Discussion Papers in Economics

Ordinal Bayesian incentive compatibility in random assignment model

Sulagna Dasgupta and Debasis Mishra

September 2020

Discussion Paper 20-06

Indian Statistical Institute, Delhi
Economics and Planning Unit
7, S. J. S. Sansanwal Marg, New Delhi, 110016, India

Ordinal Bayesian incentive compatibility in random assignment model *

Sulagna Dasgupta and Debasis Mishra †

September 29, 2020

Abstract

We explore the consequences of weakening the notion of incentive compatibility from strategy-proofness to ordinal Bayesian incentive compatibility (OBIC) in the random assignment model. If the common prior of the agents is a *uniform prior*, then a large class of random mechanisms are OBIC with respect to this prior – this includes the probabilistic serial mechanism. We then introduce a robust version of OBIC: a mechanism is *locally robust OBIC* if it is OBIC with respect *all* independent priors in some neighborhood of a given independent prior. We show that every locally robust OBIC mechanism satisfying a mild property called *elementary monotonicity* is strategy-proof. This leads to a strengthening of the impossibility result in Bogomolnaia and Moulin (2001): if there are at least four agents, there is no locally robust OBIC and ordinally efficient mechanism satisfying equal treatment of equals.

Keywords. ordinal Bayesian incentive compatibility, random assignment, probabilistic serial mechanism.

JEL Code. D47, D82

*We are grateful to Sven Seuken, Timo Mennle, Arunava Sen, Dipjyoti Majumdar, Souvik Roy, and Wonki Cho for their comments.

†Dasgupta: University of Chicago, ; Mishra: Indian Statistical Institute, Delhi, dmishra@isid.ac.in

1 Introduction

The paper explores the consequences of weakening incentive compatibility from strategy-proofness to *ordinal Bayesian incentive compatibility* in the random assignment model (one-sided matching model). Ordinal Bayesian incentive compatibility (OBIC) requires that the truth-telling *expected share vector* of an agent first-order stochastically dominates the expected share vector from reporting any other preference. This weakening of strategy-proofness in mechanism design models without transfers was proposed by d'Aspremont and Peleg (1988). We study OBIC by considering mechanisms that allow for randomization in the assignment model.

In the random assignment model, the set of mechanisms satisfying ex-post efficiency and strategy-proofness is quite rich.[1] Despite satisfying such strong incentive properties, all of them either fail to satisfy *equal treatment of equals*, a weak notion of fairness, or *ordinal efficiency*, a stronger but natural notion of efficiency than ex-post efficiency. Indeed, Bogomolnaia and Moulin (2001) propose a new mechanism, called the *probabilistic serial mechanism*, which satisfies equal treatment of equals and ordinal efficiency. However, they show that it fails strategy-proofness, and no mechanism can satisfy all these three properties simultaneously if there are at least four agents. A primary motivation for weakening the notion of incentive compatibility to OBIC is to investigate if we can escape this impossibility result.

We show two types of results. First, if the (common) prior is a uniform probability distribution over the set of possible preferences, then every *neutral* mechanism satisfying a mild property called *elementary monotonicity* is OBIC.[2] An example of such a mechanism is the probabilistic serial mechanism. This is a positive result and provides a strategic foundation for the probabilistic serial mechanism. In particular, it shows that there exists

[1] Pycia and Ünver (2017) characterize the set of deterministic, strategy-proof, Pareto efficient, and non-bossy mechanisms in this model. This includes generalizations of the top-trading-cycle mechanism.

[2] Neutrality is a standard axiom in social choice theory which requires that objects are treated symmetrically. Elementary monotonicity is a mild monotonicity requirement of a mechanism. We define it formally in Section 4.

ordinally efficient mechanisms satisfying equal treatment of equals which are OBIC with respect to the uniform prior.

Second, we explore the implications of strengthening OBIC as follows. A mechanism is *locally robust OBIC* (LROBIC) with respect to an independent prior if it is OBIC with respect to every independent prior in its "neighborhood". The motivation for such requirement of robustness in the mechanism design literature is now well-known, and referred to as the Wilson doctrine (Wilson, 1987). We show that every LROBIC mechanism satisfying elementary monotonicity is strategy-proof. An immediate corollary of this result is that the probabilistic serial mechanism is not LROBIC (though it is OBIC with respect to the uniform prior). As a corollary, we can show that when there are at least four agents, there is no LROBIC and *ordinally efficient* mechanism satisfying equal treatment of equals. This strengthens the seminal impossibility result of Bogomolnaia and Moulin (2001) by replacing strategy-proofness with LROBIC.

Both our results point to very different implications of OBIC in the presence of elementary monotonicity – if the prior is uniform, this notion of incentive compatibility is very permissive; but if we require OBIC with respect to a set of independent priors in any neighborhood of a given prior, this notion of incentive compatibility is very restrictive. As we discuss in Section 6, such implications have been known for *deterministic voting* models (Majumdar and Sen, 2004). But ours is the first paper to point this out for the *random assignment* model.

2 Model

Assignments. There are n agents and n objects.[3] Let $N := \{1, \dots, n\}$ be the set of agents and A be the set of objects. We define the notion of a feasible assignment first.

[3]All our results extend even if the number of objects is not the same as the number of agents. We assume this only to compare our results with the random assignment literature, where this assumption is common. Also, whenever we say an assignment, we mean a random assignment from now on.

