Pseudomarkets

Marek Pycia*

August 2021

1 Introduction

Pseudomarkets are allocation mechanisms for one-sided matching markets, an environment introduced in Chapter I.2.[1] Participants of pseudomarkets provide the mechanism with a representation of their cardinal utilities, and the resulting allocation is calculated as part of a Walrasian equilibrium of an auxiliary market, whose participants pay for their allocations with fictitious (token) money that has no value outside of the allocation mechanism. Pseudomarkets differ from the majority of no-transfer allocation mechanisms—e.g. Serial Dictatorships, Top Trading Cycles, Trading Cycles, Probabilistic Serial, and Deferred Acceptance (see Chapters I.2, II.6, II.9, II.11, II.12, III.2, and III.8)—in that they elicit from participants not only ordinal ranking of objects being allocated but also cardinal aspects, or intensity, of the agents' preferences. Eliciting intensities is made possible by the potentially random nature of allocations. The reliance on intensities allows pseudomarkets to allocate objects more efficiently than the ordinal mechanisms but this advantage comes at the cost of increased complexity and potentially more opportunities for gaming the mechanism (weaker incentive compatibility).

This chapter defines pseudomarket mechanisms and examines their incentive and efficiency properties. Its primary focus is on theory, particularly in single-unit demand settings such as school choice. The chapter also discusses selected experimental and empirical results and extends the analysis to multi-unit demand settings such as course allocation. Chapter III.9 discusses course-allocation applications of pseudomarkets in more depth.

*University of Zurich. Federico Echenique, Tuomas Kari, Julian Teichgräber, and Alexander Teytelboym gave me detailed and careful comments on this chapter, and I would also like to thank Yinghua He, Nicole Immorlica, Antonio Miralles, Vijay Vazirani, Jianye Yan, and, particularly, William Zame for conversations that shaped my thinking on pseudomarkets. This chapter will appear in: *Online and Matching-Based Market Design.* Federico Echenique, Nicole Immorlica and Vijay V. Vazirani, Editors. Cambridge University Press.

[1] Pseudomarket mechanisms were proposed by Hylland and Zeckhauser in 1979. The name reflects the idea of replicating a Walrasian market inside an allocation mechanism. These mechanisms are also known as Hylland-Zeckhauser mechanisms or token-money mechanisms.

2 Preliminaries: Walrasian Equilibria in Discrete Settings

We study a finite one-sided matching market with agents $i, j, k \in I = \{1, ..., |I|\}$ and indivisible objects or goods $x, y, z \in X = \{1, ..., |X|\}$. Each object x is represented by a number of identical copies $|x| \in \mathbb{N}$. If agents have outside options, we treat them as objects in X; in particular, this implies that $\sum_{x \in X} |x| \geq |I|$.

We assume that agents demand a probability distribution over objects. This assumption is known as ***single-unit demand***; we discuss its relaxation in Section 5. We denote by $q_i^x \in [0, 1]$ the probability that agent i obtains a copy of object x. Agent i's random ***assignment*** $q_i = (q_i^1, ..., q_i^{|X|})$ is a probability distribution on X; we denote the set of such probability distributions by $\Delta(X)$ and we also refer to them as individual assignments. Agents are expected-utility maximizers: agent i's utility from random assignment q_i is $u_i(q_i) = v_i \cdot q_i = \sum_{x \in X} v_i^x \cdot q_i^x$ where $v_i = (v_i^x)_{x \in X} \in [0, \infty)^{|X|}$ is the vector of agent i's von Neumann-Morgenstein valuations for objects $x \in X$.[2]

The set of economy-wide random assignments is $\Delta(X)^I$. An economy-wide assignment $q = (q_i^x)_{i \in I, x \in X}$ is ***feasible*** if $\sum_{i \in I} q_i^x \leq |x|$ for every object x. An individual assignment is deterministic if it puts mass 1 on one of the objects; an economy-wide assignment is deterministic if all individual assignments are. A mathematical result known as the Birkhoff-von Neumann Theorem states that a feasible random assignment can be expressed as a lottery over feasible deterministic assignments.

We say that a feasible economy-wide assignment q^* and a vector $p^* \in \mathbb{R}_+^X$ constitute an ***equilibrium*** for a constraint vector $w^* \in \mathbb{R}_+^I$ if $q^* = (q_i^*)_{i \in I}$ satisfies $p^* \cdot q_i^* \leq w_i^*$ for all $i \in I$ and $u_i(q_i) > u_i(q_i^*) \Longrightarrow p^* \cdot q_i > w_i^*$ for all $(q_i)_{i \in I} \in \Delta(X)^I$. This definition formally resembles the standard definition of Walrasian equilibrium.[3] The vector p^* is usually interpreted as a ***price vector*** and the constraint vector w^* is interpreted as a ***budget vector***; an interpretation we adopt. In this interpretation, agents' budgets consist of fictitious token money that allows the agents to buy probabilities but is otherwise worthless: unspent token money does not enter agents' utility. Referring to the subset of $\Delta(X)$ satisfying $p^* \cdot q_i^* \leq w_i^*$ as agent i's budget set, we can restate the first inequality as q_i^* being in the budget set, and we can restate the second inequality as no individual assignment preferred to q_i^* being in the budget set.

[2]The single-unit demand assumption imposed implies that the restriction to nonnegative utilities is immaterial; cf. the discussion of Example 1.

[3]When $\sum_{x \in X} |x| = |I|$, this definition implies that all object copies are allocated. Beyond this special case, market clearing is more subtle and is discussed later. The key differences with the standard theory of Walrasian equilibrium are as follows: (i) we study discrete objects while the objects in the standard theory are divisible, and (ii) the unspent balance of w_i^* have no value in our setting while the unspent balances are valuable in the standard theory. These two differences further imply that a key property of the standard Walrasian theory—local non-satiation—fails in our setting. While this chapter includes comparisons to the standard Walrasian theory, the chapter's substantive analysis is self-contained and presumes no background in Walrasian theory. A reader interested in such background might wish to consult e.g. Debreu (1959) or Arrow and Hahn (1971) for classical developments, and Mas-Collel, Whinston, and Green (1995) or Kreps (2013) for contemporary treatments.

As an illustration consider the following.

Example 1. There are two objects x and y such that $|x| = 1$ and $|y| = 2$, and three agents $1, 2$, and 3 such that agents 1 and 2 prefer x to y and agent 3 prefers y to x. While with three or more objects we would need to further specify cardinal utilities to fully describe an environment, with just two objects all specifications of cardinal utilities are equivalent. The reason is that our analysis is invariant to adding a constant to all utilities and to multiplying them by a positive constant. With budgets $w_1 = w_2 = w_3 = 1$ and prices $p_1 = 2$ and $p_2 = 0$, the equilibrium assignment is $q_1 = q_2 = \begin{pmatrix} .5 \\ .5 \end{pmatrix}$ and $q_3 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$.

In this example, the equilibrium happens to be unique, but in general in our model there might be multiple equilibria, just as in the standard Walrasian theory.

