Game Theory
Lecture #14 — Congestion and Potential Games

Outline:

- Congestion Games
- Identical Interest Games
- Potential Games

1 Introduction

Nash equilibrium is a simple model of emergent behavior in social systems (and distributed decision systems more generally), but we have seen already that in general, a *pure* Nash equilibrium need not exist. While existence is guaranteed if we allow *mixed* strategies, these types of equilibria are computationally challenging to find and can be less predictive of societal behavior, e.g., the mixed strategy Nash equilibrium in the Hawk Dove game. The question of this lecture is the following: Are there relevant classes of games for where a pure Nash equilibrium is guaranteed to exist?

1.1 Congestion Games

We have seen congestion games informally several times so far, but in this lecture we will address them in more depth. The defining characteristic of a congestion game is that we can express the players' strategies by considering a set of *resources* or *elements* $\mathcal{R}$ and each resource $r \in \mathcal{R}$ is associated with a *congestion function* $c_r : \mathbb{N} \to \mathbb{R}$, where $c_r(k)$ denotes the congestion on resource r with $k \geq 0$ players. A player's permissible actions are defined as various subsets of $\mathcal{R}$, which make up the agent's action set $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. Lastly, the player's individual cost function given a joint action $a = (a_i, a_{-i})$ is given by

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

where $|a|_r := |j \in N : r \in a_j|$ expresses the number of agents selecting resource r. That is, the defining characteristic of a congestion game is that a player's cost on a given resource depends on

1. The identity of the resource itself, and
2. The number of other players selecting that resource, but

not the identities of the other *players* selecting the resource. To find the player's overall individual cost, we simply sum the costs of that player's selected resources.

When considering congestion games, we will often ask how "efficient" a given action profile is with respect to the following system-level objective function:

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

The term $C(\cdot)$ is commonly referred to as the social welfare which is defined as the sum of the agents' payoff. In a congestion / transportation network the social welfare equates to the total congestion in the network.

Our questions moving forward are

- Is a PNE guaranteed to exist?
- How efficient is a PNE with respect to the system-level objective $C(a)$?
- How can we influence players to improve system-level behavior if a PNE can be inefficient?

1.2 Existence of PNE in parallel-link congestion games

First, consider a simple congestion game with only two resources $\mathcal{R} = \{r_1, r_2\}$, with each of two players selecting between either resource so that $\mathcal{A}_i = \{\{r_1\}, \{r_2\}\}$ for both players $i \in \{1, 2\}$. Each resource has a congestion function $c_r : \{1, 2\} \to \mathbb{R}$; note that *any* such function is allowable here. For an example this simple, we may write out the full (abstract) cost matrix:

Player 2

		r_1	r_2
Player 1	r_1	$c_{r_1}(2),\ c_{r_1}(2)$	$c_{r_1}(1),\ c_{r_2}(1)$
	r_2	$c_{r_2}(1),\ c_{r_1}(1)$	$c_{r_2}(2),\ c_{r_2}(2)$

Cost Matrix

Note that for this simplified example, a cost matrix captures the whole pictures regardless of the specific congestion functions. Surprisingly, for *any* choice of congestion functions c_r, a pure Nash equilibria exists. We can investigate the existence of pure Nash equilibria on this example in a completely abstract way due to its simplicity, but as we shall see, these arguments generalize quite broadly.

To prove there exists a PNE in the above game, let us assume (for the sake of contradiction) that there does not: suppose no action profile in the game is a PNE. Specifically, this means that (r_1, r_1) is not a PNE. This means one of the players has a strict incentive to deviate to resource r_2; by inspection, we see that this implies $c_{r_2}(1) < c_{r_1}(2)$.

Similarly, since (r_2, r_1) is not a PNE, then $c_{r_2}(2) < c_{r_1}(1)$. Finally, since (r_2, r_2) is not a NE, then $c_{r_1}(1) < c_{r_2}(2)$. However, this directly contradicts the statement that $c_{r_2}(2) < c_{r_1}(1)$, which must be true if there is no PNE. Therefore, there must be a PNE.

