

Lottery Pricing Equilibria

SHADDIN DUGHMI, University of Southern California, shaddin@usc.edu

ALON EDEN, Tel Aviv University, alonarden@gmail.com

MICHAL FELDMAN, Tel Aviv University, michal.feldman@cs.tau.ac.il

AMOS FIAT, Tel Aviv University, fiat@tau.ac.il

STEFANO LEONARDI, Sapienza University of Rome, leonardi@dis.uniroma.it

We extend the notion of Combinatorial Walrasian Equilibrium, as defined by Feldman et al. [2013], to settings with budgets. When agents have budgets, the maximum social welfare as traditionally defined is not a suitable benchmark since it is overly optimistic. This motivated the liquid welfare of [Dobzinski and Paes Leme 2014] as an alternative. Observing that no combinatorial Walrasian equilibrium guarantees a non-zero fraction of the maximum liquid welfare in the absence of randomization, we instead work with randomized allocations and extend the notions of liquid welfare and Combinatorial Walrasian Equilibrium accordingly. Our generalization of the Combinatorial Walrasian Equilibrium prices lotteries over bundles of items rather than bundles, and we term it a *lottery pricing equilibrium*.

Our results are two-fold. First, we exhibit an efficient algorithm which turns a randomized allocation with liquid expected welfare W into a lottery pricing equilibrium with liquid expected welfare $\frac{3-\sqrt{5}}{2} \cdot W$ ($\approx 0.3819 \cdot W$). Next, given access to a demand oracle and an α -approximate oblivious rounding algorithm for the configuration linear program for the welfare maximization problem, we show how to efficiently compute a randomized allocation which is (a) supported on polynomially-many deterministic allocations and (b) obtains [nearly] an α fraction of the optimal liquid expected welfare. In the case of subadditive valuations, combining both results yields an efficient algorithm which computes a lottery pricing equilibrium obtaining a constant fraction of the optimal liquid expected welfare.

Additional Key Words and Phrases: Walrasian equilibrium, envy-free, lotteries, combinatorial auctions, budgets

1. INTRODUCTION

In this paper, we study generalizations of the *Walrasian Equilibrium* in combinatorial markets to settings with budgets. Given a set of items, a set of agents, and prices on the items, a bundle of items is said to be in the *demand set* of an agent if the utility of the agent — *i.e.*, the difference between the agent's value of the set of items and the sum of the item prices — is maximized for this bundle. A *Walrasian pricing* [Arrow and Debreu 1954; Walras 1874, 2003] assigns prices to items so that there is an allocation of items to buyers in which (a) each buyer is assigned a bundle in her demand set, and (b) the social welfare is maximized. In general, a Walrasian pricing need not exist, but is known to exist for so-called “gross substitutes” valuations [Kelso and Crawford

The work of S. Dughmi was supported in part by NSF CAREER Award CCF-1350900. The work of M. Feldman and A. Eden was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122. The work of S. Leonardi was partially supported by Google Focused award “Algorithms for Large-scale Data analysis” and EU FET project no. 317532 “Multiplex”. The work of A. Eden and A. Fiat was partially supported by ISF 1841/14. This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'16, July 24–28, 2016, Maastricht, The Netherlands. Copyright © 2016 ACM 978-1-4503-3936-0/16/07 ...\$15.00.

<http://dx.doi.org/10.1145/2940716.2940742>

1982]. A Walrasian pricing coupled with an allocation satisfying (a) and (b) is termed a Walrasian equilibrium.

[Feldman et al. 2013] circumvent the non-existence of Walrasian Equilibrium for arbitrary valuations by relaxing the notion in two ways: by pricing bundles rather than items, and by no longer requiring the market to clear (i.e. not all items need to be sold). They term the relaxed notion a *Combinatorial Walrasian Equilibrium (CWE)*. Feldman et al. [2013] describe a process which starts from an arbitrary initial allocation A and produces a CWE with $1/2$ the social welfare of A .

Our work follows in the tradition of a number of recent works which incorporate budget constraints in economic design [Devanur et al. 2012; Dobzinski et al. 2008; Goel et al. 2014]. To the best of our knowledge, Walrasian pricing and its variants have not previously been considered in combinatorial markets with budgets. In fact, we observe that even in very simple markets with budgets, the traditional Walrasian equilibrium need not exist, and only trivial Combinatorial Walrasian Equilibria which assign no items exist. We observe that generalizing the CWE by setting prices on *lotteries over bundles* rather than deterministic bundles restores a rich set of equilibria even in the presence of budgets. We call these *Lottery Pricing Equilibria (LPE)*.

As a simple motivating example, consider one item and n buyers, with each buyer having a value n for the item and a budget of 1. If we price the item at $p \leq 1$, no allocation satisfies the demand of more than a single agent, rendering such a price “illegal”. On the other hand, if we price the item at $p > 1$, the maximum social welfare is zero. Suppose instead that we sell n lottery tickets, each of which wins the item with probability $1/n$, and each of which is priced at 1; Assuming agents are risk neutral, this guarantees a social welfare (and revenue) of n . Ergo, if we require that each agent is allocated a bundle (or lottery over bundles) in her demand set, the ratio between the social welfare attainable by pricing lotteries and that attainable by pricing (deterministic) bundles is unbounded.

In settings with budgets, the maximum social welfare, as traditionally defined, is not suitable as a benchmark. To see this, consider a single item with value 1 to agent A , and value 2 to agent B , and let both agents have a budget of zero. The social welfare is maximized at 2 by assigning the item to agent B . However, any item pricing either can not satisfy the demand of both agents (when the price is 0) or results in a social welfare of 0 (when the price exceeds 0). A similar argument reveals that the same holds for lottery pricing: every lottery pricing equilibrium must assign non-zero prices to each lottery on offer, and consequently allocates the empty bundle to each agent.

Motivated by this phenomenon, [Dobzinski and Paes Leme 2014] introduced the *liquid welfare* of a deterministic allocation, defined as the sum over all agents of the minimum between the agent’s budget and her value for her allocated bundle. In the above example with two agents each with a budget of 0, the maximum liquid welfare of an allocation is 0, making it more suitable as a benchmark. Consider also the single-item example with n buyers, each of whom has value n for the item and a budget of 1. The maximum liquid welfare of a deterministic allocation is 1, and is attained by allocating the item to an [arbitrary] buyer. The minimum between the budget and the value is 1. This is also an upper bound on the maximum revenue a seller can obtain by assigning deterministic prices. Recalling that our solution prices lotteries rather than deterministic bundles, we require a generalization of the liquid welfare to randomized allocations.

