17.810/17.811 – Game Theory

Lecture 2: Games in Strategic Form and Nash Equilibrium

Asya Magazinnik

MIT

Where We Are

We are ready to move from individual preferences to how people act on their preferences in strategic interactions with each other.

We start with static games of complete information:

- Static games take the following form:
  1. The players simultaneously choose actions
  2. They receive payoffs that depend on the combination of actions just chosen
- Complete information means everyone knows how every set of actions maps to outcomes

What We Will Learn

In this part of the class, we will learn:

1. How to translate a verbal statement of a strategic interaction to a formal game
   - Normal-form representation of a game
2. What constitutes the "solution" to a game (equilibrium)
3. How to solve for the equilibrium of a game
   - Iterated elimination of strictly dominated strategies
   - Nash equilibrium

Reading

These slides will focus on the following readings:

- Gibbons, 1.1
  - Normal Form Representation
  - Equilibrium: Iterated Elimination of Strictly Dominated Strategies
  - Equilibrium: Nash Equilibrium

Classic Example: The Prisoners' Dilemma

By way of example, let's begin with the classic Prisoners' Dilemma, which has the following verbal statement:

- Two suspects are arrested and charged with a crime. The police lack sufficient evidence to convict them unless one of them confesses.
- The police take the two suspects to separate rooms and offer each of them the same proposition:
  - If nobody confesses, both will be convicted of a minor offense and sentenced to 1 month in jail
  - If one confesses, the one who confessed is released immediately and the other receives a 9-month sentence
  - If both confess, they each serve 6 months

Elements of a Game

- A set of players $I = \{1, 2, \ldots, n\}$, with an arbitrary player i
- A strategy space S_i which is the set of actions available to each player, where $s_i \in S_i$ is one action a player might take
- For each player, preferences over the set of outcomes captured by utility functions: $\{u_1, u_2, \ldots, u_n\}$
  - where an "outcome" is fully specified by the strategies played by all players, i.e. $u_1(s_1, s_2, ..., s_n)$
- In general, we also specify players' information sets: for now, we assume "common knowledge"
- We also want to know the sequence of the game: for now, we have simultaneous play at one moment in time

The Normal-Form Representation of a Game

Definition (Normal-Form Representation)

The normal-form representation of an n-player game specifies:

1. the players' strategy spaces $S_1, S_2, ..., S_n$
2. their payoff functions $u_1, u_2, ..., u_n$

We denote this game by $G = \{S_1, S_2, ..., S_n; u_1, u_2, ..., u_n\}$.

Formalizing the Prisoners' Dilemma

Example (Prisoners' Dilemma)

		Clyde	
		Mum	Fink
Bonnie	Mum	-1,-1	-9, 0
	Fink	0,-9	-6,-6

- Players: Bonnie (1, row player) and Clyde (2, column player)
- Strategies (Actions): $S_1 = \{\text{Mum, Fink}\}$ and $S_2 = \{\text{Mum, Fink}\}$
- Payoffs: numbers (utility representations) in the matrix above
  - e.g. $u_1(s_1 = \text{Mum}, s_2 = \text{Mum}) = -1$

The Strategic Element of Choice

We've talked at length about how rational actors choose actions that maximize their utilities given constraints.

We now have our first encounter with the strategic element of choice: players are not choosing in a vacuum.

- Note, for instance, that Bonnie's payoff from choosing **Fink** critically depends on whether Clyde chooses **Fink** (-6) or **Mum** (0)

But Clyde's fate similarly depends on Bonnie's choice! How do we reason about this?

A Place to Start: Strictly Dominated Strategies

At the very least, a rational player can rule out actions that don't yield the best outcomes whatever the other players might do. Let's start with that and see if we can make any progress.

Definition (Strictly Dominated Strategy)

In the normal-form game $G = \{S_1, ..., S_n; u_1, ..., u_n\}$, let s_i' and s_i'' be feasible strategies, i.e. members of S_i.

Strategy s_i' is strictly dominated by strategy s_i'' if for each feasible combination of the other players' strategies, i's payoff from playing s_i' is strictly less than i's payoff from playing s_i'':

$$u_i(s_1, ..., s_{i-1}, s_i', s_{i+1}, ..., s_n) < u_i(s_1, ..., s_{i-1}, s_i'', s_{i+1}, ..., s_n)$$

for each $(s_1, ..., s_{i-1}, s_{i+1}, ..., s_n)$ that can be constructed from the other players' strategy spaces $S_1, ..., S_{i-1}, S_{i+1}, ..., S_n$.