The simple conclusion here is that in any parallel-link congestion game with 2 players, for any possible choice of cost functions for the links, a pure Nash equilibrium exists. However, the question remains: does this fact extend to broader classes of congestion games?

2 Potential Games

Clearly, we need more tools to successfully answer the question of equilibrium existence in general congestion games. Accordingly, in this section we will introduce the extremely powerful concept of *potential games* which will be invaluable for this direction.

2.1 Identical Interest Games

One simple class of games in which PNE are always guaranteed to exist is known as the class of "identical interest games." The defining feature of games in this class is that all players share a common utility function. Here, with player set $N = \{1, \ldots, n\}$ and player action sets $\mathcal{A}_i$, we express this common utility function simply as $U : \mathcal{A} \to \mathbb{R}$, and we assign this as the utility function for each player so that for each $i, j \in N$ and each action profile $a \in \mathcal{A}$ it holds that

$$U_i(a) = U_j(a) = U(a).$$

For concreteness, consider the following 2-player identical interest game:

Player 2

		y_1	y_2	y_3
	x_1	$1, 1$	$2, 2$	$4, 4$
Player 1	x_2	$5, 5$	$9, 9$	$-1, -1$
	x_3	$7, 7$	$8, 8$	$3, 3$

Payoff Matrix

Figure 1: An identical interest game.

Several questions arise:

1. Does there exist a PNE?
2. How "good" are the PNE?
3. Do the numbers in the matrix matter for this property?

By inspection, we see that the above game has exactly 2 PNE: (x_2, y_2) with utility 9 and (x_1, y_3) with utility 4. In this game, we have that the utility-maximizing action profile (x_2, y_2) is a Nash equilibrium; in fact, this is the case in *every* identical interest game. To see this, denote a utility-maximizing action profile as $a^* \in \arg\max_{a \in \mathcal{A}} U(a)$. Since this is a utility-maximizing action profile, it must satisfy for every i and every a'_i that

$$U_i(a_i^*, a_{-i}^*) \geq U_i(a'_i, a_{-i}^*).$$

This is simply the definition of a Nash equilibrium; that is, a^* is a PNE.

Thus, our example is not "special" — every identical interest game must have at least one PNE; namely, the action profile which maximizes the common utility function. Our example illustrates another important fact: Nash equilibria need not be unique, and there can easily be inefficient equilibria as evidenced by the action profile (x_1, y_3) with utility 4.

2.2 Shifting the Payoffs

Identical interest games lie at the core of a fascinating class of games called "potential games," which we formally introduce in the next section. The key idea of a potential game is that in a specific and detailed sense, every potential game behaves exactly like a related identical interest game. Before we introduce potential games rigorously, consider the following question. Consider the following game, constructed from the game in Figure 1 above simply by adding the constant $a \in \mathbb{R}$ to Player 1's utility in the y_2 column:

Player 2

		y_1	y_2	y_3
	x_1	$1+a, 1$	$2, 2$	$4, 4$
Player 1	x_2	$5+a, 5$	$9, 9$	$-1, -1$
	x_3	$7+a, 7$	$8, 8$	$3, 3$

Payoff Matrix

Figure 2: An identical interest game with shifted payoffs.

What are the Nash equilibria of this game? Careful inspection reveals that they are exactly the same as those of the game in Figure 1. Adding a to Player 1's payoffs in those particular action profiles changes nothing about Player 1's relative preferences between x_1, x_2, and x_3, and it clearly changes nothing about Player 2's preferences. Hence, the equilibria in Figures 1 and 2 are the same.

Unsurprisingly, there was nothing special about Player 1 or action y_1 when we performed the shift in Figure 2, and we could have done the same thing for some action of Player 2 without changing equilibria:

Here, it becomes slightly harder to see that the equilibria are unchanged, but we can convince ourselves of this fact by noticing that compared with Figure 2, only Player 2's payoffs have

		Player 2		
		y_1	y_2	y_3
Player 1	x_1	$1+a, 1$	$2, 2$	$4, 4$
	x_2	$5+a, 5+B$	$9, 9+B$	$-1, -1+B$
	x_3	$7+a, 7$	$8, 8$	$3, 3$

Payoff Matrix

Figure 3: An identical interest game with more shifted payoffs.

been shifted; thus, only Player 2's preferences have changed. However, since the same constant $B \in \mathbb{R}$ was added to each of Player 2's payoffs across the entire x_2 row, Player 2's *relative* preferences are intact and thus the set of equilibria could not have changed.

