

Constrained Pseudo-Market Equilibrium[†]

By FEDERICO ECHENIQUE, ANTONIO MIRALLES, AND JUN ZHANG*

We propose a pseudo-market solution to resource allocation problems subject to constraints. Our treatment of constraints is general: including bihierarchical constraints due to considerations of diversity in school choice, or scheduling in course allocation; and other forms of constraints needed to model, for example, the market for roommates, combinatorial assignment problems, and knapsack constraints. Constraints give rise to pecuniary externalities, which are internalized via prices. Agents pay to the extent that their purchases affect the value of relevant constraints at equilibrium prices. The result is a constrained-efficient market-equilibrium outcome. The outcome is fair to the extent that constraints treat agents symmetrically. (JEL D47, D61, D63, I11, I21)

We analyze the use of pseudo-markets for assignment problems *under constraints* in market design environments where resources are indivisible and monetary transfers are forbidden. A pseudo-market is an artificial marketplace where agents are given fixed budgets of “funny money” that is only useful within the marketplace. Agents use artificial money to buy affordable probability shares of preferred goods at market-clearing prices. By assigning probability shares to agents, it is possible to get (ex ante) efficiency and fairness among agents. The pseudo-market idea was first proposed by Hylland and Zeckhauser (HZ; 1979) to solve the allocation of indivisible goods to an equal number of agents under unit demand and unit supply constraints. By pricing goods, HZ prove that market-clearing prices exist, and the equilibrium outcome is efficient and fair.¹

*Echenique: Division of the Humanities and Social Sciences, California Institute of Technology (email: fede@hss.caltech.edu); Miralles: Department of Economics, Universitat Autònoma de Barcelona and Barcelona Graduate School of Economics and Università degli Studi di Messina (email: amirallesasensio@gmail.com); Zhang: Institute for Social and Economic Research, Nanjing Audit University (email: zhangjun404@gmail.com). Jeffrey Ely was the coeditor for this article. We thank Eric Budish, Fuhito Kojima, Andy McLennan, Hervé Moulin, and Tayfun Sönmez for comments. Echenique thanks the National Science Foundation for its support through the grants SES-1558757 and CNS-1518941, and the Simons Institute at UC Berkeley for its hospitality while part of the paper was written. Miralles acknowledges financial support from the Spanish Agencia Estatal de Investigación (AEI), through the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2019-000915-S). Zhang acknowledges financial support from National Natural Science Foundation of China (Grant 72122009 and 72033004). Echenique, Miralles, and Zhang are co-first authors.

[†]Go to <https://doi.org/10.1257/aer.20201769> to visit the article page for additional materials and author disclosure statements.

¹If there are only supply constraints, efficiency is the result of the first welfare theorem. In HZ’s model, however, there are also unit demand constraints and the first welfare theorem fails.

We substantially generalize and expand the scope of applicability of pseudo-markets. We consider assignment problems under various constraints. These problems include familiar constrained assignment problems, such as job assignments under regional “ceiling” and “floor” quotas (Kamada and Kojima 2015), and controlled school choice due to considerations of gender or demographic balance (Ehlers et al. 2014). Our approach also handles problems that had not been analyzed via markets before, such as the well-known roommate problem from matching theory, and coalition formation problems.

The key idea in our proposal is to *price* constraints. Think of two workers, Alice and Bob, in HZ’s model. Jobs are in fixed supply. If Alice buys a probability share of job J, there is less left for Bob. Such “pecuniary” externalities are handled in pseudo-markets by pricing J.² In our proposal, we think of the price of J as the price on the supply constraint for J: the constraint saying that the total demand for J cannot exceed its supply. HZ show that pecuniary externalities can be correctly internalized by equilibrium prices, and that an efficient outcome results in equilibrium. In the present paper, we interpret all sorts of constraints, well beyond supply constraints, as giving rise to pecuniary externalities. In consequence, we use prices to align agents’ choices with constraints and to find an efficient outcome. Under standard continuity, convexity, and monotonicity assumptions, we prove the existence of pseudo-market equilibria. Every equilibrium outcome is (constrained) efficient. When constraints do not single out any particular agent, every equilibrium outcome is fair.

Specifically, we differ from the literature in that we do not take as primitive a formal description of constraints. In this sense, our work is orthogonal to the question of implementability of random assignments, which is the main concern of the influential work by Budish et al. (2013) and the recent development by Akbarpour and Nikzad (2020). We start from a set of feasible *ex post* assignments, without specifying the constraints that the feasible assignments must satisfy. The set of random assignments that comply with the constraints is the convex hull of feasible *ex post* assignments, and thus implementable by construction. Our primitive is this convex hull: a polytope in finite assignment problems.³ We proceed by using linear inequalities to characterize the “upper-right” boundary of the convex hull. These linear inequalities include the standard supply constraints and additional constraints in specific applications. Each constraint is then priced. When Alice purchases one unit of good J, she will have to pay to the extent that her purchase affects other agents through different constraints. For example, if there is a ceiling constraint on how many units of J can go to a group of agents, the agents in the group will pay the price of the constraint when they buy units of J. If those agents are also involved in other constraints, then the final personalized prices they face can be different. But if two agents are always involved in the same constraints, they will face equal prices. Equal budgets then ensure that they will not envy each other in the resulting equilibrium outcome.

²We refer to pecuniary externalities as those that can be internalized through prices. It is common to use the same term to speak of the effects that agents may have on prices.

³Therefore, we can address any constraints that pin down a well-defined set of feasible *ex post* assignments.

The idea may seem familiar from the role of shadow prices in optimization with constraints, but the familiarity is deceptive. Imagine using the dual variables (or Lagrange multipliers) associated with each constraint in order to decentralize an allocation that is constrained efficient. We run into two issues. One is that some constraints that impose lower bounds on consumed quantities would lead to negative prices. The other is that decentralizing a constrained efficient allocation would require transfers, as in the second welfare theorem.⁴ With endogenous transfers, one cannot ensure a fair outcome. Our approach, in contrast, can ensure fairness because our prices constitute market equilibria. By pricing only constraints on the “upper-right” boundary, we ensure that all prices in our approach are nonnegative (Sections IE and II). Yet we prove that the equilibrium outcome satisfies all constraints. As long as the constraints do not themselves induce unfairness by treating agents differently, we can obtain a fair outcome. Finally, observe that existing results on the computational complexity of finding HZ equilibria (Vazirani and Yannakakis 2020) imply that it is impossible to obtain our results through convex programming duality because HZ’s model is a special case of ours.

We turn to a discussion of specific applications that motivate our approach, and where our results deliver new insights.

A. Motivation

Matching Jobs to Workers.—HZ illustrated the use of pseudo-markets by way of assigning jobs to workers. Each worker is to receive at most one job, which we call a *unit demand constraint*. Each job is in unit supply, so the sum of probability shares of a particular job assigned to workers cannot exceed one, which we call a *unit supply constraint*. As we shall see, constraints that only involve an individual agent, such as unit demand constraints, have no external effects and do not need to be priced.

Importantly, there is one supply constraint for each good. The price that corresponds to the supply constraint for good l is the familiar “price of good l .” Think of the constraint as capturing a *pecuniary externality*. When Alice purchases good l , the supply constraint implies that there is less good l available for Bob. By pricing supply constraints we ensure that Alice internalizes the effects that her purchase has on Bob. As we shall see, in problems with more complex constraints, we may not be able to ascribe a specific good to each specific price, but the logic of using prices to internalize the external effects induced by constraints extends.

In the jobs-to-workers application, priced constraints affect all agents in the same way. As a consequence, the prices are equal for all agents, and we can show the existence of a market equilibrium outcome that is both efficient and fair.

Finally, the methodology in our paper extends the HZ approach to combinatorial assignments; see Section IIIB (and Budish 2012 for an overview).⁵

⁴We discuss the second welfare theorem without transfers developed by Miralles and Pycia (2021) in the related literature section (Section VII). Note that the outcome of the second welfare theorem may induce envy, even in a textbook economy with only supply constraints.

⁵In combinatorial assignments, the relevant set of constraints that need to be priced does not necessarily coincide with the set of items or the set of bundles. See Section IIIB for a brief discussion.

Assigning Doctors to Hospital Positions.—The problem of assigning doctors to hospitals is similar to the jobs-to-workers example, but with an important twist. Hospitals belong to different geographical regions, and the system seeks to ensure a minimum number of doctors per region (Kamada and Kojima 2015). So we have unit demand and supply constraints as before; but there is now a lower bound—a *floor constraint*—on the number of doctors assigned to each region. Our solution turns these constraints into upper bounds.

There are then two kinds of nonindividual constraints that must be priced: supply constraints and the upper bounds derived from the original floor constraints. When Alice buys into a popular hospital position, she causes a pecuniary externality on Bob, who may have to take a position in a less-demanded regional hospital. The price on the derived upper bounds ensures that she pays more for the popular hospital than if she were only facing the supply constraint. In an equilibrium, now, prices ensure that demand spills over into less attractive regions so as to meet the lower bound for each region.

In the doctors-to-hospitals application, again, all agents are treated in the same way by the priced constraints. In consequence, all agents face the same prices, and we obtain a market equilibrium that is efficient and fair. The application to doctor-hospital matching with regional constraints is discussed in Section IIIA.

Roommates.—A set of college students need to pair up as roommates. Each student has a utility function defined over her possible roommates. We formulate the model as an assignment problem, where we assign objects to agents by treating objects as copies of agents. Each student has two roles: one as agent seeking to match to an object, and one as object that can be matched to different agents. In addition to the familiar unit demand and supply constraints, we must now impose a symmetry constraint. If agent Alice is matched to object Bob, then agent Bob must be matched to object Alice. The symmetry constraints involve more than one agents, and must therefore be priced. When Alice purchases some of the “Bob good” she is committing Bob to consume an equal amount of the “Alice good.” In our pseudo-market, this pecuniary external effect is internalized via prices.

Our result delivers an efficient equilibrium in the market for roommates, but the finding is significant because it is well known that stable matchings may not exist in the model of roommates. Market equilibria capture a different notion of stability, one that is not game theoretic in nature, ensuring that agents are optimizing at the equilibrium outcome.

The application to roommates is discussed in Section V, where we also outline how pseudo-markets can be used in more general matching and coalition formation problems.

B. Related Literature

Constrained resource allocation has received a lot of attention in recent years. The work by Kojima, Sun, and Yu (2020); Gul, Pesendorfer, and Zhang (2019); and ours seems to be the first to look at constrained allocation by way of a market mechanism. The former two papers study the role of gross substitutes in a general model of discrete allocation. Despite a similar focus on markets and constraints,

the results in our papers are very different; see Section VII for more details. In studying constraints, we are motivated by the early work of Budish et al. (2013), Ehlers et al. (2014), and Kamada and Kojima (2015). Our results apply to constraints well beyond those considered by these authors, and we differ substantially in methodology.

We provide a detailed discussion of the related literature in Section VII, once our results have been explained. We also provide a detailed comparison with other work on the pseudo-market equilibrium in that section.

I. The Model

A. Notational Conventions

For vectors $x, y \in \mathbb{R}^n$, $x \leq y$ means that $x_i \leq y_i$ for all $i = 1, \dots, n$; $x < y$ means that $x \leq y$ and $x \neq y$; and $x \ll y$ means that $x_i < y_i$ for all $i = 1, \dots, n$. The set of all $x \in \mathbb{R}^n$ with $0 \leq x$ is denoted by \mathbb{R}_+^n , and the set of all $x \in \mathbb{R}^n$ with $0 \ll x$ is denoted by \mathbb{R}_{++}^n . Inner products are denoted as $x \cdot y = \sum_i x_i y_i$.

Let $X \subseteq \mathbb{R}^n$ be convex. A function $u : X \rightarrow \mathbb{R}$ is

- *quasi-concave* if, for any $x, z \in X$ and $\lambda \in (0, 1)$,

$$\min\{u(z), u(x)\} \leq u(\lambda z + (1 - \lambda)x),$$

- *semi-strictly quasi-concave* if it is quasi-concave, and for any $x, z \in X$ and $\lambda \in (0, 1)$, $u(z) \neq u(x)$ implies that

$$\min\{u(z), u(x)\} < u(\lambda z + (1 - \lambda)x),^6$$

- *concave* if, for any $x, z \in X$ and $\lambda \in (0, 1)$,

$$\lambda u(z) + (1 - \lambda)u(x) \leq u(\lambda z + (1 - \lambda)x),$$

- *expected utility* if there exists a vector $v \in \mathbb{R}^n$ with $u(x) = v \cdot x$ for all $x \in X$,
- *strictly increasing* if $x > x'$ implies that $u(x) > u(x')$.

Given a set $A \subseteq \mathbb{R}^n$, let $\text{co}(A)$ denote the *convex hull* of A in \mathbb{R}^n : the intersection of all convex sets that contain A .

A pair (a, b) , with $a \in \mathbb{R}^n$ and $b \in \mathbb{R}$, defines a *linear inequality* $a \cdot x \leq b$. We say that a linear inequality (a, b) has *nonnegative coefficients* if $a \geq 0$ and $b \geq 0$. Any linear inequality (a, b) defines a (closed) *half-space* $\{x \in \mathbb{R}^n : a \cdot x \leq b\}$.

⁶See Avriel et al. (2010) for a discussion of semi-strictly quasi-concave functions and their applications to economics.

A *polyhedron* in \mathbb{R}^n is a set that is the intersection of a finite number of closed half-spaces. A *polytope* in \mathbb{R}^n is a bounded polyhedron. Two special polytopes are the *simplex*

$$\Delta^n = \left\{ x \in \mathbb{R}_+^n : \sum_{l=1}^n x_l = 1 \right\},$$

and the *subsimplex*

$$\Delta_-^n = \left\{ x \in \mathbb{R}_+^n : \sum_{l=1}^n x_l \leq 1 \right\}.$$

When n is understood, we use the notations Δ and Δ_- , respectively.

