Best-Response Mechanisms

Authors: Noam Nisan *et al*
Innovations in Computer Science 2011

Sergio Zumpano Arnosti
MSc Student
MO829 – Algorithmic Game Theory

Summary

- Introduction & Motivation
- The setting
- Games with incentive-compatible best-response mechanisms
- Formalization & Modeling
- Examples
- Conclusions
- References

Introduction & Motivation

- Game theory and economics
  - Often abstract how equilibrium is reached.
  - “locally rational” actions -> mysteriously the system reach a global equilibrium.
  - e.g.: repeated best-response dynamics.

Introduction & Motivation

When is it best to best-respond?

Is it rational for players to repeatedly best-respond ?

Can a long-sighted player improve, in the long run, over this repeated myopic optimization?

Games with uncoupled incomplete information

+

Repeated best-respond

Introduction & Motivation

- Attractive trait

  To best-respond each player need only to know his own utility function, as his best response does not depend on other players' utility function, but only on their actions.

- Best-response dynamics -> natural protocol

  Gradual and limited sharing of information is an effort to reach an equilibrium.

  e.g., Internet routing

The setting

- Base game
  - n-player $(1, \dots, n)$ base game G
  - Player i

    Strategy space S_i and $S = S_1 \times \dots \times S_n$

    Utility function $u_i \; such \; that \; (u_1, \dots, u_n) \in U \subseteq U_1 \times .. \times U_n$

    Only knows his own utility function (private)

  - All players' utility functions -> full-information base game
  - Desire that the outcome be an equilibrium

The setting

- Best-response mechanisms
  - Players take turns selecting strategies
  - At each discrete time step t some player i_t selects and announces strategy $s_i^t \in S_{i_t}$
  - Choose a best-response to others announced strategies
  - Fully-specification
    - (1) the starting state
    - (2) order of player activations
    - (3) rule for breaking ties among multiple best responses

The setting

- Goal
  - Identify interesting classes of (base) games for which best-response mechanisms are incentive-compatible
  - When all other players are repeatedly best-reponding, then a player is incentivized to do the same.
  - Consider games in which repeated best-response dynamics do converge to an equilibrium.

The setting

- Tie-breaking rules
  - When exists multiple best-responses
  - Tie-breaking rule must be "uncoupled" depend only on the player private information
  - For each player i
    - Fix an a-priori full order $\prec_i \, on \, S_i$
    - Instruct him to break ties between multiple best-responses according to $\prec_i$

Games with incentive-compatible best-response mechanisms

	C	D
A	2, 1	0, 0
B	3, 0	1, 2

- * Unique PNE -> (B, D)

- Best-response dynamics are guaranteed to converge to * implies the incentive-compatibility of best-resp. mechanisms

Games with incentive-compatible best-response mechanisms

	C	D
A	2, 1	0, 0
B	3, 0	1, 2

- * Unique PNE -> (B, D)

- Best-response … o converge to * implies the ince … esp. mechanisms

FALSE

Games with incentive-compatible best-response mechanisms

	C	D
A	2, 1	0, 0
B	3, 0	1, 2

- * Unique PNE -> (B, D)

- Best-response … o converge to * implies the ince … esp. mechanisms

FALSE

Repeated best-responding is not incentive compatible in this game

Games with incentive-compatible best-response mechanisms

What traits must a game have for best-response dynamics to be incentive compatible?

Games with incentive-compatible best-response mechanisms

- *Never-Best-Response-Solvable (NBR-solvable) games with clear outcome*

  - Strategies are iteratively eliminated if a best-response never leads to them.

  - Clear outcome: each player i considers the game after the other player have already eliminated strategies that can be eliminated regardless of what i does. He will not be able to do better than the outcome.