In fact, we could perform a procedure like this, change every payoff in the game, and still maintain exactly the same relative preferences and exactly the same set of equilibria:

		Player 2		
		y_1	y_2	y_3
Player 1	x_1	$1+a, 1+A$	$2+b, 2+A$	$4+c, 4+A$
	x_2	$5+a, 5+B$	$9+b, 9+B$	$-1+c, -1+B$
	x_3	$7+a, 7+C$	$8+b, 8+C$	$3+c, 3+C$

Payoff Matrix

Figure 4: An identical interest game with all payoffs shifted.

By the same arguments we made above, since we added numbers uniformly across the *columns* of Player 1's utility function and across the *rows* of Player 2's utility function, no agent's relative preferences have been modified in any way. However, by adding these numbers, we can create a new game that has no immediate or obvious connection to the original identical interest game despite the fact that it is guaranteed to have the same set of Nash equilibria. For instance, consider the game with constants selected as

$$\begin{aligned} a &= -1 \\ b &= -2 \\ c &= 1 \\ A &= -1 \\ B &= 1 \\ C &= -3, \end{aligned} \tag{1}$$

to yield the payoff matrix

Considering Figure 5, we may verify by standard techniques that the only Nash equilibria are (x_2, y_2) and (x_1, y_3) just as in the original game in Figure 1. Moreover, note that there is an exact "alignment" between the *relative* preferences in the two games. For instance, consider the action profile (x_2, y_1). In either game, if Player 1 deviates to action x_3, her

		Player 2		
		y_1	y_2	y_3
Player 1	x_1	$0, 0$	$0, 1$	$5, 3$
	x_2	$4, 6$	$7, 10$	$0, 0$
	x_3	$6, 4$	$6, 5$	$4, 0$

Payoff Matrix

Figure 5: A game derived from the identical interest game in Figure 1 by shifting payoffs according to the formula in Figure 4 and values provided in equation 1.

payoff increases by exactly 2. Likewise, if Player 2 deviates to action y_3, her payoff decreases by exactly 6.

2.3 Potential Games

In formal terms, the game in Figure 5 is called a potential game. A game with player set $N = \{1, 2, ..., n\}$, action sets $(\mathcal{A}_i)_{i \in N}$, and utility functions $U_i : \mathcal{A} \to \mathbb{R}$ is called a *potential game* if there exists a function $\phi : \mathcal{A} \to R$ such that for every action profile $a \in \mathcal{A}$, agent $i \in N$, and alternative choice $a_i' \in \mathcal{A}_i$

$$U_i(a_i', a_{-i}) - U_i(a_i, a_{-i}) = \phi(a_i', a_{-i}) - \phi(a_i, a_{-i}). \quad (2)$$

This function ϕ is called the *potential function*.

To compare with our examples above, we say that any of the games in Figures 2– 5 is a potential game with potential function equal to the common utility function $U(a)$ appearing in Figure 1.

There are several equivalent and correct ways to think about potential games:

- From a local perspective, every agent in a potential game acts as though they are playing an identical interest game in which the common utility function is ϕ. That is, the relative payoffs experienced by an agent are the same as they would be if her utility function were simply equal to ϕ.
- In a potential game, one can think of the potential function ϕ as an "imaginary" function in the background which behaves in a very particular way: any gain for an agent is a gain for the potential function.
- Every potential game is an identical interest game "in disguise," in the same way that we added/subtracted constants in Figure 4 to "disguise" the identical interest game from Figure 1.

Due to the alignment with an identical interest game, every potential game is guaranteed to have at least one pure Nash equilibrium. Specifically, the maximizer of the potential function is a pure Nash equilibrium, for precisely the same reason that in an identical interest game the maximizer of the common utility function is a pure Nash equilibrium. However, several questions remain:

- Is a NE unique in a potential game?
- How do you evaluate whether or not a game is a potential game?
- Is an identical interest game a potential game?
- Should the potential function be related to the system objective?