2.1 Market Clearing and the Existence of Equilibrium

For any budget vector w and sufficiently high prices p, the existence of an equilibrium implementing some distribution vector q follows from the standard properties of agents' demand. For any agent i, we define this agent's demand correspondence as

$$d_i(p) = \arg \max_{q_i \in \Delta(X), p \cdot q_i \leq w_i} u_i(q_i).$$

For each p, the budget set of agent i is non-empty and compact. Thus, the continuity of the objective u_i in q_i implies that the demand $d_i(p)$ is non-empty and compact. The linearity of u_i and the convexity of the maximization domain imply that $d_i(p)$ is convex. As the mapping from p to the budget set is both upper and lower hemicontinuous, Berge's maximum theorem implies that d_i is upper hemicontinuous. By definition of the demand correspondence, for any nonnegative price vector p, this price vector and any distribution vector q, such that $q_i \in d_i(p)$ for each $i \in I$, satisfy individual agents' equilibrium conditions: for each agent i the distribution q_i is in the budget set, and no assignment preferred to q_i is in the budget set. Furthermore, by setting the prices sufficiently high we can ensure that q is feasible.

More interestingly there always exists an equilibrium in which ***market clears*** in the following complementary-slackness sense: for each $x \in X$, if $p_x^* > 0$ then $\sum_{i \in I} q_i^{*x} = |x|$.

Theorem 2. ***(Existence)*** *For any budget vector $w^* \in \mathbb{R}_+^I$, there is an equilibrium (q^*, p^*) satisfying market clearing.*

The complementary slackness form of market clearing is necessary and caused by the possibility that agents may be satiated: an agent assigned probability 1 of their most favorite object can not raise their utility further. In particular, if the most favorite object is relatively inexpensive, such an agent might not spend all of their budget. As an illustration of the market clearing property, consider an ***outside option*** defined as an object o with many copies, $|o| \geq |I|$.[4] If the supply inequality is strict, $|o| > |I|$, then complementary slackness implies that the price of the outside option is 0.

Like its analogue in the standard Walrasian analysis, the proof of the existence theorem relies on the Kakutani fixed-point theorem.

Proof. In the proof, we refer to vectors $p \in \mathbb{R}^X$ as price vectors even though such vectors allow negative prices; the price vector $p^* \in \mathbb{R}_+^X$ we construct has however only nonnegative prices. Let $w = \max\left\{w_1, ..., w_{|I|}\right\}$ and let $t(p) = (\max\{0, p_x\})_{x \in X}$ be a projection of the vector $p \in \mathbb{R}^X$ onto $\mathbb{R}_+^X$. For sufficiently large $M > 0$, the price vector adjustment function

$$f(p) = t(p) + \left(\sum_{i \in I} d_i(t(p)) - (|x|)_{x \in X}\right)$$

maps the compact Cartesian-product space $\times_{x \in X}[-|x|, M]$ into itself because, if price p_x of good x is higher than $w|I|$, then the sum $\sum_{i \in I} d_i(t(p))_x$ of agents' demands for x is lower than $|x|$ and thus $f(p)_x \leq p_x$.

The properties of demand correspondences d_i established above imply that f is upper hemicontinuous and takes values that are non-empty, convex, and compact. Thus, by the Kakutani's fixed-point theorem, there exists a fixed point $\hat{p} \in f(\hat{p})$. The fixed point property of $\hat{p}$ implies that for each x

$$(|x| + (\hat{p}_x - t_x(\hat{p})))_{x \in X} \in \sum_{i \in I} d_i(t(\hat{p})),$$

and hence there is $q^* \in \times_{i \in I} d_i(p^*)$ for which

$$(|x| + (\hat{p}_x - t_x(\hat{p})))_{x \in X} = \sum_{i \in I} q_i^*.$$

By construction, q^* and $p^* = t(\hat{p})$ are in equilibrium provided q^* is feasible that is $0 \leq \sum_{i \in I} q_i^{*x} \leq |x|$ for each $x \in X$. The first inequality obtains because $q_i^* \in d_i(p^*)$, and the second inequality obtains because $t_x(\hat{p}) \geq \hat{p}_x$. The market-clearing property is satisfied because, for positive p_x^*, we have $\hat{p}_x - t_x(\hat{p}) = 0$ and hence $\sum_{i \in I} q_i^{*x} = |x| + (\hat{p}_x - t_x(\hat{p})) = |x|$. **QED**

2.2 Cheapest Distribution Selection

In our analysis of incentives and efficiency, a special role is played by equilibria satisfying the following property: an equilibrium (q^*, p^*) satisfies the ***cheapest purchase property*** if, for all $i \in I$ and $q_i \in \Delta(X)$, the utility ranking $u_i(q_i) \geq u_i(q_i^*)$ implies $p^* \cdot q_i \geq p^* \cdot q_i^*$.[5] Because an agent demand $d_i(p)$ is non-empty and compact, for any price vector p^*, there exists a distribution q_i satisfying

[4]With slightly more involved but substantially the same analysis we could introduce agent-specific outside options in a similar manner.

[5]The property of purchasing the cheapest distribution is also known as the ***least-expensive lottery property*** in single-unit demand settings and as the ***cheapest bundle property*** in multi-unit demand settings.

the above implication. More interestingly, these distributions can be selected in a way that preserves the equilibrium construction.

Remark 3. (**Existence of Cheapest Purchase Pseudomarkets**) For any budgets $w_i^* \geq 0$, $i \in I$, there is an equilibrium (q^*, p^*) satisfying the market clearing property and the cheapest purchase property. The proof follows the same steps as the proof of Theorem 2 except that we restrict individual demands $d_i(p)$ to distributions satisfying the cheapest purchase property.

The cheapest purchase property is satisfied by all pseudomarket equilibria when each agent has a unique favorite object, as then the agent either can purchase probability 1 of the most favorite object (and then this agent's demand is a singleton at equilibrium prices) or else the agent spends as little as possible on other objects in order to spend more token money on the purchase of the most favorite object. Unlike in the standard Walrasian analysis, however, the cheapest purchase property might fail for some utility profiles. Example 2 below illustrates such a possibility.

2.3 Token Money vs Trade in Endowments

Could we replace the token money with endowments of the probabilities being traded? In standard Walrasian equilibrium we could, but the failure of local non-satiation in the discrete allocation context we study breaks this possibility. We can see it by revisiting Example 1. Suppose that in this example we endow each agent with probability $\frac{1}{3}$ of object x and probability $\frac{2}{3}$ of object y. Does there exist a feasible economy-wide assignment q^* and a price vector $p^* \in \mathbb{R}_+^X$ such that $p^* \cdot q_i^* \leq p^* \begin{pmatrix} \frac{1}{3} \\ \frac{2}{3} \end{pmatrix}$ for all $i \in \{1,2,3\}$ and $u_i(q_i) > u_i(q_i^*) \Longrightarrow p^* \cdot q_i > p^* \begin{pmatrix} \frac{1}{3} \\ \frac{2}{3} \end{pmatrix}$ for all $(q_i)_{i \in \{1,2,3\}} \in \Delta(\{x,y\})^{\{1,2,3\}}$? Exercise 1 asks you to verify that such an assignment and price vector do not exist.

3 Eliciting Agents' Utilities

So far, we assumed that agents' utilities are known. In applications this information needs to be elicited. A mechanism elicits agents' information and this information determines the mechanism's outcome. In our setting, each agent's information is the agent's utility and the mechanism's outcome is the economy-wide assignment.

