6.207/14.15: Networks
Lecture 10: Introduction to Game Theory—2

Daron Acemoglu and Asu Ozdaglar
MIT

October 14, 2009

Outline

- Mixed Strategies
- Existence of Mixed Strategy Nash Equilibrium in Finite Games
- Characterizing Mixed Strategy Equilibria
- Existence of Nash Equilibrium in Infinite Games
- Extensive Form and Dynamic Games
- Subgame Perfect Nash Equilibrium
- Applications.

- Reading:
- Osborne, Chapters 3-6.

Nonexistence of Pure Strategy Nash Equilibria

- Example: Matching Pennies.

Player 1 \ Player 2	heads	tails
heads	$(-1, 1)$	$(1, -1)$
tails	$(1, -1)$	$(-1, 1)$

- No pure Nash equilibrium.
- How would you play this game?

Nonexistence of Pure Strategy Nash Equilibria

- Example: The Penalty Kick Game.

penalty taker \ goalie	left	right
left	$(-1, 1)$	$(1, -1)$
right	$(1, -1)$	$(-1, 1)$

- No pure Nash equilibrium.
- How would you play this game if you were the penalty taker?
  - Suppose you always show up left.
  - Would this be a "good strategy"?
- Empirical and experimental evidence suggests that most penalty takers "randomize" $\rightarrow$ mixed strategies.

Mixed Strategies

- Let Σ_i denote the set of probability measures over the pure strategy (action) set S_i.
  - For example, if there are two actions, S_i can be thought of simply as a number between 0 and 1, designating the probability that the first action will be played.
- We use $\sigma_i \in \Sigma_i$ to denote the mixed strategy of player i, and $\sigma \in \Sigma = \prod_{i \in \mathcal{I}} \Sigma_i$ to denote a mixed strategy profile.
- Note that this implicitly assumes that **players randomize independently**.
- We similarly define $\sigma_{-i} \in \Sigma_{-i} = \prod_{j \neq i} \Sigma_j$.
- Following von Neumann-Morgenstern expected utility theory, we extend the payoff functions u_i from S to Σ by

$$u_i(\sigma) = \int_S u_i(s) d\sigma(s).$$

Mixed Strategy Nash Equilibrium

Definition

(Mixed Nash Equilibrium): A mixed strategy profile σ^* is a (mixed strategy) Nash Equilibrium if for each player i,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*) \qquad \text{for all } \sigma_i \in \Sigma_i.$$

Proposition

Let $G = \langle \mathcal{I}, (S_i)_{i \in \mathcal{I}}, (u_i)_{i \in \mathcal{I}} \rangle$ be a finite strategic form game. Then, $\sigma^* \in \Sigma$ is a Nash equilibrium if and only if for each player $i \in \mathcal{I}$, every pure strategy in the support of σ_i^* is a best response to σ_{-i}^*.

Proof idea: If a mixed strategy profile is putting positive probability on a strategy that is not a best response, then shifting that probability to other strategies would improve expected utility.

Mixed Strategy Nash Equilibria (continued)

- It follows that **every action** in the support of any player's equilibrium mixed strategy yields the same payoff.
- **Implication:** it is sufficient to check pure strategy deviations, i.e., σ^* is a mixed Nash equilibrium if and only if for all i,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*) \qquad \text{for all } s_i \in S_i.$$

- **Note:** this characterization result extends to **infinite games**: $\sigma^* \in \Sigma$ is a Nash equilibrium if and only if for each player $i \in \mathcal{I}$, no action in S_i yields, given σ_{-i}^*, a payoff that exceeds his equilibrium payoff, the set of actions that yields, given σ_{-i}^*, a payoff less than his equilibrium payoff has σ_i^*-measure zero.

Examples

Example: Matching Pennies.

Player 1 \ Player 2	heads	tails
heads	$(-1, 1)$	$(1, -1)$
tails	$(1, -1)$	$(-1, 1)$

- Unique mixed strategy equilibrium where both players randomize with probability $1/2$ on heads.

Example: Battle of the Sexes Game.

