Serial Dictatorship Mechanisms with Reservation Prices: Heterogeneous Objects

Bettina Klaus* Alexandru Nichifor†

October 16, 2019

Abstract

We adapt a set of mechanisms introduced by Klaus and Nichifor (2019), *serial dictatorship mechanisms with (individual) reservation prices*, to the allocation of heterogeneous indivisible objects, e.g., specialist clinic appointments. We show how the characterization of serial dictatorship mechanisms with reservation prices for homogeneous indivisible objects (Klaus and Nichifor, 2019, Theorem 1) can be adapted to the allocation of heterogeneous indivisible objects by adding *neutrality*: mechanism φ satisfies *minimal tradability*, *individual rationality*, *strategy-proofness*, *consistency*, *independence of unallocated objects*, *neutrality*, and *non wasteful tie-breaking* if and only if there exists a reservation price vector r and a priority ordering $\succ$ such that φ is a *serial dictatorship mechanism with reservation prices* based on r and $\succ$.

JEL classification: C78, D47, D71

Keywords: heterogeneous objects; serial dictatorship; individual reservation prices; axiomatic characterization; strategy-proofness; consistency.

* *Corresponding author*, *e-mail*: Bettina.Klaus@unil.ch; Faculty of Business and Economics, University of Lausanne, Internef, CH-1015 Lausanne, Switzerland. Bettina Klaus gratefully acknowledges financial support from the Swiss National Science Foundation (SNFS), project 100018_162606.

† *E-mail*: nichifor@unimelb.edu.au; Faculty of Business and Economics, University of Melbourne, 111 Barry Street, Parkville, 3010, Victoria, Australia. Nichifor gratefully acknowledges financial support via the Australian Research Council's Discovery Early Career Researcher Award (DECRA), project DE170101183.

1 Introduction

We consider the problem of allocating heterogeneous indivisible objects to agents when each agent receives at most one object and pays a non-negative price. Each agent's preferences over receiving an object and his own payment are given by a general utility function that is not necessarily quasilinear. A *mechanism* selects an *outcome* for the problem by allocating an object and specifying a payment for each agent, i.e., it selects an *allotment* for each agent. We are interested in mechanisms that have desirable properties. More specifically, we consider mechanisms that: satisfy mild efficiency criteria (*minimal tradability* and *non-wasteful tie-breaking*); induce voluntarily participation (*individual rationality*); elicit agents' true valuations for the objects (*strategy-proofness*); select outcomes that are robust, in the sense of remaining invariant, when agents leave from the problem with their allotments (*consistency*) or when we remove undesired objects (*independence of unallocated objects*); and in which the names of the objects do not matter, and the objects can be relabelled while the payments are kept invariant (*neutrality*).

We show that a mechanism φ satisfies all the properties mentioned above if and only if there exists a reservation price vector r, which specifies an individual reservation price for each agent that is the same for all objects, and a priority ordering $\succ$ over the set of agents, such that φ is a *serial dictatorship mechanism with reservation prices* that is based on r and $\succ$ (Theorem 1). Our characterization is tight in the sense that each property used is indispensable.

Intuitively, a serial dictatorship mechanism with reservation prices works as follows. Agents sequentially get to choose feasible objects according to their priority, where the feasible objects are those remaining after all preceding agents made their choices: If the choosing agent's value for his most preferred object among the feasible ones exceeds his reservation price, he takes the object and he pays his reservation price; otherwise, he receives and pays nothing.

Our model, properties, and mechanisms, are well-suited for understanding real-life markets in which: wealth inequality among agents can be substantial; income redistribution is not feasible; sequential priorities as a main criteria for rationing demand are considered *fair* (Konow, 2003) and *just* (Dold and Khadjavi, 2017) – and are thus desirable, while maintaining compatibility with some payments is also required. More specifically, in Section 5 we discuss how and why our model with heterogeneous objects, and our results, help further explain and understand the allocation of specialist medical services; for concreteness, we focus on the allocation of next-available consultant-led medical appointments in Australia.

Our work is closely related to that of Klaus and Nichifor (2019). We extend their homogeneous indivisible objects model, normative properties, and the class of serial dictatorship mechanisms with reservation prices that they introduced, to heterogeneous indivisible objects. Our characterization (Theorem 1), which uses one additional key property, *neutrality*, provides a counterpart to the main result of Klaus and Nichifor (2019, Theorem 1).[1]

For settings in which agents' preferences are given by linear orders and there are no payments, several characterizations of the classical serial dictatorship mechanisms are available, e.g., Svensson (1994), Svensson (1999), Ergin (2000), Ehlers and Klaus (2007). Relative to those settings, we extend the model to allow for general utility functions, and we adapt the properties used to characterize classical serial dictatorship mechanisms, as well as the mechanism itself, in a way in which we maintain sequential priorities as the main rationing criteria (as desired), but we relax the earlier limitations that ruled out any payments. Our model and key properties end up being closer to those of Tadenuma and Thomson (1991) and Svensson and Larsson (2002), who do not characterize any mechanism as such.[2] Note that in our characterization the reservation price vector and the priority ordering are both derived from the properties, together with the serial dictatorship mechanism with reservation prices that is based on them; in this sense, our approach is related to the recent characterizations of deferred acceptance mechanisms (Kojima and Manea, 2010; Ehlers and Klaus, 2014, 2016) in which the priorities (or more generally, the choice functions) that the mechanisms are based on, and the mechanisms, are all derived together, from the properties.

2 Model

Our model extends the homogeneous objects model of Klaus and Nichifor (2019) to heterogeneous objects; to ease the comparison, our exposition and notation are closely based on that of Klaus and Nichifor (2019).

A set of heterogeneous indivisible objects are to be allocated to a set of agents; the sets of objects and the set of agents can change. Let $\mathbb{N}$ be the set of potential agents and $\mathcal{N}$ be the set of all non-empty finite subsets of $\mathbb{N}$, $\mathcal{N} \equiv \{N \subseteq \mathbb{N} : 0 < |N| < \infty\}$. Let $\mathbb{O}$ be the set of potential real objects and $\mathcal{O}$ be the set of all non-empty finite subsets of $\mathbb{O}$,

[1]In Section 4, we discuss in detail the technical similarities and the differences between our Theorem 1 and the main result of Klaus and Nichifor (2019, Theorem 1), and the role played by *neutrality*.

[2]We more precisely place each of the properties that we use within the relevant literature in Section 2, immediately after we formally introduce each property.

$\mathcal{O} \equiv \{O \subseteq \mathbb{N} : 0 < |O| < \infty\}$. We assume that $|\mathbb{O}| > 1$ and that $\mathbb{O}$ is infinite.[3] By 0 we denote the *null object*, which represents not receiving a real object in $\mathbb{O}$.

For any set of agents $N \in \mathcal{N}$ and any set of real objects $O \in \mathcal{O}$, an *allocation vector* $a = (a_i)_{i\in N} \in \big(O \cup \{0\}\big)^N$ such that [for any two agents $i, j \in N$, $i \neq j$, $a_i = a_j$ implies $a_i = a_j = 0$] describes which agent in N receives which object; we allow for the possibilities that no real objects, or only some, are allocated. We denote the *set of allocation vectors* for a set of agents $N \in \mathcal{N}$ and a set of real objects $O \in \mathcal{O}$ by $\mathcal{A}(N, O)$.

We assume that an agent $i \in \mathbb{N}$ may have to pay a non-negative *price* $p_i \in \mathbb{R}_+$, and we denote the *set of payment vectors* for a set of agents $N \in \mathcal{N}$ by $\mathcal{P}(N) \equiv \big\{p = (p_i)_{i\in N} : p \in \mathbb{R}^N_+\big\}$.