Games with incentive-compatible best-response mechanisms

Theorem (informal)

Let G be an NBR-solvable game with clear outcome. Then, for every starting point and every (finite or infinite) order of player activations with at least $T = \sum_i |S_i| - n$ "rounds" it holds that:

1. Repeated best-response dynamics converges to a pure Nash equilibrium $s^* \, of \, G$
2. Repeated best-response dynamics is incentive compatible

Formalization and modeling

Definition 2.1 (tie-breaking rules/order)

Is a full order $\prec_i \ on\ S_i$
Multiple best-resposes: player i should choose the highest (under $\prec_i$) best-response.

Definition 2.2 (never-best-response strategies)

$s_i \ \in \ S_i$ is a NBR under tie-breaking order $\prec_i \ on\ S_i$ if for all s_{-i} there exists s_i' so that:

$$u_i(s_i, s_{-i}) < u_i(s_i', s_{-i})$$

OR both

$$u_i(s_i, s_{-i}) = u_i(s_i', s_{-i})\ and\ s_i \prec s_i'$$

Formalization and modeling

Definition 2.3 (NBR-solvable games)

A game G is never-best-response-solvable under tie-breaking rules $\prec_1, \dots, \prec_n$ if there exists a sequence of eliminations of NBR strategies that results in a single strategy profile.

Definition 2.4 (shortest-elimination parameters)

Let G be an NBR-solvable game

- Exists $G_0, \dots, G_r$
- $G = G_0$, in G_r each player has only a single strategy and $\forall$i $\in \{0, \dots, r-1\}, G_{i+1}$ *is obtained from* G_i via removal of sets of NBR strategies.
- e_G: length of shortest sequence of games for G.

Formalization and modeling

Definition 2.5 (globally-optimal profiles)

$s \in S$ is globally optimal for i if $\forall t \in S, u_i(t) < u_i(s)$.

Definition 2.6 (clear outcomes)

- Let G be an NBR-solvable game under $\prec_1, \dots, \prec_n$.
- Let s^* be the unique PNE under tie-breaking of G.
- G has a clear outcome if for every player i there exists an order of elimination of NBR strategies such that s^* is globally optimal for i at the first step in the elimination sequence.
- The game obtained after the removal of all previously-eliminated strategies from G.

Formalization and modeling

Theorem 2.7 (incentive-compatible mechanisms)

- Let G be an NBR-solvable game with a clear outcome $s^* \in S$ under tie-breaking rules $\prec_1, \dots, \prec_n$.
- Let M be a best-response mechanism for G with at least $T = e_G$ "$rounds$".

1. M converges to s^*
2. M is incentive compatible.

Formalization and modeling

Proof sketch (convergence):

- G: *NBR-solvable*
- $G_0, \ldots, G_r$, $G = G_0$ *and* $\forall i \in \{0, \ldots, r-1\}, G_{i+1}$ *is obtained from* G_i *via removal of NBR strategies*.
- *Consider the first round of a best-response mechanism.*
- *Consider* $j \in [n], \exists\, s_j \in S_j$ *that is NBR in* $G = G_0$.
- *Once* j *is activated,* s_j *will never be selected thereafter. After the first round, no NBR strategy in* G_0 *will be played ever again and hence the game is effectively equivalent to* G_1.
- *Same argument for the next rounds, mimic the elimination sequence in each strategy until reach* G_r*, whose unique strategy tuple* s^* *is the unique PNE under tie-breaking of* G.

Formalization and modeling

Proof sketch (incentive compatibility):

- *Let i be a player that deviates from repeated best-response and strictly gains from doing so.*
- *G is NBR-solvable → ∃ a (player-specific) order of elimination of NBR strategies such that s^* is globally optimal for i at the first step of elimination sequence (the game obtained after the removal of all previously-eliminated strategies from G).*
- $G_0, \ldots, G_l, G = G_0$ *and* $\forall i \in \{0, \ldots, l-1\}, G_{i+1}$ *is obtained from* G_i *via removal of NBR strategies (under tie-breaking).*
- *Let t_i be the index of the first game in sequence in which i's strategies are eliminated in that order.*

