Soc Choice Welf (2009) 33:505–515
DOI 10.1007/s00355-009-0376-6

ORIGINAL PAPER

Strategy-proof, efficient, and nonbossy quota allocations

John William Hatfield

Received: 7 August 2007 / Accepted: 4 February 2009 / Published online: 20 February 2009

Abstract We consider the problem of designing a mechanism to allocate objects to agents when each agent has a quota that must be filled exactly. Agents are assumed to have responsive preferences over items. We show that the only strategy-proof, Pareto optimal, and nonbossy mechanisms are sequential dictatorships. We also show that the only strategy-proof, Pareto optimal, nonbossy, and neutral mechanisms are serial dictatorships. Since these negative results hold for responsive preferences, they hold for more general preferences as well.

1 Introduction

We consider the problem of allocating a fixed number of available items to each agent in a strategy-proof and Pareto optimal way. We show that the set of mechanisms that are strategy-proof, Pareto optimal, and nonbossy is the set of sequential dictatorships, even when we restrict preferences to be responsive.

Examples of this allocation problem include the allocating of (multiple) projects to employees, allocating items among heirs, allocating equipment time to scientists, and allocating different tutorial sessions to students. The most prominent example of this problem is allocating amateur players to professional sports teams, and the usual solution to this problem is a draft, where agents choose one item they want each round.

We are grateful to Atila Abdulkadiroğlu, Syed Nageeb Mustafa Ali, Douglas Bernheim, Navin Kartik, Jon Levin, Paul Milgrom, and Azeem Shaikh and two anonymous referees and an editor for helpful discussions. This research has been supported by the Burt and Deedee McMurtry Stanford Graduate Fellowship.

J. W. Hatfield (✉)
Graduate School of Business, Stanford University, Stanford, USA
e-mail: hatfield@stanford.edu

However, it is easy to see that such drafting mechanisms are not strategy-proof, nor are they Pareto optimal, even if agents do tell the truth. They are not strategy-proof as agents will not necessarily draft their favorite players first—if an agent believes his favorite player will still be there next round, he may choose someone he sees as more popular this round, and wait until next round to choose his favorite player. A slightly more complicated example is required to see that drafts are not necessarily Pareto optimal, even when agents reveal truthfully. Consider a draft with two agents, each of whom obtains three objects. Both agents agree that $1 \succ 2 \succ 3 \succ 4 \succ 5 \succ 6$. Hence, the first agent will get (1, 3, 5) and the second will get (2, 4, 6). However, it may well be that the first agent would prefer (2, 3, 4) to (1, 3, 5) and the second agent would prefer (1, 5, 6) to (2, 4, 6) and so the allocation generated by the draft would not be Pareto optimal. Hence, for many problems, it behooves us to choose an allocation mechanism which does not suffer from these problems. Unfortunately, the only such mechanisms are sequential dictatorships.[1]

We already see sequential dictatorships being used when only one item is being allocated. While sequential dictatorships seem reasonable for these housing assignment problems, first introduced by Shapley and Scarf (1974), here they are deeply unsatisfactory. Consider the case with two agents and four items. If the agents have exactly the same preferences, any mechanism which is strategy-proof and Pareto optimal will give one agent his two favorite items, while giving the other agent the remainders. The second agent will likely claim that the mechanism is "unfair", that it should not give him the worst items while giving the other agent the best. And he has a point—the mechanism could allocate each of them one of their favorite two items.[2] Unfortunately, any mechanism which does this will be subject to strategic manipulation.

In some sense, we can think of our results as extending the well-known results for the housing assignment problem to the case where each agent receives more than one item. Svensson (2003) shows that the only housing allocation mechanism that is strategy-proof, nonbossy, and neutral is the serial dictatorship: we will show similar results for the case when each agent obtains $Q > 1$ objects. The flavor of all of these results is similar to that of the Gibbard–Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975). However, it is important to note that these theorems do not follow directly from Gibbard and Satterthwaite's work, as agents are assumed to be indifferent over the allocations of other agents in the type of problems considered here and in Svensson's work.