B. The Economy

An *economy* is a tuple $\Gamma = (I, O, (Z_i, u_i)_{i \in I}, (q_l)_{l \in O})$, where

- I is a finite set of *agents*, with $N = |I|$;
- O is a finite set of *objects*, with $L = |O|$;
- $Z_i \subseteq \mathbb{R}_+^L$ is the *consumption space* of $i \in I$;
- $u_i : Z_i \rightarrow \mathbb{R}$ is the *utility function* of $i \in I$;
- $q_l \in \mathbb{R}_{++}$ is the amount of $l \in O$.

In an economy, $N = |I|$ is the number of agents, and $L = |O|$ is the number of different *objects*. Each object $l \in O$ is available in quantity q_l . For now we restrict attention to $Z_i = \mathbb{R}_+^L$. Agents' consumption spaces will be restricted further as we introduce constraints.

An *assignment* in Γ is a vector,

$$x = (x_{i,l})_{i \in I, l \in O} \quad \text{with } x_i \in Z_i,$$

where $x_{i,l}$ is the amount of object l received by agent i . Let \mathcal{A} denote the set of all assignments in Γ . Here we ignore the limited supply of objects. It will be imposed further as we introduce supply constraints.

In discrete allocation problems we often interpret assignments as probabilistic allocations: see Section ID. In this case, $x_{i,l}$ is the probability that agent i receives a copy of an object l .

C. Constraints

A *constrained allocation problem* is a pair (Γ, \mathcal{C}) in which Γ is an economy and \mathcal{C} is a subset of the assignments in Γ . The assignments in \mathcal{C} constitute the assignments that satisfy exogenously imposed constraints, and we call them *feasible* in (Γ, \mathcal{C}) .

Observe that the set of feasible assignments is a primitive of our model. Instead of starting from an explicit description of how assignments are constrained, we work directly with the set \mathcal{C} of feasible assignments. In any application where constraints are well-defined, the set \mathcal{C} is also well-defined. Throughout the paper we require that \mathcal{C} be a polytope.

Our model applies to environments with infinitely divisible objects. Most market design applications, however, require indivisible objects. We proceed to introduce some language that is pertinent to the indivisibility case.

D. Special Case: Discrete Allocation

In many market design applications, objects are indivisible, and randomization over deterministic assignments is used to ensure fairness. In these applications, we say that an assignment x is *deterministic* if every $x_{i,l}$ is an integer. When an assignment is not deterministic we call it a random assignment.

Constraints can take many different forms. For example, constraints are often imposed as linear inequalities on deterministic assignments. The usual *unit-demand constraints* require that $\sum_{l \in O} x_{i,l} \leq 1$ for all $i \in I$, and the *supply constraints* require that $\sum_{i \in I} x_{i,l} \leq q_l$ for all $l \in O$. Given a set of constraints, a deterministic assignment is feasible if it satisfies all constraints. A random assignment is *feasible* if it belongs to the convex hull of feasible deterministic assignments; that is, a random assignment is feasible if it can be implemented as a randomization over deterministic assignments that satisfy all constraints. As the number of feasible deterministic assignments is finite, its convex hull, we denote by \mathcal{C} , is a polytope.

The literature often directly imposes constraints on random assignments. When there are only unit demand and supply constraints, this approach is without loss of generality, because the Birkhoff-von Neumann theorem (Birkhoff 1946, von Neumann 1953) guarantees that every random assignment satisfying those constraints is a randomization over deterministic assignments that satisfy those constraints. Budish et al. (2013) generalize this approach by characterizing constraint structures that ensure the implementation of random assignments (see Section III). By taking \mathcal{C} as a primitive, we circumvent the implementation issue. That is how our approach works for “any” finite collection of constraints.

E. Preprocessing of Constraints

Our approach involves pricing constraints, but not all constraints get a price. For example, in HZ’s model, unit demand constraints are not priced; only supply constraints get a price. Here we proceed with a general constrained allocation problem (Γ, \mathcal{C}) , and “preprocess” \mathcal{C} so as to obtain the constraints that have to be assigned a price.

Recall that \mathcal{C} is a polytope. Define the **lower contour set** of \mathcal{C} to be

$$\text{lcs}(\mathcal{C}) = \{x \in \mathbb{R}_+^{NL} : \exists x' \in \mathcal{C} \text{ such that } x \leq x'\}.$$

LEMMA 1: *There exists a finite set Ω of linear inequalities with nonnegative coefficients such that*

$$\text{lcs}(\mathcal{C}) = \bigcap_{(a,b) \in \Omega} \{x \in \mathbb{R}_+^{LN} : a \cdot x \leq b\}.$$
⁷

⁷Lemma 1 is used by Balbuzanov (2019) to define a generalization of the probabilistic serial mechanism that accommodates constraints.

PROOF:

Consider

$$D = \{x' \in \mathbb{R}^{NL} : x' \leq x \text{ for some } x \in \mathcal{C}\}$$

and note that $\text{lcs}(\mathcal{C}) = D \cap \mathbb{R}_+^{NL}$. Write D as $\mathcal{C} - \mathbb{R}_+^{NL}$; thus, since \mathcal{C} is a polytope, D is finitely generated. Then by Theorem 19.1 in Rockafellar (1970) D is polyhedral, and therefore the intersection of finitely many half-spaces. Let Ω be the set of linear inequalities (a, b) defining this collection of half-spaces, so for each $(a, b) \in \Omega$ we have the half-space $\{x' \in \mathbb{R}^{NL} : a \cdot x' \leq b\}$. Since for each i and l there is $x' \in D$ with arbitrarily small $x'_{i,l}$, we must have $a \geq 0$. If $\mathcal{C} = \{0\}$ we may take $b = 0$. If there is $x \in \mathcal{C}$ with $x > 0$ then $b \geq 0$. Hence Ω defines a finite collection of linear inequalities with nonnegative coefficients.

To finish the proof, note that if $z \in D \setminus \text{lcs}(\mathcal{C})$ then $z \notin \mathbb{R}_+^{NL}$. ■

Intuitively, the set Ω of linear inequalities characterizes the “upper-right” boundary of the feasibility set \mathcal{C} . That is why every $(a, b) \in \Omega$ has nonnegative coefficients. Only pricing constraints with nonnegative coefficients is crucial for our approach. It will be clear after our equilibrium notion is defined and the existence theorem is proved.⁸

REMARK 1: *Preprocessing of constraints is needed for two reasons. First, \mathcal{C} may be described by enumerating its vertices, not by a description of its defining linear inequalities. Second, the linear inequalities that define \mathcal{C} may not be of the right form that our method requires. The preprocessing step gives us the relevant linear inequalities that must be priced.*

We sometimes use c to denote a linear inequality, and write $c = (a^c, b^c)$. For every $c \in \Omega$, define the support of c to be

$$\text{supp}(c) = \{(i, l) \in I \times O : a_{i,l}^c > 0\}.$$

Let $a_i^c = (a_{i,l}^c)_{l \in O}$ be the vector of coefficients relevant to agent i in c .

There are two types of inequalities c in Ω : those with $b^c = 0$ and those with $b^c > 0$. If $b^c = 0$, then for all $x \in \mathcal{C}$ and all $(i, l) \in \text{supp}(c)$, we must have $x_{i,l} = 0$. We can, without loss of generality, assume that there is exactly one such inequality; because if there are two inequalities $(a, 0), (a', 0) \in \Omega$ then they can be substituted by $((\max\{a_{i,l}, a'_{i,l}\})_{(i,l) \in I \times O}, 0)$, and if there is no inequality c with $b^c = 0$ then we can include the trivial inequality $(0, 0)$ in Ω . Thus we can let $(a^0, 0) \in \Omega$ be the unique inequality with $b = 0$. When $(a^0, 0)$ is nontrivial, it forbids some agents from consuming certain objects. So we say that l is a *forbidden object* for agent i if $a_{i,l}^0 > 0$.

⁸Computation of Ω can be done through the Avis-Fukuda algorithm (Avis and Fukuda 1992), which can optimize the number of relevant constraints to only include those that define a facet of $\text{lcs}(\mathcal{C})$. The consequence is fewer prices to be defined down the line. We thank an anonymous referee for pointing out this observation.

Among the remaining inequalities $\Omega \setminus \{(a^0, 0)\}$, we say c is an *individual constraint* for agent i if for all $j \neq i$ and $l \in O$, $a_{j,l}^c = 0$. In words, c only restricts i 's consumption. Let Ω^i denote the set of all individual constraints for i . We use $(a^0, 0)$ and individual constraints to refine every i 's consumption space. Let \mathcal{X}_i be the set of vectors $x_i \in Z_i$ such that $x_{i,l} = 0$ if l is a forbidden object for i and x_i satisfies all of i 's individual constraints. That is,

$$\mathcal{X}_i = \{x_i \in Z_i : a_i^0 \cdot x_i \leq 0 \text{ and } a_i^c \cdot x_i \leq b^c \text{ for all } c \in \Omega^i\}.$$

Let $\Omega^* = \Omega \setminus \{(a^0, 0)\} \cup \bigcup_{i \in I} \Omega^i$ collect all constraints that involve more than one agent, including, for example, any supply constraints. The elements of Ω^* will be “priced.” By pricing these constraints we seek to ensure that one agent’s pecuniary externalities on others, imposed via the constraints present in \mathcal{C} , are internalized.

REMARK 2: *It is important to set aside and not price individual constraints because it will allow our fairness conclusion to be stronger. For example, in the special case of HZ’s model, our approach avoids personalized prices altogether; in consequence we obtain (unqualified) envy-freeness (see the definition below).*⁹

Fairness considerations aside, recall our motivation in terms of pricing externalities. Only the inequalities in Ω^* generate externalities, so it makes sense to price these, and not individual constraints.

F. Normative Properties

Given a constrained allocation problem (Γ, \mathcal{C}) , we analyze constrained versions of efficiency and fairness: the efficiency and fairness properties that can be achieved subject to how assignments are constrained.

A feasible assignment $x \in \mathcal{C}$ is *weakly \mathcal{C} -constrained Pareto efficient* if there is no feasible assignment $y \in \mathcal{C}$ such that $u_i(y_i) > u_i(x_i)$ for all i . And $x \in \mathcal{C}$ is *\mathcal{C} -constrained Pareto efficient* if there is no feasible assignment $y \in \mathcal{C}$ such that $u_i(y_i) \geq u_i(x_i)$ for all i with a strict inequality for at least one agent.

Fairness rules out envy among agents who are treated symmetrically by the primitive constraints. We say that two agents i and j are of *equal type* if $\mathcal{X}_i = \mathcal{X}_j$ and, for all $c \in \Omega^*$, $a_i^c = a_j^c$. An agent i *envies* another agent j at an assignment x if $u_i(x_j) > u_i(x_i)$. An assignment $x \in \mathcal{C}$ is *envy-free* if no agent envies another agent at x , and *equal-type envy-free* if no agent envies another agent of equal type at x .

G. Equilibrium

For each $c \in \Omega^*$, we introduce a price p_c . When agent i purchases $x_{i,l}$, she affects other agents’ purchases through the role of $a_{i,l}^c$ in constraint c . Prices are meant to internalize such effects, just as the price of good l classically internalizes the effect

⁹That said, there is no additional difficulty in proving our main result if we were to include prices for the constraints in each Ω^i .

that i has on other agents through the supply constraint for good l . Given a price vector $p = (p_c)_{c \in \Omega^*} \in \mathbb{R}^{\Omega^*}$, the personalized price vector faced by any agent i is defined to be $p_i = (p_{i,l})_{l=1}^L$, such that

$$p_{i,l} = \sum_{c \in \Omega^*} a_{i,l}^c p_c.$$

REMARK 3: *If agents i and j are of equal type, then $p_i = p_j$. Thus prices are only personalized to the extent that constraints are personalized. We present several applications where all agents face the same prices.*

A pair (x^*, p^*) is a *pseudo-market equilibrium* for (Γ, \mathcal{C}) if

- (i) $x_i^* \in \operatorname{argmax}_{x_i \in \mathcal{X}_i} \{u_i(x_i) : p_i^* \cdot x_i \leq 1\}$;
- (ii) $x^* \in \mathcal{C}$;
- (iii) for every $c \in \Omega^*$, $a^c \cdot x^* < b^c$ implies that $p_c^* = 0$.

REMARK 4: *Condition (iii) means that “oversupplied” constraints should be free. It is a “complementary slackness” property that rules out trivial equilibria. We follow Budish et al. (2013) in imposing (iii) directly on the definition of equilibrium.*

II. Main Theorem

THEOREM 1: *Suppose that agents’ utility functions are continuous, quasi-concave, and strictly increasing.*

- *There exists a pseudo-market equilibrium (x^*, p^*) in which x^* is weakly \mathcal{C} -constrained Pareto efficient.*
- *If agents’ utility functions are semi-strictly quasi-concave, there exists a pseudo-market equilibrium (x^*, p^*) in which x^* is \mathcal{C} -constrained Pareto efficient.*
- *Every pseudo-market equilibrium assignment is equal-type envy-free.*

Theorem 1 provides conditions under which there exists a constrained Pareto efficient and fair equilibrium allocation, but it does not say that all equilibria are efficient. In fact, the first welfare theorem does not hold in our model: one can exhibit examples of Pareto inefficient pseudo-market equilibria, and even of Pareto-ranked equilibrium allocations. Crucial to Theorem 1 is the cheapest bundle property: a pair (x, p) satisfies the *cheapest-bundle property* if, for each i , x_i minimizes expenditure $p_i \cdot z_i$ among all the $z_i \in \mathcal{X}_i$ for which $u_i(z_i) = u_i(x_i)$. The cheapest bundle property, and its role in obtaining efficiency, was already established by HZ. When agents’ utility functions are semi-strictly quasi-concave, we show the existence of a pseudo-market equilibrium with the cheapest-bundle property, which in consequence is \mathcal{C} -constrained Pareto efficient. Semi-strict quasi-concavity is slightly stronger than quasi-concavity.

