Obvious Dominance and Random Priority

Marek Pycia[1] Peter Troyan[2]

[1]University of Zurich and UCLA

[2]University of Virginia

National University of Singapore
June 8th, 2018

Motivation

Social choice without transfers:

- School choice (Abdulkadiroğlu & Sönmez 2003)
- Public housing (Kaplan 1987, Bloch & Cantala 2014)
- Course allocation (Budish & Cantillon 2012)
- Military branching (Sönmez & Switzer 2013)
- Organ transplants (Roth, Sönmez, & Ünver 2014)
- Food banks (Prendergast 2017)
- Refugee resettlement (Roth 2015, Jones & Teytelboym 2016)
- Voting (Arrow 1963)
- Myriad informal settings (professors to offices, children to chores)

Motivation

- Design objectives: efficiency, fairness, **good incentives**
- "Standard" approach: apply revelation principle, use strategy-proof (SP) direct revelation mechanism
- Depends on agents *understanding* that a mechanism is SP
  - Sealed-bid vs. ascending auctions (Kagel et al. 1987, Ausubel 2004)
  - Mistakes under deferred acceptance/serial dictatorships (Chen and Sönmez 2003, Hassidim et al. 2016, Chen and Pereyra 2016, Rees-Jones 2017)
- Additional design goal: "simplicity"
- Want a definition of simplicity that:
  - gives a formal game-theoretic benchmark
  - is analytically tractable
  - is useful

Contributions

1. Introduce a new concept of simplicity for mechanisms: strong obvious strategy-proofness (SOSP)
    - Refinement of obvious strategy-proofness (OSP; Li, 2017)
2. Fully characterize the class of simple mechanisms for social choice problems without transfers
    - OSP mechanisms = "millipede games"; may be complex, require extensive foresight
    - SOSP mechanisms: need to look at most one step ahead
3. Show that there is a **unique** mechanism that is efficient, fair, and simple (SOSP): Random Priority (RP)
    - RP is widespread in practical applications
    - First general characterization that explains its popularity

Related Literature

- **Obvious strategy-proofness**: Li (2017), Ashlagi and Gonczarowski (2015), Troyan (2016), Bade and Gonczarowski (2016), Zhang and Levin (2016), Milgrom and Segal (2016) Arribillaga, Masso, and Neme (2017), Mackenzie (2017), Levin and McGee (2017)

- **Large literature on social choice without transfers**: Gibbard (1973, 1977), Sattherthwaite (1975), Shapley and Scarf (1974), Hylland and Zeckhauser (1979), Abdulkadiroglu and Sönmez (1998), Papai (2000), Pycia and Ünver (2016)

- **Random Priority as the unique incentive compatible, efficient, and fair mechanism...**
  ...in small markets ($N = 3$): Bogolmolnaia and Moulin (2001)
  ...asymptotically ($N \to \infty$): Liu and Pycia (2011)

Model

Paper: General social choice model without transfers

- $\mathcal{X} = \{x, y, z, \ldots\}$ – finite set of outcomes
- $\mathcal{N} = \{i, j, k, \ldots\}$ – finite set of agents
- $\succsim_i$ – agent i's preference ranking over outcomes
- Key assumption: "rich" preference domain

Talk: Object allocation with single-unit demand

- $\mathcal{N}$=agents, $\mathcal{O}$=objects; $|\mathcal{N}| = |\mathcal{O}|$
- $\succ_i$ – i's (strict) preference ranking over $\mathcal{O}$
- $\mathcal{X} =$ all possible allocations (care only about own assignment)
- Richness: every strict ranking of objects is possible

Notation

Γ: finite extensive game form (or "mechanism")

h: generic history in Γ

$A(h)$: set of actions available at h

S_i: strategy for agent i (complete contingent plan of action)

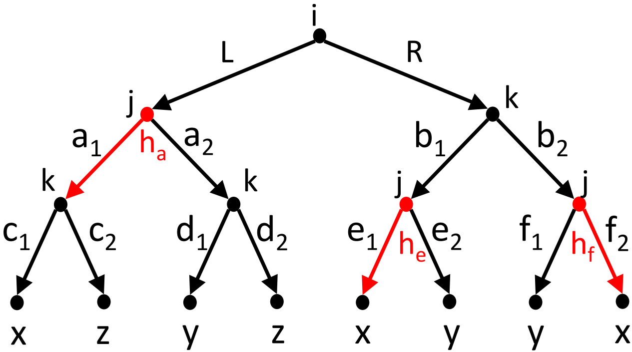

Weak vs. Obvious Dominance

$u_i(S_i, S_{-i}, \succ_i)$: type $\succ_i$'s utility when play follows (S_i, S_{-i}) in Γ

