Fictitious Play

with a Continuum of Anonymous Players

by Zinovi Rabinovich, Enrico Gerding, Maria Polukarov and Nick Jennings

University of Southampton

Agenda

- Background
  - Games in normal form
  - Fictitious Play (FP)
  - Auctions as typed games
- FP for typed games
- Continuum of Anonymous Players (CAPs)
- FP for CAPs
  - Linear structure of utilities

Games in normal form

- Game is a tuple $< N, \{A_i\}_{i=\in N}, \{u_i\}_{i \in N} >$
  - N is a set of player tags
  - A_i is a set of actions available to player $i \in N$
  - $u_i : \bigotimes_{i \in N} A_i \rightarrow \mathbf{R}$ is the utility function of player $i \in N$

Games in normal form

- Game is a tuple $< N, \{A_i\}_{i=\in N}, \{u_i\}_{i \in N} >$
  - N is a set of player tags
  - A_i is a set of actions available to player $i \in N$
  - $u_i : \bigotimes_{i \in N} A_i \rightarrow \mathbf{R}$ is the utility function of player $i \in N$
- Game play
  - All players simultaneously choose $a_i \in A_i$
  - The actions are combined into a joint action profile $\mathbf{a} = (a_i)_{i \in N}$
  - Each player receives utility $u_i(\mathbf{a})$

Games in normal form

- Game is a tuple $< N, \{A_i\}_{i=\in N}, \{u_i\}_{i \in N} >$
  - N is a set of player tags
  - A_i is a set of actions available to player $i \in N$
  - $u_i : \bigotimes_{i \in N} A_i \rightarrow \mathbf{R}$ is the utility function of player $i \in N$
- Game play
  - All players simultaneously choose $a_i \in A_i$
  - The actions are combined into a joint action profile $\mathbf{a} = (a_i)_{i \in N}$
  - Each player receives utility $u_i(\mathbf{a})$
- Game's purpose
  - Each player wishes to maximise its utility

Game's Purpose: Maximise how?

- MaxMin solution: be ready for the worst
  - $a_i^* = \arg\max_{a_i} \min_{a_{-i}} u_i(a_i, a_{-i})$

Game's Purpose: Maximise how?

- MaxMin solution: be ready for the worst
  - $a_i^* = \arg\max_{a_i} \min_{a_{-i}} u_i(a_i, a_{-i})$
- Nash equilibrium $\mathbf{a}^*$: no (single) player can do better
  - $\forall i \in N \;\; \forall a_i \in A_i \quad u_i(\mathbf{a}^*) \geq u_i(a_i, \mathbf{a}^*_{-i})$

Game's Purpose: Maximise how?

- MaxMin solution: be ready for the worst
  - $a_i^* = \arg\max_{a_i} \min_{a_{-i}} u_i(a_i, a_{-i})$
- Nash equilibrium $\mathbf{a}^*$: no (single) player can do better
  - $\forall i \in N \;\; \forall a_i \in A_i \quad u_i(\mathbf{a}^*) \geq u_i(a_i, \mathbf{a}^*_{-i})$
- Mixed Nash equilibrium π^*: choose a good (rigged) dice
  - Actions are selected at random $\pi_i \in \Delta(A_i)$
  - Joint mixed profile $\pi = (\pi_i)_{i \in N}$: $\pi(\mathbf{a}) = \prod_{i \in N} \pi_i(a_i)$
  - $\forall i \in N \;\; \forall \pi_i \in \Delta(A_i) \quad \mathbf{E}_{\pi^*}[u_i] \geq \mathbf{E}_{(\pi_i, \pi^*_{-i})}[u_i]$

Find Mixed Nash: Fictitious Play

- Idea – Statisticians' Fun
  - Assume fixed (unknown) π_{-i}
  - Play the game repeatedly
  - Form (online) history estimate $\pi_{-i} = \sum_{k=1}^{t} \delta(a_i^t)$
  - Choose actions to optimise against π_{-i}

