Implementation in Undominated Strategies via Bounded Mechanisms

"ReLax" Workshop on Games, Chennai Math Institute

Arunava Sen, Indian Statistical Institute

February 1, 2021

Classic Mechanism Design Problem

- Players/agents/voters have private information.
- They collectively want to achieve some goals.
- Equivalent reformulation: A *mechanism designer* wishes to achieve these goals.
- The goals depend on the private information of agents.
- How does the designer achieve her goals? Is it possible to achieve them?
- "Reverse Engineering" problem (Aumann).

Classic Mechanism Design Problem

- The mechanism designer designs a "mechanism" which consists of (i) a choice of message to be sent be each player and (ii) an outcome function that gives an outcome for all possible message-tuples.
- Once information is realized, players are playing a game.
- Fix an appropriate solution concept for a game.
- The mechanism is successful (given the solution concept), if equilibrium outcomes coincide with the goals of the designer (for all possible information realizations).
- Question: What goals can be achieved under these circumstances?
- Applications: Voting, Auctions etc.

Goal of the Talk

- The standard equilibrium notion for mechanism design is dominant strategies or strategy-proofness.
- Strategy-proofness is a demanding - many negative results (Gibbard-Satterthwaite Theorem in voting).
- We consider an alternative solution concept undominated strategies that has the same informational basis.
- Relatively less studied because of difficult technical issues.
- Introduce the problem and report some recent progress.
- Main message: Designer can do “better” using undominated strategies.

The Model

- Finite set of agents $N = \{1, \ldots n\}$ and set of (fixed) alternatives A.
- Each agent $i \in N$'s private information or *type* is θ_i where $\theta_i \in \Theta$.
- Neither designer nor players other than i can observe θ_i.

The Model

- In an "ordinal" model, θ_i determines an ordering over A, $R(\theta_i)$. Interpretation: $aR(\theta_i)b$ implies a is at least as good as b. Strict component $P(\theta_i)$: $aP(\theta_i)b$ implies $aR(\theta_i)b$ and not $bR(\theta_i)$.
- In a "cardinal" model, θ_i specifies a *valuation* or *utility* function $v : A \times \Theta_i \to \Re$.
- A *type profile* is an n-tuple $\theta \equiv (\theta_1, \ldots, \theta_n)$.

Social Choice Correspondences and Social Choice Functions

A *social choice correspondence SCC F* associates a non-empty subset of alternatives $F(\theta) \subset A$ for every type profile $\theta \in \Theta^n$..

A *social choice function* is a singleton- valued SCC.

Mechanisms and Implementation

A mechanism $G = (M, g)$ consists of a message space $M \equiv \times_{i \in N} M_i$ and an outcome function $g : M \to A$.

- Each player i chooses a message $m_i \in M_i$. The resulting outcome is $g(m_1, \ldots, m_n)$.
- For every $\theta \in \Theta^n$, the pair (G, θ) constitutes a game.
- Let E be a solution concept, i.e. $E(G, \theta) \subset M$ for all $\theta \in \Theta^n$.

The mechanism G implements F if $g(E(G, \theta)) = F(\theta)$ for all $\theta \in \Theta^n$.

Solution concepts: Dominant Strategies

- The most widely used solution concept is dominant strategies.
- $[\bar{m} \in E(G, \theta)] \Rightarrow [g(\bar{m}_i, m_{-i})R(\theta_i)g(m_i, m_{-i})]$ for all $m_i \in M_i, m_{-i} \in M_{-i}$, $i \in N$.
- $\bar{m}_i$ is optimal for i at θ_i irrespective of the messages sent by other players.

