Pick-an-Object Mechanisms

Inácio Bó[1] Rustamdjan Hakimov[2]

[1] University of York [2] University of Lausanne & WZB

Econometric Society - North American Winter Meetings 2020
January 5th, 2021

Summary

- We present a new family of sequential revelation mechanisms, denoted **Pick-an-Object Mechanisms** (PAO).
  - Instead of asking for agents' preferences, we ask them to pick their allocation from menus.
  - Each agent leaves with the last object they picked when the process is over.
- We characterize the allocation rules that can be "sequentialized" and implemented in a truthful equilibrium by PAO mechanisms.
- Equilibrium *behavior* is closely related to the one in Obviously Strategy-proof mechanisms, but space of implementable rules includes familiar object allocation rules, as opposed to OSP.
- Experiments indicate that PAO and OSP improve upon direct mechanisms in terms of truthful behavior.

Introduction: centralized market design

- Centralized matching mechanisms without monetary transfers: many real-life appplications.
  - School choice, college admissions,
  - Residents to hospitals,
  - Organs to patients,
  - Refugees to localities...
- Challenge: some information is known by participants but not the designer
  - Dominant strategy implementation + revelation principle $\Longrightarrow$ focus on strategy-proof, direct mechanisms.
  - Strategic simplicity.

Introduction: Behavior in strategy-proof mechanisms

Experimental evidence:

- One-shot experiments: far from universal truthful reporting. Chen & Sönmez (2002, 2006); Pais & Pinter (2008).
- Mixed evidence on learning. Chen & Kesten (2019); Zhu (2017); Ding & Schotter (2019).

Empirical evidence of dominated reporting from participants:

- Rees-Jones (2016), Chen & Pereyra (2016), Hassidim et al. (2015), Shorrer & Sóvago (2017), Artemov et al. (2017).

Obvious Strategy-proofness

- Obvious strategy proofness (OSP) (Li, 2017)
  - A refinement of strategy-proofness.
  - Accounts for the extent to which participants can easily understand its incentives.
  - Could explain why, for example, we observe more truthful behavior under a clock auction than under a sealed-bid second-price auction (Kagel et al., 1987).
  - Li, 2017 replicates second-price vs clock auction results, and also obtains similar results comparing direct RSD with OSP-RSD.

Obvious Strategy-proofness

- Pycia and Troyan, 2019
  - Characterize OSP mechanisms in the domain of object allocation as equivalent to millipede games. (Bade and Gonczarowski (2017) has similar results)
  - Decision node:
    - Leave with one of the objects in a menu,
    - At most one "continuation" action.
  - Obviously dominant strategy: choose continuation action unless the most preferred object among all available at some menu in the future is present. In that case, choose it.
- Might help explain "better" behavior in OSP mechanisms: no need to engage in counterfactual reasoning. Only the set of allocations that are still feasible.

Obvious Strategy-proofness

- OSP is, however, very restrictive. Commonly used rules are not OSP-implementable.
  - TTC (Li, 2017).
  - Stable allocations (Ashlagi and Gonczarowski, 2018).
- Bó and Hakimov (2019) and Klijn et al. (2019)
  - Iterative (sequential) DA.
  - Equilibrium behavior consists of picking the most preferred feasible option.
  - Higher rate of truthful behavior than the strategy-proof direct DA, despite not being a dominant strategy equilibrium.

Pick-an-Object Mechanisms

- We provide an alternative explanation for these results.
- Better results in OSP mechanisms could be better explained by the simple mechanics involved in the equilibrium: it is safe to simply pick your best feasible alternative whenever possible and "leave with it".
- Based on this, we introduce the class of Pick-an-Object mechanisms

Pick-an-Object Mechanisms

- **Step** 1: All agents are asked to choose an object from individualized menus.
- **Step** $k > 1$: Either
  - Agents are assigned to the last object they picked (which may include the null option) and the procedure ends, or
  - Some agents are given new menus, which are strict subsets of the previous one, not containing the last choice.

Equilibrium Strategies

- OSP Mechanisms
  - "Wait until you can pick your best feasible object."
- PAO Mechanisms
  - "Pick your best feasible object and wait to see if you can keep it."
- **Question**: By using PAO mechanisms, can we implement more familiar allocation rules while keeping (some of) the behavioral/experimental advantages of OSP mechanisms?

Model

- Finite set of object types $O = \{o_1, o_2, \ldots, o_m\} \cup \{\emptyset\}$
- Finite set of agents $A = \{a_1, a_2, \ldots, a_n\}$, each with strict preferences P_a over the set O. $\mathcal{P}$ is the set of all preference profiles.
- An allocation is a function $\mu : A \to O$. $\mathcal{M}$ is the set of all allocations.
- A rule is a function $\varphi : \mathcal{P} \to \mathcal{M}$.