Player 1 \ Player 2	ballet	football
ballet	$(1, 4)$	$(0, 0)$
football	$(0, 0)$	$(4, 1)$

- This game has two pure Nash equilibria and a mixed Nash equilibrium $\left((\frac{4}{5}, \frac{1}{5}), (\frac{1}{5}, \frac{4}{5})\right)$.

Weierstrass's Theorem

Theorem

(Weierstrass) *Let A be a nonempty compact subset of a finite dimensional Euclidean space and let $f : A \to \mathbb{R}$ be a continuous function. Then there exists an optimal solution to the optimization problem*

$$\begin{array}{ll} \textit{minimize} & f(x) \\ \textit{subject to} & x \in A. \end{array}$$

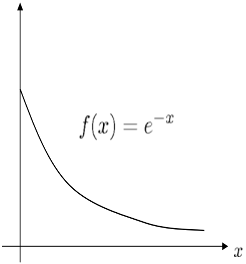

$\min_{x \geq 0} e^{-x} = 0$

There exists no optimal x that attains it

Kakutani's Fixed Point Theorem

Theorem

(Kakutani) *Let $f : A \rightrightarrows A$ be a correspondence, with $x \in A \mapsto \mathbf{f}(x) \subset A$, satisfying the following conditions:*

- *A is a compact, convex, and non-empty subset of a finite dimensional Euclidean space.*
- *$f(x)$ is non-empty for all $x \in A$.*
- *$f(x)$ is a convex-valued correspondence: for all $x \in A$, $f(x)$ is a convex set.*
- *$f(x)$ has a closed graph: that is, if $\{x^n, y^n\} \to \{x, y\}$ with $y^n \in f(x^n)$, then $y \in f(x)$.*

Then, f has a fixed point, that is, there exists some $x \in A$, such that $x \in f(x)$.

Definitions (continued)

- A set in a Euclidean space is compact if and only if it is bounded and closed.
- A set S is **convex** if for any $x, y \in S$ and any $\lambda \in [0, 1]$, $\lambda x + (1 - \lambda) y \in S$.

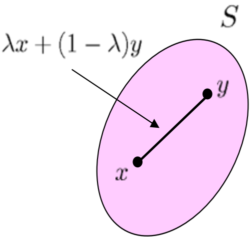

convex set

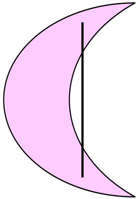

not a convex set

Kakutani's Fixed Point Theorem—Graphical Illustration

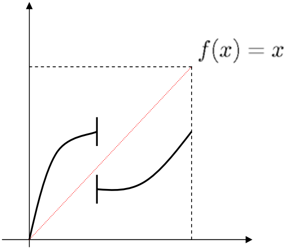

$f(x)$ is not convex-valued

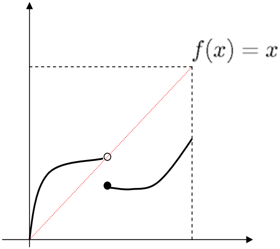

$f(x)$ does not have a closed graph

Nash's Theorem

Theorem

(Nash) *Every finite game has a mixed strategy Nash equilibrium.*

- Implication: matching pennies necessarily has a mixed strategy equilibrium.
- Why is this important?
  - Without knowing the existence of an equilibrium, it is difficult (perhaps meaningless) to try to understand its properties.
  - Armed with this theorem, we also know that every finite game has an equilibrium, and thus we can simply try to locate the equilibria.

Proof

- Recall that σ^* is a (mixed strategy) Nash Equilibrium if for each player i,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*) \qquad \text{for all } \sigma_i \in \Sigma_i.$$

- Define the best response correspondence for player i $B_i : \Sigma_{-i} \rightrightarrows \Sigma_i$ as

$$B_i(\sigma_{-i}) = \left\{\sigma_i' \in \Sigma_i \mid u_i(\sigma_i', \sigma_{-i}) \geq u_i(\sigma_i, \sigma_{-i}) \text{ for all } \sigma_i \in \Sigma_i\right\}.$$

- Define the set of best response correspondences as

$$B(\sigma) = [B_i(\sigma_{-i})]_{i \in \mathcal{I}}.$$

- Clearly

$$B : \Sigma \rightrightarrows \Sigma.$$

Proof (continued)

- The idea is to apply Kakutani's theorem to the best response correspondence $B : \Sigma \rightrightarrows \Sigma$. We show that $B(\sigma)$ satisfies the conditions of Kakutani's theorem.

