Journal of Mechanism and Institution Design
ISSN: 2399-844X (Print), 2399-8458 (Online)
DOI:10.22574/jmid.2023.12.003

A MARKET DESIGN SOLUTION FOR MULTI-CATEGORY HOUSING ALLOCATION PROBLEMS

Gian Caspari

ZEW — Leibniz Centre for European Economic Research, Germany
gian.caspari@zew.de

ABSTRACT

We study multi-category housing allocation problems: A finite set of objects, which is sorted into categories of equal size, has to be allocated to a finite set of individuals, such that everyone obtains exactly one object from each category. We show that, in the large class of category-wise neutral and non-bossy mechanisms, any strategy-proof mechanism can be constructed by simply letting individuals choose an object from each category one after another following some priority order. We refer to these mechanisms as multi-category serial dictatorships and advocate for selecting priority orders across categories as fairly as possible.

Keywords: Matching, envy-free, multi-category housing allocation

JEL Classification Numbers: D44, D50.

1. INTRODUCTION

Consider the problem of allocating $m \times n$ objects to n individuals based on the individuals' reported preference information over objects, such that every individual obtains a bundle containing m objects. Different solutions to this problem have been proposed by both researchers and practitioners.

I declare that there are no conflicts of interest. I thank the editor and the referees for their time and effort. I am especially thankful to the anonymous referee, who helped to improve the clarity of the proofs for Lemma 1 and Theorem 1. Any errors are my own.

On one hand, the theoretical literature has focused almost exclusively on strategy-proof mechanisms.[1] In this line of research (serial and sequential) dictatorship mechanisms stand out, as they are the only mechanisms that are simultaneously strategy-proof and (Pareto) efficient (Pápai, 2001; Klaus & Miyagawa, 2002; Ehlers & Klaus, 2003; Hatfield, 2009; Monte & Tumennasan, 2015):
Under a dictatorship mechanism individuals are assigned their most preferred bundle of m objects — among the remaining objects — one after another following some choosing order.

On the other hand, practitioners have proposed mechanisms that focus less on strategy-proofness and more on balancing both fairness and efficiency of the resulting allocation. In this context, an important class of fair and approximately (Pareto) efficient mechanisms are so called the draft mechanisms:
Under a draft mechanism individuals are assigned their most preferred object — among the remaining objects — one after another following some choosing order, which is reversed in each subsequent round until everyone has obtained m objects.[2]

However, how do these two classes of mechanisms compare? Using data on individuals reported as well as true preferences for the Harvard Business School course allocation, Budish & Cantillon (2012) show that draft mechanisms are indeed manipulated in practice and that these manipulations cause meaningful welfare losses. At the same time, they also find that, despite their shortcomings, draft mechanisms outperform dictatorship mechanisms in terms of welfare.[3]

[1] Under a strategy-proof mechanism truthful reporting of preferences over objects is a (weakly) dominant strategy for the individuals — allowing them to avoid costly and risky strategic behavior (Roth, 2008).

[2] Note that both draft and dictatorship mechanisms work in the same way if there are only as many objects as individuals ($m = 1$). For such single-object allocation problems (Hylland & Zeckhauser, 1979), also known as housing allocation problems, dictatorship mechanisms are the natural candidates arising, while if there are predefined property rights (Shapley & Scarf, 1974), also known as housing markets, top trading cycles mechanisms (also known as core mechanisms) are used to find an allocation. For more details, see for example Sönmez & Ünver (2011).

[3] Budish & Cantillon (2012) give the following intuition for this result: Under a dictatorship mechanism, individuals who get to pick early make their last choices independently of whether these objects would be some later-picking individuals' first choices; Individuals "callously disregard" the preferences of those who choose after them. This matters for wel-

These findings suggest that we should look at mechanisms that, akin to dictatorship mechanisms, (i) are strategy-proof while ensuring that, analogous to the allocations produced by draft mechanisms, (ii) the objects are distributed more equally (fairly) and are approximately efficient. Unfortunately, there exists a trilemma for multi-object allocation problems: Any mechanism satisfies at most two out of the three desired properties of strategy-proofness, fairness, and efficiency — even for some sensible weakenings/approximations of these properties (Caspari, 2020).[4]

In this paper, our main contribution is to show a positive result for an important special case of multi-object allocation problems — multi-category housing allocation problems: A total of $n \times m$ indivisible objects, which are sorted into m categories containing n objects each, must be allocated to a set of n individuals, based on the individuals' reported preference information over objects, such that everyone obtains exactly one object from each of the m categories. That is, we show that the presence of categories is sufficient for the existence of strategy-proof mechanisms producing fair and (approximately) efficient allocations:

Starting with one category, individuals are assigned their most preferred object — among the remaining objects — one after another following some choosing order which is reversed in each subsequent round and category until everyone has m objects — one from each category.

Intuitively, letting individuals choose objects one after another — analogous to draft mechanisms — can be implemented in a strategy-proof manner, as restricting individuals to choose from a given category each round removes any potential gains from strategic behavior. At the same time, efficiency is not lost as individuals actually want to have an object from every category. As a practical application, we consider the problem of allocating teaching-assistant positions to graduate students, where everyone has to assist exactly one spring-semester and one fall-semester course. Monte & Tumennasan (2015) analyse

fare as the benefit to the early-picking individuals from these last choices will generally be small relative to the harm these choices cause to the later-picking individuals.

[4] One notable exception concerns large markets: Budish (2011) describes the approximate competitive equilibrium from equal incomes (ACEEI) mechanisms which are approximately efficient, fair, and strategy-proof if the size of the market makes participants price takers. ACEEI has been applied to course allocation problems (Sönmez & Ünver, 2010) and was successfully implemented at Wharton Business School (Othman et al., 2010; Budish et al., 2016).

a multi-category housing allocation problem with two categories, and discuss several practical applications, including benefit and assistance programs and the allocation of new physicians in the United Kingdom. Overall, depending on the application in mind, categories can either be interpreted as containing different types of objects, containing the same set of objects for different time periods, or some combination of both.

We formally introduce multi-category housing allocation problems in Section 2. Moreover, Section 2.1 introduces the necessary framework for individuals to report rankings over categories as opposed to their full preference relation over all possible bundles. Any particularities stemming from this model choice are discussed there.[5]

Section 3 contains our main theoretical result: Any strategy-proof, category-wise neutral, and non-bossy mechanism can be obtained by specifying a choosing order — referred to as a priority order — for each category. We refer to this class of mechanisms as multi-category serial dictatorships.[6]

Section 3.1 takes a look at two ways to select priority orders for multi-category serial dictatorships: On one hand, analogous to (serial) dictatorships, one can choose an identical order for each category — referred to as the subclass of identical priority multi-category serial dictatorships. Alternatively, analogous to the draft mechanism, one can select a priority order that is reversed in every other category, referred to as the subclass of fair priority multi-category serial dictatorships. As both are strategy-proof, we can solely compare these mechanisms in terms of fairness and efficiency. We show that identical priority multi-category serial dictatorships are (Pareto) efficient but extremely unfair (Proposition 1), while identical priority multi-category serial dictatorships achieve maximal fairness while being approximately efficient (Proposition 2). Moreover, we provide a discussion, that places Proposition 1 and 2 into the broader context of the literature on dictatorship mechanisms.