A ***pseudomarket mechanism*** maps a profile of utilities $(u_i)_{i \in I}$ to an equilibrium assignment for some budget vector w^*. We allow the budget vector to depend on the profile of utilities. Even when budgets do not depend on reported utilities, there are many budget vectors and the uniqueness of equilibrium is not assured, thus there might be many pseudomarket mechanisms. Of particular importance are pseudomarket mechanisms in which budgets are equal, as in

Example 1.[6]

Can we expect agents to report their utilities truthfully? The question is usually conceptualized in terms of incentives to do so, and the usual goal—known as incentive compatibility—is to ensure that by providing us with truthful reports the agents maximize (or nearly maximize) their payoffs. Agents' incentives in general depend on what they know or believe, including what they know or believe about other agents. A gold standard of incentive compatibility is strategy-proofness because it imposes no assumptions on agents' beliefs. A mechanism is ***strategy-proof*** if reporting the true utility is a dominant strategy for each agent, that is the agent's expected utility after reporting their true utility is weakly higher than after any other report irrespective of the reports of other agents (cf. Chapter I.1).

3.1 Fixed-Price Pseudomarkets

Pseudomarkets can be strategy-proof. Suppose that there is an outside option and fix a profile of agents' budgets. Then:

Theorem 4. ***(Strategy-Proofness)*** *Any pseudomarket mechanism with fixed prices and budgets is strategy-proof.*

This result follows because with fixed prices, truthful revelation of an agent's utility implies that the mechanism assigns the agent a probability distribution that maximizes the agent's reported utility. Recall that the feasibility of assignments is assumed in any pseudomarket mechanism. This assumption implies that the prices of some objects may need to be sufficiently high so as to ensure that the demand for these object does not exceed supply. As a consequence of the high prices, the assignments generated by strategy-proof pseudomarkets might be suboptimal and the market-clearing condition from Theorem 2 might fail.

3.2 Asymptotic Incentive Compatibility

Can we still preserve agents' incentives to truthfully reveal their utilities if we endogenize the prices in a way that guarantees market clearing? The answer turns out to resemble that in the standard Walrasian analysis: in general, market-clearing pseudomarkets are not strategy-proof but they are asymptotically incentive compatible in large markets; the intuitive reason for the latter result being that, in large markets, agents become unable to substantially influence prices.[7]

[6] If $\lambda > 0$ and an assignment q^* and prices p^* are in equilibrium for budgets w^*, then the same assignment q^* and prices λp^* are in equilibrium for budgets $\frac{w^*}{\lambda}$. In particular, the set of equilibrium allocations is the same for any $w_1^* = ... = w_{|I|}^* > 0$.

[7] More is known: if there are three or more agents and three or more objects, and each agent's domain of possible utility functions equals the entire $\mathbb{R}_+^{|X|}$, then no mechanism that satisfies strategy-proofness and Pareto efficiency can give the same utility to any two agents with same utility functions. A pseudomarket mechanism with equal budgets satisfies this symmetry condition, and we will see that some such mechanisms are Pareto efficient; hence they cannot be strategy-proof.

We model a large market as a sequence of replica economies. In an n-fold replica of the base economy from Section 2, the set of objects is still X but there are $n|x|$ copies of each object and each agent i with utility u_i from the base economy is replaced by n agents with utility u_i; the set of agents is then denoted I_n.[8] To simplify the exposition, we further assume that the space $U \subseteq \mathbb{R}_+^X$ of possible utilities of an agent is finite and that the space of utility profiles is U^{I_n}. We fix a sequence of budget vectors w^n in which each agent replacing the same agent in the base economy has the same budget. Each utility profile $u_{I_n} \in U^{I_n}$ is associated with a distribution over U, in which the mass put on each $u \in U$ is equal to the fraction of agents in I_n whose utility function is u.

As in standard Walrasian analysis, we establish asymptotic incentive compatibility only for regular economies. To define them, we use the metric $\rho(\mu, \nu) = \max_{u \in U^{I_n}} |\mu(u) - \nu(u)|$ to measure the distance between two distributions μ and ν on U^{I_n}. A distribution of utilities μ^* is ***regular*** if there exists a neighborhood B of μ^* in the metric space just defined and a non-empty finite set of continuous mappings $p(\cdot)$ from distributions $\mu \in B$ to market-clearing pseudomarket price vectors $p(\mu)$ satisfying the cheapest purchase property such that, for every $\mu \in B$, (i) every market-clearing pseudomarket price vector satisfying the cheapest purchase property equals to one of the $p(\mu)$ and (ii) no two price vectors $p(\mu)$ are the same. We say that an economy is regular when the associated distribution of utilities is regular. Replica economies based on Example 1 give us an example of regular utility distributions.

Example 5. In an n-fold replica of the economy from Example 1, there are two objects x and y such that $|x| = n$ and $|y| = 2n$, and $3n$ agents such that $2n$ agents prefer x to y and the remaining agents prefers y to x. Let $U = \{1, 2, ..., 10\}^{\{x,y\}}$ be the space of possible utilities, and suppose that agents of the first type have utility $u(x) = 7 > u(y) = 3$ while agents of the second type have utility $u(x) = 2 < u(y) = 8$. Each agent's budget is $w_i = 1$. For each such replica economy, prices $p_1 = 2$ and $p_2 = 0$ and the assignment that gives the first type of agents $\begin{pmatrix} .5 \\ .5 \end{pmatrix}$ and the second type of agents $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ are in equilibrium.

For a fixed $\varepsilon > 0$, an ε-ball around the distribution putting mass $\frac{2}{3}$ on utility function $(7,3)$ and mass $\frac{1}{3}$ on $(2,8)$ consists of distributions putting mass in $\left(\frac{2}{3} - \varepsilon, \frac{2}{3} + \varepsilon\right)$ on utility function $(7,3)$, mass in $\left(\frac{1}{3} - \varepsilon, \frac{1}{3} + \varepsilon\right)$ on $(2,8)$, and mass ε or lower on other utility functions. For reasons discussed in Example 1, the equilibrium prices only depend on the mass $\mu(x \succ y)$ of agents preferring x to y, the mass $\mu(x \prec y)$ of agents preferring y to x, and the mass $\mu(x \sim y)$ of agents who are indifferent. Suppose the ε-ball we look at is sufficiently small so that object x is scarce, $\mu(x \succ y) > \frac{1}{3}$, while object y is abundant, $\mu(x \prec y) + \mu(x \sim y) < \frac{2}{3}$. Then, for any distribution in this ball, the equilibrium price vectors satisfy $p_y = 0$ and $p_x = 3(\mu(x \succ y) + \lambda\mu(x \sim y))$ for some

[8] Notice that if prices p^* and assignments $q_1^*, ..., q_{|I|}^*$ are in equilibrium for budgets $w_1, ..., w_{|I|}$ then the same prices and assignments in which each agent replacing i receives q_i^* are in equilibrium for budgets such that each agent replacing i has budget w_i.

$\lambda \in [0,1]$; note that prices $p_x > 1$ reflect the scarcity of object x. The equilibrium assignments at those prices depend on λ in the following way: agents preferring y and a fraction λ of indifferent agents obtain probabilities $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ while the remaining agents obtain probabilities $\begin{pmatrix} \frac{1}{p_x} \\ 1 - \frac{1}{p_x} \end{pmatrix}$. The cheapest purchase property is satisfied only by equilibria with price vector in which $\lambda = 0$.

In the analysis of agents' incentives in large markets we focus on pseudomarket mechanisms that select equilibria with prices that depend only on the distribution of utility functions, like in the example above, and not on which agent has which utility function. In equilibrium, the prices and the agent's budget determine the agent's resulting utility, and hence—with prices and budgets determined—the equilibrium assignment selected by the pseudomarket mechanism does not impact agents' incentives. When talking about a mapping of the utility distributions to equilibrium prices, this focus allows us to talk about one mapping instead of a sequence of pseudomarket mechanisms.