1. Σ is compact, convex, and non-empty.
   - By definition
   $$\Sigma = \prod_{i \in \mathcal{I}} \Sigma_i$$
   where each $\Sigma_i = \{x \mid \sum x_i = 1\}$ is a *simplex* of dimension $|S_i| - 1$, thus each Σ_i is closed and bounded, and thus compact. Their finite product is also compact.
2. $B(\sigma)$ is non-empty.
   - By definition,
   $$B_i(\sigma_{-i}) \in \arg\max_{x \in \Sigma_i} u_i(x, \sigma_{-i})$$
   where Σ_i is non-empty and compact, and u_i is linear in x. Hence, u_i is continuous, and by Weirstrass's theorem $B(\sigma)$ is non-empty.

Proof (continued)

3. $B(\sigma)$ is a convex-valued correspondence.
   - Equivalently, $B(\sigma) \subset \Sigma$ is convex if and only if $B_i(\sigma_{-i})$ is convex for all i. Let $\sigma_i', \sigma_i'' \in B_i(\sigma_{-i})$.
   - Then, for all $\lambda \in [0,1] \in B_i(\sigma_{-i})$, we have

$$u_i(\sigma_i', \sigma_{-i}) \geq u_i(\tau_i, \sigma_{-i}) \qquad \text{for all } \tau_i \in \Sigma_i,$$

$$u_i(\sigma_i'', \sigma_{-i}) \geq u_i(\tau_i, \sigma_{-i}) \qquad \text{for all } \tau_i \in \Sigma_i.$$

   - The preceding relations imply that for all $\lambda \in [0,1]$, we have

$$\lambda u_i(\sigma_i', \sigma_{-i}) + (1-\lambda) u_i(\sigma_i'', \sigma_{-i}) \geq u_i(\tau_i, \sigma_{-i}) \qquad \text{for all } \tau_i \in \Sigma_i.$$

   By the linearity of u_i,

$$u_i(\lambda\sigma_i' + (1-\lambda)\sigma_i'', \sigma_{-i}) \geq u_i(\tau_i, \sigma_{-i}) \qquad \text{for all } \tau_i \in \Sigma_i.$$

   Therefore, $\lambda\sigma_i' + (1-\lambda)\sigma_i'' \in B_i(\sigma_{-i})$, showing that $B(\sigma)$ is convex-valued.

Proof (continued)

4. $B(\sigma)$ has a closed graph.
   - Supposed to obtain a contradiction, that $B(\sigma)$ does not have a closed graph.
   - Then, there exists a sequence $(\sigma^n, \hat{\sigma}^n) \to (\sigma, \hat{\sigma})$ with $\hat{\sigma}^n \in B(\sigma^n)$, but $\hat{\sigma} \notin B(\sigma)$, i.e., there exists some i such that $\hat{\sigma}_i \notin B_i(\sigma_{-i})$.
   - This implies that there exists some $\sigma_i' \in \Sigma_i$ and some $\epsilon > 0$ such that

$$u_i(\sigma_i', \sigma_{-i}) > u_i(\hat{\sigma}_i, \sigma_{-i}) + 3\epsilon.$$

   - By the continuity of u_i and the fact that $\sigma_{-i}^n \to \sigma_{-i}$, we have for sufficiently large n,

$$u_i(\sigma_i', \sigma_{-i}^n) \geq u_i(\sigma_i', \sigma_{-i}) - \epsilon.$$

Proof (continued)

- [step 4 continued] Combining the preceding two relations, we obtain

$$u_i(\sigma'_i, \sigma^n_{-i}) > u_i(\hat{\sigma}_i, \sigma_{-i}) + 2\epsilon \geq u_i(\hat{\sigma}^n_i, \sigma^n_{-i}) + \epsilon,$$

  where the second relation follows from the continuity of u_i. This contradicts the assumption that $\hat{\sigma}^n_i \in B_i(\sigma^n_{-i})$, and completes the proof.