We assume that agents only care about the object they receive and their own payment. Each agent $i \in \mathbb{N}$ has preferences that are: **(i)** for any object strictly decreasing in the price paid; **(ii)** such that given the same price, receiving a real object is weakly better than receiving the null object; **(iii)** such that for any real object, either there exists a price which makes the agent indifferent between [receiving the real object at this price] and [receiving the null object and paying nothing], or he strictly prefers to [obtain the real object, whatever the price] over [receiving the null object and paying nothing] or he strictly prefers to [receive the null object and pay nothing] over [obtaining the real object, whatever the price]; and **(iv)** strict, when comparing any two real objects at the same price. Formally, for each $O \in \mathcal{O}$, we represent each agent i's preferences ($i \in \mathbb{N}$) by a utility function $u_i : (O \cup \{0\}) \times \mathbb{R}_+ \to \mathbb{R}$ that satisfies the following four properties:

(i) if $0 \leq p'_i < p_i$, then for each $o \in O$, $u_i(o, p'_i) > u_i(o, p_i)$ and $u_i(0, p'_i) > u_i(0, p_i)$;

(ii) for each $p_i \geq 0$ and each $o \in O$, $u_i(o, p_i) \geq u_i(0, p_i)$;

(iii) for each $o \in O$,

either [there exists a price $v_{i,o}$ such that $u_i(o, v_{i,o}) = u_i(0, 0)$],

or [for each $p_i \geq 0$, we have $u_i(o, p_i) > u_i(0, 0)$ and $v_{i,o} \equiv \infty$] or [for each $p_i \geq 0$, we have $u_i(0, 0) > u_i(o, p_i)$ and $v_{i,o} \equiv -\infty$];

$v_i = (v_{i,o})_{o\in O}$ is agent i's *valuation vector*[4] and we set $v_{i,0} = 0$.

(iv) Agent i's preferences over real objects are strict in the sense that for any pair of real objects $o_1, o_2 \in O$, $o_1 \neq o_2$, and any payment p_i, we have $u_i(o_1, p_i) \neq u_i(o_2, p_i)$;

[3] Finite sets of potential agents and potential real objects would not change any of our results.

[4] Requiring continuity of u_i would be a less general assumption that guarantees the existence of valuation vector v_i.

hence, for each set of real objects $O \in \mathcal{O}$, the *best real object for agent i in O* is well defined and we denote it by $\text{top}_i(O) = \arg\max\{v_{i,o}\}_{o \in O}$.

An example of an agent i's preferences with valuation vector v_i are quasilinear preferences u_i defined for each $(a_i, p_i) \in (O \cup \{0\}) \times \mathbb{R}_+$ by $u_i(a_i, p_i) = v_{i,a_i} - p_i$.

For any set of agents $N \in \mathcal{N}$ and any set of real objects $O \in \mathcal{O}$, we denote the *set of utility (function) profiles* by $\mathcal{U}(N, O)$ and the associated *set of valuation (vector) profiles* by $\mathcal{V}(N, O)$.

A *problem* γ is specified by a triple (N, O, u) such that $(N, O) \in \mathcal{N} \times \mathcal{O}$ and $u \in \mathcal{U}(N, O)$. We denote the *set of all problems* for $(N, O) \in \mathcal{N} \times \mathcal{O}$ by $\Gamma(N, O)$.

An *outcome* for any problem $\gamma \in \Gamma(N, O)$ consists of an allocation vector $a \in \mathcal{A}(N, O)$ and a payment vector $p \in \mathcal{P}(N)$. We denote the *set of outcomes* for a problem $\gamma \in \Gamma(N, O)$ by $\mathcal{O}(N, O) \equiv \mathcal{A}(N, O) \times \mathcal{P}(N)$.

A *mechanism* φ is a function that assigns an outcome to each problem. Formally, for each $(N, O) \in \mathcal{N} \times \mathcal{O}$ and each $\gamma \in \Gamma(N, O)$, $\varphi(\gamma) \in \mathcal{O}(N, O)$. Note that we can also represent a mechanism φ by its *allocation rule* α and *payment rule* π, i.e., for each $(N, O) \in \mathcal{N} \times \mathcal{O}$ and each $\gamma \in \Gamma(N, O)$, $\alpha : \Gamma(N, O) \to \mathcal{A}(N, O)$, $\pi : \Gamma(N, O) \to \mathcal{P}(N)$, and $\varphi(\gamma) = (\alpha(\gamma), \pi(\gamma))$. We denote the *allotment* of agent i at outcome $\varphi(\gamma)$ by $\varphi_i(\gamma) = (\alpha_i(\gamma), \pi_i(\gamma))$.

Given $N \in \mathcal{N}$, a vector $x \in \mathbb{R}^N$, and $M \subseteq N$, let $x_M \equiv (x_i)_{i \in M} \in \mathbb{R}^M$ be the restriction of vector x to the subset of agents M. We also use the notation $x_{-i} = x_{N \setminus \{i\}}$. For example, $(\bar{x}_i, x_{-i})$ denotes the vector obtained from x by replacing x_i with $\bar{x}_i$. We use corresponding notational conventions for utility profiles.

Properties of Mechanisms

Our first property ensures that if there are at least as many agents as objects, then there is some utility profile at which all objects are allocated.

Definition 1 (Minimal Tradability). A mechanism φ satisfies *minimal tradability* if for each $(N, O) \in \mathcal{N} \times \mathcal{O}$ such that $|O| \leq |N|$, there exists a utility profile $u \in \mathcal{U}(N, O)$ such that $\bigcup_{i \in N}\{\alpha_i(N, O, u)\} = O$.

Minimal tradability was first introduced for single-object problems by Sakai (2013), and then extended to problems with homogeneous objects by Klaus and Nichifor (2019).

Our definition of *minimal tradability* coincides with that of Sakai's (2013) for single-object problems, but it is in character less demanding than that of Klaus and Nichifor (2019).[5]

For $(N, O) \in \mathcal{N} \times \mathcal{O}$, an outcome $(a, p) \in \mathcal{O}(N, O)$ is *individually rational* for utility profile $u \in \mathcal{U}(N, O)$ with associated valuation vector $v \in \mathcal{V}(N, O)$ if for each $i \in N$, we have $u_i(a_i, p_i) \geq u_i(0, 0)$, or equivalently,

(IR1) $[a_i = 0$ implies $p_i = 0]$ and

(IR2) [for each $o \in O$, $a_i = o$ implies $p_i \leq v_{i,o}$].

By requiring a mechanism to only choose individually rational outcomes, we express the idea of voluntary participation.

Definition 2 (Individual Rationality). A mechanism φ satisfies *individual rationality* if for each $(N, O) \in \mathcal{N} \times \mathcal{O}$ and each $\gamma \in \Gamma(N, O)$, $\varphi(\gamma)$ is an *individually rational* outcome.

Strategy-proofness requires that no agent can benefit from misrepresenting his preferences.

Definition 3 (Strategy-Proofness). A mechanism φ satisfies *strategy-proofness* if for each $(N, O) \in \mathcal{N} \times \mathcal{O}$, each $(N, O, u) \in \Gamma(N, O)$, each $i \in N$, and each u'_i such that $u' \equiv (u'_i, u_{-i}) \in \mathcal{U}(N, O)$, $u_i(\varphi_i(N, O, u)) \geq u_i(\varphi_i(N, O, u'))$.