Section 4 discusses the implementation of a fair priority multi-category serial dictatorship to allocate spring-semester and fall-semester teaching positions to graduate students, while Section 5 concludes.

[5] The overall approach relates to Brams & Fishburn (2000), Brams et al. (2003), and Edelman & Fishburn (2001).

[6] The result can be seen as a generalization of a well-known characterization result by Svensson (1999) from housing allocation problems to multi-category housing allocation problems. Related, Monte & Tumennasan (2015) have shown that in multi-category housing allocation problems any strategy-proof, non-bossy, and Pareto efficient mechanism is a sequential dictatorship.

2. THE MULTI-CATEGORY HOUSING ALLOCATION PROBLEM

Let I be a finite set of $|I| = n$ **individuals**, O be a finite set of $|O| = m \times n$ **objects**, and K be a finite set of $|K| = m$ types or time periods. The set of objects can be partitioned into m different **categories** $(O^k)_{k \in K}$, each containing $|O^k| = n$ distinct objects of type k, or $|O^k| = n$ distinct objects from time period k, respectively.

Each individual has a **preference** relation $\succsim_i$ comparing all sets of objects that contain exactly one object from the same categories. Formally, let $\succsim_i$ be a partial order, such that we either have $O' \succsim O''$, $O'' \succsim O'$, or both if and only if $|O' \cap O^k| = |O'' \cap O^k| \leq 1$ for all $k \in K$.[7] Note that, we have implicitly assumed that preferences over objects (singleton sets) within each category are strict, while objects (singleton sets) from different categories cannot be directly compared with each other.

We want to distribute the available objects among the individuals such that every individual is assigned exactly one object from each category, and no two distinct individuals are assigned the same object. That is, a feasible **allocation** $A = (A_i)_{i \in I}$ assigns every individual $i \in I$ a set of objects A_i with $|A_i \cap O^k| = 1$ for all $k \in K$, and $A_i \cap A_j = \emptyset$ for all $i \in I$, $j \in I \setminus \{j\}$. Let $\mathscr{A}$ denote the set of all feasible allocations. Moreover, let $a_i^k \in O^k \cap A_i$ denote the object from allocation A in category O^k that is assigned to individual i.

We assume that preferences over allocations are **separable** in terms of categories: For all $A, A' \in \mathscr{A}$ and $i \in I$, if $\{a_i^k\} \succsim_i \{a_i'^k\}$ for all $k \in K$ then $A_i \succsim_i A_i'$, and if $\{a_i^k\} \succ_i \{a_i'^k\}$ for at least one $k \in K$ then $A_i \succ_i A_i'$. We denote the set of all separable preferences $\succsim_i$ for individual i by $\mathscr{Q}_i$.

As it is generally done, we assume there are no externalities, i.e., for all $i \in I$ we have that $A_i \succsim_i A_i'$ implies $A \succsim_i A'$. That is, any individual's preference over allocations solely depends on the object the individual is assigned.

[7] A partial order is reflexive, transitive, and antisymmetric binary relation. If either $O' \succsim O''$, $O'' \succsim O'$, or both we say that O' and O'' are comparable. Moreover, for any $O', O'' \subset O$ such that $O' \succsim_i O''$ but $O'' \not\succsim_i O'$, we say that O' is strictly preferred to O'' and write $O' \succ_i O''$. Finally, note that, anti-symmetry implies that for any $O', O'' \subset O$ such that $O' \succsim_i O''$ and $O'' \succsim_i O'$ we have $O' = O''$. In other words, preferences are strict whenever two comparable sets are not identical.

2.1. Rankings instead of Preferences

We assume that, instead of having to report their entire preferences, individuals simply need to report a separate **ranking** for each category. That is, for each $k \in K$, each individual $i \in I$ reports a transitive, asymmetric, and complete order P_i^k over O^k. We denote the associated transitive, antisymmetric, and strongly complete order by R_i^k.[8] The list containing all rankings of a single individual i is denoted by $P_i = (P_i^k)_{k \in K}$ and analogously the list containing all rankings of all individuals is denoted by $P = (P_i)_{i \in I}$. We let $P_{-i} = (P_j)_{j \in I \setminus \{i\}}$ denote the list containing all individuals' rankings, except individual i's rankings. Finally, the set of all possible lists of rankings for a single individual, all individuals, and all but one of the individuals are denoted by $\mathscr{P}_i$, $\mathscr{P}$, and $\mathscr{P}_{-i}$ respectively.

Abstracting away from truthful revelation of rankings for the moment, if P_i is reported we can narrow down the possible preference profiles i might have. That is, a preference $\succsim_i$ is **consistent** with the reported rankings P_i if it ranks objects in each category in the same way as the reported ranking. Formally, $\succsim_i \in \mathscr{Q}_i$ is consistent with P_i if for all $k \in K$, $o \in O^k$, $o' \in O^k$ we have $o \, R_i^k \, o'$ if and only if $\{o\} \succsim_i \{o'\}$. Let $\mathscr{Q}_{P_i} \subset \mathscr{Q}_i$ denote the subset of all separable preferences that are consistent with the reported rankings P_i.

Following the introduction of rankings, we now limit our attention to strategy-proof mechanisms that take a profile of rankings as input. Formally, a **ranking mechanism** $\psi : \mathscr{P} \to \mathscr{A}$ selects an allocation $A \in \mathscr{A}$ for any reported list of rankings $P \in \mathscr{P}$. Given a list of reported rankings P, we let $\psi(P)_i$ denote the set of objects obtained by individual i under mechanism ψ, and slightly abusing notation, we let $\psi(P)_i^k$ denote both the object as well as the singleton set containing the object obtained by individual i in category k under mechanism ψ. Moreover, a ranking mechanism is strategy-proof if an individual having preferences $\succsim_i \in \mathscr{Q}_{P_i}$ consistent with rankings P_i cannot benefit from reporting a different list of rankings $\hat{P}_i$ instead of P_i. That is, a ranking mechanism ψ is **strategy-proof** if for any $i \in I$, $P_i \in \mathscr{P}_i$, $\hat{P}_i \in \mathscr{P}_i$, $P_{-i} \in \mathscr{P}_{-i}$ we have

$$\psi(P)_i \succsim_i \psi(\hat{P}_i, P_{-i})_i \text{ for all } \succsim_i \in \mathscr{Q}_{P_i}.$$

[8] That is, unlike P_i^k which is asymmetric, R_i^k also compares any object in O^k with itself. Otherwise, both relations rank any two distinct objects in O^k in the same way. See, for example, Roberts (1985) for an overview on binary relations and their properties.

2.1.1. *Dominance Relation*

In this part, we discuss what we can infer about the preference of an individual based on her reported list of rankings. To fix ideas, consider any list of rankings P_i and two allocations A and A' such that for every object $a_i^k \in A_i$ of category O^k we can find a weakly lower ranked object from the same category in the other allocation $a_i'^k \in A_i'$. Note that, in this case A must be preferred to A' under any (separable) preference $\succsim_i \in \mathscr{Q}_{P_i}$ consistent with P_i.