Formalization and modeling

Proof sketch (incentive compatibility):

- *All player but* i *are repeatedly best-responding and in the* $t_i - 1$ *first steps of the elimination sequence no strategy in* S_i *is removed.*
- *The same arguments for convergence can be used to show that after* $t_i - 1$ *rounds the game is effectively equivalent to* G_{t_i}*, regardless of the actions of player* i*.*
- *However, in that game,* i *can do no better than* s^**.*

Formalization and modeling

Proof sketch (incentive compatibility):

- *All player but* i *are repeatedly best-responding and in the* $t_i - 1$ *first steps of the elimination sequence no strategy in* S_i *is removed.*
- *The same arguments for convergence can be used to show that after* $t_i - 1$ *rounds the game is effectively equivalent to* G_{t_i}*, regardless of the actions of player* i*.*
- *However, in that game,* i *can do no better than* s^**.*

A CONTRADICTION

Examples

Game	Description
Stable-roommates	Students must be paired for the purpose of sharing dorm rooms. The objective is to find a “stable matching”.
Cost-sharing	Cost of some public service must be distributed between self-interested users.
Internet routing	BGP establishes routes between the smaller networks. Abstract and prove that BGP is incentive compatible in realistic environments.
Congestion control	TCP handles congestion on the Internet. Increase the transmission rate until congestion and then decrease. Show that such behavior is in equilibrium.

Examples - Stable-roommates

- n students $1, \dots, n$
- Each has a private strict ranking of the others and prefers being matched.
- A stable matching is not guaranteed to exist in general and, if a stable matching does exist, existing algorithms for reaching it are not incentive compatible.
- The authors observed environments where a stable matching is guaranteed to exist and can be reached in an incentive compatible manner.

Examples - Stable-roommates

Two well-known special cases:

Intern-hospital matchings

- The “students” are divided into two disjoint sets (interns and hospitals).
- Hospitals have the same ranking of interns.

Correlated markets

- The “students” are vertices in a complete graph.
- Every edge has a unique “weight”.
- The “heavier” the edge connecting students the higher that student ranks the other student.

Examples - Stable-roommates

Stable-roommates games

- Players: students
- S_i: $i's$ strategy space, the set of all students $j \neq i$
- $\alpha_i(j)$: $j's$ rank in student $i's$ ranking (lowest -> rank 1)
- $\forall s = (s_1, \dots, s_n) \in S$:

$$u_i(s) = \alpha_i(j) \Leftrightarrow s_i = j \; and \; \nexists k \neq i \; such \; that$$
$$s_k = j \; and \; \alpha_j(k) > \alpha_j(i)$$

- Otherwise:

$$u_i(s) = 0$$

Examples - Stable-roommates

Theorem (stable-roommates games)

For every stable-roommates game G it holds that in both hospital-intern matchings and correlated markets

- G is NBR-solvable
- $G's$ unique PNE is a stable matchings
- $e_G \leq n$

Examples - Stable-roommates

Proof sketch:

- Cycle-free: if there is no sequence of roommates $r_1, r_2, \dots, r_k$ of length $k > 2$ such that each student r_i ranks student r_{i+1} higher than student r_{i-1} (mod k).
- Cycle-free game has an elimination sequence:
  - Start with some arbitrary student r_1
  - Construct a sequence $r_1, r_2, \dots$ in which r_{i+1} is the sudent r_i prefers
  - Number of students is finite and so the sequence must repeat. Since the game is cycle-free, the cycle must be of length 2 $\longrightarrow$ located two students that desire each other the most.
- We can eliminate for each of the two the strategies of proposing to any other student $\longrightarrow$ Maximal utility by proposing to each other.

Examples - Stable-roommates

Proof sketch:

Ramains to notice that both environments are cycle-free:

(1) Hospitals and interns

- Any cycle of players will have to include a hospital after a desired intern and before a less desired one.

