Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper 14-030

"One-Sided Matching with Limited Complementarities"

by

Thanh Nguyen, Ahmad Peivandi and Rakesh Vohra

One-Sided Matching with Limited Complementarities*

Thành Nguyen[†], Ahmad Peivandi[‡], Rakesh Vohra[§]

March 2014

Abstract

The problem of allocating *bundles* of indivisible objects without transfers arises in the assignment of courses to students, of computing resources like CPU time, memory and disk space to computing tasks and the truck loads of food to food banks. In these settings the complementarities in preferences are small compared with the size of the market. We exploit this to design mechanisms satisfying efficiency, envy-freeness and asymptotic strategy-proofness.

Informally, we assume that agents do not want bundles that are too large. There will be a parameter k such that the marginal utility of any item relative to a bundle of size k or larger is zero. We call such preferences k-demand preferences. Given this parameter we show how to represent probability shares over bundles as lotteries over *approximately* (**deterministic**) feasible integer allocations. The degree of infeasibility in these integer allocations will be controlled by the parameter k. In particular, *ex-post*, no good is over allocated by at most $k-1$ units.

1 Introduction

The problem of allocating *bundles* of indivisible objects without transfers arises in the assignment of courses to students (Budish [2011]), of computing resources like CPU time, memory and disk space to computing tasks (Gutman and Nisan [2012]), of truck loads of

*This paper is a merger of Peivandi [2012] and Nguyen and Vohra [2013].

[†]Purdue University, nguye161@purdue.edu

[‡]Northwestern University, peivandi@u.northwestern.edu

[§]University of Pennsylvania , rvohra@seas.upenn.edu

food to food banks (Houlihan [2006]), siblings to schools (Abdulkadiroğlu et al. [2006]) and couples to hospital residency positions (Kojima et al. [2013], Ashlagi et al. [2014]).

Many of the methods for allocating indivisible goods proposed and studied are limited to the case of unit demand- where each agent wishes to consume at most one object. They can be divided into two groups. In the first are mechanisms that specify a lottery over outcomes, while in the second, mechanisms specify probability shares in objects.

Examples of the first type of mechanisms are random serial dictatorship (RSD) and top-trading with random endowments (TTC) (Abdulkadiroğlu and Sönmez [1998], Hashimoto [2013]). Neither mechanism explicitly randomizes over each possible outcome given the large number of possible outcomes. Instead, they specify a procedure for assigning goods to agents from a randomly chosen starting point.[1] These methods are typically strategy-proof and Pareto optimal but lack other desirable properties like ordinal efficiency and envyfreeness. The second type of mechanism specifies probability shares in objects rather than lotteries over feasible outcomes is popular. Under the unit demand assumption, there is, by virtue of the Birkhoff-von Neuman theorem, an equivalence between probability shares and lotteries over feasible outcomes. As probability shares are in a sense 'easier' to specify, these mechanisms produce outcomes with many more desirable properties than either RSD or TTC. Examples of mechanisms in this group are Probabilistic Serial (PS) and Competitive Equilibrium with Equal Incomes (CEEI) (Bogomolnaia and Moulin [2001], Hylland and Zeckhauser [1979]).[2] Compared with the first type of mechanisms, the second type usually possess stronger efficiency and equity properties but are generally not strategy-proof. Nevertheless, as shown in Azevedo and Budish [2012] for example, the second type of mechanism are asymptotically strategy-proof. Generalizing these results to settings where agents' preferences are over bundles is difficult because the

[1] In RSD, for example, agents are randomly assigned a priority ordering which determines who gets to choose first. In TTC agents are randomly assigned a good which they can then trade with others.

[2] The PS mechanism determines probability shares by having agents nibble away at objects at a constant speed in order of preference until their supply is exhausted. The CEEI mechanism determines probability shares from a Walrasian equilibrium assuming equal incomes and that the goods are divisible.

equivalence between probability shares in objects and lotteries over outcomes does not hold.

Generalizations of the PS and CEE mechanism have been proposed (see Che and Kojima [2010]) and (Kojima [2009]). With some exceptions these generalizations rarely inherit the attractive features of their antecedents unless one rules out complementarities in preferences. Ruling out complementarities in preferences is problematic as they are an essential feature of applications such as course allocation, assigning siblings to schools. However, in many applications the 'degree' of these complementarities are small compared with the size of the market. In this setting we design mechanisms in which agents have preferences over bundles, satisfying efficiency, envyfreeness and asymptotic strategy-proofness.

Informally, we assume that agents do not want bundles that are too large. There will be a parameter k such that the marginal utility of any item relative to a bundle of size k or larger is zero. We call such preferences k-demand preferences.[3]

Given this parameter we show how to represent probability shares over bundles as lotteries over *approximately* (**deterministic**) feasible integer allocations. The degree of infeasibility in these integer allocations will be controlled by the parameter k. In particular, *ex-post*, no good is over allocated by at most $k-1$ units. Thus, these mechanisms will be relevant when k is small relative to the available supply of each good or when the resource constraints are 'soft', i.e., permit small violations. An alternative solution for hard capacity constraints is to hold $k-1$ units of each good in reserve to make up for a shortfall.[4]

One setting where k will be small relative to available supply is (University) course allocation. Each good is a course and the available supply of each good is the number of

[3]Our definition of k-demand is actually more general, which also allows for preferences over large bundles, that are additive over several small bundles.

[4]When $k = 1$ we recover the well known equivalence between probability shares in goods and lotteries over outcomes in the unit demand case.

seats in the classroom in which the course will be conducted. Typically a student is not able to take more than 5 courses in any term, so $k = 5$. This means, *ex-post*, we might over allocate at most 4 seats per class. In a classroom with 20 seats, this can easily be accommodated by adding 4 seats.

Upon this representation theorem we build a general method for obtaining mechanisms for allocating bundles of indivisible objects which have a variety of desirable properties. We describe two applications below.

1. Assuming cardinal k-demand preferences, we exhibit an *ex-ante* envy-free, asymptotically efficient and asymptotically strategy-proof mechanism where no good, *ex-post*, is overallocated by more than $k-1$ units.

2. Assuming ordinal k-demand preferences we generalize the PS mechanism to obtain a mechanism for allocating bundles that is envy-free, ordinally efficient, asymptotically strategy-proof where no good, *ex-post*, is overallocated by more than $k-1$ units.

The chief virtue of this method is that it allows the designer to specify the outcome in terms of probability shares in bundles. As noted earlier, it gives one greater control over the outcomes. Second, it allows for a succinct description of the mechanism (recall that the set of possible outcomes is significantly larger than the number of possible bundles that an agent can receive). Third, in a precise sense, it allows for efficient computation.

In the next section we introduce notation, the setting and the precise restriction on preferences we impose. Subsequently we state the main result. The following sections describe applications of this result and contrasts them with the relevant literature.

2 Notation and Approximate Implementation

As noted earlier the equivalence between probability shares and lotteries relies on the Birkhoff-von Neuman theorem. This section introduces an approximate generalization of the Birkhoff- von Neuman theorem that accommodates complementarities in preferences.

In the combinatorial assignment problem we have a set N of agents and a set G of goods. For each $j \in G$, the available supply of good j is an integer s_j. A bundle is captured by a non negative vector $S \in \mathbb{N}^{|G|}$, where the j^{th}-coordinate S_j indicates the number of copies of good j in the bundle S. The size of a bundle S, denoted as $|S|$, is defined as the total number of items in S, i.e., $\sum_{j \in G} S_j$.

Agent i is interested in obtaining at most one bundle. Here we will assume that the maximum size of a single bundle is at most k. In the course allocation problem, for example, students are agents, each good j corresponds to a course with the number of available seats being s_j. Each student requires at most 1 seat in each class. In practice students can only consume a bundle of size at most 5, so $k = 5$. In the problem of assigning couples to hospital residency positions $k = 2$. Each bundle consists of 2 positions in the same hospitals or in two different but nearby hospitals.