Further, we shall restrict the preferences of agents to be responsive: the change in utility from substituting one item for another depends only on those two items, not the other items the agent obtains. We use this model as it is the *most* restrictive set of preferences that can be regarded as realistic. Certainly, for many applications, substitutable or even broader sets of preferences may be applicable. However, since our goal is to show that every strategy-proof, Pareto optimal, and nonbossy mechanism is a sequential dictatorship, if we show this for our extremely restrictive set of

[1] We do not claim that sequential dictatorships should be used by sports leagues, which would be highly unbalancing. Rather, sports leagues use drafts, even though they have all of these problems, as the only mechanisms which do not, sequential dictatorships, would likely destroy the competitive balance each year.

[2] This sort of "fair" allocation is not possible when only allocating one object to each agent.

preferences, then it will still be true if we allow for more general preferences (since any opportunity for strategic manipulation that is available when preferences are more restricted is still available when larger classes of preferences are considered).

There are several papers which consider mechanisms that allocate multiple items to agents. The first paper, Papai (2001) considers a model where agents may have any preferences over sets of items, and looks for nonbossy, strategy-proof and Pareto optimal mechanisms; she finds the only such mechanisms are sequential dictatorships. Our paper generalizes this result as we consider a problem where the agents have a much smaller set of viable preferences; agents' preferences are responsive and they always desire Q items.[3] Hence, if we show the result for our much more restrictive set of preferences, it will still hold true for settings where more general preferences are allowed, as the opportunities for strategic manipulation have only been expanded. Ehlers and Klaus (2003) have generalized Papai's results to the case where agents may desire any number of objects and have responsive and separable preferences. As noted before, we also assume responsive preferences in this work. However, Ehlers and Klaus also assume that preferences are separable; that is, if the allocation of one particular item is preferred to nothing, then the agent prefers obtaining that item to not obtaining that item when holding any set of other items.

The next section lays out the basic model and formally defines the four conditions that we consider. The third section first proves some preliminary results, and then shows that the only strategy-proof, Pareto optimal, and nonbossy mechanism is a sequential dictatorship; adding the requirement of neutrality further restricts the set of satisfactory mechanisms to serial dictatorships. The last section concludes.

2 Model

We model the allocation problem as follows. There is a set $\mathbb{M} = \{1, \ldots, M\}$ of objects. There are N agents, each of which has a quota of $Q \geq 2$ objects which he must take.[4] Hence, for the problem to be feasible there must be $M \geq NQ$ objects available. An allocation of objects to agents will be represented by an allocation vector $\vec{x}$, where

$$x_n^m = \begin{cases} 1 & \text{if agent } n \text{ obtains object } m \\ 0 & \text{otherwise.} \end{cases}$$

Feasibility implies that

$$\sum_{n=1}^{N} x_n^m \leq 1$$

3 For many problems, this is the relevant domain: consider students choosing classes, teams drafting players, etc.

4 In principle, the quota could be agent-specific, and the results below without neutrality would still follow.

for all $m = 1, \ldots, M$ and

$$\sum_{m=1}^{M} x_n^m = Q$$

for all $n = 1, \ldots, N$. We shall refer to a particular agent n's allocation as $\vec{x}_n$, and the space of all feasible allocations for an agent as X_n. Hence, the set of feasible allocation vectors is $\prod_{n=1}^{N} X_n \equiv \mathcal{X}$.

Agents have (private) preferences over sets of objects. We shall model an agent n's preferences with a utility function of the following form:

$$u_n(\vec{x}_n) = \sum_{m=1}^{M} a_n^m x_n^m$$

where $a_n^m \in \mathbb{R}$ is the utility that agent n obtains from having object m as part of his consumption bundle. We further impose the restriction that preferences over sets of objects be strict.[5] If $u_n(\vec{x}_n) > u_n(\vec{y}_n)$ for $\vec{x}_n, \vec{y}_n \in X_n$, then we will write that $\vec{x}_n \succ_{u_n} \vec{y}_n$.[6,7] We can also represent an agent n's preferences by the vector $(a_n^1, \ldots, a_n^M)$. Let the set of admittable preferences for an agent n be defined as U_n, and then the set of admittable preference profiles for all agents is $\prod_{n=1}^{N} U_n \equiv \mathcal{U}$.

From the revelation principle (see e.g. Jackson 2003) we know that we only need to consider mechanisms where agents submit their preferences to the mechanism directly. This vastly simplifies the space of possible mechanisms.

Formally, a mechanism will be a function $F : \mathcal{U} \to \mathcal{X}$.[8] We shall also define the associated functions $f_n : \mathcal{U} \to X_n$, where

$$f_n(u) \equiv (F(u))_n$$

In words, $f_n(u)$ is the outcome for agent n when the stated preferences are u.