Sequentialization in PAO Mechanisms

Definition

An agent follows a straightforward strategy with respect to P if whenever presented with a menu $I \subseteq O$, she chooses the most preferred element of I according to P.

Definition

A Pick-an-Object mechanism sequentializes the rule φ if, for any preference profile P, the Pick-an-Object mechanism provides menus such that when each agent a_i follows the straightforward strategy with respect to P_{a_i}, the outcome $\varphi(P)$ is produced.

The allocation/information trade-off

- Information about an agent's preferences can only be obtained from choices from menus.
- Menus don't include previous choices + allocation is the last object chosen $\implies$ more information "costs" ruling out last choice as the allocation.

Monotonic Discoverability

Definition

Let μ be an allocation, and P a preference profile. We say that $\mathcal{L}(P, \mu)$ are the **continuation profiles of P at μ**.

(a) Preferences and allocation

P_{a_1}	o_1	o_2	o_3	o_4	o_5	$\emptyset$
P_{a_2}	o_3	o_1	o_4	o_3	o_5	$\emptyset$
P_{a_3}	o_2	o_5	o_4	o_1	o_3	$\emptyset$
P_{a_4}	o_5	o_4	o_2	$\emptyset$	o_1	o_3

(b) Continuation profiles

P_{a_1}	o_1	o_2	o_3			
P_{a_2}	o_3	o_1				
P_{a_3}	o_2	o_5	o_4	o_1		
P_{a_4}	o_5	o_4	o_2	$\emptyset$		

Definition

A rule φ satisfies **monotonic discoverability** if, for any allocation μ and preference profile P, either $\varphi(P) = \mu$ or there is an agent $a^* \in A$ such that $P' \in \mathcal{L}(P, \mu) \implies \mu(a^*) \neq \varphi_{a^*}(P')$.

Monotonic Discoverability

Theorem

There exists a Pick-an-Object mechanism that **sequentializes** an individually rational rule φ if and only if φ satisfies monotonic discoverability.

Generalized DA Procedures

- Rules that can be represented by a deferred acceptance procedure, in which:
  - Proposals work as in the Gale-Shapley DA, following participants' strict preferences,
  - Whether an agent is tentatively matched or rejected is a function of the **entire tentative allocation**, and **all the contemporaneous proposals**.

Remark

Generalized DA procedures include (i) Gale-Shapley Deferred Acceptance, (ii) Top Trading Cycles, and (iii) Boston Mechanism.

Generalized DA Procedures

Proposition

If φ can be described by a generalized DA procedure, then φ satisfies monotonic discoverability.

OPBE

Definition

A strategy profile σ together with a belief system ω is an **ordinal perfect Bayesian equilibrium (OPBE)** if for every $a \in A$, every $h^A \in H^A_{\mathbb{S}}$, and every strategy $\sigma_a{}'$ for agent a, the outcome under σ_a **first-order stochastically dominates** the one under $\sigma_a{}'$ for a.

Definition

A strategy profile σ is a **robust ordinal perfect Bayesian equilibrium** if for every belief system ω, σ is an OPBE.

Definition

A rule φ is **non-bossy** if
$\varphi_i(P_i, P_{-i}) = \varphi_i(P_i', P_{-i}) \implies \varphi(P_i, P_{-i}) = \varphi(P_i', P_{-i})$.

Implementation in Pick-an-Object mechanisms

Theorem

A rule is Pick-an-Object implementable in robust ordinal perfect Bayesian equilibrium if and only if it is is strategy-proof and satisfies monotonic discoverability.

Theorem

Every non-bossy OSP implementable rule is Pick-an-Object implementable in weakly dominant strategies.

Questions

1. Does a Pick-an-Object environment makes participants play the equilibrium strategies more often when compared to direct mechanisms?

2. What drives the truthful behavior in sequential OSP mechanisms - the "picking" mechanics or the stronger incentive properties?

Design of the Experiment.

- Eight objects. Eight participants.
- 21 rounds, each corresponds to a new market.
- Treatments:

Rounds	Direct	PAO	OSP
1-7	TTC cyc	TTC cyc	TTC acyc
8-14	TTC acyc	TTC acyc	TTC acyc
15-21	SD	SD	SD