To describe the set of feasible allocations of objects to agents let $0 \leq x_i(S) \leq 1$ denote the probability that agent i obtains bundle S. Here each agent is only interested in bundles of size at most k, thus we assign $x_i(S) = 0$ for all S of size larger than k. First, because each agent receives one bundle of goods (possibly empty), any allocation must satisfy the following constraint.

$$\begin{aligned} x &\geq 0 \\ x_i(S) = 0 &\text{ if } |S| > k \qquad \text{(DEMAND)} \\ \sum_S x_i(S) &= 1 \ \forall i \in N \end{aligned}$$

Feasibility requires that for each type of good j we do not allocate more than its available supply:

$$\sum_{i \in N} \sum_{S \ni j} S_j \cdot x_i(S) \leq s_j \ \forall j \in G \tag{SUPPLY}$$

Define approximate supply constraints:

$$\sum_{i \in N} \sum_{S \ni j} S_j \cdot x_i(S) \leq s_j + k - 1 \ \forall j \in G \tag{SUPPLY+k-1}$$

Call a fractional solution x to (DEMAND-SUPPLY) *implementable* if it can be expressed as a convex combination of feasible integer solutions to (DEMAND-SUPPLY). Feasible integer solutions of (DEMAND-SUPPLY) correspond to feasible allocations of indivisible bundles (integer allocations) to agents. An implementable fractional solution can be interpreted as a lottery over feasible allocations. Generally, an x satisfying (DEMAND-SUPPLY) is not implementable. Our main result, below, is that every x satisfying (DEMAND-SUPPLY) can be implemented as a lottery over integer solutions to (DEMAND) and (SUPPLY+k-1).

Our result is the following:

Theorem 2.1 *Any (fractional) solution of (DEMAND-SUPPLY) can be implemented as a lottery over integral allocations that satisfy (DEMAND) and (SUPPLY+k-1).*

Proof:

To prove Theorem 2.1 we will need the following lemma.

Lemma 2.2 *Given any (not necessarily non-negative) utility vector $u_i(S)$ and any fractional vector x satisfying (DEMAND) and (SUPPLY), we can find in polynomial time an* integral *vector x^* satisfying (DEMAND) and (SUPPLY+k-1) such that $u \cdot x^* \geq u \cdot x$.*

The proof of Lemma 2.2 is provided in Appendix B.1. It is a slight extension of a recent result in Combinatorial Optimization Király et al. [2012]. Notice that in Király et al. [2012], it is assumed that s_j is either 0 or 1. Our proof does not require such an assumption.

Given Lemma 2.2, the proof of Theorem 2.1 is as follows.

For ease of exposition, let Q be the polytope consisting of all *real* vectors satisfying (DEMAND) and (SUPPLY); and let E_k be the set of *integral* solutions to (DEMAND) and (SUPPLY+k-1).

Suppose Theorem 2.1 does not hold. Then, there is an $x \in Q$ that is not in the convex hull of E_k. Hence, there exists a hyperplane that separates x from E_k. Let u be the vector of coefficients of that hyperplane. We can choose it so that $ux > uz$ for all $z \in E_k$ which contradicts Lemma 2.2.

Notice, the proof of Theorem 2.1 is nonconstructive. However, based on the standard Ellipsoid method in convex optimization, given any $x \in Q$; one can implement x as a lottery over integral solutions of E_k in polynomial time. In Appendix B.2, we provide a practical polynomial time algorithm to construct a lottery whose expectation is arbitrarily close to the given vector x in Q.

■

3 Maximizing Social Welfare

In this section we introduce a general mechanism (called OPT) when agents have cardinal preferences over bundles. As discussed in the introduction, in the absence of transfers, identifying an integer allocation in (DEMAND-SUPPLY) that maximizes social welfare subject to fairness and incentive comparability is difficult. The absence of a numeraire good, like money, makes it difficult to discourage agents from claiming an excessively large utility for their most preferred bundle of objects. To overcome this issue, we introduce (interim) envy-freeness as linear constraints into the program (DEMAND-SUPPLY). We then show how Theorem 2.1 can be used to find an approximately feasible allocation so as to maximize weighted total social welfare. Finally, we prove that when the economy is large, this allocation is actually asymptotically strategy-proof.

3.1 k-demand Preferences

We first formally define k-demand preferences. If $u(S)$ is an agent's utility for the bundle $S \in \mathbb{N}^{|G|}$, we require $u(\emptyset) = 0$ and that $u(\cdot)$ have one of the following properties. The first is that no agent has preferences for bundles that are too large, i.e.

$$u(S) = 0 \ \forall |S| \geq k + 1. \tag{1}$$

The second is the monotone cover of the first that allows for free disposal and ensures monotonicity.

$$u(S) = \max_{A \leq S:} \{u(A) : |A| \leq k\} \ \text{ for } S \in \mathbb{N}^{|G|}. \tag{2}$$

One setting where 1 or 2 will hold is course assignment. Each good $j \in G$ is a course, s_j is the number of seats in the course and each $i \in N$ is a student. There is an upper limit, k, on the number of courses any student can take. k is usually small relative to s_j for each $j \in G$. For example, k will be at most 4 while s_j is usually at least 20 and frequently much larger.

Our analysis extends to the case where agents have preferences for large bundles. Suppose a partition $P_1, \ldots, P_t$ of G such that $|P_r| \leq k$ for all $r = 1, \ldots, t$. Then,

$$u(S) = \sum_{r=1}^{t} u(S \cap P_r). \tag{3}$$

One instance of this preferences may arise is when the objects to be allocated are bands of spectrum. Bands of spectrum that interfere have similar frequency and are located close to one another and so can be categorized in groups of small size. Utilities for interfering bands need not be additive. Bands further apart do not interfere, so utilities for non-interfering bands can be taken to be additive.

When $k = 1$, (1-2) yield unit demand preferences. When $k = 1$, under (3), utilities are additive. Thus, k-demand preferences for $k = 1$ contain the basic classes of preferences that are substitutable. When $k = |G|$, k-demand preferences impose no restriction on preferences. Intermediate values of k restrict the range of complementarities that can be expressed.

3.2 Mechanism OPT

To implement OPT, choose a positive weight w_i for each agent i and solve the linear program of maximizing $\sum_{S \subseteq G} w_i \cdot u_i(S)x_i(S)$ subject to (DEMAND-SUPPLY) and the envy-free condition defined below.

An allocation x satisfying (DEMAND-SUPPLY) is envy-free if

$$\sum_{S \subseteq G} u_i(S)x_i(S) \geq \sum_{S \subseteq G} u_i(S)x_j(S) \ \forall i \ \forall j \neq i. \tag{ENVY-FREE}$$

Formally the mechanism OPT is defined as follows.

Definition 3.1 *Given positive weights $w_i, i \in N$, let x^* be an optimal solution of*

$$\max\{\sum_{i \in N} \sum_{S \subseteq G} w_i \cdot u_i(S)x_i(S) : \ s.t. \ (DEMAND), (SUPPLY), (ENVY - FREE)\}, \tag{LIP}$$

(If there are multiple solutions, select one with a fixed tie-breaking rule.)

By Theorem 2.1, x^* can be implemented as a lottery over integral assignments satisfying (DEMAND) and (SUPPLY+k-1). The mechanism OPT takes as input a report of each agents utility function and returns the optimal (fractional) solution to program (LIP) and implements it as a lottery.

In the section below we show that if (LIP) has a unique optimal solution, one can implement x^* in a way that is asymptotically strategy-proof.[5] Thus, under this condition mechanism OPT is approximately efficient subject to *ex-ante* envy-free, and asymptotically strategy-proof.

Mechanism OPT has several attractive properties. It easy to implement (can be computed in time polynomial in $|N|$, $|G|$ and k), furthermore it gives the designer control over the outcome through selection of the weights w_i. For example, by choosing $w_i = 1$ for all i, mechanism OPT selects an allocation of maximum social welfare among the envy-free outcomes. For another choice of the w_i's, mechanism OPT can implement the CEEI mechanism of Hylland and Zeckhauser [1979] (as well as a generalization).

To see this, consider the original CEEI mechanism, where agents are restricted to unit demands. In the CEEI mechanism, agents are endowed with equal amounts of a budget of fictitious money and a competitive equilibrium of that economy is determined. Under the unit demand restriction, the Birkhoff-von Neumann theorem implies that the competitive equilibrium allocation can be implemented as a lottery. It is well known that the CEEI mechanism is pareto efficient and *ex-ante* envy-free. Thus, there exists positive weights $\{w_i\}_{i \in N}$ (Negishi weights) such that this competitive equilibrium allocation maximizes weighted social welfare subject to *ex-ante* envy-freeness. Hence, OPT with the Negishi weights can implement the CEEI mechanism.

