

5.1 Introduction

Last week we talked about computing a Nash in general two-player games, and gave exponential time algorithms but also PPAD-hardness. For the next few weeks we'll talk about other kinds of games, other kinds of equilibria, and their computability / existence. Today, instead of “zooming out” to more general notions of equilibria, we're going to “zoom in” to the special case of pure Nash equilibria. In what kinds of games *do* they exist? What about computing them?

5.2 Atomic Routing Games

We're going to spend much of today on a special class of games known as *atomic routing games*. These games can have more than two players, but there's much more structure than a general game. This will let us prove strong things about them.

An atomic routing game consists of the following.

- Directed graph $G = (V, E)$
- Edge cost functions $c_e : \mathbb{R} \rightarrow \mathbb{R}$ for each $e \in E$ (think of $c_e(x)$ as the delay along edge e if there is x traffic using it).
- k players, where each player i has some source-sink pair $(s_i, t_i) \in V \times V$ associated with it.
- Player i has strategy set $S_i = \{s_i \rightarrow t_i \text{ paths}\}$. Let $S = S_1 \times S_2 \times \dots \times S_k$.
- Let $f = (P_1, P_2, \dots, P_k) \in S$ be some strategy profile (often called a *flow*). Let $f_e = |\{i : e \in P_i\}|$ be the number of players using edge e in f . Then the cost to player i of strategy profile f is

$$C_i(f) = \sum_{e \in P_i} c_e(f_e).$$

Note that unlike the last few lectures where we've been talking about utility-maximization games, this is a cost-minimization game. But of course we could just let a player's utility be the negative of its cost and everything would work as you expect.

Theorem 5.2.1 (Rosenthal) *Every atomic routing game has at least one pure Nash equilibrium.*

Proof: We're going to define a potential function which will map strategy profiles to real numbers. For any $f = (P_1, P_2, \dots, P_k)$, let

$$\Phi(f) = \sum_{e \in E} \sum_{j=1}^{f_e} c_e(j).$$

Note that this is similar to, but not the same as, what you might naturally expect as the potential: the total cost

$$\sum_{i=1}^k C_i(f) = \sum_{i=1}^k \sum_{e \in P_i} c_e(f_e) = \sum_{e \in E} f_e c_e(f_e).$$

This potential function Φ has some nice structure (which is why we chose it). To see this, suppose that player i deviates from P_i to some other path \hat{P}_i , giving a new strategy profile $\hat{f} = (P_1, P_2, \dots, P_{i-1}, \hat{P}_i, P_{i+1}, \dots, P_k)$. Then the change in cost to player i is

$$\begin{aligned} C_i(\hat{f}) - C_i(f) &= \sum_{e \in \hat{P}_i} c_e(\hat{f}_e) - \sum_{e \in P_i} c_e(f_e) \\ &= \sum_{e \in \hat{P}_i \setminus P_i} c_e(f_e + 1) + \sum_{e \in \hat{P}_i \cap P_i} c_e(f_e) - \sum_{e \in \hat{P}_i \cap P_i} c_e(f_e) - \sum_{e \in P_i \setminus \hat{P}_i} c_e(f_e) \\ &= \sum_{e \in \hat{P}_i \setminus P_i} c_e(f_e + 1) - \sum_{e \in P_i \setminus \hat{P}_i} c_e(f_e) \end{aligned}$$

The change in potential is

$$\begin{aligned} \Phi(\hat{f}) - \Phi(f) &= \sum_{e \in E} \sum_{i=1}^{\hat{f}_e} c_e(i) - \sum_{e \in E} \sum_{i=1}^{f_e} c_e(i) \\ &= \sum_{e \in E} \left(\sum_{i=1}^{\hat{f}_e} c_e(i) - \sum_{i=1}^{f_e} c_e(i) \right) \\ &= \sum_{e \in \hat{P}_i \setminus P_i} c_e(f_e + 1) - \sum_{e \in P_i \setminus \hat{P}_i} c_e(f_e). \end{aligned}$$

So the change in potential is exactly equal to the change in cost of the player who deviates! This is extraordinarily useful, since it connects a “global” property (the potential) to a “local” decision (whether to deviate).