We can generalize this idea by defining a partial order $\geq_{P_i}$ — referred to as a **dominance relation** — over the set of allocations $\mathscr{A}$ for any list of rankings P_i: Formally, fix any P_i, then for all $A, A' \in \mathscr{A}$ we have $A_i \geq_{P_i} A_i'$ if and only if $a_i^k \; R_i^k \; a_i'^k$ for all $k \in K$.[9] In a second step, we show the following result:

Lemma 1. *Fix any P_i and $A, A' \in \mathscr{A}$. We have that $A_i \succsim_i A_i'$ for all $\succsim_i \in \mathscr{Q}_{P_i}$ if and only if $A_i \geq_{P_i} A_i'$.*

Lemma 1 states that if an individual reports a list of rankings P_i then any separable preference $\succsim_i \in \mathscr{Q}_{P_i}$ that is consistent with P_i will order any allocations in the same way as the dominance relation $\geq_{P_i}$ constructed from the same list of rankings P_i. Simultaneously, if two allocations are not comparable by the dominance relation then not all separable preferences that are consistent with P_i will rank the two allocations in the same way.

The dominance relation together with Lemma 1 is necessary for characterizing the set of strategy-proof ranking mechanisms. Moreover, we use the dominance relation later to define a weaker efficiency notion.

3. CHARACTERIZATION OF STRATEGY-PROOF MECHANISMS

We now show that, under two additional mild requirements, all strategy-proof mechanisms can be constructed by simply picking a priority order for each category, i.e., choosing an order specifying the sequence in which individuals are assigned an object following their reported rankings. Formally, a priority order is a bijection $f : I \mapsto \{1, \ldots, n\}$, with a lower number $f(i)$ indicating a higher priority. For a given list of priority orders for every category $f = (f^k)_{k \in K}$, the **multi-category serial dictatorship** mechanism is defined for $\ell \in \{1, \ldots, n\}$ as follows:

[9] We use $A_i >_{P_i} A_i'$ to denote that $A_i \geq_{P_i} A_i'$ but $A_i' \not\geq_{P_i} A_i$. Similarly we use $A_i =_{P_i} A_i'$ whenever $A_i \geq_{P_i} A_i'$ and $A_i' \geq_{P_i} A_i$— in which case $A_i = A_i'$. Note that $\geq_{P_i}$ is a partial order, i.e., a *reflexive*, *antisymmetric*, and *transitive* binary relation.

Step ℓ. *For each $k \in \{1,\ldots,m\}$, consider the ℓth highest priority individual $(f^k)^{-1}(\ell)$ according to f^k.*[10] *Then, following $P_{(f^k)^{-1}(\ell)}$, assign individual $(f^k)^{-1}(\ell)$ her most preferred object in category O^k among the remaining objects.*

For the characterization result to go through, we are left with defining two requirements: Nonbossiness and category-wise neutrality. Nonbossiness requires that no individual can influence the allocation of another individual without affecting her own allocation, while category-wise neutrality requires that the mechanism is immune to a relabeling of the object within each category. First, a ranking mechanism ψ is **nonbossy** if for all $P_i, \hat{P}_i \in \mathscr{P}_i$, and $P_{-i} \in \mathscr{P}_{-i}$ we have

$$\psi(P)_i = \psi(\hat{P}_i, P_{-i})_i \implies \psi(P) = \psi(\hat{P}_i, P_{-i}).$$

Second, let $\pi : O \to O$ be a permutations s.t. if $o \in O^k$ then $\pi[o] \in O^k$ — with Π denoting the set of all such permutations. We permute a list of simple orders P, denoted by $\pi[P]$, as follows: For all $i \in I$, $k \in K$ and $o, o \in O^k$ we have $\pi[o]\ \pi[P_i^k]\ \pi[o']$ if and only if $o\ P_i^k\ o'$. We say a ranking mechanism ψ is **category-wise neutral** if for all $k \in K$, $i \in I$, and $\pi \in \Pi$ we have

$$\pi[\psi(P)_i^k] = \psi(\pi[P])_i^k.$$

Now we go through the characterization result step-by-step. First, let us state an alternative strategy-proofness definition, which we will use to proof the next lemma. Formally, a ranking mechanism ψ is **strongly strategy-proof** if for any $i \in I$, $P_i \in \mathscr{P}_i$, $\hat{P}_i \in \mathscr{P}_i$, $P_{-i} \in \mathscr{P}_{-i}$ we have

$$\psi(P)_i \geq_{P_i} \psi(\hat{P}_i, P_{-i})_i.$$

The following statement is a corollary of Lemma 1. Additionally, note that Corollary 1 holds for ranking mechanisms but not necessarily in general.

Corollary 1. *A ranking mechanism is strategy-proof if and only if it is strongly strategy-proof.*

[10] Since f^k is a bijection, the function f^k is invertible. That is, $(f^k)^{-1} : \{1,\ldots,n\} \mapsto I$ with $f^{-1}(\ell)$ giving the ℓth highest priority individual under f^k.

Second, Lemma 2 guarantees that if an individual changes her reported list of rankings from P_i to $\hat{P}_i$, the outcome of a strategy-proof and non-bossy ranking mechanism ψ cannot change unless some objects, ranked lower than those in the same category assigned under $\psi(P)_i$, are now ranked higher under $\hat{P}_i$. In other words, one can reorder objects in category O^k without affecting the allocation of strategy-proof and nonbossy ranking mechanisms, as long as for any object $o \in O^k$ such that $\psi(P)_i^k \, P_i^k \, o$ we have $\psi(P)_i^k \, \hat{P}_i^k \, o$. Lemma 2 is adapted from Svensson (1999). Using Lemma 1, in the proof we substitute strategy-proofness for strong strategy-proofness and use the fact that $A_i =_{P_i} A_i'$ implies $A_i = A_i'$.

Lemma 2. *Let ψ be a nonbossy and strategy-proof ranking mechanism. Consider any $P_i \in \mathscr{P}_i$, $P_{-i} \in \mathscr{P}_{-i}$, and some $\hat{P}_i \in \mathscr{P}_i$ such that for all $A \in \mathscr{A}$ where $\psi(P)_i \geq_{P_i} A_i$ we have $\psi(P)_i \geq_{\hat{P}_i} A_i$. Then $\psi(P) = \psi(\hat{P}_i, P_{-i})$.*

Third, Lemma 3 establishes that for identical rankings — all individuals submit an identical ranking in each category — any category-wise neutral mechanism can be obtained through a multi-category serial dictatorship.

Formally, the set of all **identical rankings** is defined as $\mathscr{I} = \{P \in \mathscr{P} : P_j^k = P_i^k \text{ for all } i, j \in I \text{ and } k \in K\}$.

Lemma 3. *Let ψ be a ranking mechanism that is category-wise neutral. For every identical ranking $P \in \mathscr{I}$, $k \in K$, and $\ell \in \{1, \dots, n\}$ the same individual $i_\ell^k \in I$ is assigned the ℓth-highest ranked object in O^k according to ranking P_i^k.*

Finally, it remains to be shown what happens for arbitrary rankings $P \in \mathscr{P}$. We will invoke Lemma 2 to show that for any arbitrary preference profile $P \in \mathscr{P} \setminus \mathscr{I}$ there exists an identical preference profile $P \in \mathscr{I}$ leading to the same outcome.