While constrained efficiency is a property of some equilibria, fairness, in the sense of no envy among agents of equal types, holds for all. The idea is simple. Because agents have equal budgets, and agents of equal type face equal prices, agents of equal type cannot envy each other. When all agents are of equal type—we say that the constraints are *anonymous*—we obtain an unqualified notion of fairness, envy-freeness.

The proof of Theorem 1 follows standard ideas. Our main contribution is not in devising a new equilibrium proof, but rather in showing how preprocessing of constraints leads to a model that is amenable to the standard argument. So, given a price vector p , we increase the price of constraints that are violated. Consumers will tend to consume less of the goods that affect the most expensive constraints. A fixed point argument serves to establish an equilibrium price, and the corresponding assignment satisfies all constraints in Ω . Recall that the constraints in Ω only characterize the “upper-right” boundary of \mathcal{C} . Strict monotonicity of u_i ensures that x^* must actually lie in \mathcal{C} .¹⁰

We present an example of controlled school choice (Example 1) in Section IIIA to illustrate how our approach works.

III. Discrete Allocation

The main applications of our approach are to problems of assigning indivisible objects. Each object l is available in fixed integer supply, and we seek efficient and fair random assignments subject to constraints. We discuss different applications, and show how to deal with different kinds of constraints. In Section IIIA we assume that each agent demands at most one copy of any object, and \mathcal{C} is defined from explicit constraint structures that form a *bihierarchy*, as in Budish et al. (2013). We discuss examples with and without floor constraints; as well as applications to hospital-resident matching and school choice. Then in Section IIIB we discuss examples that go beyond bihierarchy constraints; and look at applications to course bidding and refugee resettlement.

A. Bihierarchy Constraints

Consider constraints of the form $\underline{q}_S \leq \sum_{(i,l) \in S} x_{i,l} \leq \bar{q}_S$, where $S \subset I \times O$ is a set of agent-object pairs termed a *constraint set*, and $q_S = (\underline{q}_S, \bar{q}_S)$ is a pair of nonnegative integers that represent *floor* and *ceiling* quotas. A *constraint structure* \mathcal{H} is a collection of constraint sets. Given a constraint structure \mathcal{H} , the set of assignments that satisfy a vector of quotas $\mathbf{q} = (q_S)_{S \in \mathcal{H}}$ is defined to be

$$\mathcal{C} = \left\{ x \in \mathbb{R}_+^{NL} : \underline{q}_S \leq \sum_{(i,l) \in S} x_{i,l} \leq \bar{q}_S \text{ for all } S \in \mathcal{H} \right\}.$$

¹⁰If $x^* \in \text{lcs}(\mathcal{C})$ but $x^* \notin \mathcal{C}$, there exists $x \in \mathcal{C}$ such that $x^* < x$. For every (i, l) such that $x_{i,l}^* < x_{i,l}$, all constraints $c \in \Omega^*$ with $(i, l) \in \text{supp}(c)$ must not be binding at x^* , and thus the price of l faced by i in p^* must be zero. So i must wish to buy more of l , contradicting that x_i^* is i 's optimal consumption.

Note that \mathcal{C} could be different from the convex hull of the deterministic assignments satisfying the quotas \mathbf{q} . When they coincide, \mathcal{C} is called *implementable*.

Budish et al. (2013) prove that if \mathcal{H} includes all singleton sets $\{(i, l)\}$, a sufficient condition for \mathcal{C} to be implementable for all possible quotas is that \mathcal{H} is a *bihierarchy*, and this is also necessary if \mathcal{H} further includes all so-called rows $\{i\} \times O$ and all so-called columns $I \times \{l\}$. A constraint structure \mathcal{H} is a *hierarchy* if for every distinct $S, S' \in \mathcal{H}$, either $S \subset S'$, or $S' \subset S$, or $S \cap S' = \emptyset$. Note, \mathcal{H} is a bihierarchy if there exist two hierarchies \mathcal{H}_1 and \mathcal{H}_2 such that $\mathcal{H} = \mathcal{H}_1 \cup \mathcal{H}_2$ and $\mathcal{H}_1 \cap \mathcal{H}_2 = \emptyset$. In applications, \mathcal{H}_1 usually consists of sub-rows, rows, and sup-rows, representing constraints from the perspective of agents, while \mathcal{H}_2 consists of sub-columns, columns, and sup-columns, representing constraints from the perspective of objects.¹¹

Now we explain how our approach deals with bihierarchy constraints. We discuss two cases.

First, if \mathcal{H} is a bihierarchy, and all floor quotas are zero, then \mathcal{C} as defined above will be the set of feasible assignments that we take as a primitive of our model. More important is that the constraints in \mathcal{H} under any (zero floor) quotas \mathbf{q} define exactly the linear inequalities that characterize $\text{lcs}(\mathcal{C})$. Indeed, in this case $\text{lcs}(\mathcal{C}) = \mathcal{C}$. So we can directly price such inequalities, and no preprocessing is needed. We discuss this case in detail in Section IIIA.

We emphasize that, even under zero floor quotas, our idea of pricing constraints is new. In their generalization of HZ's mechanism, Budish et al. assume that all floor quotas are zero, and that among the two hierarchies \mathcal{H}_1 and \mathcal{H}_2 that constitute \mathcal{H} , \mathcal{H}_1 includes only rows, sub-rows, and singleton sets, while \mathcal{H}_2 includes only columns. Their method is to use individual constraints in \mathcal{H}_1 to refine agents' consumption spaces, and price objects as HZ do. Our approach handles more kinds of zero-floor bihierarchy constraints, and does so by pricing them.

Second, if \mathcal{H} is a bihierarchy, and there are nontrivial floor quotas, then we will still take \mathcal{C} as a primitive of our model, but we will need to derive the linear inequalities that characterize $\text{lcs}(\mathcal{C})$. Preprocessing is needed. Floor constraints are common in real life, but they are well known to be difficult to deal with in market design. Our approach is to derive a new set of ceiling constraints implied by the ceiling and floor constraints present in \mathcal{H} . Pricing ceiling constraints (including the new ceiling constraints we introduce in preprocessing) will cause demand to *spill over* to satisfy the floor constraints. In Section IIIA, we discuss two concrete examples and show how the new ceiling constraints are derived. The two examples, the Japanese medical residency match and controlled school choice, are real-life problems where there are often a large number of participants, but constraints are simple. In practice, we expect that our methods could be operationalized along the lines of our discussion in Section VI.

Bihierarchy Constraints with Zero Floors.—Budish et al. suppose that \mathcal{H} is a bihierarchy, and that all floor quotas are zero. Then $\mathcal{C} = \text{lcs}(\mathcal{C})$, and we can directly price nonindividual constraints. Let \mathcal{H}^* denote the set of

¹¹ A sub-row is of the form $\{i\} \times O'$ where O' is a subset of objects, and a sup-row is of the form $I' \times O$ where I' is a subset of agents. Sub-columns and sup-columns are similar.

nonindividual constraint sets (i.e., those that are not singleton sets, sub-rows, and rows). Given ceiling quotas $(\bar{q}_S)_{S \in \mathcal{H}^*}$, the set of linear inequalities Ω^* that we will price is

$$\{(\mathbf{1}_S, \bar{q}_S) \in \mathbb{R}_+^{LN} \times \mathbb{R}_+ : S \in \mathcal{H}^*\}^{12}$$

Under our assumptions on utilities, an efficient pseudo-market equilibrium exists.

It is interesting to discuss the fairness properties of such equilibria. Two agents, i and j , are of *equal type* if $\mathcal{X}_i = \mathcal{X}_j$ and, for all $S \in \mathcal{H}^*$ and $l \in O$, $(i, l) \in S$ if and only if $(j, l) \in S$. We say that \mathcal{H} is *anonymous* if every two agents are of equal type. If \mathcal{H} is anonymous, every constraint set in \mathcal{H}^* must be a column or sup-column. Under anonymous constraint structures, every pseudo-market equilibrium is envy-free. An example is the Japanese medical residency match with regional caps studied by Kamada and Kojima (2015). Suppose that agents are doctors and objects are hospitals. Each constraint set takes the form $I \times O'$ where $O' \subseteq O$ is the set of hospitals in a geographic region (a city or a prefecture). The ceiling quota $\bar{q}_{O'}$ is the regional cap used to control the maximum number of doctors that the region O' can employ. A collection of such constraint sets is anonymous because they do not distinguish among the identities of individual doctors.

Bihierarchy Constraints with Nontrivial Floors.—When \mathcal{H} is a bihierarchy and floor quotas are nontrivial, we need to derive new ceiling constraints to capture the implications of the ceiling and floor constraints present in \mathcal{H} . We use a simple school choice example to illustrate the methodology, and how preprocessing works. Then we turn to a more general discussion of two applications: the Japanese medical residency match, and the problem of controlled school choice.

EXAMPLE 1 (An Illustration of Our Approach): *Consider an environment where five students are to be assigned seats in three schools. Three students belong to a majority group: M_1, M_2, M_3 , while two students m_1 and m_2 are minorities. There are three schools A, B, and C, each having two seats. School A aims to admit at least one minority student (a floor quota $\underline{q}_A^m = 1$), while school B aims to admit at most one majority student (a ceiling quota $\bar{q}_B^M = 1$). There are no other diversity constraints. Suppose that students have expected utility preferences, with the following von Neumann-Morgenstern utilities.*

	$u_i(A)$	$u_i(B)$	$u_i(C)$
M_1	2	3	1
M_2	2	3	1
M_3	2	3	1
m_1	1	2	3
m_2	1	2	3

¹²By $\mathbf{1}_S$ we denote the indicator vector of the set S .

To obtain the inequalities in Ω , we first calculate new ceiling quotas for the diversity constraints: $\bar{q}_A^M = 1$, $\bar{q}_B^m = 1$, and $\bar{q}_C^m = 1$; and a new ceiling quota for the union of B and C : $\bar{q}_{BC}^m = 1$.¹³ The reasoning is that, because A needs to admit at least one minority student, it can admit at most one majority student, and B and C can jointly admit at most one minority student. Then Ω will consist of the unit demand constraints, and the following inequalities:

$$(1) \quad \begin{cases} 0 \leq x_{M_1,A} + x_{M_2,A} + x_{M_3,A} + x_{m_1,A} + x_{m_2,A} \leq 2, & (e.1) \\ 0 \leq x_{M_1,B} + x_{M_2,B} + x_{M_3,B} + x_{m_1,B} + x_{m_2,B} \leq 2, & (e.2) \\ 0 \leq x_{M_1,C} + x_{M_2,C} + x_{M_3,C} + x_{m_1,C} + x_{m_2,C} \leq 2, & (e.3) \\ 0 \leq x_{M_1,A} + x_{M_2,A} + x_{M_3,A} \leq 1, & (e.4) \\ 0 \leq x_{M_1,B} + x_{M_2,B} + x_{M_3,B} \leq 1, & (e.5) \\ 0 \leq x_{m_1,B} + x_{m_2,B} \leq 1, & (e.6) \\ 0 \leq x_{m_1,C} + x_{m_2,C} \leq 1, & (e.7) \\ 0 \leq x_{m_1,B} + x_{m_2,B} + x_{m_1,C} + x_{m_2,C} \leq 1. & (e.8) \end{cases}$$

Inequalities (e.1)–(e.3) are supply constraints, while the rest arise from the diversity constraints. Importantly, after preprocessing, there are no floor constraints (other than nonnegativity), and all inequalities involve positive coefficients. The unit demand constraints go into each agent’s consumption space. The remaining inequalities must be priced: denote the price of (e. i) by p_i . Majority students will face prices $p_{M,A} = p_1 + p_4$, $p_{M,B} = p_2 + p_5$, and $p_{M,C} = p_3$; while minority students will face $p_{m,A} = p_1$, $p_{m,B} = p_2 + p_6 + p_8$, $p_{m,C} = p_3 + p_7 + p_8$.

In equilibrium, each majority student chooses the lottery $x_M^* = (1/3A, 1/3B, 1/3C)$, facing prices $p_M^* = (1, 2, 0)$; each minority student chooses the lottery $x_m^* = (1/2A, 1/2C)$, facing prices $p_m^* = (0, 2, 2)$. The equilibrium prices are such that $p_4^* = 1$, $p_5^* = 2$, $p_8^* = 2$, and every other price $p_i^* = 0$ (one may verify that (e.2) and (e.6) are not binding).

Note that A is the worst school for minority students. In equilibrium, we price B high enough for minority students to cause their demand to spill over and satisfy the minority floor quota of school A ; but this spillover is not achieved by directly pricing the supply constraint of B . In fact, B is oversupplied in equilibrium. The spillover is achieved by pricing the constraint (e.8) for the union of B and C , an inequality that was derived in our preprocessing of diversity constraints.