Definition 1 *An $n \times n$ matrix L is an* **assignment** *if*

$$\begin{aligned} L_{ia} &\in [0,1] \qquad \forall\ i \in N,\ \forall\ a \in A \\ \sum_{a \in A} L_{ia} &= 1 \qquad \forall\ i \in N \\ \sum_{i \in N} L_{ia} &= 1 \qquad \forall\ a \in A \end{aligned}$$

For any assignment L, we write L_i as the **share vector** of agent i. Formally, a share vector is a probability distribution over the set of objects. For any $i \in N$ and any $a \in A$, L_{ia} denotes the "share" of agent i of object a. The second constraint of the assignment definition requires that the total share of every agent is 1. The third constraint of the assignment requires that every object is completely assigned. Let $\mathcal{L}$ be the set of all assignments.

An assignment L is **deterministic** if $L_{ia} \in \{0,1\}$ for all $i \in N$ and for all $a \in A$. Let $\mathcal{L}^d$ be the set of all deterministic assignments. By the Birkohff-von-Neumann theorem, for every $L \in \mathcal{L}$, there exists a set of deterministic assignments in $\mathcal{L}^d$ whose convex combination generates L.

Preferences. The preference (a strict ordering) of an agent i over A will be denoted by P_i. The set of all strict preferences over A is denoted by $\mathcal{P}$. A preference profile is $\mathbf{P} \equiv (P_1, \ldots, P_n)$, and we will denote by P_{-i} the preference profile $\mathbf{P}$ excluding the preference P_i of agent i.

Prior. We assume that the preference of each agent is independently and identically drawn using a common prior μ, which is a probability distribution over $\mathcal{P}^n$. We will denote by $\mu(P_i)$ the probability with which agent i has preference P_i. With some abuse of notation, we will denote the probability with which agents in $N \setminus \{i\}$ have preference profile P_{-i} as $\mu(P_{-i})$. Note that by independence, $\mu(P_{-i}) = \times_{j \neq i} \mu(P_j)$.

3 Ordinal Bayesian incentive compatibility

Our solution concept is Bayes-Nash equilibrium but we restrict attention to ordinal mechanisms, i.e., mechanisms where we only elicit ranking over objects from each agent. Hence, whenever we say *mechanism*, we refer to such ordinal mechanisms.[4] Formally, a **mechanism** is a map $Q : \mathcal{P}^n \to \mathcal{L}$. A mechanism Q assigns a share vector $Q_i(\mathbf{P})$ to agent i at every preference profile $\mathbf{P}$.

Before discussing the notions of incentive compatibility, it is useful to think how agents compare share vectors in our model. Fix agent i with a preference P_i over the set of objects A. Denote the k-th ranked object in P_i as $P_i(k)$. Consider two share vectors π, π'. For every $a \in A$, we will denote by π_a and π'_a the share assigned to object a in π and π' respectively. We will say π **first-order-stochastically-dominates (FOSD)** π' according to P_i if

$$\sum_{k=1}^{\ell} \pi_{P_i(k)} \geq \sum_{k=1}^{\ell} \pi'_{P_i(k)} \qquad \forall\ \ell \in \{1, \ldots, n\}.$$

In this case, we will write $\pi \succ_{P_i} \pi'$. Notice that $\succ_{P_i}$ is not a complete relation over the outcomes. An equivalent (and well known) definition of $\succ_{P_i}$ relation is that for *every* von-Neumann-Morgenstern utility representation of P_i, the expected utility from π is at least as much as π'.

The most standard notion of incentive compatibility is strategy-proofness (dominant strategy incentive compatibility), which uses the FOSD relation to compare share vectors.

Definition 2 *A mechanism Q is* **strategy-proof** *if for every $i \in N$, every $P_{-i} \in \mathcal{P}^{n-1}$, and every $P_i, P'_i \in \mathcal{P}$, we have*

$$Q_i(P_i, P_{-i}) \succ_{P_i} Q_i(P'_i, P_{-i}).$$

The interpretation of this definition is that fixing the preferences of other agents, the truth-telling share vector must FOSD other share vectors that can be obtained by deviation. This definition of strategy-proofness appeared in Gibbard (1977) for voting problems, and has been the standard notion in the literature on random voting and random assignment problems.

The ordinal Bayesian incentive compatibility notion is an adaptation of this by changing the solution concept to Bayes-Nash equilibrium. It was first introduced and studied in a

[4] The restriction to not consider cardinal mechanisms is arguably arbitrary. But it is consistent with the literature on random assignment models. We do not know how the set of incentive compatible mechanisms expand if we consider cardinal mechanisms.

voting committee model in d'Aspremont and Peleg (1988), and was later used in many voting models (Majumdar and Sen, 2004). To define it formally, we introduce the notion of an interim share vector. Fix an agent i with preference P_i. Given a mechanism Q, the **interim share** of object a for agent i by reporting P'_i is:

$$q_{ia}(P'_i) = \sum_{P_{-i} \in \mathcal{P}^{n-1}} \mu(P_{-i}) Q_{ia}(P'_i, P_{-i}).$$

The **interim share vector** of agent i by reporting P'_i will be denoted as $q_i(P'_i)$.

Definition 3 *A mechanism Q is* **ordinally Bayesian incentive compatible (OBIC)** *(with respect to prior μ) if for every $i \in N$ and every $P_i, P'_i \in \mathcal{P}$, we have*

$$q_i(P_i) \succ_{P_i} q_i(P'_i).$$

It is immediate that if a mechanism Q is strategy-proof it is OBIC with respect to any prior. Conversely, if a mechanism is OBIC with respect to *all* priors (including correlated priors), then it is strategy-proof.