Theorem 1. *For any multi-category housing allocation problem a ranking mechanism ψ is strategy proof, nonbossy, and category-wise neutral if an only if ψ is a multi-category serial dictatorship.*

3.1. Two Subclasses

In the class of multi-category serial dictatorship mechanisms, two subclasses stand out. As the name suggests, the subclass of **identical priority multi-category serial dictatorships** ψ^{IPD} specify identical priority orders across all categories, i.e., $f = f^k$ for all $k \in K$. On the contrary, if priority orders are selected as fairly as possible — for all $i, j \in I$ we have $|\{k \in K : f^k(i) < f^k(j)\}| \geq \lfloor \frac{m}{2} \rfloor$ — we refer to this subclass as **fair priority multi-category serial dictatorships** ψ^{FPD}.

We show that, identical priority multi-category serial dictatorships are Pareto efficient. That is, a mechanism ψ is **Pareto efficient** if for all $P \in \mathscr{P}$,

$$\nexists A \in \mathscr{A} \text{ s.t. } A_i \succsim_i \psi(P)_i \text{ for all } i \in I \text{ and } A_i \succ_i \psi(P)_i \text{ for at least some } i \in I.$$

In comparison, fair priority multi-category serial dictatorship mechanisms are not Pareto efficient but satisfy a weaker form of efficiency: It rules out efficiency improvements that can directly be inferred from the reported rankings and is referred to as Pareto possibility (Budish, 2011). Formally, a mechanism ψ is **Pareto possible** if for all $P \in \mathscr{P}$,

$$\nexists A \in \mathscr{A} \text{ s.t. } A_i \geq_{P_i} \psi(P)_i \text{ for all } i \in I, \text{ and } A_i >_{P_i} \psi(P)_i \text{ for at least some } i \in I.$$

Next, we formulate a straightforward fairness notion to capture the trade off between fairness and efficiency when comparing identical priority with fair priority multi-category serial dictatorships. That is, for any two individuals $i \in I$ and $j \in I \setminus \{i\}$, we simply count the number of categories where j obtains a better object than i — following i's reported ranking P_i — to calculate i's envy toward j. We then say that a mechanism is ℓ envy-free if any individual i's envy toward any other individual j is at most ℓ for any possible resulting allocation. Formally, a mechanism ψ is $\boldsymbol{\ell}$ **envy-free** if for all $i \in I$, $j \in I \setminus \{i\}$, and $P \in \mathscr{P}$ we have

$$|\{k \in K : \psi(P)^k_j \, P^k_i \, \psi(P)^k_i\}| \leq \ell.$$

Two observations follow immediately: First, in the multi-category housing allocation problem, any mechanism is at best $\lceil \frac{m}{2} \rceil$ envy-free and at the very least m envy-free. Second, fair priority multi-category serial dictatorships are $\lceil \frac{m}{2} \rceil$ envy-free, while identical priority multi-category serial dictatorships fail envy-freeness for any $\ell < m$. Therefore, the minor improvement in efficiency when using identical priority multi-category serial dictatorships instead of fair priority multi-category serial dictatorships comes at the highest possible fairness cost.

Proposition 1. *Identical priority multi-category serial dictatorships are Pareto efficient (Corollary of Monte & Tumennasan (2015) Theorem 2 for $m = 2$) and not ℓ envy-free for any $\ell < m$.*

Proposition 2. *Fair priority multi-category serial dictatorships are Pareto possible and* $\lceil \frac{m}{2} \rceil$ *envy-free.*

Next, we discuss Propositions 1 and 2 in the context of the literature on dictatorship mechanisms, designed for allocation problems with more objects than individuals (Pápai, 2001; Klaus & Miyagawa, 2002; Ehlers & Klaus, 2003; Hatfield, 2009; Monte & Tumennasan, 2015). With two exceptions, the class of multi-category serial dictatorships cannot be directly compared with other definitions in this literature, as it is specifically designed with multi-category housing allocation problems in mind — although, as do the other definitions, the class of multi-category serial dictatorships generalizes the class of serial dictatorships for single-object assignment problems, e.g., as defined in Svensson (1999).

One exception is Monte & Tumennasan (2015) who discus a class of sequential dictatorships that generalize the subclass of identical priority multi-category serial dictatorships for the case of two categories. In that sense, Pareto efficiency of identical priority multi-category serial dictatorships for two categories can be seen as a corollary of Theorem 2 in Monte & Tumennasan (2015) — with the caveat that this paper analyzes ranking mechanisms while they look at direct mechanisms, and therefore their analysis does not specify how sequential dictatorships would work with rankings instead of preferences as inputs.

The other exception is Caspari (2020) who discusses the class of booster draft (ranking) mechanisms. This class of mechanisms generalizes the idea of fair priority multi-category serial dictatorships to multi-object allocation problems, by creating an arbitrary partition of the objects into categories — referred to as boosters. Moreover, if we specify the same priority order for every booster, we can also generalize the idea of identical priority multi-category dictatorships to multi-object allocation problems. Note that for this class of generalized multi-category dictatorship mechanisms to be strategy-proof, we would have to specify the partition into boosters/categories prior to the elicitation of preferences. Furthermore, for any given specification of priority orders, even if we could choose the partition into boosters/categories after observing the reported rankings, this class of mechanisms will not satisfy Pareto possibility when preferences are separable — as most of the time, there will not exist a partition, such that every individual will prefer any subset, containing exactly one object from each subset of the created partition, to any other every other subset. When it comes to the fairness notion, while envy-freeness for multi-object allocation problems (Budish, 2011; Budish & Cantillon, 2012; Caspari, 2020) is strongly related to our fairness notion discussed here, they are not identical. This relates to the fact that objects are not directly comparable across different categories, while in the more general problem individuals can directly compare all the available objects with each other. As a consequence, in multi-object allocation problems we can find mechanisms that are 1-envy free, while for multi-category housing allocation problems, the best achievable fairness for any class of mechanisms is $\lceil \frac{m}{2} \rceil$ envy-freeness. Therefore, even though booster draft mechanisms are shown to be $\lceil \frac{m}{2} \rceil$ envy-free (Theorem 4 in Caspari (2020)), this does not directly imply that fair priority multi-category serial dictatorships are $\lceil \frac{m}{2} \rceil$ envy-free. Finally, readers interested in an example that contrasts the two classes of mechanisms which illustrate our theoretical framework can find Example 1 in the appendix.

4. AN APPLICATION: TEACHING ASSIGNMENTS FOR GRADUATE STUDENTS

In this section, we examine the allocation of teaching positions to graduate students at the economics department of Boston College: From 2019 until the present, following the proposal of this paper, a fair priority multi-category serial dictatorship has been in place and thus replaced the previous allocation system — kick-started by multiple complaints from graduate students over their final assignments in 2018.