The idea extends to the case of of k-demand preferences. Namely, if we allow for the fractional assignment of bundles of objects, and give each agent the same fictitious budget, market-clearing prices exist (see Appendix A for a short proof based on an appeal to the celebrated results of Arrow-Debreu-McKenzie). This fractional assignment while feasible may not be implementable. Under Theorem 2.1 it can be implemented so that it over allocates each good by at most $k - 1$ units. This is a generalization of the CEEI mechanism, which we call the bundled competitive equilibrium from equal income **B-CEEI**

[5] The assumption that (LIP) has a unique optimal solution is a mild one because we can always guarantee this by perturbing w_i slightly.

mechanism. Notice that an equilibrium allocation of bundles can always be obtained by maximizing a suitable weighted sum of utilities subject to (DEMAND), (SUPPLY), furthermore an outcome of a competitive equilibrium with equal budgets is envy-free. Thus, our mechanism OPT with the proper weights w_i can implement the outcome of the B-CEEI mechanism. From this construction, it is easy to see that the B-CEEI mechanism is approximately efficient, *ex-ante* envy-free. Furthermore, as we show later, it is also asymptotically strategy-proof.

It is instructive to compare the B-CEEI mechanism to Budish's (Budish [2011]) generalization of the CEEI mechanism, call it the A-CEEI mechanism. It is a *deterministic* mechanism for the combinatorial assignment problem based on computing an approximate competitive equilibrium from approximately equal incomes. Thus, the preference information required of agents is ordinal rather than cardinal as in our case. Budish's mechanism returns an allocation that is approximately efficient, approximately envy-free in an *ex-post* sense, asymptotically strategy-proof and, like ours, violates the resource constraints. Budish bounds the violation in terms of the Euclidean distance $(O(\sqrt{\min\{2k, |G|\}|G|}))$ between the supply vector and the vector of the number of goods allocated, unlike the bound on the maximum violation in each type of goods considered in this paper. The two bounds are not comparable.

Recently, Akbarpour and Nikzad [2014] also considered an approximate implementation for problems, where several quotas are imposed on agents and objects. Their model restricts agents to linear utilities over objects, and the approximation results only hold probabilistically rather than ex-post.

Hashimoto [2013] proposes a generalization of the RSD mechanism that also yields an outcome that is approximately feasible. The error bound is multiplicative and so weaker than ours which is additive.

In the next subsection we introduce notation to define precisely what is meant by asymptotic strategy-proofness as well as prove that the mechanism just described is asymptotically strategy-proof.

3.3 Asymptotic Strategy-Proofness

We assume each agent $i \in N$ has one of a finite number of types Θ. For each type $\theta \in \Theta$, let n_θ be the number of agents of type θ. The type of an agent encodes their preferences which are represented by a von-Neuman Morgenstern **utility function** defined on bundles of goods. In our main applications (course allocation) we assume, for simplicity only, that each agent wishes to consume at most one copy of each good.

For an agent of type $\theta \in \Theta$, let $u^\theta(S) \geq 0$ be his utility for bundle $S \subseteq \mathbb{N}^{|G|}$. We will also use the notation $u_i^\theta(S)$ (or for short $u_i(S)$) for the utility of agent i for bundle S when his type is θ. We assume that an agent's utility depends exclusively on his type and outcome. Furthermore, we assume for each type of agent, the utility function satisfies either (1), (2) or (3). Without loss of generality, we also assume $u^\theta(\emptyset) = 0$ for all type θ. Given a **lottery** (a probability distribution) over a set of bundles, an agent's utility is his expected utility from the lottery.

Given a type profile an allocation is **envy-free** if all agents weakly prefer the lottery assigned to them to any lottery assigned to another. That is,

$$u^{\theta_i}(x_i) \geq u^{\theta_i}(x_j).$$

Let A denote the set of (approximately) feasible allocations. Recall that N is the set of agents. For every $|N| > 0$, a mechanism $\Phi^{(N)}$ is a mapping from a profile of agents' types to a lottery over (approximately) feasible allocations. More precisely,

$$\Phi^{(N)} : \Theta^N \to \Delta(A).$$

Without ambiguity, we sometimes use Φ instead of $\Phi^{(N)}$ for short.

It will be useful to consider a mechanism from the perspective of an agent i. Let

$$\Phi_i^{(N)} : \Theta \times \Theta^{N-1} \to \Delta(A_i),$$

where A_i denotes the possible bundles that agent i obtains, and $\Phi_i(\theta_i, \theta_{-i})$ denotes the lottery over bundles that agent i receives when he reports θ_i and other agents report θ_{-i}.

A mechanism is Φ is **strategy-proof** if it is optimal for each agent to truthfully report their type given any vector of type reports of the other agents, that is

$$u^{\theta_i}(\Phi_i(\theta_i, \theta_{-i})) \geq u^{\theta_i}(\Phi_i({\theta'}_i, \theta_{-i})).$$

A mechanism is ϵ**- strategy-proof** if it is "almost'' optimal for each agent to report truthfully given any vector of reports by the other agents, that is

$$u^{\theta_i}(\Phi_i(\theta_i, \theta_{-i})) \geq u^{\theta_i}(\Phi_i({\theta'}_i, \theta_{-i})) - \epsilon.$$

Finally, we define asymptomatic strategy-proofness. Φ is **asymptotically strategy-proof** if for any $\epsilon > 0$ there exists a constant n_0 such that Φ is ϵ-strategy-proof whenever there are at least n_0 agents reporting r to Φ for each every type $r \in \Theta$.

Definition 3.2 Φ *is* **asymptotically strategy-proof** *if for any* $\epsilon > 0$ *there exists a constant* n_0 *such that if* $\vec{\theta}$ *satisfies* $|\{i|\theta_i = r\}| \geq n_0$ *for all* $r \in \Theta$*, then*

$$u^{\theta_i}(\Phi_i(\theta_i, \theta_{-i})) \geq u^{\theta_i}(\Phi_i({\theta'}_i, \theta_{-i})) - \epsilon.$$

Our definition of asymptotic strategy-proofness is similar in spirit to the notion of 'strategy-proofness in the large' introduced by Azevedo and Budish [2012]. To define this notion assume agents' reports are drawn *independently* from a distribution over the type set Θ with full support. A mechanism is strategy-proof in the large if it is ϵ strategy-proof when the number of agents is large enough. In fact, any mechanism that is asymptotically strategy-proof in our sense will also be strategy-proof in the large.

Theorem 3.3 *Set w_i to be the same for each type of agents and suppose (LIP) has an unique optimal solution x^* implementable by a lottery $\bar{x}$. Then, the mechanism OPT that takes as input a report of each agents type and returns $\bar{x}$ is asymptotically strategy-proof.*

Proof: Recall that the set of types, Θ, is finite. Let θ_i be the type reported by agent i, and let n_θ be the number of agents reporting type θ. Also let w_{θ_i} be the weight that the mechanism chooses for an agent of type θ_i.

Consider the following program for finding a utilitarian allocation that is envy-free:

$$\max \sum_{i \in N} \sum_{S \subseteq G} w_{\theta_i} \cdot u^{\theta_i}(S) x_i^{\theta_i}(S) \tag{4}$$

$$\sum_{S \subseteq G} x_i^{\theta_i}(S) \leq 1 \ \forall i \in N \tag{5}$$

$$\sum_{i \in N} \sum_{S \ni j} x_i^{\theta_i}(S) \leq s_j \ \forall j \in G \tag{6}$$

$$\sum_{S \subseteq G} u^{\theta_i}(S) x_i^{\theta_i}(S) \geq \sum_{S \subseteq G} u^{\theta_i}(S) x_j^{\theta_j}(S) \ \forall i, j \tag{7}$$

Recall, that we can set $x_i^{\theta_i}(S) = 0$ whenever $|S| > k$. Call (4-7) the disaggregate formulation.