3.1 A motivating example

We investigate a simple example to understand the implications of strategy-proofness and OBIC for the probabilistic serial mechanism. Suppose $n = 3$ with three objects $\{a, b, c\}$. Consider the preference profiles (P_1, P_2, P_3) and (P'_1, P_2, P_3) shown in Table 3.1 – the table also shows the share vector of each agent in the probabilistic serial mechanism of Bogomolnaia and Moulin (2001). In the probabilistic serial mechanism, each agent starts "eating" her favorite object simultaneously till the object is finished. Then, she moves to the best available object according to her preference and so on. Table 3.1 shows the output of the probabilistic serial mechanism for preference profiles (P_1, P_2, P_3) and (P'_1, P_2, P_3). Since $Q_{1a}(P_1, P_2, P_3) + Q_{1c}(P_1, P_2, P_3) > Q_{1a}(P'_1, P_2, P_3) + Q_{1c}(P'_1, P_2, P_3)$, we conclude that $Q_1(P'_1, P_2, P_3) \not\succ_{P'_1} Q_1(P_1, P_2, P_3)$. Hence, agent 1 can manipulate from P'_1 to P_1, when agents 2 and 3 have preferences (P_2, P_3).

P_1	P_2	P_3	P'_1	P_2	P_3
$a;\ \frac{1}{2}$	$a;\ \frac{1}{2}$	$c;\ \frac{3}{4}$	$c;\ \frac{1}{2}$	$a;\ \frac{2}{3}$	$c;\ \frac{1}{2}$
$c;\ \frac{1}{4}$	$b;\ \frac{1}{2}$	$a;\ 0$	$a;\ \frac{1}{6}$	$b;\ \frac{1}{3}$	$a;\ \frac{1}{6}$
$b;\ \frac{1}{4}$	$c;\ 0$	$b;\ \frac{1}{4}$	$b;\ \frac{1}{3}$	$c;\ 0$	$b;\ \frac{1}{3}$

Table 1: Manipulation by agent 1.

When can such a manipulation be prevented by OBIC? Note that P_1 is generated from P'_1 by permuting a and c. Suppose we permute P_2 and P_3 also to get P'_2 and P'_3 respectively:

$$c\ P'_2\ b\ P'_2\ a \text{ and } a\ P'_3\ c\ P'_3\ b.$$

Since the probabilistic serial mechanism is *neutral* (with respect to objects), the share vector of agent 1 at (P'_1, P_2, P_3) is an (a, c) permutation of its share vector at (P_1, P'_2, P'_3). Further, when all the preferences are equally likely, the probability of (P_2, P_3) is equal to the probability of (P'_2, P'_3). So, the total expected probability of a and c for agent 1 at P_1 and P'_1 is the same (where expectation is taken over (P_2, P_3) and (P'_2, P'_3)). As we show below, this argument generalizes and the expected share vector at P'_1 first-order-stochastic-dominates the expected share vector at P_1 when the true preference is P'_1 and prior is uniform.

4 Uniform prior and possibilities

In this section, we present our first result which shows that the set of OBIC mechanisms is much larger than the set of strategy-proof mechanisms if the prior is the uniform prior. A prior μ is the **uniform prior** if $\mu(P_i) = \frac{1}{|\mathcal{P}|} = \frac{1}{n!}$ for each $P_i \in \mathcal{P}$. Uniform prior puts equal probability on each of the possible preferences. We call a mechanism **U-OBIC** if it is OBIC with respect to the uniform prior.

We show that there is a large class of mechanisms which are U-OBIC - this will include some well-known mechanisms which are known to be not strategy-proof. This class is characterized by two axioms, neutrality and elementary monotonicity, which we define next. To define neutrality, consider any permutation $\sigma : A \to A$ of the set of objects. For every preference P_i, let P_i^{σ} be the preference generated when the permutation σ is applied to P_i. Let $\mathbf{P}^{\sigma}$ be the preference profile generated by permuting each preference in the preference profile $\mathbf{P}$ by the permutation σ.

Definition 4 *A mechanism Q is* **neutral** *if for every $\mathbf{P}$ and every permutation σ,*

$$Q_{ia}(\mathbf{P}) = Q_{i\sigma(a)}(\mathbf{P}^{\sigma}) \qquad \forall\ i \in N,\ \forall\ a \in A.$$

Neutrality requires that objects be treated symmetrically by the mechanism. Any mechanism which does not use the "names" of the objects is neutral – this includes the random priority mechanisms, the simultaneous eating algorithm mechanisms (including the probabilistic serial mechanism) in Bogomolnaia and Moulin (2001).

Our next axiom is elementary monotonicity, an axiom which requires a mild form of monotonicity. To define it, we need the notion of "adjacency" of preferences. We say preferences P_i and P'_i are **adjacent** if there exists a $k \in \{1, \ldots, n-1\}$ such that

$$P_i(k) = P'_i(k+1), P_i(k+1) = P'_i(k),\ \text{and}\ P_i(k') = P'_i(k')\ \forall\ k' \notin \{k, k+1\}.$$

In other words, P'_i is obtained by swapping consecutively ranked objects in P_i. Here, if $P_i(k) = a$ and $P_i(k+1) = b$, we say that P'_i is an (a, b)-swap of P_i.

