UTMD Working Paper
The University of Tokyo
Market Design Center

UTMD-026

A Planner-Optimal Matching Mechanism and Its Incentive Compatibility in a Restricted Domain

Shunya Noda
University of British Columbia

First Draft: October 7, 2019
This Version: October 12, 2019

A Planner-Optimal Matching Mechanism and Its Incentive Compatibility in a Restricted Domain*

Shunya Noda†

First Draft: October 7, 2019 This Version: October 12, 2019

Abstract

In many random assignment problems, the central planner has their own policy objective, such as matching size and minimum quota fulfillment. A number of practically important policy objectives are not aligned with agents' preferences and known to be incompatible with strategy-proofness. This paper proves that such policy objectives can be achieved by mechanisms that satisfy Bayesian incentive compatibility in a restricted domain of von Neumann Morgenstern utilities. We prove that if a mechanism satisfies the three axioms of swap monotonicity, lower invariance, and interior upper invariance, then the mechanism satisfies Bayesian incentive compatibility in an *inverse-bounded-indifference* (IBI) domain. We apply this axiomatic characterization to analyze the incentive property of a novel mechanism, the *constrained serial dictatorship mechanism* (CRSD). CRSD is designed to generate an individually rational assignment that maximizes the central planner's policy objective function. As CRSD satisfies these axioms, CRSD is Bayesian incentive compatible in an IBI domain.

JEL Codes: C61, C78, D47, D82

Keywords: Random Assignment, Ordinal Mechanisms, Strategy-Proofness, Maximum Matching, Minimum Quota

*First Presentation: August 9th, 2017. We are grateful to Itai Ashlagi, Gabriel Carroll, Atsushi Iwasaki, Atsushi Kajii, Yuichiro Kamada, Fuhito Kojima, Simon Loertscher, Takanori Maehara, Delong Meng, Alvin Roth, Sven Seuken, Ashutosh Thakur, and all the seminar participants in Stanford University, the RIKEN Center for Advanced Intelligence Project, and the Summer Economic Workshop on Theory 2017 in Otaru.

†Vancouver School of Economics, University of British Columbia, 6000 Iona Dr, Vancouver, BC V6T 1L4 Canada. E-mail: shunya.noda@gmail.com Noda has been supported by the Funai Overseas Scholarship and the E. K. Potter fellowship.

1 Introduction

In an assignment problem (also known as a one-sided matching problem), the central planner allocates a set of indivisible objects to a set of agents, who can each consume at most one object. The central planner cannot directly observe agents' preferences. Furthermore, some agents may declare that some objects are unacceptable to them, and the central planner is prohibited from allocating an unacceptable object to the agent (i.e., the assignment must be *individually rational*). To respect agents' preferences and generate a desirable assignment, the central planner must use an allocation mechanism that incentivizes agents to report their preferences truthfully. We focus on a situation where monetary transfers are precluded, and therefore the mechanism must induce truthtelling through the variation in the probability of assigning each object.

In many applications, the central planner not only must respect agents' preferences but also has her own policy objective. In refugee resettlement and daycare assignment problems, the government wants to maximize the *matching size*, defined as the total number of agents who are assigned to some objects, because it is costly to keep many people in the waitlist (Andersson and Ehlers 2016, Delacrétaz et al. 2016, Kamada and Kojima 2018, Noda 2018a). Institutions responsible for refugee resettlement are also concerned with refugees' total *job employment rate*, predicted from the characteristics of refugee families, e.g., the language they speak (Bansak et al. 2018). In doctor-hospital and student-laboratory matching problems, the institution often wants to set *minimum quota* constraints to make sure that doctors are allocated to rural areas and students are allocated to all laboratories (Goto et al. 2014, Fragiadakis et al. 2016, Fragiadakis and Troyan 2017). Maximum and minimum quotas can also be used for maintaining a diversity of student types in school choice problems (Ehlers et al. 2014, Tomoeda 2018).

These objectives are not always aligned with agents' preferences; thus, agents may want to deviate from the assignment that the central planner prefers. Accordingly, if no additional condition is imposed, it is generally impossible (i) to achieve a central-planner-optimal assignment, and (ii) to achieve *strategy-proofness*, i.e., to make truthful reporting of preference order a weakly dominant strategy, simultaneously. However, such mechanisms may satisfy *Bayesian incentive compatibility* (BIC) (i.e., truthtelling may maximize each agent's interim expected utility) in a restricted domain of von Neumann Morgenstern (vNM) utility functions underlying preference orders.

The contribution of this paper is twofold. First, we extend the analysis of Mennle and Seuken (2018) to obtain a new characterization of mechanisms satisfying BIC in two new domain classes of vNM utility functions. If an *interim mechanism*, defined as a mapping from an agent's preference report to his interim probability of obtaining each object, satisfies *swap monotonicity* and *lower invariance*, then the interim mechanism satisfies BIC in an *inverse-uniformly-relatively-bounded-indifference* (IURBI) domain. If an interim mechanism additionally satisfies *interior upper invariance*, then the interim mechanism satisfies BIC in an *inverse-bounded-indifference* (IBI) domain.

Second, we propose a novel mechanism, which is designed for generating a planner-optimal assignment. We establish the *constrained random serial dictatorship mechanism* (CRSD), which maximizes the central planner's objective function. Then, we analyze its incentive property by applying the axiomatic characterization. CRSD is parameterized by the central planner's objective function, which is used for evaluating assignments. For example, if the central planner wants to achieve a maximum matching, she should take the matching size as her objective. If she wants to satisfy minimum quota constraints, she should penalize assignments that do not satisfy minimum quota by setting a low value for such assignments. CRSD first identifies the set of all individually rational assignments from the reported preference profile and then computes the set of individually rational assignments that maximizes the objective function. After that, just like the standard *random serial dictatorship mechanism* (RSD) (Abdulkadirolu and Sonmez 1998), the mechanism determines the priority order uniformly at random, and each agent chooses his favorite object from his choice set sequentially. CRSD imposes an additional requirement on each agent's choice set:

to make sure that the returned assignment is one of the maximizers of the objective function, each agent is prohibited from choosing an object that is not consistent with any maximizer (fixing earlier movers' choices).

This paper proves that for all objective functions, CRSD satisfies lower invariance and interior upper invariance. Furthermore, for certain objective functions, CRSD also satisfies swap monotonicity. Accordingly, CRSD satisfies BIC in an IBI domain, implying that if agents do not want to tolerate the risk of becoming unassigned in the hope of obtaining their favorite objects, then truthful reporting is induced.

2 Related Literature

In matching theory literature, strategy-proofness is known to be incompatible with various important policy objectives, including ordinal efficiency, matching size, and minimum quota. Some previous studies have attempted to indirectly achieve such difficult goals by imposing relatively tractable constraints on assignments. For example, Kamada and Kojima (2015) set the maximum quota to decrease the number of doctors matched to urban areas so as to allocate more doctors to rural areas. Some works have designed mechanisms for environments in which preferences are partially observable, e.g., the set of objects acceptable to each agent is public information (Goto et al. 2014, Fragiadakis et al. 2016, Fragiadakis and Troyan 2017, Ashlagi et al. 2019). Some papers have evaluated the performance of strategy-proof mechanisms for achieving a given policy objective. For example, Krysta et al. (2014), Afacan and Dur (2018), and Noda (2018a,b) evaluate the matching size achieved by strategy-proof mechanisms. Some previous studies have investigated special environments where non-strategy-proof mechanisms become strategy-proof. For example, the deferred acceptance mechanism (DA) (Gale and Shapley 1962) and the probabilistic serial mechanism (PS) (Bogomolnaia and Moulin 2001) are not strategy-proof in general but becomes strategy-proof in large markets (DA: Kojima and Pathak 2009, PS: Kojima and Manea 2010). Other studies consider the domain of preferences with which strategy-proofness is compatible with difficult goals. Bogomolnaia and Moulin (2004) consider a two-sided matching problems where agents have dichotomous preferences and study strategy-proof mechanisms in such an environment. Balbuzanov (2016) considers convex strategy-proofness, which only requires the existence of a vNM utility function with which truthtelling is a weakly dominant strategy, and proves that PS is convex strategy-proof. Mennle and Seuken (2018) provide an axiomatic characterization of the set of vNM utility functions with which truthtelling becomes a dominant strategy.

Our axiomatic characterization is an extension of Mennle and Seuken (2018). They prove that strategy-proofness in the unrestricted domain can be decomposed into three axioms: swap monotonicity, upper invariance, and lower invariance. They also demonstrate that, if a mechanism satisfies swap monotonicity and upper invariance, a mechanism is strategy-proof in a domain called the *uniformly-relatively-bounded-indifference* domain. Although various mechanisms satisfy swap monotonicity, upper invariance, and certain other desirable properties (e.g., PS is swap monotonic, upper invariant, and ordinally efficient), no interesting mechanism satisfying both swap monotonicity and lower invariance has been documented in the literature.[1] However, the new mechanism proposed in this paper (CRSD) satisfies swap monotonicity and lower invariance, though it does not satisfy upper invariance. Given this, we establish a new characterization of incentive-compatible mechanisms in a restricted domain based on swap monotonicity and lower invariance.