1.1. Our contribution

In this work we introduce lottery pricing equilibria, and show that approximate lottery pricing equilibria achieve near optimal liquid welfare, for a natural generalization of the liquid welfare to randomized allocations. We show how to efficiently compute such

approximate equilibria from an initial allocation of lotteries, while losing only a constant fraction of the liquid welfare. We also show how to compute an initial allocation of lotteries with near-optimal liquid welfare under mild conditions which so happen to be satisfied by subadditive valuations. We now outline our main modeling and technical contributions.

Liquid welfare and lotteries. Given a randomized allocation of items to agents, we define its liquid expected welfare as the sum, over all agents, of the minimum between the agent’s budget and her expected value for her bundle. We use the maximum liquid expected welfare as our benchmark, and note that it also serves as an upper bound on the maximum revenue which can be obtained by a seller from risk neutral agents. As illustrated by our previous single-item n -agent example with valuations of n and budgets of 1, the gap in liquid expected welfare between the best randomized and the best deterministic allocations can be as large as n , and it is easy to see that this is tight.

Our pricing equilibria will assign prices to a collection of *lotteries*, where each lottery is a distribution over bundles of items. A collection of m lotteries is mutually compatible — or *feasible* — if there exists an *implementation* of these lotteries as the parts of a random partition of the items into m sets. We can thus interpret a randomized allocation of items to agents as a feasible collection of lotteries, one per agent. More generally, we consider many-to-one allocations of a collection of lotteries to agents, as needed to reason about our lottery pricing equilibria. When an agent is allocated multiple lotteries in a collection, the implementation of the collection is of paramount consequence in determining the resulting randomized allocation of items. Therefore the implementation must be fixed before we can calculate an agent’s expected value or the liquid expected welfare.¹

Lottery Pricing Equilibria. We define the Lottery Pricing Equilibrium (LPE), generalizing the Combinatorial Walrasian Equilibrium to lotteries. For $\epsilon \geq 0$, a ϵ -LPE consists of jointly distributed (feasible) collection of lotteries (and an associated implementation), a price for each lottery, and an allocation of those lotteries to the agents so that each agent receives a bundle of lotteries (with total price at most his budget) approximately maximizing his expected utility to within ϵ . For each $\epsilon > 0$, we show the existence of a ϵ -LPE for budgeted agents with arbitrary valuations that achieves a constant fraction of the liquid expected welfare of the best randomized allocation.

Our proof of existence is constructive, by way of an efficient algorithm analogous to that of [Feldman et al. 2013]. Specifically, our algorithm takes as input an arbitrary randomized allocation consisting of one lottery per agent, and terminates with a ϵ -LPE obtaining at least a $\frac{3-\sqrt{5}}{2}$ fraction of the liquid expected welfare of the initial randomized allocation.

The algorithm starts by setting an initial price for each lottery equal to a constant fraction of its contribution to the liquid expected welfare, and then proceed in rounds. Each round involves either “merging” a set of lotteries in the collection (and pricing the merged lottery at the sum of the two prices), or “resolving contention” between two agents demanding the same lottery. Merging a set of lotteries is the natural generalization of taking the union of a set of bundles: specifically, the random bundle associated

¹To see this, consider two lotteries L_1 and L_2 in a collection, assume that L_1 uniformly randomizes between the bundles $\{a\}$ and \emptyset , and L_2 uniformly randomizes between $\{b\}$ and \emptyset . Consider two implementations: one in which drawing the empty set from both lotteries has probability 0.5 (i.e. the lotteries are “positively correlated”), and another in which at least one of the lotteries draws a non-empty set (i.e., the lotteries are “negatively correlated”). If both lotteries are assigned to single agent who values items a and b as substitutes, then such an agent prefers the second implementation.

with the merged lottery is the union of the set of random bundles associated with the original lotteries. The implementation is adapted appropriately.

The main difference between our algorithm and that of [Feldman et al. 2013] — necessitated by the presence of budgets — is in the contention resolution procedure: instead of increasing the price of the contentious lottery, we carefully scale down the probability of each non-empty bundle in its support (i.e., we mix the lottery with the empty bundle).

Computing the Initial Allocation. The above-described procedure reduces finding a lottery pricing equilibrium with near-optimal liquid welfare to the problem of computing a randomized allocation with near-optimal liquid welfare. Consider the latter problem. By adapting the configuration linear program for combinatorial welfare maximization to the liquid expected welfare objective, we show how to sample a randomized allocation with liquid expected welfare within an α factor of optimal whenever two conditions are satisfied: (a) We are given demand oracle access to agents’ valuations; (b) There is a rounding algorithm which, when applied to any fractional solution of the configuration LP, outputs a (possibly random) allocation preserving each player’s valuation to within a factor of α . All rounding algorithms for the configuration LP we are aware of satisfy the latter player-by-player welfare guarantee; in particular, for submodular (and, more generally, XOS) valuations we get $\alpha = \frac{e}{e-1}$ using the rounding technique of [Dobzinski and Schapira 2006], and for subadditive valuations we get $\alpha = 2$ using the rounding technique of [Feige 2009].

Finally, we show how to produce an explicitly-described randomized allocation — supported on polynomially-many deterministic allocations — with nearly the same player-by-player guarantee. This is accomplished by Monte-Carlo sampling a near-optimal randomized allocation mixed with the allocation assigning all items to a random agent; the mixing is required in order to handle (exponentially) low-probability yet high-value events in the near-optimal randomized allocation.

Organization of the paper. Section 3 defines the basics of lotteries, collections of lotteries and their implementations, and our generalization of the liquid welfare to randomized allocations. Section 4 describes our procedure for turning a randomized allocation to a lottery pricing equilibrium at a loss of a constant factor in the liquid welfare. Finally, Section 5 examines the problem of computing a randomized allocation with high liquid expected welfare.

2. RELATED WORK

Starting with the work of Walras [1874, 2003], the study of market equilibria has been a cornerstone of economic theory. The classical equilibrium concept there is the Walrasian equilibrium (a.k.a. competitive equilibrium). When goods are divisible, Walrasian equilibria exist in quite general settings, with many buyers and sellers, as well as budgets (see [Mas-Colell et al. 1995]). In combinatorial markets involving indivisible goods, on the other hand, the situation is more nuanced. In this paper, we focus on combinatorial markets consisting of a single seller and multiple buyers, in which case the Walrasian equilibrium reduces to a pricing of the goods coupled with a market-clearing allocation satisfying each buyer’s demand. In the absence of budgets, it was shown by [Kelso and Crawford 1982] that such equilibria exist when buyers have Gross Substitutes valuations, and subsequently [Gul and Stacchetti 1999a] showed that this is essentially necessary. In each Walrasian equilibrium, the associated allocation of goods is welfare maximizing.