Definition 5 *A mechanism Q satisfies* **elementary monotonicity** *if for every $i \in N$, every $P_{-i} \in \mathcal{P}^{n-1}$, and every $P_i, P'_i \in \mathcal{P}$ such that P'_i is an (a, b)-swap of P_i, we have*

$$Q_{ib}(P'_i, P_{-i}) \geq Q_{ib}(P_i, P_{-i})$$
$$Q_{ia}(P'_i, P_{-i}) \leq Q_{ib}(P_i, P_{-i}).$$

In other words, as agent i lifts alternative b in ranking by one position by swapping it with a (and keeping the ranking of every other object the same), elementary monotonicity requires that the share of object b should weakly increase for agent i, while share of object a should weakly decrease. It is not difficult to see that elementary monotonicity is a necessary condition for strategy-proofness. As we show later, elementary monotonicity is satisfied by a variety of mechanisms - including those which are not strategy-proof. However, every neutral mechanism satisfying elementary monotonicity is U-OBIC.

Theorem 1 *Every neutral mechanism satisfying elementary monotonicity is U-OBIC.*

Proof: Fix a neutral mechanism Q satisfying elementary monotonicity. The proof goes in various steps.

Step 1. Pick an agent i and two preferences P_i and P_i'. Pick any $k \in \{1, \ldots, n\}$ and suppose $P_i(k) = a$ and $P_i'(k) = b$. We show that $q_{ia}(P_i) = q_{ib}(P_i')$. This is a consequence of uniform prior and neutrality. To see this, let $P_i' = P_i^\sigma$ for some permutation σ of objects in A. Then, $b = \sigma(a)$ and hence, for every P_{-i}, we have

$$Q_{ia}(P_i, P_{-i}) = Q_{i\sigma(a)}(P_i^\sigma, P_{-i}^\sigma) = Q_{ib}(P_i', P_{-i}^\sigma).$$

Due to uniform prior and using the above expression,

$$q_{ia}(P_i) = \frac{1}{n!}\sum_{P_{-i}} Q_{ia}(P_i, P_{-i}) = \frac{1}{n!}\sum_{P_{-i}} Q_{ib}(P_i', P_{-i}^\sigma) = \frac{1}{n!}\sum_{P_{-i}} Q_{ib}(P_i', P_{-i}) = q_{ib}(P_i'),$$

where the third equality follows from the fact that $\{P_{-i} : P_{-i} \in \mathcal{P}^{n-1}\} = \{P_{-i}^\sigma : P_{-i} \in \mathcal{P}^{n-1}\}$.

In view of step 1, with some abuse of notation, we write q_{ik} to denote the interim share of the object at rank k in the preference. We call q_i the interim **rank vector** of agent i.

Step 2. Pick an agent i and a preference P_i. We show that $q_{ik} \geq q_{i(k+1)}$ for all $k \in \{1, \ldots, n-1\}$ (i.e., interim shares are non-decreasing with rank). Fix a k and let $P_i(k) = a$ and $P_i(k+1) = b$. Then, consider the preference P_i', which is an (a, b)-swap of P_i. For every P_{-i}, elementary monotonicity implies $Q_{ia}(P_i, P_{-i}) \geq Q_{ia}(P_i', P_{-i})$. Due to uniform prior, $q_{ia}(P_i) \geq q_{ia}(P_i')$. But by Step 1,

$$q_{ik} = q_{ia}(P_i) \geq q_{ia}(P_i') = q_{i(k+1)}.$$

Step 3. We show that Q is OBIC with respect to uniform priors. Suppose agent i has preference P_i. By Steps 1 and 2, he gets interim rank vector $(q_{i1}, \ldots, q_{in})$ by reporting P_i with $q_{ij} \geq q_{ij+1}$ for all $j \in \{1, \ldots, n-1\}$. Suppose she reports $P_i' = P_i^\sigma$, where σ is some permutation of set of objects. By Steps 1 and 2, the interim share vector is a permutation of interim rank vector q_i . Using non-decreasingness of this interim share vector with respect to ranks, we get $q_i(P_i) \succ_{P_i} q_i(P_i')$. Hence, Q is OBIC with respect to uniform prior. ■

Theorem 1 significantly generalizes, an analogous result in Majumdar and Sen (2004), who consider the voting problem and *only* deterministic mechanisms. They arrive at the same conclusion as Theorem 1 in their model. Theorem 1 shows that their result holds even in the *random* assignment problem.

4.1 Probabilistic serial mechanism and U-OBIC

Bogomolnaia and Moulin (2001) define a family of mechanisms, which they call the **simultaneous eating algorithms** (SEA). Though the SEAs are not strategy-proof, they satisfy compelling efficiency and fairness properties, which we discuss in Section 5. We informally introduce the SEAs – for a formal discussion, see Bogomolnaia and Moulin (2001).

Each SEA is defined by a (possibly time-varying) *eating speed* function for each agent. At every preference profile, agents simultaneously start "eating" their favorite objects at a rate equal to their eating speed. Once an object is completely eaten (i.e., the entire share of 1 is consumed), the amount eaten by each agent is the share of that agent of that object. Once an object completely eaten, agents go to their next preferred object and so on.