We shall consider four restrictions on our mechanism. The first is strategy-proofness, i.e. that no matter what preferences other agents submit, agent n has no incentive to lie about his preferences. Formally, then, a mechanism F is *strategy-proof*

[5] This restriction is not without loss: we do not consider the problem where students are contesting for finite numbers of identical seats within different classes. For a discussion of this problem, see Sönmez and Ünver (2007).

[6] We shall also occasionally abuse notation and write $i \succ_{u_n} j$ to mean that agent n prefers item i to item j under preference profile u , i.e. $a_n^i > a_n^j$.

[7] Since preferences over any two distinct bundles are strict, $\vec{x}_n \succsim_{u_n} \vec{y}_n$ can be read as stating that $\vec{x}_n$ is the same as or strictly preferred to $\vec{y}_n$.

[8] Note that we shall only consider deterministic mechanisms. Considering random mechanisms would lead to similar results, in that for any strategy-proof, Pareto optimal and nonbossy mechanism there would always exist one agent who recieved his favorite set of objects. While these mechanisms are not dictatorial in a strict sense, it does not solve the essential problem that one agent, even if randomly chosen, obtains all his favorite items while other agents are offered only what the first agent does not want.

if for all $u \in \mathcal{U}$, for all $n = 1, \ldots, N$, and for all $v_n \in U_n$, $f_n(u) \succsim_{u_n} f_n(v_n, u_{-n})$.[9] The second condition we will impose is Pareto optimality, that whatever allocation the mechanism produces, it should not be possible to find another, different, allocation where all agents are weakly better off. Formally, a mechanism F is *Pareto optimal* if for all $u \in \mathcal{U}$, there does not exist an allocation $\vec{x}$ such that $\vec{x}_n \succsim_{u_n} f_n(u)$ for all $n = 1, \ldots, N$ and $\vec{x}_n \succ_{u_n} f_n(u)$ for some $n = 1, \ldots, N$. The third condition is nonbossiness; a mechanism is nonbossy if when agent n submits different preferences, but his allocation does not change, then the overall allocation does not change.[10] Formally, a mechanism F is *nonbossy* if for all $u \in \mathcal{U}$, for all $n = 1, \ldots, N$, and all $v_n \in U_n$, if $f_n(u) = f_n(v_n, u_{-n})$, then $F(u) = F(v_n, u_{-n})$.

Finally, we shall prove a corollary for neutral mechanisms. Consider a permutation π of $1, \ldots, M$. Let $\pi(u)$ be the preference profile such that all agents permute their given scores according to π, i.e. $a_n^m \to a_n^{\pi(m)}$, and let $\pi(\vec{x})$ be the allocation vector where the allocation of all agents is permuted according to π, i.e. $\pi(x_n^m) = x_n^{\pi(m)}$. Formally, a mechanism F is *neutral* if for all permutations π of $1, \ldots, M$ and all $u \in \mathcal{U}$, $F(\pi(u)) = \pi(F(u))$.

3 Results

We first state some preliminary results. Lemma 1 guarantees that strategy-proof and nonbossy mechanisms satisfy monotonicity: i.e., when an agent changes his preferences from $u_n \in U_n$ to $v_n \in U_n$, the outcome of the mechanism does not change unless he states that some set of objects, dispreferred to $f_n(u)$ under u_n, is now preferred under v_n. If this is true for all agents when we switch from the preference profile u to the preference profile v, we have that by induction the allocation is unchanged.

Lemma 1 *Let F be a strategy-proof and nonbossy mechanism, and let $u, v \in \mathcal{U}$. Suppose that for all $\vec{x}_n \in X_n$ and all $n = 1, \ldots, N$, such that $f_n(u) \succsim_{u_n} \vec{x}_n$, we have that $f_n(u) \succsim_{v_n} \vec{x}_n$. Then $F(u) = F(v)$.*[11]

Our next lemma shows that any strategy-proof, Pareto optimal, nonbossy mechanism must, when agents submit identical preferences, allocate objects to the agents in a manner that resembles the outcome of a sequential dictatorship.