Find Mixed Nash: Fictitious Play

- Repeat until convergence
  - Policy estimates at time t are $\{\pi_i^t\}_{i \in N}$
  - Choose $a_i^{t+1} = \arg\max_{a_i \in A_i} u_i(a_i, \pi_{-i})$
  - Update $\pi_i^{t+1} = \frac{t-1}{t}\pi_i^t + \frac{1}{t}\delta(a_i^{t+1})$

Find Mixed Nash: Fictitious Play

- Repeat until convergence
  - Policy estimates at time t are $\{\pi_i^t\}_{i \in N}$
  - Choose $a_i^{t+1} = \arg\max_{a_i \in A_i} u_i(a_i, \pi_{-i})$
  - Update $\pi_i^{t+1} = \frac{t-1}{t}\pi_i^t + \frac{1}{t}\delta(a_i^{t+1})$
- Twofold convergence
  - In action selection to Nash $\mathbf{a}^*$:
    $\exists T > 0 \;\; and \;\; \mathbf{a}^* \;\; s.t. \;\; \forall t > T \;\; \mathbf{a}^t = \mathbf{a}^*$
  - In beliefs to mixed Nash π^*: $\exists \pi^* = \lim_{t \to \infty} \pi^t$

FP Convergence: mixed Nash

- Two players
  - In zero sum games
  - In $2xK$ games
- N players
  - Identical interest games

Sealed Second Price Auctions

- Protocol:
  - N players wish to acquire an item
  - Each submits a bid $b_i \in A_i$
  - The winner is $i = \arg\max_{i \in N} b_i$
  - The looser pay nothing
  - The winner pays $c = \max_{j \neq i \in N} b_j$

Sealed Second Price Auctions

- Protocol:
  - N players wish to acquire an item
  - Each submits a bid $b_i \in A_i$
  - The winner is $i = \arg\max_{i \in N} b_i$
  - The looser pay nothing
  - The winner pays $c = \max_{j \neq i \in N} b_j$
- Questions:
  - Why would players $i, j \in N$ bid $b_i \neq b_j$?
  - Is there a significance to the identity of a bid?

Sealed Second Price Auctions

- Protocol:
  - N players wish to acquire an item
  - Each submits a bid $b_i \in A_i$
  - The winner is $i = \arg\max_{i \in N} b_i$
  - The looser pay nothing
  - The winner pays $c = \max_{j \neq i \in N} b_j$
- Properties:
  - Why would players $i, j \in N$ bid $b_i \neq b_j$?
  - Is there a significance to the identity of a bid?

Sealed Second Price Auctions

- Protocol:
  - N players wish to acquire an item
  - Each submits a bid $b_i \in A_i$
  - The winner is $i = \arg\max_{i \in N} b_i$
  - The looser pay nothing
  - The winner pays $c = \max_{j \neq i \in N} b_j$
- Properties:
  - Players have types $\alpha_i \in T$: $u_{gen}(\alpha_i, \mathbf{a}) \rightarrow \mathbf{R}$
    - A distribution over T is given
  - Is there a significance to the identity of a bid?

Sealed Second Price Auctions

- Protocol:
  - N players wish to acquire an item
  - Each submits a bid $b_i \in A_i$
  - The winner is $i = \arg\max_{i \in N} b_i$
  - The looser pay nothing
  - The winner pays $c = \max_{j \neq i \in N} b_j$
- Properties:
  - Players have types $\alpha_i \in T$: $u_{gen}(\alpha_i, \mathbf{a}) \rightarrow \mathbf{R}$
  - Utility depends on the set of anonymous bids

Typed games

- Game setting $< N, \{A_i\}_{i=\in N}, \{u_i\}_{i\in N}, T, \sigma >$
  - T is the space of private types
  - σ is a distribution over T
  - $u_i : T \times \bigotimes_{i\in N} A_i \rightarrow \mathbf{R}$