If the eating speed of each agent is the same, then the simultaneous eating algorithm is *anonymous*. Bogomolnaia and Moulin (2001) call the unique anonymous SEA, the **probabilistic serial** mechanism. [5]

Corollary 1 *Every simultaneous eating algorithm mechanism is U-OBIC.*

Proof: Clearly, the SEAs are neutral. The SEAs also satisfy elementary monotonicity (Cho, 2018; Mennle and Seuken, 2018).[6] Hence, by Theorem 1, we are done. ■

[5] For axiomatic characterization of the PS mechanism, see Bogomolnaia and Heo (2012) and Hashimoto et al. (2014).

[6] See Theorem 3 and the discussions following Theorem 3 in Cho (2018)

5 Locally robust OBIC

While the uniform prior is an important prior in decision theory, it is natural to ask if Theorem 1 extends to other "generic" priors. Though we do not have a full answer to this question, we have been able to answer this question in negative under a natural robustness requirement. Our robustness requirement is *local*. Take any independent prior μ, and let μ' be any independent prior in the ϵ-radius ball around μ (where $\epsilon > 0$), i.e., $||\mu(P) - \mu'(P)|| < \epsilon$ for all $P \in \mathcal{P}$. In this case, we write $\mu' \in B_\epsilon(\mu)$. Our local robustness requirement is the following.

Definition **6** *A mechanism Q is* **locally robust OBIC (LROBIC)** *with respect to an independent prior μ if there exists an $\epsilon > 0$ such that for every independent prior $\mu' \in B_\epsilon(\mu)$, Q is OBIC with respect to μ'.*

It is well known that Bayesian incentive compatibility with respect to all priors lead to strategy-proofness (Ledyard, 1978). Here, we require OBIC with respect to all independent priors in the ϵ-neighborhood of an independent prior. Bhargava et al. (2015) study a version of LROBIC with respect to *uniform* prior but their robustness also allows the mechanism to be OBIC with respect to correlated priors. They show that a large class of voting rules satisfy their notion of OBIC. We show that in the random assignment model, LROBIC with respect to *any* independent prior has a very different implication.

Theorem **2** *A mechanism is LROBIC with respect to an independent prior and satisfies elementary monotonicity if and only if it is strategy-proof.*

Proof: Every strategy-proof mechanism is OBIC with respect to any prior. A strategy-proof mechanism satisfies elementary monotonicity (Mishra, 2016). So, we now focus on the other direction of the proof. Let Q be an LROBIC mechanism with respect to an independent prior μ. Suppose Q satisfies elementary monotonicity. We do the proof in two steps.

Step 1. In this step, we decompose OBIC into three conditions. This decomposition is similar to the decomposition of strategy-proofness in Mennle and Seuken (2018) – there are some minor differences in axioms and we look at interim share vectors whereas they look at ex-post share vectors.

Our decomposition of OBIC uses the following three axioms.

Definition 7 *A mechanism Q satisfies* **interim elementary monotonicity** *if for every $i \in N$ and every P_i, P_i' such that P_i' is an (a,b)-swap of P_i, we have*

$$q_{ib}(P_i') \geq q_{ib}(P_i)$$
$$q_{ia}(P_i') \leq q_{ia}(P_i).$$

Give a preference ordering P_i of agent i and an object $a \in A$, define $U(a, P_i) := \{x \in A : x \; P_i \; a\}$ and $L(a, P_i) := \{x \in A : a \; P_i \; x\}$.

Definition 8 *A mechanism Q satisfies* **interim upper invariance** *if for every $i \in N$ and every P_i, P_i' such that P_i' is an (a,b)-swap of P_i, and for every $x \in U(a, P_i)$, we have*

$$q_{ix}(P_i') = q_{ix}(P_i).$$

Definition 9 *A mechanism Q satisfies* **interim lower invariance** *if for every $i \in N$ and every P_i, P_i' such that P_i' is an (a,b)-swap of P_i, and for every $x \in L(b, P_i)$, we have*

$$q_{ix}(P_i') = q_{ix}(P_i).$$

The following proposition characterizes OBIC using these axioms.

Proposition 1 *A mechanism Q is OBIC with respect to a common independent prior if and only if it satisfies interim elementary monotonicity, interim upper invariance, and interim lower invariance.*

Since the proof of Proposition 1 is similar to the characteration of strategy-proofness in Mennle and Seuken (2018), we postpone its proof to the Supplementary Appendix A.

Step 2. Using Proposition 1, we now complete the proof. Pick an agent $i \in N$ and $P_i, P'_i \in \mathcal{P}$ such that P'_i is an (a, b)-swap of P_i. By Proposition 1, Q satisfies interim upper invariance and interim lower invariance. Hence, we know that for all $c \notin \{a, b\}$, we get

$$\sum_{P_{-i}} \mu(P_{-i}) \Big[Q_{ic}(P_i, P_{-i}) - Q_{ic}(P'_i, P_{-i}) \Big] = 0. \tag{1}$$

Since μ is a probability distribution over $\mathcal{P}$, we can treat it as a vector in $\mathbb{R}^{n!-1}$. Using $\mu(P_{-i}) \equiv \times_{j \neq i} \mu(P_j)$, we note that the LHS of the Equation 1 is a polynomial function of $\{\mu(P)\}_{P \in \mathcal{P}}$. The equation describes the zero set of this polynomial function, which has measure zero (Caron and Traynor, 2005). Hence, given any independent prior μ^* and $\epsilon > 0$, if Equation 1 has to hold for *all* $\mu \in B_\epsilon(\mu^*)$ (which has non-zero measure), then $Q_{ic}(P_i, P_{-i}) = Q_{ic}(P'_i, P_{-i})$ for all $c \notin \{a, b\}$.