The work most closely related to ours is by [Feldman et al. 2013], who are motivated by the non-existence of the Walrasian Equilibrium for more general classes of valua-

tions in combinatorial markets. Their notion of Combinatorial Walrasian Equilibrium (CWE) prices bundles rather than items, and relaxes the market clearing property. They show the existence of CWE obtaining half the optimal social welfare. Since [Feldman et al. 2013] prices only the bundles assigned to buyers in the associated allocations (with all other bundles implicitly priced at infinity), their work is related in spirit to much of the earlier literature on envy-free pricing. The literature on envy-free pricing is vast (e.g., [Aumann 1975; Bikhchandani and Mamer 1997; Cohen et al. 2012; Fiat and Wingarten 2009; Foley 1967; Gul and Stacchetti 1999b; Leonard 1983; Mu’Alem 2009; Shapley and Shubik 1971; Varian 1974]). More recent work in the computer science community has focused on approximately optimal envy-free pricings as well as computational considerations (see e.g. [Balcan et al. 2008; Briest 2006; Cheung and Swamy 2008; Feldman et al. 2012; Guruswami et al. 2005; Hartline and Yan 2011]).

Since near-optimal CWE exist for general valuations, the work of [Feldman et al. 2013] partially bridges the gap between settings with divisible goods — in which Walrasian equilibria exist generally — and settings with indivisible goods. Our work continues in this spirit by seeking to remedy the second main deficiency of Walrasian equilibria in combinatorial markets: budgets. As previously described, doing so requires a new metric other than the social welfare which incorporates budgets, and our starting point for that is the liquid welfare of [Dobzinski and Paes Leme 2014].

Finally, we remark that pricing lotteries has been studied at length in mechanism design (see e.g. [Briest et al. 2010; Chen et al. 2015; Hart and Nisan 2013; Thanassoulis 2004]), where it has long been known that selling lotteries is required to maximize revenue (with or without budgets).

3. PRELIMINARIES

3.1. Lotteries and Collections

Given a set M of items, a *lottery* L is simply a distribution supported on 2^M . For a bundle of items $S \subseteq M$, we use $L(S)$ to denote the probability of S in lottery L .

We also consider *collections* of lotteries. Given a collection $\mathcal{L} = (L_1, \dots, L_n)$ of n lotteries, we say this collection is *feasible* if it is consistent with a random partial allocation of the items to n bins. Formally, the collection \mathcal{L} is feasible if there is a random collection of disjoint bundles $S_1, \dots, S_n \subseteq M$ so that each S_i is distributed as the lottery L_i . An algorithm I which samples such a random allocation $\mathcal{S} = (S_1, \dots, S_n)$ is said to be an *implementation* of \mathcal{L} . We refer to the range of outputs of I — a set of deterministic allocations — as the *support* of the implementation. We use the random variable $I(\mathcal{L})$ to denote the (random) allocation output by the implementation. More generally, given a sub-collection $\mathcal{L}' \subseteq \mathcal{L}$ of lotteries, we use $I(\mathcal{L}')$ to denote the projection of the allocation $I(\mathcal{L})$ onto the the lotteries in \mathcal{L}' . For example, if $\mathcal{L} = (L_1, L_2, L_3)$ with $\mathcal{S} = (S_1, S_2, S_3) \sim I(\mathcal{L})$, and if \mathcal{L}' consists of L_1 and L_3 , then $I(\mathcal{L}')$ is the implementation of \mathcal{L}' which samples the random variable $\mathcal{S}' = (S_1, S_3)$.

Consider an agent with a valuation function $v : 2^M \rightarrow \mathbb{R}$ defined on bundles. We use the natural extension of v to lotteries: $v(L) = E_{S \sim L} v(S)$. Similarly, given a collection of lotteries $\mathcal{L} = \{L_1, \dots, L_n\}$ and an implementation I of \mathcal{L} , we extend v to sub-collections \mathcal{L}' of \mathcal{L} as follows:

$$v(\mathcal{L}', I) = E_{\mathcal{S} \sim I(\mathcal{L}')} v(\cup_{S \in \mathcal{S}} S).$$

In other words, the agent’s value for a sub-collection of lotteries equals her expected value for the union of the associated bundles.

3.2. Operations on Lotteries

We now describe two operations on lotteries and their collections, which will be useful when describing our algorithms.

Merging Lotteries. Given a collection $\mathcal{L} = (L_1, \dots, L_n)$ of lotteries with implementation I , the merger of \mathcal{L} with respect to I , denoted by $\text{Merge}(\mathcal{L}, I)$, is the lottery which samples the random variable $\cup_i S_i$ for $(S_1, \dots, S_n) \sim I(\mathcal{L})$. In other words, merging a jointly distributed family of lotteries simply merges their associated (random) bundles.

Similarly, we often merge a subset of the lotteries in a collection to form a new collection. Given a feasible collection $\mathcal{L} = (L_1, \dots, L_n)$ with implementation I and a sub-collection $\mathcal{L}' = (L_{i_1}, \dots, L_{i_k})$, merging \mathcal{L}' yields the new collection $\mathcal{L} \cup \{\text{Merge}(\mathcal{L}')\} \setminus \mathcal{L}'$. Naturally, this new collection of lotteries is feasible, and admits the implementation $\text{newimp}(I, \mathcal{L}, \mathcal{L}')$ which simply samples the collection $(S_1, \dots, S_n) \sim I(\mathcal{L})$, removes from it S_{i_1}, \dots, S_{i_k} , and then adds their union.

We note that our extension of an agent's valuation v to collections of lotteries can be alternately described as follows. Given a collection \mathcal{L} with implementation I , and a sub-collection \mathcal{L}' of \mathcal{L} , the agent's value for \mathcal{L}' can be written as $v(\mathcal{L}', I) = v(\text{Merge}(\mathcal{L}', I(\mathcal{L}))) = E_{S \sim \text{Merge}(\mathcal{L}', I(\mathcal{L}))} v(S)$.

Scaling Lotteries. Our final operation on lotteries is *scaling*. Given a lottery L and a scaling factor $q \in [0, 1]$, we define $\text{Decrease-Probability}(L, q)$ as the lottery \tilde{L} with $\tilde{L}(S) = qL(S)$ for each non-empty bundle S . In other words, scaling lottery L by q simply mixes L with the empty bundle in proportions q and $1 - q$.