Lemma 2 *Let F be a strategy-proof, Pareto optimal, and nonbossy mechanism. Then, for each preference profile $u \in \mathcal{U}$ such that agents' preferences are identical, there exists for each $i = 1, \ldots, N$ an agent n_i such that this agent receives his* $[(i-1)Q+1]$th *through* iQth *choices.*

Proof Suppose not. Then there exists an identical preference profile $u \in \mathcal{U}$ and two agents, A(lice) and B(ob), who receive arbitrary objects a, b, c such that according to u, $a \succ_{u_i} b \succ_{u_i} c$ and A receives a and c while B receives b. We will now construct

[9] $(v_n, u_{-n}) \equiv (u_1, \ldots, u_{n-1}, v_n, u_{n+1}, \ldots, u_N)$.

[10] This condition was first introduced by Satterthwaite and Sonnenschein (1981).

[11] This lemma is a generalization of Lemma 1 in Svensson (2003). The proof follows exactly as in Svensson (2003).

alternative preferences v. For all agents $n \neq A, B$ (if any) their allocation under u is their favorite set of objects according to v. Furthermore, for A and B, let the other objects they receive besides a, b, c be their highest ranked objects, and let all other objects be ranked lower than a, b, c, and let a, b, c still be ranked $a \succ_{v_A} b \succ_{v_A} c$ and $a \succ_{v_B} b \succ_{v_B} c$. Hence, by Lemma 1, the allocation does not change; $F(u) = F(v)$.

We can represent their preferences over these three objects in a list as follows,

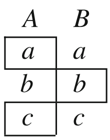

where higher listed objects are ranked higher. Boxed objects are ones the agent now obtains. Now

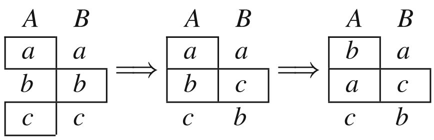

where the first implication follows from Pareto optimality (as A must get b) and strategy-proofness (as if B obtained a by submitting the second set of preferences, he would have an incentive to lie when his preferences were actually represented by the first list). The second implication follows from Lemma 1. However,

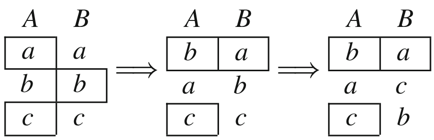

where the first implication follows from Pareto optimality (as A must get b) and strategy-proofness (as if A obtained a and b by submitting the second set of preferences, he would have an incentive to lie when his preferences were actually represented by the first list). The second implication follows from Lemma 1. However, this contradicts the outcome proscribed before. Hence, such agents can not exist.[12] □

We now consider how the order changes as a function of which identical preference profiles the agents submit. It is important to consider nonneutral mechanisms for this problem; neutrality is a much more restrictive assumption than in the housing assignment problem. For the housing assignment problem, it is known that allocating objects to agents randomly and then running a top trading cycles mechanism (a nonneutral mechanism) is equivalent to running a random serial dictatorship (see Sönmez and Ünver 2005). However, this is not obvious for our problem, as one could imagine mechanisms that assign "property rights" which have a better chance of allocating to both agents something they want. For instance, consider an example with

[12] The author thanks an anonymous referee for a simplifying suggestion for the proof.

$N = Q = 2$, and $M = 4$. Then a mechanism could allocate “property rights” to the first two objects to the first agent, and “property rights” to the last two objects to the second agent. Then, if both agents submitted preferences such that $1 \succ 3 \succ 2 \succ 4$, the first agent would get 1 and 2, and the second agent 3 and 4. This mechanism is fairer, in some sense, as no agent, even if preferences align exactly, always receives his favorite set of objects. However, any of these “property rights” mechanisms are subject to strategic manipulation, as we shall see.

We now consider the question of how the order of dictators can change when we consider different elements of $\mathcal{U}$, although we still only consider elements such that the preference profiles of the various agents are identical. We find that the order of dictators can change, but not in a very satisfactory way; the order can only rely on the objects already chosen. In particular, identity of the initial dictator is the same regardless of the choice of u, while the succeeding dictators' identities may depend on the choices of earlier dictators.