- The existence of the fixed point then follows from Kakutani's theorem.
- If $\sigma^* \in B(\sigma^*)$, then by definition σ^* is a mixed strategy equilibrium.

Existence of Equilibria for Infinite Games

- A similar theorem applies for pure strategy existence in infinite games.

Theorem

(Debreu, Glicksberg, Fan) *Consider an infinite strategic form game* $\langle \mathcal{I}, (S_i)_{i \in \mathcal{I}}, (u_i)_{i \in \mathcal{I}} \rangle$ *such that for each* $i \in \mathcal{I}$

1. S_i *is compact and convex;*
2. $u_i(s_i, s_{-i})$ *is continuous in* s_{-i}*;*
3. $u_i(s_i, s_{-i})$ *is continuous and concave in* s_i *[in fact quasi-concavity suffices].*

Then a pure strategy Nash equilibrium exists.

Definitions

- Suppose S is a convex set. Then a function $f : S \to \mathbb{R}$ is **concave** if for any $x, y \in S$ and any $\lambda \in [0, 1]$, we have

$$f(\lambda x + (1-\lambda)y) \geq \lambda f(x) + (1-\lambda) f(y).$$

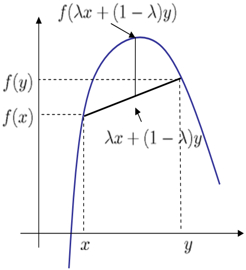

concave function

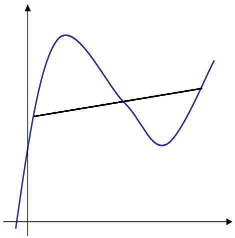

not a concave function

Proof

- Now define the best response correspondence for player i, $B_i : S_{-i} \rightrightarrows S_i$,

$$B_i\left(s_{-i}\right) = \left\{s_i' \in S_i \mid u_i(s_i', s_{-i}) \geq u_i(s_i, s_{-i}) \text{ for all } s_i \in S_i\right\}.$$

  Thus restriction to pure strategies.

- Define the set of best response correspondences as

$$B\left(s\right) = \left[B_i\left(s_{-i}\right)\right]_{i \in \mathcal{I}}.$$

  and

$$B : S \rightrightarrows S.$$

Proof (continued)

- We will again apply Kakutani's theorem to the best response correspondence $B : S \rightrightarrows S$ by showing that $B(s)$ satisfies the conditions of Kakutani's theorem.

1. S is compact, convex, and non-empty.
   - By definition
     $$S = \prod_{i \in \mathcal{I}} S_i$$
     since each S_i is compact [convex, nonempty] and finite product of compact [convex, nonempty] sets is compact [convex, nonempty].
2. $B(s)$ is non-empty.
   - By definition,
     $$B_i(s_{-i}) \in \arg\max_{s \in S_i} u_i(s, s_{-i})$$
     where S_i is non-empty and compact, and u_i is continuous in s by assumption. Then by Weirstrass's theorem $B(s)$ is non-empty.

Proof (continued)

3. $B(s)$ is a convex-valued correspondence.
   -
     - This follows from the fact that $u_i(s_i, s_{-i})$ is concave [or quasi-concave] in s_i. Suppose not, then there exists some i and some $s_{-i} \in S_{-i}$ such that $B_i(s_{-i}) \in \arg\max_{s \in S_i} u_i(s, s_{-i})$ is not convex.
     - This implies that there exists $s_i', s_i'' \in S_i$ such that $s_i', s_i'' \in B_i(s_{-i})$ and $\lambda s_i' + (1-\lambda)s_i'' \notin B_i(s_{-i})$. In other words,

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

       But this violates the concavity of $u_i(s_i, s_{-i})$ in s_i [recall that for a concave function $f(\lambda x + (1-\lambda)y) \geq \lambda f(x) + (1-\lambda) f(y)$].
     - Therefore $B(s)$ is convex value.