- Within-subjects: TTC-cyc; TTC-acyc; SD
- Between-subjects: Direct; PAO; OSP

Results

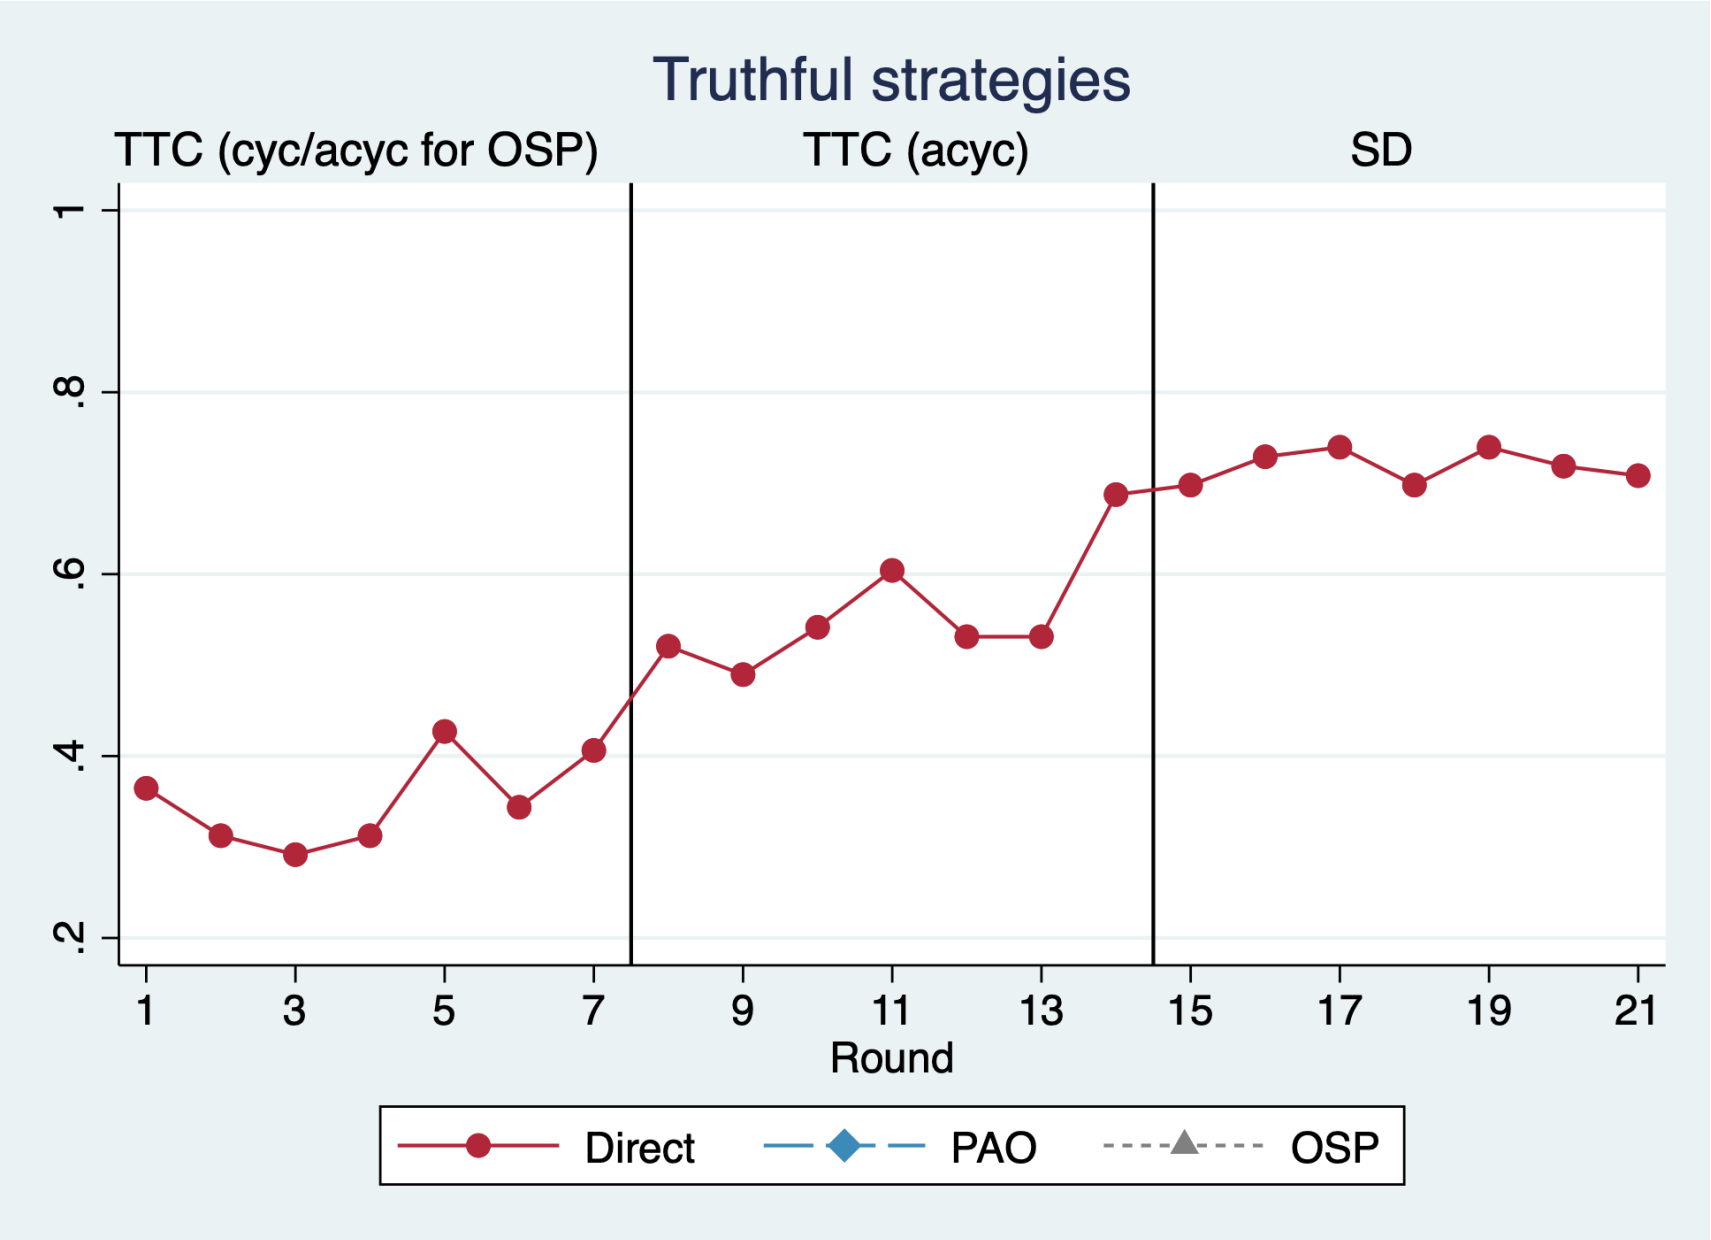
Truthful strategies
TTC (cyc/acyc for OSP)
TTC (acyc)
SD
1
.8
.6
.4
.2
1
3
5
7
9
11
13
15
17
19
21
Round
Direct
PAO
OSP

