scheme, i.e., $t_r : \{1, 2, ...\} \to R$ for each resource $r \in \mathcal{R}$ such that the new congestion game associated with the player's cost functions

$$\tilde{J}_i(a_i = r, a_{-i}) = c_r(|a|_r) + t_r(|a|_r)$$

is a potential game with potential function $C(\cdot)$.

HINT: We looked into this problem for non-atomic congestion games in lecture.

2. Consider the following non-atomic routing game where there is one unit of mass seeking to traverse across the network.

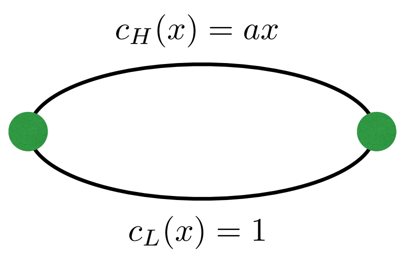

(a) Characterize the Nash flow as a function of $a \geq 0$.

(b) Characterize the optimal flow as a function of $a \geq 0$.

(c) Characterize the price of anarchy as a function of $a \geq 0$.

(d) Characterize the price of anarchy of the set of games where $a \geq 0$.

(e) Characterize the price of anarchy of the set of games where $a \geq 2$.

3. Consider the following non-atomic routing game where there is one unit of mass seeking to traverse across the network.

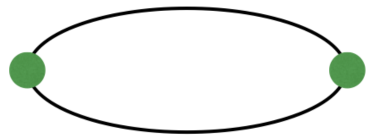

In class, we discussed how in any nonatomic congestion game with linear-affine latency functions (i.e., latency functions equal to $\ell_i(x_i) = a_i x_i + b_i$, where a_i and b_i are nonnegative constants), if all agents *act like* the cost functions are the *marginal cost functions* $\ell_i(x_i) = 2a_i x_i + b_i$ that their Nash flows are optimal; that is, their Nash flows minimize the *social cost* $C(x) = \sum_i x_i \ell_i(x_i)$. Then we used this fact to argue that road tolls may be one way to make people "think" their cost functions are the marginal cost functions, by charging them an *optimal toll* equal to $a_i x_i$ on each network link. In this homework, you will explore a related idea; here, rather than charging people a toll for using roads with a high marginal cost, we will *pay people* to use roads with a relatively low marginal cost.

Accordingly, define an *optimal subsidy* as paying each player a price of $b_i/2$ on each link (where the link has as a latency function of $\ell_i(x_i) = a_i x_i + b_i$). In other words, the cost an agent experiences on link i under an optimal subsidy is

$$J_i(x_i) = a_i x_i + b_i - \frac{b_i}{2}.$$

(a) As defined above, what are the optimal subsidies for each link of the depicted network, as a function of $a \geq 0$?

(b) What are the Nash flows on the above network *with optimal subsidies*, as a function of $a \geq 0$? Verify that these flows are optimal.

(c) What are the optimal *tolls* (i.e., each link is charged a toll of $a_i x_i$) for each link of the depicted network, as a function of $a \geq 0$? Note that the optimal tolls are actually *tolling functions* — they may depend on the amount of traffic.

(d) What are the Nash flows on the above network *with optimal tolls*? Verify that these flows are the same as what you derived in part 2.

(e) Explain why the subsidy policy of $b_i/2$ results in the same Nash flows as a toll policy of $a_i x_i$.