That is, a mechanism is *strategy-proof* if (in the associated direct revelation game) it is a weakly dominant strategy for each agent to report his utility truthfully.

To introduce our next property, *consistency*, which is a key requirement in many frameworks with variable population, we first define what a *reduced problem* is.

Let $(N, O) \in \mathcal{N} \times \mathcal{O}$, $\gamma = (N, O, u) \in \Gamma(N, O)$, and $M \subsetneq N$. When the set of agents M leaves problem γ with their $\bigcup_{i \in M} \alpha_i(\gamma)$ allocated objects, the set of remaining objects is $O_{N \setminus M} = O \setminus \bigcup_{i \in M} \alpha_i(\gamma)$. Hence, the *reduced problem* is $\gamma_{N \setminus M} = (N \setminus M, O_{N \setminus M}, u_{N \setminus M})$, where $u_{N \setminus M} \in \mathcal{U}(N \setminus M, O_{N \setminus M})$ is obtained from u by deleting the utilities of agents in M as well as the utility information of removed objects.[6]

[5]Klaus and Nichifor (2019) define *minimal tradability* by requiring that (i) if there are at least as many agents as objects, then there is some utility profile for which all objects are allocated and (ii) if there are more objects than agents, then there is some utility profile at which each agent receives an object. Our definition of *minimal tradability* is in character strictly weaker than that of Klaus and Nichifor (2019) because we drop their second requirement.

[6]Strictly speaking, the notation $u_{N \setminus M}$ has only been introduced for the removal of agents $N \setminus M$ from utility profile u but in the context of *consistency*, the additional adjustment to a smaller set of objects should neither be confusing nor require additional notation.

Consistency is an invariance requirement of the solution if some agents leave together with their allotments. That is, *consistency* requires that if some agents leave with their allotments, then in the resulting reduced problem the allocation and the payment for all remaining agents should not change.

Definition 4 (Consistency). A mechanism φ satisfies *consistency* if for each $(N, O) \in \mathcal{N} \times \mathcal{O}$, each $\gamma \in \Gamma(N, O)$, and each $M \subseteq N$, we have $\varphi(\gamma_{N \setminus M}) = \varphi(\gamma)_{N \setminus M}$.

Consistency, first introduced by Thomson (1983), is one of the key properties in many frameworks with variable populations (see Thomson, 2015). We use a similar notion of *consistency* as Tadenuma and Thomson (1991) do (since our models are similar) but adapt it to apply to functions (they allow for correspondences) and we decompose it into two properties: our *consistency* together with our next property, *independence of unallocated objects*, corresponds to the direct adaptation of Tadenuma and Thomson's *consistency* to our model.

Next, we require that if not all objects are allocated, removing some of the unallocated objects leaves the outcome unchanged.

Definition 5 (Independence of Unallocated Objects). A mechanism φ satisfies *independence of unallocated objects* if for each $(N, O) \in \mathcal{N} \times \mathcal{O}$, each $(N, O, u) \in \Gamma(N, O)$, and each $O' \subseteq O \setminus \bigcup_{i \in N} \alpha_i(N, O, u)$, we have $\varphi(N, O, u) = \varphi(N, O \setminus O', u_{O \setminus O'})$, where $u_{O \setminus O'} \in \mathcal{U}(N, O \setminus O')$ is obtained from u by deleting the utility information of removed objects.

Independence of unallocated objects was introduced for a homogeneous goods model by Klaus and Nichifor (2019) who required that removing all unallocated goods leaves an outcome unchanged. We extend their property to allow for heterogeneous goods, and we require that only *some*, but not necessary *all*, unallocated goods be removed.

To introduce our next property, *neutrality*, which is a key requirement in many frameworks in which the names of the objects should not matter in the allocation process, we first define what a *relabelling of the objects* is.

For each $O \in \mathcal{O}$, a *relabelling of the objects* is given by a permutation function $\sigma : O \cup \{0\} \to O \cup \{0\}$ with $\sigma(0) = 0$, i.e., under σ the names of the real objects are exchanged, e.g., object $o \in O$ becomes object $\sigma(o) \in O$, while the naming of the null object remains unchanged. We denote the *set of relabellings* for a set of real objects $O \in \mathcal{O}$ by $\mathcal{S}(O)$.

Let $(N, O) \in \mathcal{N} \times \mathcal{O}$ and $\sigma \in \mathcal{S}(O)$. Then, for each utility profile $u \in \mathcal{U}(N, O)$ with associated valuation vector $v \in \mathcal{V}(N, O)$, a *relabelling of the utility profile* $u^\sigma \in \mathcal{U}(N, O)$ with associated *relabelling of valuation vector profile* $v^\sigma \in \mathcal{V}(N,O)$ is such that for each $i \in N$ and each $o \in O$, we have $u^\sigma_{i,o} = u_{i,\sigma^{-1}(o)}$ and $v^\sigma_{i,o} = v_{i,\sigma^{-1}(o)}$.

For example, consider $N = \{1,2,3\}$, $O = \{a,b,c\}$, utility profile $u \in \mathcal{U}(N,O)$ with associated valuation vector $v \in \mathcal{U}(N,O)$, and relabeling $\sigma(a) = b$, $\sigma(b) = c$, $\sigma(c) = a$, and $\sigma(0) = 0$. Then, for each agent $i \in N$, $u_i = (u_{i,a}, u_{i,b}, u_{i,c})$ and $v_i = (v_{i,a}, v_{i,b}, v_{i,c})$ are relabelled by σ to $u^\sigma_i = (u^\sigma_{i,a}, u^\sigma_{i,b}, u^\sigma_{i,c}) = (u_{i,c}, u_{i,a}, u_{i,b})$ and $v^\sigma_i = (v^\sigma_{i,a}, v^\sigma_{i,b}, v^\sigma_{i,c}) = (v_{i,c}, v_{i,a}, v_{i,b})$.

A mechanism is *neutral* if a *relabelling of the objects* results in each agent being allocated the object that is the relabelled version of the object that he was previously allocated, while the payments for all agents remain the same as before.

Definition 6 (Neutrality). A mechanism φ satisfies *neutrality* if for each $(N,O) \in \mathcal{N} \times \mathcal{O}$, each $\sigma \in \mathcal{S}(O)$, and for each $i \in N$, we have

$$\alpha_i(N,O,u^\sigma) = \sigma(\alpha_i(N,O,u)) \text{ and } \pi_i(N,O,u^\sigma) = \pi_i(N,O,u).$$

For our three agent and three object example above, if say $\alpha(N,O,u) = (a,b,c)$ and $\pi(N,O,u) = (5,0,1)$, then *neutrality* would imply that $\alpha(N,O,u^\sigma) = (b,c,a)$ and $\pi(N,O,u^\sigma) = (5,0,1)$.

Neutrality was first introduced by Smith (1973) in a voting context. We use the same notion of *neutrality* as Svensson and Larsson (2002) do, and our models are similar, except that we allow for more general preferences than Svensson and Larsson who require quasilinear preferences.

Our last property requires that a mechanism does not select an outcome where an agent is indifferent between [receiving a real object at some price] and [not receiving an object at and not paying anything, i.e., withdrawing from the market].

Definition 7 (Non Wasteful Tie-Breaking). A mechanism φ satisfies *non wasteful tie-breaking* if for each $(N,O) \in \mathcal{N} \times \mathcal{O}$, each $\gamma \in \Gamma(N,O)$, each $i \in N$, and each $o \in O$, $\alpha_i(\gamma) = o$ implies that $u_i(o, \pi_i(\gamma)) \neq u_i(0,0)$.