Introduce variables $y^\theta(S)$ to denote the 'aggregate' amount of bundle S that all agents reporting θ get. Namely if we consider an anonymous solution of (4-7), that is $x_i^{\theta_i}(S) = x_j^{\theta_j}(S)$ whenever $\theta_i = \theta_j = \theta$, then $y^\theta(S) = n_\theta x_i^{\theta_i}(S)$. Now consider the following 'aggregate' formulation.

$$\max \sum_{\theta \in \Theta} \sum_{S \subseteq G} w_\theta \cdot u^\theta(S) y^\theta(S) \tag{8}$$

$$\sum_{S \subseteq G} \frac{1}{n_\theta} y^\theta(S) \leq 1 \ \forall \theta \in \Theta \tag{9}$$

$$\sum_{t\in\Theta}\sum_{S\ni j} y^\theta(S) \leq s_j \ \forall j \in G \tag{10}$$

$$\frac{1}{n_\theta}\sum_{S\subseteq G} u^\theta(S)y^\theta(S) \geq \frac{1}{n_{\theta'}}\sum_{S\subseteq G} u^\theta(S)y^{\theta'}(S) \ \forall \theta,\theta' \in \Theta. \tag{11}$$

To show that our mechanism is asymptotically strategy-proof, we need to prove that for every $\epsilon > 0$, there exists n_o such that if $n_\theta \geq n_0$ for all $\theta \in \Theta$, then no agent can improve his utility by more than ϵ.

Suppose agent i of type p pretends to be of type q. We will show that the impact on the allocations of the other agents from this misreport can be computed by solving (8-11) with a perturbed right hand side.

If agent i of type p pretends to be of type q then the number of agents reporting p is decreased by one and the number of agents reporting q is increased by 1. Let n'_θ be the number of agents reporting type θ in this case, then

$$n'_p = n_p - 1; n'_q = n_q + 1; n'_\theta = n_\theta \quad \forall \theta \neq p, q$$

Thus, the aggregate program becomes

$$\max \sum_{\theta\in\Theta}\sum_{S\subseteq G} w_\theta \cdot u^\theta(S)y^\theta(S) \tag{12}$$

$$\sum_{S\subseteq G} \frac{1}{n'_\theta} y^\theta(S) \leq 1 \ \forall \theta \in \Theta \tag{13}$$

$$\sum_{\theta\in\Theta}\sum_{S\ni j} y^\theta(S) \leq s_j \ \forall j \in G \tag{14}$$

$$\frac{1}{n'_\theta}\sum_{S\subseteq G} u^\theta(S)y^\theta(S) \geq \frac{1}{n'_{\theta'}}\sum_{S\subseteq G} u^\theta(S)y^{\theta'}(S) \ \forall \theta,\theta' \in \Theta. \tag{15}$$

Compare program (12)-(15) to program (8)-(11). If both n_p and n_q are large enough then the objective function and the constraints of both program are close to each other.

Thus, as n_p and n_q go to infinity, the maximum value of (12)-(15) converges to the maximum value of (8)-(11). Furthermore, because of the assumption that (8)-(11) has an **unique** maximizer, the solution of (12)-(15) will converge to that unique maximizer, otherwise the maximum value of (12)-(15) would not converge to the maximum value of (8)-(11).

Thus, there exists n_0 such that if both n_p and n_q are at least n_0 then (12)-(15) also has a unique solution. Furthermore, as n_0 increases the solution of (12)-(15) is converging to the solution of (8)-(11). In other words, if n_0 is large enough, the agent who misreports their type can only change their allocation by $O(\epsilon)$. Thus, by the envy-free constraint, their utility changes by at most $O(\epsilon)$. This shows that the mechanism is asymptotically strategy-proof according to Definition 3.2. ■

4 Generalizing the Probabilistic Serial Mechanism

Mechanism OPT required that agents communicate cardinal preferences. This is sometimes criticized as impractical. Hence, in this section we turn our attention to mechanisms that rely on ordinal information alone. Our goal is to generalize the well known Probabilistic Serial PS mechanism for allocating indivisible goods when agents have strict preferences and unit demands (introduced by Bogomolnaia and Moulin [2001]). The PS mechanism begins with each agent consuming, at the same constant rate, their most preferred object. When the supply of an object is exhausted, agents consuming that object switch to consuming the next available object on their preference list. At termination, the fraction of each object an agent has consumed determines their probability shares in the relevant object. These probability shares can be implemented as a lottery over feasible allocations. It is well known that the PS mechanism is envy-free, ordinally efficient and asymptotically strategy-proof. We define ordinal efficiency for the case when agents have preferences over bundles rather than single objects.

Assume, for this section only, that agents have strict preferences over bundles. Let $\prec_i$ be agent i's ordinal preference ranking over acceptable bundles. As each agent receives a lottery over allocations we extend a preference ordering over bundles to a *partial* ordering over lotteries of bundles via stochastic dominance. Recall that a lottery over allocations induces probability shares x over bundles that satisfy (DEMAND) and (SUPPLY). Thus, we may identify each lottery with a solution of (DEMAND) and (SUPPLY).[6] An allocation x satisfying (DEMAND) and (SUPPLY) *weakly* stochastically dominates an allocation y for agent i, if for all $B \subseteq G$:

$$\sum_{S \succ_i B} x_i(S) \geq \sum_{S \succ_i B} y_i(s).$$

Allocation x stochastically dominates y for agent i, if the above inequality holds strictly for some bundles S. A mechanism is *ordinally efficient* if there is no other random assignment that weakly stochastically dominates the mechanism's allocation with respect to all agents preferences over bundles.

As preferences in this section are ordinal, the notion of strategy-proofness and envy-freeness from Section 3.3 must be modified. An ordinal mechanism is *strategy-proof* if for any agent, the allocation resulting from misreporting is stochastically dominated by the allocation from truthfull reporting, with respect to agent's true preference. A mechanism is envy-free if for all agents, the allocation assigned to him stochastically dominates all other agents' assignments, with respect to his preference. A mechanism is *weakly strategy-proof* if for each agent, his allocation from truthful reporting is not stochastically dominated by the allocation produced by a misreport, with respect to his true preference. A mechanism is *weakly envy-free* if no agent's allocation is stochastically dominated by the allocation of another agent.

Ours is not the first paper to extend the PS mechanism beyond the unit demand case.

[6]Note that each solution of (DEMAND) and (SUPPLY) does not correspond to a lottery over allocations.

See for example, Kojima [2009] and Budish et al. [2013]. Our generalization of the PS mechanism differs from these papers in the kind of complementarities in preferences we allow. Those papers assume that agents rank lotteries over assignments based on first order stochastic dominance on *single* objects. As the example below shows, in their setting, an agent with responsive cardinal preferences may prefer a utility-dominated (as defined in those papers) lottery. This assumption allows these papers to abstract away from the implementability problem caused by complementarities in ordinal preferences.

Example: There are two objects a and b each with two copies. Agent i has the following cardinal preference for bundles:

$$u_i(\{a,a\}) = 6, u_i(\{a,b\}) = 5, u_i(\{b,b\}) = 2, u_i(\{a\}) = 1, u_i(\{b\}) = 0.5$$

The ordinal preference associated with this cardinal preference is responsive. Consider the following two lotteries:

A: agent i receives bundle $\{a,a\}$ with probability half and bundle $\{b,b\}$ with probability half.

B: agent i receives bundle $\{a,b\}$ with probability 0.99 and $\{b,b\}$, with probability 0.01.

Under the preferences defined in Kojima [2009] and Pycia [2011], agent i prefers lottery A to lottery B, since under lottery A agent i has a higher chance of receiving copies of object a. However, agent i has a higher expected utility for lottery B.

4.1 Natural Probabilistic Serial Mechanism

A natural generalization of the PS mechanism for non-unit demand is to have agents consume bundles rather than individual objects. When the supply of an object is exhausted agents switch to their most preferred bundle not containing objects whose supply has been exhausted. This mechanism, which we call the Natural Probabilistic Serial (NPS) mechanism, returns probability shares in bundles. Therefore, it produces outcomes that, in general, cannot be implemented. Kojima [2009] and Budish et al. [2013] circumvent this difficulty by assuming preferences satisfy first order stochastic dominance on *single* objects. In this restricted case, the NPS mechanism reduces to the PS mechanism. However, if we assume preferences satisfy the k-demand condition, we may invoke Theorem 2.1. Under these conditions the NPS mechanism is envy-free, ordinally efficient, asymptotically strategy-proof and overallocates each good by at most $k-1$ units. These results along with a formal description of the NPS mechanism are stated below.