This operation naturally extends to feasible collections of lotteries, where we may scale down a lottery in a collection by some factor while maintaining feasibility. Given a collection of lotteries $\mathcal{L} = (L_1, \dots, L_n)$ with implementation I , a particular lottery L_i , and a scaling factor $q \in [0, 1]$, we implement the new collection $(\mathcal{L}_{-i}, \text{Decrease-Probability}(L, q))$ canonically by first sampling $(S_1, \dots, S_n) \sim I(\mathcal{L})$, then — using an independent biased coin — replacing S_i with the empty set with probability $1 - q$.

3.3. Liquid Expected Welfare

Given an agent i with valuation $v_i : 2^M \rightarrow \mathbb{R}_+$, recall that the agent's expected value for a draw from lottery L is denoted by $v_i(L)$. When the agent has a budget of b_i , we also define her *liquid expected value* for a lottery L as the minimum of $v_i(L)$ and b_i . When the budget b_i is clear from context, we use $\bar{v}_i(L) = \min(v_i(L), b_i)$ to denote the agent's liquid expected value for a lottery L . We note that an agent's liquid expected value for a lottery serves as an upperbound on the price an individually-rational and risk-neutral agent is willing to pay for a draw from that lottery.

Given a randomized allocation of items to budgeted agents, we define the *liquid expected welfare* as the sum, over all agents, of the agent's liquid expected value for her assigned random bundle (equivalently, lottery). More formally, given n agents with valuations v_1, \dots, v_n and budgets b_1, \dots, b_n , if the randomized allocation implements a collection of lotteries $\mathcal{L} = (L_1, \dots, L_n)$ — i.e., the bundle assigned to agent i is distributed as a draw from L_i — then the liquid expected welfare of the allocation equals $\sum_{i=1}^n \bar{v}_i(L_i)$. We denote the maximum possible liquid expected welfare — taken over all randomized allocations — by OPT. Note that, assuming risk neutrality and individual rationality, the liquid expected welfare of a randomized allocation serves as an upperbound on the total revenue the seller may generate using that allocation, and OPT serves as an upperbound on the optimal revenue. As illustrated in the introduction, attaining a liquid expected welfare close to OPT often require randomization in the allocation.

4. LOTTERY PRICING EQUILIBRIA

In this section we are given a collection $\mathcal{L} = (L_1, \dots, L_n)$ of feasible lotteries, where LW denotes $\sum_i \bar{v}(L_i)$. We aim at computing a tuple $(\mathcal{L}', \mathbf{p})$ such that \mathcal{L}' are feasible lotteries, $(\mathcal{L}', \mathbf{p})$ form an ϵ -LPE, and the liquid expected welfare, $\sum_i \bar{v}_i(L'_i)$, is a constant fraction of the liquid expected welfare of the initial lotteries.

4.1. Definition of Lottery Pricing Equilibrium

Let us first present a formal definition of LPE.

Definition 4.1.

- (1) Given a set of lotteries \mathcal{L} , implementation I , and price function $p : \mathcal{L} \mapsto \mathbb{R}^{\geq 0}$, the utility of buyer i on a collection of lotteries $\mathcal{L}' \subseteq \mathcal{L}$, $\mathcal{L}' = L_{\ell_1}, L_{\ell_2}, \dots, L_{\ell_k}$ is

$$u_i(\mathcal{L}', p, I) = \begin{cases} -\infty & \text{if } \sum_{j=1}^k p(L_{\ell_j}) > b_i \\ v_i(\mathcal{L}', I) - \sum_{j=1}^k p(L_{\ell_j}) & \text{otherwise} \end{cases}.$$

- (2) Given a set of lotteries \mathcal{L} , an implementation I , and a price function $p : \mathcal{L} \mapsto \mathbb{R}^{\geq 0}$, the demand set of agent i , $D_i(\mathcal{L}, p, I)$ includes all $\mathcal{L}' \subseteq \mathcal{L}$ that maximize the utility for agent i , $u_i(\text{Merge}(\mathcal{L}', I), p, I)$. Note that the demand set may be of exponential size.
- (3) A Lottery Price Equilibrium (LPE) for bidders with Budgets consists of
- A set of n budget limited buyers, where buyer i has valuation function v_i and budget b_i .
 - A collection of n feasible lotteries, $\mathcal{L} = (L_1, L_2, \dots, L_n)$, some of which may be the empty lottery (assigns nothing deterministically).
 - An implementation, I of \mathcal{L} .
 - A price function $p : \mathcal{L} \mapsto \mathbb{R}^{\geq 0}$ defined on the set of all lotteries onto the reals. Buyer i maximizes her utility by choosing lottery L_i , over every other subset of lotteries.

We next give the definition of ϵ -LPE.

Definition 4.2.

- We say that f is ϵ -maximized at $x \in D$ over a domain D if $f(x) \geq f(y) - \epsilon$ for all $y \in D$.
- Given a collection of lotteries \mathcal{L} , an implementation I , and a price function $p : \mathcal{L} \mapsto \mathbb{R}^{\geq 0}$, the ϵ -demand of agent i , $\epsilon\text{-}D(\mathcal{L}, p, I)$, includes all sets of lotteries for which the utility of agent i is ϵ -maximized; i.e., if $\mathcal{L}' \subseteq \mathcal{L}$ is in $\epsilon\text{-}D_i(\mathcal{L}, p, I)$, then for all $\hat{\mathcal{L}} \subseteq \mathcal{L}$, $u_i(\mathcal{L}', p, I) \geq u_i(\hat{\mathcal{L}}, p, I) - \epsilon$.
- An ϵ -approximate Lottery Price Equilibrium for bidders with budgets differs from a Lottery Price Equilibrium for bidders with budgets in item 3 above. The difference is that the buyer need not maximize her utility at L_i but only ϵ -maximized by L_i , over all other combinations of lotteries on sale.

4.2. Computing Lottery Pricing Equilibria

We present in Figure 1 the process that outputs a ϵ -LPE.

The algorithm will start from an initial collection of lotteries $\mathcal{L} = (L_1, \dots, L_n)$, one for each agent, of high liquid expected welfare. The initial pricing of lottery L_i is set to $p(L_i) = \varphi \cdot \bar{v}_i(L_i)$ for some $\varphi \in (0, 1)$ to be determined later. However, for an agent a , L_a may not belong to $\epsilon\text{-}D_a(\mathcal{L}, p, I)$. If this is the case, we consider a set S of lotteries in the (true) demand of buyer a . If S is not a singleton, we merge the lotteries in S and we update the implementation accordingly. The pricing of the merged lottery is equal to the sum of the prices of the lotteries in S . If S is formed by one single lottery, say L' ,

before to allocate L' to agent a , the algorithm checks whether L' is in the true demand set of some other agent b . If this is the case, the algorithm $\text{resolve}(L', a, b)$ given in Figure 2 solves the conflict by decreasing the probability of lottery L' till it is no longer in the true demand set of a or b . We are still guaranteed that L' is in the ϵ -demand of the other agent which is given the lottery.