Results

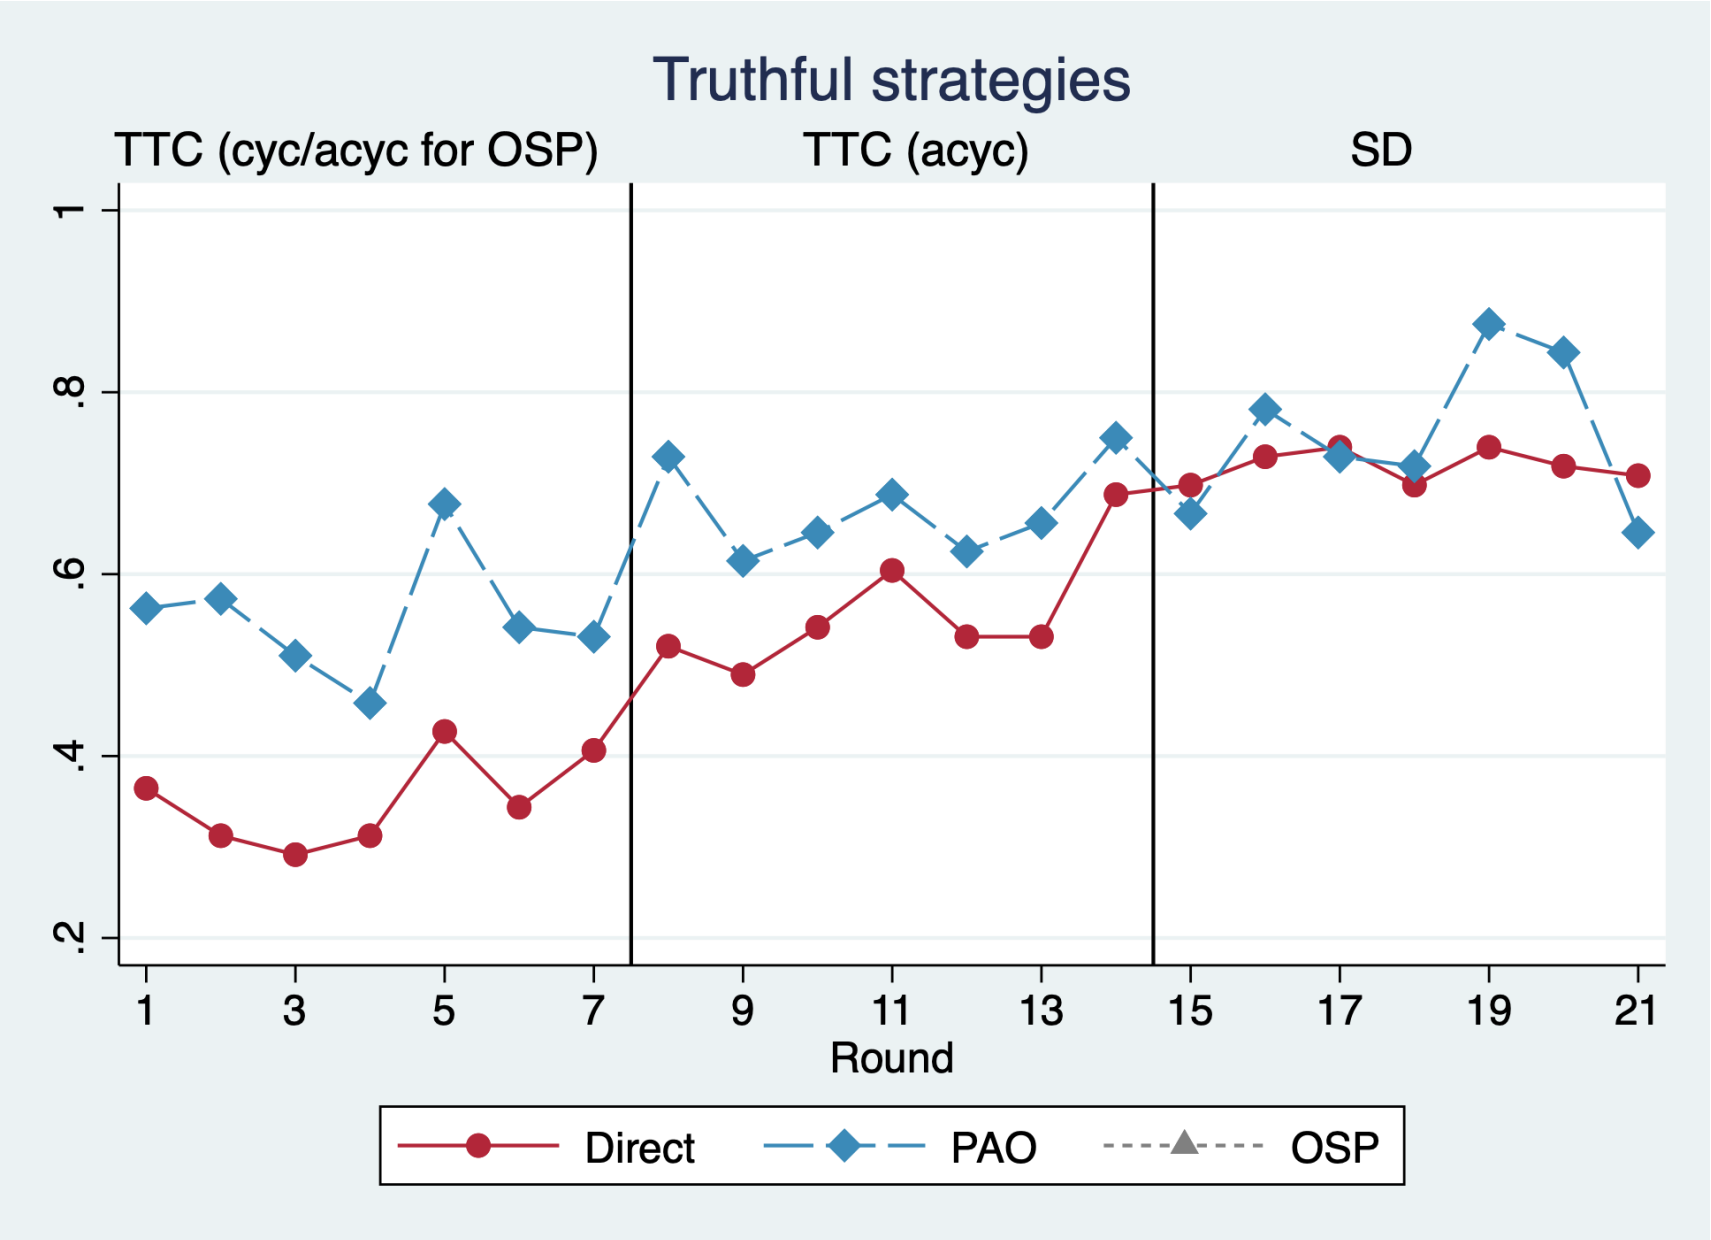
Truthful strategies
TTC (cyc/acyc for OSP)
TTC (acyc)
SD
1
.8
.6
.4
.2
1
3
5
7
9
11
13
15
17
19
21
Round
Direct
PAO
OSP