Does the Prisoners' Dilemma have strictly dominated strategies?

Example (Prisoners' Dilemma)

		Clyde	
		Mum	Fink
Bonnie	Mum	-1,-1	-9, 0
	Fink	0,-9	-6,-6

- Let's look at Bonnie's decision; the game is symmetric, so the same goes for Clyde
- If Clyde stays Mum: Fink (0) is strictly better than Mum (-1)
- If Clyde Finks: Fink (-6) is strictly better than Mum (-9)
- So we have found a strictly dominated strategy: **Mum**

Iterated Elimination of Strictly Dominated Strategies

Example (Prisoners' Dilemma)

		Clyde	
		Mum	Fink
Bonnie	Mum	-1,-1	-9, 0
	Fink	0,-9	-6,-6

We have eliminated **Mum** for Bonnie and assert that she will choose **Fink**

This simplifies the decision for Clyde: knowing Bonnie will **Fink**, he too will **Fink** (-6) rather than stay **Mum** (-9)

Some Important Notes on the Prisoners' Dilemma

Example (Prisoners' Dilemma)

		Clyde	
		Mum	Fink
Bonnie	Mum	-1,-1	-9, 0
	Fink	0,-9	-6,-6

We have just solved for our first equilibrium! What did we learn?

- The equilibrium (**Fink**, **Fink**) is rational, but not socially optimal
- The solution would have been the same with any numbers that preserved the structure: Temptation $>$ Reward $>$ Punishment $>$ Sucker ($0 > -1 > -6 > -9$), a reminder about ordinal utilities
- This structure is common to many social problems: price collusion, pandering, nuclear proliferation

But What If There Are No Strictly Dominated Strategies?

Example

		Player 2 L	C	R
Player 1	T	0,4	4,0	5,3
	M	4,0	0,4	5,3
	B	3,5	3,5	6,6

We can quickly see that we cannot eliminate any strategies:

- Player 1: **T** best when **C**, **M** best when **L**, **B** best when **R**
- Player 2: **L** best when **T**, **C** best when **M**, **R** best when **B**

(If it's the best in some circumstance, it can't be strictly dominated.)

We define this equilibrium concept by what it looks like when you're there (and return later to the matter of *how* we get there):

Definition (Nash Equilibrium)

In the n-player normal-form game $G = \{S_1, ..., S_n; u_1, ..., u_n\}$, the strategies $(s_1^*, ..., s_n^*)$ are a Nash equilibrium if, for each player i, s_i^* is (at least tied for) player i's best response to the strategies specified for the $n-1$ other players, $(s_1^*, ..., s_{i-1}^*, s_{i+1}^*, ..., s_n^*)$:

$$u_i(s_1^*, ..., s_{i-1}^*, s_i^*, s_{i+1}^*, ..., s_n^*) \geq u_i(s_1^*, ..., s_{i-1}^*, s_i, s_{i+1}^*, ..., s_n^*)$$

for every feasible strategy $s_i \in S_i$. In other words, s_i^* solves:

$$\max_{s_i \in S_i} \; u_i(s_1^*, ..., s_{i-1}^*, s_i, s_{i+1}^*, ..., s_n^*)$$

Nash Equilibrium

Leaving aside the thorny issue of how the *players* got there, let's see how an analyst might get there (let's solve for a N.E.).

Strategy:

1. We'll first need to identify Player 1's best response to every action Player 2 might take
2. Then we identify Player 2's best response to every action Player 1 might take
3. Then we'll look for any cells where both players are simultaneously best-responding

Solving for a Nash Equilibrium in a Normal Form Game

Let's revisit the game form that we failed to solve using iterated elimination of strictly dominated strategies.

Example

		Player 2		
		L	C	R
Player 1	T	0,4	4,0	5,3
	M	4,0	0,4	5,3
	B	3,5	3,5	6,6

1. Start with Player 1

Solving for a Nash Equilibrium in a Normal Form Game

Let's revisit the game form that we failed to solve using iterated elimination of strictly dominated strategies.

Example

		Player 2		
		L	C	R
Player 1	T	0,4	4,0	5,3
	M	4,0	0,4	5,3
	B	3,5	3,5	6,6

1. Start with Player 1

Solving for a Nash Equilibrium in a Normal Form Game

Let's revisit the game form that we failed to solve using iterated elimination of strictly dominated strategies.