Typed games

- Game setting $< N, \{A_i\}_{i=\in N}, \{u_i\}_{i\in N}, T, \sigma >$
  - T is the space of private types
  - σ is a distribution over T
  - $u_i : T \times \bigotimes_{i\in N} A_i \rightarrow \mathbf{R}$
- Game play:
  - $\alpha_i \in T$ is sampled from σ for each player,
  - Utilities are fixed $u_i(\mathbf{a}) = u_i(\alpha_i, \mathbf{a})$
  - Normal form game starts

Typed games

- Game setting $< N, \{A_i\}_{i=\in N}, \{u_i\}_{i\in N}, T, \sigma >$
  - T is the space of private types
  - σ is a distribution over T
  - $u_i : T \times \bigotimes_{i\in N} A_i \rightarrow \mathbf{R}$
- Game play:
  - $\alpha_i \in T$ is sampled from σ for each player,
  - Utilities are fixed $u_i(\mathbf{a}) = u_i(\alpha_i, \mathbf{a})$
  - Normal form game starts
- Solution concept:
  - Policy is conditioned on type $\pi_i : T \rightarrow \Delta(A_i)$
  - Equilibrium is now in expectation over T as well $\mathbf{E}_{\pi^*,\sigma}[u_i] \geq \mathbf{E}_{(\pi_i,\pi^*_{-i}),\sigma}[u_i]$

Auctions vs. Typed games

- Auctions are typed games

Auctions vs. Typed games

- Auctions are typed games, but also commonly:
  - Anonymous
  - A_i is discrete and finite
    - Money is finite and discrete
  - T is continuous
    - Evaluations may be in fractions of a pence

Auctions vs. Typed games

- Auctions are typed games, but also commonly:
  - Anonymous
  - A_i is discrete and finite
    - Money is finite and discrete
  - T is continuous
    - Evaluations may be in fractions of a pence
- Recall how play starts:
  - $\alpha_i \in T$ is sampled from σ for each player,
  - Utilities are fixed $u_i(\mathbf{a}) = u_i(\alpha_i, \mathbf{a})$

Auctions vs. Typed games

- Auctions are typed games, but also commonly:
  - Anonymous
  - A_i is discrete and finite
    - Money is finite and discrete
  - T is continuous
    - Evaluations may be in fractions of a pence
- Recall how play starts:
  - $\alpha_i \in T$ is sampled from σ for each player,
  - Utilities are fixed $u_i(\mathbf{a}) = u_i(\alpha_i, \mathbf{a})$
- Type and player tag are a name for an utility function

Auctions and CAPs

- Continuum of Anonymous Players (CAPs) captures a vast range of auction problems
  - First and second price auctions
  - Multiple simultaneous or sequential

Auctions and CAPs

- Continuum of Anonymous Players (CAPs) captures a vast range of auction problems
  - First and second price auctions
  - Multiple simultaneous or sequential
- Algorithm to compute an equilibrium in CAPs is a generic solution
  - Previous algorithms have narrow specialisation

Auctions and CAPs

- Continuum of Anonymous Players (CAPs) captures a vast range of auction problems
  - First and second price auctions
  - Multiple simultaneous or sequential
- Algorithm to compute an equilibrium in CAPs is a generic solution
  - Previous algorithms have narrow specialisation
- Fictitious Play algorithm for CAPs

CAPs

- Anonymous games
  - Given a set of actions $\{a_i\}_{i \in I}$ taken by the opponents
  - A players utility depends on the composition of the set, but not on the identity of opponents
- Continuum of players
  - The set of players in the game is infinite, the cardinality of the set is $\aleph$
  - A finite set of players is randomly selected from a cardinality $\aleph$ set.

Example: Private Value Auctions

- An item is auctioned to n players
- Each player has a private value for the item
  - The value is not known to opponents
  - Value obtained by sampling from a common distribution
    - E.g. uniform over a $[0, 1]$ range.