Our CRSD bridges the literature on incentive-compatible mechanism design and optimization algorithms. In the literature on operations research and computer science, a variety of algorithms for solving assignment problems have been established. If preferences were public information, we could run such algorithms to maximize the central planner's objective function. Many previous studies indicate the usefulness of such algorithms in social

[1] Mennle and Seuken (2017) also prove that there is no lower-invariant counterpart of the probabilistic serial mechanism in the sense that no mechanism can be swap monotonic, lower invariant, ordinally efficient, anonymous, neutral, and non-bossy.

science problems, but the extent to which incentives for truthful reporting are compatible with optimization is not yet known. For example, Delacrétaz et al. (2016) posit that, if refugee families' preferences were known, the total number of resettlements could be maximized by the algorithm developed by Song et al. (2008). However, how the use of Song et al.'s algorithm breaks down strategy-proofness is ambiguous. Bansak et al. (2018) show that refugees' total job employment rate after resettlement can be improved by using a data-driven algorithm, though they disregard refugees' preferences. Afacan et al. (2018) propose the *efficient assignment maximizing mechanism* that is not strategy-proof but always returns a maximum matching with respect to the reported preference profile. However, given any vNM utility functions, the equilibrium assignment of their mechanism is identical to the serial dictatorship mechanism. By contrast, CRSD satisfies BIC in a restricted domain of vNM utility functions. In addition, CRSD can integrate not only the matching size but also general policy objective functions.

3 Model

There is a finite set of *agents* N and a finite set of *objects* M. Each agent consumes at most one object from M, and each object has a unit capacity. The outside option is denoted by $\perp$, and the agent is assigned to $\perp$ if he is not assigned to any $k \in M$. We assume that the outside option has an unlimited capacity.

Each agent i has a strict *preference order* P_i over objects and outside option $M \cup \{\perp\}$, where $P_i : a \succ b$ indicates that agent i strictly prefers object a to b. We denote the set of all possible preference orders by $\mathcal{P}$. A *preference profile* $P = (P_i)_{i \in N} \in \mathcal{P}^N$ is a profile of preference orders of all agents, and $P_{-i} \in \mathcal{P}^{N \setminus \{i\}}$ is a profile of preference orders of all agents except i.

The *neighborhood of a preference order* P_i, denoted by $\mathcal{N}(P_i)$, is the set of all preference orders that differ from P_i by a swap of two consecutively ranked objects. The *upper contour* *set of* k *at* P_i is the set of objects that i strictly prefers to k, denoted $U(k, P_i)$. Conversely, the *lower contour set of* k *at* P_i is the set of objects that i strictly disprefers to k, denoted by $L(k, P_i)$.

Agent i's (von Neumann Morgenstern, vNM) *utility function* is denoted by $u_i : M \cup \{\perp\} \to \mathbb{R}$. To make every agent non-redundant, we assume that every agent has at least one object that is strictly preferred to the outside option. Without loss of generality, we always normalize $u_i(\perp) = 0$ and $\max_{k \in M} u_i(k) = 1$. A utility function u_i is *consistent with* P_i if $u_i(a) > u_i(b)$ whenever $P_i : a \succ b$. We denote the set of all possible utility functions by $\mathcal{U}$. We refer to a subset of $\mathcal{U}$ as the *utility subdomain*. We denote the set of all utility functions that are consistent with preference order P_i by $\mathcal{U}(P_i)$. Conversely, we denote the preference order consistent with the utility function u_i by $p(u_i)$.

Let F_i be the marginal distribution of agent i's utility. We assume that agents' utilities are independently distributed: the joint probability distribution of the utility profile can be written as a product of the marginal distributions of agents' utility functions.

A *probabilistic assignment* (or matching) is represented by the matrix, $x = (x_{i,k})_{i \in N, k \in M \cup \{\perp\}}$ such that (i) $\sum_{k \in M \cup \{\perp\}} x_{i,k} = 1$ for every $i \in N$, (ii) $\sum_{i \in N} x_{i,k} \leq 1$ for every $k \in M$, and $x_{i,k} \in [0, 1]$ for every $(i, k) \in N \times M \cup \{\perp\}$. The value $x_{i,k}$ represents the probability that agent i obtains object k. The ith row $x_i = (x_{i,k})_{k \in M \cup \{\perp\}}$ of x is called the probabilistic assignment of i. When agent i's assignment is x_i, his expected utility is $\mathbb{E}_{x_i}[u_i] = \sum_{k \in M} x_{i,k} u_i(k)$. An assignment x is called *deterministic* if $x_{i,k} \in \{0, 1\}$ for all $i \in N$, $k \in M \cup \{\perp\}$. Note that by the Birkhoff-von Neumann theorem (Birkhoff 1946, Von Neumann 1953), every probabilistic assignment can be decomposed into a convex combination of deterministic assignments. We denote X and $\Delta(X)$ the set of all deterministic and probabilistic assignments, respectively.

We consider ordinal assignment mechanisms, which take preference orders as their inputs. A *mechanism* is a mapping $\phi : \mathcal{P}^N \to \Delta(X)$ that selects a probabilistic assignment based on a preference profile. If the returned assignment is always deterministic, we refer to it as a deterministic mechanism. Throughout the paper, we focus on *individually rational* mechanisms that assign an object to an agent only if the agent prefers the object to the outside option; i.e., $P_i : \perp \succ a$ implies $\phi_{i,a}(P_i, P_{-i}) = 0$ for all $P_{-i} \in \mathcal{P}^{N\setminus\{i\}}$.

4 Interim Reporting Problem

4.1 Setup

We consider each agent's *interim reporting problem.* In an interim stage, each agent i knows his own preference P_i but is not aware of other agents' preferences P_{-i}. The interim probability that agent i is assigned to object k when he reports preference P'_i is given by

$$\Phi_{i,k}(P'_i) \coloneqq \int_{u_{-i}} \phi_{i,k}\left(P'_i, (p(u_j))_{j\in N\setminus\{i\}}\right) dF_{-i}. \tag{1}$$

Note that the above formula is independent of agent i's true type (u_i) because we assume independent types. We refer to $\Phi_i = (\Phi_{i,k})_{k\in M\cup\{\perp\}}$ as *agent i's interim mechanism.*

The (ordinal) *Bayesian incentive compatibility* requires that a truthful reporting of preference provides the largest expected payoff to the agent. In a standard definition, the optimality of truthtelling is required for all possible utility functions $\mathcal{U}$. Here, we explicitly define the utility subdomain with which an agent has an incentive for truthtelling.

Definition 1 (Bayesian Incentive Compatibility on the Utility Subdomain)**.** Agent i's interim mechanism Φ_i satisfies *Bayesian incentive compatibility (BIC) on utility subdomain U* if, for all utility functions $u_i \in U$ for all misreports $P'_i \in \mathcal{P}$, we have $\mathbb{E}_{\Phi_i(p(u_i))}[u_i] \geq \mathbb{E}_{\Phi_i(P'_i)}[u_i]$.

Note that if we take $U = \mathcal{U}$, our definition of BIC coincides with the standard one.

Remark 1. Alternatively, we can require the mechanism to be *strategy-proof* (i.e., incentive compatible in weakly dominant strategies) by assuming that agent i knows other agents' reporting P'_{-i}. To obtain a corresponding condition, we should replace Φ_i with $\phi_i(\cdot, P_{-i})$ and impose each condition for all $P_{-i} \in \mathcal{P}^{N\setminus\{i\}}$. Although the property required on mechanism ϕ

becomes stronger, the characterizations of strategy-proof mechanisms on a restricted domain can be obtained in a similar manner.

4.2 Preliminary Results

Here, we review three axioms proposed by Mennle and Seuken (2018).

Axiom 1 (Swap Monotonicity)**.** Agent i's interim mechanism Φ_i is *swap monotonic* if, for all preferences $P_i \in \mathcal{P}$, all misreports $P_i' \in \mathcal{N}(P_i)$, and objects or outside options $a, b \in M \cup \{\perp\}$ with $P_i : a \succ b$ but $P_i' : b \succ a$, one of the following holds: either (i) $\Phi_i(P_i) = \Phi_i(P_i')$ or (ii) $\Phi_{i,a}(P_i) > \Phi_{i,a}(P_i')$, and $\Phi_{i,b}(P_i) < \Phi_{i,b}(P_i')$.

Swap monotonicity requires that if an agent swaps his preference order over two objects (a and b), either one of the following outcomes must hold: (i) it does not affect the resultant probabilistic assignment at all, or (ii) it *strictly* increases the probability to get the object brought forward and *strictly* decreases the probability to get the object carried down.

Axiom 2 (Upper Invariance)**.** Agent i's interim mechanism Φ_i is *upper invariant* if for all preferences $P_i \in \mathcal{P}$, all misreports $P_i' \in \mathcal{N}(P_i)$, and objects or outside options $a, b \in M \cup \{\perp\}$ with $P_i : a \succ b$ but $P_i' : b \succ a$, we have that $\Phi_{i,k}(P_i) = \Phi_{i,k}(P_i')$ for all $k \in U(a, P_i)$.

