Approximate Nash Equilibrium for Spatial Line Competition

Hussin Suleiman

Acknowledgements:

This winter, under the mentorship of Louis-Roy Langevin, my partner Henry Nguyen and I have discovered a fascinating game theoretic problem, which is simple in its statement yet has proven to be very challenging to solve. I am grateful for having had the chance to work and share ideas with Louis and Henry, who have been wonderful to work with throughout the semester. I also thank the DRP organizers for making such an experience possible. I have thoroughly enjoyed working on the problem and discussing ideas, while having a glimpse of what the life of a graduate student doing research looks like. In this write-up, I will share what we have learned and what we have tried to approach the problem. I hope that you, dear reader, will be as fascinated by the problem as we have been by the time you finish reading this write-up.

Abstract

We started our journey by reading the paper *“Nearly Tight Bounds on Approximate Equilibria in Spatial Competition on the Line”* by U. Bhaskar and S. Pyne from the Tata Institute of Fundamental Research, published only last year. This paper introduces a game for which one cannot find a pure Nash-equilibrium in general and proves a lower bound for the best possible *“approximation”* of a Nash-equilibrium (we will formally define what is meant by approximation when introducing the problem). The paper also shows that this lower bound can be attained for three players, but whether it can be attained for more than three players remains an open question. So, we investigated whether the lower bound can be attained for any number of players. We first introduce the problem and the results from the paper, before presenting ideas and approaches to solve the problem.

Introduction

Consider three political parties trying to win the presidential election. Each party, based on the political convictions of the population, wants to choose the position that will maximize its proportion of the votes. For example, if most people in the population favor a left political ideology, the three parties would have the incentive to adopt a more extreme left position to please the largest number of voters. However, if the most left-extremist party, call it A, is too extreme, the second most extremist party, call it B, would have incentive to adopt a position close to A to steal as many voters from him as possible, which would be bad for party A. On the other hand, if party A is not left extremist, then the other two parties would have an incentive to be more extremist than party A to convince more people to vote for them, which would also be bad for party A. So, as we can see, choosing a strategic position for each player is more delicate than one might think.

A natural question is whether this problem can be modeled and solved using mathematical tools. In 1929, Hotelling introduced a framework representing voters as a mass distribution over the unit interval [0,1] and analyzed where political parties — now viewed as players — should position themselves to attract the most voters. Focusing on the case of two players, Hotelling observed that each player has an incentive to move closer to the other and attract the voters between the players. If both players adopt the same position, they split the voter base evenly and neither has a reason to move, since any deviation would result in losing votes. This situation is a **Nash equilibrium** — a state where no player can improve their outcome unilaterally. However, political parties rarely adopt identical positions. To reflect this, we assume that no two players occupy the exact same position. Under this condition, a Nash equilibrium does not always exist for two players, but one can get arbitrarily close to an equilibrium by moving the two players towards each other. This insight leads to a broader question: how can we approximate Nash equilibria in this model when more than two players are involved?

Let us now formalize the problem statement. We are given a positive and integrable function f over the interval $[0,1]$ with $\int_0^1 f(x)dx = 1$ to represent the distribution of voters. Then, let n be the number of players, and for $1 \leq i \leq n$, define x_i to be the position of the i-th player on the interval $[0,1]$, when reading from left to right. Since we assume that no two players coincide, we have that $0 \leq x_1 < x_2 <\ldots< x_n \leq 1$. Then, the proportion of voters for a given player, which we now refer to as utility, is the area under the curve of f and above the points on the interval $[0,1]$ closest to the player. More formally, for any $1 \leq i < j \leq n$, we let $x_{i,j} = \frac{x_i+x_j}{2}$ be the midpoint of the positions of players i and j, and we denote by $U(i,(x_1,x_2,\ldots,x_n))$ the utility of player i given the positions $x_1, x_2, \ldots, x_n$ of the players. Then, we have that:

- $U\big(1,(x_1,x_2,\ldots,x_n)\big) = \int_0^{x_{1,2}} f(x)dx$
- $U\big(n,(x_1,x_2,\ldots,x_n)\big) = \int_{x_{(n-1),n}}^1 f(x)dx$
- $U\big(i,(x_1,x_2,\ldots,x_n)\big) = \int_{x_{(i-1),i}}^{x_{i,(i+1)}} f(x)dx \ \forall 2 \leq i < n$

Now, we need to define what it means to *"approximate"* a Nash-equilibrium. Formally speaking, given some $\varepsilon \geq 0$, we define an *additive ε Nash-equilibrium* as a strategy profile $(x_1, x_2, \ldots, x_n)$ for the players such that no player can increase his utility by more than ε if all other players maintain their positions. Note that if $\varepsilon = 0$, then we have a perfect Nash-equilibrium. The smaller ε is, the closer to a Nash-equilibrium we are. The problem statement is therefore the following: What is the smallest possible $\varepsilon \geq 0$ such that for any distribution f of the voters, there exists an additive ε Nash-equilibrium?