Let $f = \arg \min_{f' \in S} \Phi(f')$, i.e., a flow with minimum potential (such a flow must exist since S is finite). I claim that f is a pure Nash equilibrium. Based on the above discussion, this is easy to see. Suppose that some player i deviates from the path P_i it uses in f to some other path \hat{P}_i , giving a strategy profile \hat{f} . Then

$$C_i(\hat{f}) - C_i(f) = \Phi(\hat{f}) - \Phi(f) \geq 0,$$

and thus player i has no incentive to deviate. Thus f is a pure Nash equilibrium. ■

5.2.1 Potential Games

The only thing that we actually used in the previous argument was the existence of a potential function. So consider an arbitrary game with k players, strategies S_i for each $i \in [k]$ with $S = S_1 \times S_2 \times \dots \times S_k$, and costs $C_i : S \rightarrow \mathbb{R}$. This is a *potential game* if there exists some function $\Phi : S \rightarrow \mathbb{R}$ such that

$$\Phi(s_{-i}, s'_i) - \Phi(s) = C_i(s_{-i}, s'_i) - C_i(s)$$

for all $i \in [k]$, for all $s \in S$, and for all $s'_i \in S_i$ (where recall from the first lecture that (s_{-i}, s'_i) refers to the strategy profile obtained by replacing the strategy used by player i in s with s'_i).

Atomic routing games are potential games (we just proved this!), but so are many other interesting games.

Theorem 5.2.2 *Every potential game has a pure Nash equilibrium.*

Proof: This is the same argument as for atomic routing games. The strategy profile $s = \arg \min_{s \in S} \Phi(s)$ is a pure Nash equilibrium, since if any player had incentive to deviate there would be a strategy profile with smaller potential. ■

Interestingly, potential games also have the property that a simple and “natural” algorithm actually finds a pure Nash equilibrium (albeit possibly in exponential time). Consider the following process: while s is not a pure Nash equilibrium, pick some arbitrary player i and arbitrary beneficial deviation s'_i and let $s \leftarrow (s_{-i}, s'_i)$. This process is known as “best-response dynamics” (at least in this context).

It’s obvious that best-response dynamics end at a pure Nash equilibrium *if* they halt. If there is no pure Nash equilibrium then clearly they will never halt, and in general games they might not halt even if there is a pure Nash equilibrium.

Theorem 5.2.3 *In every potential game, best response dynamics always terminate (at a pure Nash equilibrium).*

Proof: Every iteration involved a beneficial deviation, so by the definition of the potential function Φ we know that Φ decreases in every iteration. Thus (since S is finite) we will eventually halt at a minimizer of Φ . ■

Note that this might take exponential time, but at least a natural algorithm will get to a pure Nash!

5.3 Hierarchy of Equilibria

So now we have the following picture: pure Nash equilibria might not exist, but in at least some interesting games they do exist and can be found via a natural algorithm. Containing pure Nash are mixed Nash equilibria, which always exist but can be difficult to find (PPAD-complete). Is there any notion of equilibrium which always exists, *and* can always be found? And what if we only allow “simple” or “natural” algorithms?

Surprisingly, the answer is yes! We’re going to see two additional notions of equilibria: *correlated equilibria* contain all mixed Nash equilibria (so always exist) and can be computed using “somewhat simple” algorithms, and *coarse correlated equilibria* contain all correlated equilibria (so always exist) and can be computed using pretty simple algorithms! I’ll spend the rest of the day defining and discussing these equilibria, and then we’ll talk about computing them on Thursday (and possible next Tuesday).

5.3.1 Correlated Equilibria

Consider a cost-minimization game with k players $[k]$ and strategy sets $\{S_i\}_{i \in [k]}$, with $S = S_1 \times S_2 \times \dots \times S_k$ and cost functions $c_i : S \rightarrow \mathbb{R}$ for each player $i \in [k]$. Let’s define a mixed Nash again (with slightly different notation) to compare and contrast with these other equilibria. I’m going to define these with respect to cost-minimization, but of course we could switch to utility maximization.

Definition 5.3.1 *Let σ_i be a distribution over S_i for all $i \in [k]$. Let $\sigma = \sigma_1 \times \sigma_2 \times \dots \times \sigma_k$ be the product distribution over S defined by the individual player distributions. Then σ is a mixed Nash*

equilibrium if

$$\mathbf{E}_{s \sim \sigma}[c_i(s)] \leq \mathbf{E}_{s \sim \sigma}[c_i(s_{-i}, s'_i)]$$

for all $i \in [k]$ and for all $s'_i \in S_i$.

Now we can define correlated equilibria.