We have used the rankings submitted by graduate students for the 2018 academic year, to compare fair priority multi-category serial dictatorships with identical priority multi-category serial dictatorships as well as the actual allocation made that year. That is, analogous to our theoretical part, there are as many students as teaching positions in each semester, and everyone submits a separate ranking for both the fall and spring semester. Moreover, as there are multiple versions of the same position, students end up having to rank only seven different options for each semester.[11] Then, based on 10,000 randomly

[11] The same positions were available in each semester (category), with the number in brackets giving the total number available for each semester: ta (teaching assistant) principles (12,12), ta statistics (3,3), ta econometrics (3,3), lab (laboratory) stats (5,5), lab econometrics (4,4), tf (teaching fellow) principles (4,4), tf statistics (1,1), and “special arrangements” (5,5). In our data 5 out of 37 students made “special arrangements” outside the available positions,e.g., having received a fellowship that freed them of work for one semester. These 5 students

generated priority orders, we have simulated the resulting allocations under both fair and identical priority multi-category serial dictatorships.

First, under the actual allocation, the number of graduate students envying both assignments of at least one other graduate student was roughly 29% — providing a potential reason for the complaints following the 2018 allocation. Surprisingly, even under identical priority multi-category serial dictatorships, on average, only 15% of graduate students would have envied both assignments of at least one other graduate student — while obviously amounting to 0% under any fair priority multi-category serial dictatorship.

Second, to obtain a grasp of how well graduate students like their assignment, we simply took the average rank of their assignment as a proxy — with the best value being 2 (first choice in both semesters) and the worst value being 14 (last choice in both semesters). We found that both classes of mechanisms lead to an expected average rank of 3.21, which is a stark improvement over the 4.72 of the actual 2018 allocation. We note that this measure does not capture the existence of potential inefficiencies under a fair priority multi-category serial dictatorship, where two graduate students would like to trade their bundles with each other. However, even though knowledge of the assignments is publicly available and a cohort of economic graduate students is generally aware of the concept of Pareto improving trades, no one has come forth suggesting a trade of assignments. This suggests that these trades are not particularly relevant in this application.

Third, the standard deviation in the average rank is roughly 2.2 for identical priority multi-category serial dictatorships compared to 1.6 for fair priority multi-category serial dictatorships. That is, while under an identical priority multi-category serial dictatorship, graduate students have a better chance to get their first two choices compared to a fair priority multi-category serial dictatorship, they also have a higher probability of ending up with a much worse average rank — with the worst possible assignment (assignment with positive probability to realize) under the former being 12 and under the latter being 9. Assuming that graduate students are at least mildly risk averse, the last two

---

automatically rank their special arrangement first, in the respective semester, while all other students rank them last — ensuring these students end up with their "special arrangements." Apart from this exception, students then had to simply rank the seven positions for each semester. As there was a new director of graduate studies in charge of the allocation for 2018 it was therefore unknown how the reported rankings would translate into the final allocation, one would reasonably expect that graduate students reported their rankings truthfully.

points imply that fair priorities lead to more preferable lotteries than identical priorities, i.e., lotteries with the same expected average rank but a lower variance.

5. CONCLUSION

We consider the problem of allocating a set of objects, which is sorted into categories of equal size, to a set of individuals, such that everyone obtains exactly one object from each category. Our main theoretical result shows that, in the large class of category-wise neutral and non-bossy mechanisms, any strategy-proof mechanism can be constructed by simply letting individuals choose an object from each category one after another, following some priority order. In this class of mechanisms two ways of selecting priority orders stick out: Either choose an identical priority order for each category or select a priority order that is reversed in every other category. Both intuition and the sparse empirical literature (Budish & Cantillon, 2012) seem to suggest that the second variant should lead to better results. This research also aligns with the discussion for future research of Monte & Tumennasan (2015), suggesting a need to look into solution concepts other than Pareto efficiency, due to its restrictiveness when applied to multi-category housing allocation problems.

A. MATHEMATICAL APPENDIX

Proof of Lemma 1. **If.** Fix any P_i, A, A' and suppose that $A_i \geq_{P_i} A'_i$.
Given the definition of the dominance relation, $A_i \geq_{P_i} A'_i$ implies $a_i^k \; R_i^k \; a_i^{k\prime}$ for all $k \in K$.
Pick any $\succsim_i \in \mathscr{Q}_{P_i}$, we have that $a_i^k \succsim_i a_i^{k\prime}$ for all $k \in K$.
Finally, by separability it follows that $A_i \succsim_i A'_i$ for all $\succsim_i \in \mathscr{Q}_{P_i}$, concluding the proof.
Only if. Fix any P_i, A, A' and suppose that $A_i \not\geq_{P_i} A'_i$.
Note that, $A_i \not\geq_{P_i} A'_i$ implies that A_i and A'_i are distinct allocations. Therefore, for any $\succsim_i \in \mathscr{Q}_{P_i}$ such that $A_i \succsim_i A'_i$, we also have $A'_i \not\succsim_i A_i$ as $\succsim_i$ is antisymmetric — that is, $A_i \succsim_i A'_i$ implies $A_i \succ_i A'_i$ and $A'_i \succsim_i A_i$ implies $A'_i \succ_i A_i$.
We want to show that there exists at least one $\succsim_i \in \mathscr{Q}_{P_i}$ such that $A'_i \succ_i A_i$. Start by randomly selecting a preference $\succsim_i \in \mathscr{Q}_{P_i}$. If $A'_i \succsim_i A_i$ — which implies $A'_i \succ_i A_i$, as A_i and A'_i are distinct and $\succsim_i$ is antisymmetric — we are done. Otherwise, consider a preference $\succsim'_i$ constructed as follows. First, recall that

any two sets of objects O' and O'' are comparable under any preference $\succsim_i \in \mathscr{Q}$ if and only if $|O' \cap O^k| = |O'' \cap O^k| \leq 1$ for all $k \in K$. Then, for any comparable O' and O'' such that $O' = O''$ let $O' \succsim'_i O''$ and $O'' \succsim'_i O'$. More importantly, for any two distinct and comparable sets of objects O' and O'', let $O' \succ'_i O''$ if $O' \geq_{P_i} O''$, let $O'' \succ'_i O'$ if $O'' \geq_{P_i} O'$, and otherwise let $O' \succ'_i O''$ if $O'' \succ_i O'$. By definition we have $A'_i \not\geq_{P_i} A_i$ and $A_i \succ_i A'_i$, and hence $A'_i \succ_i A_i$. It remains to be shown that $\succsim'_i \in Q_{P_i}$. First, for any two comparable, singleton sets $O' = \{o\}$, $O'' = \{o'\}$ — where by definition $\{o\}$ and $\{o'\}$ are in the same category — we have $\{o\} \succsim'_i \{o'\}$ if and only if $\{o\} \succsim_i \{o'\}$. That is, since $\succsim_i$ is consistent with P_i, $\succsim'_i$ is also consistent with P_i.

It remains to be shown that $\succsim'_i$ is separable. For any O' and O'' such that $O' = O''$ this is trivially satisfied. Now, suppose by contradiction that $\succsim'_i$ violates separability for two distinct comparable sets of objects O' and O''. That is, $O'' \succsim'_i O'$ but $o^{k\prime} \succsim'_i o^{k\prime\prime}$ for all $k \in \{k \in K : |O' \cap O^k| = |O'' \cap O^k|\}$. As $\succsim'_i$ is consistent with P_i, we have $o^{k\prime} P_i o^{k\prime\prime}$ for all $k \in \{k \in K : |O' \cap O^k| = |O'' \cap O^k|\}$ and therefore $O' \geq_{P_i} O''$. By construction $O' \geq_{P_i} O''$ implies $O' \succ'_i O''$ — a contradiction with $O'' \succsim'_i O'$.