Non wasteful tie-breaking, first introduced by Klaus and Nichifor (2019), is a mild *efficiency* requirement: The *tie-breaking*, which rules out allocating a real object at some price to an agent who is indifferent between such an allotment and withdrawing from the market, is *non-wasteful* in that it keeps the object available, because another agent might strictly prefer to receive it.

3 Serial Dictatorships with Reservation Prices

We adapt the class of *serial dictatorships with reservation prices* introduced by Klaus and Nichifor (2019) for the allocation of *homogeneous* indivisible objects to our model with *heterogeneous* objects as follows.

First, we need to define and fix *reservation prices* and a *priority ordering*.

We assume that for each agent $i \in \mathbb{N}$ a *(fixed) reservation price* $f_i \geq 0$ exists. We interpret f_i as the price at which a real object can be allocated to agent i. Note that the reservation price f_i is the same for different real objects. We denote a *vector of (fixed) reservation prices* for the set of potential agents $\mathbb{N}$ by $f = (f_i)_{i \in \mathbb{N}}$ and by $\mathcal{F}$ we denote the *set of all (fixed) reservation price vectors* for $\mathbb{N}$.

A *priority ordering* $\rhd$ over the set of potential agents $\mathbb{N}$ is a complete, asymmetric, and transitive binary relation, with the interpretation that for any two distinct agents $i, j \in \mathbb{N}$, $i \rhd j$ means that i has a higher priority than j. Note that the priority ordering $\rhd$ is the same for different real objects. Let $\mathcal{P}$ denote the *set of all priority orderings* over $\mathbb{N}$.

Given a reservation price vector $f \in \mathcal{F}$ and a priority ordering $\rhd \in \mathcal{P}$, the *serial dictatorship mechanism with reservation prices* based on f and $\rhd$ is denoted by $\psi^{(f,\rhd)}$ and determines an outcome for each problem $\gamma = (N, O, u) \in \Gamma(N, O)$ with associated valuation vector $v \in \mathcal{V}(N, O)$ as follows.

Step 0: If there are no real objects to allocate, then stop and all agents receive and pay nothing. Otherwise, continue.

Step 1: The agent with the highest priority in N is considered. Let $i \in N$ be this agent.

- If for his most preferred object $\text{top}_i(O)$ in O, $v_{i,\text{top}_i(O)} > f_i$, then agent i obtains $\text{top}_i(O)$ and pays f_i. Set $O_1 := O \setminus \{\text{top}_i(O)\}$. If $O_1 = \emptyset$, then we stop and all remaining agents receive and pay nothing. Otherwise, continue.
- If $v_{i,\text{top}_i(O)} \leq f_i$, then agent i receives and pays nothing. Set $O_1 := O$ and continue.

Step *l*: The agent with the l^{th} highest priority in N is considered. Let $j \in N$ be this agent.

- If for his most preferred object $\text{top}_j(O_{l-1})$ in O_{l-1}, $v_{j,\text{top}_j(O_{l-1})} > f_j$, then agent j obtains $\text{top}_j(O_{l-1})$ and pays f_j. Set $O_l := O_{l-1} \setminus \{\text{top}_j(O_{l-1})\}$. If $O_l = \emptyset$, then we stop and all remaining agents receive and pay nothing. Otherwise, continue.
- If $v_{j,\text{top}_j(O_{l-1})} \leq f_j$, then agent j receives and pays nothing. Set $O_l := O_{l-1}$ and continue.

We continue until either all real objects are allocated or all agents have been considered. We denote the resulting outcome by $\psi^{(f,\rhd)}(N, O, u)$.[7]

[7] Note that if the reservation prices are zero for all agents, we obtain the classical serial dictatorship

4 Characterization

Theorem 1. *A mechanism φ satisfies minimal tradability, individual rationality, strategy-proofness, consistency, independence of unallocated objects, neutrality, and non wasteful tie-breaking if and only if there exist a reservation price vector $r \in \mathcal{F}$ and a priority ordering $\succ \in \mathcal{P}$ such that φ is a serial dictatorship mechanism with reservation prices based on r and $\succ$, i.e., $\varphi = \psi^{(r,\succ)}$.*

We formally prove Theorem 1 in Appendix A. Below, we give a proof sketch to develop some intuition, and we discuss how our result relates to that of Klaus and Nichifor (2019, Theorem 1).

Recall that we extended the homogeneous indivisible objects model of Klaus and Nichifor (2019), their normative properties, and their class of serial dictatorship mechanisms with reservation prices to heterogeneous indivisible objects. For our characterization (Theorem 1), we use all the properties that are used by Klaus and Nichifor (2019, Theorem 1), suitably adapted from homogeneous to heterogeneous objects, to which we add one new key property: *neutrality*. In both results, the uniqueness proof consists of four parts; next, we sketch these parts, highlighting the role played by *neutrality*.

Proof Sketch (Uniqueness). We assume that φ satisfies all the properties in the theorem, and then proceed as follows:

1. we construct the individual reservation price vector $r \in \mathcal{F}$ (*neutrality* here implies that each agent's reservation price is the same for any real object);
2. we construct the priority ordering $\succ \in \mathcal{P}$ over $\mathbb{N}$ (*neutrality* here implies that the priority ordering is the same for any real object);
3. for single-object problems, by Klaus and Nichifor (2019, Proof of Theorem 1, Part 3) it follows that $\varphi = \psi^{(r,\succ)}$; and
4. we extend the single-object result that $\varphi = \psi^{(r,\succ)}$ to any set of real objects $O \in \mathcal{O}$ via an induction argument.

Parts 1 and 2 bear some resemblance to the corresponding proofs of the main result in Klaus and Nichifor (2019, Proof of Theorem 1, Parts 1 and 2). Some work has to be done to make sure that these proofs still work for the allocation of heterogeneous objects; the

mechanism. That is, given the reservation price vector $\mathbf{0} = (0, 0, \ldots) \in \mathcal{F}$ and a priority ordering $\triangleright \in \mathcal{P}$, $\psi^{(\mathbf{0},\triangleright)}$ is a *serial dictatorship mechanism.*

additional proof steps that require *neutrality* in each part are key, and entirely new. Part 4 is very different from the corresponding proof part in Klaus and Nichifor (2019, Proof of Theorem 1, Part 4) due to the fact that we deal with heterogeneous instead of homogeneous indivisible objects. □

The following examples present mechanisms that satisfy all the properties in Theorem 1, except for the one in the title of the example.

Example 1 (***Minimal Tradability***)
The *no-trade mechanism* never allocates any real object and no payments are made.

Example 2 (***Individual Rationality***)
Fix a positive price $P > 0$ and assign objects sequentially at price $P > 0$ to the agents with the lowest indices within the set of agents who in this manner can receive a real object with a valuation larger than P, until we run out of objects or agent; all remaining agents receive the null object and, except agent 1, pay nothing; if agent 1 is present, he pays price P (even if he did not receive a real object).

Example 3 (***Strategy-Proofness***)
We assign objects sequentially to the agents with the lowest indices within the set of agents who in this manner can receive an object with a positive valuation, until we run out of objects or agents; agents who obtain an object pay half their valuation, all remaining agents receive the null object and pay nothing.

Example 4 (***Consistency***)
Let $f \in \mathcal{F}$ and $\rhd, \rhd' \in \mathcal{P}$ such that $\rhd \neq \rhd'$. We apply $\psi^{(f,\rhd)}$ to problems $\gamma \in \Gamma(N, O)$ where the set of agents N has cardinality 2 and $\psi^{(f,\rhd')}$ otherwise.