Solution concepts: Bayes-Nash Equilibrium

- Another popular solution concept is Bayes-Nash equilibrium.
- Appropriate for the ordinal model but can be adapted suitably to the ordinal model.
- Player have a "belief" μ over types, i.e. $\mu(\theta) \geq 0$ and $\sum_\theta \mu(\theta) = 1$.
- $[\bar{m} \in E(G, \theta)] \Rightarrow [\int_{\theta_{-i}} v(g(\bar{m}_i, \bar{m}_{-i}), \theta_i) d(\mu(\theta_{-i}|\theta_i)) \geq \int_{\theta_{-i}} v(g(m_i, \bar{m}_{-i}), \theta_i) d(\mu(\theta_{-i}|\theta_i)]$ for all $m_i \in M_i$ and $i \in N$.
- $\bar{m}_i$ is optimal for i at θ_i in expectation, assuming other players play according to $\bar{m}$.

Solution Concepts: Undominated Strategies

- We focus on the solution concept undominated strategies.
- $[\bar{m} \in E(G,\theta)] \Rightarrow [\nexists i \in N, \nexists m_i' \in M_i$ such that $g(m_i', m_{-i})R(\theta_i)g(\bar{m}_i, m_{-i})$ for all $m_{-i} \in M_{-i}$ and $g(m_i', \hat{m}_{-i})P(\theta_i)g(\bar{m}_i, \hat{m}_{-i})$ for some $\hat{m}_{-i} \in M_{-i}]$.
- $\bar{m}_i$ is undominated for player i of type θ_i, i.e. there does not exist another message for i that does at least as well as $\bar{m}_i$ for any possible message of the other players and does strictly better for some message.
- Set-valued notion unlike dominant strategies.

Solution concepts: Remarks

- In the solution concepts, equilibrium message for i at profile θ depends *only* on her type.
- In dominant strategies and undominated strategies, the following is true: For every $\theta \in \Theta^n$, $E(G,\theta) = (E_1(G,\theta_1) \times \ldots \times E_n(G,\theta_n))$ where $E_i(G,\theta_i)$ is the set of equilibrium messages for voter i of type θ_i.
- Consistent with the private information environment.
- Implementation in Bayes-Nash equilibrium depends on beliefs - on the other hand, both dominant strategies and undominated are "detail-free" in that respect.
- In dominant strategy implementation, there is *a* message $\bar{m}_i(\theta_i)$ that weakly dominates every other message $m_i' \in M_i$.
- Not true in undominated implementation.
- There can be *different* messages that do not dominate each other.

The Revelation Principle

- Let f be a SCF.
- A direct mechanism is the mechanism $D_f \equiv (\Theta, f)$ i.e. players announce their types and f is the outcome function.
- The SCF f can be truthfully implemented according to E if $\theta_i \in E_i(D_f, \theta_i)$, for all $\theta_i \in \Theta_i$ and i, i.e. truth-telling is an equilibrium for every player of every type. If dominant strategy is the solution concept, then f is strategy-proof.
- Note that we are not insisting $\theta_i = E_i(D_f, \theta_i)$., i.e. there could be non truth-telling equilibrium type announcements as well.

The Revelation Principle

Theorem

(Revelation Principle): Suppose F can be implemented in dominant strategies (or Bayes-Nash equilibrium). Then any SCF $f \in F$ can be truthfully implemented.

Proof.

Suppose $G = (M, f)$ implements F. Let $f \in F$. For every $\theta \in \Theta$, there exists $\bar{m}_i(\theta_i) \in M_i$ such that $\bar{m}_i(\theta_i) \in E_i(G, \theta_i)$ for all $i \in N$ and $g(\bar{m}) = f(\theta)$. Construct direct mechanism D_f. Note θ_i in D_f is a message that corresponds to $m_i(\theta_i)$ in M_i. Messages that are not equilibrium messages for some θ_i are eliminated in D_f. If E is either dominant strategies or BNE, then θ_i will be an equilibrium in D_f.

□

The Revelation Principle

- Without loss of generality $\Theta_i \subset M_i$ where $G = (M, g)$ is the implementing mechanism.
- Critical feature of dominant strategies and BNE that make the Revelation Principle work: if m_i is an equilibrium for type θ_i, it continues to remain an equilibrium when "redundant" strategies of other players are eliminated.
- Revelation Principle enormously simplifies mechanism design for solution concepts that satisfy it - the direct mechanism is a "canonical mechanism".
- $\theta_i \in E(D_f, \theta_i)$ is also called the incentive-compatibility requirement. Necessary for implementing f.