We have shown that, if $A_i \not\geq_{P_i} A'_i$, then for any $\succsim_i \in Q_{P_i}$ with $A_i \succsim_i A'_i$ — which implies $A_i \succ_i A'_i$ — we can construct another preference $\succsim'_i \in Q_{P_i}$ such that $A'_i \succ'_i A_i$, concluding the proof.

□

Proof of Lemma 2. By Lemma 1 we can substitute strategy-proofness for strong strategy-proofness.

By strong strategy-proofness we have $\psi(P)_i \geq_{P_i} \psi(\hat{P}_i, P_{-i})_i$.

By the assumption of the lemma we have $\psi(P)_i \geq_{\hat{P}_i} \psi(\hat{P}_i, P_{-i})_i$.

Using strong strategy-proofness again we get $\psi(\hat{P}_i, P_{-i})_i \geq_{\hat{P}_i} \psi(P)_i$.

Combining the second and third line we get $\psi(\hat{P}_i, P_{-i})_i \geq_{\hat{P}_i} \psi(P)_i$ which implies that $\psi(\hat{P}_i, P_{-i})_i = \psi(P)_i$.

By nonbossiness it directly follows that $\psi(P) = \psi(\hat{P}_i, P_{-i})$ — if i's outcome did not change no-ones outcome changes. □

Proof of Lemma 3. Consider the outcome of any category-wise neutral ranking mechanism ψ for any two identical preference profiles $P \in \mathscr{I}$ and $\hat{P} \in \mathscr{I}$. Let us define the ℓth best choice in O^k under the identical preference profile P as well as $\hat{P}$: For all $\ell \in \{1, \ldots, n\}$ and $k \in K$, let o^k_ℓ denote $o \in O^k$ s.t. $|\{o' \in O^k : o' \, R^k_i \, o\}| = \ell$ respectively $\hat{o}^k_\ell$ denote $o \in O^k$ s.t. $|\{o' \in O^k : o' \, \hat{R}^k_i \, o\}| = \ell$.

Consider the individual i_l^k that is assigned o_l^k under P, i.e. $\psi(P)_{i_\ell^k}^k = o_\ell^k$. We want to show that the same individual gets the ℓth best choice in O^k under any other identical preference profile $\psi(\hat{P})_{i_\ell^k}^k = \hat{o}_\ell^k$. Consider the following permutation $\hat{\pi}$ defined for all $k \in \{1,\dots,m\}$ and $\ell \in \{1,\dots,n\}$ as $\hat{\pi}[o_\ell^k] = \hat{o}_\ell^k$. By construction, for this particular permutation we have that $\hat{\pi}[P^k] = \hat{P}^k$ for all $k \in \{1,\dots,m\}$. In other words, we have $P^k : o_1^k - o_2^k - \cdots - o_n^k$ and $\hat{\pi}[P^k]$: $\hat{\pi}[o_1^k] - \hat{\pi}[o_2^k] - \cdots - \hat{\pi}[o_n^k]$ which is nothing else than $\hat{\pi}[P^k] : \hat{o}_1^k - \hat{o}_2^k - \cdots - \hat{o}_n^k$, so $\hat{P}^k = \hat{\pi}[P^k]$ for all $k \in \{1,\dots,m\}$.
By neutrality and the construction above, we get $\hat{\pi}[\psi(P)_{i_\ell^k}^k] = \psi((\hat{\pi}[P]))_{i_\ell^k}^k = \psi(\hat{P})_{i_\ell^k}^k$. Moreover, by the definition of the permutation $\hat{\pi}$ we have $\hat{\pi}[\psi(P)_{i_\ell^k}^k] = \hat{\pi}[o_\ell^k] = \hat{o}_\ell^k$. Combining both leads to the desired conclusion that the same individual gets the ℓth best object in set O^k for any two identical preference profiles $\psi(\hat{P})_{i_\ell^k}^k = \hat{o}_l^k$ — both $\hat{o}_l^k$ and o_l^k are assigned to the same individual i_ℓ^k.
□

Proof of Theorem 1. **If.** It is obvious that any multi-category serial dictatorship is category-wise neutral and nonbossy. For (strong) strategy-proofness, suppose by contradiction that there exists $\psi(P)_i \not\succsim_{P_i} \psi(P_i', P_{-i})_i$. Then there exists at least one category O^k such that $\psi(P_i', P_{-1})_i^k \; P_i \; \psi(P)_i^k$. However, as P_{-i} is fixed, all individuals with higher priority will pick identical items in category k independent of i reporting P_i or P_i', so i gets to choose from the same set of remaining objects. Hence, we have that the obtained item under P_i is weakly preferred to any item obtained by reporting another ranking, i.e. $\psi(P)_i^k \; R_i \; \psi(P_i', P_{-i})_i^k$ for all $k \in K$ contradicting the initial statement.
Only if. We now show that any (strongly) strategy-proof, nonbossy, and category-wise neutral mechanism ψ is a multi-category serial dictatorship.
Start by randomly selecting any identical preference profile $P \in \mathscr{I}$, and consider any strategy-proof, nonbossy, and category-wise neutral ranking mechanism ψ. Then, construct a priority order f^k over individuals I for each category $k \in K$ as follows:

$$f^k(i) = |o \in O^k : o \; R_i^k \; \psi(P)_i^k|$$

That is, the individual with the best object in category k under ψ has priority 1 in this category, the individual with the second best object has priority 2 in this category, and so on. Let ψ_f^{FP} denote the multi-category serial dictatorship mechanism with priority orders $f = (f^k)_{k\in K}$ as constructed above. By Lemma 3 the mechanism ψ assigns the same individual $i^k_\ell \in I$ the ℓth best object in O^k according to ranking P^k across every identical ranking $P \in \mathscr{I}$. It is therefore easy to check that, $\psi^{FP}_f(P) = \psi(P)$ for any $P \in \mathscr{I}$ — as i^k_ℓ, the uniquely identifiable individual with the ℓth highest priority in category k under mechanism ψ, is also the individual with the ℓth highest priority in category k under ψ^{FP}_f, i.e., $i^k_\ell = (f^k)^{-1}(\ell)$. That is, for each strategy-proof, nonbossy, and category-wise neutral ranking mechanism ψ, we can construct a unique multi-category serial dictatorship mechanism ψ^{FP}_f, such that $\psi^{FP}_f(P) = \psi(P)$ for all $P \in \mathscr{I}$. Given ψ, it remains to be shown that ψ^{FP}_f gives the same allocation as ψ for any arbitrary preference profile. Start by randomly selecting any preference profile $P \in \mathscr{P}$ and construct an identical preference profile $\hat{P} \in \mathscr{I}$ based on P as follows. For each category $k \in K$, let $\hat{P}^k$ rank object $\psi^{FP}_f(P)^k_{i^k_1} = \psi(P)^k_{i^k_1}$ first, object $\psi^{FP}_f(P)^k_{i^k_2} = \psi(P)^k_{i^k_2}$ second, and so on, with object $\psi^{FP}_f(P)^k_{i^k_n} = \psi(P)^k_{i^k_n}$ ranked last. Note that, for any $A \in \mathscr{A}$ and $i \in I$, such that $\psi(P)_i \geq_{\hat{P}_i} A$ we also have $\psi(P)_i \geq_{P_i} A$. Therefore, by Lemma 2, we can change i's ranking from $\hat{P}_i$ to P_i without changing the outcome of ψ. Recursively applying Lemma 2 for each $i \in I$ we get that $\psi(P) = \psi(\hat{P})$. In a similar fashion, it is easy to check that, under the two preference profiles P and $\hat{P}$ we have $\psi^{FP}_f(P) = \psi^{FP}_f(\hat{P})$. Combining these two observations, we have $\psi^{FP}_f(P) = \psi^{FP}_f(\hat{P}) = \psi(\hat{P}) = \psi(P)$, and therefore $\psi^{FP}_f(P) = \psi(P)$ for all preference profiles $P \in \mathscr{P}$, concluding the proof. □