Example 5 (***Independence of Unallocated Objects***)
Let $f \in \mathcal{F}$ and $\rhd \in \mathcal{P}$. We apply $\psi^{(f,\rhd)}$ to each problem in which there are weakly less objects than agents who want them, and a no-trade mechanism under which no objects are allocated and no payments are made, otherwise.[8]

[8] We say that an agent wants an object if his valuation is higher than his reservation price for it. The cases (i) weakly less objects than agents who want them, and (ii) more objects than agents who want them, are mutually exclusive and exhaustive. Note that removing agents together with their allotments cannot switch a problem between cases (i) and (ii); thus, *consistency* is satisfied. Meanwhile, removing unallocated objects may switch a problem from (ii) to (i); thus, *independence of unallocated objects* is violated.

Example 6 (***Neutrality***)
Let o be one of the objects and let i and j be two distinct agents. Let $f \in \mathcal{F}$ be a reservation price vector. Let $\rhd, \rhd' \in \mathcal{P}$ be two distinct priority orderings such that: agent i has the highest priory under $\rhd$ and the second highest priority under $\rhd'$; agent j has the highest priory under $\rhd'$ and the second highest priority under $\rhd$; for all agents other than i and j, their priorities under $\rhd, \rhd'$ are the same. We apply $\psi^{(f,\rhd)}$ to problems in which agent i and object o are available, and object o is i's best real object $\text{top}_i(O)$; and $\psi^{(f,\rhd')}$, otherwise.[9]

Example 7 (***Non Wasteful Tie-Breaking***)
Consider a modification of our serial dictatorship mechanism with reservation prices in which an agent who is indifferent between [not receiving the null object and not paying anything] and [receiving his most preferred available real object and paying his reservation price] receives his most preferred available real object.

5 Discussion

We illustrate how our model, properties, and results can be used for the problem of allocating specialist-led medical appointments by means of a simple stylized example that aims to capture the main features of the provision of specialist-led medical services in Australia.

Consider a health care system in which seeing a general practitioner (GP) is free, but a specialist-led consultation or procedure in a hospital (SLCP) may attract an out-of-pocket cost; also known as "gap payment," it is the difference between the price of the SLCP and what the insurance covers. Although the cost C charged by a hospital for some well-defined standard SLCP (e.g., a mole removal surgery) is required to be the same for any two patients i and j, differences in agents' insurance cover, e.g., $c_i \leq C$ and $c_j \leq C$, translate to differences in out-of-pocket costs. That is, the prices to be paid by i and j are $p_i = C - c_i$ and $p_j = C - c_j$, and we interpret them as agents' idiosyncratic non-negative *reservation prices*.[10]

[9] To see that *strategy-proofness* is satisfied, note that the different priorities $\rhd$ and $\rhd'$ only matter when agents i and j have the same best (acceptable) object: if that best object is o, then agent i receives it and he has no incentive to misrepresent his preferences; otherwise, if that best object is different from o, then agent i receives his second best acceptable object or the null object (and again, he has no incentive to misrepresent his preferences). The case of agents i and j having the same best (acceptable) object, object o versus an object $o' \neq o$, also illustrates why *neutrality* is violated.

[10] Since under normal circumstances it is never the case that a patient is paid to receive medical treatment, requiring non-negative prices is natural. More generally, our non-negativity assumption for prices is the equivalent of the common requirement in the mechanism design and auction literature that the mechanism generates no deficit for the seller, i.e., there are no subsidies.

Unless there is an immediate life threatening emergency, patients cannot directly go to the hospital; instead, they first have to see a GP. When a GP is unable to treat a patient because a SLCP is required, the GP establishes the patient's clinical need, and based on it refers him to the hospital (e.g., a GP who determines that i's mole is potentially cancerous, while j's is "only" a cosmetic issue, refers both i and j for a mole-removal surgery SLCP, but assigns to i a higher *priority* than to j). In turn, the hospital is required to interpret the referrals received from the GPs as exogenously given *priorities* that the patients have for SLCPs, and to allocate scarce SLCPs based on these priorities.

Due to capacity constraints (e.g., a SLCP requires suitably qualified medical staff and adequately equipped rooms), a hospital offers similar SLCPs at different times, and patients typically view any two such offers as distinct. For instance, for many patients a SLCP is usually more valuable tomorrow than in a month from now; additionally, each patient's work and family schedule induces preferences over time slots. To capture this idea, we consider k similar SLCPs and a set of discrete time slots T and we associate SLCPs to time-slots to induce a set of heterogeneous objects $\mathcal{O} \subseteq k \times T$; each patient i may receive at most one object $o \in \mathcal{O}$, and if he does receive it he pays his reservation price p_i for it. As due to wealth inequality agents may not value money equally and utility comparisons across agents may not be possible, agents' utility functions over objects and payments are general, not necessarily quasilinear. Since income redistribution is not feasible, and since each SLCP when allocated specifies the name of the patient, we require that agents cannot trade objects or make transfers among themselves.

The mechanisms that a hospital may use to allocate scarce appointments are constrained by normative criteria set in charters that govern the provision of medical services. Allowable mechanisms are required to be *just* and *fair* by respecting the *priorities* induced by clinical prioritization, while still maintaining compatibility with some *payments*.[11] Mechanisms that avoid "sophisticated" patients being able to misreport their utility for an object and in doing so gaining an advantage over "unsophisticated" patients are interpreted as leveling the playing field and ensuring "equality of access" (*strategy-proofness*). If some patients who did not receive an SLCP leave the queue, possibly because they no longer need it, or if some patients take their earlier SLCP, then the outcome for everyone else has to remain unchanged (*consistency*).[12] Finally, it is natural to require that a relabelling of the objects (e.g., through

[11] Note that such requirements rule out any auction-like mechanisms that would end up allocating SLCPs to the highest bidders.

[12] *Consistency* is required for the provision of specialist-led medical services in public hospitals in Australia (and in the UK), but it is of course not a universally valid requirement. For instance, a profit maximizing

a change in the appointment scheduling software) results in each agent being allocated the object which is the relabelled version of the object that he was previously allocated; however, since the out-of-pocket costs of a SLCP does not depend on when it is offered, the price paid has to remain the same (*neutrality*).[13]

In practice, the mechanism that the hospitals arrived at for allocating objects (i.e., the available SLCPs at various time-slots) is equivalent to the following procedure. The patient with the highest priority gets to choose his most preferred object in $\mathcal{O}$: If he takes it, he pays his reservation price; otherwise, he receives and pays nothing. The patient with the second highest priority gets to choose his most preferred object among the remaining ones: If he takes it, he pays his reservation price; otherwise, he receives and pays nothing. The hospital continues offering objects in this fashion until they run out, or until all referred patients have been considered, whichever comes first.

While our example above is stylized, our properties aim to capture the main requirements set in the charters that govern the provision of specialist-led consultation or procedures in hospitals in Australia.[14] From these properties, we show that a *priority order* over the set of patients and an *individual reservation price* for each patient can be derived, together with the *serial dictatorship mechanism with reservation prices* that is based on them. Perhaps not surprisingly, this mechanism mimics the procedure that the hospitals arrived at for the allocation of SLCPs in Australia: *first* the mechanism is intuitive and easy to implement; *second*, as our characterization is tight, this mechanism is the only possible one that satisfies all properties.

private hospital may not be required to respect any clinically determined priority order, or to charge everyone the same fee for the same procedure; instead, it may allocate the next-available appointment to j, not to i, if j can pay more – in such cases, it is of course natural to assume that payments might depend on other patients on the waiting list.