Example

		Player 2		
		L	C	R
Player 1	T	0,4	4,0	5,3
	M	4,0	0,4	5,3
	B	3,5	3,5	6,6

1. Start with Player 1

Solving for a Nash Equilibrium in a Normal Form Game

Let's revisit the game form that we failed to solve using iterated elimination of strictly dominated strategies.

Example

			Player 2	
		L	C	R
Player 1	T	0,4	4,0	5,3
	M	4,0	0,4	5,3
	B	3,5	3,5	6,6

2. Go on to Player 2

Solving for a Nash Equilibrium in a Normal Form Game

Let's revisit the game form that we failed to solve using iterated elimination of strictly dominated strategies.

Example

		Player 2		
		L	C	R
Player 1	T	0,4	4,0	5,3
	M	4,0	0,4	5,3
	B	3,5	3,5	6,6

2. Go on to Player 2

Solving for a Nash Equilibrium in a Normal Form Game

Let's revisit the game form that we failed to solve using iterated elimination of strictly dominated strategies.

Example

		Player 2		
		L	C	R
Player 1	T	0,4	4,0	5,3
	M	4,0	0,4	5,3
	B	3,5	3,5	6,6

2. Go on to Player 2

Solving for a Nash Equilibrium in a Normal Form Game

Let's revisit the game form that we failed to solve using iterated elimination of strictly dominated strategies.

Example

		Player 2		
		L	C	R
Player 1	T	0,4	4,0	5,3
	M	4,0	0,4	5,3
	B	3,5	3,5	6,6

3. The only cell where both players are simultaneously best-responding to one another is (**B**,**C**) and that is our only Nash equilibrium.

N.E. vs. I.E.S.D.S.

How can we relate the two solution concepts we've seen thus far?

- If I.E.S.D.S. eliminates all but the strategies $s_1^*, s_2^*, ..., s_n^*$, then these strategies are also the unique Nash equilibrium of the game

- However, some strategies might survive I.E.S.D.S. but not constitute a Nash equilibrium (as in our example)

Thus Nash equilibrium is a “stronger” solution concept.

A Few Important Notes on Nash Equilibrium

The Nash equilibrium concept is silent about how the players ended up there, but here are some helpful ideas:

- Suppose we start elsewhere and offer any one of the players the opportunity to deviate. Will any player do so? Yes.
- Suppose we start at a Nash equilibrium and offer any one of the players the opportunity to deviate. Will any player do so? No.
- Note that this logic only makes sense for one player at a time; get used to fixing other players' strategies.

Thus Nash equilibria can be said to be stable and self-reinforcing.

What If There Are Multiple Nash Equilibria?

Example (Bach or Stravinsky)

		Pat	
		Bach	Stravinsky
Chris	Bach	2,1	0,0
	Stravinsky	0,0	1,2

- Chris and Pat prefer to spend the evening together than apart
- But they have different musical preferences
- Note that (**Bach**, **Bach**) and (**Stravinsky**, **Stravinsky**) are both Nash equilibria because Chris & Pat love being together
- But we have no idea which music they'll settle on

An Example: Interest Group Contributions

- The U.S. Supreme Court recognizes donating to political campaigns and causes as a form of political speech (Citizens United vs. FEC)
- Given the vast opportunities to influence policy through campaign contributions, what determines when and how much different groups contribute?

Contribution Game: The Problem

- Suppose there are two interest groups who seek to influence government policy (healthcare)
  1. PAC 1 (American Federation of State County and Municipal Workers)
  2. PAC 2 (American Medical Association)
- Both groups know the final policy is a function of how much campaign support they give to the governing party
- Policy preferences can be represented on a line
- PAC 1's most preferred policy is 0, PAC 2's most preferred policy is 1
- The government favors a policy of $\frac{1}{2}$, but may be influenced by campaign contributions

Utility functions

- Each group chooses to contribute an amount $s_i \in \mathbb{R}^+$ and the final policy is $\Gamma(s_1, s_2) = \frac{1}{2} - s_1 + s_2$
- The groups make their choices simultaneously and the government keeps all the contributions regardless of the policy outcome
- Suppose the interest groups each have utility functions over policies and contributions of the form:

$$u_1(s_1, s_2) = -(\Gamma(s_1, s_2))^2 - s_1$$
$$u_2(s_1, s_2) = -(1 - \Gamma(s_1, s_2))^2 - s_2$$

Utility functions

If we substitute the final policy into these utility functions, we get:

$$u_1(s_1, s_2) = -\left(\frac{1}{2} - s_1 + s_2\right)^2 - s_1$$
$$u_2(s_1, s_2) = -\left(1 - \left(\frac{1}{2} - s_1 + s_2\right)\right)^2 - s_2$$

Finding the best response: optimization

To calculate Nash equilibrium, we want to write down both players' best response functions.