U. Bhaskar and S. Pyne have proven that for $n \geq 3$ players, for any $\varepsilon < \frac{1}{n+3}$, one can find a distribution for which there is no additive ε Nash-equilibrium. However, the question of whether one can find an $\varepsilon = \frac{1}{n+3}$ equilibrium for $n \geq 3$ remains open. In fact, U. Bhaskar and S. Pyne have only proven the existence of an $\varepsilon = \frac{1}{n+1}$ equilibrium for any distribution f of the voters. In this write-up, we will first present the results shown by U. Bhaskar and S. Pyne, before exploring ideas to improve the results shown in their paper.

Results by U. Bhaskar and S. Pyne

Now, we present three results from the paper published by U. Bhaskar and S. Pyne. The first theorem explores the case of three players, while the next two theorems provide more general results for n players. The second theorem provides an ε-equilibrium, while the third shows a lower bound for the best possible ε. Without further ado, here are the theorems:

Theorem 1: When $n = 3$, for any distribution of the voters, there exists an $\varepsilon = \frac{1}{6} + \gamma$ equilibrium for any $\gamma > 0$.

Theorem 2: When $n \geq 3$, for any distribution of the voters, there exists an $\varepsilon = \frac{1}{n+1}$ equilibrium.

Theorem 3: For any $n \geq 3$, we have that for any $\varepsilon < \frac{1}{n+3}$, there exists a distribution f for which there exists no ε-equilibrium.

We will present proofs for Theorems 1 and 2 below. As we will see, the proofs of Theorems 1 and 2 are quite straightforward but involve a lot of case analysis, highlighting the complexity of the problem in general. We will omit the proof of Theorem 3, which is quite technical. Here are the proofs for Theorems 1 and 2:

Proof of Theorem 1: Let $\gamma > 0$ be arbitrarily small. Then, for a given $x \in [0,1]$, let $Cut(x)$ denote the point $y \in [0,1]$ such that $\int_0^y f(t)dt = x$. Set $x_1 = Cut\left(\frac{1}{3}\right)$ and $x_3 = Cut\left(\frac{2}{3}\right)$.

Then, $\int_{x_1}^{x_3} f(t)dt = \frac{1}{3}$, so by additivity of the integral, we have:

$$\int_{x_1}^{x_{1,3}} f(t)dt + \int_{x_{1,3}}^{x_3} f(t)dt = \frac{1}{3}$$

Then, it follows that one term of the sum is at least $\frac{1}{6}$, assume WLOG that $\int_{x_1}^{x_{1,3}} f(t)dt \geq \frac{1}{6}$.

Then, note that $0 = \int_{x_3}^{x_{3,3}} f(t)dt < \frac{1}{6} - \frac{\gamma}{2} < \int_{x_1}^{x_{1,3}} f(t)dt$, and the integral is continuous, so by the intermediate value theorem, $\exists z \in (x_1, x_3)$ such that $\int_{z}^{\frac{z+x_3}{2}} f(t)dt = \frac{1}{6} - \frac{\gamma}{2}$. Let $x_2 = z$.

Then, we claim that the strategy profile (x_1, x_2, x_3) is an $\varepsilon = \frac{1}{6} + \gamma$ equilibrium.

First, assume that players 2 and 3 keep their respective positions x_2, x_3. Then, by design, $\int_{x_{1,2}}^{x_2} f(t)dt < \frac{1}{6} + \gamma$, so player 1 cannot improve his utility by more than ε by moving closer to the left of x_2. Moreover, since $\int_{x_2}^{x_3} f(t)dt$ and $\int_{x_3}^{1} f(t)dt$ are less than $\frac{1}{2} + \gamma$, player 1 cannot improve his utility by more than ε by going between the players or right of player 3 (note that player 1 already has $\frac{1}{3}$ utility on his left). So, player 1 has no incentive to move assuming the other players keep their positions.