[13]Note that the remaining properties in our characterization are very mild and are satisfied by virtually all desirable mechanisms.

[14]For a closely related discussion and more details about the provision of SLCPs in Australia, see also Klaus and Nichifor (2019, Section 5 and 6).

Appendix

A Proof of Theorem 1

It is easy to see that any serial dictatorship mechanism with reservation prices induced by some reservation price vector $f \in \mathcal{F}$ and some priority ordering $\succ \in \mathcal{P}$ satisfies all the properties in the theorem.

For the uniqueness proof we assume that φ satisfies all the properties in the theorem; and, as announced in our proof sketch, we split the proof into four parts: first, we construct the individual reservation price vector $r \in \mathcal{F}$; second, we construct the priority ordering $\succ \in \mathcal{P}$ over $\mathbb{N}$; third, we prove that $\varphi = \psi^{(r,\succ)}$ for single object problems; fourth, we extend the result that $\varphi = \psi^{(r,\succ)}$ to any set of real objects $O \in \mathcal{O}$ via an induction argument.

Part 1: Individual Reservation Prices

For object allocation problems with one real object, Klaus and Nichifor (2019) have established the following lemma.

Lemma 1 (Klaus and Nichifor, 2019, Lemma 3)**.** *Assume that mechanism φ satisfies minimal tradability, individual rationality, and strategy-proofness. Consider a real object $o \in \mathbb{O}$. Then, for each agent $i \in \mathbb{N}$, there exists an individual reservation price $r_{i,o} \geq 0$ such that for each utility function $u_i \in \mathcal{U}(\{i\},\{o\})$ with associated valuation vector $v_i \in \mathcal{V}(\{i\},\{o\})$:*

(i) $v_{i,o} > r_{i,o}$ *implies* $\varphi_i(\{i\},\{o\},u_i) = (o, r_{i,o})$,

(ii) $v_{i,o} = r_{i,o}$ *implies* $\varphi_i(\{i\},\{o\},u_i) \in \{(0,0),(o,r_{i,o})\}$, *and*

(iii) $v_{i,o} < r_{i,o}$ *implies* $\varphi_i(\{i\},\{o\},u_i) = (0,0)$.

Next, we show that the individual reservation prices do not depend on which real object $o \in \mathbb{O}$ is used in the above lemma.

Lemma 2. *Assume that mechanism φ satisfies minimal tradability, individual rationality, and strategy-proofness. Then, for each agent $i \in \mathbb{N}$ and any two real objects $o, \hat{o} \in \mathbb{O}$,*

$$r_i := r_{i,o} = r_{i,\hat{o}},$$

where r_i is agents i's reservation price.

Proof. Assume that mechanism φ satisfies all the properties in the lemma. Let $i \in \mathbb{N}$ and consider $o, \hat{o} \in \mathbb{O}$. If $o = \hat{o}$, then $r_{i,o} = r_{i,\hat{o}}$ follows trivially. Hence assume that $o \neq \hat{o}$. Next, choose an (auxiliary) agent $j \in \mathbb{N}$, $i \neq j$. We will specify both agents allotment but note that only that of agent i matters for our proof.

Let $N = \{i, j\}$ and $O = \{o, \hat{o}\}$. By *minimal tradability*, there exist $u = (u_i, u_j) \in \mathcal{U}(N, O)$ with associated valuation vector $v = (v_i, v_j) \in \mathcal{V}(N, O)$ such that both real objects o and $\hat{o}$ are allocated. Without loss of generality, assume $\alpha_i(N, O, u) = o$ and $\alpha_j(N, O, u) = \hat{o}$. Then, by *consistency* and Lemma 1 (i) and (ii),

$$\varphi_i(\{i, j\}, \{o, \hat{o}\}, u) = \varphi_i(\{i\}, \{o\}, u_i) = (o, r_{i,o})$$

and

$$\varphi_j(\{i, j\}, \{o, \hat{o}\}, u) = \varphi_j(\{j\}, \{\hat{o}\}, u_j) = (\hat{o}, r_{j,\hat{o}}),$$

respectively. In particular, note that

$$\pi_i(N, O, u) = r_{i,o} \text{ and } \pi_j(N, O, u) = r_{j,\hat{o}}. \tag{1}$$

Consider the relabelling $\sigma \in \mathcal{S}(O)$ such that $\sigma(o) = \hat{o}$, $\sigma(\hat{o}) = o$, and $\sigma(0) = 0$; note that since there are only two real objects, σ is the only possible relabelling. By *neutrality*,

$$\alpha_i(N, O, u^\sigma) = \hat{o} = \sigma(o) = \sigma(\alpha_i(N, O, u))$$

and

$$\alpha_j(N, O, u^\sigma) = o = \sigma(\hat{o}) = \sigma(\alpha_j(N, O, u)).$$

Then, by *consistency* and Lemma 1 (i) and (ii),

$$\varphi_i(\{i, j\}, \{o, \hat{o}\}, u^\sigma) = \varphi_i(\{i\}, \{\hat{o}\}, u_i^\sigma) = (\hat{o}, r_{i,\hat{o}})$$

and

$$\varphi_j(\{i, j\}, \{o, \hat{o}\}, u^\sigma) = \varphi_j(\{j\}, \{o\}, u_j^\sigma) = (o, r_{j,o}),$$

respectively. In particular, note that

$$\pi_i(N, O, u^\sigma) = r_{i,\hat{o}} \text{ and } \pi_j(N, O, u^\sigma) = r_{j,o}. \tag{2}$$

By *neutrality*, we must also have

$$\pi_i(N,O,u^\sigma)=\pi_i(N,O,u) \text{ and } \pi_j(N,O,u^\sigma)=\pi_j(N,O,u),$$

which by (1) and (2) implies that

$$r_{i,o}=r_{i,\hat{o}} \text{ and } r_{j,\hat{o}}=r_{j,o}.$$

Since agent j and real objects o and $\hat{o}$ were arbitrarily chosen, it follows that agent i has a unique reservation price $r_i(=r_{i,o}=r_{i,\hat{o}})$. □

By our next lemma, for any problem, if an agent receives a real object, then his valuation has to be weakly larger than his individual reservation price (which also equals his payment); otherwise, his payment is necessarily null.

Lemma 3. *Assume that mechanism φ satisfies minimal tradability, individual rationality, strategy-proofness, consistency, independence of unallocated objects, and neutrality. Then, for each $(N,O)\in\mathcal{N}\times\mathcal{O}$, each $\gamma\in\Gamma(N,O)$ with associated valuation vector $v\in\mathcal{V}(N,O)$, each $i\in N$, and each $o\in O$, if $\alpha_i(\gamma)=o$, then $\pi_i(\gamma)=r_i\leq v_{i,o}$ (with r_i as obtained in Lemma 2).*

Proof. Assume that mechanism φ satisfies all the properties in the lemma. Let $(N,O)\in\mathcal{N}\times\mathcal{O}$, and $\gamma=(N,O,u)\in\Gamma(N,O)$ with associated valuation vector $v\in\mathcal{V}(N,O)$. Let $i\in N$, $o\in O$, and $\alpha_i(\gamma)=o$. If all agents but agent i leave with their allotments, then the reduced problem is $\gamma_{\{i\}}=(\{i\},O_{\{i\}},u_i)$ where $O_{\{i\}}=O\setminus\bigcup_{j\in N\setminus\{i\}}\alpha_j(\gamma)$. By *consistency*, $\varphi_i(\gamma_{\{i\}})=\varphi_i(\gamma)$ and $\alpha_i(\gamma_{\{i\}})=\alpha_i(\gamma)=o$. If $O_{\{i\}}=\{o\}$, then $\gamma_{\{i\}}=(\{i\},\{o\},u_i)$. If $\{o\}\subsetneq O_{\{i\}}$, then using *independence of unallocated objects*, we obtain $\varphi_i(\gamma_{\{i\}})=\varphi_i(\{i\},\{o\},u_i)$.