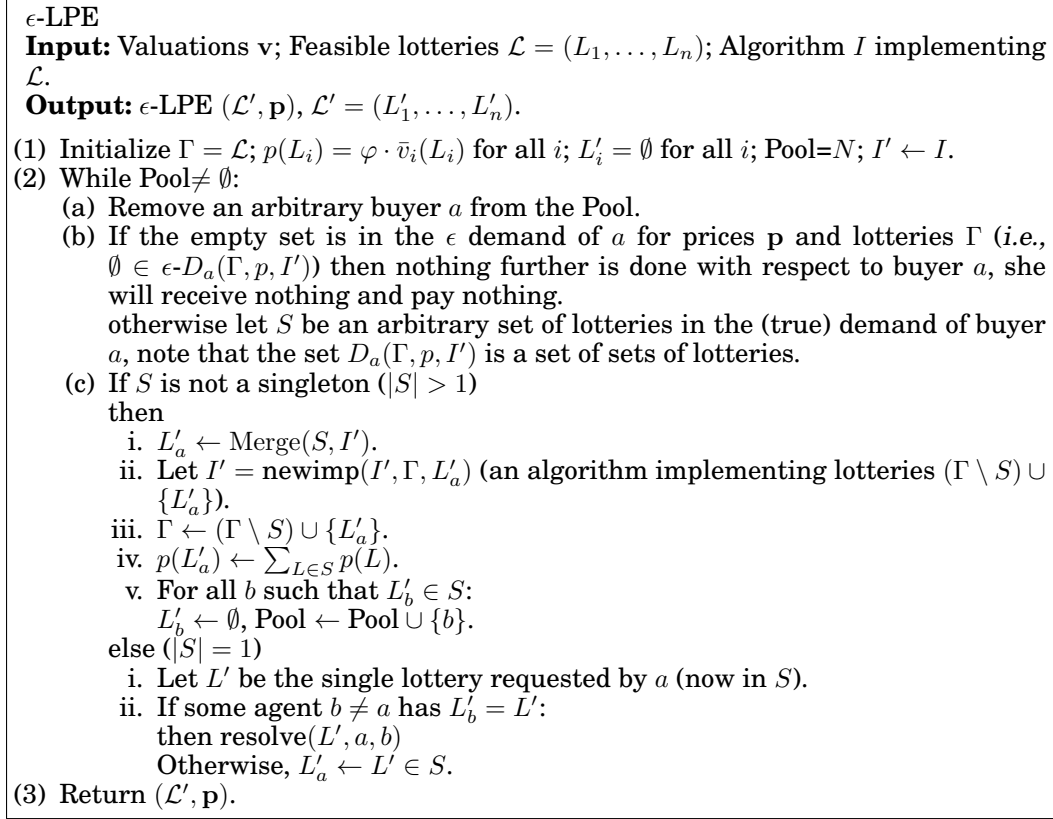


Fig. 1: The ϵ -LPE process.

The following invariants are used to prove convergence of the process to an ϵ -LPE.

OBSERVATION 4.3. *Let $i \in \{a, b\}$ be the agent that gets the lottery at the end of $\text{resolve}(L', a, b)$. We have that:*

- (1) $v_i(L') \geq p(L')$.
- (2) $L' \in \epsilon\text{-}D_i(\Gamma, p, I)$.

PROOF. At the start of the run of resolve , if one of the invariants does not hold for some agent i , then $L' \notin D_i(\Gamma, p, I)$, and agent i does not get the lottery, therefore, both invariants hold at first.

We show that after every iteration of the loop in step 2 of resolve , both invariants hold. Let $L' : 2^M \rightarrow \mathbb{R}^{\geq 0}$ be the lottery before calling $\text{Decrease-Probability}(L', q)$ in some iteration, and $\tilde{L} : 2^M \rightarrow \mathbb{R}^{\geq 0}$ and \tilde{I} be the the lottery and implementation afterwards. Now, $v_i(\tilde{L}) = \sum_{S \in 2^M \setminus \{\emptyset\}} \tilde{L}(S) \cdot v_i(S) = \sum_{S \in 2^M \setminus \{\emptyset\}} q \cdot L'(S) \cdot v_i(S) = q \cdot v_i(L')$. Since

<p>resolve(L', a, b)</p> <p>(1) $q \leftarrow 0$.</p> <p>(2) While $\{L'\} \in D_a(\Gamma, p, I') \cap D_b(\Gamma, p, I')$ and $q \neq 1$:</p> <p> (a) $v_{\min} \leftarrow \min(v_a(L'), v_b(L'))$; $v_{\max} \leftarrow \max(v_a(L'), v_b(L'))$.</p> <p> (b) $q \leftarrow \max\left(1 - \frac{\epsilon}{v_{\max}}, \frac{p(L')}{v_{\min}}\right)$.</p> <p> (c) Decrease-Probability(L', q), adjust the implementation accordingly.</p> <p>(3) If $\{L'\} \notin D_a(\Gamma, p, I')$ then</p> <p> — $L'_b \leftarrow L'$; $L'_a \leftarrow \emptyset$; Pool \leftarrow Pool \cup $\{a\}$.</p> <p> Else:</p> <p> — $L'_a \leftarrow L'$; $L'_b \leftarrow \emptyset$; Pool \leftarrow Pool \cup $\{b\}$.</p>
--

Fig. 2: The algorithm that resolves the conflict of attribution of a lottery.

$q \geq 1 - \frac{\epsilon}{v_{\max}}$, we have that $v_i(\tilde{L}) \geq \left(1 - \frac{\epsilon}{v_{\max}}\right) v_i(L') \geq v_i(L') - \epsilon$. Therefore, if $L' \in D_i(p, \Gamma, I')$, then $\tilde{L} \in \epsilon\text{-}D_i(\Gamma, p, \tilde{I})$.

Finally, it follows from $q \geq \frac{p(L')}{v_{\min}}$ that $v_i(\tilde{L}) \geq \frac{p(L')}{v_{\min}} \cdot v_i(L') \geq \frac{p(L')}{v_{\min}} \cdot v_{\min} = p(L') = p(\tilde{L})$, which implies that the utility of agent i from \tilde{L} is non-negative. \square

The following implies that for an agent who is not in the pool, her current allocation is in her ϵ demand.