Example: Private Value Auctions

- An item is auctioned to n players
- Utility from winning
  - Depends on bids composition, not bidders identity
- Probability of winning
  - Engendered by probability distributions of opponent bids
  - Not specific choices of bidders

CAPs (cont)

- A is a compact metric space
- $\mathcal{M}$ is the space of probability distributions over A
- $\mathcal{U}$ is the space of continuous functions $u : A \times \mathcal{M} \rightarrow \mathbf{R}$
  - This is the set of possible players
  - For auctions this is utility given private value
    - Winning by placing a bid $b \in A$
    - Against opponent bids distributed by $ob \in \mathcal{M}$
- Anonymous game with continuum of players is then characterised by a distribution μ over the space $\mathcal{U}$.
  - For auctions this corresponds to the distribution of private values

Cournot-Nash Equilibrium (CNE)

- A distribution τ over $\mathcal{U} \times A$ so that
  - $\tau_{\mathcal{U}} = \mu$
  - $\tau\{(u, a) | u(a, \tau_A) \geq u(A, \tau_A)\} = 1$
- If there is indeed a continuum of players
  - Single player action choice will not modify τ_A
  - Any modification away from the CNE can only reduce utility

Cournot-Nash Equilibrium (CNE)

- A distribution τ over $\mathcal{U} \times A$ so that
  - $\tau_{\mathcal{U}} = \mu$
  - $\tau\{(u, a) | u(a, \tau_A) \geq u(A, \tau_A)\} = 1$
- For auctions the continuum of players is virtual
  - Player faces <u>unknown</u> values of opponents
  - Player's response is computed with respect to
    - <u>continuum</u> of private values
    - <u>distribution</u> over the range of values

Existence of CNE

- **Theorem** (Mas-Colell, 84):
  If μ is atomless and A is discrete and finite, then there exists a pure CNE.
- Pure CNE is characterised by a function $h : \mathcal{U} \to A$
  - Inherently symmetric: players with the same type choose same action

Fictitious Play

- Iterative process of adaptation
  - Given a history of opponents play
  - Compute a frequency of action for each player
  - Compute and apply best response to that frequency
- Properties:
  - If best response converges, then to a pure Nash
  - If frequency estimate converges, then to a mixed Nash equilibrium

Fictitious Play

- Iterative process of adaptation
  - Given a history of opponents play
  - Compute a frequency of action for each player
  - Compute and apply best response to that frequency
- Properties:
  - For anonymous continuum games the procedure becomes symmetric
  - If frequency converges, a pure CNE can be obtained from it

FP in CAPs

- Frequency of actions (for all types)
  - A distribution τ over $\mathcal{U} \times A$
- FP computes best response
  - A function $h : \mathcal{U} \rightarrow A$
  - Inherent $\tau_h(\mathbf{u}, a) = \mu(h^{-1}(a) \cap \mathbf{u})$
- Updating frequency estimate
  - $\tau = \alpha * \tau + (1 - \alpha) * \tau_h$

FP in CAPs

- Best response computation depends only on τ_A, the marginal distribution
- The algorithm can be simplified
  - Frequency of actions τ_A
  - Compute best response $h : \mathcal{U} \rightarrow A$
  - Compute $\tau_{h,A}(a) = \mu(h^{-1}(a))$
  - Update $\tau_A = \alpha * \tau_A + (1 - \alpha) * \tau_{h,A}$
- τ_A can be transformed into a complete equilibrium

FP in CAPs

- Best response computation depends only on τ_A, the marginal distribution
- The algorithm can be simplified
  - Frequency of actions τ_A
  - Compute best response $h : \mathcal{U} \rightarrow A$
  - Compute $\tau_{h,A}(a) = \mu(h^{-1}(a))$
  - Update $\tau_A = \alpha * \tau_A + (1 - \alpha) * \tau_{h,A}$
- τ_A can be transformed into a complete equilibrium