The Revelation Principle

- Implementation via undominated strategies *does not* satisfy the "stability with respect to the deletion of redundant messages" condition satisfied by implementation by dominant strategies or BNE.
- Consider messages $\bar{m}_i, m'_i \in M_i$ such that $g(m'_i, m_{-i})P_i(\theta_i)g(\bar{m}_i, m_{-i})$ for all $m_{-i} \in M_{-i} \setminus \hat{m}_{-i}$ and $g(\bar{m}_i, \hat{m}_{-i})P_i(\theta_i)g(m'_i, \hat{m}_{-i})$.
- Both $\bar{m}_i$ and m'_i are undominated at θ_i. However, $\bar{m}_i$ is dominated by m'_i if $\hat{m}_{-i}$ is deleted.
- The Revelation does not hold for implementation in undominated strategies as the next Example shows.

Example (Jackson 1992)

- $N = \{1, 2\}$.
- $A = \{a, b\}$.
- $\Theta_1 = \{\theta_1, \theta_1'\}$, $\Theta_2 = \{\theta_2\}$.
- $aP(\theta_1)b$, $bP(\theta_1')a$ and $aP(\theta_2)b$.
- $f(\theta_1, \theta_2) = b$, $f(\theta_1', \theta_2) = a$.
- Player 1's worse alternative is selected in both states.
- In the direct mechanism, lying is a dominant strategy for player 1 in both states.
- However, f can be implemented in undominated strategies by the following crazy mechanism!

Example (Jackson, 1992)

		M_2									
		m_2									
	m_1	b	a	a	a	$\dots$	a	a	a	a	$\dots$
		b	a	a	a	$\dots$	b	b	b	b	$\dots$
		b	b	a	a	$\dots$	b	b	b	b	$\dots$
		b	b	b	a	$\dots$	b	b	b	b	$\dots$
M_1		$\vdots$	$\vdots$	$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$	$\vdots$	
	$\tilde{m}_1$	a	b	b	b	$\dots$	b	b	b	b	$\dots$
		a	a	a	a	$\dots$	a	b	b	b	$\dots$
		a	a	a	a	$\dots$	a	a	b	b	$\dots$
		a	a	a	a	$\dots$	a	a	a	b	$\dots$
		$\vdots$	$\vdots$	$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$	$\vdots$	

Example

- m_1 is the only message undominated at θ_1.
- $\tilde{m}_1$ is the only message undominated at θ_1'.
- m_2 is the only undominated message at θ_2.
- $g(m_1, m_2) = b$ and $g(\tilde{m}_1, m_2) = a$ as required for implementing f.

The Jackson (1992) Result

- Recall that the incentive-compatibility compatibility condition ensured that truth-telling was an equilibrium in the direct mechanism.
- Since direct mechanisms are not useful here, what is the "appropriate" incentive-compatibility condition?
- NONE!

Jackson (1992) proved the following amazing result!

Theorem (Jackson 1992): Every SCF defined on an arbitrary domain (satisfying very weak conditions) is implementable in undominated strategies.

Every SCF defined on the complete domain $\mathcal{P}$ is implementable in undominated strategies.

A Difficulty

- The mechanism in the example works by using the following trick: infinite string of messages each dominating the one preceding it.
- This is unsatisfactory. For example, when player 1's (row player) preference is $aP(\theta_1)b$, she has no best-response among the "lower block" of messages.
- Jackson proposed further restrictions on mechanisms in oder to avoid this difficulty.

Bounded Mechanisms

The mechanism (M, g) is bounded if the following property is true: for all $i \in N$ and $\theta_i \in \Theta_i$:

If message $m_i \in M_i$ is weakly dominated at θ_i, there exists a message $m_i' \in M_i$ that weakly dominates m_i at θ_i and is undominated at θ_i.