Proof of Proposition 1. **Identical priority multi-category serial dictatorships are Pareto efficient.**

Consider any Identical priority multi-category serial dictatorship ψ^{IPD}, and let $\mathscr{A}^1 = \mathscr{A} \setminus \{\psi^{IPD}\}$ be the set of allocations potentially Pareto dominating allocation ψ^{IPD}. Note that, the highest priority individual $i_1 = f^{-1}(1)$ gets her m best objects, i.e., for all $P \in \mathscr{P}$ and $k \in K$ we have $\psi^{IPD}(P)^k_{i_1}\ P^k_{i_1}\ o$ for all $o \in O^k \setminus \{\psi^{IPD}(P)^k_{i_1}\}$.
By the definition of the dominance relation we get $\psi^{IPD}(P)_{i_1} \geq_{P_{i_1}} A_i$ for all $A \in \mathscr{A}^1$.
By lemma 1 it follows that $\psi^{IPD}(P)_{i_1} \succsim_{i_1} A_{i_1}$ for all $A \in \mathscr{A}^1$ and for all $\succsim_{i_1} \in \mathscr{Q}_{P_{i_1}}$.
It follows that any allocation Pareto dominating ψ^{IPD} must assign $\psi^{IPD}(P)_{i_1}$

to i_1. That is, the set of allocations potentially Pareto dominating allocation ψ^{IPD} becomes $\mathscr{A}^2 = \{A \in \mathscr{A} : A_{i_1} = \psi^{IPD}(P)_{i_1}\} \setminus \{\psi^{IPD}\}$.
Invoking an analogous argument for $i_2 = f^{-1}(2)$, we get that $\psi^{IPD}(P)_{i_2} \succsim_{i_2} A_{i_2}$ for all $A \in \mathscr{A}^2$ and for all $\succsim_{i_2} \in \mathscr{Q}_{P_{i_2}}$, and therefore the set of allocations potentially Pareto dominating allocation ψ^{IPD} becomes $\mathscr{A}^3 = \{A \in \mathscr{A} : A_{i_1} = \psi^{IPD}(P)_{i_1}$ and $A_{i_2} = \psi^{IPD}(P)_{i_2}\}$.
Iterative applying an analogous argument for individuals $i_3 = f^{-1}(3)$ to $i_{n-1} = f^{-1}(n-1)$ we get that the the set of allocations potentially Pareto dominating allocation ψ^{IPD} becomes $\mathscr{A}^{n-1} = \{A \in \mathscr{A} : A_{i_1} = \psi^{IPD}(P)_{i_1}$ and $\dots$ and $A_{i_{n-1}} = \psi^{IPD}(P)_{i_{n-1}}\} \setminus \{\psi^{IPD}\} = \emptyset$, concluding the proof. □

Proof of Proposition 1. **Identical priority multi-category serial dictatorships are not ℓ envy-free for any $\ell < m$.**

Consider any identical priority multi-category serial dictatorship ψ^{IPD} and some $\ell < m$. Pick any reported list of rankings in the set of identical rankings $P \in \mathscr{I}$. Consider $i_1 = f^{-1}(1)$ and any $j \in I \setminus \{i_1\}$. By the definition of the identical multicategory serial dictatorship, it immediately follows that $|\{k \in K : \psi^{IPD}(P)^k_{i_1}\ P^k_j\ \psi^{IPD}(P)^k_j\}| = m > \ell$, concluding the proof. □

Proof of Proposition 2. **Fair priority multi-category serial dictatorships are Pareto possible.**

Consider any fair priority multi-category serial dictatorship ψ^{FPD} and suppose by contradiction there exists $A \in \mathscr{A}$ such that $A_i \geq_{P_i} \psi^{FPD}(P)_i$ for all $i \in I$ holding strictly for at least one individual. By the definition of the dominance relation $A_i \geq_{P_i} \psi^{FPD}(P)_i$ implies $a^k_i\ R^k_i\ \psi^{FPD}(P)^k_i$ for all $k \in K$ and for all $i \in I$ holding strictly for at least some k and i. Now, pick the highest priority individual i in the first category O^k such that $a^k_i P^k_i \psi^{FPD}(P)^k_i$. As individuals report strict rankings over categories, all higher priority individuals in that category get the same object as before, i.e., $a^k_j = \psi^{FPD}(P)^k_j$ for all $j \in \{j \in I : f^k(j) > f^k(i)\}$. It follows that a^k_i is still available when its i's turn to choose an object form category O^k, and thus $\psi^{FPD}(P)^k_i\ R^k_i\ a^k_i$ contradicting $a^k_i\ P^k_i\ \psi^{FPD}(P)^k_i$. □

Proof of Proposition 2. **Fair priority multi-category serial dictatorships are $\lceil \frac{m}{2} \rceil$ envy-free.**

Consider any fair priority multi-category serial dictatorship ψ^{FPD}. By the definition of any fair priority multi-category serial dictatorship, for any $i \in I$ and $j \in J \setminus \{i\}$ we have that $|\{k \in K : f^k(i) < f^k(j)\}| \geq \lfloor \frac{m}{2} \rfloor$ and therefore

$|\{k \in K : \psi^{FPD}(P)^k_i \, P^k_i \, \psi^{FPD}(P)^k_j\}| \geq \lfloor \frac{m}{2} \rfloor$ for any $P \in \mathscr{P}$. It follows that the maximum envy any $i \in I$ can have is $m - |\{k \in K : \psi^{FPD}(P)^k_i \, P^k_i \, \psi^{FPD}(P)^k_j\}| \leq |\{k \in K : \psi^{FPD}(P)^k_i \, P^k_j \, \psi^{FPD}(P)^k_i\}| \leq \lceil \frac{m}{2} \rceil$ for all $P \in \mathscr{P}$, concluding the proof. □