When Player 2 is playing his best response (call it s_2^*), Player 1's utility function is:

$$u_1(s_1, s_2^*) = -\left(\frac{1}{2} - s_1 + s_2^*\right)^2 - s_1$$

What is the s_1 that maximizes Player 1's utility? It's:

$$\arg\max_{s_1 \in \mathbb{R}^+} u_1(s_1, s_2^*)$$

Finding the best response: optimization

Take the first derivative:

$$\frac{\partial}{\partial s_1} u_1(s_1, s_2^*) = -2\left(\frac{1}{2} - s_1 + s_2^*\right)(-1) - 1 = -2s_1 + 2s_2^*$$

And do the same for Player 2, holding s_1 fixed at Player 1's best response, s_1^*:

$$\frac{\partial}{\partial s_2} u_2(s_1^*, s_2) = 2\left(1 - \left(\frac{1}{2} - s_1^* + s_2\right)\right) - 1 = 2s_1^* - 2s_2$$

The maximum is reached when these first derivatives are equal to zero. We call these the first order conditions:

$$FOC_1 : -2s_1^* + 2s_2^* = 0$$
$$FOC_2 : 2s_1^* - 2s_2^* = 0$$

Finding the best response: optimization

Each player's best response function is given by solving for their optimal strategy (rearranging the FOC):

$$BR_1 : s_1^* = s_2^*$$
$$BR_2 : s_2^* = s_1^*$$

So this problem has many solutions: any (s_1, s_2) such that $s_1 = s_2$ constitutes a Nash equilibrium.

Note: You can learn/refresh these skills in the following sections of *Mathematics for Economists*:

- 2.4, Computing Derivatives
- 4.1, Composite Functions and the Chain Rule
- 17.1–17.3, Unconstrained Optimization

Finding the best response: optimization

We have learned that matching your opponent's contribution is a best response.

For example, each could donate \$1, \$10, or \$1,000,000 — all are equilibria.

All contribution schedules get the same policy outcome, $\frac{1}{2}$.

Does this accord with reality?

- Senate Democrat PAC contributions in 2010: \$44,091,038
- Senate Republicans PAC contributions in 2010: \$45,725,692

War of Attrition

Consider the following model of conflict between two polities over a piece of territory.

- Each player chooses a length of time $t \geq 0$ that they will fight.
- When one player gives up, the other gets the territory.
- If both players give up at the same time, then each has an equal chance of getting the territory.
- Fighting is assumed to be costly: each player prefers as short a fight as possible.

Let's model this game formally

- Time is a continuous variable that starts at 0 and runs indefinitely
- Assume that player i's valuation of the territory is $v_i > 0$
- Moreover, she values the 50-50 lottery of winning the territory at $v_i/2$
- Assume that each unit of time that passes costs 1 unit
- Payoffs:
  - If a player concedes at time t_i and does not win the territory her payoff is $-t_i$
  - If at time t_j a player's opponent concedes, then her payoff is $v_i - t_j$
  - If there is a tie, then the payoff is $v_i/2 - t_i$, i.e., the expected value of the lottery

Let's model this game formally

- Players: Two political actors (states, insurgents, etc)
- Actions: The set of possible concession times, $t_i \in \mathbb{R}^+$
- Preferences: Player i's payoffs are represented by the utility function:

$$u_i(t_1, t_2) = \begin{cases} -t_i & \text{if } t_i < t_j \\ \frac{1}{2}v_i - t_i & \text{if } t_i = t_j \\ v_i - t_j & \text{if } t_i > t_j. \end{cases}$$

We can find the NE in two ways: by best response and by direct argument.