Lemma 3 *Let F be a strategy-proof, Pareto optimal, and nonbossy mechanism. Then there exists an agent* n_1*, and for each preference profile* $u \in \mathcal{U}$ *such that agents' preferences are identical, an ordering of remaining agents* $n_2 \equiv n_2(f_{n_1}(u)), \dots, n_N \equiv n_N(f_{n_1}(u), \dots, f_{n_{N-1}}(u))$ *such that for all* $i = 1, \dots, N$*,* $f_{n_i}(u)$ *consists of agent* n_i*'s* $[(i-1)Q+1]$th *through* iQth *choices according to the preference profile u.*

Proof We shall show that for any identical preference ordering, the same agent always chooses first. Consider the agent who gets his favorite set of objects for an identical ordering with preferences

$$1 \succ 2 \succ \cdots \succ Q-1 \succ Q \succ Q+1 \succ \cdots \succ M$$

(who must exist according to Lemma 2). Now consider the permutation of these preferences to

$$Q \succ 1 \succ 2 \succ \cdots \succ Q-1 \succ Q+1 \succ \cdots \succ M$$

for all agents. Then, by Lemma 1, the outcome does not change, and hence, under this permutation of preferences, the same agent receives his favorite objects. Now consider any object $m > Q$. The identical preference ordering

$$1 \succ 2 \succ \cdots \succ Q-1 \succ m \succ Q \succ \cdots \succ m-1 \succ m+1 \succ \cdots \succ M$$

has some agent n who receives his top Q choices (by Lemma 2). Now, let all other agents change their preferences back to the ordering $1 \succ 2 \succ \cdots \succ M$. Since they do not receive object m, by Lemma 1 their allocations must remain the same. Now, let agent n change his preferences back to the ordering $1 \succ 2 \succ \cdots \succ M$. By strategy-proofness, he can not lose objects $1, 2, \dots, Q-1$. Hence, since we know some agent obtains objects $1, 2, \dots, Q$, it must be this agent.

The above arguments show that inserting the Qth objects into the first slot in the identical preference ordering does not change the agent who receives the top Q objects, nor does inserting the mth ($m > Q$) object into the Qth position in the identical preference ordering. Hence, putting these together, we can raise any object to the 1st position in the preference ordering (leaving other elements fixed relative to each other). Applying this multiple times, we can create any identical preference ordering, and hence we have shown that there exists an agent that, for all identical preference orderings, receives his top Q choices.

By similar reasoning, holding the top Q elements of the preferences fixed (and hence the choices of n_1), there exists an agent who always receives his $[Q+1]$th through $2Q$th choices, and so on. □

We now characterize the set of mechanisms which are strategy-proof, Pareto optimal, and nonbossy. We shall show that all such mechanisms are sequential dictatorships. Formally, a mechanism F is a *sequential dictatorship* if there exists an agent n_1, and, for each preference profile $u \in \mathcal{U}$, an ordering of the remaining agents $n_2(f_{n_1}(u)), \ldots, n_N(f_{n_1}(u), \ldots, f_{n_{N-1}(f_{n_1}(u),\ldots,f_{n_{N-2}(\cdot)}(u))}(u))$ such that

$$f_{n_1}(u) = \arg \max_{\substack{Z\subseteq\mathbb{M}\\ |Z|=Q}} \left\{ \sum_{m\in Z} a^m_{n_1} \right\},$$

$$f_{n_2(f_{n_1}(u))}(u) = \arg \max_{\substack{Z\subseteq\mathbb{M}\setminus f_{n_1}(u)\\ |Z|=Q}} \left\{ \sum_{m\in Z} a^m_{n_2(f_{n_1}(u))} \right\},$$

and, recursively, for $i = 3, \ldots, N$ (if any)

$$f_{n_i\left(f_{n_1}(u),\ldots,f_{n_{i-1}(\cdot)}(u)\right)}(u)$$
$$= \arg \max_{\substack{Z\subseteq\mathbb{M}\setminus\left[f_{n_1}(u)\cup\bigcup_{j=2}^{i-1} f_{n_j\left(f_{n_1}(u),\ldots,f_{n_{j-1}(\cdot)}(u)\right)}(u)\right]\\ |Z|=Q}} \left\{ \sum_{m\in Z} a^m_{n_i\left(f_{n_1}(u),\ldots,f_{n_{i-1}(\cdot)}(u)\right)} \right\}.$$

In other words, for every preference profile $u \in \mathcal{U}$, agent n_1 (whose identity is fixed and does not depend on u) receives his Q most favorite objects, $n_2(f_{n_1}(u))$ (whose identity may depend on $f_{n_1}(u)$ and hence on u) receives his Q most favored objects from those remaining, etc.