Example 1. *Consider two individuals (graduate students) $I = \{i_1, i_2\}$. Suppose that they have to work as teaching assistants for a spring semester course $O^1 = \{micro^1, macro^1\}$ and a fall semester course $O^2 = \{micro^2, stats^2\}$. Note that, if both individuals want different teaching assignments within each category, the chosen priority order does not matter. Therefore, more interesting cases are those where both compete for the same objects. In particular, assume both individuals are interested in microeconomics and thus report identical rankings, i.e., for $i \in I$ we have*

$$P^1_i : micro^1 - macro^1, and$$
$$P^2_i : micro^2 - stats^2.$$

The dominance relation $\geq_{P_i}$ tells us that both individuals $i \in I$ (strictly) prefer — under any preference consistent with the reported ranking $\succsim_i \in \mathscr{Q}_{P_i}$ — $\{micro^1, micro^2\}$ to both $\{micro^1, stats^2\}$ and $\{macro^1, micro^2\}$ which they in turn prefer to $\{macro^1, stats^2\}$. Observe that the rankings give no insight into how either one compares $\{micro^1, stats^2\}$ to $\{macro^1, micro^2\}$, i.e., whether $\{micro^1, stats^2\} \succ_i \{macro^1, micro^2\}$, $\{macro^1, micro^2\} \succ_i \{micro^1, stats^2\}$, or $\{micro^1, stats^2\} \sim_i \{macro^1, micro^2\}$.

Now, consider any fair priority multi-category serial dictatorship ψ^{FPD} where $i \in I$ gets to choose first from the spring assignments O^1 and $j \in I \setminus \{i\}$ gets to choose first from the fall assignments O^2. This leads to an assignment of $\psi^{FPD}_i = \{micro^1, stats^2\}$ to i and $\psi^{FPD}_i = \{macro^1, micro^2\}$ to j — which in some sense is a natural way to allocate these objects.

In contrast, it has been suggested that one ought to use identical priority multi-category serial dictatorships — or similar mechanisms like sequential dictatorships — due to them being Pareto efficient, while the weaker notion of Pareto possibility does not rule out all possible inefficiencies. In particular, in the example there is one possible inefficiency where $\{macro^1, micro^2\} \succsim_i \{micro^1, stats^2\}$ and $\{micro^1, stats^2\} \succsim_j \{macro^1, micro^2\}$ with at least one preference holding strictly. However, the way this inefficiency is resolved under an identical priority multi-category serial dictatorship strikes us as unsatisfactory: That is, the only strategy-proof way to avoid this inefficiency is to assign one of the two individuals their absolute best bundle leaving the other to pick up the remains — which gets even more problematic the more categories there are. In this case, the identical priority multi-category serial dictatorship assigns $\psi_i^{IPD} = \{micro^1, micro^2\}$ to i $\psi_i^{FPD} = \{macro^1, stats^2\}$ to j or vice versa.

From an economic perspective, why one might find identical priority multi-category serial dictatorship mechanisms problematic in the example above, stems from the observation that envy-freeness combined with Pareto possibility might very well be a better proxy for the welfare of the resulting allocation than Pareto efficiency. To gain an intuition, consider the following utilities, where the discussed inefficiency occurs and nonetheless fair priority multi-category serial dictatorship leads to higher welfare:[12]

	$\{micro^1, micro^2\}$	$\{micro^1, stats^2\}$	$\{macro^1, micro^2\}$	$\{macro^1, stats^2\}$
u_i	60	40	45	10
u_j	60	45	40	10

Here, the fair priority multi-category serial dictatorship (in its worst case) leads to a welfare of 80 and best case to a welfare of 90, while the identical priority multi-category serial dictatorship leads to a welfare of 70. Moreover, if the market designer wants to maximize Rawlsian welfare (Rawls, 1971), the clear winner is the fair priority multi-category serial dictatorship leading to either 40 or 45 while the identical priority multi-category serial dictatorship leads to a Rawlsian welfare of 10. The utilities in the example reflect the intuition provided by Budish et al. (2016), i.e., that moving from a "bad bundle" to a "medium bundle" leads to higher utility gains compared to moving from a "medium bundle" to a "good bundle". This provides a reasonable explanation as to why mechanisms ensuring that individuals' realized resources are roughly equal might strike us as more appealing and seem to outperform their counterparts in practice.

References

Brams, S. J., Edelman, P. H., & Fishburn, P. C. (2003). Fair division of indivisible

[12] Utilities $u_i : \mathscr{A} \mapsto \mathbb{R}$ represent the underlying worth of an assignment for a given individual — representing cardinal values as opposed to the ordinal interpretation of preferences. The welfare of an assignment is then simply the total sum of the resulting utilities.

items. *Theory and Decision*, *55*(2), 147–180.
Brams, S. J., & Fishburn, P. C. (2000). Fair division of indivisible items between two people with identical preferences: Envy-freeness, pareto-optimality, and equity. *Social Choice and Welfare*, *17*(2), 247–267.
Budish, E. (2011). The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, *119*(6), 1061–1103.
Budish, E., Cachon, G. P., Kessler, J. B., & Othman, A. (2016). Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research*, *65*(2), 314–336.
Budish, E., & Cantillon, E. (2012). The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *American Economic Review*, *102*(5), 2237–71.
Caspari, G. (2020). Booster draft mechanism for multi-object assignment. *ZEW-Centre for European Economic Research Discussion Paper*(20-074).
Edelman, P., & Fishburn, P. (2001). Fair division of indivisible items among people with similar preferences. *Mathematical Social Sciences*, *41*(3), 327–347.
Ehlers, L., & Klaus, B. (2003). Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems. *Social Choice and Welfare*, *21*(2), 265–280.
Hatfield, J. W. (2009). Strategy-proof, efficient, and nonbossy quota allocations. *Social Choice and Welfare*, *33*(3), 505–515.
Hylland, A., & Zeckhauser, R. (1979). The efficient allocation of individuals to positions. *Journal of Political Economy*, *87*(2), 293–314.
Klaus, B., & Miyagawa, E. (2002). Strategy-proofness, solidarity, and consistency for multiple assignment problems. *International Journal of Game Theory*, *30*(3), 421–435.
Monte, D., & Tumennasan, N. (2015). Centralized allocation in multiple markets. *Journal of Mathematical Economics*, *61*, 74–85.
Othman, A., Sandholm, T., & Budish, E. (2010). Finding approximate competitive equilibria: Efficient and fair course allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1* (pp. 873–880).
Pápai, S. (2001). Strategyproof and nonbossy multiple assignments. *Journal of Public Economic Theory*, *3*(3), 257–271.
Rawls, J. (1971). *A Theory of Justice*. Harvard university press.
Roberts, F. S. (1985). *Measurement Theory*. Cambridge University Press.
Roth, A. E. (2008). What have we learned from market design? *Economic Journal*, *118*(527), 285–310.

Shapley, L., & Scarf, H. (1974). On cores and indivisibility. *Journal of Mathematical Economics*, *1*(1), 23–37.

Sönmez, T., & Ünver, M. U. (2010). Course bidding at business schools. *International Economic Review*, *51*(1), 99–123.

Sönmez, T., & Ünver, M. U. (2011). Matching, allocation, and exchange of discrete resources. In *Handbook of Social Economics* (Vol. 1, pp. 781–852). Elsevier.

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