Thus, $\varphi_i(\{i\},\{o\},u_i)=\varphi_i(\gamma)$ and $\alpha_i(\{i\},\{o\},u_i)=\alpha_i(\gamma)=o$. By Lemma 1, $v_{i,o}\geq r_i$ and $\varphi_i(\{i\},\{o\},u_i)=(o,r_i)=\varphi_i(\gamma)$. In particular, $\pi_i(\gamma)=r_i\leq v_{i,o}$. □

Part 2: Priority Ordering

For object allocation problems with one real object, Klaus and Nichifor (2019, Proof of Theorem 1, Part 2) have shown that for any $o\in\mathbb{O}$, there exists a (transitive) priority ordering $\succ_o\in\mathcal{P}$ over $\mathbb{N}$. More specifically, for any two agents $i,j\in\mathbb{N}$ and any real object $o\in\mathbb{O}$, let $N=\{i,j\}$ and fix a utility profile $(\bar{u}_i,\bar{u}_j)\in\mathcal{U}(N,\{o\})$ such that the associated valuation vector $\bar{v} = (r_i + 1, r_j + 1)$. Then,

$$i \succ_o j \text{ if and only if } \alpha_i(N, \{o\}, (\bar{u}_i, \bar{u}_j)) = o. \quad (3)$$

Next, we show that the priority ordering over agents does not depend on which real object $o \in \mathbb{O}$ is considered, i.e., we show that for two distinct real objects $o, \hat{o} \in \mathbb{O}$,

$$\succ_o = \succ_{\hat{o}} .$$

Specifically, we show that for any two agents $i, j \in \mathbb{N}$, $i \neq j$, $i \succ_o j$ if and only if $i \succ_{\hat{o}} j$. So assume that $i \succ_o j$.

Let $N = \{i, j\}$ and consider the problem where both objects are available, i.e., $O = \{o, \hat{o}\}$. Assume further that both agents would like to have object o while they are not interested in object $\hat{o}$, i.e., the utility profile $u \in \mathcal{U}(N, O)$ is such that $u = (u_i, u_j) = \big((\bar{u}_{i,o}, u^0_{i,\hat{o}}), (\bar{u}_{j,o}, u^0_{j,\hat{o}})\big)$ with associated valuation vector $v = ((r_i + 1, 0), (r_j + 1, 0)) \in \mathcal{V}(N, O)$. We consider problem (N, O, u).

First, we show that object $\hat{o}$ is not allocated. Suppose, by contradiction, that for $x \in N$, $\alpha_x(N, O, u) = \hat{o}$. Since $v_{x,\hat{o}} = 0$ and prices are non-negative, by *individual-rationality* (IR2), $\pi_x(N, O, u) = 0$. By *non-wasteful tie-breaking*, $\alpha_x(N, O, u) = 0$ (otherwise, if $\alpha_x(N, O, u) = \hat{o}$, we would have $u_x(\hat{o}, \pi_x(N, O, u)) = u_x(0, 0)$; a contradiction). Hence, for both agents $x \in N$, $\alpha_x(N, O, u) \neq \hat{o}$.

Second, we show that object o has to be allocated to agent i. When removing unallocated object $\hat{o}$ from problem (N, O, u) we obtain problem $(N, \{o\}, u_{\{o\}})$; by *independence of unallocated objects*, we have that $\alpha_i(N, O, u) = \alpha_i(N, \{o\}, u_{\{o\}})$. Since problem $(N, \{o\}, u_{\{o\}}) = (N, \{o\}, (\bar{u}_i, \bar{u}_j))$, by (3), $i \succ_o j$ implies $\alpha_i(N, O, u) = o$.

Third, consider the relabelling $\sigma \in \mathcal{S}(O)$ such that $\sigma(o) = \hat{o}$, $\sigma(\hat{o}) = o$, and $\sigma(0) = 0$; note that since there are only two real objects, σ is the only possible relabelling. The relabelling of the utility profile is $u^\sigma = (u^\sigma_i, u^\sigma_j) = \big((u^0_{i,\hat{o}}, \bar{u}_{i,o}), (u^0_{j,\hat{o}}, \bar{u}_{j,o})\big) \in \mathcal{U}(N, O)$ with associated valuation vector $((0, r_i + 1), (0, r_j + 1)) \in \mathcal{V}(N, O)$. By *neutrality*,

$$\alpha_i(N, O, u^\sigma) = \hat{o} = \sigma(o) = \sigma(\alpha_i(N, O, u))$$

and

$$\alpha_j(N, O, u^\sigma) = 0 = \sigma(0) = \sigma(\alpha_j(N, O, u)),$$

which by (3) implies $i \succ_{\hat{o}} j$.

Thus, for any distinct agents $i, j \in \mathbb{N}$ and any distinct real objects $o, \hat{o} \in \mathbb{O}$, if $i \succ_o j$, then $i \succ_{\hat{o}} j$, which implies that $\succ_o = \succ_{\hat{o}}$.

We denote the unique priority ordering over agents by $\succ$.

Part 3: Single-Object Problems

For object allocation problems with one real object, Klaus and Nichifor (2019, Proof of Theorem 1, Part 3) have shown that φ always assigns the object and payments as if it is a serial dictatorship mechanism based on $r \in \mathcal{F}$ (from Part 1) and $\succ \in \mathcal{P}$ (from Part 2).

Part 4: An Arbitrary Set of Real Objects

We now show by induction on the number of objects that $\varphi = \psi^{(r,\succ)}$ for the general domain of all problems.

Induction Basis $|O| = 0, 1$: Let $N \in \mathcal{N}$, $O \in \mathcal{O}$, and $\gamma = (N, O, u) \in \Gamma(N, O)$ such that $|O| = 0, 1$. Then, $\varphi(\gamma) = \psi^{(r,\succ)}(\gamma)$ follows for $|O| = 0$ by *individual rationality* (IR1) and for $|O| = 1$ by Part 3.

Induction Hypothesis $|O| \leq k$: On the subdomain of problems where at most $k \geq 1$ real objects are available, we assume $\varphi = \psi^{(r,\succ)}$.

Induction Step $|O| = k + 1$: We show that for problems where $k + 1$ real objects are available, we have $\varphi = \psi^{(r,\succ)}$. Let $\varphi = (\alpha, \pi)$ and $\psi^{(r,\succ)} = (\alpha', \pi')$.

Consider a set of agents $N \in \mathcal{N}$, a set of real objects $O \in \mathcal{O}$ such that $|O| = k + 1$, and a utility profile $u \in \mathcal{U}(N, O)$ with associated valuation vector $v \in \mathcal{V}(N, O)$. If no agent in N would like to receive a real object at problem (N, O, u), i.e., if for all $i \in N$ and $o \in O$, $v_{i,o} \leq r_i$, then by Lemma 3 and *non-wasteful tie-breaking*, for all $i \in N$, $\varphi_i(N, O, u) = (0, 0) = \psi^{(r,\succ)}(N, O, u)$. Hence, assume that for some agent $i \in N$ and some real object $o \in O$, $v_{i,o} > r_i$.