Theorem 1 *The mechanism F is a strategy-proof, Pareto optimal, and nonbossy mechanism if and only if F is a sequential dictatorship.*

Proof It is clear that a sequential dictatorship is strategy-proof, Pareto optimal, and nonbossy.

We now show that any strategy-proof, Pareto optimal, and nonbossy mechanism F is a sequential dictatorship. Apply F to the subset of identical preference profiles $\mathcal{V} \subset \mathcal{U}$ to obtain, by Lemma 3, agent n_1 and for each $v \in \mathcal{V}$ and each $i = 2, \ldots, N$, $n_i \equiv n_i(f_{n_1}(v), \ldots, f_{n_{i-1}}(v))$ where, $f_{n_i}(v)$ consists of agent n_i's $[(i-1)Q+1]$th through iQth choices according to the identical preference profile v. Thus, given F we can identify uniquely an agent n_1, and, given any set of objects $Z_1 \subset \mathbb{M}$ where $|Z_1| = Q$ we can uniquely identify an agent

$$n_2 (Z_1) \tag{1}$$

and, in general, for $i = 3, \ldots, N$ (if any) given any sequence $Z_1, \ldots, Z_{i-1} \subset \mathbb{M}$ where $\left|Z_j\right| = Q$ for all $j = 1, \ldots, i-1$, we can uniquely identify an agent

$$n_i (Z_1, \ldots, Z_{i-1}) . \tag{2}$$

Now, let $u \in \mathcal{U}$ be an arbitrary (and not necessarily identical) preference profile. To show that there exists an ordering of the agents with the desirable properties we will first construct from u an identical preference profile $v^u \in \mathcal{V}$. The first best Q objects of v^u are the first best Q objects according to u_{n_1}, where n_1 was identified above from F; these objects are denoted by Z_1. The next Q objects (from $Q+1$ to $2Q$) are the first best Q objects from $\mathbb{M}\backslash Z_1$ according to $u_{n_2(Z_1)}$, where $n_2(Z_1)$ is defined in (1); these objects are denoted by Z_2. Finally, in general, for each $i = 3, \ldots, N$ (if any), given $Z_1, \ldots, Z_{i-1}$, the next Q objects (from $(i-1)Q+1$ to iQ) are the first best Q objects from the set $\mathbb{M}\backslash \bigcup_{j=1}^{i-1} Z_j$ according to $u_{n_i(Z_1,\ldots,Z_{i-1})}$, where $n_i(Z_1, \ldots, Z_{i-1})$ is identified in (2); these objects are denoted by Z_i. Now,

$$F\left(v^u\right) = F(u) .$$

The equality follows from Lemma 1, as for agent n_1, if $f_{n_1}(v^u) \succsim_{v^u_{n_1}} \vec{x}_{n_1}$ for any $\vec{x}_{n_1} \in X_{n_1}$, then $f_{n_1}(v^u) \succsim_{u_{n_1}} \vec{x}_{n_1}$; and for each $n_i(Z_1, \ldots, Z_{i-1})$ with $i = 2, \ldots, N$, if $f_{n_i(Z_1,\ldots,Z_{i-1})}(v^u) \succsim_{v^u_{n_i(Z_1,\ldots,Z_{i-1})}} \vec{x}_{n_i(Z_1,\ldots,Z_{i-1})}$ for any $\vec{x}_{n_i(Z_1,\ldots,Z_{i-1})} \in X_{n_i(Z_1,\ldots,Z_{i-1})}$, then $f_{n_i(Z_1,\ldots,Z_{i-1})}(v^u) \succsim_{u_{n_i(Z_1,\ldots,Z_{i-1})}} \vec{x}_{n_i(Z_1,\ldots,Z_{i-1})}$. Thus F is a sequential dictatorship. □

This shows that when preferences are strict and responsive, any strategy-proof, Pareto optimal and nonbossy mechanism must be a sequential dictatorship. Since this is a negative result, it also must hold for more general sets of preferences, such as when agents have more general preferences over objects.