Axiom 3 (Lower Invariance)**.** Agent i's interim mechanism Φ_i is *lower invariant* if for all preferences $P_i \in \mathcal{P}$, all misreports $P_i' \in \mathcal{N}(P_i)$, and objects or outside options $a, b \in M \cup \{\perp\}$ with $P_i : a \succ b$ but $P_i' : b \succ a$, we have that $\Phi_{i,k}(P_i) = \Phi_{i,k}(P_i')$ for all $k \in L(b, P_i)$.

Upper invariance[2] requires that an agent cannot manipulate their probability of obtaining a more-preferred object by changing the order of less-preferred objects. Conversely, lower invariance requires that an agent cannot manipulate their probability of obtaining less-preferred object by changing the order of more-preferred objects.

[2] Upper invariance is originally introduced by Hashimoto et al. (2014), and it is called *weak invariance* in their paper.

An interim mechanism satisfies BIC in the whole utility space $\mathcal{U}$ if and only if it satisfies swap monotonicity, upper invariance, and lower invariance. Mennle and Seuken (2018) also show that when the interim mechanism satisfies swap monotonicity and upper invariance, if an agent's utility function is in a uniformly-relatively-bounded-indifference (IRBI) subdomain, then truthful reporting is optimal for the agent.

Definition 2 (URBI)**.** A utility function u_i satisfies *uniformly relatively bounded indifference (URBI) with respect to* $r \in [0, 1]$ if, for all objects $a, b \in M \cup \{\perp\}$ with $u_i(a) > u_i(b)$, we have

$$r \cdot (u_i(a) - u_i(\perp)) \geq u_i(b) - u_i(\perp). \tag{2}$$

Since we normalize $u_i(\perp) = 0$, the inequality (2) simplifies to

$$ru_i(a) \geq u_i(b). \tag{3}$$

We denote the set of all utility functions satisfying URBI with respect to r by $URBI(r)$.

Theorem 1 (Theorem 1 and 2 of Mennle and Seuken, 2018)**.**

1. *Agent* i*'s interim mechanism* Φ_i *satisfies BIC in* $\mathcal{U}$ *if and only if* Φ_i *satisfies swap monotonicity, upper invariance, and lower invariance.*

2. *Agent* i*'s interim mechanism* Φ_i *satisfies BIC in* $URBI(r)$ *for some* $r > 0$ *if and only if* Φ_i *satisfies swap monotonicity and upper invariance.*

The utility subdomain $URBI(r)$ is increasing in r in the sense that $r' > r$ implies $URBI(r') \supset URBI(r)$. When $r = 1$, the utility subdomain coincides with the whole domain: $URBI(1) = \mathcal{U}$. Conversely, $URBI(0)$ only contains utility functions that accept only one object ($u_i(k) = 1$ for some k, and $u_i(l) < 0$ for all $l \in M \setminus \{k\}$). An example of a utility function satisfying $u_i \in URBI(0.5)$ is depicted in Figure 1a.

4.3 Inverse URBI Domain

We study a class of mechanisms that satisfy swap monotonicity and *lower* invariance, rather than upper invariance. For such mechanisms, we characterize a new utility subdomain in which an interim mechanism satisfies BIC.

Definition 3 (IURBI)**.** A utility function u_i satisfies *inverse uniformly relatively bounded indifference (IURBI) with respect to* $r \in [0,1]$ if, for all objects $a, b \in M$ with $u_i(a) > u_i(b)$, we have

$$r \cdot \left(\max_{k \in M} u_i(k) - u_i(b) \right) \geq \max_{k \in M} u_i(k) - u_i(a). \tag{4}$$

Since we normalize $\max_{k \in M} u_i(k) = 1$, (4) simplifies to

$$r(1 - u_i(b)) \geq 1 - u_i(a). \tag{5}$$

We denote the set of all utility functions satisfying IURBI with respect to r by $IURBI(r)$.

Both URBI and IURBI domains require that when $P_i : a \succ b$, the underlying utility function u_i must sufficiently differentiate these two objects. However, their "reference points" are different. $URBI(r)$ requires the following inequality condition: whenever $a \succ b$,

$$r \cdot u_i(a) + (1 - r) \cdot 0 \geq u_i(b). \tag{6}$$

The URBI domain measures the intensity of preferences attached to a and b by using the agent's preference for the outside option. A utility function is more likely to belong to the URBI domain if he has a more extreme attachment to a more-preferred object. Indeed, a utility function that accepts only one (favorite) object belongs to $URBI(r)$ for all r.

By contrast, the IURBI domain compares utility using the favorite object as a reference point. Formally, (5) can be rearranged to produce the following inequality:

$$u_i(a) \geq (1 - r) \cdot 1 + r \cdot u_i(b). \tag{7}$$

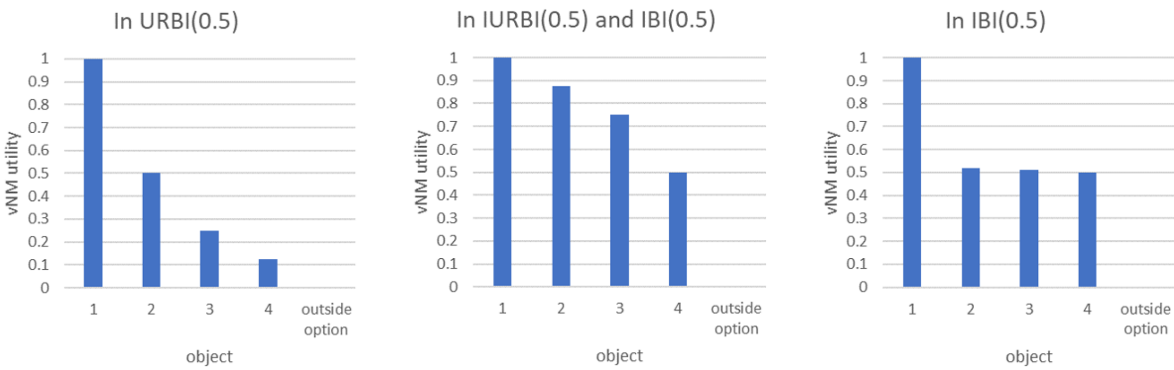

(a) $u_i \in URBI(0.5)$ (b) $u_i \in IRUBI(0.5), IBI(0.5)$ (c) $u_i \in IBI(0.5)$

Figure 1: Utility functions in each subdomain. All utility functions satisfy $p(u_i) : 1 \succ 2 \succ 3 \succ 4 \succ \perp$. All utility functions are normalized to satisfy $u_{i,1} = 1$ and $u_{i,\perp} = 0$. In the following formulas, $k = 5$ denotes $\perp$.
Figure 1a: $u_{i,k} \coloneqq 0.5 \cdot u_{i,k-1}$ for $k = 2, 3, 4$. $u_i \in URBI(0.5)$.
Figure 1b: $u_{i,k} \coloneqq 1 - 0.5 \cdot (1 - u_{i,k+1})$ for $k = 2, 3, 4$. $u_i \in IURBI(0.5) \subset IBI(0.5)$.
Figure 1c: $u_{i,4} \coloneqq 1 - 0.5 \cdot (1 - u_{i,\perp}) = 0.5$, $u_{i,k} = u_{i,k+1} + 0.01$ for $k = 2, 3$. $u_i \in IBI(0.5)$.

This inequality requires that the agent still prefers to take object a even if he would have a small chance (with probability $1 - r$) to get his favorite object. A utility function is more likely to belong to the IURBI domain if he has *less* extreme attachment to his favorite object. If an agent is (nearly) indifferent with regards to all acceptable objects, then the utility function is in $IURBI(r)$ for (nearly) all r.

Parallel to the URBI domain, the IURBI domain is increasing in r in the sense that $r' > r$ implies $IURBI(r') \supset IURBI(r)$. Furthermore, these definitions coincide when $r = 1$: $IURBI(1) = \mathcal{U}$. In this sense, the value of r measures the "degree of BIC." An example of a utility function satisfying $u_i \in IURBI(0.5)$ is depicted in Figure 1b.

The following theorem states that an interim mechanism is BIC in $IURBI(r)$ with $r > 0$ if and only if it satisfies swap monotonicity and lower invariance.

Theorem 2. *Agent i's interim mechanism Φ_i satisfies BIC in $IURBI(r)$ for some $r > 0$ if and only if Φ_i satisfies swap monotonicity and lower invariance.*

Proofs are provided in the appendix.

The intuition is as follows. Let a and b be two objects (or an object and the outside option), where agent i prefers a to b. Suppose that an interim mechanism satisfies swap monotonicity and lower invariance. By swap monotonicity, if the agent swaps the preference order of a and b and his assignment is changed, then he will incur a strict loss by decreasing the probability of obtaining a and increasing the probability of obtaining b. For a misreporting to be profitable, the agent must obtain some gain from a change in the assignment probability of other objects. By lower invariance, he cannot benefit by decreasing the probability of obtaining objects less preferred to b. Hence, all the (potential) deviation gains are from the increment of the probability of obtaining objects more preferred to a. In the best case, he can increase the probability of obtaining his favorite object; thus, to identify whether the swap is profitable, we compare the value of $\max_{k \in M} u_i(k)$, $u_i(a)$, and $u_i(b)$. The swap is profitable only if the latter effect is larger than the former effect, and the IURBI domain rules out such preferences. Accordingly, when the agent's utility function belongs to the IURBI domain, it is optimal to report his preference truthfully.