Hence, for any LROBIC mechanism Q, for every $i \in N$, for every P_{-i}, if we consider two preferences P_i and P'_i such that P'_i is an (a, b) swap of P_i, we see that for all $c \notin \{a, b\}$, we have $Q_{ic}(P_i, P_{-i}) = Q_{ic}(P'_i, P_{-i})$. Hence, we have shown that Q satisfies *ex-post* versions of our interim lower invariance and interim upper invariance. Mennle and Seuken (2018) refer to these properties as *upper invariance* and *lower invariance* (see also Cho (2018)). They show that upper invariance, lower invariance, and elementary monotonicity are equivalent to strategy-proofness. Since Q satisfies elementary monotonicity, it is strategy-proof. ■

We now explore the compatibility of LOBIC and ordinal efficiency.

Definition 10 *A mechansim Q is* **ordinally efficient** *if at every preference profile* $\mathbf{P}$ *there exists no assignment L such that*

$$L_i \succ_{P_i} Q_i(\mathbf{P}) \qquad \forall\, i \in N,$$

with $L_i \neq Q_i(\mathbf{P})$ for some i.

Bogomolnaia and Moulin (2001) show that every ordinally efficient mechanism is ex-post efficient but the converse is not true if $n \geq 4$. In fact, for $n \geq 4$, strategy-proofness is incompatible with ordinally efficiency along with the following weak fairness criterion.

Definition 11 *A mechanism Q satisfies* **equal treatment of equals** *if at every preference profile* $\mathbf{P}$ *and for every* $i, j \in N$*, we have*

$$\Big[P_i = P_j\Big] \Rightarrow \Big[Q_i(\mathbf{P}) = Q_j(\mathbf{P})\Big]$$

Due to Theorem 2, we can strengthen the impossibility results in Bogomolnaia and Moulin (2001) and Mennle and Seuken (2017) as follows.

Corollary 2 *Suppose $n \geq 4$. Then, there is no locally robust OBIC and ordinally efficient mechanism satisfying equal treatment of equals.*

Proof: By Theorem 2, a locally robust OBIC mechanism satisfies *ex-post* versions of upper invariance and lower invariance. Mennle and Seuken (2017) show that the proof in Bogomolnaia and Moulin (2001) can be adapted by replacing strategy-proofness with ex-post versions of upper invariance and lower invariance. Hence, these two properties are incompatible with ordinal efficiency and equal treatment of equals for $n \geq 4$, and we are done. ■

6 Relation to the literature

There is fairly large literature on random assignment problems. We briefly summarize them and relate them to our results. We also discuss the literature on ordinal Bayesian incentive compatibility.

Ordinal Bayesian incentive compatibility. Strategy-proofness is a demanding requirement in the voting models. Specially, when the domain of preferences is unrestricted, strategy-proofness and unanimity lead to dictatorship (Gibbard, 1973; Satterthwaite, 1975; Gibbard, 1977). Majumdar and Sen (2004) weaken strategy-proofness to OBIC in the voting models but restrict attention to deterministic mechanisms. They show that every deterministic neutral voting mechanism satisfying elementary monotonicity is OBIC with respect to uniform priors. Our Theorem 1 shows that this result generalizes to the random assignment model. Majumdar and Sen (2004) also show that with "generic" priors, every deterministic

OBIC mechanism satisfying unanimity is a dictatorship in the unrestricted domain. Mishra (2016) generalizes this result to some restricted domains of voting (like the single peaked domain). He shows that in the deterministic voting model, elementary monotonicity and OBIC with respect to "generic" prior is equivalent to strategy-proofness in a variety of restricted domains – see also Hong and Kim (2018) for a strengthening of this result. Though these results are similar to our Theorem 2, there are significant differences. First, we consider randomization while these results are only for deterministic mechanisms. Our notion of locally robust OBIC is stronger than OBIC with respect to generic priors used in these papers. Second, ours is a model of private good allocation (random assignment), while these papers deal with the voting model.

Random assignment model. The deterministic assignment model stands on the pillars of two mechanisms: (a) serial dictatorship and (b) top trading cycle. These mechanisms are ex-post efficient and strategy-proof. In fact, these mechanisms satisfy a stronger incentive property called *group strategy-proofness*, which is equivalent to strategy-proofness and *non-bossiness*. For characterizations of deterministic mechanisms satisfying Pareto efficiency and group strategy-proofness, see Pycia and Ünver (2017); Pápai (2000); Svensson (1999).