To define the NPS mechanism formally, let $t(0) = 0 \leq t(1) \leq t(2) \leq t(v) \leq t(v+1)...$ be the instances in time when the supply of at least one good falls to zero. At time $t(v)$, the set of goods in non-zero-supply is denoted $G(v)$. Initially, $G(0) = G$. Let $m_j(v)$ be the total number of copies of object j in all agents' most preferred bundles (from among bundles with objects from $G(v)$). Moreover, let z^v be a deterministic allocation where each agent is allocated his best most preferred bundle of objects in $G(v)$.[7] Use z_j^v to denote the number of copies of object j consumed in the allocation z^v. The NPS mechanism can be represented by the following steps:

- Starting with available supply of $G(v-1)$, the latest time at which the current supply of good j would be exhausted is

$$t_j(v) = sup\{t \in [0,1] | z_j^{v-1} + m_j(v-1)(t - t(v-1)) \leq n_j\}.$$

- Therefore, the first instance at which any good is exhausted is $t(v) = \min_{j \in G(v-1)} t_j(v)$.

- At time $t(v)$, the set of goods with non-zero supply is

$$G(v) = G(v-1) \setminus \{j \in G(v-1) | t_j(v) = t(v)\}.$$

[7] Agents are allowed to receive multiple copies of an object in $G(v)$.

- $$z_j^v = z_i^{v-1} + m_j(v-1)(t_j(v) - t_j(v-1))$$

- The allocation returned by the NPS mechanism at time $t(v)$ is

$$NPS(v) = NPS(v-1) + (t(v) - t(v-1))z^{v-1}.$$

- NPS terminates at time $t(v)$ where v is the smallest index such that $t(v) = 1$.

Theorem 4.1 *The NPS mechanism is envy-free, weakly strategy-proof and under k-demand preferences can be implemented so that it overallocates each good by at most* $k-1$ *units.*

Proof: The first two items admit a proof similar to the proof of Theorem 1 and Proposition 1 in Bogomolnaia and Moulin [2001]. The last item follows from Theorem 2.1 as the probability shares produced by NPS satisfy (DEMAND) and (SUPPLY). ■

Though the NPS mechanism is not in general strategy-proof, it possesses this property asymptotically. To discuss the behavior of the NPS mechanism in large economies, we formally define large economies. The definition given here of the $q-$economy is similar to the definition in Che and Kojima [2010]. For each $q \in \mathbb{N}$, the set of objects in the $q-$economy is G and the set of agents is N_q. Each object $j \in G$ in the $q-$ economy has $s_j^q \geq k$ copies. Furthermore, $\lim_{q\to\infty} s_j^q = \infty$ for all $j \in G$. The set of agents, N_q, is partitioned into r subsets, Π_θ^q for $1 \leq \theta \leq r$. Agents in the set Π_θ^q are said to have type θ. Agents with the same type θ have the same preference ranking over bundles in the $q-$ economy.

We assume that the number of copies of each object and the number of agents of each type grow at the same rate as q.

Assumption 4.1 *For some positive real numbers* $(n_j^*)_{j\in G}$ *and non-negative real numbers* $(m_\theta^*)_{1\le\theta\le r}$ *the following holds:* $\lim_{q\to\infty}\frac{n_j^q}{q}=n_j^*$ *and* $\lim_{q\to\infty}\frac{|\Pi_q^\theta|}{q}=m_\theta^*\in\mathbb{R}$, $\forall 1\le\theta\le r$.

Definition 4.2 *Allocation* x $\epsilon-$*stochastically dominates* y *for agents* i *with respect to ordinal preference* $\succ_i$ *if,*

$$\sum_{S\succeq_i B} x_i(S)+\epsilon\sum_{S\succeq_i B} y_i(S).$$

An ordinal mechanism is asymptotically strategy-proof if for all $\epsilon>0$ *there exists* $Q>0$ *large enough such that: for all agents and* $q>Q$*, the allocation that is returned from truthful reporting* $\epsilon-$*stochastically dominates any misreport in the* $q-$*economy.*

Theorem 4.3 *Under k-demand preferences the NPS mechanism has the following properties:*

1. *It is asymptotically implementable.*

2. *It produces ordinally efficient probabilistic allocations.*

3. *It is weakly strategy-proof.*

4. *It is envy-free.*

5. *It is asymptotically strategy-proof.*

Proof: Propostition 4.1 states the first three properties. The fourth property follows from an asymptotic equivalence between the NPS mechansim and the RSD mechanism that generalizes Che and Kojima [2010]. To prove it, we first define a continuum economy as the limit of the q-economy as $q\to\infty$. Let NPS* be the NPS mechanism in the continuum economy. We show that the NPS and RSD mechanisms converge to NPS*. The difficulty in our case, compared with Che and Kojima [2010], arises because of the complementarities in agents' preferences. In the unit demand case, once consumption of an objects begins, it continues until exhaustion. In our case, consumption of an object occurs in fits and starts. Therefore, a simple adaptation of their proof is not possible. We show equivalence by proving that the available supply of each object at each step of the mechanism converges. For a complete proof, see Appendix C .[8] ∎

References

A. Abdulkadiroğlu and T. Sönmez. Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems. *Econometrica*, 66(3):689–702, May 1998.

A. Abdulkadiroğlu, P. Pathak, A. E. Roth, and T. Sönmez. Changing the boston school choice mechanism. Nber working papers, National Bureau of Economic Research, Inc, Jan. 2006.

M. Akbarpour and A. Nikzad. Approximate random allocation mechanisms. *Working Paper*, 2014.

I. Ashlagi, M. Braverman, and A. Hassidim. Stability in large matching markets with complementarities. *Operations Research, forthcoming*, 2014.

E. M. Azevedo and E. Budish. Strategyproofness in the large as a desideratum for market design. In *ACM Conference on Electronic Commerce*, page 55, 2012.

A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, October 2001.

[8]In the unit demand case, the PS and RSD mechanisms were shown to be asymptotically equivalent by Che and Kojima [2010]. Liu and Pycia [2013] generalized this result to show that any two mechanisms that satisfy ordinal efficiency, symmetry and asymptotic strategy-proofness coincide asymptotically. PS and RSD are two examples of mechanisms in this class. Pycia [2011] generalized this result to the case of agents with multi-unit demand assuming that agents rank lotteries over assignments based on first order stochastic dominance on *single* objects.

E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061 – 1103, 2011.

E. Budish, Y.-K. Che, F. Kojima, and P. Milgrom. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2):585–623, September 2013.

Y.-K. Che and F. Kojima. Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms. *Econometrica*, 78(5):1625–1672, 09 2010.

A. Gutman and N. Nisan. Fair allocation without trade. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '12, pages 719–728, 2012.

T. Hashimoto. The generalized random priority mechanism with budgets. *Working Paper*, 2013.

P. Houlihan. How food banks came to love the free market. *Chicago Booth Magazine*, pages 34–39, 2006.

A. Hylland and R. Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political Economy*, 87(2):293–314, April 1979.

T. Király, L. Lau, and M. Singh. Degree bounded matroids and submodular flows. *Combinatorica*, 32(6):703–720, 2012. ISSN 0209-9683. doi: 10.1007/s00493-012-2760-6.

F. Kojima. Random assignment of multiple indivisible objects. *Mathematical Social Sciences*, 57(1):134–142, 2009.

F. Kojima, P. A. Pathak, and A. E. Roth. Matching with couples: Stability and incentives in large markets. *The Quarterly Journal of Economics*, 128(4):1585–1632, 2013.

Q. Liu and M. Pycia. Ordinal efficiency, fairness, and incentives in large markets. *Working Paper*, 2013.

T. Nguyen and R. Vohra. The allocation of indivisible objects via rounding. *Working Paper*, 2013.

A. Peivandi. Random allocation of bundles. *Working paper, Northwestern University*, 2012.

M. Pycia. Ordinal efficiency, fairness, and incentives in large multi-unit-demand assignments. *Working Paper*, 2011.

A Appendix A

A.1 Equilibrium Existence in B-CEEI

In this section we show how to use standard Arrow-Debreu-Mckenzie arguments to establish the existence of market clearing prices in the Hylland and Zeckhauser setting when agents have non-unit demands. Recall that an agent $i \in N$ who consumes a fraction $x_i(S)$ of bundle $S \subseteq G$ enjoys a utility of $\sum_{S \subseteq G} u_i(S)x_i(S)$ where $u_i(S)$ is the utility derived from bundle S. Equivalently, if we give agent i the vector $z \in [0,1]^{|G|}$, where z_j is interpreted to be the fraction of good j, agent i's utility can be represented as

$$U_i(z) = \max \sum_{S \subseteq G} u_i(S)x_i(S)$$

$$\text{s.t.} \sum_{S \ni j} x_i(S) \leq z_j \ \forall j \in G$$

$$x_i(S) \geq 0 \ \forall i \in N, \forall S \subseteq G$$

Notice, $U_i(z)$ is concave in z. It follows immediately from the usual arguments that there must exists prices that clear the market.