A sequence of pseudomarket mechanisms—and the corresponding mapping from the utility distribution to prices—is ***asymptotically incentive compatible*** on a sequence of replica economies if the maximum utility gain of an agent from submitting a utility profile different from the truth vanishes along the sequence. That is, for every $\varepsilon > 0$, there exists n^* such that $n > n^*$ implies that the utility gain from unilateral misreporting for every agent in the n-fold replica is bounded by ε when everyone else is truth-telling. Notice that the utility gain from unilateral misreporting is uniformly bounded for all agents but the bound might depend on the utility distribution.

The pseudomarket mechanisms constructed in Example 5 are asymptotically incentive compatible. Exercise 2 asks you to extend the price construction beyond the ball discussed in Example 5 in such a way that these pseudomarkets are strategy-proof.[9]

Theorem 6. ***(Asymptotic Incentive Compatibility)*** *For every fixed budget vector on the base economy, there exists a pseudomarket mechanism that is asymptotically incentive compatible on any sequence of replica economies with regular base economy. Furthermore, we may require the pseudomarket mechanism to clear the market and satisfy the cheapest purchase property for regular economies.*

Proof. In view of the discussion above, to prove Theorem 6, it is sufficient to construct a mapping from utility distributions to market-clearing equilibrium prices. Take any regular distribution μ^* and notice that its regularity implies that there is an open neighborhood B of μ^* and a finite set of distinct continuous mappings $p(\cdot)$ from distributions $\mu \in B$ to market-clearing pseudomarket price vectors $p(\mu)$ satisfying the cheapest purchase property and that

[9] The strategy-proofness is not a general property but a by-product of there being just two objects in this example.

every market-clearing pseudomarket price vector satisfying the cheapest purchase property lies on one of these price mappings. Let $B(\mu^*)$ be the union of all open sets containing μ^* for which this regularity property holds. The regularity property extends to the entire $B(\mu^*)$. Because any regular distribution has an open neighborhood of regular distributions, any distribution from the topological boundary of $B(\mu^*)$ is not regular. Sets $B(\mu^*)$ constitute thus an open partition of the set of regular distributions. On each $B(\mu^*)$ we can thus separately pick a continuous mapping $p(\cdot)$ from distributions $\mu \in B(\mu)$ to market-clearing pseudomarket price vectors $p(\mu)$ satisfying the cheapest purchase property. The union of these mappings maps all regular distributions to such equilibrium prices and is continuous at any regular distribution. For non-regular distributions we can set any market-clearing equilibrium prices, whose existence is guaranteed by Theorem 2 and Remark 3.

To conclude the proof, consider the n-fold replica of an economy with regular utility distribution μ. Let $\mu_{(i,u)}$ be the utility distribution if some agent i submits a report u instead of u_i while other agents report truthfully. The ρ-distance between μ and $\mu_{(i,u)}$ is $\frac{1}{n|I|}$. By taking sufficiently large n we can ensure that this distance is arbitrarily small. As price function p^* is continuous at μ we can further infer that for large n the impact of agent i's deviation on prices is arbitrarily small. Because agents' utilities are continuous in prices, the theorem follows. **QED**

Note that the constructed pseudomarket mechanism satisfies market clearing and the cheapest purchase property on all economies, not only on regular ones.[10] Some of the assumptions made simplified the setting but played no substantive role in the argument. For instance, instead of studying replica economies at any sequence of economies such that the utility distribution converges to a regular distribution, we can also allow the budget vector to depend continuously on the distribution of utilities and we can relax the assumption that U is finite. For general U, we would use the Prohorov metric to measure the distance between two distributions, μ and ν.

3.3 Preference Reporting

The analysis so far assumed that agents are able to report their preferences, a standard assumption in mechanism design. In the context of no-transfer allocation, the competing mechanisms, mentioned at the beginning of the chapter, only require reporting ordinal preferences, which is prima facie easier than reporting cardinal utilities. Not only cardinal utilities contain more information—and hence may require more effort to learn them for mechanism participants—but they might be harder to conceptualize and communicate than ordinal ranking. How a market participant is to conceptualize their value of an object such as a medical transplant or a school assignment in the context in which these objects are not usually evaluated in terms of money? Utilities in our model reflect

[10]The restriction to regular economies plays a role in the extensions of Theorem 6 discussed later.

comparisons among lotteries, and the mechanism might elicit these comparisons directly. The list of all lotteries is however much longer than the list of all sure outcomes.

The issue was experimentally evaluated in the context of course assignment at the Wharton School at the University of Pennsylvania. The conclusion of the experiment was that reporting cardinal preferences entails more errors than reporting ordinal ones, but that despite these errors pseudomarkets performed better than status quo mechanism used back then by Wharton.[11] In effect, the course assignment at Wharton is now run via a pseudomarket mechanism.

4 Efficiency

Pseudomarkets have good efficiency properties and their efficiency is the key reason one might want to use them to assign objects. A feasible assignment q is ex-ante Pareto efficient—or, simply, ***efficient***—if no other feasible assignment is weakly preferred by all agents and strictly preferred by some agents. Note that when an efficient assignment is represented as a lottery over feasible deterministic assignments, then all deterministic assignments in the support of the lottery are efficient as well.

4.1 Efficiency of Pseudomarkets

Not all pseudomarkets are efficient. The high-price pseudomarkets from Theorem 2 do not need to be. However, by selecting equilibria satisfying market clearing and the cheapest purchase property, we can ensure that the pseudomarket is efficient, obtaining an analogue of the First Welfare Theorem familiar from the Walrasian analysis.

Theorem 7. ***(Efficiency of Pseudomarkets)*** *Let (q^*, p^*) be an equilibrium for some budgets $w^* \in \mathbb{R}_+^{|I|}$ and suppose that the market clearing and cheapest purchase properties hold. Then, q^* is efficient.*

The proof follows the steps of the standard Walrasian analysis with the cheapest purchase property replacing the assumption that agents prefer to have more money to less money.

Proof. By way of contradiction, suppose $q^* \in \Delta(X)^I$ is not efficient. Then there is an allocation $q \in \Delta(X)^I$ such that $\sum_{i \in I} q_i^x \leq |x|$ and $u_i(q_i) \geq u_i(q_i^*)$ for all $i \in I$, with at least one inequality strict. If agent i is satiated under q_i^*—that is the agent receives probability 1 of her most preferred object—then the cheapest purchase property implies that $p^* \cdot q_i \geq p^* \cdot q_i^*$. If an agent i is not satiated then $p^* \cdot q_i \geq p^* \cdot q_i^*$ by the same argument that works in standard competitive equilibrium theory with non-satiated agents. Indeed, suppose $p^* \cdot$

[11] The preferences over course assignments are harder to report than preferences over schools because the course preferences are over bundles of objects (courses) as opposed to single objects (schools). We should hence be careful in making inferences for single-unit demand problems, like school choice, from a multi-unit demand problem like course assignment. See also the discussion of multi-unit demand pseudomarkets in Section 5.2 and Chapter III.9).