Furthermore, all three of these conditions are necessary for the characterization. Consider an example with $Q = 2$, four objects $\{1, 2, 3, 4\}$, and two agents $\{A, B\}$. Consider a "draft", i.e. a mechanism that allocates A his most favored object, then gives B his most favored object of the remaining three, then A his most favored object of the remaining two, and B the last object. A may wish to change his ordering if his most favored object is not desired by B, but his second most favorite object is; hence the mechanism is not strategy-proof. However, this mechanism is Pareto optimal (with respect to the submitted preferences) and nonbossy.

An allocation mechanism which assigns a fixed allocation is strategy-proof and nonbossy, but clearly not Pareto optimal.

Finally, consider an example with $Q = 2$, eight objects $\{1, 2, 3, 4, 5, 6, 7, 8\}$, and four agents $\{A, B, C, D\}$. The mechanism gives A his two most favored objects, and

B his two most favored objects out of the six remaining when A is done. Then the mechanism allocates C his two most favored objects of the remaining four, unless $u_A = u_B$; in that case, the mechanism allocates D his two most favored objects of the remaining four. This mechanism is strategy-proof and Pareto optimal, but not nonbossy; B could change his preferences while not changing his allocation, but still change the allocation of C and D.

We can also show that neutrality further restricts the set of mechanisms to serial dictatorships. Formally, a mechanism F is a *serial dictatorship* if there exists an ordering of the agents $n_1, n_2, \ldots, n_N$ such that for all preferences $u \in \mathcal{U}$,

$$f_{n_1}(u) = \arg \max_{\substack{Z \subseteq \mathbb{M} \\ |Z|=Q}} \left\{ \sum_{m \in Z} a^m_{n_1} \right\}$$

and recursively, for $i = 2, \ldots, N$,

$$f_{n_i}(u) = \arg \max_{\substack{Z \subseteq \mathbb{M} \setminus \bigcup_{j=1}^{i-1} f_{n_j}(u) \\ |Z|=Q}} \left\{ \sum_{m \in Z} a^m_{n_i} \right\}.$$

In other words, for every preference profile $u \in \mathcal{U}$, agent n_1 receives his Q most favorite objects, agent n_2 receives his Q most favored objects from those remaining, etc. The key difference between serial dictatorships and sequential dictatorships is that under a serial dictatorship, the identity of the second, third, fourth, etc. dictators no longer relies on the choices of earlier dictators.

Corollary 1 *The mechanism F is a strategy-proof, Pareto optimal, nonbossy, and neutral mechanism if and only if F is a serial dictatorship.*

Proof It is clear that a serial dictatorship is strategy-proof, Pareto optimal, nonbossy and neutral. The reverse implication follows from Theorem 1. □

4 Conclusion

The above results have shown that the only acceptable mechanism for allocation problems of this sort is a sequential dictatorship, even when we restrict preferences to be responsive. Hence, broader classes of preferences, where agents may have substitutable or more complex preferences, or may demand varying numbers of objects, will allow an even greater chance for strategic manipulation if a sequential dictatorship is not used. (It is obvious that sequential dictatorships are strategy-proof and Pareto optimal even for the most general set of preferences.) There have been a number of studies of the difficulties of mechanism design for these types of problems, including recent work by Sönmez and Ünver (2007) on the difficulty of allocating business students places in oversubscribed classes. Although unfortunate, it seems that in many of these applications, the best procedure (even if it is not considered "fair") may well be a random serial dictatorship.

References

Ehlers L, Klaus B (2003) Coalitional strategy-proof and resource monotonic solutions for multiple assignment problems. Soc Choice Welfare 21:265–280

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

Jackson MO (2003) Mechanism theory. In: Ulrich D (ed) Optimization and operations research, in Encyclopedia of Life Support Systems. EOLSS Publishers, Oxford, http://eolss.net

Papai S (2001) Strategy-proof and nonbossy multiple assignments. J Public Econ Theory 3:257–271

Satterthwaite MA (1975) Strategy proofness and arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. J Econ Theory 10:187–216

Satterthwaite MA, Sonnenschein H (1981) Strategy-proof at differentiable points. Rev Econ Stud 48: 587–597

Shapley L, Scarf H (1974) On cores and indivisibility. J Math Econ 1:23–37

Sönmez T, Ünver MU (2005) House allocation with existing tenants: an equivalence. Games Econ Behav 52:153–185

Sönmez T, Ünver MU (2007) Course bidding at business schools. Int Econ Rev (forthcoming)

Svensson L-G (2003) Strategy-proof allocation of indivisible goods. Soc Choice Welf 16:557–567