4.4 Inverse Bounded Indifference Domain

Some mechanisms do not satisfy upper invariance because they treat the outside option differently from objects. Even in such a case, a swap of two *objects* ($a, b \in M$, i.e., $a, b \neq \perp$) may not change the probability of assigning more-preferred objects.

Axiom 4 (Interior Upper Invariance)**.** Agent i's interim mechanism Φ_i is *interior upper invariant* if for all preferences $P_i \in \mathcal{P}$, all misreports $P'_i \in \mathcal{N}(P_i)$, and objects $a, b \in M$ with $P_i : a \succ b$ but $P'_i : b \succ a$, we have that $\Phi_{i,k}(P_i) = \Phi_{i,k}(P'_i)$ for all $k \in U(a, P_i)$.

Likewise, we say that swap monotonicity or lower invariance is satisfied *in the interior* if the respective condition is satisfied whenever the two objects a, b are chosen from M. Conversely, we say that swap monotonicity or lower invariance is satisfied *on the boundary* if the respective condition is satisfied if one of a, b is an object (belongs to M) and another is the outside option ($\perp$).

Interior upper invariance is similar to upper invariance but imposes no requirement on the interim mechanism when the set of acceptable objects is changed; i.e., whenever the preference order of $\perp$ is changed. Clearly, interior upper invariance is implied by upper invariance.

If an interim mechanism satisfies swap monotonicity, lower invariance, and interior upper invariance, the agent has no incentive to manipulate the preference order over objects that are declared to be acceptable.

Theorem 3. *Suppose that agent i's interim mechanism Φ_i satisfies swap monotonicity, lower invariance, and interior upper invariance. Consider any preference $P_i \in \mathcal{P}$ and any misreport $P_i' \in \mathcal{P}$. Construct another report P_i'' by the following rule: (i) $P_i'' : \perp \succ k$ whenever $P_i' : \perp \succ k$, and (ii) for all $a, b \in M$ such that $P_i' : a \succ \perp$ and $P_i' : b \succ \perp$, we have $P_i'' : a \succ b$ if and only if $P_i : a \succ b$. Then, for all $u_i \in \mathcal{U}(P_i)$, we have $\mathbb{E}_{\Phi_i(P_i'')}[u_i] \geq \mathbb{E}_{\Phi_i(P_i')}[u_i]$.*

Hashimoto et al. (2014) show that a mechanism is weakly truncation robust (i.e., the agent cannot obtain a deviation gain by truncating or extending the set of acceptable objects) if and only if it satisfies upper invariance. Conversely, we show that other axioms characterizing BIC are crucial for a truthful reporting of the preference order over objects that are declared to be acceptable.

If an interim mechanism satisfies interior upper invariance in addition to swap monotonicity and lower invariance, BIC is guaranteed to be satisfied on a larger utility subdomain.

Definition 4 (IBI)**.** A utility function u_i satisfies *inverse bounded indifference (IBI) with respect to* $r \in [0, 1]$ if, for all objects $a \in M$ such that $u_i(a) > u_i(\perp)$, we have

$$r \cdot \left(\max_{k \in M} u_i(k) - u_i(\perp) \right) \geq \max_{k \in M} u_i(k) - u_i(a). \tag{8}$$

Since we normalize $\max_{k \in M} u_i(k) = 1$ and $u_i(\perp) = 0$, (8) simplifies to

$$u_i(a) \geq r. \tag{9}$$

We denote the set of all utility functions satisfying IBI with respect to r by $IBI(r)$.

It is clear that $IBI(r) \supset IURBI(r)$ for all $r \in (0,1)$. $IBI(r)$ is increasing in r in the sense that $r' > r$ implies $IBI(r') \supset IBI(r)$. Similar to the URBI and IURBI domains, when $r = 1$, the IBI domain coincides with the full domain: $IBI(1) = \mathcal{U}$. An example of a utility function satisfying $u_i \in IBI(0.5)$ is depicted in Figure 1c.

Suppose that interior upper invariance is satisfied in addition to swap monotonicity and lower invariance. Then, to increase the probability of obtaining the favorite object, an agent must declare an object to be unacceptable. Accordingly, to determine whether dropping an object from the agent's preference order is profitable, we compare (i) the gain of increasing the probability of obtaining his favorite object, (ii) the loss of decreasing the probability of obtaining each object, and (iii) the loss of increasing the probability of being unassigned. Hence, to determine the utility subdomain in which BIC is satisfied, we do not have to impose a condition on relative utilities between acceptable objects (as the IURBI domain does). Instead, we need a condition that guarantees that the agent prefers each acceptable object significantly more than the outside option, as the IBI domain does.

Theorem 4. *Agent i's interim mechanism Φ_i satisfies BIC in $IBI(r)$ with some $r > 0$ if Φ_i satisfies swap monotonicity, lower invariance, and interior upper invariance.*

5 Constrained Random Serial Dictatorship

5.1 Mechanism

In this section, we introduce the *constrained random serial dictatorship mechanism*, designed for optimizing the central planner's policy objective. This mechanism satisfies swap monotonicity, lower invariance, and interior upper invariance, but does not satisfy upper invariance.

Definition 5. A *constrained random serial dictatorship mechanism* (CRSD), parameterized by an objective function $g : X \to \mathbb{R}$, generates a probabilistic assignment in the following way.

1. Based on agents' preference reports $P = (P_i)_{i \in N}$, identify the set of *individually rational (deterministic) assignments* $X^*(P) \subseteq X$ such that for every $x \in X^*$, we have $x_{i,k} = 0$ whenever $P_i : \perp \succ k$.

2. Compute the set of maximizers of g. Let $X^{**}(P) \coloneqq \arg\max_{x \in X^*(P)} g(x)$. We define $X^{**}(P)$ as the set of *approved (deterministic) assignments*.

3. We run the random serial dictatorship mechanism (RSD) to choose some $x \in X^{**}(P)$. Draw agents' priority order π uniformly at random. According to this priority order, each agent sequentially chooses his favorite object from his choice set. Let $\pi(n)$ be the nth agent who makes a choice, and $\psi(n)$ be the object agent $\pi(n)$ chooses. Agent $\pi(n)$ can choose any object that is consistent with some approved assignment, given earlier movers' choices; i.e., any object k such that there exists $x \in X^{**}(P)$ such that $x_{\pi(1),\psi(1)} = 1, x_{\pi(2),\psi(2)} = 1, \ldots, x_{\pi(n-1),\psi(n-1)} = 1, x_{\pi(n),k} = 1$. Agent $\pi(n)$ chooses his favorite object or outside option among such a choice set (with respect to reported preference $P_{\pi(n)}$) and make it $\psi(n)$. Iterate this procedure until the last agent, $\pi(|N|)$, makes his choice. Return the generated deterministic assignment x such that $x_{\pi(n),\psi(n)} = 1$ for $n = 1, 2, \ldots, |N|$ and $x_{i,k} = 0$ otherwise.

Note that the algorithm described above implements CRSD naïvely. Later, we show that CRSD is implementable in polynomial time whenever the objective function g can be optimized in polynomial time (see Subsection 5.4).

CRSD prioritizes the central planner's preference over agents' preferences on objects. Agents are only allowed to choose an object that is consistent with some deterministic assignment that maximizes the objective function g. A deterministic assignment returned by CRSD is always a maximizer of g among all individually rational assignments with respect to the reported preference profile.

The objective function g can capture both (i) the hard constraint on the assignment ($g(x) = -\infty$ if x is infeasible) and (ii) the central planner's preference. Various policy objectives that are known to be incompatible with BIC (and strategy-proofness) can be represented as the objective function g.

Example 1 (Maximum Matching)**.** In many applications (refugee resettlement, daycare assignment, etc.), the *matching size*, defined as the expected number of agents who are assigned to some objects, is one of the primary policy objectives of the central planner. A number of previous studies have analyzed the matching size achieved by existing mechanisms (Krysta et al. 2014, Bogomolnaia and Moulin 2015, Afacan and Dur 2018) and proposed mechanisms that generate a large matching (Andersson and Ehlers 2016, Delacrétaz et al. 2016, Afacan et al. 2018, Kamada and Kojima 2018, Noda 2018a, Ashlagi et al. 2019). The previous studies have shown that strategy-proofness (or BIC), the maximum size, and individual rationality are incompatible (Krysta et al. 2014, Noda 2018b).

If we set $g(x) = \sum_{i \in N} \sum_{k \in M} x_{i,k}$, the objective function g becomes the matching size. If we run CRSD with this objective function g, the central planner first computes the set of maximum matchings, and "breaks a tie" in favor of agents' preferences. Accordingly, the returned matching is always a maximum individually rational matching with respect to the reported preference profile.