$q_i < p^* \cdot q_i^*$ and let $r_i \in \Delta(X)^I$ put probability 1 on a most preferred object of agent i. We can then find a small weight $\alpha > 0$ such that $u_i(\alpha r_i + (1-\alpha) q_i) > u_i(q_i^*)$ and $p^* \cdot (\alpha r_i + (1-\alpha) q_i) \leq p^* \cdot q_i^* \leq w_i^*$, contradicting that q_i^* was an optimal choice in i's budget set. Agents i for whom the inequality $u_i(q_i) \geq u_i(q_i^*)$ is strict are not satiated and the strict-preference analogue of the above argument gives us $p^* \cdot q_i > p^* \cdot q_i^*$.

Summing up the inequalities we derived above, we conclude that $\sum_{i \in I} p^* \cdot q_i > \sum_{i \in I} p^* \cdot q_i^*$. There is thus an object x with positive price $p_x^* > 0$ and such that $\sum_{i \in I} q_i^x > \sum_{i \in I} q_i^{*x}$; for this object, market clearing implies that $\sum_{i \in I} q_i^{*x} = |x|$. Thus $\sum_{i \in I} q_i^x > |x|$ contradicting the assumption that $\sum_{i \in I} q_i \leq |x|$. **QED**

Market clearing and the cheapest purchase property are both needed to guarantee the efficiency of a pseudomarket assignment. Without market clearing, assigning some of the unassigned supply might raise an agent's utility and hence lead to a Pareto dominant assignment. Without the cheapest purchase property, an indifferent agent buying an unnecessarily expensive distribution might impose a negative externality on other agents.[12]

4.2 Pseudomarkets' Efficiency Edge over Ordinal Mechanisms

Efficiency is the main advantage of pseudomarkets over ordinal mechanisms. Consider the following comparison between pseudomarkets and Random Priority, a standard ordinal mechanisms also known as Random Serial Dictatorship (cf. Chapters I.2, II.6, II.9, and II.11).

Example 8. Consider three agents $1, 2, 3$ and three objects with one copy each. Take $\epsilon \in \left(0, \frac{1}{4}\right)$ and suppose that agent 2 valuations are $v_2^3 = 1$, $v_2^2 = 1 - \epsilon$, and $v_2^1 = \epsilon$ while agents $i = 1, 3$ valuations are $v_i^3 = 1$, $v_i^2 = 2\epsilon$, and $v_i^1 = \epsilon$. We may think of these objects as schools: school 3 is the most popular and agent 2 lives in the neighborhood of school 2 while the other agents do not.

Random Priority assigns to each agent probability $\frac{1}{3}$ of each school. This assignment is not efficient. Indeed, each agent would strictly prefer the following assignment: agent 2 receives probability 1 of school 2 and each of the remaining two agents receives probability $\frac{1}{2}$ of school 1 and probability $\frac{1}{2}$ of school 3.

The latter—Pareto dominant—assignment can be implemented as the outcome of a pseudomarket mechanism with prices $p_1 = 2$, $p_2 = 1$, and $p_3 = 0$ and budgets $w_1 = w_2 = w_3 = 1$.

The efficiency advantage illustrated in this example remains true for ordinal mechanisms other than Random Priority. For ordinal mechanisms that treat

[12] Consider, for instance, an economy with one unit of object x, one unit of object y, and two agents: one agent strictly prefers x to y, while the other agent is indifferent between the two. There is then an inefficient equilibrium—that violates the cheapest purchase property—in which each agent has budget of 1, the price of good x is 2, the price of good y is 0, and both agents buy probability .5 in good x and probability .5 in good y. See also Example 5 where the pseudomarket equilibria with index $\lambda > 0$ violate the cheapest purchase property and are inefficient.

agents symmetrically, the advantage is easy to see in the setting of the example, but the per-capita advantage hinges neither on the symmetry of Random Priority nor on the other particulars of the environment. Similar examples can be constructed for assignment markets of any size and schools with arbitrary number of seats. Indeed, the average utility gain of pseudomarket mechanisms over ordinal ones—even the per-capita utility gains—can be arbitrarily large in sufficiently large markets.

What are the measures of the utility gains? We could measure the gains by looking at differences in willingness to pay. The school assignment problem gives us also other measures of utility gain: a difference in an agent's utility from being assigned to one school versus another can be evaluated in terms of how much longer the travel time, or distance, to the preferred school would need to be in order to make an agent indifferent between the two schools.

In terms of the distance measure, the per-capita efficiency advantage of eliciting cardinal information in the data from the New York City school district was estimated to be 3.11 miles; that is, the improvement in moving from assignments based on only ordinal rankings to assignments based on cardinal information is equivalent to shrinking the distances so that each school is on average 3.11 miles closer to each student attending it.

4.3 Pseudomarket Representation of Efficient Assignments

Can we implement all efficient assignments through pseudomarkets? The positive answer resembles the Second Welfare Theorem of Walrasian analysis and it facilitates discrete mechanism design.

Theorem 9. ***(Pseudomarket Representation of Efficient Assignments).*** *If feasible assignment q^* is efficient, then there exists a budget vector $w^* \in \mathbb{R}_+^{|I|}$ and a price vector $p^* \in \mathbb{R}_+^{|X|}$ such that q^* and p^* constitute an equilibrium with budgets w^*.*

The standard Walrasian argument breaks down in the discrete setting we study. Referring to the sum of individual assignments as the corresponding ***aggregate assignment***, we can recapitulate the standard proof as follows. Let Y be the set of aggregate feasible assignments and Z be the set of aggregate assignments that Pareto dominate a fixed efficient assignment $q^* = (q_i^*)_{i \in I}$ that we want to implement in equilibrium; these sets are disjoint, the concavity of utility assures that Z is convex while the feasibility condition assures that Y is convex. If now some agent $i \in I$ strictly prefers some q_i to q_i^*, then $q = \left(q_i, q_{-i}^*\right)$ Pareto dominates q^* and partial separation gives us $p^* \cdot \left(q_i + \sum_{j \in I \setminus \{i\}} q_j^*\right) \geq w \geq p^* \cdot \sum_{j \in I} q_j^*$. The second inequality can be shown to be an equality, allowing us to set $w_i^* = p^* \cdot q_i^*$.

The next step of the argument relies on the separating hyperplane theorem: for any two disjoint convex sets $Y, Z \subseteq \mathbb{R}^n$ there exists a price vector $p^* \in \mathbb{R}^n$ and budget $w \in \mathbb{R}$ such that $p^* \cdot z \geq w \geq p^* \cdot y$ for each $z \in Z$ and $y \in Y$. We

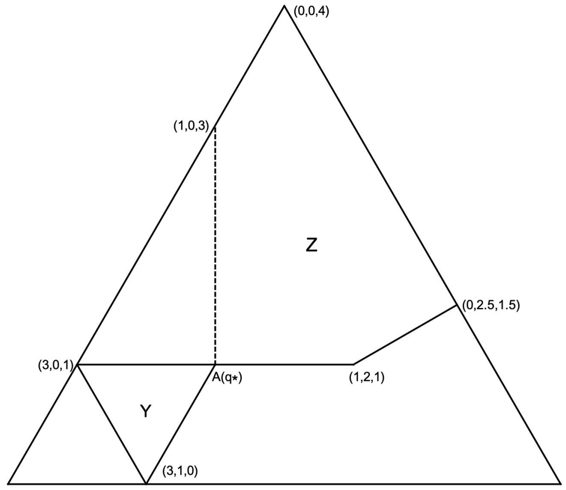

Figure 1: The simplex of (potentially infeasible) aggregate assignments in Example 10. The aggregate assignment $A(q^*)$ is on the intersection of the boundaries of sets Y and Z.

