

Lecture 3

Lecturer: Bo Waggoner

Scribe: Aaron Roth, Bo Waggoner

When do Best Response Dynamics Converge?

In this lecture, we ask how far we can go beyond congestion games while still being certain that best response dynamics will converge to pure strategy Nash equilibria. (Recall that if BRD converges, it is necessarily to a pure strategy Nash equilibrium, so the question is really just when do they converge.) We'll do this by studying a couple of games (outside of the class of congestion games), and proving that best response dynamics converges, and then back out exactly what we needed to make the proof work.

The first game we study will look *almost* like a congestion game (in that there are still players and facilities), but is not¹ since the costs of each facility depend not just on how many people are playing on it, but on *which* players are playing on it.

Definition 1 *In a load balancing game on identical machines, there are n players $i = 1, \dots, n$. Each player i has a job of size $w_i > 0$. These jobs must each be scheduled on one of m identical machines F . The action space of the game is $A_i = F$ for each player.*

After each player i chooses a machine a_i , the load of machine $j \in F$ is $\ell_j(a) = \sum_{i:a_i=j} w_i$. The cost of player i is the load of the machine he plays on: $c_i(a) = \ell_{a_i}(a)$.

Theorem 2 *Best response dynamics converge in load balancing games on identical machines.*

Corollary 3 *Load balancing games on identical machines have pure strategy Nash equilibria*

Proof We use a variation of our potential function argument (but need a new potential function).

Define $\phi(a) = \frac{1}{2} \sum_{j=1}^m \ell_j(a)^2$. Suppose player i switches from machine j to machine j' . Then we have:

$$\begin{aligned} \Delta c_i(a) &\equiv c_i(j', a_{-i}) - c_i(j, a_{-i}) \\ &= \ell_{j'}(a) + w_i - \ell_j(a) \\ &< 0 \end{aligned}$$

Similarly, we have:

$$\begin{aligned} \Delta \phi(a) &\equiv \phi(j', a_{-i}) - \phi(j, a_{-i}) \\ &= \frac{1}{2} ((\ell_{j'}(a) + w_i)^2 + (\ell_j(a) - w_i)^2 - \ell_{j'}(a)^2 - \ell_j(a)^2) \\ &= \frac{1}{2} (2w_i \ell_{j'}(a) + w_i^2 - 2w_i \ell_j(a) + w_i^2) \\ &= w_i (\ell_{j'}(a) + w_i - \ell_j(a)) \\ &= w_i \cdot \Delta c_i(a) \\ &< 0 \end{aligned}$$

Note that unlike in congestion games, the change in potential function is not equal to the change in player cost when player's make unilateral deviations. Nevertheless, it decreases with every better-response deviation, and because it is always non-negative (and because there are only finitely many action profiles), this process must eventually halt. ■

Not all games in which best response dynamics converge need to "look like" congestion games.

¹It is a *weighted* congestion game, which in general have very different properties.

Definition 4 The *Red State/Blue State game* is played on a graph $G = (V, E)$.

1. The players are vertices $P = V$.
2. Each edge $e = (i, j) \in E$ has weight w_e
3. Actions $A_i = \{-1, 1\}$ (read {red, blue})
4. $u_i(a) = \sum_{e=(i,j) \in E} w_e \cdot a_i \cdot a_j$.

In other words, everyone picks an affiliation, and you obtain utility equal to the weight of partners who pick the same affiliation as you, and disutility equal to the weight of partners who pick a different affiliation as you.

Theorem 5 The Red-State/Blue-State game always has a pure strategy Nash equilibrium.