OBSERVATION 4.4. *For any agent i , at any point of the ϵ -LPE process, if $i \notin \text{Pool}$, then L'_i is in the ϵ -demand of agent i .*

PROOF. This holds trivially at the start of the process. Whenever an agent takes a single lottery (or the empty set) in her demand, this holds. Whenever an agent takes a set of lotteries in her demand, and this set is merged at step 2(c)i of the process, this is true since the adapted implementation returns a random bundle which is the union of the random bundles received by the merged lotteries. Since the price of the new lottery is the sum of prices of the merged lotteries, this insures that the agent has the same utility for the new lottery as she had for the original set of lotteries. Finally, whenever two agent compete over a lottery in resolve, Observation 4.3 implies that for the agent who gets the lottery, this lottery is in her ϵ -demand. \square

Finally, the following holds simply since no lottery is dropped during the process.

OBSERVATION 4.5. *Let L be some lottery that is unassigned at the end of the ϵ -LPE process. It follows that this lottery was unassigned throughout the process.*

We now show that the process converges to an ϵ -LPE in weakly polynomial time.

LEMMA 4.6. *The process described above outputs an ϵ -LPE, and runs in time (weakly) polynomial in the input size and $1/\epsilon$.*

PROOF. We first claim that the process terminates. At every iteration, one of the following can occur:

- (1) A merger of two or more lotteries occurs.
- (2) The probability of a lottery is reduced until one of the buyers has a 0 utility.
- (3) The probability of a lottery is reduced until the reduced lottery is no longer in the demand set of one of the buyers.

(1) can happen at most n times, as when lotteries are merged, the overall number of lotteries decreases by 1. For every lottery, (2) can happen once for every agent until the merger of the lottery with another lottery, and therefore, (2) can happen at most $2n \cdot n = 2n^2$ times.

Let $\hat{v} = \max_i v_i(M)$. Let L' be a lottery whose probability has not yet been decreased, and let $\hat{v}_{L'} = \max_i v_i(L') \leq \hat{v}$. Every time the lottery's probability is decreased, the expected value of the lottery to the agent who values the lottery the most decreases by at least ϵ . This can happen until the expected value of the agent for the lottery is equal to its price p' . At that point, the utility of all agents for this lottery is at most 0, and it won't be demanded again. Therefore, Step 2c of resolve can happen at most $\frac{\hat{v}_{L'} - p'}{\epsilon} \leq \frac{\hat{v}}{\epsilon}$ times, which is weakly polynomial. Since there are no more than $2n$ lotteries throughout the process (every time a lottery is created, the overall number of lotteries decreases by 1), this implies that the number of times (3) can happen is weakly polynomial as well.

According to Observation 4.4, when the process terminates, for every agent, her given bundle is in her ϵ -demand. \square

Finally, we show that the computed ϵ -LPE has liquid expected welfare at least $\frac{3-\sqrt{5}}{2}$ of the liquid expected welfare of the initial allocation.

THEOREM 4.7. *Let $\mathcal{L}' = (L'_1, \dots, L'_n)$ be the final assignment of lotteries to agents, where agent i receives lottery L'_i , and let $\mathcal{L} = (L_1, \dots, L_n)$ be the initial set of lotteries. Setting $\varphi = \frac{\sqrt{5}-1}{2}$ in process ϵ -LPE, we get*

$$\sum_i \bar{v}_i(L'_i) \geq \frac{3-\sqrt{5}}{2} \sum_i \bar{v}_i(L_i) \approx 0.3819 \sum_i \bar{v}_i(L_i).$$

PROOF. Let $\text{Untaken} = \{i : L_i \text{ remains unassigned at the end of the process}\}$, and let $\text{Taken} = \{1, \dots, n\} \setminus \text{Untaken}$. We define the following set:

$$L(i) = \{j : L_j \text{ was merged in the process of creating } L'_i\}.$$

Notice that $p(L'_i) = \sum_{j \in L(i)} p(L_j) = \varphi \cdot \sum_{j \in L(i)} \bar{v}_j(L_j)$, and that for every $i \neq j$, $L(i) \cap L(j) = \emptyset$. Therefore,

$$\sum_i p(L'_i) = \varphi \cdot \sum_{j \in \text{Taken}} \bar{v}_j(L_j). \quad (1)$$

For every agent i , since her expected utility is non-negative at the end of the ϵ -LPE process, we have

$$\bar{v}_i(L'_i) \geq p(L'_i). \quad (2)$$

For every agent i such that $i \in \text{Untaken}$, we show that

$$\bar{v}_i(L'_i) - \varphi \cdot p(L'_i) \geq (1 - \varphi) \cdot \bar{v}_i(L_i). \quad (3)$$

Since L_i was left untouched during the run of ϵ -LPE, agent i 's utility from the lottery she received is at least $v_i(L_i) - p(L_i) - \epsilon$. Although this ϵ can be arbitrarily small, and does not change our statement by much, we observe that we can get rid of this ϵ by a slight change in the ϵ -LPE process, as follows. When two agents compete over a lottery in resolve, if the lottery drops out of both agents' demand at once, and one of the agents' initial lotteries is still unassigned, we allocated the lottery to the other agent. If both agents' initial lottery is unassigned, we make sure that the probability of the lottery is reduced in such a way that one of the agents does not strictly prefer her initial lottery, and assign the lottery to this agent. This change ensures that if an

agent's initial lottery is left untouched, her utility from the lottery she received is at least the utility from her initial lottery, *i.e.*,

$$v_i(L_i) - p(L_i) \geq \bar{v}_i(L_i) - \varphi \cdot \bar{v}_i(L_i) = (1 - \varphi) \cdot \bar{v}_i(L_i), \quad (4)$$

where the first inequality follows by $v_i(L_i) \geq \bar{v}_i(L_i)$ and $p(L_i) = \varphi \cdot \bar{v}_i(L_i)$.

For every agent $i \in \text{Untaken}$, either

— $\bar{v}_i(L'_i) = v_i(L'_i)$: In this case,

$$\bar{v}_i(L'_i) - \varphi \cdot p(L'_i) = v_i(L'_i) - \varphi \cdot p(L'_i) > v_i(L'_i) - p(L'_i) = u_i(L'_i) \geq (1 - \varphi) \cdot \bar{v}_i(L_i),$$

where the last inequality follows by Eq. (4), or,

— $\bar{v}_i(L'_i) = b_i$: Since the agent has a non-negative utility from L'_i , $b_i \geq p(L'_i)$, and therefore, $\bar{v}_i(L'_i) - \varphi \cdot p(L'_i) \geq b_i - \varphi \cdot b_i \geq (1 - \varphi) \cdot \bar{v}_i(L_i)$, where the first inequality follows since the agent has a non-negative utility from her assigned lottery (and therefore $b_i \geq p(L'_i)$), and last inequality follows by the definition of liquid welfare.