Weak vs. Obvious Dominance

$u_i(S_i, S_{-i}, \succ_i)$: type $\succ_i$'s utility when play follows (S_i, S_{-i}) in Γ

- S_i **weakly dominates** S'_i (for type $\succ_i$) if:

$$u_i(S_i, S_{-i}; \succ_i) \geq u_i(S'_i, S_{-i}; \succ_i) \qquad \text{for all } S_{-i}$$

Weak vs. Obvious Dominance

$u_i(S_i, S_{-i}, \succ_i)$: type $\succ_i$'s utility when play follows (S_i, S_{-i}) in Γ

- S_i **weakly dominates** S'_i (for type $\succ_i$) if:

$$u_i(S_i, S_{-i}; \succ_i) \geq u_i(S'_i, S_{-i}; \succ_i) \qquad \text{for all } S_{-i}$$

- S_i **obviously dominates** S'_i (for type $\succ_i$) if, starting from any earliest h at which S_i and S'_i differ:

$$\min_{S_{-i}} u_i(S_i, S_{-i}; \succ_i) \geq \max_{S_{-i}} u_i(S'_i, S_{-i}; \succ_i)$$

Weak vs. Obvious Dominance

$u_i(S_i, S_{-i}, \succ_i)$: type $\succ_i$'s utility when play follows (S_i, S_{-i}) in Γ

- S_i **weakly dominates** S_i' (for type $\succ_i$) if:

$$u_i(S_i, S_{-i}; \succ_i) \geq u_i(S_i', S_{-i}; \succ_i) \qquad \text{for all } S_{-i}$$

- S_i **obviously dominates** S_i' (for type $\succ_i$) if, starting from any earliest h at which S_i and S_i' differ:

$$\min_{S_{-i}} u_i(S_i, S_{-i}; \succ_i) \geq \max_{S_{-i}} u_i(S_i', S_{-i}; \succ_i)$$

- A mechanism Γ is **(obviously) strategy-proof** if every type $\succ_i$ has an (obviously) dominant strategy

Weak vs. Obvious Dominance

Say $x_1 \succ_i x_2 \succ_i x_3 \succ_i x_4$

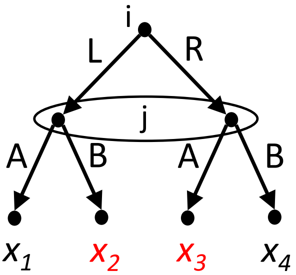

- Here, L is obviously dominant (and also weakly dominant)
  - Worst case from L: x_2
  - Best case from R: x_3

Weak vs. Obvious Dominance

Say $x_1 \succ_i x_2 \succ_i x_3 \succ_i x_4$

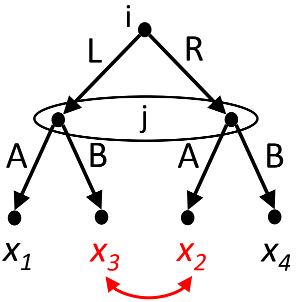

- L is **not** obviously dominant, but **is** weakly dominant
  - Requires contingent reasoning ("If j plays A, I should play L b/c $x_1 \succ_i x_2$; if j plays B, I should play L b/c $x_3 \succ_i x_4$")
  - Often difficult for people, even in single-agent decision problems (Charness and Levin 2009; Esponda and Vespa 2014)

Classifying Mechanisms

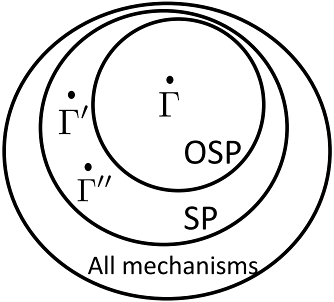

In our no-transfer setting, e.g.:

Γ = extensive-form serial dictatorship

Γ' = normal-form serial dictatorship

Γ'' = top trading cycles, deferred acceptance, ...