Now, assume that players 1 and 3 keep their respective positions x_1, x_3. Then, since player 2 already has $\frac{1}{6} - \frac{\gamma}{2}$ utility on his right, he cannot increase his utility by more than ε by moving left of x_1 or right of x_3, as $\frac{1}{6} - \frac{\gamma}{2} + \varepsilon > \frac{1}{3}$. Also, since $\int_{x_1}^{x_3} f(t)dt = \frac{1}{3} < \frac{1}{6} - \frac{\gamma}{2} + \varepsilon$, player 2 has no incentive to move anywhere else between players 1 and 3. So, player 2 has no incentive to move.

Finally, assume that players 1 and 2 keep their respective positions x_1, x_2. Then, since $\int_{0}^{x_1} f(t)dt$ and $\int_{x_1}^{x_2} f(t)dt$ are less than $\frac{1}{2} + \gamma$, player 3 has no incentive to move between players 1 and 2 or left of player 1. Moreover, since $\int_{x_2}^{x_{2,3}} f(t)dt = \frac{1}{6} - \frac{\gamma}{2} < \varepsilon$, player 3 has no incentive to move closer to the right of x_2. So, player 3 also has no incentive to change his position.

So, no player has a way to improve his utility by more than ε if the other players keep their positions, which means that the strategy profile (x_1, x_2, x_3) is an $\varepsilon = \frac{1}{6} + \gamma$ equilibrium. □

Proof of Theorem 2: Let $Cut(x)$ be defined as in the proof of Theorem 1 for $x \in [0,1]$.

Then, let $x_i = Cut(\frac{i}{n+1})\ \forall 1 \leq i \leq n$.

We claim that this strategy profile is an $\varepsilon = \frac{1}{n+1}$ equilibrium.

First observe that no player has any incentive to go to the left of his left neighbor (if it exists) or right of his right neighbor (if it exists), since by design, we have:

$$\int_0^{x_1} f(t)dt = \int_{x_1}^{x_2} f(t)dt = \cdots = \int_{x_n}^{1} f(t)dt = \frac{1}{n+1} = \varepsilon$$

Also, since $\int_{x_1}^{x_2} f(t)dt = \int_{x_{n-1}}^{x_n} f(t)dt = \varepsilon$, players 1 and n have no incentive to move closer to their neighbor.

Now, we consider the case of some player i, with $2 \leq i < n$, moving between his neighbors. Assume that such player i moves to position $x_{i'}$, where $x_{i-1} < x_{i'} < x_{i+1}$. WLOG, we can assume that $x_{i'} < x_i$. We denote by $x_{i',i+1}$ the position of the midpoint of $x_{i'}$ and x_{i+1}.

Let $F(x) = \int_0^x f(t)dt$, let $S = (x_1, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_n)$ denote the original strategy profile, and $S' = (x_1, \ldots, x_{i-1}, x_{i'}, x_{i+1}, \ldots, x_n)$ denote the new strategy. Then, we have the following:

$$U(i; S') - U(i; S) = \left[F\left(x_{i',i+1}\right) - F\left(x_{i-1,i'}\right)\right] - \left[F(x_{i,i+1}) - F(x_{i-1,i})\right]$$

By assumption, $x_{i'} < x_i$, so $F\left(x_{i,i+1}\right) > F\left(x_{i',i+1}\right)$, and we have:

$$U(i; S') - U(i; S) < F(x_{i-1,i}) - F\left(x_{i-1,i'}\right)$$

Clearly, $F(x_{i-1,i}) < F(x_i)$ and $F\left(x_{i-1,i'}\right) > F(x_{i-1})$.

Thus, it follows that:

$$U(i; S') - U(i; S) < F(x_i) - F(x_{i-1}) = \varepsilon$$

So, no player has any incentive to move anywhere between his neighbors.

Hence, we have found an $\varepsilon = \frac{1}{n+1}$ equilibrium. □

In the next section of the write-up, we will explain what makes the problem of finding an equilibrium for $\varepsilon < \frac{1}{n+1}$ so challenging.

Difficulties of Finding ε-Equilibria

As we have seen, proving that a strategy profile is an ε-equilibrium requires a lot of casework even for simple strategies, as we need to show that no player can improve his utility by more than ε assuming all others maintain their positions. Specifically, when $n \geq 3$, two main difficulties arise:

i) One must ensure that players 2 up to $n-1$ have enough utility not to have any incentive to move left of player 1 or right of player n.

ii) One must also ensure that the players have no incentive to move closer to their neighbors.