Player i's utility function in the War of Attrition

Let's start by deriving the players' best response functions. To do so, let's first look at their utilities:

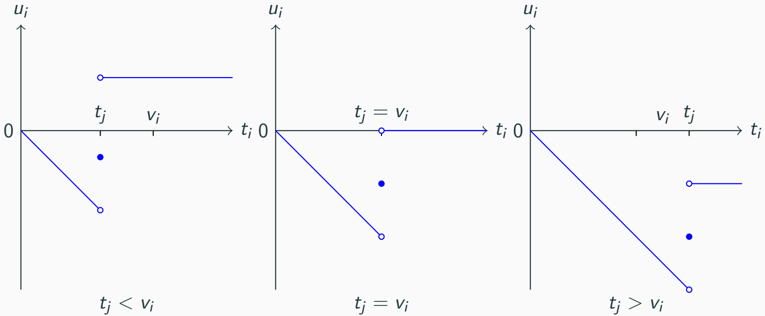

Best Response Functions

The best response is simply the utility-maximizing strategy in every case:

1. If j's intended concession time is early enough, then it is optimal for i to wait for player j to concede.
   - Anytime t_i beyond t_j is as good as any other (why?)
   - Because i will get the same payoff: $v_i - t_j$ if $t_j < t_i$
2. If j intends to hold out for a long time (t_j is high) i should concede immediately.

Formally, we have the following best response function:

$$B_i(t_j) = \begin{cases} \{t_i : t_i > t_j\} & \text{if } t_j < v_i \\ \{t_i : t_i = 0 \text{ or } t_i > t_j\} & \text{if } t_j = v_i \\ \{0\} & \text{if } t_j > v_i. \end{cases}$$

The players' best response functions in the War of Attrition

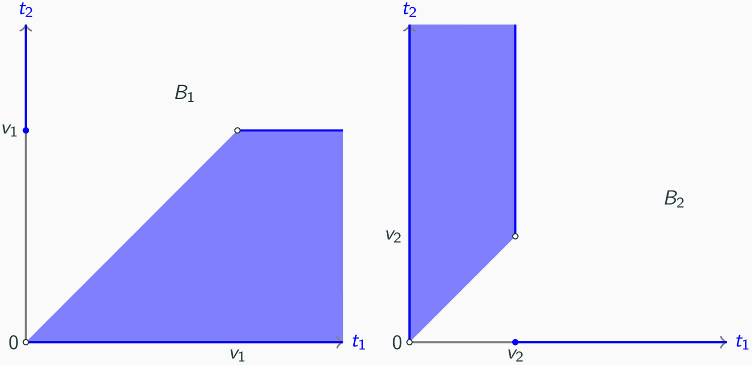

For a case in which $v_1 > v_2$. Player 1's best response is in the left panel; Player 2's is in the right panel. (The sloping edges are excluded.)

The players' best response functions in the War of Attrition

The Nash equilibria are simply the areas of overlap between the two panels: where both players are simultaneously best-responding.

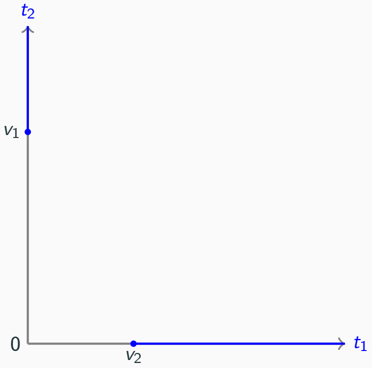

1. $t_1 = 0, t_2 \geq v_1$
2. $t_2 = 0, t_1 \geq v_2$

Direct argument for Nash equilibrium

Here we consider the Nash equilibrium directly, by cases.

1. Suppose $t_1 = t_2$
   - Then either player can increase her payoff by conceding slightly later, at $t_i + \epsilon$, and get the object for sure, making them better off
   - Therefore, this cannot be an equilibrium strategy

2. Suppose $0 < t_i < t_j$
   - Here player i loses for sure and pays a cost t_i
   - If she were to concede at $t = 0$ she would not pay this cost and be better off
   - Therefore, these strategies cannot be part of a NE

Direct argument for Nash equilibrium

3. Suppose that $0 = t_i < t_j < v_i$
   - Then player i can increase her payoff from 0 to $v_i - t_j > 0$ by conceding slightly after t_j
   - So this is not a Nash equilibrium either

Therefore, none of the above strategies can be a Nash equilibrium.

Direct argument for Nash Equilibrium

4. The remaining possibility is that $0 = t_i < t_j$ and $t_j > v_i$.
    - Here there is no possibility for any player to improve their payoffs unilaterally

The Nash equilibria are:

1. $t_1 = 0, t_2 \geq v_1$
2. $t_2 = 0, t_1 \geq v_2$

These points are the intersections of the best response pictures.