- According to revelation principle: $\Gamma = \Gamma'$.
- Real-world agents much more likely to play dominant strategies in Γ (Li 2017, Chen and Pereyra 2016)

Classifying Mechanisms

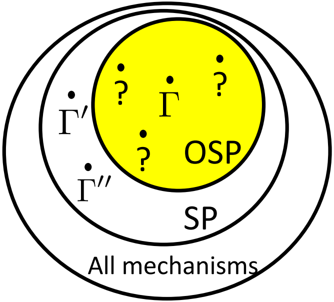

In our no-transfer setting, e.g.:

Γ = extensive-form serial dictatorship

Γ' = normal-form serial dictatorship

Γ'' = top trading cycles, deferred acceptance, ...

Question: What else is in the shaded area? Of interest because:

1. Might discover new simple mechanisms
2. "Stress test" of OSP - does it conform with our intuitive understanding of simplicity?

Millipede Games

Q1: Can we find **anything** else that is OSP?

Millipede Games

Q1: Can we find **anything** else that is OSP?

A: Yes. Consider the following game (x, y, z are objects):

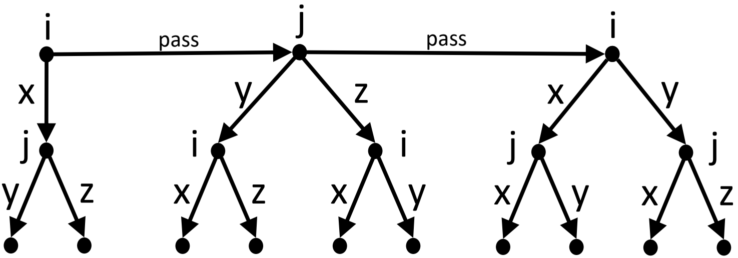

Two types of actions: **clinching** and **passing**

Clinch-or-pass structure analogous to famous centipede game, but more "legs" $\longrightarrow$ **millipede game**

Characterizing OSP

Q2 (harder): Can we find **everything** else that is OSP?

Characterizing OSP

Q2 (harder): Can we find **everything** else that is OSP?

Theorem

A mechanism Γ is obviously strategy-proof if and only if it is equivalent to a millipede game.

More formal definitions will follow.

Detour: Equivalence

Two mechanisms Γ and $\hat{\Gamma}$ are **equivalent** if there exist (obviously dominant) strategies $(S_{\succ_i})_{i\in\mathcal{N}}$ and $(\hat{S}_{\succ_i})_{i\in\mathcal{N}}$ such that for all type profiles $(\succ_i)_{i\in\mathcal{N}}$, the outcome when agents play $(S_{\succ_i})_{i\in\mathcal{N}}$ in Γ is the same as when agents play $(\hat{S}_{\succ_i})_{i\in\mathcal{N}}$ in $\hat{\Gamma}$

A warm-up result:

Lemma (Ashlagi and Gonczarowski, 2016)

Every obviously strategy-proof game of imperfect-information is equivalent to an obviously strategy-proof game of perfect information.

Clinching vs. Passing

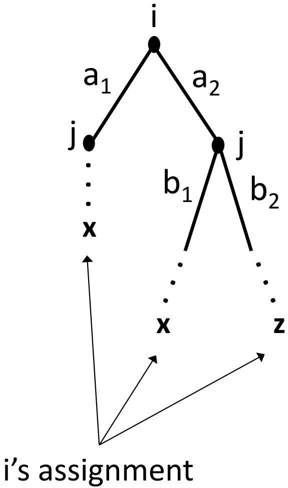

- a_1 = clinching action
- a_2 = passing action (i's outcome depends on j's choice)

Clinching vs. Passing

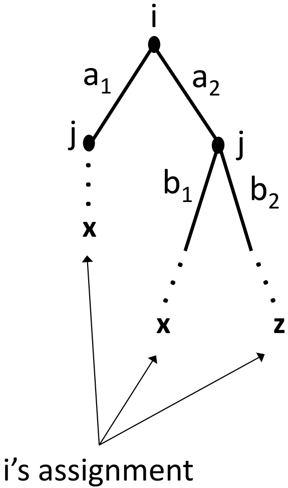

- a_1 = clinching action
- a_2 = passing action (i's outcome depends on j's choice)
- Both types ($x \succ_i z$ and $z \succ_i x$) have a obviously dominant strategy:
  - $x \succ_i z$: a_1 obviously dominant
  - $z \succ_i x$: a_2 obviously dominant

Multiple Passing Actions?