If M_i is finite for all i, $G = (M, R)$ must be bounded. However non-finite mechanisms can also be bounded.

What are the SCCs that are implementable in undominated strategies by bounded mechanisms (IUSBM)?

Not a very easy question to answer. We will have a sufficient condition later.

A Restriction Imposed by Boundedness

- Suppose F can be IUSBM implemented. Let $G = (M, g)$ implement it.
- Let θ be a type-profile and $a \in F(\theta)$.
- Then there exists an (undominated) message profile at θ, $\tilde{m}$ such that $g(\tilde{m}) = a$.
- Let i be a player and θ_i' be another type for the player.
- Either $\tilde{m}_i$ is undominated at θ_i' or there exists another message m_i' that is *undominated* (by boundedness) at θ_i' which dominates $\tilde{m}_i$.

Strategy-Resistance

- In the former case $a \in F(\theta_i', \theta_{-i})$.
- In the latter case $b = g(m_i', \tilde{m}_{-i}) R(\theta_i') a$.
- In each case, there exists $b \in F(\theta_i', \theta_{-i})$ such that $bR(\theta_i')a$.
- Jackson (1992) refers to this condition as strategy-resistance and is necessary for IUSBM.
- Strategy-resistance is not sufficient for IUSBM (Ohseto (1994)).

Börgers (1991)

- Börgers (1991) raised the issue of implementing Pareto-efficient outcomes.
- Implementing specific sub-correspondences of the Pareto correspondence is not a problem. Dictatorship can be obviously implemented.
- However the correspondence of best-ranked alternatives can also be implemented.
  $F^T(\theta) = \{a | a \text{ is } P(\theta_i) \text{ maximal for some } i \in N\}$.

Implementing F^T: "pseudo-random" dictatorship

- Each player i announces an integer s_i in the set $\{1, \ldots, n\}$ and θ_i.
- Let r be the residue of $\sum_i s_i$ mod n.
- Outcome is the maximal element according to the ordering of the $r + 1$th player, i.e. the maximal element of $P(\theta_{r+1})$.
- Suppose i's true type is θ_i. Suppose θ_i' is such that $\max P(\theta_i) \neq \max P(\theta_i')$. Then the strategy (θ_i, s_i) weakly dominates (θ_i', s_i).
- Moreover, announcing true ordering and arbitrary integer is undominated.

Implementing Compromises

- Can one implement SCCs that contain outcomes that are not first-ranked by some player?
- Difficulty: Making sure that announcing false orderings is dominated seems to require agents to have a strategy that gives their maximal alternatives according to their true preferences. But this strategy may also weakly dominate the strategy that gives a compromise.

Implementing Compromises

- Börgers (1991) shows that a variant of approval voting implements the Pareto correspondence in the case of $m = 3$. Fails for $m \geq 4$.
- According to the paper "....if agents play "undominated strategies" it is not obvious how to ensure Pareto efficient collective decisions for all possible preference profiles without excluding compromises."
- Proves an impossibility result (for three alternatives or two players) by imposing an additional axiom: there exists a profile where choosing a top-ranked alternative is not allowed.

The Pareto-Correspondence: Mukherjee, Muto, Ramaekers, Sen (2019)

- MMRS(2019) the entire Pareto Correspondence *can* be implemented.
- Idea: Augmented modulo game.
- Players have strategies that give compromise outcomes (playing Green). They also have a strategy that gives the maximal element at any ordering (playing Blue). However there are situations where playing Green does strictly better than Blue.
- This can be done without violating Pareto-efficiency.

The Pareto Correspondence Result

Theorem

Consider the domain of strict orderings. The Pareto Correspondence can be implemented in IUSBM.

Mukherjee, Muto and Sen (2021)

Provide a sufficient condition for implementation in *finite* mechanims. involving three properties:

1. Strategy-resistance of SCC F.
2. Strategy-proofness of "range-top selections".
3. The "Flip Condition."

Provide several applications.