Now we turn to a more general discussion of two applications. The first application is the Japanese medical residency match. By introducing regional caps to restrict the number of doctors assigned to urban hospitals, the Japanese government wants to increase the number of doctors assigned to rural hospitals. We can interpret this procedure as an application of our approach: the government’s ideal distribution of doctors can be described by constraints with floor and ceiling quotas both. Specifically, in the Japanese medical residency match, the hospitals O are located in K disjoint regions. Accordingly, there is a partition of hospitals $O = R_1 \cup R_2 \cup \dots \cup R_K$ such that each R_k denotes the set of hospitals in a

¹³We omit the ceiling quotas for other unions of schools because they are implied by the inequalities in (1).

region. So we simply refer to each R_k as a region. For each region R_k , there is a constraint

$$\underline{q}_{R_k} \leq \sum_{l \in R_k} \sum_{i \in I} x_{i,l} \leq \bar{q}_{R_k}.$$

Assume that there are enough hospital positions to assign each doctor a position (we can always add extra “null” hospitals when that is not the case), and that there are enough doctors to meet all floor constraints (that is, $N \geq \sum_{k=1}^K \underline{q}_{R_k}$). We also assume that for each R_k , $\underline{q}_{R_k} \leq \sum_{l \in R_k} q_l$. That is, the floor quota for each R_k does not exceed the total capacity of the hospitals in R_k . Below we derive the inequalities in Ω that characterize $\text{lcs}(\mathcal{C})$, and show that they are anonymous. Theorem 1 delivers a pseudo-market equilibrium that satisfies all constraints, and is efficient and envy-free.

Let $\mathcal{R} = \{R_1, R_2, \dots, R_K\}$ denote the set of regions. For each $\ell \in \{1, 2, \dots, K\}$, let \mathcal{R}_ℓ be the collection of sets that are the union of ℓ distinct regions; that is,

$$\mathcal{R}_\ell = \{R_{k_1} \cup R_{k_2} \cup \dots \cup R_{k_\ell} : \{k_1, k_2, \dots, k_\ell\} \subset \{1, 2, \dots, K\}\}.$$

In particular, $\mathcal{R}_1 = \mathcal{R}$.

Consider the following inequalities:

$$(2) \quad \begin{cases} 0 \leq \sum_{l \in O} x_{i,l} \leq 1, & \text{for all } i \in I; \\ 0 \leq \sum_{i \in I} x_{i,l} \leq q_l, & \text{for all } l \in O; \\ 0 \leq \sum_{i \in I, l \in R} x_{i,l} \leq \bar{q}_R^*, & \text{for all } \ell \in \{1, \dots, K\} \text{ and } R \in \mathcal{R}_\ell, \end{cases}$$

where \bar{q}_R^* is defined according to the following procedure.

- For every $R \in \mathcal{R}_1$, redefine its ceiling quota to be

$$\bar{q}_R^* := \min \left\{ \bar{q}_R, N - \sum_{R' \in \mathcal{R} \setminus \{R\}} \underline{q}_{R'} \right\}.$$

Note that $\bar{q}_R^* \geq \underline{q}_R$ because $N \geq \sum_{R' \in \mathcal{R}} \underline{q}_{R'}$, and \bar{q}_R^* is weakly smaller than the original ceiling quota \bar{q}_R .

- For every $R = R_{k_1} \cup R_{k_2} \in \mathcal{R}_2$, define its ceiling quota to be

$$\bar{q}_R^* := \min \left\{ \bar{q}_{R_{k_1}}^* + \bar{q}_{R_{k_2}}^*, N - \sum_{R' \in \mathcal{R} \setminus \{R_{k_1}, R_{k_2}\}} \underline{q}_{R'} \right\}.$$

- Inductively, for every $R = R_{k_1} \cup R_{k_2} \cup \dots \cup R_{k_\ell} \in \mathcal{R}_\ell$, define its ceiling quota to be

$$\bar{q}_R^* := \min \left\{ \bar{q}_{R \setminus R_{k_x}}^* + \bar{q}_{R_{k_x}}^* \quad \forall x \in \{1, 2, \dots, \ell\}, N - \sum_{R' \in \mathcal{R} \setminus \{R_{k_1}, \dots, R_{k_\ell}\}} \underline{q}_{R'} \right\}.$$

Our next result states the needed preprocessing.

PROPOSITION 1: $\text{lcs}(\mathcal{C}) = \{x \in \mathbb{R}_+^{NL} : x \text{ satisfies the inequalities in (2)}\}$.

Note that among the inequalities in (2), the first two lines are unit demand and supply constraints. What is new is the third line, where we redefine ceiling quotas for each region, and define new ceiling quotas for unions of regions. The new ceiling quotas capture the implications of the existing ceiling and floor constraints in \mathcal{H} .

Aside from unit demand constraints, which are not priced, the remaining inequalities in (2) do not distinguish among doctors' identities. So all doctors face equal prices, and every pseudo-market equilibrium is envy-free.

Our second application is controlled school choice, generalizing our discussion in Example 1. When implementing a school choice program, many school districts care about diversity. We present a model in which each school uses floor and ceiling quotas to control the number of minority and majority students it admits. We assume that the students I are classified into minority I^m and majority I^M , with I^k having cardinality N^k . Each school l has a pair of quotas $(\bar{q}_l^m, \underline{q}_l^m)$ for minority, and a pair of quotas $(\bar{q}_l^M, \underline{q}_l^M)$ for majority students. So, aside from supply constraints, each school l must meet the constraints

$$\begin{cases} \underline{q}_l^m \leq \sum_{i \in I^m} x_{i,l} \leq \bar{q}_l^m, \\ \underline{q}_l^M \leq \sum_{i \in I^M} x_{i,l} \leq \bar{q}_l^M. \end{cases}$$

Of course, we assume that $\underline{q}_l^m + \underline{q}_l^M \leq q_l$.

The inequalities that characterize $\text{lcs}(\mathcal{C})$ can be derived similarly to how we dealt with regional hospitals above. The only difference is that we need to take into account the interaction between the quotas for the two student types within each school. After that, we can deal with the assignments for two types separately. Formally, consider the following inequalities:

$$(3) \quad \begin{cases} 0 \leq \sum_{l \in O} x_{i,l} \leq 1, & \text{for all } i \in I; \\ 0 \leq \sum_{i \in I} x_{i,l} \leq q_l, & \text{for all } l \in O; \\ 0 \leq \sum_{i \in I^m, l \in O'} x_{i,l} \leq \bar{q}_{O'}^{m*}, & \text{for all nonempty } O' \subset O; \\ 0 \leq \sum_{i \in I^M, l \in O'} x_{i,l} \leq \bar{q}_{O'}^{M*}, & \text{for all nonempty } O' \subset O, \end{cases}$$

where $\bar{q}_{O'}^{m*}$ and $\bar{q}_{O'}^{M*}$ are defined as follows.

- For every $l \in O$, redefine its ceiling quotas to be

$$\begin{aligned} \bar{q}_l^{m*} &:= \min \left\{ \bar{q}_l^m, q_l - \underline{q}_l^M, N^m - \sum_{l' \in O \setminus \{l\}} \underline{q}_{l'}^m \right\}, \\ \bar{q}_l^{M*} &:= \min \left\{ \bar{q}_l^M, q_l - \underline{q}_l^m, N^M - \sum_{l' \in O \setminus \{l\}} \underline{q}_{l'}^M \right\}. \end{aligned}$$

- Inductively, for every nonsingleton $O' \subset O$, define its ceiling quotas to be

$$\begin{aligned} \bar{q}_{O'}^{m*} &:= \min \left\{ \bar{q}_{O' \setminus \{l\}}^{m*} + \bar{q}_l^{m*} \quad \forall l \in O', N^m - \sum_{l' \in O \setminus O'} \underline{q}_{l'}^m \right\}, \\ \bar{q}_{O'}^{M*} &:= \min \left\{ \bar{q}_{O' \setminus \{l\}}^{M*} + \bar{q}_l^{M*} \quad \forall l \in O', N^M - \sum_{l' \in O \setminus O'} \underline{q}_{l'}^M \right\}. \end{aligned}$$

Now, preprocessing characterizes $\text{lcs}(\mathcal{C})$ by the inequalities in (3).

PROPOSITION 2: $\text{lcs}(\mathcal{C}) = \{x \in \mathbb{R}_+^{NL} : x \text{ satisfies the inequalities in (3)}\}$.

The proof of Proposition 2 is similar to that of Proposition 1 and thus omitted.

Aside from unit demand constraints, the inequalities in (3) do not distinguish among the identities of students within each type. So in every pseudo-market equilibrium, minority students will not envy other minority students, and majority students will not envy other majority students.

B. Non-bihierarchy Constraints

Our approach works with constrained allocation problems where the constraints do not arise from a bihierarchy. We discuss two specific examples: course allocation and refugee resettlement. Another example is in Section V. Aside from the bihierarchy issue, the applications we have discussed so far impose a unit demand constraint, our next application does not.¹⁴

Course Bidding: Course allocation is a combinatorial assignment problem in which agents demand bundles of objects. The agents are students and the objects are courses. There are obvious supply constraints, stemming from course capacities, but there may exist additional, and more problematic, constraints. For example, if two courses l and l' are complements for some students, so that they must take both of them or neither, then we have the constraint $x_{i,l} = x_{i,l'}$ for every such student i . These constraints cannot be described by inequalities of the form studied by Budish et al. (2013). Our next example illustrates that one cannot solve the problem by using bundles as the primitive objects.

EXAMPLE 2: *Suppose there are three students (1, 2, 3) and three courses (a, b, c). Each student has to take a bundle of two courses or nothing, and each course has one seat. Let $O = \{ab, ac, bc\}$. The following random assignment looks feasible because it satisfies unit demand constraints for students and supply constraints for courses. But it is actually not feasible because bundles are not independent objects. When a bundle is assigned, the other two bundles become unavailable.*

i	ab	ac	bc
1	1/2	0	0
2	0	1/2	0
3	0	0	1/2

In our approach, we can describe \mathcal{C} by considering its vertices: the deterministic allocations that satisfy the relevant constraints. Let A be the basic set of “items,” each of which has a number of copies. Let $O \subset 2^A$ be the set of bundles under consideration. A deterministic assignment is represented by a matrix $x \in \{0, 1\}^{NL}$

¹⁴In the absence of unit demand constraints, the assumption that u_i only depends on i 's marginal distribution over assignments may be problematic, and may require u_i to be linear. This happens if there are multiple implementations of a random assignment, and an agent could obtain different numbers of objects in different ex post assignments. The problem is avoided in cases such as Example 2.

such that $x_{i,l} = 1$ if and only if i obtains the bundle $l \in O$. Now, let \mathcal{C} be the convex hull of the set of deterministic assignments. Starting from \mathcal{C} , one needs to preprocess $\text{lcs}(\mathcal{C})$ and our theorem will deliver a pseudo-market equilibrium with the desirable normative properties.

Refugee Resettlement: The agents are refugee families, and objects are geographical “localities,” places where they can be assigned a living space. Refugee families differ in size. When they are resettled to a new locality, a larger family will occupy more resources than a smaller one. We discuss the refugee resettlement model due to Delacrétaz et al. (2019), and explain how the resulting *knapsack constraints* will result in new difficulties for random assignment problems.

In the model, I is a set of families and O is a set of localities. To accommodate a family, every locality $l \in O$ must provide facilities of multiple dimensions, the set of which is denoted by D . The demand size of every family i is a vector $(v_i^d)_{d \in D} \in \mathbb{Z}_{\geq 0}^{|D|}$. The capacity of every locality l is a vector $(q_l^d)_{d \in D} \in \mathbb{Z}_{\geq 0}^{|D|}$. A deterministic assignment x assigns every family to at most one locality without violating the capacity of any locality in any dimension. So x has to satisfy the constraints

$$(4) \quad \begin{cases} 0 \leq \sum_{l \in O} x_{i,l} \leq 1, & \text{for all } i \in I; \\ 0 \leq \sum_{i \in I} v_i^d \cdot x_{i,l} \leq q_l^d, & \text{for all } l \in O, d \in D. \end{cases}$$

These are simply the unit demand constraints for families and supply constraints for localities, but supply constraints take the form of “knapsack inequalities,” different from those studied by Budish et al. (2013). So their implementation theorem does not apply. Below is a simple example with a single knapsack inequality that illustrates the problem.

EXAMPLE 3 (Knapsack Constraints): *There is only one locality l , and it provides only one facility (i.e., $|D| = 1$). We assume $q_l = 3$. There are two individuals (families of size one) 1, 2 and one family 3 with two members. Consider a random assignment x in which $x_{1,l} = x_{2,l} = 1$, and $x_{3,l} = 1/2$. It is clear that x satisfies unit demand constraints, and the only knapsack constraint $0 \leq x_{1,l} + x_{2,l} + 2x_{3,l} \leq 3$. But x is not implementable, because when 1, 2 are accommodated by l , the remaining capacity of l is not enough to accommodate family 3.*

Despite the difficulties in implementing an allocation, our approach can be applied to the refugee resettlement problem. The feasible integer allocations can be described, and \mathcal{C} defined as their convex hull.

In the course allocation, and refugee resettlement applications, obtaining a complete characterization of \mathcal{C} can be difficult. With knapsack constraints a complete characterization is only known in certain special cases.¹⁵ But characterizing the

¹⁵See Pisinger and Toth (2013). When there is only one knapsack inequality, Weismantel (1996) obtains a complete characterization when the weights v_i^d take only two values. The obtained inequalities are complicated. Dahl and Foldnes (2003) obtain a simpler characterization in a special case where every weight is either one or a fixed larger integer.

linear inequalities that define $\text{lcs}(\mathcal{C})$ can be easier than working with \mathcal{C} , and as a practical matter, one can obtain the set of inequalities Ω from the vertex description of \mathcal{C} by an application of the Avis-Fukuda algorithm (Avis and Fukuda 1992). As we note in Section VI, open source software is readily available. With Ω in hand, our theorem provides a pseudo-market solution that is both efficient and fair.

IV. A Pseudo-Market for “Bads”

So far we have assumed that objects are “goods,” in the sense that agents’ utility functions are strictly increasing. In some applications, however, objects are “bads.” They could represent duties, or tasks, that agents dislike. Other examples are waste disposal, or pollution. Suppose that a certain minimum amount of such bads have to be allocated: the question is to whom, and in which quantities.