(2) Correlated markets

- Any cycle of nodes in the graph must include an edge with a lower weight that appears after an edge with a higher one.

(1), (2) the preferences do not induce a cycle in the matching graph

∴ The mechanism is a best-response mechanism for stable-roommates games
Theorem 2.7 ⟶ implements incentive-compatibility

Examples - Internet routing

- The network is an undirected graph $G = (V, E)$
- $V: n$ source nodes $1, \dots, n$ and a unique destination node d
- Each has a private strict ranking of all simple (loop-free) routes between itself and the dest. d.
- Under BGP, each source repeatedly examines its neighboring nodes' most recent announcements. Forwards through the neighbor whose route it likes the most, and announces its newly chosen route.
- BGP converges to a "stable" tree is the subject of networking research.

Examples - Internet routing

Theorem (Levin *et al* [2])

BGP is incentive-compatible in ex-post Nash in networks for which No Dispute Wheel holds.

"No Dispute Wheel" is the condition that no Dispute Wheel exist in the network.

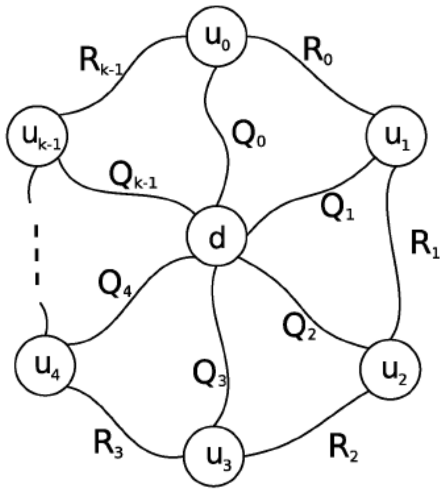

Each pivot node u_i would rather route clockwise through pivot node u_{i+1} than through the direct route Q_i.

Examples - Internet routing

BGP games

- Players: source nodes in V
- S_i: $i's$ outgoing edges in E
- $\vec{f} = (f_1, \dots, f_n)$: vector of source nodes' traffic forwarding decisions (strategies)
- $u_i(\vec{f} = (f_1, \dots, f_n))$: $i's$ rank for the simple route from i to d under $\vec{f}$ (lowest -> rank 1)

- Otherwise:

$$u_i\left(\vec{f}\right) = 0$$

Examples - Internet routing

Theorem (BGP games)

For a BGP game G it holds that:

- G is NBR-solvable
- $G's$ unique PNE is a stable routing tree
- $e_G \leq n$.

Examples - Internet routing

Proof sketch:

- Elimination order: locate a node that can guarantee its most preferred route (current subgame) and eliminate all other routing actions for it.
  - Start with an arbitrary node a_0 with at least 2 action.
  - Let R_0 be a_0's most preferred existing route to d.
  - Let a_1 be the vertex closest to d on R_0 , with two available actions in the current subgame, such that a_1 prefers some other route R_1 that leads a_1 to d.
  - Choose a_2 closest to d, R_2 *that is* a_2's most preferred.
  - Continue to choose $a_3, a_4, \ldots$ (finite number of vertices).

∴ We are able to find a node that can guarantee its most preferred route and continue with the elimination, until there are no more nodes with actions.

Examples - Congestion control

- Handled via combination of transmission-rate-adjustment protocols at the sender-receiver level (TCP) and queueing management policies.

- TCP is notoriously not incentive compatible.

- Godfrey *et al* [3] analyses incentives in TCP-inspired environments.

Examples - Congestion control

- The network is an undirected graph $G = (V, E)$.
- $c(e)$: capacity function that specifies the capacity for each edge $e \in E$.
- n source-target pairs of vertices (α_i, β_i) that aims to send traffic along a fixed route R_i in G.
- α_i can select transmission rates in the interval $[0, M_i]$.
- M_i is α_i's private information and wishes to maximize its achieved throughput.
- Congestion: sum of incoming flows exceeds edge's capacity, excess traffic must be discarded.