Definition 5.3.2 Let σ be a distribution over $S = S_1 \times \dots \times S_k$. Then σ is a correlated equilibrium if

$$\mathbf{E}_{s \sim \sigma}[c_i(s)|s_i] \leq \mathbf{E}_{s \sim \sigma}[c_i(s_{-i}, s'_i)|s_i]$$

for all $i \in [k]$ and for all $s_i, s'_i \in S_i$.

This might be a little hard to interpret, in part because the notation isn't great (but is traditional, so I'm sticking with it). At a very high level, it allows for σ to be an arbitrary distribution over S rather than a product distribution (as in Nash), so players pick actions in a "correlated" way rather than independently. The conditioning inside of the expectation basically says that each player i has no incentive to deviate *even if* I know that the particular draw from σ means that I play some action s_i (which, since there are correlations between players, means that I might now know a lot about what other players probably drew in this draw from σ).

The usual interpretation (due to Robert Aumann), which allows for more intuition about these objects, involves changing how we think of a game being played. Until now, we've talked about every player knowing all of the distributions, but then "closing our eyes" and simultaneously revealing our actions to all of the players. But let's think of a different setup.

- There is some "trusted third party" U , and σ is publicly known.
- U samples some strategy profile s from σ (but keeps s secret).
- For each player $i \in [k]$, U *privately* tells s_i to i .
- Now each player i decides whether to play s_i or to deviate to some other strategy.

Note that each player i knows σ and knows s_i , so can figure out their expected cost if no one deviates: $\mathbf{E}_{s \sim \sigma}[c_i(s)|s_i]$. On the other hand, player i can also figure out their expected cost if they deviated to s'_i but no one else deviates: $\mathbf{E}_{s \sim \sigma}[c_i(s_{-i}, s'_i)|s_i]$. So the correlated equilibrium condition says that no one has incentive to deviate when the game is played this way!

Note that every Nash equilibrium is a correlated equilibrium, since if σ is a product distribution then the conditioning does not affect anything. But there can be correlated equilibria which are not Nash.

Let's see a classical and intuitive example of this to see why it makes sense. Consider the *stoplight game*, modeling the intersection of two roads (note that since the numbers are costs, negative is good and positive is bad)

	stop	go
stop	(0,0)	(0,-1)
go	(-1,0)	(20, 20)

Clearly there are two pure Nash equilibria: (go, stop) and (stop, go). But these are both kind of unfair: if we stay at one of these equilibria (which is the whole point of the notion of equilibrium, after all) then all the traffic on one road will get to go through the intersection, but none of the traffic on the other road will get to go through. In order to be fair, we'd like for each player to basically spend 1/2 the time stopping and 1/2 the time going. If we try to create a Nash equilibrium that looks like this, then we fail: each player randomizing 50/50 independently means that there's a 1/4 chance of a collision! This clearly gives enough negative utility that it is not a Nash equilibrium: both players actually have incentive to deviate to "stop".

In other words, we want the distribution

$$\sigma(\text{stop}, \text{go}) = 1/2 \qquad \sigma(\text{go}, \text{stop}) = 1/2.$$

This is definitely not a product distribution, and so is not a Nash equilibrium. But it is a correlated equilibrium! Suppose we're one of the players. If the stoplight (the trusted third party) tells us to "go", then we *know* (since we know σ) that the other player was told to "stop". And if we're told to "stop", then we know that the other player was told to "go". In either case, if the other player does what they're told then the best thing for us to do is what we're told.

5.3.2 Coarse Correlated Equilibria

We're going to see one more definition of equilibrium.

Definition 5.3.3 *Let σ be a distribution over $S = S_1 \times \dots \times S_k$. Then σ is a coarse correlated equilibrium if*

$$\mathbf{E}_{s \sim \sigma} [c_i(s)] \leq \mathbf{E}_{s \sim \sigma} [c_i(s_{-i}, s'_i)]$$

for all $i \in [k]$ and for all $s'_i \in S_i$.

Like correlated equilibria we are now allowed to have non-product distributions, but like Nash equilibria there is no conditioning on our action. The standard interpretation is the same setup as correlated equilibria with their being a trusted third party, but now each player i has to decide whether to deviate *before* being told s_i .

It's not hard to prove (good exercise to do at home!) that any correlated equilibrium is a coarse correlated equilibrium. Intuitively, if we have no incentive to deviate no matter what we're told to play, then we don't have incentive to deviate even if we're not told what to play. The converse is definitely not true, though: there are coarse correlated equilibria that are not correlated equilibria. Even if we don't have incentive to deviate if we're not told what to play, we might have incentive to deviate if there are particular actions that we're told to play.