Range-top selection

- Fix SCC F. Let t^i be the following SCF: such for each θ, $t^i(\theta)$ is the most-preferred alternative among $F(\theta)$ w.r.t. $R(\theta_i)$.
- t^i is a range-top selection for i from F.
- F is generated by adding to t^i worse alternatives w.r.t. agent i's preference.

Range-top selection

- Suppose F satisfies strategy-resistance.
- Agent i cannot gain by misrepresentation of preference in t^i (actually, equivalent to strategy-resistance).
- However t^i is not strategy-proof.
- Agent $j \neq i$ could manipulate by affecting the range of F.

Key to our Approach

- Assume the existence of strategy-proof t^i for each i.
- At every profile add alternatives that are worse than t^i for each i but may be preferred by the designer.
- Implement the resulting SCC.

Examples

- Social Choice environment: t^i is dictatorial. Add "compromise" or other Pareto efficient alternatives.
- Add the full-surplus extraction outcome to the second-price auction outcome at every valuation profile.
- In public-good provision model add an efficient budget-balanced outcome to the VCG outcome.
- Add all other stable matchings to the man-optimal stable matching.
- In each case, the range-top selection is strategy-proof.

Extended Strategy-Resistance

Definition (Extended Strategy-Resistance, ESR)

The SCC F satisfies Extended Strategy-Resistance (ESR) if (i) F satisfies strategy-resistance and (ii) for each i, there exists a range-top selection t^i such that t^i is strategy-proof.

- Condition (ii) does not imply (i) in the definition of ESR.

The Flip Condition

Definition

The SCC F satisfies the Flip Condition if for each $i \in N$ and each θ_i, θ_i' there exist $x, y \in A$ such that

1. $xP(\theta_i)y$ and $yR(\theta_i')x$, and
2. for each each agent $j \neq i$, there exists a range-top selection t^j of j from F such that for all preference profiles $\bar{\theta}$, we have $t^j(\bar{\theta})R(\bar{\theta}_j)x$ and $t^j(\bar{\theta})R(\bar{\theta}_j)y$.

- x and y "flip" between θ_i and θ_i'.
- all agents $j \neq i$ "dislike" x, y. Need not be in the range of F at any profile.
- E.g., x and y involve a large monetary payment for $\forall j \neq i$.
- The Flip Condition holds in most "economic" environments.

A sufficient condition for implementation

Theorem

A SCC satisfying ESR and the Flip Condition is implementable in undominated strategies by a finite mechanism.

Proof uses a refinement of the modulo game called the extended modulo game. Similar ideas to MMRS.

Applications

- We apply this result to auctions, public good provision and stability.
- In each case that the designer can do “better” than in dominant strategies.
- Better in what sense?
- Involves comparisons between and implementable SCF (in dominant strategies) and an implementable SCC.

Ranking SCCs

- Börgers and Smith (2012)
- A SCC F *outperforms* another SCC G if at every θ, the planner weakly prefers any outcome given by F to that by G, and strictly prefers some outcome in $F(\theta)$ to that in $G(\theta)$.
- If a SCC F is constructed by adding new outcomes to a SCF G at each θ such that these outcomes are "more desirable" for the planner (at θ), then F outperforms G.
- F strictly outperforms G if F outperforms G and G does not outperform F.

Example: Auctions

- Single indivisible object, private values.
- Let $N = \{1, \ldots, n\}$ be the set of bidders.
- Each bidder i's valuation is θ_i. Payoff is $\theta_i - p_i$ if she gets object and pays p_i.
- $\Theta := \{\theta^k \in \Re_+ \mid k = 1, \ldots, K\}$ where $0 \leq \theta^1 < \theta^2 < \cdots < \theta^K$.
- $\emptyset$: seller keeps object.
- Outcome is a pair consisting of a bidder and a payment or $\emptyset$.
- $f^{\mathrm{II},r}$ is the SCF where for each θ, the highest θ_i gets object and pays second-highest θ_j. Seller's bid is r.