The presence of bads gives rise to floor constraints, but we cannot use our previous methods directly as all agents will choose zero consumption from their consumption space. We can, however, borrow an idea from the standard model of labor markets: labor supply is often described as consumption of leisure. We endow every agent with a copy of every bad, and allow them to buy the options of not consuming a bad. Such options become “goods,” and our previous methods apply.

Specifically, for every $l \in O$, q_l denotes the minimum number of copies of l that have to be assigned. Every agent can be assigned at most one object (unit demand). If $\sum_{l \in O} q_l = N$, then every agent must obtain an object so that the problem becomes the one studied by HZ (1979). So we assume that $\sum_{l \in O} q_l < N$. For every $x \in \Delta_-$ and every $i \in I$, $u_i(x)$ is strictly decreasing in x : if $x' > x$, then $u_i(x') < u_i(x)$.

We consider a dual problem $(I, \tilde{O}, \tilde{\Delta}_-, (\tilde{u}_i)_{i \in I}, (q_{\tilde{l}})_{\tilde{l} \in \tilde{O}})$ in which

- The set of objects is $\tilde{O} = \{\tilde{l}\}_{l \in O}$ where every \tilde{l} is an artificial object dual to $l \in O$, and its supply is $q_{\tilde{l}} = N - q_l$. When an agent i consumes an amount z of \tilde{l} , it is understood that i consumes $1 - z$ of l . Because at least q_l of l need to be assigned, the number of copies of \tilde{l} is $N - q_l$.
- The consumption space for every agent is $\tilde{\Delta}_- = \{x \in \mathbb{R}_+^L : x_{\tilde{l}} \in [0, 1] \text{ for every } l \in \tilde{O}, \sum_{\tilde{l} \in \tilde{O}} x_{\tilde{l}} \in [L - 1, L]\}$. So the amount of objects in O that i will consume is $L - \sum_{\tilde{l} \in \tilde{O}} x_{\tilde{l}} \in [0, 1]$.
- Every agent i has the utility function \tilde{u}_i such that for every $x \in \tilde{\Delta}_-$, $\tilde{u}_i(x) = u_i(\mathbf{1} - x)$. When u_i is (semi-strictly) quasi-concave and strictly decreasing, \tilde{u}_i is (semi-strictly) quasi-concave and strictly increasing.

In the dual problem, agents can consume multiple artificial objects. We impose floor constraints on individual consumption, and can derive the inequalities to characterize $\text{lcs}(\mathcal{C})$ as in Section IIIA. Then Theorem 1 applies to give a desirable outcome. We omit the details.

V. A Pseudo-Market for Roommates and Coalition Formation

Our model accommodates very general assignment problems with constraints. Some problems that are usually not described as assignment problems can be

expressed as special cases of our model. In this section we discuss the roommates problem as an illustration of the power of our approach.

The roommates problem is arguably the best-known example in matching theory where game-theoretic stability solutions fail to exist. As a corollary of our main theorem, we obtain the existence of efficient pseudo-market equilibrium assignments. Equilibria embody a form of stability: optimizing agents do not want to change their behavior in the market. In this sense, our results offer a possible way out of the nonexistence of stable matchings. In contrast with some of the other applications in our paper, the main point is not that the allocations are normatively appealing, but rather that, in a model where game-theoretic stability is impossible, market stability (understood as individual optimizing and market clearing) provides a solution.

Consider a set of agents I that constitute the potential roommates or partners. Let O be a copy of I , so $N = L$, and think of $i \in O$ as the alter ego of agent $i \in I$. If x is an assignment, interpret $x_{i,j} = 1$ as agents i and j forming a partnership, or becoming roommates. When i is alone without a roommate, we have $x_{i,i} = 1$. In consequence, we restrict attention to assignments x where $x_{i,j} = x_{j,i}$, meaning that the matrix $(x_{i,j})_{i \in I, j \in I}$ is symmetric.

We say that an assignment x is a *matching* if (i) $x_{i,j} \in \{0, 1\}$ for all $(i, j) \in I \times I$, (ii) x is symmetric, and (iii) x satisfies the unit demand constraints with equality ($\sum_j x_{i,j} = 1$). Define \mathcal{C} to be the convex hull of all matchings.

Note that \mathcal{C} is not equal to the set of symmetric assignments that satisfy the unit demand constraints, dropping the integrality constraints $x_{i,j} \in \{0, 1\}$. Katz (1970) proves that the latter set is the convex hull of all matrices of the form $(1/2)(P + P')$ (P' is the transpose of P) where P is a permutation matrix with no even cycles greater than two. A celebrated result of Edmonds (1965) provides a characterization of \mathcal{C} , which we will use in the proof of Proposition 3 below.

To operationalize our approach, we need to work out the set of inequalities Ω for the roommates problem. To this end, let \mathcal{F} be the set of subsets $F \subseteq I \times I$ such that (i) for all i , $(i, i) \notin F$ and (ii) for every $(i, j) \in F$, $(j, i) \notin F$. For each $F \in \mathcal{F}$, let G_F be the graph with vertex set I and edge set $\{(i, j) : (i, j) \in F \text{ or } (j, i) \in F\}$. Denote the cardinality of the maximum independent edge set of G_F by k_F . For every $i \in I$, let \mathcal{J}_i be the set of subsets $J \subset (\{i\} \times I) \cup (I \times \{i\})$ such that $(i, i) \in J$ and for every $j \neq i$, either $(i, j) \in J$ or $(j, i) \in J$ but not both. Then $\text{lcs}(\mathcal{C})$ is characterized by the following inequalities.

PROPOSITION 3:

$$\begin{aligned} \text{lcs}(\mathcal{C}) = & \left(\bigcap_{\emptyset \neq F \in \mathcal{F}} \left\{ x \in \mathbb{R}_+^{I \times I} : \sum_{(i,j) \in F} x_{i,j} \leq k_F \right\} \right) \\ & \cap \left(\bigcap_{i \in I, J \in \mathcal{J}_i} \left\{ x \in \mathbb{R}_+^{I \times I} : \sum_{(i',j') \in J} x_{i',j'} \leq 1 \right\} \right). \end{aligned}$$

A pseudo-market equilibrium implies a random matching x^* (a probability distribution over matchings) that is Pareto efficient. Of course, x^* needs not be stable in the game theoretic sense, but it corresponds to individual agents' optimizing behavior, as long as these agents take prices as given. Price taking behavior is a plausible assumption in a large centrally run market for partnerships, like for example a

market for roommates in college dormitories. A pseudo-market could be set up by the college, and equilibrium prices could be enforced.

We finalize with a numerical example where stable matchings fail to exist, but where our results deliver an efficient equilibrium.

EXAMPLE 4 (A Market for Roommates): Let $I = \{1, 2, 3\}$. For each agent i , consuming object l is the same as having agent l as her roommate. Suppose that the agents' utilities are as the following table.

	$u_i(1)$	$u_i(2)$	$u_i(3)$
1	1	2	3
2	3	1	2
3	2	3	1

In words, 1 prefers 3 to 2, and prefers 2 to herself (the utilities are respectively 3, 2, and 1).

With these preferences, there are no stable matchings. However, there is a pseudo-market equilibrium. In the equilibrium, the price of the following constraint is two,

$$x_{2,1} + x_{1,3} + x_{3,2} \leq 1;$$

the price of the following constraint is one,

$$x_{1,2} + x_{2,3} + x_{3,1} \leq 1;$$

and the price of every other constraint is zero. Then, agent 1's personalized price vector is $(0, 1, 2)$, 2's personalized price vector is $(2, 0, 1)$, and 3's personalized price vector is $(1, 2, 0)$. All of them choose the consumption $(1/3, 1/3, 1/3)$, and this is the equilibrium assignment.

The application to the roommates problem can be adapted to a general *coalition formation* problem. Given a set of agents I and an integer k , let O be the set of all size- k *coalitions* from I , that is, all subsets of I of cardinality k . A deterministic assignment is a partition of agents into coalitions, and can be represented by a matrix $x \in \{0, 1\}^{NL}$ such that $x_{i,l} = 1$ if and only if i joins the coalition $l \in O$. Unit demand constraints will imply that agents are members of a single coalition. We let \mathcal{C} be the convex hull of the set of deterministic assignments. Then there exists a pseudo-market equilibrium, and the equilibrium assignment is a probability distribution over coalitions.

VI. Practical Implementation

The pseudo-market approach raises a number of practical issues. The first is incentives. Agents should be price takers, and not seek to manipulate market outcomes by misreporting their preferences. We do not include a detailed result on incentive compatibility in our paper, but one can use standard arguments to show that, when there are many agents, and under certain regularity assumptions, agents

have no significant incentives to misreport their preferences. Price taking is important in practice because it simplifies agents' strategies, makes the mechanism's performance robust, and provides policy makers with information about the agents' true preferences. Such information is useful for welfare analysis and policy evaluation.

Because our pseudo-market equilibria are generally not unique, we define the *pseudo-market mechanism* (PMM) as a mechanism that makes an arbitrary selection from the set of equilibria. It is well known that PMM is manipulable, even in the simple HZ environment.¹⁶ In our model, agents who are treated selectively by the constraints may have particularly strong reasons to manipulate prices. But if there are enough agents of each type (classes within which agents are treated symmetrically by the constraints), then any gains from manipulation will be small, and vanish as the market grows. This can be formalized as the *strategy proofness in the large* (SP-L) property proposed by Azevedo and Budish (2019). SP-L is a notion of approximate (interim) incentive compatibility, requiring that truthful reporting be approximately optimal for every agent against any full-support, independent and identical distribution of other agents' reports. SP-L is stronger than approximate Bayes-Nash incentive compatibility used by the literature. Azevedo and Budish prove that any anonymous and envy-free mechanism is SP-L, and extend this result to mechanisms that classify agents into groups and satisfy anonymity and envy-freeness within each group. Because PMM is anonymous and envy-free within agents of equal type, it is SP-L. So we do not worry about incentives in PMM when the market is large and rich in agents' types.¹⁷

A second practical issue is computation, and has several specific aspects. It is often infeasible to ask agents to report their preferences over all possible assignments, so it is important to design a concise but expressive preference-reporting language. In many problems, it is acceptable to assume that agents have additively separable utilities, and therefore ask agents to report utilities over individual objects. Additive separability makes preference reporting relatively easy for agents. Next, a practical implementation of pseudo-markets should leave all the computational burden to computer algorithms. This means that, after agents submit their utilities, an algorithm computes their demand functions, and adjusts prices to find an equilibrium. The problem of computing equilibria is the focus of an extensive literature in computer science, and we shall not go into any details here. Suffice it to say that, while computing equilibria is understood to be computationally hard (see Vazirani and Yannakakis 2020), there are algorithms that work in practice for real-world implementations of pseudo-market mechanisms. One such algorithm is described in Othman, Sandholm, and Budish (2010), and implemented in course bidding at the Wharton School of Business (Budish et al. 2017).

A third practical issue is the computation of the inequalities in Ω . Computing these inequalities is specific to the pseudo-market approach we propose in this paper, and the bulk of our preceding discussion has dealt with how Ω is determined in various applications. In applications where the inequalities are directly available (e.g., bihierarchy constraints with zero floors as in Section IIIA), the task of computing Ω

¹⁶This issue is discussed by HZ, who also anticipate the resolution when there are many agents.

¹⁷Alternatively, we could follow the approach of He et al. (2018) to prove that PMM is asymptotically strategy proof.

is immediate. In other applications, such as the Japanese medical residency match problem and the controlled school choice problem, this task is also easy. When constraints are complex, the task of computing Ω can be difficult, but it can also be automated: we have already referred to the Avis-Fukuda algorithm several times. There is existing software to determine the facets of a polytope from a description of its vertices.¹⁸

We finish our discussion of practical issues by describing a specific application for which we can easily access the practically relevant numerical magnitudes. The point is to illustrate how our methods fit in with an existing real-world implementation of pseudo-markets. The application is to course bidding in business schools, as first studied by Krishna and Ünver (2008) and Sönmez and Ünver (2010). We use the implementation of course bidding at the Wharton School of Business as a realistic guide to the number of students and courses present in the problem.

The number of students and courses in a school can be large, but the set of permissible courses for a specific student in any given semester is not large, and a student can freely choose only a few courses (especially when some core courses are exogenously chosen by the school for the student). Let O^i denote the set of permissible courses for student i , and let d_i denote the maximum number of courses student i can choose from O^i . The most important constraint students face in real life is that they cannot choose courses that have an overlapping meeting time. This constraint can be described through a linear inequality. Partition courses O into disjoint subsets O_1, O_2, \dots, O_K such that every O_k consists of courses with overlapping meeting times. Then we require that, for every student i , $0 \leq \sum_{l \in O_k} x_{i,l} \leq 1$ for every O_k . The constraints imposed on courses are also simple. Aside from capacity constraints, the school may desire to control the student composition of each class. It can have diversity objectives, or want a balance between first- and second-year students. Suppose that the set of students I is partitioned into disjoint types I_1, I_2, \dots, I_T . Every course l sets an upper bound q_l^t on the type I_t students who can take course l . Then we obtain the constraint $0 \leq \sum_{i \in I_t} x_{i,l} \leq q_l^t$.