We get that

$$\begin{aligned} \sum_i \bar{v}_i(L'_i) &= \sum_i (\bar{v}_i(L'_i) - \varphi \cdot p(L'_i)) + \varphi \sum_i p(L'_i) \\ &\geq \sum_{i \in \text{Untaken}} (\bar{v}_i(L'_i) - \varphi \cdot p(L'_i)) + \varphi \left(\varphi \sum_{i \in \text{Taken}} \bar{v}_i(L_i) \right) \\ &\geq (1 - \varphi) \sum_{i \in \text{Untaken}} \bar{v}_i(L_i) + \varphi^2 \sum_{i \in \text{Taken}} \bar{v}_i(L_i) \\ &\geq \min(1 - \varphi, \varphi^2) \cdot \sum_i \bar{v}_i(L_i), \end{aligned}$$

where the first inequality is due to (1) and (2), and the second inequality is due to (3). Solving for $1 - \varphi = \varphi^2$, we get $\varphi = \frac{\sqrt{5}-1}{2}$, and therefore

$$\sum_i \bar{v}_i(L'_i) \geq \frac{3 - \sqrt{5}}{2} \sum_i \bar{v}_i(L_i) \approx 0.3819 \sum_i \bar{v}_i(L_i),$$

as desired. \square

In [Feldman et al. 2013] a combinatorial algorithm giving an exact CWE is presented. Unfortunately, our algorithm and a straightforward adaptation of the combinatorial algorithm of [Feldman et al. 2013] fails to give an exact LPE. A concrete example appears in Appendix A.

5. COMPUTING THE INITIAL ALLOCATION

5.1. The Liquid Welfare LP

In this section we show how to compute the initial randomized allocations used for the LPE process. We first present an LP for finding an assignment that maximizes a relaxed version of our problem, where agents can receive fractional bundles.

Although this LP has an exponential $2^m \cdot n$ variables, the dual can be solved using a separation oracle. In the dual there are exponentially many constraints. Nevertheless, we can solve it, assuming we have access the demand oracle of each player. Given a solution $\{u_i, z_i\}_{i \in N}, \{p_j\}_{j \in M}$, we want to verify that $u_i + (z_i - 1) \cdot v_i(S) + \sum_{j \in S} p_j \geq 0$ holds for every agent i and set S . For agent i , consider 2 cases. If $z_i \geq 1$, then the LHS of the inequality is non-negative since all variables are non-negative. Therefore, all constraints hold for agent i . If $z_i < 1$, then consider the equivalent set of constraints

LW-LP:

$$\begin{aligned}
\max \quad & \sum_{i,S} y_{i,S} \cdot v_i(S) \\
\text{s.t.} \quad & \sum_S y_{i,S} \cdot v_i(S) \leq b_i \quad \text{for every } i \in N \\
& \sum_S y_{i,S} \leq 1 \quad \text{for every } i \in N \\
& \sum_{i,S \ni j} y_{i,S} \leq 1 \quad \text{for every } j \in M \\
& y_{i,S} \in [0, 1] \quad \text{for every } i \in N, S \subseteq M
\end{aligned}$$

Fig. 3: An LP that maximizes the liquid welfare **fractionally**.

Dual:

$$\begin{aligned}
\min \quad & \sum_{i=1}^n (u_i + b_i \cdot z_i) + \sum_{j=1}^m p_j \\
\text{s.t.} \quad & u_i + v_i(S) \cdot z_i + \sum_{j \in S} p_j \geq v_i(S) \quad \text{for every } i \in N, S \subseteq M \\
& u_i, z_i \geq 0 \quad \text{for every } i \in N \\
& p_j \geq 0 \quad \text{for every } j \in M
\end{aligned}$$

Fig. 4: The dual of LW-LP.

$v_i(S) - \sum_{j \in S} \frac{p_j}{1-z_i} \leq \frac{u_i}{1-z_i}$ for every set S ; This can be verified by checking if the constraint holds for the set that maximizes the agent's utility at prices $\{p_j/(1-z_i)\}_{j \in M}$, which can be done given agent i 's demand oracle.

5.2. Producing Lotteries from the LP solution

Since we want to produce lotteries for the agents, we must use a rounding algorithm R that takes a solution \mathbf{y} of LW-LP and returns a feasible allocation $\mathcal{S} = (S_1, \dots, S_n)$. Let $LP = \sum_{i,S} y_{i,S} \cdot v_i(S)$, and for every agent i , let $LP_i = \sum_S y_{i,S} \cdot v_i(S)$. Let R be a randomized polynomial time rounding algorithm such that:

- (1) For every $\mathcal{S} \sim R(\mathbf{x})$ and for every $i \neq j$, $S_i \cap S_j = \emptyset$.
- (2) For every i , $\mathbb{E}_{\mathcal{S} \sim R(\mathbf{x})}[v_i(S_i)] \geq \alpha \cdot LP_i$.

If R satisfies both properties, we say R is an α -**guarantee** rounding algorithm.

We note that the existence of such an algorithm already guarantees the existence of feasible lotteries $\mathcal{L} = (L_1, \dots, L_n)$ such that $\sum_i v_i(L_i) \geq \alpha \cdot OPT$. This can be done by setting $L_i(S) = \Pr_{\mathcal{S} \sim R(\mathbf{x})}[S_i = S]$. The algorithm that implements \mathcal{L} is simply $R(\mathbf{x})$.

We get the following:

$$\begin{aligned}
\sum_i \bar{v}_i(L_i) &= \sum_i \min(v_i(L_i), b_i) \geq \sum_i \min(\alpha \cdot LP_i, b_i) \\
&= \sum_i \alpha \cdot LP_i = \alpha \cdot LP,
\end{aligned} \tag{5}$$

where the inequality holds since for every i , $v_i(L_i) = \mathbb{E}_{S \sim R(x)}[v_i(S_i)] \geq \alpha \cdot LP_i$, and the following equality holds since $LP_i \leq b_i$ for every i .

We next claim that $LP \geq OPT$. To see this, let $\mathcal{L}^* = \{L_1^*, \dots, L_n^*\}$ be a feasible set of lotteries that yields optimal liquid expected welfare. One can transform \mathcal{L}^* into a solution to the LW-LP by setting $y_{i,S}$ to be the probability of S in lottery L_i^* . By the feasibility of \mathcal{L}^* , all the feasibility constraints in LW-LP are satisfied. In addition, one can assume, without loss of generality, that in an optimal set of lotteries, the expected welfare of every agent does not exceed her budget (otherwise, decrease probabilities of lottery outcomes without affecting liquid welfare or feasibility). Since $\sum_S y_{i,S} \cdot v_i(S)$ is exactly the expected welfare of agent i from L_i^* , this means that the LW-LP constraints of type $\sum_S y_{i,S} \cdot v_i(S) \leq b_i$ are also satisfied. We conclude that the transformation of \mathcal{L}^* as described above gives a feasible assignment to LW-LP with objective value of OPT . Therefore, $LP \geq OPT$. Combining this with Equation (5) yields $\sum_i \hat{v}_i(L_i) \geq \alpha \cdot OPT$.