Example: Auctions

- Myerson (1981): assume i.i.d distribution there exists $r \geq 0$ such that $f^{\mathrm{II},r}$ maximizes expected revenue.
- $f^{\mathrm{II},r}$ is strategy-proof.
- Let f^{SE} be the "full extraction" SCF - at each θ, highest θ_i gets object and pays θ_i.
- Let $F(\theta) = f^{\mathrm{II},r} \cup f^{\mathrm{SE}}$.

Example: Auctions

- For every i, $t^i(\theta) = f^{\mathrm{II},r}(\theta)$.
- If i has the highest valuation, she wins the object but pays a lower amount in $f^{\mathrm{II},r}$.
- We know $f^{\mathrm{II},r}$ is strategy-proof.
- For a revenue-maximizing auctioneer, F outperforms $f^{\mathrm{II},r}$

Proposition

The SCC F can be implemented in undominated strategies by a finite mechanism.

Example: Public good provision

- Let $N = \{1, \ldots, n\}$ be the set of agents. These agents jointly decide whether to provide an indivisible public good. This decision is denoted by $g \in \{0, 1\}$, where $g = 1$ if the public good is produced, and $g = 0$ if not.
- Each agent $i \in N$ has a valuation θ_i (this is a private information for i) on the public good. Let $p_i \in \Re$ be a monetary transfer from i: her utility is $\theta_i g - p_i$.
- (Finite environment:) we assume that the valuations are discrete, and the set of valuations is $\Theta := \{\theta^k \in \Re_+ \mid k = 1, \ldots, K\}$, where $0 \leq \underline{\theta} = \theta^1 < \theta^2 < \cdots < \theta^K = \overline{\theta}$.
- A generic outcome is denoted by $(g, p_1, \ldots, p_n) \in A$.

Example: Public goods provision

- The cost of provision is $c > 0$; the cost function: gc.
- Assume $n\underline{\theta} < c < (n-1)\overline{\theta} + \underline{\theta}$;
  - the public good is still socially beneficial when only one agent has the lowest valuation and the others have the highest (no veto).
- Also for each $\theta \in \Theta^n$, $\sum_{i \in N} \theta_i \neq c$;
  - rules out complications caused by tie-breaking rules. Generically true.

Example: Public goods provision

For each valuation profile $\theta = (\theta_1, \ldots, \theta_n) \in \Theta^n$, let $g^*(\theta) \in \{0, 1\}$ be the socially optimal decision (maximizing the total surplus)

$$g^*(\theta) = \begin{cases} 1 & \text{if } \sum_{i \in N} \theta_i > c, \\ 0 & \text{otherwise.} \end{cases}$$

- The *pivotal* or *Vickrey–Clarke–Groves* (*VCG*) SCF is given by the following transfer function $p^{\mathrm{VCG}}(\theta)$ together with the socially optimal decision g^*:

  for each $i \in N$ and each $\theta \in \Theta^n$,

  $$p_i^{\mathrm{VCG}}(\theta) = g^*(\underline{\theta}, \theta_{-i})\underline{\theta} + \left(g^*(\theta) - g^*(\underline{\theta}, \theta_{-i})\right)\left(c - \sum_{j \in N \setminus \{i\}} \theta_j\right).$$

Example: public goods provision

- Well-known that (g^*, p^{VCG}) is strategy-proof but not budget balanced - leads to deficits that have to financed from outside.
- Define first-best transfers as follows:

$$p_i^{\mathrm{FB}}(\theta) = g^*(\theta)\left(p_i^{\mathrm{VCG}}(\theta) + \frac{\theta_i - p_i^{\mathrm{VCG}}(\theta)}{\sum_{j\in N}(\theta_i - p_i^{\mathrm{VCG}}(\theta))}\left(c - \sum_{j\in N} p_j^{\mathrm{VCG}}(\theta\right.\right.$$

- Note $\sum_{j\in N} p_i^{\mathrm{FB}}(\theta) = 0$ if $g^*(\theta) = 0$ and c if $g^*(\theta) = 1$.
- Payments adjusted proportionally to VCG payments to cover deficits - budget balanced.