To sum up, we will have the following constraints in a course allocation problem:

$$(5) \quad \begin{cases} 0 \leq \sum_{l \in O} x_{i,l} \leq d_i \text{ for every } i, & \text{(demand constraints)} \\ 0 \leq \sum_{l \in O_k} x_{i,l} \leq 1 \text{ for every } i \text{ and } O_k, & \text{(schedule constraints)} \\ 0 \leq \sum_{i \in I} x_{i,l} \leq q_l & \text{(supply constraints)} \\ 0 \leq \sum_{i \in I_t} x_{i,l} \leq q_l^t \text{ for every } l \text{ and } I_t. & \text{(diversity constraints)} \end{cases}$$

We also have the feasibility constraints that, for every i and l , $0 \leq x_{i,l} \leq 1$ if $l \in O^i$ and $x_{i,l} = 0$ if $l \notin O^i$. Such constraints, together with the demand and schedule constraints, are individual and do not need to be priced. The resulting constraint structure is a bihierarchy and all floor quotas are zero. So we can directly price the above supply and diversity constraints in our implementation. See the discussion in Section IIIA.

Next, the PMM requires the following steps.

¹⁸See the open source software polymake: <https://polymake.org/doku.php>. We note that our approach only needs to characterize the “upper-right” boundary of C , which can be easier than characterizing C .

- (i) Students report utilities for the courses they wish to take, say assigning positive integers between 1 and 100 to desirable courses, and 0 to courses that they do not wish to take.
- (ii) Compute and select a pseudo-market equilibrium (x^*, p^*) .
- (iii) Announce the equilibrium (x^*, p^*) .

Now we argue that PMM would work well in practice in a situation like Wharton's. Wharton has about 1,700 students and offers up to 350 courses. In every semester, every student takes four to six credit units of courses, which often means no more than ten courses. The system in Wharton implements a competitive equilibrium outcome from equal incomes, as documented in Budish et al. (2017). There are about 14 meeting slots, meaning that $K = 14$. If Wharton wants to introduce diversity constraints, then the price vector will have the dimension of $350(1 + T)$, where $T = 2$ if students are grouped by seniority or gender. Our approach, adding constraints and pricing them, does not lead to a significantly more complicated problem (either computationally, or for students to comprehend) than the existing implementation in Wharton Business School. The complexity of the students preference submission stage is similar, as is demand computation. The problem of finding an equilibrium is similar, and does not involve substantially more prices than in Wharton.

VII. Related Literature

Constrained resource allocation has received a lot of attention in recent years. Budish et al. (2013) identify the bihierarchy structure of constraint blocks in the assignment matrix as the sufficient and necessary condition for implementation. Akbarpour and Nikzad (2020) extend this result by relaxing some constraints and considering approximate implementation. We circumvent the implementation issue by taking the set of implementable assignments as the primitive. Budish et al. allow for floor constraints in implementation but rule out them in their applications. In their extension of the pseudo-market mechanism, they consider column constraints, row constraints and sub-row constraints. By incorporating all row and sub-row constraints into agents' consumption spaces, they prove the existence of equilibria much like HZ's. Their extension is a special case of ours. We can deal with more general constraints on both rows and columns, and allow for floor constraints. When there are no floor constraints, we directly price ceiling constraints, and when there are floor constraints, we translate floor constraints into a different set of ceiling constraints.

Ehlers et al. (2014) focus on the problem of controlled school choice, whereby children have to be assigned seats at different schools to satisfy some diversity objective.¹⁹ Kamada and Kojima (2015) are mainly (but not exclusively) motivated

¹⁹Controlled school choice is introduced by Abdulkadiroğlu and Sönmez (2003), and also investigated by, among others, Ehlers (2010); Hafalir, Yenmez, and Yildirim (2013); Kominers and Sönmez (2013); Westkamp (2013); Echenique and Yenmez (2015); Fragiadakis and Troyan (2017); Aygün and Bó (2021); and Nguyen and Vohra (2019).

by the problem of allocating doctors to hospitals to satisfy geographic quotas. The objective of the quotas is to avoid an excessive concentration of doctors in urban areas.²⁰ Both papers proceed by adapting the notion of stability to capture the presence of constraints, and to add structure to the constraints being considered. To address more general constraints, Kamada and Kojima (forthcoming) relax stability and focus on feasible, individually rational, and fair assignments. They demonstrate that the class of general upper-bound constraints on individual schools are the most permissive constraints under which a student-optimal fair matching exists. That class rules out floor constraints. Our paper can deal with the same kinds of constraints in the above papers, but we follow a different methodological tradition. Instead of a two-sided game-theoretic matching model, we consider object allocation and propose a competitive equilibrium solution. The papers above also investigate the role of incentives in their mechanisms. Our pseudo-market mechanism will be incentive compatible in large markets, as discussed in Section VI, but we choose to focus on existence, efficiency, and fairness in the paper.

The recent work of Balbuzanov (2019) considers a version of the probabilistic serial mechanism for object allocation subject to constraints. Like us, he works on a one-sided object allocation model, but the focus on probabilistic serial makes his analysis clearly distinct from ours. We borrow from this paper the idea, expressed in Lemma 1, allowing us to focus on nonnegative linear inequalities.

The use of markets over lottery shares to solve centralized allocation problems was first proposed by HZ (1979). They assume no constraints other than unit demands and limited supply. They emphasize that equilibrium may not be efficient, and introduce the “cheapest bundle” property that we employ as well in our version of the first welfare theorem. Many other papers have followed HZ in analyzing competitive equilibria as solutions in market design; see for instance, Budish (2011), Ashlagi and Shi (2015), Hafalir and Miralles (2015), and He et al. (2018). Miralles and Pycia (2021) establish the second welfare theorem for the market with satiated preferences and token money: every Pareto efficient assignment may be supported in a Walrasian equilibrium with properly chosen budgets. None of these papers consider constrained allocation problems.²¹

Kojima, Sun, and Yu (2020) and Gul, Pesendorfer, and Zhang (2019) consider market equilibrium in economies with gross substitutes utilities and constraints. Kojima et. al characterize the constraints that preserve the gross substitutes property of firms’ demands in a transferable utility model (Kelso and Crawford’s 1982 job matching model). Gross substitutes ensure equilibrium existence, and Kojima et. al show that the constraint structures have to take the form of “interval constraints.” Gul et. al prove the existence of equilibrium in discrete resource allocation, with limited transfers or no transfers. They show that equilibrium requires random allocations and can be approached by the equilibrium with full transfers. They also show that equilibrium allocations satisfying certain constraints can be constructed by building these constraints into utility functions or into a production technology.

²⁰ See Kamada and Kojima (2017) for an overview.

²¹ He et al. (2018) consider priority-based constraints, which are different from the class of constraints studied here. Miralles and Pycia (2021) do not focus on constraints, but can accommodate linear and individual constraints.

Different from them, we price constraints and can accommodate more general preferences and constraint structures.

Related to our applications, Manjunath (2016) proposes a competitive equilibrium notion for a two-sided fractional matching market. The double-indexed price system in his notion resembles our personalized price system, but he needs to deal with both sides' preferences. As a consequence, his equilibrium exists when there are transfers, but only approximately exists when transfers are forbidden. Bogomolnaia et al. (2017, 2019) study the competitive equilibrium allocation of a mixed manna that contains "goods" and "bads." They prove that an equilibrium always exists. Our model is different than theirs in that agents have unit-demand constraints. So their existence result does not hold in our paper.

Finally, the recent work by Root and Ahn (2020) looks at constrained allocation from a mechanism design perspective. They allow for very general constraint sets, and prove a characterization of group strategy-proof mechanisms that assign objects in a deterministic way. This makes their result distinct from ours.

VIII. Proof of Theorem 1

We first prove Theorem 1 by assuming that all utility functions are semi-strictly quasi-concave. We then explain in Remark 5 the differences when utility functions are only quasi-concave.

We define a price ceiling

$$\bar{p} = \frac{NL}{\min\{b^c : c \in \Omega^*\}} + 1,$$

and a price space $\mathcal{P} = [0, \bar{p}]^{\Omega^*}$. We may assume wlog that each \mathcal{X}_i is bounded.²²

We write $\sum_{l \in O} p_{i,l} x_{i,l}$ as $p_i \cdot x_i$. For every $p \in \mathcal{P}$, we define

$$\begin{aligned} v_i &= \max\{u_i(x_i) : x_i \in \mathcal{X}_i\}, \\ B_i(p) &= \{x_i \in \mathcal{X}_i : p_i \cdot x_i \leq 1\}, \\ d_i(p) &= \operatorname{argmax}\{u_i(x_i) : x_i \in B_i(p)\}, \\ \underline{d}_i(p) &= \operatorname{argmin}\{p \cdot x_i : x_i \in d_i(p)\}, \\ V_i(p) &= \max\{u_i(x_i) : x_i \in B_i(p)\}. \end{aligned}$$

LEMMA 2: *If $V_i(p) < v_i$ then $d_i(p) = \underline{d}_i(p)$.*

PROOF:

Let $x_i \in d_i(p)$. We shall prove that $p_i \cdot x_i = 1$, which means we are done because it implies that all bundles in $d_i(p)$ cost the same at prices p . Let $z_i \in \mathcal{X}_i$ be such that

²²The reason is that C is bounded, so we can add the individual constraint that $x_{i,l} \leq M$ for each agent i and l , for a large enough M . Such a constraint will not bind in equilibrium.

$u_i(z_i) = v_i > u_i(x_i)$. For any $\varepsilon \in (0, 1)$, since \mathcal{X}_i is convex, $\varepsilon z_i + (1 - \varepsilon)x_i \in \mathcal{X}_i$. By the semi-strict quasi-concavity of u_i , $u_i(\varepsilon z_i + (1 - \varepsilon)x_i) > u_i(x_i)$. This means that, for any $\varepsilon \in (0, 1)$,

$$\varepsilon p_i \cdot z_i + (1 - \varepsilon)p_i \cdot x_i > 1.$$

But this is only possible, for arbitrarily small ε , if $p_i \cdot x_i \geq 1$. Since $x_i \in B_i(p)$, we have $p_i \cdot x_i = 1$. ■

LEMMA 3: *If $V_i(p) = v_i$, then*

$$\underline{d}_i(p) = \operatorname{argmin}\{p_i \cdot x_i : u_i(x_i) = v_i \text{ and } x_i \in \mathcal{X}_i\}.$$

PROOF:

Let $x_i \in \underline{d}_i(p)$. Then for any $z_i \in \mathcal{X}_i$ with $p_i \cdot z_i < p_i \cdot x_i$, we have $z_i \in B_i(p)$. So $u_i(z_i) < v_i$ by definition of \underline{d}_i . Therefore, if $z_i \in \operatorname{argmin}\{p_i \cdot x_i : u_i(x_i) = v_i \text{ and } x_i \in \mathcal{X}_i\}$, then

$$p_i \cdot z_i = p_i \cdot x_i,$$

and therefore

$$\underline{d}_i(p) \supseteq \operatorname{argmin}\{p_i \cdot x_i : u_i(x_i) = v_i \text{ and } x_i \in \mathcal{X}_i\}.$$

The converse set inclusion follows similarly because if x_i is not in the right-hand set, there would exist $z_i \in \mathcal{X}_i$ with $p_i \cdot z_i < p_i \cdot x_i$ and $u_i(z_i) = v_i$, which is not possible as such z_i would be in $B_i(p)$. ■

LEMMA 4: *d_i is upper hemi-continuous.*

PROOF:

To prove upper hemi-continuity, we shall prove that d_i has a closed graph. Let $(x_i^n, p^n) \rightarrow (x_i, p)$, with $x_i^n \in d_i(p^n)$ for all n . We prove that $x_i \in d_i(p)$. Since $x_i^n \in B_i(p^n)$, in the limit we have $x_i \in B_i(p)$. Suppose that there is $x_i' \in B_i(p)$ with $u_i(x_i') > u_i(x_i)$. If $p_i \cdot x_i' < 1$, then this strict inequality will be true for p^n with n large enough; a contradiction, as u_i is continuous. If $p_i \cdot x_i' = 1$, then there is $\lambda \in (0, 1)$ large enough that $u_i(\lambda x_i') > u_i(x_i)$, $p_i \cdot (\lambda x_i') < 1$, and $\lambda x_i' \in \mathcal{X}_i$ (recall that the construction of \mathcal{X}_i ensures that this is the case)). Then the argument above applies again. ■

LEMMA 5: *$\underline{d}_i(p)$ is upper hemi-continuous.*

PROOF:

Let $(x_i^n, p^n) \rightarrow (x_i, p)$ with $x_i^n \in \underline{d}_i(p^n)$ for all n . We prove that $x_i \in \underline{d}_i(p)$.

First, consider the case where $V_i(p) < v_i$. By the maximum theorem, V_i is continuous, so $V_i(p^n) < v_i$ for all large enough n . Then Lemma 2 implies that $x_i \in \underline{d}_i(p)$ as d_i is upper hemi-continuous.

Second, consider the case where $V_i(p) = v_i$. We know that $x_i \in d_i(p)$ as d_i is upper hemi-continuous. Suppose (towards a contradiction) that $x_i \notin \underline{d}_i(p)$. Then there is $y_i \in d_i(p)$ with

$$p_i \cdot y_i < p_i \cdot x_i \leq 1.$$

Then for all n large enough,

$$p_i^n \cdot y_i < 1.$$

Since $y_i \in d_i(p)$ and $V_i(p) = v_i$, $u_i(y_i) = v_i$. This means that $V_i(p^n) = v_i$ for all n large enough, as $y_i \in B_i(p^n)$. Then, by Lemma 3, $x_i^n \in \operatorname{argmin}\{p_i^n \cdot x_i : u_i(x_i) = v_i\}$ and $x_i \in \mathcal{X}_i$ for all n large enough. But the correspondence

$$p \mapsto \operatorname{argmin}\{p_i \cdot x : u_i(x) = v_i \text{ and } x \in \mathcal{X}_i\}$$

is upper hemi-continuous, by the maximum theorem. So

$$x_i \in \operatorname{argmin}\{p_i \cdot x : u_i(x) = v_i \text{ and } x \in \mathcal{X}_i\},$$

which, by Lemma 3, implies $x_i \in \underline{d}_i(p)$; a contradiction. ■

It is easy to see that $d_i(p)$ is nonempty, compact-valued, and convex-valued. So $\underline{d}_i(p)$ is also nonempty, compact-valued, and convex-valued. For every $c \in \Omega^*$, define the aggregate demand on c by

$$D_c(p) = \sum_{(i,l) \in \operatorname{supp}(c)} a_{i,l}^c \underline{d}_{i,l}(p) = \cup \{a^c \cdot x : x \in \times_i d_i(p)\}.$$

Define the aggregate demand correspondence by

$$D(p) = (D_c(p))_{c \in \Omega^*},$$

and the excess demand correspondence by

$$z(p) = D(p) - \{\mathbf{b}\},$$

where $\mathbf{b} = (b^c)_{c \in \Omega^*}$.