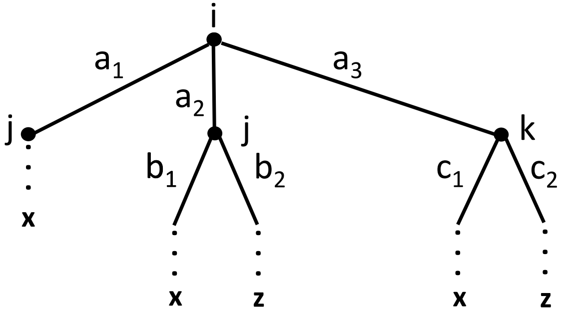

- a_2 does not obviously dominate a_3 for type $z \succ_i x$ (and vice-versa)
- Implication: OSP games can have **at most** one passing action at each history

When Is It Obviously Dominant to Pass?

Example: Say $y \succ_i x \succ_i z \succ_i w$ for i.

- Clinching x is not obviously dominant
- Worst case from passing depends on j's choice

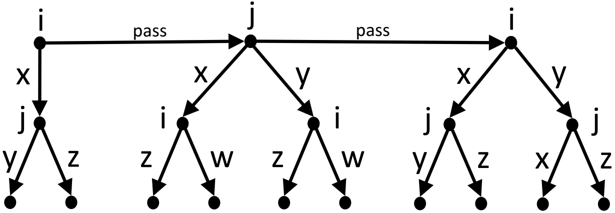

When Is It Obviously Dominant to Pass?

Example: Say $y \succ_i x \succ_i z \succ_i w$ for i.

- Clinching x is not obviously dominant
- Worst case from passing depends on $j's$ choice

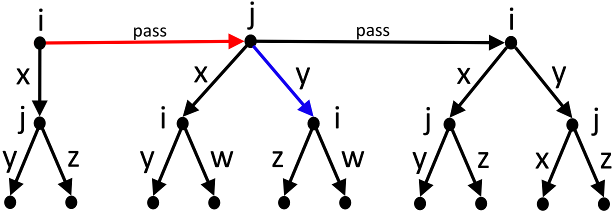

When Is It Obviously Dominant to Pass?

Example: Say $y \succ_i x \succ_i z \succ_i w$ for i.

- Clinching x is not obviously dominant
- Worst case from passing depends on $j's$ choice

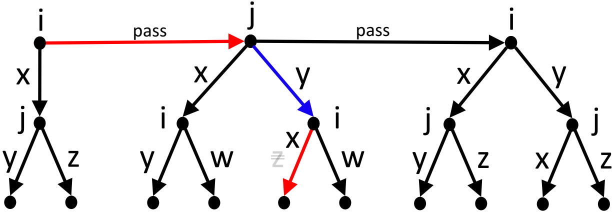

More generally, following a pass, must guarantee:

- If something that was possible, but not clinchable, disappears, must offer everything that was clinchable
- If something that was clinchable disappears, must offer everything that was previously possible, but not clinchable

Millipede Games: Formal Definition

Definition

A **millipede game** is a finite game of perfect information such that, at any history h

(a) At most one action in $A(h)$ is a passing action.

(b) For all x, one of the following holds:

(i) $x \in P_i(h)$

(ii) $x \notin P_i(\tilde{h})$ for some $\tilde{h} \in \mathcal{H}_i(h)$

(iii) $x \in \cup_{\tilde{h} \in \mathcal{H}_i(h)} C_i(\tilde{h})$

(iv) $\cup_{\tilde{h} \in \mathcal{H}_i(h)} C_i(\tilde{h}) \subseteq C_i(h)$

where

- $P_i(h)$: outcomes that are possible for i at h
- $C_i(h)$: outcomes that i can clinch at h
- $\mathcal{H}_i(h)$: histories prior to h where i moves.

Characterizing Obvious Strategy-Proofness

Theorem

Every millipede game is obviously strategy-proof. If a mechanism Γ is obviously strategy-proof, then it is equivalent to a millipede game.

- A "revelation principle" for obvious dominance

How “Obvious” is Obvious Dominance?

Examples of games with obviously dominant strategies:

- (Extensive-form) serial dictatorships, Random Priority
- The game below (e.g., obviously dominant for type $o_1 \succ_i o_2 \cdots \succ_i o_{100}$ to always pass):

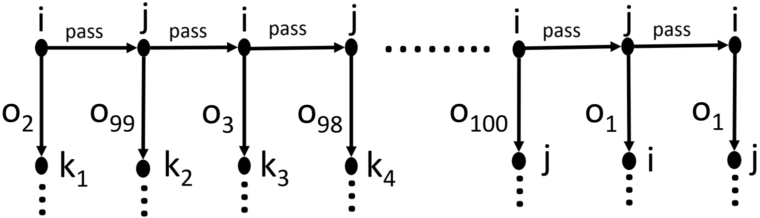

- Chess

How “Obvious” is Obvious Dominance?