Example 2 (Minimum Quota)**.** The minimum quota constraints are also relevant in many real-world settings. For example, the government may want to assign doctors to rural hospitals, and an academic department at a university may want to assign students to all laboratory sections. Many previous studies (e.g. Ehlers et al. 2014, Goto et al. 2014, Fragiadakis et al. 2016, Tomoeda 2018) have studied the incentive properties of the matching mechanisms in settings with minimum quota constraints. In general, when the central planner cannot observe the set of acceptable objects for each agent, minimum quota constraints strongly incentivize agents to truncate their preferences. This is because if an agent declares an object to be acceptable, then he might be forced to take it so as to fulfill the minimum quota requirement.

Let $\mathcal{C}$ be the set of minimum quota constraints. Each $c \in \mathcal{C}$ specifies (i) a subset of objects $M(c) \subseteq M$, and (ii) the minimum quota $q(c) \in \mathbb{Z}_{++}$. The constraint c is satisfied if at least $q(c)$ agents are assigned to a subset of objects $M(c)$. The central planner's objective function g is to maximize the number of minimum quota constraints satisfied by the assignment: $g(x) = \sum_{c \in \mathcal{C}} \mathbf{1}\left\{\sum_{i \in N} \sum_{k \in M(c)} x_{i,k} \geq q(c)\right\}$, where $\mathbf{1}\{\cdot\}$ is an indicator function. With this objective function, each deterministic assignment returned probabilistically by CRSD always satisfies the minimum quota requirement whenever there exists an individually rational assignment that satisfies the minimum quotas. Note that a more general class of constraints (e.g., type-specific minimum quota constraints, maximum quota constraints, proportional constraints) can be represented by an objective function in a similar manner.

Example 3 (Hospital-Optimal Stable Matching)**.** Kojima et al. (2018) consider a doctor-hospital matching with two-sided preferences and distributional constraints, and propose a condition on g under which the generalized deferred acceptance mechanism becomes strategy-proof and generates a *doctor-optimal* stable matching. To represent the hospitals' (joint) preferences and feasibility constraints, they also introduced an objective function. According to their definition, a deterministic assignment is stable if, for every agent, it is impossible to improve the agent's payoff and the hospitals' objective function, simultaneously. If we regard g as a hospitals' payoff function, our CRSD always returns a *hospital-optimal* stable matching. For any input, the resultant assignment maximizes the hospitals' payoff function subject to the individual rationality constraint; thus, it is unimprovable.

5.2 Efficiency

CRSD can be viewed as a variant of serial dictatorship in which the central planner moves as the first dictator. Since the central planner has a preference over the assignments, rather than objects, she first selects a subset of assignments she prefers. After that, just like a standard RSD, agents choose their favorite objects sequentially. Hence, the resultant deterministic assignment is ex post Pareto efficient *for agents and the central planner* in the sense that it is impossible to improve all agents and the central planner's utility simultaneously.

If g represents the matching size, the returned assignment is also ex post Pareto efficient *for agents*; i.e., it is impossible to improve all agents' utilities simultaneously. When g is the matching size, the set of approved assignments coincides with the set of maximum matching. As long as an assignment is individually rational, a maximum matching is never Pareto dominated by non-maximum matching. A deterministic assignment returned by CRSD is not Pareto dominated by another maximum matching either, because CRSD selects a maximum assignment in the manner of serial dictatorship.

Theorem 5.

1. *For any g, CRSD is ex post Pareto efficient for the central planner and agents in the sense that for all reported preference profile P and for all deterministic assignments x that are possibly returned from CRSD, there is no other deterministic assignment $x' \in X \setminus \{x\}$ that satisfies the following two conditions simultaneously: (i) $g(x') \geq g(x)$, and (ii) either $P_i : x'_{i,k} \succ x_{i,k}$ or $x'_{i,k} = x_{i,k}$ holds for all $i \in N$.*

2. *If g represents the matching size, i.e., $g(x) = \sum_{i \in N} \sum_{k \in M} x_{i,k}$, then CRSD is ex post Pareto efficient for agents in the sense that for all deterministic assignments x that are possibly returned from CRSD, there is no other deterministic assignment $x' \in X \setminus \{x\}$ such that either $P_i : x'_{i,k} \succ x_{i,k}$ or $x'_{i,k} = x_{i,k}$ holds for all $i \in N$.*

Proofs are straightforward; thus, they are omitted.

In this sense, although CRSD prioritizes the central planner's objective over agents' preferences, CRSD also respects agents' preferences.

5.3 Incentive Property

CRSD does not satisfy BIC (or strategy-proofness) in general. By pretending some objects to be unacceptable, agents can change the set of maximizers $X^{**}(P)$ so as to obtain a more-preferred object. Since each agent may be able to increase the probability of obtaining a more-preferred object by dropping a less-preferred object from the set of acceptable objects, CRSD does not satisfy upper invariance.

However, CRSD generally satisfies interior upper invariance and lower invariance. The set of individually rational assignments given P, $X^*(P)$ depends only on the set of acceptable objects with respect to P. Hence, the set of approved assignments $X^{**}(P)$ also depends only on the set of acceptable objects. Once the set of approved assignments is fixed, the process is identical to RSD. Accordingly, CRSD is interior upper invariant. It also easy to see that CRSD is lower invariant in the interior. Furthermore, any object that belongs to a lower contour set of the outside option is unacceptable; thus, CRSD never allocates such objects to the agent because CRSD is an individually rational mechanism. For this reason, lower invariance on the boundary is trivial. Accordingly, CRSD is lower invariant.

Theorem 6. *For any objective function g and preference distribution f, CRSD satisfies lower invariance and interior upper invariance.*

We can also show that CRSD satisfies swap monotonicity in the interior in a similar manner to interior upper invariance. Swap monotonicity on the boundary is less straightforward. First, we show the following lemma.

Lemma 1. *Fix any preference profile $P_{-i} \in \mathcal{P}^{N\setminus\{i\}}$ and priority order π. Take any preference $P_i \in \mathcal{P}$ and any neighboring preference $P_i \in \mathcal{N}(P_i)$ such that $P_i : \perp \succ k$ but $P_i' : k \succ \perp$. Let $l \in M \cup \{\perp\}$ be the object assigned to agent i if he reports P_i. Then, if agent i reports preference P_i', he obtains either k or l.*

The proof idea is as follows. Since agent i accepts more objects in P_i', $X^*(P_i, P_{-i}) \subset X^*(P_i', P_{-i})$; thus, the maximized value of g under P_i' is not smaller than that under P_i. If g

takes a strictly larger value, a maximizer should locate in $X^*(P_i', P_{-i}) \setminus X^*(P_i, P_{-i})$. Hence, object k is the only choice available to agent i, and he is forced to take it. Otherwise, the value of g is not changed. Hence, all the approved assignments under P_i are also approved under P_i': $X^{**}(P_i, P_{-i}) \subset X^{**}(P_i', P_{-i})$. If there is an earlier mover who makes a different choice, we refer to the first mover among such agents as agent j. Since agent j's preference is fixed, the change in his choice must be due to the change in his choice set. Since all the agents before agent j are making the same choice, the change in his choice set is caused by the change in agent i's preference report. Since agent j is choosing an object that is not allowed under P_i but allowed under P_i', agent i is forced to take object k, which is newly added to the preference list under P_i'. If we cannot find any earlier movers who change their choice, agent i's choice set under P_i' is that under P_i and object k. Hence, agent i will take either object l, which is the choice under P_i, or object k.

Lemma 1 indicates that CRSD is "weakly" swap monotonic on the boundary in the sense that if agent i drops object k from the preference list (i.e., declares P_i' such that $P_i' : \perp \succ k$ instead of $P_i : k \succ \perp$), then (i) whenever agent i's assignment differs between P_i and P_i', agent i obtains object k, and (ii) whenever agent i obtains some object at P_i, he also obtains some object at P_i'. From the interim perspective, (i) implies $\Phi_i(P_i) \neq \Phi_i(P_i') \Rightarrow \Phi_{i,k}(P_i) > \Phi_{i,k}(P_i')$, and (ii) implies $\Phi_{i,\perp}(P_i) \leq \Phi_{i,\perp}(P_i')$ (the hypothesis is redundant, as we only have a weak inequality here).

To have swap monotonicity on the boundary, whenever the probabilistic assignment is changed (i.e., $\Phi_i(P_i) \neq \Phi_i(P_i')$), the probability of being unmatched must be *strictly* decreased (i.e., $\Phi_{i,\perp}(P_i) < \Phi_{i,\perp}(P_i')$ must be the case). This property is not satisfied with general objective function g. To see this, consider an objective function such that its value is solely determined by agent i's assignment, and its value becomes larger if agent i is allocated to an object with a larger index. In such a case, agent i is guaranteed to be assigned: agent i has at least one acceptable object, and by reporting it to the central planner, he can obtain it. Accordingly, for any P_i and P_i', we have $\Phi_{i,\perp}(P_i') = \Phi_{i,\perp}(P_i) = 0$. However, agent i's assignment changes when he adds an object that has a larger index than any object in the current preference list. Hence, for any preference distribution F, CRSD with this objective function does not satisfy swap monotonicity. Indeed, for any (strict) utility u_i, it is a weakly dominant strategy of agent i to report only his favorite object to be acceptable.