Example:Public Good Provision

- Define F as follows:

$$F(\theta) = \left\{(g^*(\theta), p^{\mathrm{VCG}}(\theta)), (g^*(\theta), p^{\mathrm{FB}}(\theta))\right\}.$$

- F is the union of VCG and FB.
- For each i, $t^i(\theta)$ is given by $(g^*(\theta), p^{\mathrm{VCG}}(\theta))$ - FB payments are higher.
- Hence t^i is strategy-proof.

Public Good Provision

- Suppose the designer has the following lexicographic preferences: "first" care about maximizing social surplus, "second" care about minimizing budget deficit.
- F also outperforms $(g^*(\theta), p^{\mathrm{VCG}}(\theta))$ for such a designer.

Proposition

The SCC F can be implemented in undominated strategies by a finite mechanism.

Example: Two-sided matching

- Marriage Problem of Gale-Shapley.
- $N = \{1, \ldots, n\}$- set of men, and $W = \{w_1, \ldots, w_m\}$ -set of women. Let $\emptyset$ be the alternative "single".
- Every $i \in N$ has a strict preference over $W \cup \{\varnothing\}$: $\succ_i$ (private information to i). The set of all strict preferences over $W \cup \{\varnothing\}$ is denoted by $\mathcal{P}$.
- Every woman $w_j \in W$, she has a strict preference over $M \cup \{\emptyset\}$ denoted by $\succ_{w_j}$.
- Assume that each woman's preference $\succ_{w_j}$ is known to everyone, and fixed.

Example: Two-sided matchings

For each preference profile $\succ = (\succ_i)_{i \in N}$, a matching μ is *stable* at $\succ$ if

- there exists no man $i \in N$ such that $\varnothing \succ_i \mu(i)$,
- there exists no woman $w_j \in W$ such that $\varnothing \succ_{w_j} \mu^{-1}(w_j)$, and
- there exists no pair of man $i \in N$ and woman $w \in W$ such that $w \succ_i \mu(i)$ and $i \succ_w \mu^{-1}(w)$.

Let $\mathcal{S}(\succ) \subseteq A$ (set of all possible matchings) be the set of all stable matchings at $\succ$.

For each $\succ$, there exists a unique stable matching $\mu \in \mathcal{S}(\succ)$ such that for each $\mu' \in \mathcal{S}(\succ)$ and each $i \in N$, either $\mu(i) \succ_i \mu'(i)$ or $\mu(i) = \mu'(i)$: *man-optimal stable matching* at $\succ$.

In μ' each woman has the worst partner that she could have in any stable matching.

Example: Two-sided matchings

- Let $f^{\mathrm{MO}} : \mathcal{P}^n \to A$ be the SCF such that for each $\succ$, $f^{\mathrm{MO}}(\succ)$ is the man-optimal stable matching at $\succ$.

- $t^i \equiv f^{\mathrm{MO}}$. Known to be strategy-proof.
- If the planner considers women's preferences, S outperforms f^{MO}.

Proposition

The SCC S is implementable in undominated strategies by a finite mechanism.

Literature

- Ohseto (1994) proves impossibility for the plurality correspondence. Satisfies strategy-resistance. Hence strategy-resistance is not sufficient.
- Yamashita (2012) provides a necessity condition (chain dominance) stronger than strategy resistance.
- Yamashita (2015) gives a performance bound for mechanisms when players play undominated strategies.
- Caroll (2014) proves a complexity result.
- Mukherjee, Muto and Ramaekers (2016) provide a characterization of implementable SCCs when players satisfy the additional behavioural assumption of *partial honesty*. The condition is a stronger version of Yamashita's (2012) necessary condition.
- Li and Dworczak (2020) investigate cases where undominated implementation outperforms particular strategy-proof SCFs.