- Recall the definition of obvious dominance:

$$\min_{S_{-i}} u_i(S_i, S_{-i}; \succ_i) \geq \max_{S_{-i}} u_i(S_i', S_{-i}; \succ_i)$$

- Min/max taken over S_{-i}, *taking S_i as given*
- Allows for agents who cannot contingently reason about actions of others
- Presumes they can with regard to their own future actions
- May still require significant foresight

Strong Obvious Dominance

- Natural refinement: take max/min over all possible future strategies of ALL agents
  - Includes all $j \neq i$ and i's "future self"
- We call this **strong obvious dominance**
- If every type $\succ_i$ of every agent has a strongly obviously dominant strategy, then Γ is **strongly obviously strategy-proof (SOSP)**
- Strongly obviously dominant strategies are those that can be recognized as weakly dominant by a cognitively limited agent who cannot distinguish between outcome-set equivalent games

Characterizing SOSP Mechanisms

- History h is **payoff-irrelevant** for i if either (i) $|A(h)| = 1$ or (ii) $|P_i(h)| = 1$. Otherwise, h is **payoff-relevant**

Theorem

Along any path of a SOSP mechanism, there is at most one payoff-relevant history for each agent.

- The unique payoff-relevant history (if it exists) is first time an agent is called to choose from among more than two actions
- May be called later, but choice cannot affect his own payoff
  - Might affect payoffs of others
- Eliminates the more complex examples of millipede games

Characterizing SOSP Mechanisms

A **curated dictatorship** is a perfect-information game in which:

- Agents are called to play sequentially, with each agent called at most once
- The next agent to move and the set of objects offered to her are determined by the actions taken by prior agents
- If 3 or more objects are possible for agent i:
  - She is offered the opportunity to clinch any possible object
- If only 2 objects (say $\{x, y\}$) are possible for agent i, either:
  - She is offered the opportunity to clinch either x or y
  - She is offered the opportunity clinch one (say x) or pass. If she passes, outcome (x or y) determined by future agents

Characterizing SOSP Mechanisms

Theorem

A mechanism Γ is strongly obviously strategy-proof if and only if it is equivalent to a curated dictatorship.

Proof sketch:

- Take a SOSP game Γ.
- At most one payoff-relevant history for each i, denoted h_0^i
- If i is called to move again, construct equivalent Γ' where i is asked to make all future choices at h_0^i
- Since future choices payoff-irrelevant, SOSP preserved

Other desiderata

- Thus far, have focused on incentives
- Two other important desiderata when designing mechanisms: efficiency and fairness
- A mechanism Γ and strategy profile $(S_{\succ_i})_{i \in \mathcal{N}}$ are **efficient** if the final outcome is Pareto efficient for every type profile
- For fairness, we use **equal treatment of equals (ETE)**: if $\succ_i = \succ_j$, then i and j receive the same (distribution over) outcomes
- Curated dictatorships may violate both of these properties

SOSP and Efficient Mechanisms

- Curated dictatorships allow agents to be “exogenously” denied some objects: may be inefficient

SOSP and Efficient Mechanisms

- Curated dictatorships allow agents to be “exogenously” denied some objects: may be inefficient
- An **almost-sequential dictatorship** is a curated dictatorship in which, at each history, every still-available object is possible for the agent called to act

SOSP and Efficient Mechanisms

- Curated dictatorships allow agents to be "exogenously" denied some objects: may be inefficient
- An **almost-sequential dictatorship** is a curated dictatorship in which, at each history, every still-available object is possible for the agent called to act

Theorem

A mechanism Γ is strongly obviously strategy-proof and efficient if and only if it is equivalent to an almost sequential dictatorship.

Fairness

- Serial dictatorships are a special case of almost-sequential dictatorships in which the agent ordering is fixed in advance
- In a serial dictatorship, the first mover always gets her favorite; last mover gets whatever is left
- Violates equal treatment of equals
- Standard solution: Randomization

Random Priority

Random Priority (RP)

1. Nature selects an ordering of the agents uniformly at random from all possible permutations of $\mathcal{N}$
2. Each agent moves once, in this order. At i's turn, she is offered all still-available objects, and selects one

Equivalent formulation:

1. Nature randomly selects an agent from $\mathcal{N}$, say i_1.
2. i_1 chooses an object from $\mathcal{O}$, say o_1
3. Nature randomly selects another agent from $\mathcal{N} - \{i_1\}$, say i_2
4. i_2 chooses an object from $\mathcal{O} - \{o_1\}$, say o_2
5. ...

