Game Theory Lecture 11 & 12

Fredrik Magnusson

Department of Automatic Control
Faculty of Engineering
Lund University, Sweden

October 10, 2014

Outline

1. Lecture 11: Learning in games
2. Lecture 12: Extensive form games

Outline

Fictitious play

- Consider a finite game that is played repeatedly in discrete time.
- The basic idea of fictitious play is that each player assumes that his opponents are using a fixed mixed strategy, and updates his beliefs about these strategies at each time step.
- Players choose actions in each time step to maximize that step's expected payoff given their belief of the opponents's strategies.
- The belief is formed as the empirical frequency distribution of the opponents's previously played strategies.

Properties of fictitious play

- Myopic, because players are trying to maximize current payoff without considering their future payoffs. They are also not learning the “true model” generating the empirical frequencies (that is, how their opponent is actually playing the game).
- Not a unique rule due to multiple best responses. Traditional analysis assumes player chooses any of the pure best responses.
- Players do not need to know about their opponents’s payoff; they only form beliefs about how their opponents will play.

Convergence

It is of interest to study when fictitious play converges, and if so, to what. Let $\{s^t\}_{t=0}^{\infty}$ be a sequence in S generated by fictitious play. Each player i needs some initial belief s_i^0 is needed to get started.

Definition (Convergence to pure strategy)

The sequence $\{s^t\}$ converges to $s^\infty \in S$ if

$$\exists T \in \mathbb{N} : \forall t \geq T, s^t = s^\infty.$$

Theorem (Convergence and Nash equilibria)

1. *If s^∞ exists, then it is a Nash equilibrium.*
2. *If s^* is a Nash equilibrium and $s^T = s^*$, then $s^t = s^*$ for all $t \geq T$.*

Convergence to mixed strategies

Fictitious play will sometimes converge to a mixed strategy rather than a pure one.

Definition (Convergence to mixed strategy)

The sequence $\{s^t\}$ converges to $\sigma \in \Sigma$ if

$$\forall i \in \mathcal{I}, \forall s_i \in S_i, \lim_{T \to \infty} \frac{\sum_{t=0}^{T-1} \delta(s_i^t - s_i)}{T} = \sigma_i(s_i).$$

Theorem (Convergence and Nash equilibria)

If $\{s^t\}$ converges to $\sigma \in \Sigma$, then σ is a Nash equilibrium.

Convergence conditions

Theorem (Convergence conditions)

Fictitious play converges if any of the following holds.

1. G *is a two-player zero-sum game.*
2. G *is a two-player game with* $|S_i| = 2$.
3. G *is solvable by iterated strict dominance.*
4. $\forall i, j \in \mathcal{I}, u_i = u_j$.
5. G *is a potential game.*

Miscoordination

Fictitious play may converge to a mixed Nash equilibrium, but continued execution of fictitious play may yield very different payoffs than simply repeatedly playing the strategies in the corresponding Nash equilibrium.

See lecture notes for illustrative example.

Continuous time fictitious play

Let $p_i^t : S_i \to \mathbb{R}_+$ be the empirical probability distribution of the plays by player i up until time t. This is determined by a differential equation involving the best response to the opponents's play up until time t, denoted by $BR_i(p_{-i}^t)$, by two variants (CTFP and perturbed CTFP):

1. $\frac{\mathrm{d}p_i^t}{\mathrm{d}t} \in BR_i(p_{-i}^t) - p_i^t,$

   $BR_i(p_{-i}^t) = \arg\max_{\sigma_i \in \Sigma_i} u_i(\sigma_i, p_{-i}^t).$

2. $\frac{\mathrm{d}p_i^t}{\mathrm{d}t} = C_i(p_{-i}^t) - p_i^t,$

   $C_i(p_{-i}^t) = \arg\max_{\sigma_i \in \Sigma_i} \left(u_i(\sigma_i, p_{-i}^t) - V_i(\sigma_i)\right),$

   $V_i : \Sigma_i \to \mathbb{R} \wedge$ strictly convex $\wedge$ boundary condition.

Perturbation V_i is used to make the best response unique.

Questions

- Do the proofs for CTFP carry over to DTFP?
- Is (perturbed) CTFP just a theoretical tool, or can it be employed in practice? Would you want to (e.g. to avoid miscoordination)?

  Does p_i^t actually tell you what strategy to play? How to update your beliefs about the opponents?
- Can DTFP be extended to infinite games? Can CTFP?

$$\frac{\partial}{\partial t} p_i(s,t) = C_i(s, p_{-i}^t) - p_i(s,t)$$

$$C_i(s, p_{-i}^t) = \arg\max_{\sigma_i \in \Sigma_i} \Big(u_i(\sigma_i, p_{-i}^t) - V_i(\sigma_i) \Big)\Big|_s$$

- For perturbed CTFP, p^∞ is seemingly a Nash equilibrium iff for all i, $V_i(p_i^\infty)$ is very small compared to $u_i(p_i^\infty, p_{-i}^\infty)$. Is this the case? If so, is that a problem (i.e., is it hard to find a useful V)?