B Appendix B

B.1 Proof of Lemma 2.2

Recall that we have the following constraints.

$$\begin{aligned} x &\geq 0 \\ x_i(S) &= 0 \text{ if } |S| > k \\ \sum_S x_i(S) &= 1 \ \forall i \in N \end{aligned} \tag{DEMAND}$$

Feasibility requires that for each type of good j we do not allocate more than its available supply:

$$\sum_{i \in N} \sum_{S \ni j} S_j \cdot x_i(S) \leq s_j \ \forall j \in G \tag{SUPPLY}$$

Define approximate supply constraints:

$$\sum_{i \in N} \sum_{S \ni j} S_j \cdot x_i(S) \leq s_j + k - 1 \ \forall j \in G \tag{SUPPLY+k-1}$$

The proof uses the following algorithm called the Iterative Rounding Algorithm (IRA). Let P be the polytope defined by $(DEMAND)$ and $(SUPPLY)$. The IRA takes as input an extreme point, $x^* \in \arg\max\{u \cdot x : x \in P\}$ where $u \geq 0$ and $u_i(S) = 0$ for all $i \in N$ and $S \subseteq G$ such that $|S| > k$. It then rounds x^* into a 0-1 vector $\bar{x}$ that satisfies (DEMAND) and (SUPPLY+k-1).

Beginning with x^*, we remove from (DEMAND-SUPPLY) all variables $x_i(S)$ for which

$x_i^*(S) = 0$. In other words, a variable that is zero in x^* will be rounded down to zero and fixed at that value in all subsequent iterations. Similarly, remove from (DEMAND-SUPPLY) all variables $x_i(S)$ for which $x_i^*(S) = 1$ and adjust the right hand sides of (SUPPLY) accordingly. In other words, a variable set to 1 (or 0) by x^* is fixed at 1 (or 0) in all subsequent iterations. In the system that remains pick a non-negative extreme point that optimizes the vector u and repeat. At some iteration, when the remaining supply of good j is s'_j, we may obtain an extreme point with no variable set to 1. Call it y.

The main observation here is that, in this case there must exist a $j \in G$ such that

$$\sum_{i \in N} \sum_{S \ni j} S_j \cdot \lceil y_i^*(S) \rceil \leq s'_j + k - 1.$$

For each such j, remove the corresponding constraint (SUPPLY) and in the relaxed system find an extreme point that optimizes u and repeat. Stop once all variables have been fixed at either 0 or 1 and denote the resulting 0-1 vector by $\bar{x}$.

There are three observations to be made about $\bar{x}$.

1. At each iteration, inequality (DEMAND) holds. Thus, $\bar{x}$ satisfies (DEMAND).

2. At each iteration, the original program is (possibly) relaxed. Thus, $u \cdot \bar{x} \geq u \cdot x^*$.

3. Because $\bar{x}_i(S) = 1$ only if $x_i^*(S) > 0$, it follows that for the inequalities in (DEMAND) thrown away, $\sum_{i \in N} S_j \cdot \sum_{S \ni j} \bar{x}_i(S) \leq s_j + k - 1$.

We need to show that if we have not found an integral solution in the iterative process, then we will find a constraint j that can be thrown away. In particular, we need to prove the following.

Lemma B.1 *Let $u_i(S)$ be any utility function. Let $\mathcal{S}_0(i), \mathcal{S}_1(i)$ be the set of bundles that $x_i(S)$ forall $S \in \mathcal{S}_0(i)$ have been fixed to be 0, and $x_i(S)$ for all $S \in \mathcal{S}_1(i)$ have been fixed to be 1, respectively.*

Let x^* be an extreme point the linear program

$$\max\{u{\cdot}x : x \text{ satisfies (DEMAND) and (SUPPLY)}, x_i(S) = 0 \forall S \in \mathcal{S}_0(i); x_i(S) = 1 \forall S \in \mathcal{S}_1(i)\}.$$

Assume that $x_i^*(S) < 1$ for all $i \in N$ and S such that $S \notin \mathcal{S}_0(i) \cup \mathcal{S}_1(i)$. (In other words, $x_i^*(S)$ has not been fixed). Then, there exists a $j \in G$ such that

$$\sum_{i \in N} \sum_{S: j \in S} S_j \lceil x_i^*(S) \rceil \leq s_j + k - 1.$$

This lemma will guarantee that if we cannot round all the variable, then we can find a constraint to throw away and resolve the linear program. To prove Lemma B.1, we will use the following property of an extreme point of a linear program:

> The number of non-zero variables in an extreme point x^* is equal to the number of linearly independent and binding constraints in (DEMAND) and (SUPPLY).

To prove the lemma, assume the contradiction that for all $j \in G$ $0 < x_i^*(S) < 1$ for all i, S that $x_i(S)$ has not been fixed to be 0 or 1, and $\sum_{i \in N} \sum_{S \in \mathcal{S}: j \in S} S_j \lceil x_i^*(S) \rceil > s_j + k - 1$. Because $\sum_{i \in N} \sum_{S: j \in S} S_j \lceil x_i^*(S) \rceil$ is an integral value, thus we have

$$\sum_{i \in N} \sum_{S: j \in S} S_j \lceil x_i^*(S) \rceil \geq s_j + k. \tag{16}$$

we will derive a contradiction to the property of the extreme point above.

Given the extreme point x^*, where we credit each non-zero variable $x_i^*(S)$ with a single token. We then redistribute these tokens to the binding, linearly independent constraints in a particular way. We show that if (16) holds then each binding constraint will get at least one token, and there is one token left over. This shows that the number of non-zero variable $x_i^*(S)$ is larger than the number of binding, linearly independent constraints, which is a contradiction.

We redistribute the tokens given as follows. Credit $x_i^*(S)$ fraction of the tokens to the constraint corresponding to agent i (DEMAND). Credit $S_j \frac{1-x_i^*(S)}{k}$ to each constraint corresponding to each good $j \in S$. Notice that this is feasible because the size of each bundle is $\sum_{j \in G} S_j \leq k$.

If the constraint corresponding to agent i binds then the number of tokens this constraint is credited with is $\sum_S x_i(S) = 1$. Now, consider a binding constraint corresponding to good j, we have.

$$\sum_{i \in N} \sum_{S \ni j} S_j x_i(S) = s_j. \tag{17}$$

The total tokens that this constraint obtains is

$$\sum_{S, i \in N : x_i^*(S) > 0} S_j \frac{1 - x_i^*(S)}{k} = \frac{1}{k} \sum_{S, i \in N : x_i^*(S) > 0} S_j - \frac{1}{k} \sum_{i \in N} \sum_{S \ni j} S_j x_i(S).$$

From (16) and (17) this number of tokens is at least

$$\frac{1}{k}(s_j + k - s_j) = 1.$$

Thus, any binding constraint j (DEMAND) is credited with at least 1 token.

Hence, we have shown that the amount of tokens given at the beginning (which is the number of non-zero x^* variables) has been redistributed to the binding constraints, so that each is credited with at least 1 token. Thus the number of non-zero x^* variables is at least the number of binding constraints.

Now, the equality obtains only if for every nonzero $x_i^*(S)$, the size of bundle S is exactly k, that is $\sum_{j \in G} S_j = k$. Furthermore, the constraint corresponding to agent i as well as all the constraint corresponding to for all $j \in S$ bind. However, this case one can show that the set of binding constraints is not linearly independent. To see this, consider the sum of all the binding constraints in SUPPLY:

$$\sum_{j \in G} \sum_{i,S:S \ni j} S_j x_i^*(S) = \sum_{j \in G} s_j.$$

Because for each $x_i^*(S) > 0$, $\sum_{j \in S} S_j = k$, this sum can be rewritten as

$$k \cdot \sum_{i,S} x_i^*(S) = \sum_{j} s_j.$$

This last expression is the sum of all the constraints in (DEMAND), contradicting linear independence of the binding constraints. By this we have shown that the number of nonzero variables in an extreme point solution is larger than the number of linearly independent binding constraints. ■

B.2 An Algorithm To Construct a Lottery

Recall that Theorem 2.1 shows that any $x \in Q$ can be expressed as a convex combination of points in E_k. In this section we show how to (approximately) decompose any $x \in Q$ into a convex combination of points in E_k.