2.4 Identifying Potential Games

To determine whether a given game is a potential game, one must determine whether that game has a potential function. That is, does some function (a potential function) exist which is aligned with all possible agent unilateral deviations? If a potential function exists, then the game is a potential game — otherwise, it is not.

2.4.1 Prisoner's Dilemma

As a simple initial example, we may ask if the Prisoner's Dilemma is a potential game?

	L	R
T	$3, 3$	$0, 4$
B	$4, 0$	$1, 1$

If this is a potential game, then a potential function exists — and the utility functions provide enough information to derive it. However, if we attempt to derive a potential function and then fail, it must be that no potential function exists — and it is thus not a potential game. But how do we derive a potential function?

The first thing to note is that a potential function is never unique: since all that matters about the potential function is agent deviations, we can add any constant to a potential function and the resulting function is still a potential function. Thus, when we attempt to derive a potential function for a given game, we can "seed" the potential function at any number we like. We demonstrate this on the Prisoner's Dilemma, with the payoff matrix on the left and the candidate potential function on the right; note that we have "seeded" the potential function with the value 0 for the (T, L) action profile:

	L	R
T	$3, 3$	$0, 4$
B	$4, 0$	$1, 1$

Prisoner's Dilemma

	L	R
T	0	?
B	?	?

Candidate Potential Function

Since changes in the potential function must exactly match changes in individual agent utility functions, we can now derive the necessary potential function values for both the (B, L) and (T, R) action profiles. That is, according to (2), we solve $\phi(B, L) - \phi(T, L) = U_1(B, L) - U_1(T, L)$ for $\phi(B, L)$ and we solve $\phi(T, R) - \phi(T, L) = U_2(T, R) - U_2(T, L)$ for $\phi(T, R)$:

	L	R
T	3, 3	0, 4
B	4, 0	1, 1

Prisoner's Dilemma

	L	R
T	0	1
B	1	?

Candidate Potential Function

Having filled in all but 1 box in the potential function grid, it now remains to be seen if it is possible to define $\phi(B, R)$ in a way that is consistent with both agent utility functions. That is, we seek a single value for the lower-right box which satisfies *both* equations

$$\phi(B, R) - \phi(T, R) = U_1(B, R) - U_1(T, R) \quad (3)$$
$$\phi(B, R) - \phi(B, L) = U_2(B, R) - U_2(B, L). \quad (4)$$

Happily, by setting $\phi(B, R) = 2$, both equations are satisfied, which verifies that this is a potential game with potential function equal to:

	L	R
T	3, 3	0, 4
B	4, 0	1, 1

Prisoner's Dilemma

	L	R
T	0	1
B	1	2

Potential Function

2.4.2 If a game has a PNE, is it a potential game?

As we stated earlier, every potential game has at least one PNE. Is the converse true? Let us consider the following game which has one PNE (T, L), and check if it is a potential game.

	L	R
T	1, 1	3, 0
B	0, 0	0, 1

	L	R
T	0	?
B	?	?

Candidate Potential Function

As with the Prisoner's Dilemma, we can seed the potential function with any number we want, so in the above, we have chosen $\phi(T, L) = 0$. Given this seed, we compute the values of $\phi(B, L) = U_1(B, L) - U_1(T, L) = -1$ and $\phi(T, R) = U_2(T, R) - U_2(T, L) = -1$:

	L	R
T	1, 1	3, 0
B	0, 0	0, 1

	L	R
T	0	−1
B	−1	?

Candidate Potential Function

Finally, we can check if there exists a value for $\phi(B, R)$ which is consistent with unilateral deviations of *both* Player 1 and Player 2:

	L	R
T	0	−1
B	−1	−4

Consistent with Player 1

	L	R
T	0	−1
B	−1	0

Consistent with Player 2

Since we cannot select a single value for $\phi(B, R)$ that is consistent with both players deviating (B, R), there definitively exists no potential function, and this is thus not a potential game.