In both proofs, difficulty i) is overcome by setting $\int_0^{x_1} f(t)dt = \int_{x_n}^1 f(t)dt = A$ and giving the $n-2$ middle players at least $A-\varepsilon$ utility. Note that for the case $n=3$, there is only one middle player, so by the argument in the proof of Theorem 1, one can always ensure that the middle player gets at least $\frac{1}{2} - A - \gamma$ for any $\gamma > 0$. Then, for $n=3$, it suffices to have:

$$\frac{1}{2} - A > A - \varepsilon \Leftrightarrow A < \frac{1}{4} + \frac{1}{2}\varepsilon$$

For the proof of Theorem 2, difficulty i) is overcome by setting $A = \varepsilon$, so even if the middle players get no utility, they will not have an incentive to move left of player 1 or right of player n.

The proof of Theorem 1 overcomes difficulty ii) by ensuring that at least $1-\varepsilon$ utility is allocated by the strategy. Then, for $n=3$, it suffices to have:

$$2A + \left(\frac{1}{2} - A - \gamma\right) = 1 - \varepsilon \Leftrightarrow A = \frac{1}{2} + \gamma - \varepsilon$$

Then, since $A < \frac{1}{4} + \frac{1}{2}\varepsilon$, it follows that:

$$\frac{1}{2} + \gamma - \varepsilon < \frac{1}{4} + \frac{1}{2}\varepsilon \Leftrightarrow \frac{3}{2}\varepsilon > \frac{1}{4} + \gamma \Leftrightarrow \varepsilon > \frac{1}{6} + \frac{2}{3}\gamma$$

This argument already suggests that for $\varepsilon < \frac{1}{6}$, one will not be able to find a general strategy to find an ε equilibrium for $n = 3$ players. Note that this is a special case of Theorem 3, which is proven by U. Bhaskar and S. Pyne.

The proof of Theorem 2 overcomes difficulty ii) by ensuring that the area between any two players is no more than ε. So, the strategy adopted for the proof of Theorem 2 can only be applied when we have:

$$2A + (n-1)\varepsilon \geq 1 \Leftrightarrow (n+1)\varepsilon \geq 1 \Leftrightarrow \varepsilon \geq \frac{1}{n+1}$$

So, when $\varepsilon < \frac{1}{n+1}$, we need to adopt a different strategy to overcome both difficulties i) and ii).

In the next section of the write-up, we will present how we attempted to overcome those difficulties and get an ε-equilibrium for $\varepsilon < \frac{1}{n+1}$.

Our Strategy for Finding ε-Equilibria

During the semester, we mainly focused on the case of $n = 4$ players, hoping to find a pattern and generalize it for all values of n. However, the case $n = 4$ has already proven itself to be quite difficult, and we have not been able to find a strategy profile to find an $\varepsilon < \frac{1}{5}$ equilibrium. As explained in the above section, when $n \geq 4$, finding a strategy to deal with both difficulties i) and ii) when $\varepsilon < \frac{1}{n+1}$ is a real challenge. Nonetheless, we will present our train of thought, hoping to inspire further work and ideas.

The first observation we made is that to meet difficulty i) above, we need the utility of each of the middle players to be at least $max(\int_0^{x_1} f(x)dx\,, \int_{x_n}^1 f(x)dx) - \varepsilon$. Thus, setting $\int_0^{x_1} f(x)dx = \int_{x_n}^1 f(x)dx$ provides us with the most flexibility for our strategy and maintains symmetry in the repartition of the unit mass. Moreover, we observed from the proof of Theorem 2 that by making the utilities between the players no more than ε, no player could ever improve his utility by more than ε by moving closer to his neighbors, thus overcoming difficulty ii).

So, for the case $n = 4$, our observations have motivated us to look for an $\varepsilon = \frac{1}{5} - \gamma$ equilibrium for some $\gamma > 0$ by first setting the areas between the players equal to ε and setting $\int_0^{x_1} f(x)dx = \int_{x_4}^1 f(x)dx$, giving us that:

- $\int_0^{x_1} f(x)dx = \int_{x_4}^1 f(x)dx = \frac{1}{5} + \frac{3}{2}\gamma$
- $\int_{x_1}^{x_2} f(x)dx = \int_{x_2}^{x_3} f(x)dx = \int_{x_3}^{x_4} f(x)dx = \varepsilon = \frac{1}{5} - \gamma$

Note that in the current configuration, we overcome difficulty ii), but we also need players 2 and 3 to have their utility be at least $\frac{1}{5}+\frac{3}{2}\gamma-\varepsilon=\frac{5}{2}\gamma$ to deal with difficulty i), which is not guaranteed for all distributions by our strategy.