Results

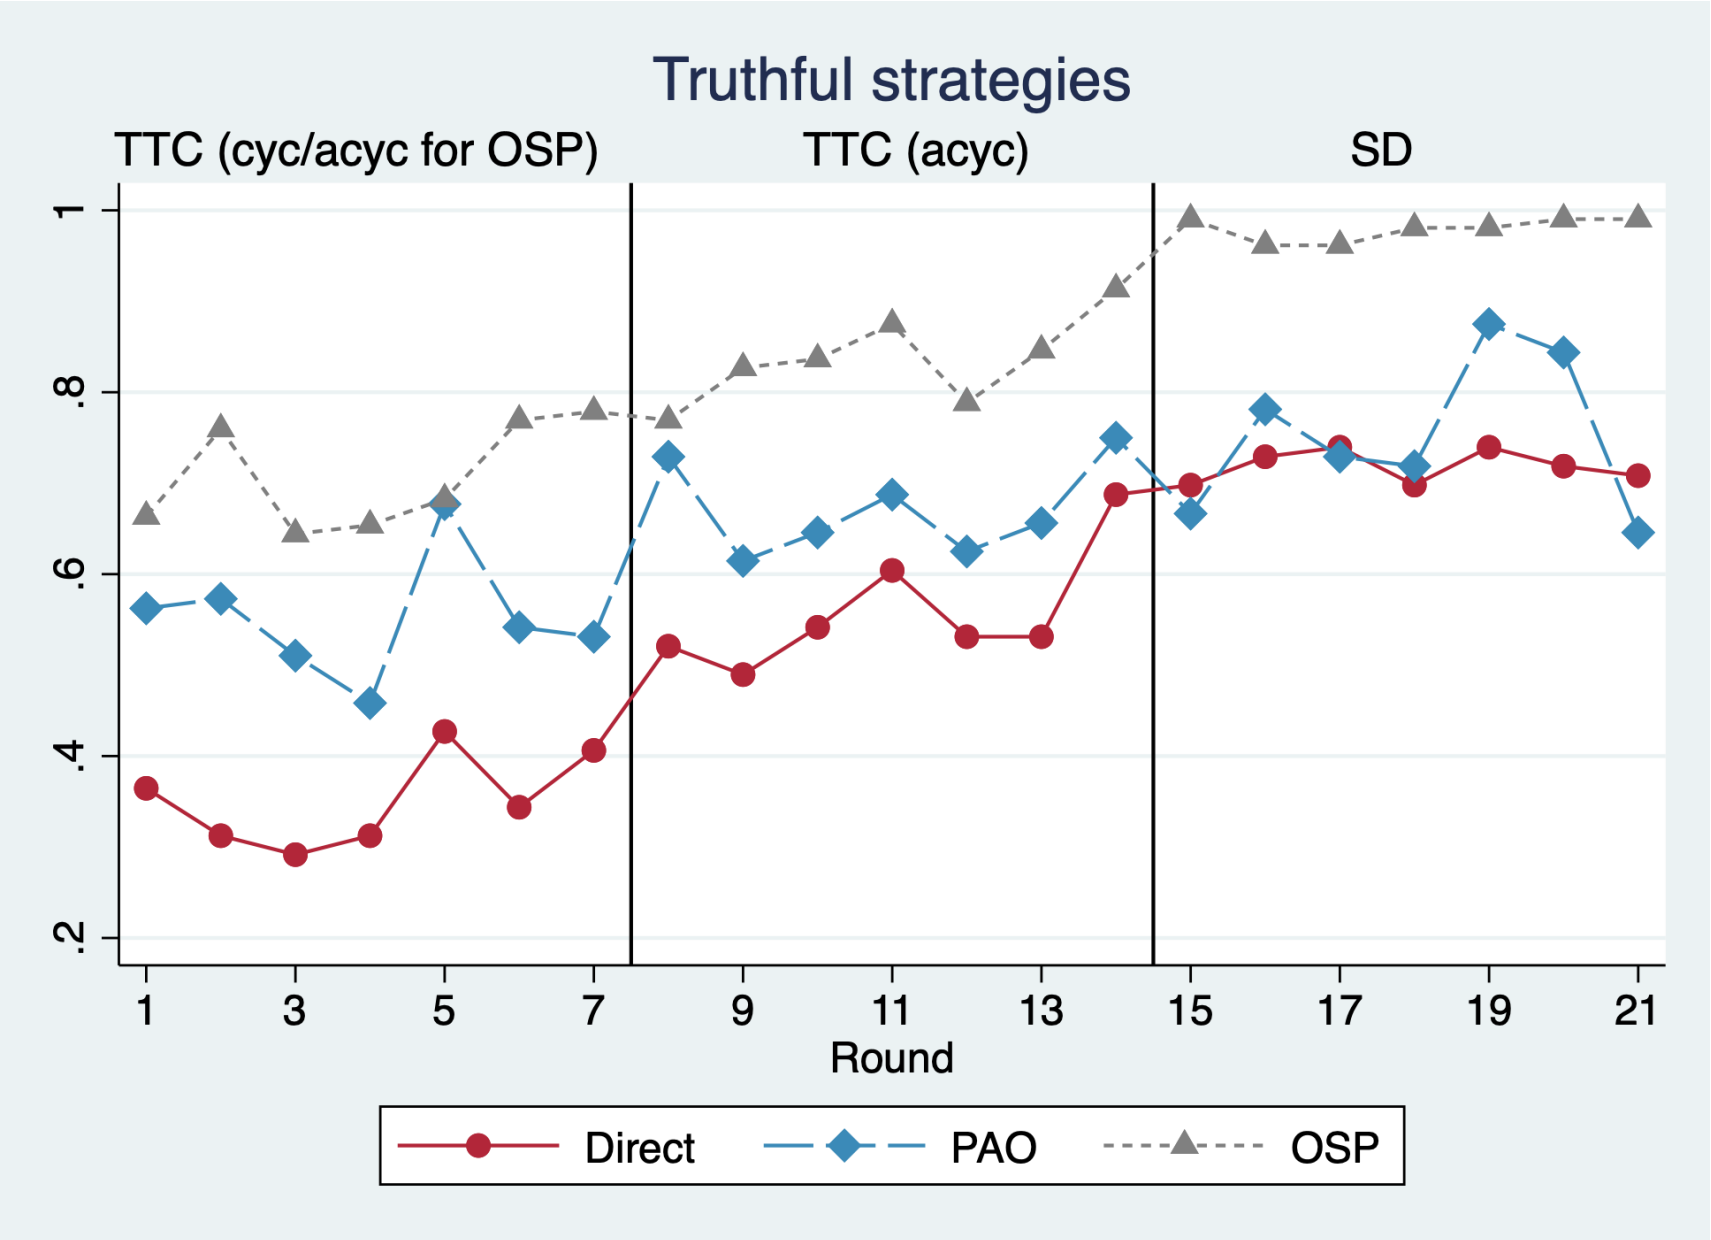
Truthful strategies
TTC (cyc/acyc for OSP)
TTC (acyc)
SD
1
.8
.6
.4
.2
1
3
5
7
9
11
13
15
17
19
21
Round
Direct
PAO
OSP