Characterizing Random Priority

- RP is well-known to be efficient and fair (ETE)
- Easy to show that RP is SOSP
- In fact, it is the **unique** mechanism that satisfies these properties:

Theorem

A mechanism Γ is strongly obviously strategy-proof, efficient, and satisfies equal treatment of equals if and only if it is equivalent to Random Priority.

Proof sketch

- SOSP + efficiency implies Γ equivalent to almost-sequential dictatorship (earlier result)

Proof sketch

- SOSP + efficiency implies Γ equivalent to almost-sequential dictatorship (earlier result)
- First: all i have equal chance of being first mover
  - If not, consider $o_1 \succ_i o_2 \succ_i \cdots$ for all i. Some j has higher prob. to receive $o_1 \rightarrow$ violates ETE

Proof sketch

- SOSP + efficiency implies Γ equivalent to almost-sequential dictatorship (earlier result)
- First: all i have equal chance of being first mover
  - If not, consider $o_1 \succ_i o_2 \succ_i \cdots$ for all i. Some j has higher prob. to receive $o_1 \rightarrow$ violates ETE
- Conditional on chosen i_1, all $\mathcal{N} - \{i_1\}$ have equal chance of being second mover

Proof sketch

- SOSP + efficiency implies Γ equivalent to almost-sequential dictatorship (earlier result)
- First: all i have equal chance of being first mover
  - If not, consider $o_1 \succ_i o_2 \succ_i \cdots$ for all i. Some j has higher prob. to receive $o_1 \rightarrow$ violates ETE
- Conditional on chosen i_1, all $\mathcal{N} - \{i_1\}$ have equal chance of being second mover
  - Consider $\succ_{i_1} = o_1 \succ_{i_1} x \cdots$ and $\succ_j = x \succ \cdots$ for all $j \neq i_1$

$$\underbrace{\Pr(j \text{ gets } x)}_{(A)} = \underbrace{\Pr(j \text{ moves 1st})}_{(B)} + \underbrace{\Pr(j \text{ moves 2nd} | i_1 \text{ moves 1st})}_{(C)}$$

  - (A) & (B) are the same for all $j \neq i_1$ (by ETE and step 1)
  - Therefore, (C) is the same for all $j \neq i_1$
- ...

Transfers

- With transfers (e.g., combinatorial auctions), SOSP mechanisms are **(personalized) posted price mechanisms**
- Einav et al. (2018) document a dramatic shift from auctions to posted prices on eBay
- In computer science: computational complexity of combinatorial auctions has led to interest in posted price mechanisms (e.g., Chawla et al. 2010, and Feldman et al. 2014)
- Our work provides an additional reason for the ubiquity of posted prices: besides being computationally simple, they are also strategically simple (in the sense of SOSP)

Conclusion

- In environments without transfers, even OSP mechanisms may not necessarily be “simple”
- Strong OSP: at most one payoff-relevant choice for each agent
  - An explanation for why some mechanisms used more than others, despite being equivalent according to revelation principle and OSP
  - Random Priority is the unique SOSP, efficient, and fair mechanism
- More generally, taking “simplicity” seriously as constraint is a new and interesting research agenda
  - Rapid expansion of OSP literature following Li (2017)
  - Borgers and Li (2017): introduce alternative notion of simplicity that is weaker than SP
  - ‘Strength’ needed depends on application: good to have a variety of definitions

Behavioral Agents and Strong Obvious Dominance

- Say Γ and Γ' are **outcome-set equivalent** if there exists a bijection of histories $\phi : \mathcal{H} \to \mathcal{H}'$ such that

$$P_i(h) = P'_i(\phi(h)) \text{ for all } i, h$$

Theorem

For all i, $\succ_i$, strategy S_i is strongly obviously dominant in Γ if and only if the corresponding strategy S'_i is weakly dominant in any outcome-set equivalent game Γ'.

- Consider an agent i who knows all possible outcomes, but cannot contingently reason about how they depend on her opponents or her own future actions
- Even such cognitively-limited agents will be able to recognize strongly obviously dominant strategies

Back