Consider the correspondence $\varphi : \mathcal{P} \rightarrow \mathcal{P}$ defined by

$$\varphi_c(p) = \left\{ \min\{\max\{0, z_c + p_c\}, \bar{p}\} : z \in z(p) \right\} \quad \text{for all } c \in \Omega^*.$$

Note that $D(p)$, and therefore $z(p)$, are upper hemi-continuous, convex-valued, and compact-valued. Thus, φ is upper hemi-continuous, convex-valued, and compact-valued. By Kakutani's fixed point theorem, there exists $p^* \in \mathcal{P}$ with $p^* \in \varphi(p^*)$.

Note that there exists $z^* \in z(p^*)$ such that

$$(6) \quad p_c^* = \min\{\max\{0, z_c^* + p_c^*\}, \bar{p}\} \quad \text{for all } c \in \Omega^*.$$

Choose $x^* \in \mathbb{R}_+^{NL}$ such that $x_i^* \in d_i(p^*)$ for all i and $a^c \cdot x^* - b^c = z_c^*$ for all $c \in \Omega^*$. We shall prove that (x^*, p^*) is a pseudo-market equilibrium.

LEMMA 6: $p^* \cdot z^* \geq 0$.

PROOF:

If $p^* \cdot z^* < 0$, then there is some $c \in \Omega^*$ with $p_c^* > 0$ and $z_c^* < 0$. By equation (6), then, either $p_c^* = p_c^* + z_c^* \leq \bar{p}$, or $p_c^* = \bar{p} \leq p_c^* + z_c^*$, both of which are not possible as $z_c^* < 0$. ■

LEMMA 7: $p_c^* < \bar{p}$ for all $c \in \Omega^*$.

PROOF:

Suppose towards a contradiction that there exists $c \in \Omega^*$ for which $p_c^* = \bar{p}$. Then, $z_c^* + p_c^* \geq \bar{p}$ implies that $z_c^* \geq 0$. So $\sum_{(i,l) \in \text{supp}(c)} a_{i,l}^c x_{i,l}^* \geq b^c$, and for every $(i, l) \in \text{supp}(c)$, $p_{i,l}^* \geq a_{i,l}^c p_c^*$. However, this is impossible because

$$\sum_{(i,l) \in \text{supp}(c)} a_{i,l}^c x_{i,l}^* \leq \sum_{(i,l) \in \text{supp}(c)} a_{i,l}^c \frac{1}{p_{i,l}^*} \leq \sum_{(i,l) \in \text{supp}(c)} \frac{1}{p_c^*} \leq \frac{NL}{p_c^*} < b^c. \blacksquare$$

PROOF OF THEOREM 1:

We claim that (x^*, p^*) is a pseudo-market equilibrium. For every $c \in \Omega^*$, if $p_c^* > 0$, since $p_c^* < \bar{p}$, then $p_c^* = z_c^* + p_c^*$, which implies $z_c^* = 0$. If $p_c^* = 0$, then $z_c^* + p_c^* \leq 0$, which implies $z_c^* \leq 0$. Recall that $z_c^* = a^c \cdot x^* - b^c$. So this means that x^* satisfies all inequalities in Ω^* . Because for every i , $x_i^* \in \mathcal{X}_i$, x^* satisfies then all inequalities in Ω . Hence,

$$x^* \in \text{lcs}(\mathcal{C}).$$

Moreover, if $z_c^* < 0$, it must be that $p_c^* = 0$, as $p_c^* > 0$ implies $z_c^* = 0$.

It remains to show that $x^* \in \mathcal{C}$. Suppose to the contrary that $x^* \notin \mathcal{C}$. Since $x^* \in \text{lcs}(\mathcal{C})$, there exists $x' \in \mathcal{C}$ such that $x^* \leq x'$. Then $x^* \neq x'$. So there is $(i^*, l^*) \in I \times O$ with $x_{i^*, l^*}^* < x'_{i^*, l^*}$. By definition of \mathcal{C} , $x_{i^*}^* \in \mathcal{X}_{i^*}$.

Consider y_{i^*} defined as $y_{i^*, l} = x_{i^*, l}^*$ for all $l \neq l^*$, and $y_{i^*, l^*} = x'_{i^*, l^*}$. Since $x' \in \mathcal{C}$ and $x'_{i^*, l^*} > 0$, l^* cannot be a forbidden object for i^* . Hence, $y_{i^*} \in \mathcal{X}_{i^*}$.

Moreover, for any $c \in \Omega^*$, if $(i^*, l^*) \in \text{supp}(c)$, then $a^c \cdot x^* < a^c \cdot x' \leq b^c$ and therefore $z_c^* < 0$ (c must not be binding at x^*). Hence, $p_c^* = 0$. In consequence,

$$\begin{aligned} p_{i^*}^* \cdot y_{i^*} &= \sum_{l \neq l^*} p_{i^*, l}^* y_{i^*, l} + \left(\sum_{c \in \Omega^*} \underbrace{p_c^* a_{i^*, l^*}^c}_{=0} \right) y_{i^*, l^*} \\ &= \sum_{l \neq l^*} p_{i^*, l}^* x_{i^*, l}^* \\ &\leq 1. \end{aligned}$$

Thus $y_{i^*} \in B_{i^*}(p^*)$ and $x_{i^*}^* < y_{i^*}$, contradicting the strict monotonicity of u_{i^*} and that $x_{i^*}^* \in d_{i^*}(p^*)$.

We next prove that x^* is \mathcal{C} -constrained Pareto efficient. Suppose towards a contradiction that x is an feasible assignment that Pareto dominates x^* . Given that $x \in \mathcal{C}$, $x_i \in \mathcal{X}_i$. Then, for all $i \in I$, $u_i(x_i) \geq u_i(x_i^*)$, and for some $j \in I$, $u_j(x_j) > u_j(x_j^*)$. So by definition of \underline{d}_i we have that

$$p_i^* \cdot x_i \geq p_i^* \cdot x_i^* \text{ for all } i \in I, \text{ and } p_j^* \cdot x_j > p_j^* \cdot x_j^*.$$

Thus,

$$\sum_{i \in I} p_i^* \cdot x_i > \sum_{i \in I} p_i^* \cdot x_i^*.$$

This is equivalent to

$$\sum_{c \in \Omega^*} p_c^* \left(\sum_{(i,l) \in \text{supp}(c)} a_{i,l}^c x_{i,l} \right) > \sum_{c \in \Omega^*} p_c^* \left(\sum_{(i,l) \in \text{supp}(c)} a_{i,l}^c x_{i,l}^* \right).$$

So there must exist $c \in \Omega^*$ such that $p_c^* > 0$ and

$$\sum_{(i,l) \in \text{supp}(c)} a_{i,l}^c x_{i,l} > \sum_{(i,l) \in \text{supp}(c)} a_{i,l}^c x_{i,l}^*.$$

However, $p_c^* > 0$ implies that $z_c^* = 0$ (i.e., c is binding at x^*), and thus x violates c and is not feasible, which is a contradiction.

Equal-type envy-freeness follows the fact that agents of equal type have equal consumption spaces and equal budgets, and face equal personalized prices. ■

REMARK 5: *The proof uses semi-strict quasi-concavity only in the proof of upper hemi-continuity of \underline{d}_i . To prove existence of an equilibrium without imposing the cheapest-bundle property, observe that continuity and quasi-concavity of u_i is enough to ensure that d_i is upper hemi-continuous, convex-valued, and compact-valued. If z is defined from d_i in place of \underline{d}_i , the proof can be written same as above. To prove that every pseudo-market equilibrium assignment x^* is weakly \mathcal{C} -constrained Pareto efficient, suppose towards a contradiction that there exists a feasible assignment x such that for all $i \in I$, $u_i(x_i) > u_i(x_i^*)$. By utility maximization, for all $i \in I$,*

$$p_i^* \cdot x_i > p_i^* \cdot x_i^*.$$

Thus,

$$\sum_{i \in I} p_i^* \cdot x_i > \sum_{i \in I} p_i^* \cdot x_i^*.$$

So we obtain a contradiction as before.

IX. Proofs of Propositions 1 and 3

PROOF OF PROPOSITION 1:

Let A denote the set on the right-hand side of the proposition. It is easy to see that $A = \text{lcs}(A)$. By the procedure to define \bar{q}_R^* , all elements of \mathcal{C} satisfy (2). So $\mathcal{C} \subset A$ and thus $\text{lcs}(\mathcal{C}) \subset A$. To prove $A \subset \text{lcs}(\mathcal{C})$, we first prove a claim.

CLAIM: For every $\ell \in \{2, \dots, K\}$, every $R = R_{k_1} \cup R_{k_2} \cup \dots \cup R_{k_\ell} \in \mathcal{R}_\ell$, and every $x \in \{1, \dots, \ell\}$,

$$\bar{q}_R^* \geq \underline{q}_{R_{k_x}} + \bar{q}_{R \setminus R_{k_x}}^*.$$

PROOF OF THE CLAIM:

Base case $\ell = 2$: For every $R = R_{k_1} \cup R_{k_2} \in \mathcal{R}_2$, if $\bar{q}_R^* = \bar{q}_{R_{k_1}}^* + \bar{q}_{R_{k_2}}^*$, then the claim holds obviously. Otherwise, $\bar{q}_R^* = N - \sum_{R' \in \mathcal{R} \setminus \{R_{k_1}, R_{k_2}\}} \underline{q}_{R'}$. By definition, $N - \sum_{R' \in \mathcal{R} \setminus \{R_{k_1}\}} \underline{q}_{R'} \geq \bar{q}_{R_{k_1}}^*$. So $\bar{q}_R^* = N - \sum_{R' \in \mathcal{R} \setminus \{R_{k_1}, R_{k_2}\}} \underline{q}_{R'} \geq \bar{q}_{R_{k_1}}^* + \sum_{R' \in \mathcal{R} \setminus \{R_{k_1}\}} \underline{q}_{R'} - \sum_{R' \in \mathcal{R} \setminus \{R_{k_1}, R_{k_2}\}} \underline{q}_{R'} = \bar{q}_{R_{k_1}}^* + \underline{q}_{R_{k_2}}$. Similarly, we prove that $\bar{q}_R^* \geq \underline{q}_{R_{k_1}} + \bar{q}_{R_{k_2}}^*$.

Induction step: Suppose the claim is true for $2, \dots, \ell$. We then prove that it is also true for $\ell + 1$. For any $R = R_{k_1} \cup R_{k_2} \cup \dots \cup R_{k_{\ell+1}} \in \mathcal{R}_{\ell+1}$, if $\bar{q}_R^* = \bar{q}_{R \setminus R_{k_x}}^* + \bar{q}_{R_{k_x}}^*$ for some $x \in \{1, 2, \dots, \ell + 1\}$, then it is obvious that $\bar{q}_R^* \geq \bar{q}_{R \setminus R_{k_x}}^* + \underline{q}_{R_{k_x}}$. By the induction assumption, for every $y \neq x$, $\bar{q}_{R \setminus R_{k_x}}^* \geq \underline{q}_{R_{k_y}} + \bar{q}_{R \setminus (R_{k_x} \cup R_{k_y})}^*$. So $\bar{q}_R^* \geq \underline{q}_{R_{k_y}} + \bar{q}_{R \setminus (R_{k_x} \cup R_{k_y})}^* + \bar{q}_{R_{k_x}}^* \geq \underline{q}_{R_{k_y}} + \bar{q}_{R \setminus R_{k_x}}^*$.

Otherwise, we must have $\bar{q}_R^* = N - \sum_{R' \in \mathcal{R} \setminus \{R_{k_1}, \dots, R_{k_{\ell+1}}\}} \underline{q}_{R'}$. By definition, for every $x \in \{1, 2, \dots, \ell + 1\}$, $N - \sum_{R' \in \mathcal{R} \setminus \{R_{k_1}, \dots, R_{k_{\ell+1}}\} \cup \{R_{k_x}\}} \underline{q}_{R'} \geq \bar{q}_{R \setminus R_{k_x}}^*$. So $\bar{q}_R^* \geq \bar{q}_{R \setminus R_{k_x}}^* + \underline{q}_{R_{k_x}}$.

By induction, we complete the proof of the claim.