To exclude the counterexample described above, we impose an additional condition on the objective function g with which it is risky to truncate a preference list. For CRSD to be swap monotonic, under some circumstances, CRSD must make the truncator more likely to be unassigned. Although we can consider a number of sufficient conditions, we propose only one simple condition that fits our main applications (maximum matching, minimum quota) well.

Definition 6. Objective function g is *anonymous* if for all permutations of agents, $\mu : N \to N$, we have

$$g(x) = g(x^{\pi}), \tag{10}$$

where x^{π} is defined by $x_i^{\pi} = x_{\pi(i)}$.

Definition 7. Objective function g is *non-decreasing* if whenever $x'_{i,k} \geq x_{i,k}$ for all $(i,k) \in N \times M$, we have $g(x') \geq g(x)$.

It is clear that g is anonymous and non-decreasing when g represents the matching size (Example 1) or the number of minimum quota fulfillments (Example 2).

The anonymity condition requires that no agent is special, and therefore that no agent is guaranteed to be assigned to some object. We assume that g is anonymous and non-decreasing, F has a full support in the sense that any preference order $P_j \in \mathcal{P}$ may realize with a positive probability (this assumption does *not* require that any utility function $u_j \in \mathcal{U}$ may realize with a positive probability), and there is a shortage of objects, i.e., $|N| \geq |M|$. With these assumptions, (i) for every object k, there is a preference profile P_{-i} and priority order π with which agent i can only possibly obtain object k, and (ii) such a situation occurs with a positive probability. Together with Lemma 1, we obtain swap monotonicity of CRSD.

Theorem 7. *Suppose that g is anonymous and non-decreasing, $|N| \geq |M|$, and for every $j \in N \setminus \{i\}$, for every P_j, $F_j(\{u_j \in \mathcal{U} : p(u_j) = P_j\}) > 0$. Then, agent i's interim mechanism implied by CRSD and F is swap monotonic.*

Corollary 1. *Suppose that g is anonymous and non-decreasing, $|N| \geq |M|$, and for every $j \in N \setminus \{i\}$, for every P_j, $F_j(\{u_j \in \mathcal{U} : p(u_j) = P_j\}) > 0$. Then, there exists $r > 0$ such that agent i's interim mechanism implied by CRSD and F satisfies BIC in $IBI(r)$.*

5.4 Computational Complexity

The algorithm presented in Definition 5 is computationally slow. In the second step, the algorithm identifies the set of all the approved assignments $X^{**}(P)$. The number of elements in the set of deterministic assignments grows super-exponentially; thus, when $|N|$ or $|M|$ is large, it is infeasible to compute and save the set of all approved assignments. However, to run CRSD, we do not have to identify the set of all the approved assignments: we should only check whether the assignment agents are going to take is in $X^{**}(P)$.

We denote the set of acceptable objects of agent i by $A_i \coloneqq \{k \in M : P_i : \perp \succ k\}$. First, we define **Unconstrained**(A, g) as the unconstrained optimization problem for deriving the optimized value of g.

$$
\begin{aligned}
& \max_{x \in \mathbb{R}^{N \times M}} g(x) && \textbf{Unconstrained}(A, g) \\
\text{s.t. } & x_{i,k} = 0 \text{ for all } i \in N \text{ and } k \in M \setminus A_i && \text{(Individual Rationality)} \\
& \sum_{k \in M} x_{i,k} \leq 1 \text{ for all } i \in N && \text{(Agent Capacity)} \\
& \sum_{i \in N} x_{i,k} \leq 1 \text{ for all } k \in K && \text{(Object Capacity)} \\
& x_{i,k} \in \{0, 1\} \text{ for all } (i, k) \in N \times M. && \text{(Integer)}
\end{aligned}
$$

The constraint set of **Unconstrained**(A, g) is identical to $X^*(P)$; thus, its value is equal to $\max_{x \in X^*(P)} g(x)$.

Next, we define **Constrained**$(n^*, k^*, (\psi(n))_{n=1}^{n^*-1}, A, g)$ as the constrained optimization problem for checking whether n^*th mover with respect to the priority order π is allowed to take object k^*.

$$\max_{x \in \mathbb{R}^{N \times M}} g(x) \qquad \textbf{Constrained}(n^*, k^*, (\psi(n))_{n=1}^{n^*-1}, A, g)$$

$$\text{s.t. } x_{\pi(n^*),k^*} = 1 \qquad \text{(Choice)}$$

$$x_{\pi(n),\psi(n)} = 1 \text{ for } n = 1, \dots, n^* - 1 \qquad \text{(Consistency)}$$

(Individual Rationality), (Agent Capacity), (Object Capacity), and (Integer).

We regard agent $\pi(i^*)$'s choice (to take object k^*) and the sequence of earlier movers' choices $x_{\pi(n),\psi(n)} = 1$ as constraints. If the value of **Constrained**$(n^*, k^*, (\psi(n))_{n=1}^{n^*-1}, A, g)$ is equal to **Unconstrained**(A, g), then agent $\pi(n^*)$'s choice is consistent with some maximizer of g, and therefore, agent $\pi(n^*)$ is allowed to choose it, and the algorithm finalizes the assignment for agent $\pi(n^*)$. Otherwise, agent $\pi(n^*)$ is forced to change his choice to maintain the value of g. The whole procedure is described as Algorithm 1.

We solve **Unconstrained** only one time, and **Constrained** at most $|N| \cdot |M|$ times. Accordingly, if **Constrained** (and **Unconstrained**) can be run in polynomial time, then CRSD can also be run in polynomial time.

Remark 2. In general, optimization of possibly non-linear function g under the presence of integer constraints ($x_{i,k} \in \{0, 1\}$) might be computationally difficult. However, for some special cases, **Constrained** and **Unconstrained** can be solved efficiently. For example, when the objective is either the matching size or weighted matching size, we may apply the Hopcroft–Karp algorithm (Hopcroft and Karp 1973) or the Hungarian algorithm (Kuhn 1955) to obtain a solution.

Algorithm 1 Faster Implementation of CRSD

Input N, M, g, $P = (P_i)_{i \in N}$.

Output A deterministic assignment $x \in X$.

Initialize $\psi(n) \leftarrow \perp$ for all $n \leftarrow 1, \ldots, |N|$, $x_{i,k} \leftarrow 0$ for all $(i,k) \in N \times M$, and $x_{i,\perp} \leftarrow 1$ for all $i \in N$.

Compute $A_i \coloneqq \{k \in M : P_i : \perp \succ k\}$.

Solve **Unconstrained**(A, g) to obtain its value, $\bar{g}$

Draw a priority order π uniformly at random.

for $n^* = 1, \ldots, N$ **do**

    for $k \in M \cup \{\perp\}$ in a descending order of $P_{\pi(n)}$ **do**

        if $k^* = \perp$ **then** break

        else

            Solve **Constrained**$(n^*, k^*, (\psi(n))_{n=1}^{n^*-1}, A, g)$. Let g' be its value.

            if $g' = \bar{g}$ **then** $\psi(n^*) \leftarrow k^*$, $x_{\pi(n^*),k^*} \leftarrow 1$, $x_{\pi(n^*),\perp} \leftarrow 0$, and break

            end if

        end if

    end for

end for

6 Conclusion

We study an assignment problem in which the central planner has an objective that is not directly associated with agents' welfare. If a mechanism is designed to maximize the central planner's objective function, it has been widely observed that strategy-proofness or BIC cannot be satisfied for an unrestricted domain.

Given this motivation, we extend an axiomatic characterization of Mennle and Seuken (2018) to obtain a domain of utility functions with which we can construct a mechanism that (i) generates an assignment the central planner prefers, and (ii) satisfies BIC in a restricted domain. We show that, if a mechanism satisfies swap monotonicity and lower invariance, then it satisfies BIC in an IURBI domain. If a mechanism further satisfies interior upper invariance, then it satisfies BIC in an IBI domain.

We further construct a new mechanism, CRSD, that (i) returns an assignment that optimizes the central planner's objective among all individually rational assignments (with respect to the reported preference profile), (ii) returns an ex post Pareto efficient assignment in that there is no assignment that improves agents' and the central planner's welfare simultaneously, and (iii) satisfies swap monotonicity, lower invariance, and interior upper invariance; thus, it satisfies BIC in an IBI domain.

We cannot immediately conclude that CRSD can be used in practical situations. CRSD satisfies BIC if agents' vNM utility functions belong to an IBI domain. However, whether agents' utilities are actually in the demanded domain crucially depends on the detail of the situation. Agents' utility attached for acceptable objects might be bounded away from that for the outside option in some applications, but we cannot hope for this property for general problems. Furthermore, the largeness of the IBI domain (parameterized by r) depends on the policy objective g, preference distribution F, and the number of agents and objects $|N|$ and $|M|$, etc. When and whether it is practically possible to implement CRSD is still an open question.