interpret these inequalities as a ***partial separation*** of Y and Z; the ***separation is full*** if one of the inequalities can be assumed to be strict. The separating hyperplane theorem tells us that

$$u_i(q_i) > u_i(q_i^*) \Longrightarrow p^* \cdot q_i \geq w_i^*.$$

This implication is weaker than what's required in equilibrium and it needs to be strengthened to

$$u_i(q_i) > u_i(q_i^*) \Longrightarrow p^* \cdot q_i > w_i^*$$

for all $i \in I$ and for all $(q_i)_{i\in I} \in \Delta(X)^I$. This last step of the standard proof is by contradiction: we take an assignment $q = (q_i)_{i\in I}$ that Pareto dominates q^* while there is an agent i for whom q_i costs the same as q_i^*; in a small open neighborhood of q we then find an assignment that still Pareto dominates q^* while being cheaper than it, contrary to the weak implication above.[13]

In our setting, the standard separating hyperplane theorem partially separates the Pareto dominating aggregate assignments from the feasible ones and we can obtain the weak implication above, but the last step of the above proof fails. The reason is that the last step of the proof relies on the local non-satiation of standard Walrasian agents: for each assignment there is a nearby assignment that they strictly prefer. In effect, the set of aggregate assignments that strictly Pareto dominate an assignment is an open set. Both of these properties fail in our setting. Local non-satiation fails when agents obtain their most preferred object. The failure of openness is illustrated by the following.

[13]The prices p^*, distributions q^*, and budgets w^* satisfying the system of weak implications are called a quasi-equilibrium.

Example 10. There are three objects: object 1 with three copies, and objects 2 and 3 with one copy each. We are assigning these objects to four agents: the odd agents 1 and 3 have von Neumann-Morgenstern utility vector $v = (1, 0, 2)$, and the even agents 2 and 4 have the utility vector $v' = (0, 2, 1)$. Assigning the odd agents the distribution $q_1^* = q_3^* = (\frac{1}{2}, 0, \frac{1}{2})$ and the even agents the distribution $q_2^* = q_4^* = (\frac{1}{2}, \frac{1}{2}, 0)$ is efficient. In particular, the aggregate assignment of q^* is $A(q^*) = \sum_i q_i^* = (2, 1, 1)$. Figure 1 places this point in the barycentric simplex of aggregate assignments. Set Y represents feasible aggregate assignments in the simplex; it is the triangle spanned by $A(q^*), (3, 0, 1)$ and $(3, 1, 0)$. Set Z represents aggregate assignments corresponding to some assignment q in which all agents are weakly better off than under q^* and at least one agent is strictly better off (because q^* is efficient, these assignments q are not feasible). Set Z is a pentagon spanned by five points:

- $(2, 1, 1)$, the aggregate assignment corresponding to q^*,
- $(1, 2, 1)$, the aggregate assignment when the odd agents obtain q^* and the even agents obtain $(0, 1, 0)$,
- $(0, 2\frac{1}{2}, 1\frac{1}{2})$, the aggregate assignment when the odd agents obtain $\left(0, \frac{1}{4}, \frac{3}{4}\right)$ and the even agents obtain $(0, 1, 0)$,
- $(0, 0, 4)$, the aggregate assignment when each agent obtains good 3,
- $(1, 0, 3)$, the aggregate assignment when the odd agents obtain q^* and the even agents obtain $(0, 0, 1)$.

Only the middle three of these five points belong to Z, and one of the borders of Z, the dashed line, is disjoint with Z. Thus, the set Z is neither open nor closed.

The failure of openness of the set of Pareto-dominant assignments breaks the standard proof. In particular in the above example, there is a horizontal hyperplane that partially, but not fully, separates Y and Z; a hyperplane that could not exist in a standard Walrasian analysis. At the same time, in this example there are many hyperplanes fully separating Y and Z.

It turns out that in discrete settings that we discuss, fully separating hyperplanes always exist even though the equivalence of full and partial separation fails. We provide a proof that relies on the following mathematical lemma. In this lemma, a polyhedron is any intersection of closed half-spaces.

Lemma 11. ***(Full Separation Lemma).*** *Let $Y \subset \mathbb{R}^n$ be a closed and convex polyhedron. Let $Z \subset \mathbb{R}^n$ be convex, non-empty, and such that its closure $\bar{Z} \subset \mathbb{R}^n$ is a closed and convex polyhedron. Suppose that $Z \cap Y = \emptyset$ and that, for all $y \in Y \cap \bar{Z}$, $\delta \in \mathbb{R}^n$, and $\varepsilon > 0$, if $y + \delta \in Z$ then $y - \varepsilon\delta \notin \bar{Z}$. Then, there exists a price vector $p^* \in \mathbb{R}^n_+$ and a budget $w \in \mathbb{R}$ such that, for any $z \in Z$ and $y \in Y$, we have $p^* \cdot z > w \geq p^* \cdot y$ and, for any $\bar{z} \in \bar{Z}$ and $y \in Y$, we have $p^* \cdot \bar{z} \geq w \geq p^* \cdot y$.*

In the above example, both sets Y and Z are polyhedra, and other conditions of the lemma are satisfied as well. Given the lemma, the proof of the second welfare theorem hinges on showing that the assumptions of the lemma are satisfied.

Proof of Theorem 9. We denote by $A(q)$ the aggregate assignment $A(q)$ associated with q, that is $A(q) = \sum_{i \in I} q_i$. We write $q \succ q^*$ when q Pareto dominates q^*, that is, when $u_i(q_i) \geq u_i(q_i^*)$ for every $i \in I$ with at least one strict inequality. Let $Z = \{A(q) : q \succ q^*, q \in \Delta(X)^I\}$. If Z is empty then we can support q^* as equilibrium assignment by setting all prices to zero and giving agents arbitrary budgets. Suppose thus that Z is non-empty and notice that Z is convex. Let $\bar{Z}$ be the topological closure of Z, and notice that $\bar{Z}$ is a non-empty convex polytope. Let Y be the set of aggregate assignments $A(q)$ corresponding to feasible assignments q. This set is a closed and convex polytope, and the efficiency of q^* implies that the intersection of Z and Y is empty. To use the full separation lemma, we need the following

Claim. For any $y \in Y \cap \bar{Z}$, $\delta \in \mathbb{R}^{|X|}$ and $\varepsilon > 0$, if $y + \delta \in Z$ then $y - \varepsilon\delta \notin \bar{Z}$.

Proof of the claim: If $y + \delta \in Z$ then there is a $q \succ q^*$ such that $A(q) = y + \delta$. By way of contradiction, assume $y - \varepsilon\delta \in \bar{Z}$. Thus, there is a $\tilde{q} = (\tilde{q}_i)_{i \in I}$ such that $u_i(\tilde{q}_i) \geq u_i(q_i^*)$ for every $i \in I$ and $A(\tilde{q}) = y - \varepsilon\delta$. Then, the random assignment $\bar{q} = \frac{\varepsilon}{1+\varepsilon} q + \frac{1}{1+\varepsilon} \tilde{q}$ is feasible, and the linearity of utilities in probabilities imply that $\bar{q} \succ q^*$. But this contradicts efficiency of q^*, proving the claim.