War of Attrition

- Notice that there is no fighting in the equilibrium of the war of attrition!
- This result appears many times in mathematical models of politics.
- Distributional conflict is not enough to generate confrontation between rational actors

Theory of Candidate Competition: Hotelling-Downs Model

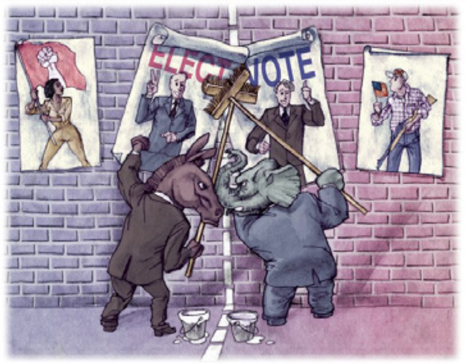

- The "Nixon Two-step": Run to the right in the primary, then back to the middle in the general election
- The incentives in majority rule elections push political actors toward the "median" voter

Voter and Preferences on a Line: The Median

- What is the median voter?
- First, we need a policy space: let's use the number line ($\mathbb{R}^1$)
- Each point on the line represents a policy that the government can choose
- There are N voters, each of whom has a most preferred policy (ideal point)
- The voter with $\frac{1}{2}$ of the voters to their "right" and $\frac{1}{2}$ voters to their "left" is the median voter

The Median Voter

Let's line these voters up on our policy space:

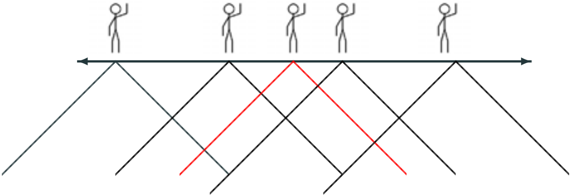

Utilities

- Each voter has a utility over policy outcomes represented by:

$$u_i(x_i^*, x) = -|x_i^* - x|$$

  where x_i^* is voter i's ideal point

- Two candidates, *D*emocrat and *R*epublican
  - Each candidate is office motivated and has utilities:

$$u_c(x) = \begin{cases} 1 & \text{if they win} \\ 0 & \text{if they lose} \end{cases}$$

  - Note that they don't care directly about policy; it only matters insofar as it gets them into office

Sequence of the Game

1. Each candidate simultaneously chooses a policy platform in the policy space.
2. Voters simultaneously and sincerely vote for their most preferred candidate.
3. The candidate with the most votes wins. If there is a tie, then each candidate wins with probability $\frac{1}{2}$.

Best Response Functions

- For an office motivated candidate, the median is key
- If the median voter (call them m) votes for you, then so does everyone on one side of m, giving you just over half the voters (you win)
- How do you capture the median voter?

$$B_D(x_R) = \begin{cases} \{x_D : x_R < x_D < 2m - x_R\} & \text{if } x_R < m \\ \{m\} & \text{if } x_R = m \\ \{x_D : 2m - x_R < x_D < x_R\} & \text{if } x_R > m. \end{cases}$$

Best Response Functions

If R chooses $x_R < m$, D will win the median with any x_D in the red area ($x_R < x_D < 2m - x_R$).

Argument for Nash Equilibrium

- Can (m, m) be a Nash equilibrium? Yes.

  1. Suppose both players are playing (m, m). They each win with probability $\frac{1}{2}$. At any other x, the non-deviating player wins for sure and the person choosing $x \neq m$ gets a payoff of zero.
  2. Therefore, (m, m) is an equilibrium.

- Is anything else an equilibrium?

Argument for Nash Equilibrium

1. Only 1 candidate at the median? No.
   - The candidate not at the median loses for sure. If she moves to m, she ties and has expected utility $\frac{1}{2} > 0$.
2. Two candidates to the left of the median? No.
   - The candidate closest to the median wins, or they tie if they are at the same place.
   - At least one candidate to can move to m and get more than $\frac{1}{2}$ the votes and win for sure. This payoff (1) is better than a tie ($\frac{1}{2}$), so this is not an equilibrium.
3. Two candidates to the right of the median? No.
   - Same argument as to the left.
4. One candidate on either side of the median? No.
   - In this case, one candidate loses or they tie if they are the same distance from the median. If a candidate moves to the median, they win for sure and get a payoff of $1 > \max\{\frac{1}{2}, 0\}$.

Median Voter Theorem

In majority rule elections with two parties and one dimension of competition, office motivated candidates propose policy platforms that appeal to the median voter.

This result persists under many different assumptions!

- Further reading: Scott Gehlbach's *Formal Models of Domestic Politics*, Chapters 1 and 2