A natural motivation for studying random mechanisms is fairness. This motivates the study of *random serial dictatorship* or, as they are commonly called, the *random priority* mechanisms. The random priority mechanisms are fair but they fail a natural efficiency criteria called ordinal efficiency. Bogomolnaia and Moulin (2001) introduce a family of mechanisms called simultaneous eating algorithms which generate ordianally efficient random assignments.[7] The probabilistic serial mechanism belongs to this family and it is anonymous. However, it is not strategy-proof. In fact, Bogomolnaia and Moulin (2001) show that there is no ordinally efficient and strategy-proof mechanism satisfying equal treatment of equals when there are at least four agents.[8]

There is a small literature that provides strategic foundations to the probabilistic se-

[7] Katta and Sethuraman (2006) extend the simultaneous eating algorithm to allow for ties in preferences.
[8] With three agent, the random priority mechanism satisfies these properties.

rial (PS) mechanism. Bogomolnaia and Moulin (2001) show that the PS mechanism satisfies *weak*-strategy-proofness. Their notion of weak strategy-proofness requires that the manipulation share vector *cannot* first-order-stochastic-dominate the truth-telling share vector. Bogomolnaia and Moulin (2002) study a problem where agents have an *outside option*. When agents have the same ordinal ranking over objects but the position of outside option in the ranking of objects is the only private information, they show that the PS mechanism is strategy-proof. In fact, they show that the PS mechanism is characterized by strategy-proofness, ordinal efficiency, and equal treatment of equals. This can be viewed as a domain restriction to achieve strategy-proofness of the PS mechanism. Other contributions in this direction include Liu (2019); Liu and Zeng (2019), who identifies domains where the probabilistic serial mechanism is strategy-proof. Che and Kojima (2010) show that the PS mechanism and the random priority mechanism (which is strategy-proof) are asymptotically equivalent. Similarly, Kojima and Manea (2010) show that when sufficiently many copies of an object are present, then the PS mechanism is strategy-proof. Thus, in large economies, the PS mechanism is strategy-proof. Balbuzanov (2016) introduce a notion of strategy-proofness which is stronger than weak strategy-proofness and show that the PS mechanism satisfies it. His notion of strategy-proofness is based on the "convex" domination of lotteries, and hence, called *convex strategy-proofness*. Mennle and Seuken (2018) define a notion called *partial strategy-proofness*, which is weaker than strategy-proofness and show that the PS mechanism satisfies it. They show that strategy-proofness is equivalent to upper invariant, lower invariant and elementary monotonicity (they call it swap monotonicity). Their notion of partial strategy-proofness is equivalent to upper invariance and elementary monotonicity, and hence, it is weaker than strategy-proofness. The main difference between these weakenings of strategy-proofness and ours is that OBIC is a *prior-based* notion of incentive compatibility. It is the natural analogue of Bayesian incentive compatibility in an ordinal environment. Ehlers and Massó (2007) study OBIC in a *two-sided matching* problem. Their main focus is on *stable* mechanisms. They characterize the beliefs for which a stable mechanism is OBIC.

References

Balbuzanov, I. (2016): "Convex strategyproofness with an application to the probabilistic serial mechanism," *Social Choice and Welfare*, 46, 511–520.

Bhargava, M., D. Majumdar, and A. Sen (2015): "Incentive-compatible voting rules with positively correlated beliefs," *Theoretical Economics*, 10, 867–885.

Bogomolnaia, A. and E. J. Heo (2012): "Probabilistic assignment of objects: Characterizing the serial rule," *Journal of Economic Theory*, 147, 2072–2082.

Bogomolnaia, A. and H. Moulin (2001): "A new solution to the random assignment problem," *Journal of Economic theory*, 100, 295–328.

——— (2002): "A simple random assignment problem with a unique solution," *Economic Theory*, 19, 623–636.

Caron, R. and T. Traynor (2005): "The zero set of a polynomial," University of Windsor, Technical Report.

Carroll, G. (2012): "When are local incentive constraints sufficient?" *Econometrica*, 80, 661–686.

Che, Y.-K. and F. Kojima (2010): "Asymptotic equivalence of probabilistic serial and random priority mechanisms," *Econometrica*, 78, 1625–1672.

Cho, W. J. (2016): "Incentive properties for ordinal mechanisms," *Games and Economic Behavior*, 95, 168–177.

——— (2018): "Probabilistic assignment: an extension approach," *Social Choice and Welfare*, 51, 137–162.

d'Aspremont, C. and B. Peleg (1988): "Ordinal Bayesian incentive compatible representations of committees," *Social Choice and Welfare*, 5, 261–279.

Ehlers, L. and J. Massó (2007): "Incomplete information and singleton cores in matching markets," *Journal of Economic Theory*, 136, 587–600.

Gibbard, A. (1973): "Manipulation of voting schemes: a general result," *Econometrica*, 41, 587–601.

——— (1977): "Manipulation of schemes that mix voting with chance," *Econometrica*, 45, 665–681.

Hashimoto, T., D. Hirata, O. Kesten, M. Kurino, and M. U. Ünver (2014): "Two axiomatic approaches to the probabilistic serial mechanism," *Theoretical Economics*, 9, 253–277.

Hong, M. and S. Kim (2018): "Unanimity and local incentive compatibility," Working paper, Yonsei University.

Katta, A.-K. and J. Sethuraman (2006): "A solution to the random assignment problem on the full preference domain," *Journal of Economic theory*, 131, 231–250.

Kojima, F. and M. Manea (2010): "Incentives in the probabilistic serial mechanism," *Journal of Economic Theory*, 145, 106–123.

Ledyard, J. O. (1978): "Incentive compatibility and incomplete information," *Journal of Economic Theory*, 18, 171 – 189.

Liu, P. (2019): "Random assignments on sequentially dichotomous domains," *Games and Economic Behavior*.

Liu, P. and H. Zeng (2019): "Random assignments on preference domains with a tier structure," *Journal of Mathematical Economics*, 84, 176–194.

Majumdar, D. and A. Sen (2004): "Ordinally Bayesian incentive compatible voting rules," *Econometrica*, 72, 523–540.