Assume E_k is bounded with diameter D. Denote by $|x - y|$ the Euclidean distance between x and y. Recall that we have a subroutine that will for any fractional $x \in Q$ and any cost vector c, return an integral $\bar{x} \in E_k$ such that $c\bar{x} \geq cx$.

Given this subroutine, we exhibit a polynomial time algorithm that for a given point $x \in Q$, finds at most $d+1$ integral points in E_k whose convex hull is arbitrarily close to x. The algorithm also returns a lottery over these $d+1$ integral vectors whose expectation is close to x.

Given a fractional solution $x \in Q$. Let

$$B(x, \delta) = \{z : \text{ satisfying } (DEMAND) \text{ and } |z - x| \leq \delta\}$$

We assume there exists $\delta > 0$ such that $B(x, \delta) \subset Q$. Notice that for our purpose, this assumption is without loss of generality, because otherwise we can always choose x' in the interior of Q close to x.

Given an allowable error $\epsilon > 0$, the algorithm is the following.

Algorithm In each step maintain a subset S of points in E_k. Each iteration consists of the following steps.

1. Compute $y \in conv(S)$ that is closest to x. If $|y - x| < \epsilon$, the algorithm terminates.
2. Otherwise, because y is the closest point to x in S, y lies in a hyperplane of $conv(S)$. Thus, there exists a subset $S' \subset S$ of size at most d such that $y \in conv(S')$. (Recall d is the dimension).

   Consider $z = x + \delta \frac{x-y}{|x-y|}$. Notice, $z \in Q$ because $B(x, \delta) \in Q$. Use the rounding algorithm to find an integral $z' \in E_k$, such that

   $$< z, x - y > \leq < z', x - y > .$$

3. Update $S := S' \cup \{z'\}$; and repeat.

To show that the algorithm terminates in polynomial time, we show that after each iteration, the distance $|x - y|$ is reduced by at least a constant factor. To prove this, let y' be the point in the interval $[z', y]$ that is closest to x. We will prove the following.

Claim B.2 *There exists $0 < \gamma < 1$ that depends on D and δ such that $|x - y'| < (1 - \gamma)|x - y|$*

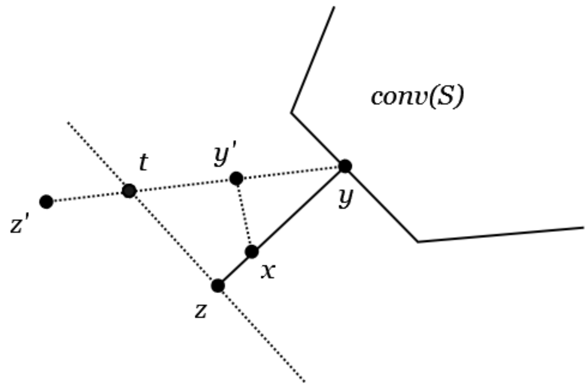

Figure 1:

Proof: Let t be the point in the interval (z', y) such that $< t - z, x - y >= 0$. Because $< z, x - y > \leq < z', x - y >$, such a t exists. See Figure 1.

Now,

$$\frac{|x-y|^2}{|x-y'|^2} = \frac{|t-y|^2}{|t-z|^2} = \frac{|t-z|^2 + (|x-z| + |x-y|)^2}{|t-z|^2} \geq \frac{|t-z|^2 + \delta^2}{|t-z|^2}$$

We have $|t - z| \leq |z' - z|$. Furthermore, because the diameter of E_k is D, $|z' - z| \leq D$. Thus, $|t - z| \leq D$.

Hence, we obtain

$$\frac{|x-y|^2}{|x-y'|^2} \geq \frac{D^2 + \delta^2}{D^2}$$

Thus, there exists $0 < \gamma < 1$, depending on D and δ such that

$$|x - y'| < (1 - \gamma)|x - y|,$$

which is what we need to prove. ∎

The claim above shows that after each iteration the distance between x and y is reduced by at least a factor of $(1-\gamma)$. Consider $K = \frac{\ln(D/\epsilon)}{\gamma}$, we have

$$D(1-\gamma)^K \leq \epsilon,$$

Thus, after at most K iterations, the algorithm will terminate.

C Appendix C

The proof is adapted from Che and Kojima [2010]. First we define a *continuum* economy, which is the natural candidate for what the limit of a q-economy might be as $q \to \infty$. For each object $j \in G$, there is a mass s_j^* of this object. The set of agents, N^*, is an interval of real numbers partitioned into d intervals $(\Pi_\theta^*)_{1 \leq \theta \leq d}$. Each point in N^* corresponds to an agent. For each $1 \leq \theta \leq d$, the length of Π_θ^* is m_θ^*. Agents find bundles of size more than k unacceptable. The set of all bundles of size no more than k is denoted $\mathbf{B}^*$. For each $1 \leq \theta \leq d$, agents in Π_θ^* have the same preference ranking of bundles in $\mathbf{B}^*$ as do agents with type θ in a $q-$ economy.

An allocation in the continuum economy is a function $x : N^* \times \mathbf{B}^* \to [0,1]$. The allocation is deterministic if the range of x is $\{0,1\}$. A deterministic allocation is implementable if (i) agents are allocated to at most one bundle, and (ii) for all objects $j \in G$ the measure of agents allocated to bundles that include object j (including all copies of a in all bundles) does not exceed s_j^*. That is, $\int_{N^*} \sum_{S \in \mathbf{B}^*} x_i(S)|S|_j di \leq s_j^*$ for all $j \in G$. Here $|S|_j$ is the number of copies of object j in the set S.

An allocation in the continuum economy is implementable if it can be represented as a probability distribution over implementable deterministic allocations. Similarly, we can define feasibility for allocations: x is feasible if

$$\int_{N^*} \sum_{S \in \mathbf{B}^*} x_i(S)|S|_j di \leq s_j^*$$

for all $j \in G$ and for all $i \in N^*$, $\sum_{S \in \mathbf{B}^*} x_i(S) \leq 1$. If agents with the same type receive the same assignment, the assignment is called symmetric. Note that in a continuum economy, feasibility and implementability are equivalent.

The extension of the NPS mechanism to the continuum economy is called NPS^* and is defined as follows: for step $v = 0$, let $G^*(0) = G^*,\ t^*(0) = z^v = 0$. Let $m^i(v)$ denote agent i's most preferred bundle from objects in $G^*(v)$. Let $m_j(v) = \int_{i \in N^*} |m^i(v)|_j di$ be the mass of object j in agents' most prefered bundle. Given $G^*(v-1), t^*(v-1), z^{v-1}$ for all $j \in G$:

1. $t_j^*(v) = sup\{t \in [0,1] | z_j^{v-1} + m_j(v-1)(t - t(v-1)) \leq s_j^*\}$.

2. The first instance at which any good is exhausted is $t^*(v) = \min_{j \in G^*(v-1)} t_j^*(v)$.

3. At time $t^*(v)$, the set of goods with non-zero supply is

$$G^*(v) = G^*(v-1) \setminus \{j \in G^*(v-1) | t_j^*(v) = t^*(v)\}.$$

4. $$z_j^v = z_j^{v-1} + m_j(v-1)(t^*(v) - t^*(v-1))$$

5. The allocation returned by the NPS* mechanism at time $t(v)$ is

$$NPS^*(v) = NPS^*(v-1) + (t^*(v) - t^*(v-1))z^{v-1}.$$

6. NPS terminates at time $t^*(v)$ where v is the smallest index such that $t^*(v) = 1$ and outputs $NPS^*(v)$ as the final allocation.

NPS in the $q-$economy, denoted by NPS^q, is defined similarly. For any object $j \in G$ and any subset of goods $S \subseteq G$, let $m^i(S)$ be agent i's most preferred bundle, consisting of objects in S. Let $m^i_j(S)$ be the number of copies of object j in $m^i(S)$, that is, $m^i_j(S) = (m^i(S))_j$. Set $m^q_j(S) = \sum_{i \in N_q} m^i_j(S)$which is the total number of copies of object j in all agents' preferred bundles.