2.5 Approaches for identifying potential game structure

In general, identifying whether or not a game is a potential game is hard and there are no "easy" ways to verify this property. The two most commonly applied approaches include the following:

(i) Guess and Check: In this first approach, you merely need to guess the potential function and verify whether or not the property holds. For smaller games, as discussed above, you can provide an intelligent guess where you initialize the potential function value of some states and iteratively define the value of the potential function for other action profiles.

(ii) Inspect Unilateral Deviation Cycles: In this section approach, we consider cycles of action profiles of the form

$$(a_i, a_j, \cdot) \to (a_i', a_j, \cdot) \to (a_i', a_j', \cdot) \to (a_i, a_j', \cdot) \to (a_i, a_j, \cdot)$$

where each subsequent joint action is the result of a unilateral deviation, i.e., the first transition results in player i transitioning from a_i to a_i' while all other players stay fixed. If we explore the change in utility of the deviating player for each of these transitions we have

$$\begin{aligned}
\Delta_1 &= U_i(a_i', a_j, \cdot) - U_i(a_i, a_j, \cdot) = \phi(a_i', a_j, \cdot) - \phi(a_i, a_j, \cdot) \\
\Delta_2 &= U_j(a_i', a_j', \cdot) - U_j(a_i', a_j, \cdot) = \phi(a_i', a_j', \cdot) - \phi(a_i', a_j, \cdot) \\
\Delta_3 &= U_i(a_i, a_j', \cdot) - U_i(a_i', a_j', \cdot) = \phi(a_i, a_j', \cdot) - \phi(a_i', a_j', \cdot) \\
\Delta_4 &= U_j(a_i, a_j, \cdot) - U_j(a_i, a_j', \cdot) = \phi(a_i, a_j, \cdot) - \phi(a_i, a_j', \cdot)
\end{aligned}$$

where the second equality holds if the game is a potential game with potential function ϕ. Accordingly, what is $\Delta_1 + \Delta_2 + \Delta_3 + \Delta_4 =$? Well, if the game is a potential game then the answer must be 0 for any cycle of unilateral deviations. Checking that all cycles of length four satisfies this constraint proves the result. Recall example:

	L	R
T	1, 1	3, 0
B	0, 0	0, 1

Note that counterclockwise cycle starting at (T, L) has sum

$$(0-1)+(1-0)+(3-0)+(1-0) \neq 0.$$

Hence, this game is not a potential game.

While these represent two common approaches to evaluate whether or not a game is a potential game, in this class we will typically rely on the guess and check approach to verify whether or not a potential function exists.

3 Revisiting Congestion Games

We started this lecture by reviewing the class of congestion games, which has application to both traffic management and control and numerous other resource allocation problems relevant to engineering based systems. The goal of this lecture was to build tools to evaluate whether or not an interesting class of games have guaranteed existence of a pure Nash equilibrium. While we showed that for a class of two-player congestion games, it remained an upon question as to whether or not that property held true to arbitrary congestion games. The following theorem resolves this question.

Theorem 3.1 (Monderer and Shapley, 1996) *Consider any congestion game as defined in Section 1.1. The congestion game is a potential game with potential function*

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

Proof 3.1 *In order to show that a congestion game is in fact a potential game, we need to show that for any action profile $a \in \mathcal{A}$, agent $i \in N$, and alternative choice $a_i' \in \mathcal{A}_i$ the following equality is satisfied*

$$J_i(a_i', a_{-i}) - J_i(a_i, a_{-i}) = \phi(a_i', a_{-i}) - \phi(a_i, a_{-i}).$$

For simplicity, assume $\mathcal{A}_i = \mathcal{R}$, i.e., we have a parallel network where any agent can select any single edge from the set $\mathcal{R}$. Let $a_i = r$, $a_i' = r'$, where $r' \neq r$. Given any a_{-i}, the cost function of player i satisfies

$$\begin{aligned} J_i(a_i = r, a_{-i}) &= c_r(|a|_r) \\ J_i(a_i' = r', a_{-i}) &= c_{r'}(|a|_{r'} + 1) \end{aligned}$$

where the $+1$ comes from the fact that there is now 1 extra player on the resource r' when player i switches to that one. If we inspect our potential function, we have that