4. The proof that $B(s)$ has a closed graph is identical to the previous proof in is left for the homework.

Existence of Nash Equilibria

- Can we relax concavity?
- Example: Consider the game where two players pick a location $s_1, s_2 \in \mathbb{R}^2$ on the circle. The payoffs are $u_1(s_1, s_2) = -u_2(s_1, s_2) = d(s_1, s_2)$, where $d(s_1, s_2)$ denotes the Euclidean distance between $s_1, s_2 \in \mathbb{R}^2$.
- No pure Nash equilibrium.
- However, it can be shown that the strategy profile where both mix uniformly on the circle is a mixed Nash equilibrium.

A More Powerful Theorem

Theorem

(Glicksberg) *Consider an infinite strategic form game* $\langle \mathcal{I}, (S_i)_{i \in \mathcal{I}}, (u_i)_{i \in \mathcal{I}} \rangle$ *such that for each* $i \in \mathcal{I}$

1. S_i *is compact and convex;*
2. $u_i(s_i, s_{-i})$ *is continuous in* s_i *and* s_{-i}*.*

Then a mixed strategy Nash equilibrium exists.

- The proof of this theorem is harder and we will not discuss it.
- In fact, finding mixed strategies in continuous games is more challenging and is beyond the scope of this course.

Extensive Form Games

- Extensive-form games model multi-agent sequential decision making.
- For now, we will focus is on multi-stage games with observed actions.
- Extensive form represented by **game trees**.
- Additional component of the model, histories (i.e., sequences of action profiles).
- Extensive form games will be useful when we analyze **dynamic** games, in particular, to understand issues of cooperation and trust in groups.

Histories

- Let H^k denote the set of all possible stage-k histories
- Strategies are maps from all possible histories into actions: $s_i^k : H^k \rightarrow S_i$

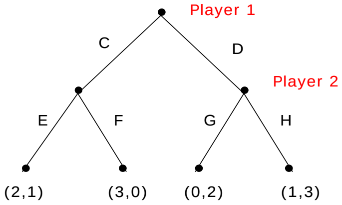

- Example:
  - Player 1's strategies: $s_1 : H^0 = \emptyset \rightarrow S_1$; two possible strategies: C,D
  - Player 2's strategies: $s^2 : H^1 = \{C, D\} \rightarrow S_2$; four possible strategies: EG,EH,FG, FH

Strategies in Extensive Form Games

- Consider the following two-stage extensive form version of matching pennies.

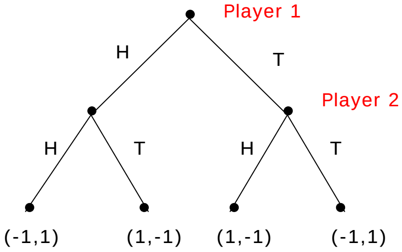

- How many strategies does player 2 have?

Strategies in Extensive Form Games (continued)

- Recall: strategy should be a *complete contingency plan.*
- Therefore: player 2 has four strategies:
  1. heads following heads, heads following tails (HH,HT);
  2. heads following heads, tails following tails (HH, TT);
  3. tails following heads, tails following tails (TH, TT);
  4. tails following heads, heads following tails (TH, HT).

Strategies in Extensive Form Games (continued)

- Therefore, from the extensive form game we can go to the strategic form representation. For example:

Player 1/Player 2	(HH, HT)	(HH, TT)	(TH, TT)	(TH, HT)
heads	$(-1, 1)$	$(-1, 1)$	$(1, -1)$	$(1, -1)$
tails	$(1, -1)$	$(-1, 1)$	$(-1, 1)$	$(1, -1)$

- So what will happen in this game?