Proof As before, we prove that best response dynamics converges by exhibiting a potential function. Define:

$$\phi(a) = - \sum_{e=(i,j) \in E} w_e a_i a_j$$

Now consider a change in action made by player i from a_i to $-a_i$. We have:

$$\begin{aligned} \Delta u_i &= \sum_{j:(i,j) \in E} w_{i,j} \cdot (-a_i) \cdot a_j - \sum_{j:(i,j) \in E} w_{i,j} \cdot a_i \cdot a_j \\ &= -2 \sum_{j:(i,j) \in E} w_{i,j} \cdot a_i \cdot a_j \end{aligned}$$

Similarly, consider how the potential changes when player i changes her action:

$$\begin{aligned} \Delta \phi &= - \sum_{j:(i,j) \in E} w_{i,j} \cdot (-a_i) \cdot a_j - \left(- \sum_{j:(i,j) \in E} w_{i,j} \cdot a_i \cdot a_j \right) \\ &= 2 \sum_{j:(i,j) \in E} w_{i,j} \cdot a_i \cdot a_j \\ &= -\Delta u_i \end{aligned}$$

As before, this is sufficient to prove convergence: Each time a player improves her utility, the potential function strictly decreases, so best-response dynamics can never cycle (if it cycled back to a previous action profile, the potential would have to decrease and then increase again). There are only finitely many action profiles, so it must halt at some profile where no player can improve: a pure-strategy equilibrium. ■

Now that we have seen three proofs in this template, we can ask: can we characterize exactly those games in which best response dynamics converge? We have seen this class extends beyond congestion games, and similarly, beyond games in which we can define a potential function that changes *exactly* has the best response player's utility changes (e.g. the load balancing game). What do we need out of a potential function to make the proof work?

Definition 6 A function $\phi : A \rightarrow \mathbb{R}_{\geq 0}$ is an **exact potential function** for a game G if for all $a \in A$, all i , and all $a_i, b_i \in A_i$:

$$\phi(b_i, a_{-i}) - \phi(a_i, a_{-i}) = c_i(b_i, a_{-i}) - c_i(a_i, a_{-i})$$

As we saw, congestion games and the red-state blue-state games have exact potential functions, and having an exact potential function is sufficient for best response dynamics to converge. However, it's not clear that it is necessary – in particular, we did not exhibit one for the load balancing game.

Definition 7 $\phi : A \rightarrow \mathbb{R}_{\geq 0}$ is an **ordinal potential function** for a game G if for all $a \in A$, all i , and all $a_i, b_i \in A_i$:

$$(c_i(b_i, a_{-i}) - c_i(a_i, a_{-i}) < 0) \Rightarrow (\phi(b_i, a_{-i}) - \phi(a_i, a_{-i}) < 0)$$

i.e. the change in cost is always equal in sign to the change in potential.

Note: If the game is defined with a utility function u_i for each player instead of a cost c_i , then we can just take $c_i = -u_i$ in the above definitions.

We will now show that ordinal potential functions exactly characterize those games in which best response dynamics is guaranteed to converge – i.e. our proof technique is without loss of generality.

Theorem 8 *Best response dynamics is guaranteed to converge in a game G if and only if the game has an ordinal potential function.*

Proof We already know how to show that having an ordinal potential function is sufficient for best response dynamics converges exactly – this is the template proof we have applied 3 times now.

To prove that its existence is necessary, consider the following graph $G = (V, E)$:

1. Let each $a \in A$ be a vertex in the graph: i.e. $V = A$.
2. For each pair of vertices $a, b \in V$, add a directed edge (a, b) if it is possible to get to get from b to a via a best response move – i.e. if there is some index i such that $b = (b_i, a_{-i})$, and $c_i(b_i, a_{-i}) < c_i(a)$.

Note that best response dynamics can be viewed as traversing this graph, starting at some arbitrary vertex a , and then traversing the graph along its edges (which it can do breaking ties arbitrarily). The Nash equilibria are exactly the sinks in this graph (in which no player can make a best response move). If best response dynamics always converges, it must be that the graph has no cycles.

In this case, we construct an ordinal potential function ϕ as follows. Since the graph has no cycles, there must be reachable from every state a some some sink s (i.e. a pure strategy Nash equilibrium). For each state a , let $\phi(a)$ denote the length of the longest finite path in G from a to any sink s . The property we require is that $\phi(b) < \phi(a)$ for any pair of vertices (a, b) with an edge $a \rightarrow b$. But observe that by definition, if there is an edge $a \rightarrow b$, then $\phi(a) \geq \phi(b) + 1$ (since there is at the very least a path that first goes to b , and then takes the longest path from b to a sink), which completes the proof. ■