However, we observed that since $\int_{x_2}^{x_3} f(x)dx=\frac{1}{5}-\gamma$, one of the utilities $\int_{x_2}^{x_{2,3}} f(x)dx$ and $\int_{x_{2,3}}^{x_3} f(x)dx$ must be at least $\frac{1}{2}\left(\frac{1}{5}-\gamma\right)=\frac{1}{10}-\frac{\gamma}{2}$. Namely, at least one of players 2 or 3 already has at least $\frac{1}{10}-\frac{\gamma}{2}$ utility. Assume WLOG that $\int_{x_{2,3}}^{x_3} f(x)dx \geq \frac{1}{10}-\frac{\gamma}{2}$, namely, player 3 has at least $\frac{1}{10}-\frac{\gamma}{2}$ utility. Then, recall that player 3 needs $\frac{5}{2}\gamma$ utility to have an ε-equilibrium, so it suffices to have:

$$\frac{1}{10}-\frac{\gamma}{2}\geq\frac{5}{2}\gamma \Leftrightarrow \frac{1}{10}\geq 3\gamma \Leftrightarrow \gamma\leq\frac{1}{30}$$

This observation motivates us to focus our attention on the case $\gamma=\frac{1}{30}$ (as we want to maximize γ to minimize ε). Note that for $\gamma=\frac{1}{30}$, we have $\varepsilon=\frac{1}{5}-\frac{1}{30}=\frac{1}{6}$, and our strategy profile gives us the following:

- $\int_0^{x_1} f(x)dx=\int_{x_4}^{1} f(x)dx=\frac{1}{4}$
- $\int_{x_1}^{x_2} f(x)dx=\int_{x_2}^{x_3} f(x)dx=\int_{x_3}^{x_4} f(x)dx=\frac{1}{6}$

Then, WLOG, we can assume that player 3 has at least $\frac{1}{12}$ utility and thus does not have an incentive to move left of player 1 or right of player 4. So, we only need to ensure that player 2 gets $\frac{1}{12}$ utility while making sure that the players do not have an incentive to move closer to each other.

Note that if indeed player 2 happens to have at least $\frac{1}{12}$ utility for the given function, then we have successfully found an $\varepsilon = \frac{1}{6}$ equilibrium. However, if the current utility of player 2 is less than $\frac{1}{12}$, then we need to update the position of player 2 so that he does not have any incentive to move left of x_1 or right of x_4. One way to achieve this is to observe that because the current utility of player 2 is less than $\frac{1}{12}$, it follows that $\int_{x_1}^{x_{1,2}} f(x)dx > \frac{1}{12}$, so by the intermediate value theorem, $\exists z \in (x_1, x_2)$ such that $\int_{z}^{\frac{z+x_2}{2}} f(x)dx = \frac{1}{12}$. Then, by updating the position of player 2 to be z, we ensure that player 2 has at least $\frac{1}{12}$ utility. But now, the utility between players 2 and 3 has become more than $\frac{1}{6}$, so we encounter difficulty ii) again.

And this is where we are stuck in our reasoning. Whenever we try to deal with difficulty ii), difficulty i) arises, and when overcoming difficulty i), difficulty ii) becomes an issue. This duality explains why this problem is very challenging. We hope nonetheless that our ideas and struggles will inspire future researchers and enthusiasts to find an $\varepsilon < \frac{1}{n+1}$ equilibrium for $n \geq 4$, or show that such equilibrium does not exist.

Conclusion

For our Winter 2025 DRP project, we have discovered a fascinating game theory problem that involves finding an additive ε-Nash equilibrium for any given distribution of voters over the interval $[0,1]$ given n political parties, so that no party has an incentive to deviate from his position. Despite the simplicity of its statement, the problem has proven to be very difficult to solve, even for as few as $n = 4$ parties. It remains that the $\varepsilon = \frac{1}{n+1}$ equilibrium found by U. Bhaskar and S. Pyne is nearly optimal for large values of n, since for any $n \geq 3$ and $\varepsilon < \frac{1}{n+3}$, there exists a distribution of voters such that there is no additive ε-Nash equilibrium. However, the fact that one can find an $\varepsilon = \frac{1}{n+3} + \gamma$ for any $\gamma > 0$ when $n = 3$ leaves open the possibility of improving on the $\varepsilon = \frac{1}{n+1}$ equilibrium.

Reference

Bhaskar, Umang, and Soumyajit Pyne. "Nearly Tight Bounds on Approximate Equilibria in Spatial Competition on the Line." *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 13, Apr. 2025, pp. 13641–48. https://doi.org/10.1609/aaai.v39i13.33490.