Best Response Computation

- If $\mathcal{U}$ is linear, single dimensional (e.g. in auctions)
  - One can associate a function $u_{gen} : T \times A \times \mathcal{M} \rightarrow R$ linear in type T
  - Given a fixed τ_A, utility of an action $a \in A$ becomes linear in type
- Best response is:

$$h(\lambda \in T) = \arg\max_{a \in A} u_{gen}(\lambda, a, \tau_A)$$

Best Response Computation

- Best response is: $h(\lambda \in T) = \arg\max_{a \in A} u_{gen}(\lambda, a, \tau_A)$

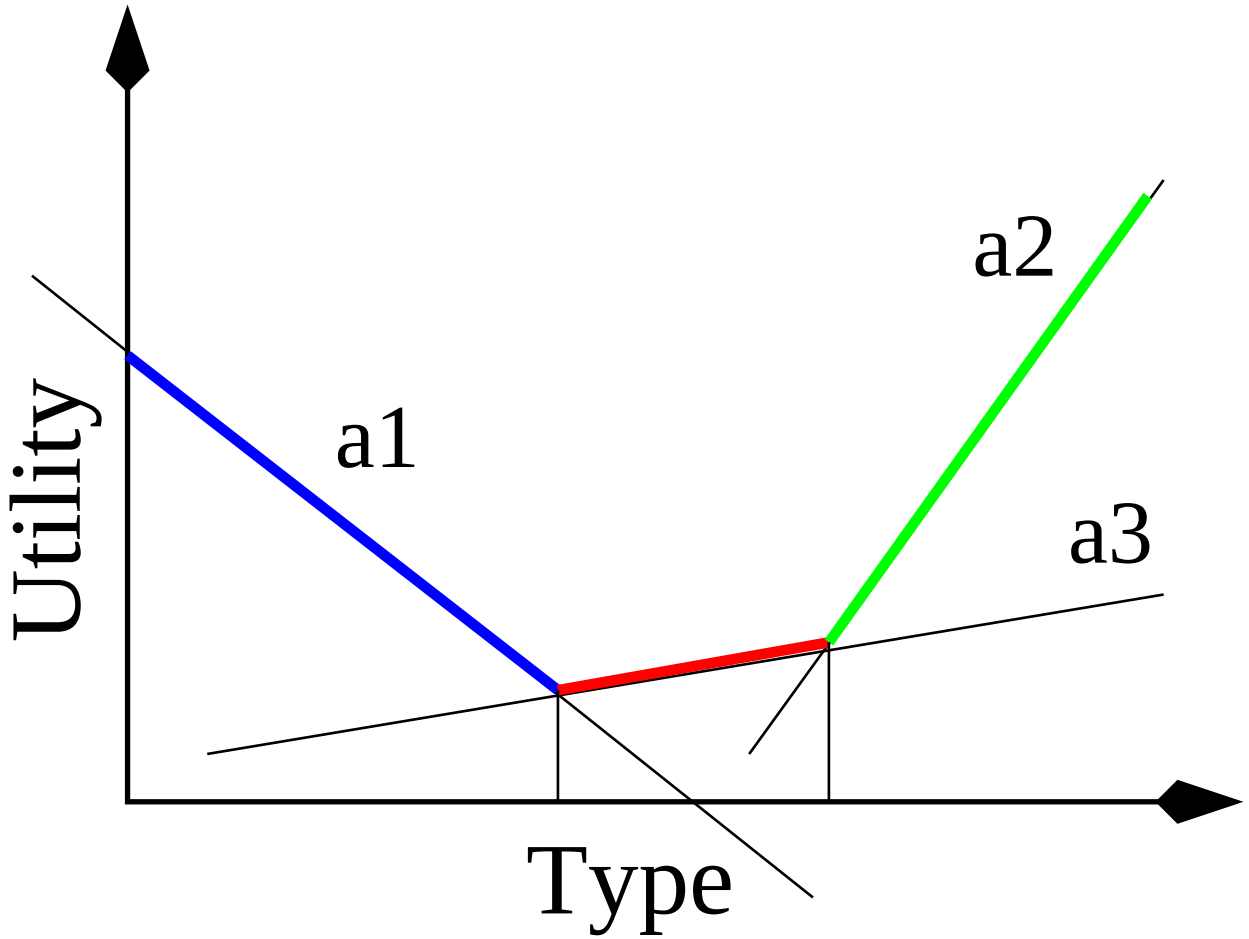

Best Response Computation

- Best response is: $h(\lambda \in T) = \arg\max_{a \in A} u_{gen}(\lambda, a, \tau_A)$

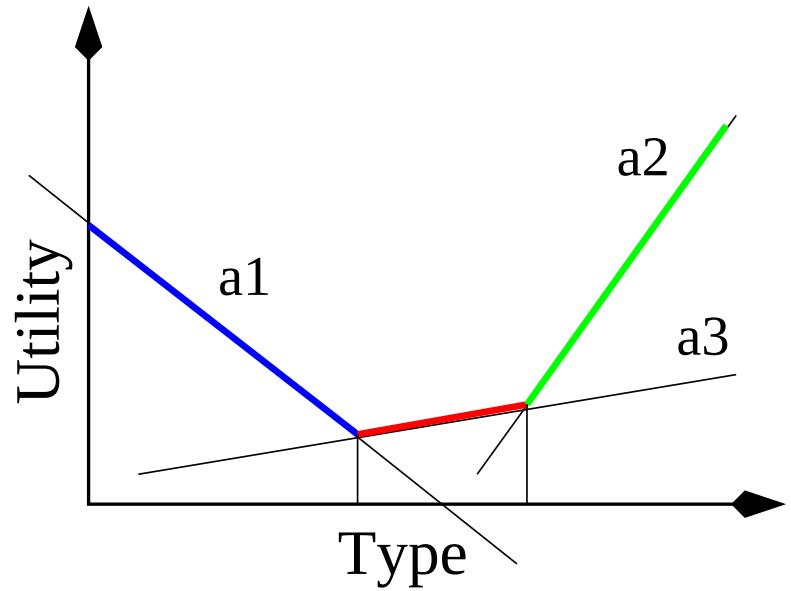

  - Set of distinct intervals I
  - For any $\alpha \in I$, if $\lambda_1, \lambda_2 \in \alpha \subset T$ then $h(\lambda_1) = h(\lambda_2)$
  - For any $\alpha_1 \neq \alpha_2 \in I$, if $\lambda_i \in \alpha_i$ then $h(\lambda_1) \neq h(\lambda_2)$

Best Response Computation

- Best response is: $h(\lambda \in T) = \arg\max_{a \in A} u_{gen}(\lambda, a, \tau_A)$
  - Set of distinct intervals I
  - For any $\alpha \in I$, if $\lambda_1, \lambda_2 \in \alpha \subset T$ then $h(\lambda_1) = h(\lambda_2)$
  - For any $\alpha_1 \neq \alpha_2 \in I$, if $\lambda_i \in \alpha_i$ then $h(\lambda_1) \neq h(\lambda_2)$
- Best response is piece-wise linear
  - Efficient computation of $\tau_{h,A}$

FP in CAPs: Results

- Empirically converges in all auction experiments
  - Equilibrium for simultaneous first price auctions with continuous private type
- Empirical convergence rate appears to be exponential
  - Surprise: in finite games FP converges slowly

Present and Future Work

- Efficient best response computation for manifold $\mathcal{U}$
  - Linearity of u_{gen} is type is not essential
- Counterexample exists for general convergence of FP
  - In auction based experiments it converges – why?
  - Necessary and sufficient conditions for FP to converge in CAPs

Questions