We also note that, when the central planner has an objective, she may secretly manipulate the assignment to make it more preferable or advantageous for his own aims. Akbarpour and Li 2019 study such secret manipulations in an auction problem. Even if the central planner does not manipulate the assignment, agents might doubt it and manipulate their preference reports. Indeed, Rees-Jones (2018) reports that many participants manipulate their preferences in a strategy-proof mechanism. In this sense, mechanism design researchers should pay attention to a situation where the central planner has her own assignment preferences, even if such objectives are not explicitly incorporated into the implemented mechanism.

Appendix

A Proofs

A.1 Proof of Theorem 2

We use the following lemma. The proof is similar to the one for Lemma 1 of Mennle and Seuken (2018), thus it is omitted.

Lemma 2. *The following are equivalent:*

A. *For all preferences $P_i, P_i' \in \mathcal{P}$ with $\Phi_i(P_i) \neq \Phi_i(P_i')$, there exists an object $b \in M$ such that $\Phi_{i,b}(P_i) < \Phi_{i,b}(P_i')$ and $\Phi_{i,k}(P_i) = \Phi_{i,k}(P_i')$ for all $k \in L(b, P_i)$.*

B. *Φ is swap monotonic and lower invariant.*

Theorem 2 claims that Statement B and Statement C, defined as follows, are equivalent.

C. There exists $r > 0$ such that Φ_i satisfies BIC in $IURBI(r)$.

We will prove the equivalence of Statements A and C. Fixing Φ_i, let δ be the smallest non-zero variation in the assignment resulting from any change of report by agent;

$$\delta = \min \left\{ |\Phi_{i,k}(P_i) - \Phi_{i,k}(P_i')| \;\middle|\; \begin{array}{c} i \in N, k \in M, P_i, P_i' \in \mathcal{P}, \\ \text{s.t. } |\Phi_i(P_i) \neq \Phi_i(P_i')| \neq 0 \end{array} \right\} \tag{11}$$

Whenever Φ_i is non-constant, δ is strictly positive.

(Statement C $\Leftarrow$ Statement A) Consider any preferences $P_i, P_i' \in \mathcal{P}$. Statement A implies that whenever $\Phi_i(P_i) \neq \Phi_i(P_i')$, by reporting P_i', agent i becomes more likely to obtain b by δ. Let a be the object that i ranks directly above b in P_i. Then, i's expected deviation gain from misreporting is greatest if the following two conditions are satisfied: (i) when being truthful, agent i obtains object a for sure, and (ii) when misreporting, i obtains his favorite object with probability $1-\delta$ and object b with probability δ.

The deviation gain is thus bounded from above by

$$\left((1-\delta)\max_{k\in M} u_i(k) + \delta u_i(b)\right) - u_i(a) \tag{12}$$

which is non-positive if

$$\delta\left(\max_{k\in M} u_i(k) - u_i(b)\right) \geq \max_{k\in M} u_i(k) - u_i(a). \tag{13}$$

Inequality (13) holds for all utility functions in $\mathcal{U}(P_i)\cap IURBI(\delta)$.

(Statement C $\Rightarrow$ Statement A) We will show the contraposition. Suppose that Statement A is violated. Then, whenever $\Phi_i(P_i)\neq\Phi_i(P_i')$, there exist $P_i, P_i'\in\mathcal{P}$ such that $\Phi_{i,b}(P_i) > \Phi_{i,b}(P_i')$ and $\Phi_{i,k}(P_i)=\Phi_{i,k}(P_i')$ for all $k\in L(b,P_i)$. Again, let a be the object that i ranks directly above b in P_i. Since i's assignment for b decreases, it must decrease by at least δ. Then, i's expected deviation gain from misreporting is smallest if the following two conditions are satisfied: (i) when being truthful, i receives b with probability δ and his favorite object with probability $1-\delta$, and (ii) i receives a for sure when misreporting. Hence,

$$u_i(a) - \left((1-\delta)\max_{k\in M} u_i(k) + \delta u_i(b)\right) \tag{14}$$

is the smallest possible deviation gain from misreporting. This bound is strictly positive if

$$\delta\left(\max_{k\in M} u_i(k) - u_i(b)\right) > \max_{k\in M} u_i(k) - u_i(a), \tag{15}$$

which holds for all utility functions in $\mathcal{U}(P_i)\cap IURBI(r)$ for $r<\delta$. Accordingly, Statement C is not satisfied. ■

A.2 Proof of Theorem 3

We start from any misreport $P_i' \in \mathcal{P}$. P_i'' is obtained by applying the bubble sort algorithm to P_i'. In each step of the bubble sort, we swap the order of two neighboring objects a and b such that $P_i : a \succ b$, which are both more preferred to the outside option. By interior upper invariance and lower invariance, the probability of assigning objects other than a and b are unchanged by this operation. Furthermore, by swap monotonicity, the probability that agent i obtains a can only increase, and the probability that agent i obtains b can only decrease. Accordingly, in each step, agent i's expected utility can only be improved. ■

A.3 Proof of Theorem 4

By Theorem 3, agent i cannot improve his expected payoff only by swapping the preference order of objects that are declared to be acceptable. Accordingly, for a misreporting to be a profitable, it must drop at least one object that is acceptable with respect to the truthful preference. By swap monotonicity, when misreporting, agent i must become unassigned with probability at least δ, where δ is defined in (11). Hence, Agent i's expected deviation gain from misreporting is greatest if the following two conditions are satisfied: (i) when being truthful, agent i obtains his least favorite acceptable object for sure, and (ii) when misreporting, agent i obtains his favorite object with probability $1 - \delta$ and becomes unassigned with probability δ.

The deviation gain is thus bounded above by

$$\left((1-\delta)\max_{k \in M} u_i(k) + \delta u_i(\perp)\right) - \min_{l \in M : u_i(l) > 0} u_i(l) \tag{16}$$

which is non-positive if

$$\delta\left(\max_{k \in M} u_i(k) - u_i(\perp)\right) \geq \max_{k \in M} u_i(k) - \min_{l \in M : u_i(l) > 0} u_i(l). \tag{17}$$

Inequality (17) holds for all utility functions in $\mathcal{U}(P_i) \cap IBI(\delta)$. ■

A.4 Proof of Theorem 6

First, we simultaneously show that CRSD satisfy swap monotonicity, upper invariance, lower invariance in the interior.

Fix an other agents' preferences P_{-i} and the priority order π arbitrarily. Take any $P_i \in \mathcal{P}$, $P_i' \in \mathcal{N}(P_i)$, and objects $a, b \in M$ with $P_i : a \succ b$ but $P_i' : b \succ a$.

Since P_i and P_i' have an identical set of acceptable objects, $X^*(P_i, P_{-i}) = X^*(P_i', P_{-i})$. Accordingly, $X^{**}(P_i, P_{-i}) = X^{**}(P_i', P_{-i})$. Hence, the choice set of the first mover with respect to π, $\pi(1)$ is also identical: he can choose any $k \in M$ such that there exists $x \in X^{**}(P_i, P_{-i}) = X^{**}(P_i', P_{-i})$ such that $x_{\pi(1),k} = 1$. Hence, agent $\pi(1)$ will make an identical choice. Similarly, we can verify that, for any agent $\pi(j)$ who moves earlier than agent i, given that all the agents who move earlier than $\pi(j)$ make an identical choice, agent $\pi(j)$ also makes an identical choice. Accordingly, agent i's choice set is also identical.

Swap Monotonicity in the Interior If agent i obtains object a when he reports P_i, then he obtains b by reporting P_i'. Accordingly, the probability of obtaining a is decreased strictly, and probability of obtaining b is increased strictly. Otherwise, agent i obtains an identical object or outside option. Hence, the assignment is unchanged.

Interior Upper Invariance, Lower Invariance in the Interior By the above argument, whenever agent i's assignment is changed, he obtains object a at P_i and object b at P_{-i}. Accordingly, his probability of (i) obtaining an object more preferred to a and b, and (ii) obtaining an object less preferred to a and b, are unchanged.

Since the above argument holds for every P_{-i} and π, we have interior upper invariance, swap monotonicity in the interior, and lower invariance in the interior.

Finally, we show that lower invariance is also satisfied on the boundary. Again, fix an other agents' preferences P_{-i} and the priority order π arbitrarily. Take any $P_i \in \mathcal{P}$, $P_i' \in \mathcal{N}(P_i)$, and objects $a \in M$ with $P_i : k \succ \perp$ but $P_i' : \perp \succ k$. Since CRSD always selects an individually rational assignment, for all $l \in L(\perp, P_i)$, we have $\Phi_{i,l}(P_i) = 0$. Similarly, for all $l \in L(\perp, P_i')$, we have $\Phi_{i,l}(P_i') = 0$. Since $L(\perp, P_i') = L(\perp, P_i) \cup \{k\} \subset L(\perp, P_i)$, we have $\Phi_{i,l}(P_i) = \Phi_{i,l}(P_i')$ $(= 0)$ for all $l \in L(\perp, P_i)$, as desired. ■

A.5 Proof of Lemma 1

Since P_i' accepts a larger set of objects than P_i does, we have $X^*(P_i', P_{-i}) \supset X^*(P_i, P_{-i})$. Hence, we have

$$\max_{x \in X^*(P_i', P_{-i})} g(x) \geq \max_{x \in X^*(P_i, P_{-i})} g(x). \tag{18}$$

When (18) holds with strict inequality, we have $X^{**}(P_i', P_{-i}) \subseteq X^*(P_i', P_{-i}) \setminus X^*(P_i, P_{-i})$. By construction, for all $x \in X^*(P_i', P_{-i}) \setminus X^*(P_i, P_{-i})$, agent i is assigned to object k.