Outline

1. Lecture 11: Learning in games

2. Lecture 12: Extensive form games

Extensive-form games

- We have so far considered games where each player i chooses a single (possibly mixed) strategy out of their action set S_i (or Σ_i).
- We have also considered situations where the same game is played repeatedly.
- Extensive-form games:
  - Divided into K (possibly infinite) stages.
  - In each stage k, only a subset of the players choose an action a^k.
  - The available actions depend on k and the history of previous actions $h^k = (a^0, a^1, \ldots, a^{k-1})$.

Strategies

- Strategies are contingency plans for every possible history h^k.
- That is, a pure strategy is a set of maps $s_i := \{s_i^k\}_{k=0}^K$ from all possible histories to available actions; that is,

$$s_i^k : H^k \to S_i(H^k),$$

where $H^k = \{h_j^k\}$.
- A mixed strategy (behavior strategy) is a set of maps from all possible histories to a probability measure on the available actions.

Utility functions

Utility function for player i can be defined

- on the set of possible strategies; that is

$$u_i : S \to \mathbb{R}.$$

- or on the set of final outcomes of the game (since each pure strategy determines an outcome); that is

$$u_i : H^{K+1} \to \mathbb{R}.$$

Tree

Extensive-form games are conveniently represented using trees.
Example shamelessly stolen from lecture slides:

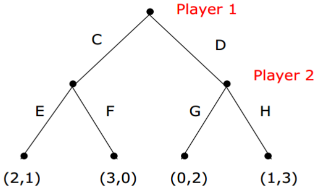

Player 1's strategies (two): $s_1 : H^0 = \emptyset \rightarrow S_1(H^0) = \{C, D\}$

Player 2's strategies (four):

$$s_2 : H^1 = \{\{C\}, \{D\}\} \rightarrow S_2(H^1) = \{E, F, G, H\}$$
$$\wedge \, \forall h_j^1 \in H^1, s_2(h^1) \in S_2(h^1)$$

The four strategies for player 2 are EG, EH, FG, and FH.

History information

- A game in which the players that need to choose an action at stage k have full information about h^k is said to have perfect information.
- Information sets $h \in H$ are introduced to model games without perfect information.
- Details omitted, but this is used to e.g. model that players choose action simultaneously: player i acts first and then player j acts without knowing what player i did.

Exogenous events

- Book also introduces exogenous events (moves by ”Nature”), which is not covered in the lecture slides.
- In some stages, no player chooses an action, but rather some random event occurs which will affect future stages.

Strategic- vs. extensive-form games

My interpretation of Section 3.4 Strategies and Equilibria in Extensive-Form Games:

- Strategic-form games can be transformed to extensive-form games of perfect recall (players never lose information that they once had) and vice versa.
- Most important part of proving this is

Theorem (3.1)

In a game of perfect recall, mixed strategies in the strategic form are equivalent to mixed strategies in the extensive form.

- We can thus use our theory from strategic form games to analyze extensive form games.

Equilibria

Theorem (3.2)

A finite game of perfect information has a pure-strategy Nash equilibrium.

- For finite strategic form games, we can only guarantee the existence of a mixed-strategy Nash equilibrium.
- As we already know, some Nash equilibria are unlikely to be played in a real situation (e.g. if they are not evolutionarily stable).
- The equilibrium concept for extensive-form games can be refined to get rid of (some) “unlikely” equilibria: subgame-perfect equilibria.

Subgame-perfect equilibrium

- Loosely speaking, a subgame of an extensive-form game G is a subtree of G in which every player knows that they stay in the subtree once they reach it.
- Thus, in a game of perfect information, every node in the tree corresponds to a subgame.

Definition (Subgame-perfect equilibrium)

A mixed strategy σ of an extensive-form game is a subgame-perfect equilibrium if for every subgame G the restriction of σ to G is a Nash equilibrium of G.

Finding subgame-perfect equilibria

- For finite games, all subgame-perfect equilibria can be found using backward induction.
- Backward induction is essentially dynamic programming: find the Nash equilibria of the last (smallest) subgame, associate the strategies and payoffs of the equilibria with the top node of the subgame, and proceed to the next smallest subgame.

Existence of subgame-perfect equilibria

Theorem (3.2, refined)

A finite game of perfect information has a pure-strategy subgame-perfect equilibrium.

Theorem (Existence without perfect information)

A finite game has a mixed-strategy subgame-perfect equilibrium.

One-stage deviation principle

- If $K = \infty$, we have no algorithm for finding all subgame-perfect equilibria.
- We do however have a useful characterization if the game has perfect information, which is essentially the principle of optimality:

Theorem (One-stage deviation principle)

If the game is continuous at infinity and has perfect information, then s^* is a subgame-perfect equilibrium if and only if

$$\forall i \in \mathcal{I}, \forall k \in \mathbb{Z}_+, \forall l \in \mathbb{Z}_{++}, \forall h^k \in H^k, \forall h^{k+l} \in H^{k+l},$$
$$\forall s_i \in S_i : s_i(h^k) \neq s_i^*(h^k) \wedge s_i(h^{k+l}) = s_i^*(h^{k+l}),$$
$$u_i(s_i^*, s_{-i}^* | h^k) \geq u_i(s_i, s_{-i}^* | h^k).$$

- Loosely speaking, a strategy is a subgame-perfect equilibrium if and only if no deviation in a single stage yields a higher payoff.