1. $t^q_j(v) = sup\{t \in [0,1] | z^{v-1}_j(q) + \frac{m^q_j(G(v-1))}{q}(t - t^q(v-1)) \leq \frac{s_j}{q}\}$.
2. $t^q(v) = \min_{j \in G(v-1)} t^q_j(v)$.
3. $G(v) = G(v-1) \setminus \{j \in G(v-1) | t^q_j(v) = t^q(v)\}$
4. $z^v_j(q) = z^{v-1}_j(q) + \frac{m^q_j(G(v-1))}{q}(t^q(v) - t^q(v-1))$.
5. $NPS^q(v) = NPS^q(v-1) + (t^q(v) - t^q(v-1))z^{v-1}(q)$
6. The terminal step occurs at the smallest v that satisfies $t^q(v) = 1$.

C.1 Convergence of NPS^q to NPS^*

Let t^q_j and t^*_j be the exhaustion date of object j in NPS^q and NPS^*, respectively. To prove that NPS^q converges to NPS^*, it is enough to show that the stock of object j at time t^*_j in NPS^q is of order of $o(q)$ and before that it is of the order of $O(q)$, this implies that the exhaustion dates of objects in the $q-$economy converges to their corresponding dates in NPS^*.

As the supply of multiple objects can become exhausted at the same time, the set $\{t^*(1), t^*(2), t^*(3), ..., t^*(v^*)\}$, will contain some duplicates. Let the distinct values in the set $\{t^*(1), t^*(2), t^*(3), ..., t^*(v^*)\}$ be $\{t_1, t_2, ..., t_g\}$ and ordered so that $t_1 < t_2 < t_3 < ... < t_g$. Set $t^*(v_i) = t_i$. Note g may not be equal to v^*. Let $A_r \subseteq G$ be the set of objects whose supply is exhausted at time t_r in the NPS^* mechanism. Let $s^q_j(t)$ and $s^*_j(t)$ be the supply of object j at time t in NPS^q and NPS^*, respectively. Also, assume $NPS^*(t)$ and

$NPS^q(t)$ are the allocation of the NPS algorithm at time t in the continuum economy and the $q-$economy respectively. We prove the following by induction:

1. For all $r \leq g$ and all $j \in G$, $\lim_{q\to\infty} \frac{s_j^q(t_r)}{q} = s_j^*(t_r)$.

2. For all $r \leq g$, $\lim_{q\to\infty} NPS^q(t_r) = NPS^*(t_r)$.

For $r = 1$, note that at the beginng since all objects are available, all agents are allocated their most preferred bundle. Therefore, $t_1 = \min_{j\in G} \frac{s_j^*(0)}{m_j(G)}$. Note that for all $j \in G$ the supply of object j at time t_1 in NPS^q is $s_j^q(0) - t_1 m_j^q(G)$. Hence,

$$\lim_{q\to\infty} \frac{s_j^q(0) - t_1 m_j^q(G)}{q} = s_j^*(0) - t_1 m_j(G) = s_j^*(t_1).$$

Assume statements 1 and 2 are true for $r-1$, we prove them for r. For all $\epsilon > 0$, we show there exists large enough $Q > 0$ such that for all $q > Q$, $|\frac{s_j^q(t_r)}{q} - s_j^*(t_r)| < \epsilon$.

Given ϵ_1, let Q_1 be such that

$$|\frac{s_j^q(t_{r-1})}{q} - s_j^*(t_{r-1})| < \epsilon_1 \tag{18}$$

for all $q > Q_1$ and $j \in G$. If $j \in A_\tau$ for some $\tau \leq r-1$, then the availability of object j at time t_{r-1} is at most $\epsilon_1 q$. Therefore, objects in A_τ will be allocated for at most $\frac{|G|\epsilon_1 q}{\vartheta^q}$ period of time, where $\vartheta^q = \min_{A\subset G} \min_{j\in A}\{m_j^q(A) | m_j^*(A) \neq 0\}$. Note that if $m_j^*(A) \neq 0$, then $m_j^q(A) = O(q)$. Therefore, $\vartheta^q = O(q)$. Let $\vartheta^* = \lim_{q\to\infty} \frac{q|G|}{\vartheta^q}$. Let Q_2 be such that $\frac{q|G|\epsilon_1}{\vartheta^q} \leq 2\vartheta^*\epsilon_1 = \epsilon_2$ for all $q > Q_2$. Choose ϵ_1 such that $\epsilon_1 + \epsilon_2 = \epsilon_1(2\vartheta^* + 1) < t_r - t_{r-1}$. For all $q > \max\{Q_1, Q_2\}$, the NPS mechanism in the $q-$economy at $t_{r-1} + \epsilon_1 + \epsilon_2$ would allocate the same bundles as NPS^*. This allocation would end when the supply of one of the objects is exhausted, call this date τ_r^q. This date is within the $\epsilon_1 + \epsilon_2$ neighbourhood of $\min_{a\in G^*(v_{r-1})} \frac{s_a^q(t_{r-1})}{m_a^q(G^*(v_{r-1}))} + t_{r-1}$. Note that $\lim_{q\to\infty} \frac{s_j^q(t_{r-1})}{m_a^q(G^*(v_{r-1}))} = \frac{s_a^*(t_{r-1})}{m_j^*(G^*(v_{r-1}))}$. For $\epsilon_3 > 0$

let Q_3 be such that for all $q > Q_3$, $|\frac{s_j^q(t_{r-1})}{m_j^q(G^*(v_{r-1}))} - \frac{s_j^*(t_{r-1})}{m_j^*(G^*(v_{r-1}))}| < \epsilon_3$. Therefore if $q > max\{Q_1, Q_2, Q_3\}$, then $|\tau_r^q - t_r| < \epsilon_1 + \epsilon_2 + \epsilon_3$. Note that if $q > max\{Q_1, Q_2, Q_3\}$, then NPS^q and NPS^* allocate the same bundles in the interval of $[t_{r-1}, t_r]$ except for a period with length at most $\epsilon_1 + \epsilon_2 + \epsilon_1 + \epsilon_2 + \epsilon_3$. Hence,

$$|\frac{s_j^q(t_r)}{q} - s_j^*(t_r)| \leq \frac{k(\epsilon_1 + \epsilon_2 + \epsilon_1 + \epsilon_2 + \epsilon_3)|N^q|}{q} \tag{19}$$

Since $\lim_{q\to\infty} \frac{|N^q|}{q}$ exists, choosing ϵ_is small enough for this case proves the inductive steps.

C.2 Convergence of RSD^q to NPS^*

Let $f_i \in [0, 1]$ be the draw of agent i. Consider all agents whose draws are in the interval $(t', t]$. Let $\hat{m}_j^q(A, t', t)$ be the total number of copies of object j in all agents' preferred bundles who's random draw is in $(t', t]$. Following Che and Kojima [2010], the RSD^q mechanism can be represented by the following step.

- $\hat{t}_j^q(v) = sup\{t \in [0, 1] | \hat{z}_j^q(v-1) + \frac{\hat{m}_j^q(\hat{G}^q(v-1), \hat{t}^q(v-1), t)}{q}(t - \hat{t}^q(v-1)) \leq \frac{s_j^q}{q}\}$.
- $\hat{t}^q(v) = \min_{j\in\hat{G}^q(v-1)} \hat{t}_j^q(v)$.
- $\hat{G}^q(v) = \hat{G}^q(v-1) \setminus \{j \in \hat{G}^q(v-1) | \hat{t}_j^q(v) = \hat{t}^q(v)\}$
- $\hat{z}_j^q(v) = \hat{z}_j^q(v-1) + \frac{\hat{m}_j^q(\hat{G}(v-1), \hat{t}^q(v-1), \hat{t}^q(v))}{q}(\hat{t}^q(v) - \hat{t}^q(v-1))$.
- With the terminal step the $\min_v$ that satisfies $\hat{t}^q(v) = 1$.

Since f_is are uniformly distributed, the weak law of large number implies that $\lim_{q\to\infty} \frac{\hat{m}_j^q(A, t', t)}{q} = m_j^*(A)(t' - t)$. Similar arguments as in the previous case establishes the asymptotic equivalence.