Results

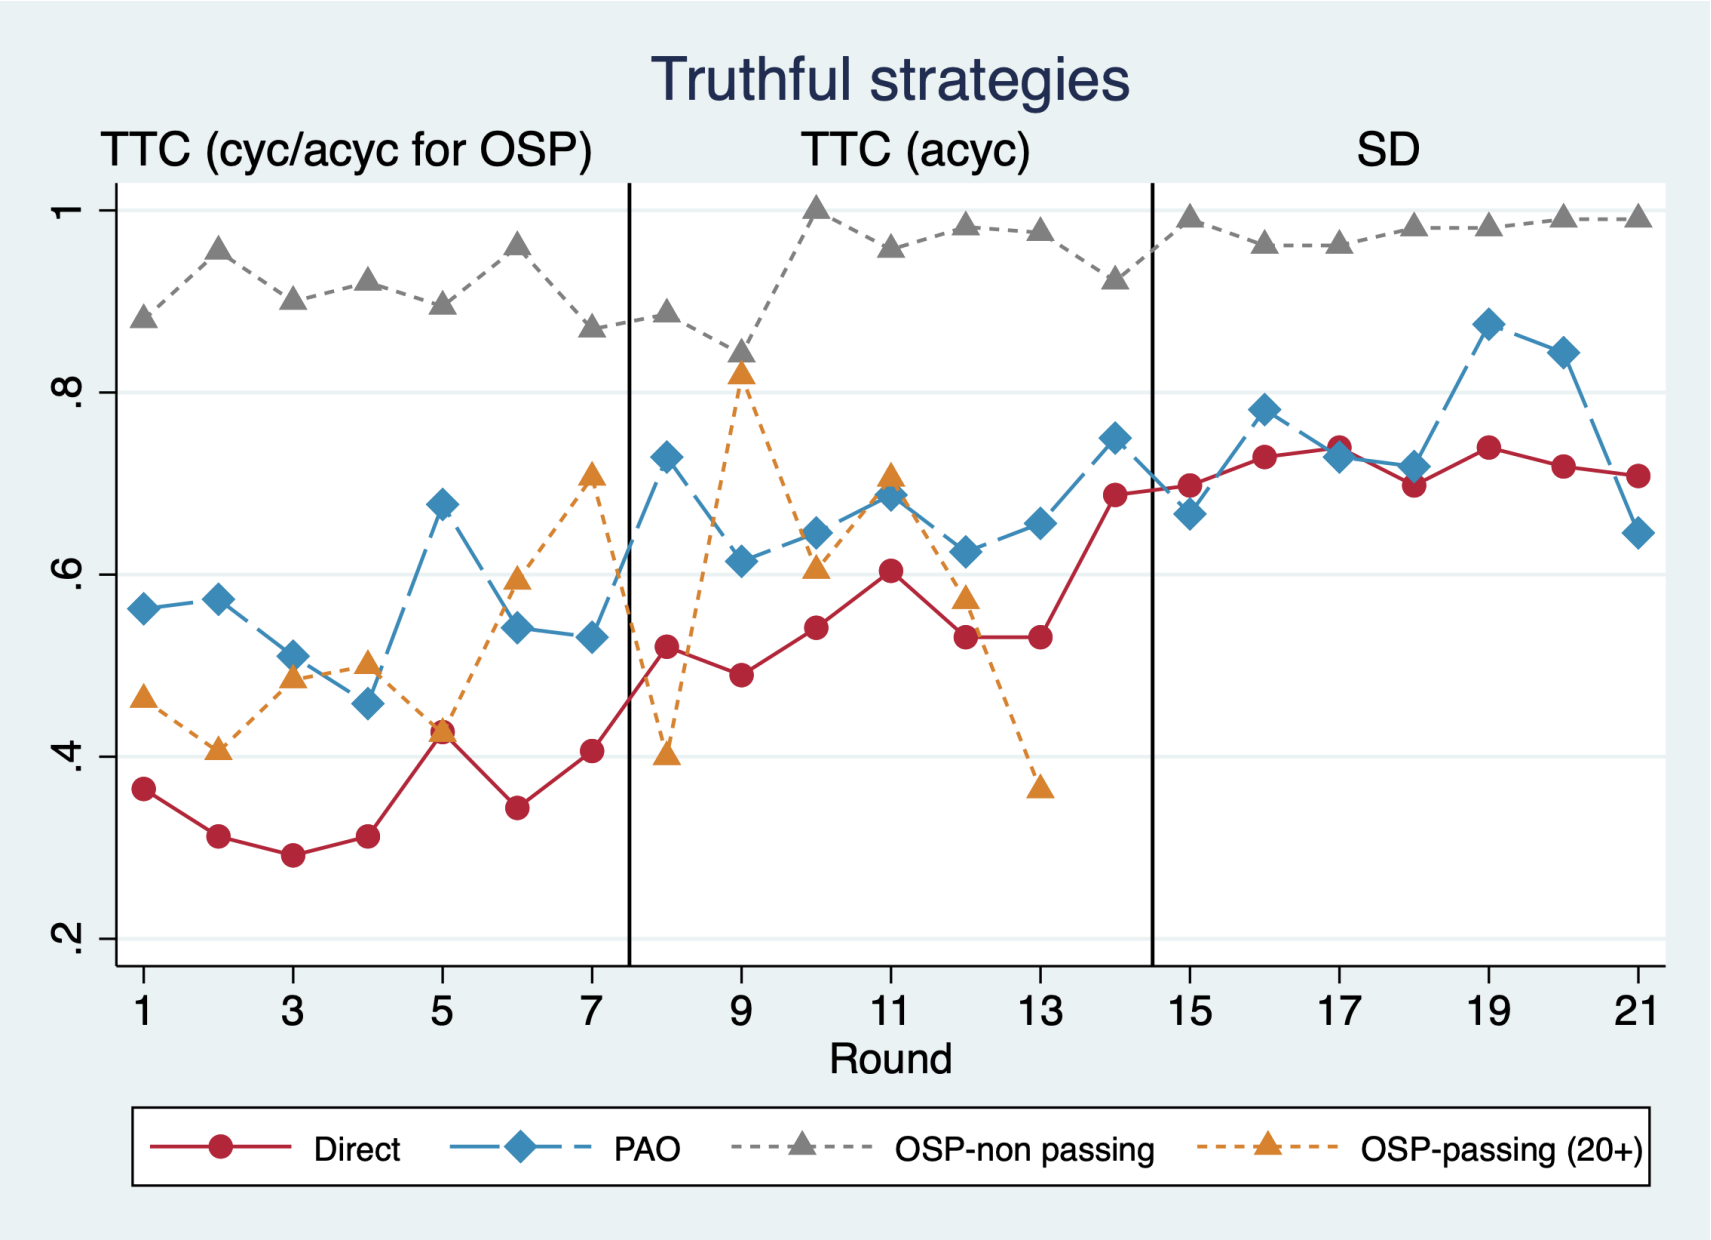
Truthful strategies
TTC (cyc/acyc for OSP)
TTC (acyc)
SD
1
.8
.6
.4
.2
1
3
5
7
9
11
13
15
17
19
21
Round
Direct
PAO
OSP-non passing
OSP-passing (20+)

Conclusions

- We propose the family of Pick-an-Object mechanisms for implementing object allocation rules.
  - Monotonic discoverability characterizes sequentializable allocation rules.
  - MD+SP characterizes allocation rules implementable in robust OPBE.
- One can implement many OSP allocation rules via PAO mechanisms, while following similar equilibrium strategies.
- Experiments show that PAO environment increases the rates of truthful reporting relative to direct mechanism, especially in TTC, but cannot reach the levels of OSP mechanisms.
- In OSP mechanisms the rates are almost universal for "clinching" decisions, but much lower for decisions of "passing".