When (18) holds with equality, we have $X^{**}(P_i, P_{-i}) \subseteq X^{**}(P_i', P_{-i})$. Hence, whenever the first mover, agent $\pi(1)$, can take object k given P_i, he can also take it given P_i'. If the choice of agent $\pi(1)$ given P_i' is different from that given P_i, agent $\pi(1)$ is choosing an object that is (i) not available under P_i, but (ii) available under P_i'. Such an assignment must belong to $X^{**}(P_i', P_{-i}) \setminus X^{**}(P_i, P_{-i})$; thus, in such a case, agent i is forced to take object k.

We look at each of earlier movers' choices sequentially. If one of earlier movers changes his choice, by the above argument, we can conclude that agent i is forced to take object k. Otherwise, all agents who move earlier than agent i make an identical choice between P_i and P_i'. Then, agent i's choice set given P_i' is the union of the choice set given P_i and a singleton of object k. Accordingly, agent i will take either (i) object l, which is the choice under P_i, or (ii) object k.

In all the cases above, agent i ends up with taking either object k or l. ■

A.6 Proof of Theorem 7

In the proof of Theorem 6, we showed that CRSD is swap monotonic in the interior. Furthermore, by Lemma 1, for any $P_i \in \mathcal{P}$, $P_i' \in \mathcal{N}(P_i)$ such that $P_i' : k \succ \perp$ but $P_i : \perp \succ k$, we have (i) $\Phi_i(P_i') \neq \Phi_i(P_i)$ implies $\Phi_{i,k}(P_i') > \Phi_{i,k}(P_i)$, and (ii) $\Phi_{i,\perp}(P_i') \leq \Phi_{i,\perp}(P_i)$. Given that we assume that any P_{-i} is realized with a positive probability, it suffices to show that, for every $P_i \in \mathcal{P}_i$, there exists $P_{-i} \in \mathcal{P}^{N \setminus \{i\}}$ such that $\phi_{i,\perp}(P_i', P_{-i}) < \phi_{i,\perp}(P_i, P_{-i})$.

Take any $P_i \in \mathcal{P}$, $P_i' \in \mathcal{N}(P_i)$ such that $P_i' : k \succ \perp$ but $P_i : \perp \succ k$. Define P_j by $P_j = P_i$ for all $j \in N \setminus \{i\}$. Take a priority order such that agent i moves at the very end. Since P_i declares at least one object (k) to be unacceptable, P_i accepts at most $|M| - 1$ objects. Since g is anonymous and $|N| \geq |M|$, if agent i reports P_i, all the objects accepted by P_i are taken by earlier movers; thus, he is assigned to $\perp$. However, since g is non-decreasing and no $j \in N \setminus \{i\}$ declares that object k is acceptable, if agent i reports P_i', agent i is assigned to object k. Since such a priority order π occurs with a positive probability, we have $\phi_{i,\perp}(P_i', P_{-i}) < \phi_{i,\perp}(P_i, P_{-i})$, as desired. ■

References

Abdulkadirolu, A. and T. Sonmez (1998): "Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems," *Econometrica*, 66, 689–701.

Afacan, M. O., I. Bó, and B. Turhan (2018): "Assignment Maximization," Working Paper.

Afacan, M. O. and U. M. Dur (2018): "Strategy-Proof Size Improvement: Is It Possible?" Working Paper.

Akbarpour, M. and S. Li (2019): "Credible Mechanisms," Working Paper.

Andersson, T. and L. Ehlers (2016): "Assigning Refugees to Landlords in Sweden: Stable Maximum Matchings," Working Paper.

Ashlagi, I., A. Saberi, and A. Shameli (2019): "Assignment Mechanisms under Distributional Constraints," in *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, Society for Industrial and Applied Mathematics, 229–240.

Balbuzanov, I. (2016): "Convex Strategyproofness with an Application to the Probabilistic Serial Mechanism," *Social Choice and Welfare*, 46, 511–520.

Bansak, K., J. Ferwerda, J. Hainmueller, A. Dillon, D. Hangartner, D. Lawrence, and J. Weinstein (2018): "Improving Refugee Integration through Data-Driven Algorithmic Assignment," *Science*, 359, 325–329.

Birkhoff, G. (1946): "Three Observations on Linear Algebra," *Univ. Nac. Tacuman, Rev. Ser. A*, 5, 147–151.

Bogomolnaia, A. and H. Moulin (2001): "A New Solution to the Random Assignment Problem," *Journal of Economic Theory*, 100, 295–328.

——— (2004): "Random Matching under Dichotomous Preferences," *Econometrica*, 72, 257–279.

——— (2015): "Size Versus Fairness in the Assignment Problem," *Games and Economic Behavior*, 90, 119–127.

Delacrétaz, D., S. D. Kominers, and A. Teytelboym (2016): "Refugee Resettlement," Working Paper.

Ehlers, L., I. E. Hafalir, M. B. Yenmez, and M. A. Yildirim (2014): "School Choice with Controlled Choice Constraints: Hard Bounds Versus Soft Bounds," *Journal of Economic Theory*, 153, 648 – 683.

Fragiadakis, D., A. Iwasaki, P. Troyan, S. Ueda, and M. Yokoo (2016): "Strategyproof Matching with Minimum Quotas," *ACM Transactions on Economics and Computation*, 4, 6.

Fragiadakis, D. and P. Troyan (2017): "Improving Matching under Hard Distributional Constraints," *Theoretical Economics*, 12, 863–908.

Gale, D. and L. Shapley (1962): "College Admissions and the Stability of Marriage," *The American Mathematical Monthly*, 69, 9–15.

Goto, M., N. Hashimoto, A. Iwasaki, Y. Kawasaki, S. Ueda, Y. Yasuda, and M. Yokoo (2014): "Strategy-Proof Matching with Regional Minimum Quotas," in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, AAMAS '14, 1225–1232.

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.

Hopcroft, J. E. and R. M. Karp (1973): "An $n^{5/2}$ Algorithm for Maximum Matchings in Bipartite Graphs," *SIAM Journal on computing*, 2, 225–231.

Kamada, Y. and F. Kojima (2015): "Efficient Matching under Distributional Constraints: Theory and Applications," *American Economic Review*, 105, 67–99.

——— (2018): "Fair Matching under Constraints: Theory and Applications," Working Paper.

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

Kojima, F. and P. A. Pathak (2009): "Incentives and Stability in Large Two-Sided Matching Markets," *American Economic Review*, 99, 608–27.

Kojima, F., A. Tamura, and M. Yokoo (2018): "Designing Matching Mechanisms under Constraints: An Approach from Discrete Convex Analysis," *Journal of Economic Theory*, 176, 803 – 833.

Krysta, P., D. Manlove, B. Rastegari, and J. Zhang (2014): "Size Versus Truthfulness in the House Allocation Problem," in *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, ACM, 453–470.

Kuhn, H. W. (1955): "The Hungarian Method for the Assignment Problem," *Naval Research Logistics Quarterly*, 2, 83–97.

Mennle, T. and S. Seuken (2017): "Two New Impossibility Results for the Random Assignment Problem," Workign Paper.

——— (2018): "Partial Strategyproofness : Relaxing Strategyproofness for the Random Assignment Problem," Working Paper.

Noda, S. (2018a): "Large Matchings in Large Markets with Flexible Supply," Working Paper.

——— (2018b): "Size versus Truncation Robustness in the Assignment Problem," Working Paper.

Rees-Jones, A. (2018): "Suboptimal Behavior in Strategy-Proof Mechanisms: Evidence from the Residency Match," *Games and Economic Behavior*, 108, 317 – 330, special Issue in Honor of Lloyd Shapley: Seven Topics in Game Theory.

Song, Y., C. Zhang, and Y. Fang (2008): "Multiple Multidimensional Knapsack Problem and Its Applications in Cognitive Radio Networks," in *MILCOM 2008-2008 IEEE Military Communications Conference*, IEEE, 1–7.

Tomoeda, K. (2018): "Finding a Stable Matching under Type-Specific Minimum Quotas," *Journal of Economic Theory*, 176, 81 – 117.

Von Neumann, J. (1953): "A Certain Zero-Sum Two-Person Game Equivalent to the Optimal Assignment Problem," in *Contributions to the Theory of Games*, ed. by W. Kuhn and A. W. Tucker, Princeton: Princeton University Press, vol. 2, 5–12.