Examples - Congestion control

Two capacity-allocation schemes:

Strict-Priority-Queuing (SPQ)

- $\forall e \in E$ there is an edge-specific order over source nodes.
- Sharing: the most highly ranked source whose route traverses the edge gets its entire flow sent along the edge (up to $c(e)$); unused capacity is allocated to the second most highly ranked.

Weighted-Fair-Queueing (WFQ)

- $\forall e \in E$, each source node α_i has weight $w_i(e)$ at e.
- Allocated capacity: $\frac{w_i}{\Sigma_j w_j} c(e)$.
- Special case: $\forall e \in E, \forall i \in [n], w_i(e) = 1$ is called "fair queuing" (FQ).

Examples - Congestion control

- Godfrey *et al* [3] considers a TCP-like protocol called Probing-Increase-Educate-Decrease (PIED). PIED is shown to be incentive compatible in SPQ and WFQ.

Theorem 3.7 (Godfrey *et al* [3])

- PIED is incentive compatible in networks in which all edges use SPQ with coordinated priorities.

Theorem 3.8 (Godfrey *et al* [3])

- PIED is incentive compatible in networks in which all edges use WFQ with coordinated weights (and so if all edges use FQ then PIED is incentive compatible)

Examples - Congestion control

TCP games

- Players: source nodes
- $S_i = [0, M_i]$: $i's$ strategy space
- $\vec{r} = (r_1, \dots, r_n)$: vector of source nodes' transmission rates (strategies)
- $u_i(\vec{r})$: is α_i's achieved throughput in the unique traffic-flow equilibrium point of the network for $\vec{r}$.
- Godfrey *et al* [3] shows that such a unique point exists for SPQ and WFQ.
- Tie-breaking rules:

$$\forall s, t \in S_i, \; s \prec_i t \iff s > t$$

Examples - Congestion control

Theorem (TCP games)

For every TCP game G such that all edges use SPQ with coortinated priorities, or all edges use WFQ with coordinated weights, it holds that:

- G is NBR-solvable under tie-breaking rules.
- G's unique PNE under these tie-breaking is a stable flow pattern.
- $e_G \leq n$.

The proof will be only for the case of Weighted-Fair-Queuing, with equal weights.

Examples - Congestion control

Proof sketch:

- For each edge e, the share of each flow as $\beta_e = c_e / k_e$.
- Construct an elimination sequence:
  - Let e^* be the edge with the minimal β.
  - Each flow on this edge is guaranteed β_{e^*} traffic and at least that amount on all other edges.
  - Therefore is possible to eliminate all actions of transmitting less than β_{e^*}.
  - If player i eliminates actions below β_{e^*} last among players that go through e^*, then he does so in game in which the final profile is optimal for him.

∴ The game has a clear outcome. Theorem 2.7 implies a result that is similar in spirit to the two theorems of Godfrey *et al* [3].

Conclusions

- It was possible to explore when such locally-rational dynamics are also globally rational.
- Results along the article give an incentive to think in new structures of existing protocols/mechanisms and provide new insights into the design of them.
- It was interesting to see that, in some specific conditions, real environments with repeated best-response mechanism can be incentive compatible.

References

[1] N. Nisan, M. Schapira, G. Valiant, A. Zohar "Best-Response Mechanisms", Innovations in Computer Science - ICS 2011, pages 155-165, 2011.

[2] H. Levin, M. Schapira, and A. Zohar. Interdomain routing and games. In STOC, pages 57-66, 2008.

[3] P. B. Godfrey, M. Schapira, A. Zohar, and S. Shenker. Incentive compatibility and dynamics of congestion control. SIGMETRICS Perform. Eval. Rev., 38(1):95-106, 2010.