This claim and the full separation lemma imply that there exists a price vector $p^* \in \mathbb{R}_+^{|X|}$ and a budget $w \in \mathbb{R}$ such that $p^* \cdot z > w \geq p^* \cdot y$, for any $z \in Z$ and $y \in Y$. Since q^* is feasible $\sum_{i \in I} q_i^* \in Y$ and thus $p^* \cdot \sum_{i \in I} q_i^* \leq w$. Furthermore, $p^* \cdot \sum_{i \in I} q_i^* \geq w$ because $q^* \in \bar{Z}$. We conclude $p^* \cdot \sum_{i \in I} q_i^* = w$. Now, if we take some q_i that some agent $i \in I$ strictly prefers to q_i^*, then $q_i + \sum_{j \in I \setminus \{i\}} q_j^* \in Z$, and we have $p^* \cdot \left(q_i + \sum_{j \in I \setminus \{i\}} q_j^*\right) > w = p^* \cdot \left(q_i^* + \sum_{j \in I \setminus \{i\}} q_j^*\right)$. Consequently we have $p^* \cdot q_i > p^* \cdot q_i^*$, proving that p^* and q^* constitute an equilibrium for budgets $w_i^* = p \cdot q_i^*$. **QED**

5 Fairness, Multiple-Unit Demand, Priorities, and Constraints

5.1 Fairness

In addition to efficiency, a key objective in mechanism design is fairness. A frequently invoked fairness objective is envy-freeness. A mechanism is ***envy-free*** if no participant strictly prefers a distribution of another participant to their own. Envy-freeness implies that if two agents have the same utility function then they are indifferent between their respective distributions.

Pseudomarkets can achieve envy-freeness:

Theorem 12. *(**Envy-freeness**). The pseudomarket mechanism in which all agents are endowed with equal budgets is envy-free.*

Indeed, if an agent strictly preferred another agent's distribution, then the outcome would not be in equilibrium because the cost of the preferred distribution would be within the first agent's budget.

Exercise 4 asks you to verify that the converse statement is not true: there are envy-free efficient assignments that cannot be implemented via equal-budget pseudomarkets, even though, as we have seen in Theorem 9, they can be implemented by other pseudomarkets.

5.2 Multi-unit Demand

So far we focused on the single-unit demand environment in which agents demand a probability distribution over objects. With two important exceptions—equilibrium existence and the implementability of random assignments as lotteries over deterministic ones—our analysis can be easily extended to the multi-unit demand environments.

In ***multi-unit demand*** environments, we continue to have a set of agents I and a set of objects X, with each object x represented by a finite number of copies $|x|$. Agents now demand probability distributions over bundles: a ***bundle*** might contain multiple copies of multiple objects. By $B_i \subseteq \times_{x\in X} \{0, 1, ..., |x|\}$ we denote the finite set of admissible individual bundles for agent i. For instance, setting $B_i = \left\{b \in \times_{x\in X} \{0, 1, ..., |x|\} \mid \sum_{x\in X} b^x = 1\right\}$ we can embed the single-unit demand environment as a special case of the multi-unit demand one. A deterministic assignment of bundles $(b_i)_{i\in I} \in \times_{i\in I} B_i$ is feasible if $\sum_{i\in I} b_i^x \leq |x|$. A (random) assignment is feasible if it can be represented as a lottery over feasible deterministic assignments. The definition of pseudomarket equilibrium is the same as before except that distributions over bundles play the role of distributions over objects.

In this multi-unit demand setting we encounter two subtleties. First, in the single-unit demand setting, the Birkhoff-von Neumman Theorem guaranteed the feasibility of every assignment such that the expectation of the number of copies of each object x assigned is weakly lower than the supply of this object, $|x|$. This key property is not guaranteed in the general multi-unit assignment setting.

Example 13. There are three agents $1, 2, 3$ and four objects $x_1, ...x_4$ such that $|x_1| = |x_2| = |x_3| = 1$ and $|x_4| = 2$. Consider the assignment in which: agent i receives probability $\frac{1}{2}$ of bundle $\{x_i, x_{i+1 \mod 3}\}$ and probability $\frac{1}{2}$ of bundle $\{x_4\}$.[14] In expectation $\frac{3}{2}$ copy of object x_4 and one copy of each of the objects $x_1, ..., x_3$ is assigned. However, we can verify that the above assignment cannot be expressed as a lottery over feasible deterministic assignments.

In the special case in which agents' utilities over bundles are additively separable in terms of utilities of objects in the bundle, the non-implementability of

[14] The term $i + 1 \mod 3$ equals $i + 1$ for $i = 1, 2$, and it equals 1 for $i = 3$.

Example 13 can be overcome by assigning to agents distributions over copies of individual objects instead of distributions over bundles; the sum of an agent's probabilities then does not need to equal 1 and hence we refer to them as quantities. Breaking the bundles in this way in the above example leads to each agent i receiving quantity $\frac{1}{2}$ of object x_i, quantity $\frac{1}{2}$ of object $x_{i+1 \mod 3}$, quantity $\frac{1}{4}$ of object x_4, and quantity $\frac{1}{4}$ of object x_5. This broken-up assignment can be implemented as lottery that puts probability $\frac{1}{2}$ on giving each agent i half of object x_i and half of object x_4 and puts probability $\frac{1}{2}$ on giving each agent i half of object $x_{i+1 \mod 3}$ and half of object x_4.

Second, for some budget vectors—including the equal-budget one—the equilibrium existence is not guaranteed.

Example 14. There are four agents $1, 2, 3, 4$ and four objects $x_1, ..., x_4$, each with two copies. Agent i's most preferred bundle consists of the three objects different from x_i and the second most preferred bundle is $\{x_i\}$. We may check that if agents have equal budgets then no pseudomarket equilibrium satisfies market clearing even if we require the assignments to be only feasible in expectation.

The non-existence of market-clearing equilibria can be addressed by adjusting the budgets and allowing for limited failures of market clearing.

The rest of our analysis transfers more easily to the multi-unit setting. The analogues of the incentive Theorems 4 and 6, and the fairness Theorem 12 hold true. The conclusion of the efficiency Theorem 7 can be strengthened: equilibrium assignments are not only efficient but they are also undominated by assignments that are merely feasible in expectation; in turn, for such assignments, the efficiency Theorem 9 also holds true.[15]

5.3 Priorities and Constraints

Assignments may be subject to constraints. A popular type of constraint in single-unit demand assignment is based on endowing objects (e.g. schools) with priorities over agents (cf. Chapter III.2) and requiring the underlying deterministic assignments to ***honor*** these ***priorities*** in the following sense: any object that an agent i prefers to the object assigned has all its copies assigned to agents whose priority at the object is at least as high as that of i. This constraint is known as stability (for more on stability and the related concept of the lack of justified envy see Chapters II.12 and III.2).

Pseudomarkets have been adapted to guarantee that priorities are honored. The adaptation allows for prices of objects to differ across agents in such a way that an agent with weakly higher priority faces a weakly lower price. In equilibrium, there is at most one priority level at which the price can belong to $(0, \infty)$; for higher priority levels the price is 0 and for lower ones the price

[15]The multi-unit versions of these results formally hold true irrespective of whether market-clearing equilibria exist. However, Theorems 6, 12, and 7 are only useful when the relevant equilibria exist.

is prohibitively high, e.g. infinite. The existence and incentive-compatibility results continue to hold for such priority-adjusted pseudomarkets.