Strategies in Extensive Form Games (continued)

- Can we go from strategic form representation to an extensive form representation as well?
- To do this, we need to introduce information sets. If two nodes are in the same information set, then the player making a decision at that point cannot tell them apart. The following two extensive form games are representations of the simultaneous-move matching pennies. These are imperfect information games.
- Note: For consistency, first number is still player 1's payoff.

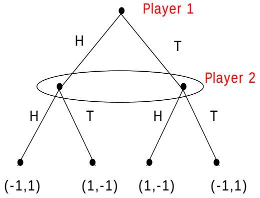

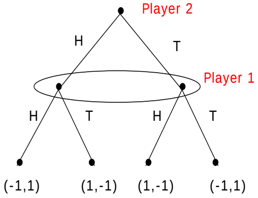

Entry Deterrence Game

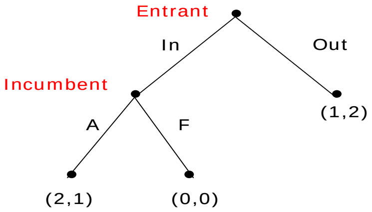

- Equivalent strategic form representation

Entrant\Incumbent	Accommodate	Fight
In	$(2,1)$	$(0,0)$
Out	$(1,2)$	$(1,2)$

- Two pure Nash equilibria: (In,A) and (Out,F).

Are These Equilibria Reasonable?

- The equilibrium (Out,F) is sustained by a **noncredible threat** of the monopolist
- Equilibrium notion for extensive form games: Subgame Perfect (Nash) Equilibrium
- It requires each player's strategy to be "optimal" not only at the start of the game, but also after every history
- For finite horizon games, found by *backward induction*
- For infinite horizon games, characterization in terms of **one-stage deviation principle.**

Subgames

- Recall that a game G is represented by a game tree. Denote the set of nodes of G by V_G.
- A game has **perfect information** if all its information sets are singletons (i.e., all nodes are in their own information set).
- Recall that history h^k denotes the play of a game after k stages. In a perfect information game, each node $v \in V_G$ corresponds to a unique history h^k and vice versa. This is not necessarily the case in imperfect or incomplete information games.
- We say that a node $x \in V_G$ is a successor of node $y \in V_G$, or $y \succ x$, if in the game tree we reach y through x.

Subgames (continued)

Definition

(Subgames) *A subgame G' of game G is given by the set of nodes $V_G^x \subset V_G$ in the game tree of G that are successors of some node $x \in V_G^x$ [i.e., for all $y \in V_G^x$, we have $y \succ x$] and are not successors of any node $y \notin V_G^x$ [i.e., for any $z \in V_G^x$ if there exists y such that $z \succ y$, then $y \succ x$].*

- A restriction of a strategy s subgame G', $s_{|G'}$ is the action profile implied by s in the subgame G'.

Subgames: Examples

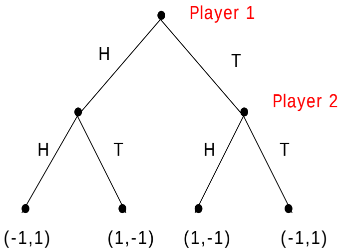

- Recall the two-stage extensive-form version of the matching pennies game
- In this game, there are two proper subgames and the game itself which is also a subgame, and thus a total of three subgames.

Subgame Perfect Equilibrium

Definition

(Subgame Perfect Equilibrium) *A strategy profile s^* is a Subgame Perfect Nash equilibrium (SPE) in game G if for any subgame G' of G, ${s^*}_{|G'}$ is Nash equilibrium of G'.*

- Loosely speaking, subgame perfection will remove noncredible threats, since these will not be Nash equilibria in the appropriate subgames.
- In the entry deterrence game, following entry, F is not a best response, and thus not a Nash equilibrium of the corresponding subgame. Therefore, (Out,F) is not a SPE.
- How to find SPE? One could find all of the Nash equilibria, for example as in the entry deterrence game, then eliminate those that are not subgame perfect.
- But there are more economical ways of doing it.