Mennle, T. and S. Seuken (2017): "Two New Impossibility Results for the Random Assignment Problem," Working paper, University of Zurich, https://www.ifi.uzh.ch/ce/publications/IPS.pdf.

——— (2018): "Partial Strategyproofness: Relaxing Strategyproofness for the Random Assignment Problem," Working paper, University of Zurich, https://www.ifi.uzh.ch/ce/publications/PSP.pdf.

Mishra, D. (2016): "Ordinal Bayesian incentive compatibility in restricted domains," *Journal of Economic Theory*, 163, 925–954.

Pápai, S. (2000): "Strategyproof assignment by hierarchical exchange," *Econometrica*, 68, 1403–1433.

Pycia, M. and M. U. Ünver (2017): "Incentive compatible allocation and exchange of discrete resources," *Theoretical Economics*, 12, 287–329.

Satterthwaite, M. A. (1975): "Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions," *Journal of economic theory*, 10, 187–217.

Svensson, L.-G. (1999): "Strategy-proof allocation of indivisible goods," *Social Choice and Welfare*, 16, 557–567.

Wilson, R. (1987): "Game theoretic approaches to trading processess," in *Advances in Economic Theory: Fifth World Congress of Econometric Society*, ed. by T. Bewley, Cambridge, UK: Cambridge University Press, 201–213.

A Supplementary Appendix

Proof of Proposition 1

Proof: The proof uses *locally* OBIC.

Definition 12 *A mechanism Q is* **locally ordinally Bayesian incentive compatible (locally OBIC)** *with respect to an independent common prior μ if for every $i \in N$ and every $P_i, P_i' \in \mathcal{P}$ such that P_i' is an (a,b)-swap of P_i, we have*

$$q_i(P_i) \succ_{P_i} q_i(P_i').$$

Locally OBIC only considers a subset of incentive constraints. Using Carroll (2012); Cho (2016), we conclude that locally OBIC implies OBIC in our model. Hence, it is enough to show that interim elementary monotonicity, interim upper invariance, and interim lower invariance are equivalent to locally OBIC.

Suppose Q is locally OBIC. Fix agent $i \in N$ and consider P_i and P_i' such that P_i' is an (a,b)-swap of P_i.

interim upper invariance. If $U(a,P_i) = \emptyset$, then there is nothing to show. Else, let $U(a,P_i) = \{a_1, \ldots, a_k\}$ with $a_1 \; P_i \; a_2 \; P_i \; \ldots \; P_i \; a_k$. Note that since P_i' is an (a,b)-swap of P_i, we have $a_1 \; P_i' \; a_2 \; P_i' \; \ldots \; P_i' \; a_k$. Using OBIC, we get

$$q_{ia_1}(P_i) \geq q_{ia_1}(P_i') \;\text{ and }\; q_{ia_1}(P_i') \geq q_{ia_1}(P_i). \tag{2}$$

Hence, $q_{ia_1}(P_i) = q_{ia_1}(P_i')$. Now, suppose $q_{ia_j}(P_i) = q_{ia_j}(P_i')$ for all $j \in \{1, \ldots, \ell - 1\}$, where $\ell \leq k$. Then, using OBIC again, we get

$$\sum_{h=1}^{\ell} q_{ia_h}(P_i) \geq \sum_{h=1}^{\ell} q_{ia_h}(P_i') \;\text{ and }\; \sum_{h=1}^{\ell} q_{ia_h}(P_i') \geq \sum_{h=1}^{\ell} q_{ia_h}(P_i).$$

Hence, we get $\sum_{h=1}^{\ell} q_{ia_h}(P_i) = \sum_{h=1}^{\ell} q_{ia_h}(P_i')$. But $q_{ia_j}(P_i) = q_{ia_j}(P_i')$ for all $j \in \{1, \ldots, \ell-1\}$ implies that $q_{ia_\ell}(P_i) = q_{ia_\ell}(P_i')$. By induction, we conclude that Q satisfies interim upper invariance.

interim elementary monotonicity. Interim elementary monotonicity follows by considering OBIC from P_i' to P_i:

$$\sum_{x \in U(a,P_i')} q_{ix}(P_i') + q_{ib}(P_i') \geq \sum_{x \in U(a,P_i')} q_{ix}(P_i) + q_{ib}(P_i).$$

But interim upper invariance implies that $q_{ib}(P_i') \geq q_{ib}(P_i)$ as desired. A similar proof shows $q_{ia}(P_i') \leq q_{ia}(P_i)$.

INTERIM LOWER INVARIANCE. By OBIC,

$$\sum_{x \in U(a,P_i)} q_{ix}(P_i) + q_{ia}(P_i) + q_{ib}(P_i) \geq \sum_{x \in U(a,P_i)} q_{ix}(P_i') + q_{ia}(P_i') + q_{ib}(P_i')$$

Using interim upper invariance of Q, we get

$$q_{ia}(P_i) + q_{ib}(P_i) \geq q_{ia}(P_i') + q_{ib}(P_i').$$

The OBIC constraint from P_i' to P_i is similar and gives the other inequality and we conclude

$$q_{ia}(P_i) + q_{ib}(P_i) = q_{ia}(P_i') + q_{ib}(P_i').$$

With this, we can repeat the argument of interim upper invariance to get interim lower invariance.

The converse statement that any Q satisfying interim elementary monotonicity, interim upper invariance, and interim lower invariance is locally OBIC is straightforward, and left to the reader. ■