Multi-unit demand environments allow for a wide variety of constraints. By allowing for individualized sets of feasible deterministic assignments, the model introduced above already allows for many constraints. E.g. it allows for some seats in a school to be reserved for some type of applicants, while allowing all applicants to compete for the remaining seats. To model such a constraint we create an auxiliary object "reserved seats," which are feasible only for the selected applicants.

In addition, the conjunctions of ***linear constraints***—requiring that a weighted sum of probabilities that agents receive specific bundles is weakly lower or higher than some baseline—are particularly tractable. Under linear constraints, the above incentive and efficiency results continue to hold true. Existence is more subtle: it has been established for some linear constraints and it may require allowing different prices for agents differently impacted by constraints, just like in the case of priority-based constraints above.

6 Exercises

1. Consider the modification of Example 1 discussed in Section 2.3 and show that no assignment q^* and price vector $p^* \in \mathbb{R}_+^X$ satisfies $p^* \cdot q_i^* \leq p^* \begin{pmatrix} \frac{1}{3} \\ \frac{2}{3} \end{pmatrix}$ for all $i \in \{1,2,3\}$ and $u_i(q_i) > u_i(q_i^*) \implies p^* \cdot q_i > p^* \begin{pmatrix} \frac{1}{3} \\ \frac{2}{3} \end{pmatrix}$ for all $(q_i)_{i\in\{1,2,3\}} \in \Delta(\{x,y\})^{\{1,2,3\}}$.

2. In the environment of Example 5, find a strategy-proof pseudomarket mechanism such that the equilibrium price vectors are such that $p_y = 0$ and $p_x = 3\mu(x \succ y)$ for any distribution in the small ball studied therein.

3. Verify the construction of sets Y and Z in Example 10.

4. (i) Construct a single-unit demand environment and an efficient and envy-free assignment that cannot be implemented via an equal budget pseudomarket. (ii) Construct a non-equal budget pseudomarket that implements the assignment from part (i).

5. Verify that the assignment constructed in Example 13 is not feasible.

6. Verify that there does not exist pseudomarket equilibrium in which all agents have equal budgets in the environment of Example 14 even if we require the assignments to be only feasible in expectation.

7. Extend Theorem 7 to the multi-unit demand environment.

7 Notes

Section 2 is based on Hylland and Zeckhauser (1979). Section 3.1 is based on Pycia (2014). Section 3.2 is based on He, Miralles, Pycia, and Yan (2018), except for the impossibility result in footnote 7, which comes from Zhou (1990). For an alternative take on large market incentives, the reader may also want to consult Azevedo and Budish (2019). Section 3.3 is based on Budish and Kessler (2020). Section 4.1 is based on Hylland and Zeckhauser (1979). The example of Section 4.2 is based on Abdulkadiroglu, Che and Yasuda (2011) and Pycia (2014), and the following gain discussion additionally draws on Pycia (2019). The discussion of utility measurement draws on Calsamiglia, Martínez-Mora, and Miralles (2019), and, primarily, on Abdulkadiroglu, Agarwal, and Pathak (2017), who estimated the gain potential in New York City. Section 4.3 is based on Miralles and Pycia (2021); the proof of Lemma 11 might be found in this paper. The theorem of Section 5.1 comes from Hylland and Zeckhauser (1979) while the discussion of the converse comes from Miralles and Pycia (2015). This section focuses on ex-ante envy-freeness (before the realization of lotteries), and for a discussion of ex-post envy-freeness the reader might want to consult Budish (2011). In Section 5.2, the model comes from Budish (2011), the pseudomarket definition and the efficiency results from Miralles and Pycia (2021), and the incentive-compatibility results from author's notes. The discussion of implementability counterexample is based on Budish, Che, Kojima, and Milgrom (2013) and Nguyen, Peivandi, and Vohra (2016); the resolution for additively separable utilities is from Hylland and Zeckhauser (1979) and Budish, Che, Kojima, and Milgrom (2013). The non-existence example comes from Budish (2011) as does its approximation-based resolution. In section 5.3, the discussion of priorities follows He, Miralles, Pycia, and Yan (2018). The discussion of efficiency under linear constraints follows Miralles and Pycia (2021), the discussion of incentives is based on author's notes, and the discussion of existence follows Echenique, Miralles, and Zhang (2021).

References

[1] Abdulkadiroglu, Atila, Nikhil Agarwal, and Parag A. Pathak (2017). "The Welfare Effects of Coordinated Assignment: Evidence from the New York City High School Match." *American Economic Review* 107(12): 3635–3689.

[2] Abdulkadiroglu, Atila, Yeon-Koo Che, and Yasuke Yasuda. (2011) "Resolving Conflicting Preferences in School Choice: the Boston Mechanism Reconsidered." *American Economic Review* 101(1), 399-410.

[3] Arrow, Kenneth J., and Frank H. Hahn (1971). *General Competitive Analysis.* Holden-Day, Oakland.

[4] Azevedo, Eduardo M., and Eric Budish (2019) "Strategyproofness in the Large." *Review of Economic Studies* 86, 81-116.

[5] Budish, Eric (2011). "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes," *Journal of Political Economy* 119, 1061-1103.

[6] Budish, Eric, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom (2013). "Designing Random Allocation Mechanisms: Theory and Applications," *American Economic Review* 103, 585-623.

[7] Budish, Eric, and Judd B. Kessler (2020). "Can Market Participants Report their Preferences Accurately (Enough)?" Working Paper.

[8] Calsamiglia, Caterina, Francisco Martínez-Mora, and Antonio Miralles (2919). "Cardinal Assignment Mechanisms: Money Matters More Than It Should." Working Paper.

[9] Debreu, Gerard (1959). *Theory of Value. An Axiomatic Analysis of Economic Equilibrium.* Yale University Press.

[10] Echenique, Federico, Antonio Miralles, and Jun Zhang (2021). "Constrained Pseudo-Market Equilibrium." *American Economic Review* (Forthcoming).

[11] He, Yinghua, Antonio Miralles, Marek Pycia, and Jianye Yan (2018). "A Pseudo-Market Approach to Assignment with Priorities." *American Economic Journal: Microeconomics* 10(3), 272-314.

[12] Hylland, Aanund, and Richard Zeckhauser (1979). "The Efficient Allocation of Individuals to Positions," *Journal of Political Economy* 87, 293-314.

[13] Kreps, David M. (2013). *Microeconomic Foundations I: Choice and Competitive Markets.* Princeton University Press.

[14] Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995). *Microeconomic Theory.* Oxford University Press.

[15] Miralles, Antonio and Marek Pycia (2015). "Large vs. Continuum Assignment Economies." Working Paper.

[16] Miralles, Antonio and Marek Pycia (2021). "Foundations of Pseudomarkets: Walrasian Equilibria for Discrete Resources." *Journal of Economic Theory* 196 (105303), 1-24.

[17] Nguyen, Thành, Ahmad Peivandi, and Rakesh Vohra (2016). "Assignment problems with complementarities." *Journal of Economic Theory* 165, 209-241.

[18] Pycia, Marek (2014). "The Cost of Ordinality." Working Paper.

[19] Pycia Marek (2019). "Evaluating with Statistics: Which Outcome Measures Differentiate Among Matching Mechanisms?" Working Paper.

[20] Zhou, Lin (1990). "On a Conjecture by Gale a One-Sided Matching Problems." *Journal of Economic Theory* 52, 123-135.