Backward Induction

- Backward induction refers to starting from the last subgames of a finite game, then finding the Nash equilibria or best response strategy profiles in the subgames, then assigning these strategies profiles to be subgames, and moving successively towards the beginning of the game.

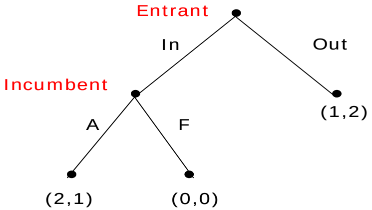

Backward Induction (continued)

Theorem

Backward induction gives the entire set of SPE.

Proof: backward induction makes sure that in the restriction of the strategy profile in question to any subgame is a Nash equilibrium.

Existence of Subgame Perfect Equilibria

Theorem

Every finite perfect information extensive form game G has a pure strategy SPE.

Proof: Start from the end by backward induction and at each step one strategy is best response.

Theorem

Every finite extensive form game G has a SPE.

Proof: Same argument as the previous theorem, except that some subgames need not have perfect information and may have mixed strategy equilibria.

Examples: Value of Commitment

- Consider the entry deterrence game, but with a different timing as shown in the next figure.

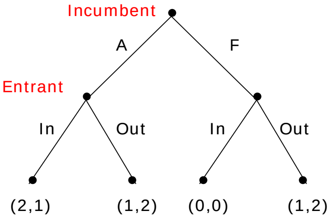

- Note: For consistency, first number is still the entrant's payoff.
- This implies that the incumbent can now commit to fighting (how could it do that?).
- It is straightforward to see that the unique SPE now involves the incumbent committing to fighting and the entrant not entering.

Examples: Stackleberg Model of Competition

- Consider a variant of the Cournot model where player 1 chooses its quantity q_1 first, and player 2 chooses its quantity q_2 after observing q_1. Here, player 1 is the Stackleberg leader.
- Suppose again that both firms have marginal cost c and the inverse demand function is given by $P(Q) = \alpha - \beta Q$, where $Q = q_1 + q_2$, where $\alpha > c$.
- This is a dynamic game, so we should look for SPE. How to do this?
- Backward induction—this is not a finite game, but all we have seen so far applies to infinite games as well.
- Look at a subgame indexed by player 1 quantity choice, q_1. Then player 2's maximization problem is essentially the same as before

$$\begin{aligned} \max_{q_2 \geq 0} \pi_2(q_1, q_2) &= [P(Q) - c]\, q_2 \\ &= [\alpha - \beta(q_1 + q_2) - c]\, q_2. \end{aligned}$$

Stackleberg Competition (continued)

- This gives best response

$$q_2 = \frac{\alpha - c - \beta q_1}{2\beta}.$$

- Now the difference is that player 1 will choose q_1 recognizing that player 2 will respond with the above best response function.
- Player 1 is the Stackleberg leader and player 2 is the follower.

Stackleberg Competition (continued)

- This means player 1's problem is

$$\begin{aligned} \text{maximize}_{q_1 \geq 0} \quad & \pi_1(q_1, q_2) = [P(Q) - c]\, q_1 \\ \text{subject to} \quad & q_2 = \frac{\alpha - c - \beta q_1}{2\beta}. \end{aligned}$$

- Or

$$\max_{q_1 \geq 0} \left[\alpha - \beta\left(q_1 + \frac{\alpha - c - \beta q_1}{2\beta}\right) - c\right] q_1.$$

Stackleberg Competition (continued)

- The first-order condition is

$$\left[\alpha - \beta\left(q_1 + \frac{\alpha - c - \beta q_1}{2\beta}\right) - c\right] - \frac{\beta}{2} q_1 = 0,$$

which gives

$$q_1^S = \frac{\alpha - c}{2\beta}.$$

- And thus

$$q_2^S = \frac{\alpha - c}{4\beta} < q_1^S$$

- Why lower output for the follower?
- Total output is

$$Q^S = q_1^S + q_2^S = \frac{3(\alpha - c)}{4\beta},$$

which is greater than Cournot output. Why?