Without loss of generality assume that agent 1 is the highest priority agent in N according to $\succ$ such that for some real object $o \in O$, $v_{1,o} > r_1$. Let $\hat{o} := \text{top}_1(O)$. Thus, $\psi_1^{(r,\succ)}(N, O, u) = (\hat{o}, r_1)$ and in particular, $\alpha'_1(N, O, u) = \hat{o}$. Assume, by contradiction, that $\alpha_1(N, O, u) \neq \hat{o}$.

Consider a utility function $u'_1 = \big(u_{1,\hat{o}}, (u^{-\infty}_{1,o'})_{o' \in O \setminus \{\hat{o}\}}\big) \in \mathcal{U}(\{1\}, O)$ with valuation vector $v'_1 = (v_{1,\hat{o}}, -\infty, \ldots, -\infty)$, i.e., $v'_1 = v_{1,\hat{o}}$ and for all $o \in O \setminus \{\hat{o}\}$, $v'_{1,\hat{o}} = -\infty$. Let $u' = (u'_1, u_{-1})$.

Then, by *strategy-proofness*, $\alpha_1(N, O, u') \neq \hat{o}$ and $\alpha'_1(N, O, u') = \hat{o}$. Additionally, using Lemma 3, $\alpha_1(N, O, u') = 0$.

Case 1. There exists an agent $i \in N \setminus \{1\}$ such that $\alpha_i(N, O, u') = \hat{o}$.

Recall that $\alpha_1(N, O, u') = 0$. Hence, when all agents except agents 1 and i leave with their allotments, we obtain the reduced problem $(\{1, i\}, O', u'_{\{1,i\}})$ where $\hat{o} \in O' \subseteq O$. By *consistency*, we then have

$$\alpha_1(N, O, u') = \alpha_1(\{1, i\}, O', u'_{\{1,i\}}).$$

Note that only object $\hat{o}$ is allocated. Hence, when removing all unallocated objects from reduced problem $(\{1, i\}, O', u'_{\{1,i\}})$, we obtain the problem $(\{1, i\}, \{\hat{o}\}, u'_{\{1,i\}})$. By *independence of unallocated objects*, we then have

$$\alpha_1(\{1, i\}, O', u'_{\{1,i\}}) = \alpha_1(\{1, i\}, \{\hat{o}\}, u'_{\{1,i\}}).$$

In particular in follows that

$$\alpha_1(\{1, i\}, \{\hat{o}\}, u'_{\{1,i\}}) = \alpha_1(N, O, u') = 0,$$

contradicting the Induction Basis (since for problems with one real object, agent 1 as the highest priority agent who wants the object should receive it).

Case 2. For all agents $i \in N \setminus \{1\}$, $\alpha_i(N, O, u') \neq \hat{o}$.

Recall that $\alpha_1(N, O, u') = 0$. Hence, when all agents except agent 1 leave with their allotments, we obtain the reduced problem $(\{1\}, O'', u'_1)$ where $\hat{o} \in O'' \subseteq O$. By *consistency*, we then have

$$\alpha_1(N, O, u') = \alpha_1(\{1\}, O'', u'_1).$$

We now remove the set of unallocated real objects $O'' \setminus \{\hat{o}\}$ and obtain the problem $(\{1\}, \{\hat{o}\}, u'_1)$. By *independence of unallocated objects*, we then have

$$\alpha_1(\{1\}, O'', u'_1) = \alpha_1(\{1\}, \{\hat{o}\}, u'_1).$$

In particular in follows that

$$\alpha_1(\{1\}, \{\hat{o}\}, u'_1) = \alpha_1(N, O, u') = 0,$$

contradicting the Induction Basis (since for problems with one real object agent 1 as the highest priority agent who wants the object should receive it).

Cases 1 and *2* now imply that $\alpha_1(N, O, u) = \alpha_1(N, O, u') = \hat{o}$. Recall that $\alpha'_1(N, O, u) = \hat{o}$. Hence, when agent 1 leaves problem (N, O, u) with his allotment under both mechanisms φ and $\psi^{(r,\succ)}$, we obtain the reduced problem $(N \setminus \{1\}, O_{N\setminus\{1\}}, u_{N\setminus\{1\}})$. By *consistency*, for all $i \in N \setminus \{1\}$, we then have

$$\varphi_i(N, O, u) = \varphi_i(N \setminus \{1\}, O_{N\setminus\{1\}}, u_{N\setminus\{1\}})$$

and

$$\psi_i^{(r,\succ)}(N, O, u) = \psi_i^{(r,\succ)}(N \setminus \{1\}, O_{N\setminus\{1\}}, u_{N\setminus\{1\}}).$$

By the Induction Hypothesis, for all $i \in N \setminus \{1\}$, we have

$$\varphi_i(N \setminus \{1\}, O_{N\setminus\{1\}}, u_{N\setminus\{1\}}) = \psi_i^{(r,\succ)}(N \setminus \{1\}, O_{N\setminus\{1\}}, u_{N\setminus\{1\}}).$$

Together with $\varphi_1(N, O, u) = \psi_1^{(r,\succ)}(N, O, u) = (\hat{o}, r_1)$ (by Lemma 3), this completes the proof that $\varphi(N, O, u) = \psi^{(r,\succ)}(N, O, u)$. □

References

Dold, M. and M. Khadjavi (2017). Jumping the queue: An experiment on procedural preferences. *Games and Economic Behavior 102*, 127–137.

Ehlers, L. and B. Klaus (2007). Consistent house allocation. *Economic Theory 30*(3), 561–574.

Ehlers, L. and B. Klaus (2014). Strategy-proofness makes the difference: Deferred-acceptance with responsive priorities. *Mathematics of Operations Research 39*(4), 949–966.

Ehlers, L. and B. Klaus (2016). Object allocation via deferred-acceptance: Strategy-proofness and comparative statics. *Games and Economic Behavior 97*, 128–146.

Ergin, H. (2000). Consistency in house allocation problems. *Journal of Mathematical Economics 34*(1), 77–97.

Klaus, B. and A. Nichifor (2019). Serial dictatorship mechanisms with reservation prices. *Economic Theory*, forthcoming, `https://ideas.repec.org/p/lau/crdeep/19.04.html`.

Kojima, F. and M. Manea (2010). Axioms for deferred acceptance. *Econometrica 78*(2), 633–653.

Konow, J. (2003). Which is the fairest one of all? A positive analysis of justice theories. *Journal of Economic Literature 41*(4), 1188–1239.

Sakai, T. (2013). Axiomatizations of second price auctions with a reserve price. *International Journal of Economic Theory 9*(3), 255–265.

Smith, J. H. (1973). Aggregation of preferences with variable electorate. *Econometrica 46*(6), 1027–1041.

Svensson, L.-G. (1994). Queue allocation of indivisible goods. *Social Choice and Welfare 11*(4), 323–330.

Svensson, L.-G. (1999). Strategy-proof allocation of indivisible goods. *Social Choice and Welfare 16*(4), 557–567.

Svensson, L.-G. and B. Larsson (2002). Strategy-proof and nonbossy allocation of indivisible goods and money. *Economic Theory 20*(3), 483–502.

Tadenuma, K. and W. Thomson (1991). No-envy and consistency in economies with indivisible goods. *Econometrica 59*(6), 1755–1767.

Thomson, W. (1983). The fair division of a fixed supply among a growing population. *Mathematics of Operations Research 8*(3), 319–326.

Thomson, W. (2015). *Consistent allocation rules.* Cambridge University Press. forthcoming.