$$\begin{aligned}
\phi(a'_i = r', a_{-i}) - \phi(a_i = r, a_{-i}) &= \sum_{r\in\mathcal{R}} \sum_{k=1}^{|a'|_r} c_r(k) - \sum_{r\in\mathcal{R}} \sum_{k=1}^{|a|_r} c_r(k) \\
&= \sum_{k=1}^{|a'|_{r'}} c_{r'}(k) - \sum_{k=1}^{|a|_{r'}} c_{r'}(k) + \sum_{k=1}^{|a'|_r} c_r(k) - \sum_{k=1}^{|a|_r} c_r(k) \\
&= c_{r'}(|a|_r + 1) - c_r(|a|_r) \\
&= J_i(a') - J_i(a)
\end{aligned}$$

Hence the game is a potential game with potential function ϕ.

The above 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 latency functions, or the agents' action sets. This is the first significant class of games where we demonstrated the existence of pure Nash equilibria.

4 Conclusions

This 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, note the difference in structure between the global objective and the potential function, i.e.,

$$\begin{aligned}
C(a) &= \sum_{r\in\mathcal{R}} |a|_r \cdot c_r(|a|_r), \\
\phi(a) &= \sum_{r\in\mathcal{R}} \sum_{k=1}^{|a|_r} c_r(k).
\end{aligned}$$

What is the implication of these functions being different? Are there mechanisms a system operator can employ to make these functions more aligned? These questions will be the focus of the next lecture.

5 Exercises

1. Consider the following games from a previous lecture

    - BoS:

	B	S
B	$2, 1$	$0, 0$
S	$0, 0$	$1, 2$

- Stag hunt:

	Stag	Hare
Stag	$2, 2$	$0, 1$
Hare	$1, 0$	$1, 1$

- Typewriter:

	Alt	Std
Alt	$3, 3$	$0, 0$
Std	$0, 0$	$1, 1$

Are any of the three games a potential game? If so, which one(s)? Further, what is the potential function?

2. Complete the payoffs in the following game so that it is a potential game:

	L	C	R
T	1,?	2,?	1,?
M	?,4	5,3	2,?
B	1,?	6,?	7,?

3. This question studies a learning algorithm known as Fictitious Play. At any given time $t \in \{1, 2, \dots\}$, the Fictitious Play algorithm proceeds as follows:

   - **Bookkeeping:** Each player $i \in N$ maintains the empirical frequency of each player's previous actions. Let $q_j(t) \in \Delta(\mathcal{A}_j)$ represent the empirical frequency of player j's action over the stages $k = 1, 2, ..., t-1$. For example, if player j selected $a_j(1) = H$, $a_j(2) = H$, $a_j(3) = T$, then $q_j(4) = (2/3, 1/3)$ where the $(2/3)$ corresponds to empirical frequency of H and the $(1/3)$ corresponds to the empirical frequency of T.
   - **Assumptions:** At any given time $t > 0$, each player assumes that all other players $j \neq i$ will select an action *independently* using the strategy $q_j(t)$.
   - **Best Response:** Each player selects an action at time t seeking to maximize his expected utility given his belief that each player $j \neq i$ will be selecting her action independently according to strategy $q_j(t)$, i.e.,

   $$a_i(t) = BR_i(\{q_j(t)\}_{j \neq i})$$

   - **Repeat:**

   Fictitious play has several nice convergence results. Two such results are as follows:

   - Empirical frequency of play converges to a Nash equilibrium in any potential game;

- Empirical frequency of play converges to the optimal security strategies in any zero-sum game.

This question will explore this result on the following two player game:

	a_2	b_2
a_1	$5, 2$	$0, 4$
b_1	$0, 4$	$2, 6$

(a) Does a pure Nash equilibrium exist? If so, what are the pure Nash equilibria?

(b) Is the above game a potential game? If so, what is the potential function?

(c) Suppose the game defined above is repeated each period $t = 1, 2, ...$ and each player uses the learning algorithm Fictitious Play to select their respective action at each time t. If the initial action is $a(1) = (b_1, a_2)$ ***derive*** the ensuing action trajectory $a(2), a(3), ..., a(k)$ for any $k \geq 1$.