Define $A' = \{x \in A : \nexists x' \in A \text{ such that } x < x'\}$. It is clear that $A = \text{lcs}(A')$. We prove that $A' \subset \mathcal{C}$. Suppose there exists $x \in A'$ such that $x \notin \mathcal{C}$. Because x satisfies all original ceiling constraints that define \mathcal{C} , x must violate the floor constraint for some R_k . That is, $\sum_{i \in I, l \in R_k} x_{i,l} < \underline{q}_{R_k}$. Then there must exist some doctor i such that $\sum_{l \in O} x_{i,l} < 1$, since otherwise $\sum_{i \in I, l \in O \setminus R_k} x_{i,l} = N - \sum_{i \in I, l \in R_k} x_{i,l} > N - \underline{q}_{R_k} \geq \bar{q}_{O \setminus R_k}^*$, which contradicts the assumption that $x \in A' \subset A$. Because $\underline{q}_{R_k} \leq \sum_{l \in R_k} q_l$, there must exist $l \in R_k$ such that $\sum_{i \in I} x_{i,l} < q_l$. Now consider a new assignment x' such that $x'_{i,l} = x_{i,l} + \epsilon$ where $0 < \epsilon < \min\{1 - \sum_{l \in O} x_{i,l}, q_l - \sum_{i \in I} x_{i,l}, \underline{q}_{R_k} - \sum_{i \in I, l \in R_k} x_{i,l}\}$, and x' coincides with x in the other cells. So $x < x'$. Below we prove that $x' \in A$, which contradicts the assumption that $x \in A'$.

Suppose towards a contradiction that $x' \notin A$. Let $\ell > 1$ be the smallest index such that there exists $R \in \mathcal{R}_\ell$ with $\sum_{i \in I, l \in R} x'_{i,l} > \bar{q}_R^*$. It is clear that $R_k \subset R$. By Claim, $\bar{q}_R^* \geq \underline{q}_{R_k} + \bar{q}_{R \setminus R_k}^*$. So

$$\sum_{i \in I, l \in R} x'_{i,l} > \underline{q}_{R_k} + \bar{q}_{R \setminus R_k}^*.$$

Because ϵ is chosen such that $\sum_{i \in I, l \in R_k} x'_{i,l} < \underline{q}_{R_k}$. So

$$\sum_{i \in I, l \in R \setminus R_k} x'_{i,l} > \bar{q}_{R \setminus R_k}^*.$$

But it means that $\sum_{i \in I, l \in R \setminus R_k} x_{i,l} > \bar{q}_{R \setminus R_k}^*$, which contradicts that $x \in A$. So $x' \in A$. ■

PROOF OF PROPOSITION 3:

Let D denote the set on the right-hand side of the proposition. We first prove that $D \subset \text{lcs}(\mathcal{C})$. For every $x \in D$, consider the matrix x' obtained by letting $x'_{i,j} = \max\{x_{i,j}, x_{j,i}\}$ for all $(i,j) \in I \times I$. Then x' is symmetric and $x \leq x'$. We prove that $x' \in D$. For any $\emptyset \neq F \in \mathcal{F}$, suppose towards a contradiction that $\sum_{(i,j) \in F} x'_{i,j} > k_F$. Then we define $F' \subset I \times I$ such that for every $(i,j) \in F$, if $x_{i,j} \geq x_{j,i}$, let $(i,j) \in F'$, and otherwise let $(j,i) \in F'$. So $G_{F'}$ and G_F have the same (undirected) edge set, and thus $k_F = k_{F'}$. However, $\sum_{(i,j) \in F'} x_{i,j} = \sum_{(i,j) \in F} x'_{i,j} > k_F$, which contradicts that $x \in D$. Similarly we can prove that for every i and every $J \in \mathcal{J}_i$, $\sum_{(i',j) \in J} x'_{i',j} \leq 1$. Thus, $x' \in D$.

Now define another matrix y by (i) for every $(i,j) \in I \times I$ with $i \neq j$, let $y_{i,j} = x'_{i,j}$, and (ii) for every $i \in I$, let $y_{i,i} = 1 - \sum_{j \neq i} x'_{i,j}$. It is clear that y is symmetric and that $x' \leq y$ (as $\{i\} \times I \in \mathcal{J}_i$). For any $F \in \mathcal{F}$, $(i,i) \notin F$; hence $\sum_{(i,j) \in F} y_{i,j} = \sum_{(i,j) \in F} x'_{i,j} \leq k_F$. Since x' is symmetric and $x' \in D$, for every i and every $J \in \mathcal{J}_i$, $\sum_{(i',j) \in J} y_{i',j} = 1$. So $y \in D$ and it is a bistochastic matrix.

Now we prove that $y \in \mathcal{C}$. Edmonds (1965) proves that a symmetric bistochastic matrix z belongs to \mathcal{C} if and only if, for every $r \in \mathbb{N}$ and every $I' \subset I$ with $|I'| = 2r + 1$, $\sum_{(i,j) \in F} z_{i,j} \leq r$, where $F \subset I' \times I'$ is such that there does not exist $(i,i) \in F$ and for every $(i,j) \in I' \times I'$ with $i \neq j$, either $(i,j) \in F$ or $(j,i) \in F$ but not both. For any such F , $k_F = r$ because I' is odd and we can form r pairs among the $2r$ elements of I' that can be paired. Since $F \in \mathcal{F}$, then, y satisfies Edmonds' inequalities and thus $y \in \mathcal{C}$. Since $x \leq x' \leq y$, $x \in \text{lcs}(\mathcal{C})$. This means that $D \subset \text{lcs}(\mathcal{C})$.

To prove $\text{lcs}(\mathcal{C}) \subset D$, consider any $x \in \mathcal{C}$. Then x is the convex combination of deterministic matchings x^k . For each $\emptyset \neq F \in \mathcal{F}$ and each i , there is at most one j with $x^k_{i,j} = 1$. By the definition of independent edge set, then $\sum_{(i,j) \in F} x^k_{i,j} \leq k_F$. So $\sum_{(i,j) \in F} x_{i,j} \leq k_F$. It is clear that x satisfies the other inequalities related to every \mathcal{J}_i . So $x \in D$. Then it means that $\text{lcs}(\mathcal{C}) \subset D$. ■

REFERENCES

- Abdulkadiroğlu, Atila, and Tayfun Sönmez. 2003. "School Choice: A Mechanism Design Approach." *American Economic Review* 93 (3): 729–47.
- Akbarpour, Mohammad, and Afshin Nikzad. 2020. "Approximate Random Allocation Mechanisms." *Review of Economic Studies* 87 (6): 2473–2510.
- Ashlagi, Itai, and Peng Shi. 2016. "Optimal Allocation without Money: An Engineering Approach." *Management Science* 62 (4): 1078–97.
- Avis, David, and Komei Fukuda. 1992. "A Pivoting Algorithm for Convex Hulls and Vertex Enumeration of Arrangements and Polyhedra." *Discrete and Computational Geometry* 8: 295–313.
- Avriel, Mordecai, Walter E. Diewert, Siegfried Schaible, and Israel Zang. 2010. *Generalized Concavity*. Philadelphia: SIAM.
- Aygiün, Orhan, and Inácio Bó. 2021. "College Admission with Multidimensional Privileges: The Brazilian Affirmative Action Case." *American Economic Journal: Microeconomics* 13 (3): 1–28.
- Azevedo, Eduardo M., and Eric Budish. 2019. "Strategy-Proofness in the Large." *Review of Economic Studies* 86 (1): 81–116.
- Balbusanov, Ivan. 2019. "Constrained Random Matching." Unpublished.
- Birkhoff, Garrett. 1946. "Three Observations on Linear Algebra." *Universidad Nacional de Tucumán Revista, Ser. A* 5: 147–51.
- Bogomolnaia, Anna, Herve Moulin, Fedor Sandomirskiy, and Elena Yanovskaya. 2017. "Competitive Division of a Mixed Manna." *Econometrica* 85 (6): 1847–71.

- Bogomolnaia, Anna, Herve Moulin, Fedor Sandomirskiy, and Elena Yanovskaia.** 2019. "Dividing Bads under Additive Utilities." *Social Choice and Welfare* 52 (3): 395–417.
- Budish, Eric.** 2011. "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes." *Journal of Political Economy* 119 (6): 1061–1103.
- Budish, Eric.** 2012. "Matching 'versus' Mechanism Design." *ACM SIGecom Exchanges* 11 (2): 4–15.
- Budish, Eric, Gerard P. Cachon, Judd B. Kessler, and Abraham Othman.** 2017. "Course Match: A Large-Scale Implementation of Approximate Competitive Equilibrium from Equal Incomes for Combinatorial Allocation." *Operations Research* 65 (2): 314–36.
- Budish, Eric, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom.** 2013. "Designing Random Allocation Mechanisms: Theory and Applications." *American Economic Review* 103 (2): 585–623.
- Dahl, Geir, and Njål Foldnes.** 2003. "Complete Description of a Class of Knapsack Polytopes." *Operations Research Letters* 31 (5): 335–40.
- Delacrétaz, David, Scott Duke Kominers, and Alexander Teytelboym.** 2019. "Matching Mechanisms for Refugee Resettlement." Unpublished.
- Echenique, Federico, and M. Bumin Yenmez.** 2015. "How to Control Controlled School Choice." *American Economic Review* 105 (8): 2679–94.
- Edmonds, Jack.** 1965. "Maximum Matching and a Polyhedron with 0,1-Vertices." *Journal of Research of the National Bureau of Standards B* 69B (1–2): 125–30.
- Ehlers, Lars.** 2010. "School Choice with Control." Unpublished.
- Ehlers, Lars, Isa E. Hafalir, M. Bumin Yenmez, and Muhammed A. Yildirim.** 2014. "School Choice with Controlled Choice Constraints: Hard Bounds versus Soft Bounds." *Journal of Economic Theory* 153: 648–83.
- Fragiadakis, Daniel, and Peter Troyan.** 2017. "Improving Matching under Hard Distributional Constraints." *Theoretical Economics* 12 (2): 863–908.
- Gul, Faruk, Wolfgang Pesendorfer, and Mu Zhang.** 2019. "Market Design and Walrasian Equilibrium." Unpublished.
- Hafalir, Isa, and Antonio Miralles.** 2015. "Welfare-Maximizing Assignment of Agents to Hierarchical Positions." *Journal of Mathematical Economics* 61: 253–70.
- Hafalir, Isa E., M. Bumin Yenmez, and Muhammed A. Yildirim.** 2013. "Effective Affirmative Action in School Choice." *Theoretical Economics* 8 (2): 325–63.
- He, Yinghua, Antonio Miralles, Marek Pycia, and Jianye Yan.** 2018. "A Pseudo-Market Approach to Allocation with Priorities." *American Economic Journal: Microeconomics* 10 (3): 272–314.
- Hylland, Aanund, and Richard Zeckhauser.** 1979. "The Efficient Allocation of Individuals to Positions." *Journal of Political Economy* 87 (2): 293–314.
- Kamada, Yuichiro, and Fuhito Kojima.** 2015. "Efficient Matching under Distributional Constraints: Theory and Applications." *American Economic Review* 105 (1): 67–99.
- Kamada, Yuichiro, and Fuhito Kojima.** 2017. "Recent Developments in Matching with Constraints." *American Economic Review* 107 (5): 200–204.
- Kamada, Yuichiro, and Fuhito Kojima.** Forthcoming. "Fair Matching under Constraints: Theory and Applications." *Review of Economic Studies*.
- Katz, M.** 1970. "On the Extreme Points of a Certain Convex Polytope." *Journal of Combinatorial Theory* 8 (4): 417–23.
- Kelso, Alexander S., Jr., and Vincent P. Crawford.** 1982. "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica* 50 (6): 1483–504.
- Kojima, Fuhito, Ning Sun, and Ning Neil Yu.** 2020. "Job Matching under Constraints." *American Economic Review* 110 (9): 2935–47.
- Kominers, Scott Duke, and Tayfun Sönmez.** 2013. "Designing for Diversity in Matching." *EC '13: Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, 603–04. New York: Association for Computing Machinery.
- Krishna, Aradhna, and M. Utku Ünver.** 2008. "Improving the Efficiency of Course Bidding at Business Schools: Field and Laboratory Studies." *Marketing Science* 27 (2): 262–82.
- Manjunath, Vikram.** 2016. "Fractional Matching Markets." *Games and Economic Behavior* 100: 321–36.
- Miralles, Antonio, and Marek Pycia.** 2021. "Foundations of Pseudomarkets: Walrasian Equilibria for Discrete Resources." *Journal of Economic Theory* 105303.
- Nguyen, Thanh, and Rakesh Vohra.** 2019. "Stable Matching with Proportionality Constraints." *Operations Research* 67 (6): 1503–19.
- Othman, Abraham, Tuomas Sandholm, and Eric Budish.** 2010. "Finding Approximate Competitive Equilibria: Efficient and Fair Course Allocation." *AAMAS '10: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, Vol. 1, 873–80. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

- Pisinger, D., and P. Toth.** 2013. “Knapsack Problems.” In *Handbook of Combinatorial Optimization*, Vol. 1., edited by Ding-Zhu Du and Panos M. Pardalos, 299–428. Boston: Kluwer Academic Publishers.
- Rockafellar, R. Tyrrell.** 1970. *Convex Analysis*. Princeton, NJ: Princeton University Press.
- Root, Joseph, and David S. Ahn.** 2020. “Incentives and Efficiency in Constrained Allocation Mechanisms.” Unpublished.
- Sönmez, Tayfun, and M. Utku Ünver.** 2010. “Course Bidding at Business Schools.” *International Economic Review* 51 (1): 99–123.
- Vazirani, Vijay V., and Mihalis Yannakakis.** 2020. “Computational Complexity of the Hylland-Zeckhauser Scheme for One-Sided Matching Markets.” Unpublished.
- von Neumann, John.** 1953. “A Certain Zero-sum Two-person Game Equivalent to the Optimal Assignment Problem.” In *Contributions to the Theory of Games*, Vol. 2, edited by Harold William Kuhn and Albert William Tucker, 5–12. Princeton, NJ: Princeton University Press.
- Weismantel, Robert.** 1996. “Hilbert Bases and the Facets of Special Knapsack Polytopes.” *Mathematics and Operations Research* 21 (4): 886–904.
- Westkamp, Alexander.** 2013. “An Analysis of the German University Admissions System.” *Economic Theory* 53 (3): 561–89.