Plugging these lotteries to the LPE process yields a LPE with a liquid expected welfare which is within a $\frac{3-\sqrt{5}}{2} \cdot \alpha$ factor from the optimal one. From [Feige 2009] we can use a method that computes a feasible allocation in polynomial time with $\alpha = 1/2$ for subadditive valuations and $\alpha = (e-1)/e$ for submodular valuations. However, L_i might be a lottery over an exponential number of sets (although R is a polynomial time algorithm). We devote the next section to show how to produce a set of feasible lotteries over a polynomial number of sets with approximately the same guarantees.

5.3. Reducing support size to be polynomial

In this section, we are given a set of feasible lotteries $\mathcal{L} = (L_1, \dots, L_n)$ with a possibly super-polynomial support size, and an implementation of these lotteries I . For every i , let $\hat{v}_i = v_i(L_i)$. Our goal is to output a set of lotteries with a polynomial support size, that are easily implemented, with approximately the same liquid expected welfare guarantee \hat{v}_i for each agent i . Since the lotteries might be of super-polynomial support size, they cannot necessarily be explicitly described. For this reason, we assume a sample access to the implementation of the lotteries \mathcal{O}^I ; *i.e.*, \mathcal{O}^I returns an allocation $S = (S_1, \dots, S_n)$, where $\Pr_{S \sim \mathcal{O}^I}[S_i = T] = L_i(T)^2$.

A straightforward approach to produce a set of lotteries with the same welfare is the following – Let S^1, \dots, S^ℓ be the output produced in ℓ samples from \mathcal{O}^I ($S^t = (S_1^t, \dots, S_n^t)$ is the output of the t th sample). A set $\hat{\mathcal{L}}$ of feasible lotteries can be produced the following way – \hat{L}_i is a lottery which with probability $\frac{1}{\ell}$ outputs S_i^k for $k \in \{1, \dots, \ell\}$. The implementation of $\hat{\mathcal{L}}$ is simply to choose S^k with probability $\frac{1}{\ell}$ for $k \in \{1, \dots, \ell\}$ and allocated S_i^k for the owner of \hat{L}_i . For a sufficiently large ℓ , this process guarantees that with high probability $v_i(\hat{L}_i)$ is close to \hat{v}_i . Unfortunately, this ℓ can be exponential in n .

To show this, consider the following example — a single item a is up for sale. Agents $i = \{1, \dots, n-1\}$ are with value 2^n for the item and a budget of 1. agent n is with value 1 and budget $1 - \frac{n-1}{2^n}$ for the item. The unique optimal solution for LW-LP in this case is to set $x_i = \frac{1}{2^n}$ for $i = \{1, \dots, n-1\}$ and $x_n = 1 - \frac{n-1}{2^n}$. Let \mathcal{L} be the set of lotteries representing this solution; *i.e.*, $L_i(a) = x_i$ for every i . An implementation I of \mathcal{L} can simply give the item a to L_i with probability x_i . However, for any $\ell = p(n)$ runs, with an astonishingly high probability, none of agents $1, \dots, n-1$ will receive an item. Therefore, using such process above to produce lotteries will not give us the welfare guarantee we want, unless the support of each lottery is exponential.

²In [Feige 2009], he shows how to implement such an oracle for subadditive valuations in polynomial time.

	\hat{L}_1	\hat{L}_2	\hat{L}_3	...	\hat{L}_i	...	\hat{L}_n
1	S_1^1	S_2^1	S_3^1	...	S_i^1	...	S_n^1
2	S_1^2	S_2^2	S_3^2	...	S_i^2	...	S_n^2
3	S_1^3	S_2^3	S_3^3	...	S_i^3	...	S_n^3
...
k	S_1^k	S_2^k	S_3^k	...	S_i^k	...	S_n^k
...
ℓ	S_1^ℓ	S_2^ℓ	S_3^ℓ	...	S_i^ℓ	...	S_n^ℓ

Fig. 5: An illustration of the process of creating the new set of lotteries. Every column i represents a lottery \hat{L}_i created by the process, where each entry S_i^k is a set in the support of \hat{L}_i . The entries in the k th row, $S = (S_1^k, S_2^k, \dots, S_i^k, \dots, S_n^k)$ represents disjoint sets of items returned by the k th sample from $\hat{\mathcal{O}}^I$. For every $k \in [\ell]$, the implementation of the lotteries chooses the k th row with probability $1/\ell$, and for each lottery \hat{L}_i , returns the corresponding S_i^k .

In the above example, some agents receive most of their welfare only in an event of a very low probability. To overcome this, we introduce the following change in the process. At each iteration, with probability $1 - \epsilon$, we sample \mathcal{O}^I as before to produce the sets of items; with probability ϵ , we pick a random agent i , and give this agent all the items. This modification reduces the expected welfare of each agent by a factor of at most $1 - \epsilon$, and for the agents who only receive a substantial welfare from low probability events, increases the probability of such events. The modification is captured by the sample oracle $\hat{\mathcal{O}}^I$ (Figure 7) in use by the process. The process is presented by Algorithm ProduceLotteries (Figure 6). See Figure 5 for a visualization of the resulting lotteries obtained by the process.

ProduceLotteries

Input: \mathcal{O}^I : A an oracle access to the implementation I of a set \mathcal{L} of lotteries.

Output: Feasible lotteries $\hat{\mathcal{L}} = (\hat{L}_1, \dots, \hat{L}_n)$.

- (1) For $t \leftarrow 1 \dots \ell$:
 - $S^t = (S_1^t, \dots, S_n^t) \sim \hat{\mathcal{O}}^I(\mathcal{O}^I)$.
- (2) For every i , let \hat{L}_i be the algorithm that picks $k \in [\ell]$ uniformly at random and returns S_i^k .
- (3) Return $\hat{\mathcal{L}} = (\hat{L}_1, \dots, \hat{L}_n)$.

Fig. 6: Producing initial lotteries for the LPE process.




Fig. 7: An $\alpha \cdot (1 - \epsilon)$ -guarantee oblivious rounding algorithm.

We finally show that using poly time iterations of this process, we produce a set $\widehat{\mathcal{L}}$ of feasible lotteries with the desired properties.

LEMMA 5.1. *Letting $\ell = \Theta\left(\frac{n \log(n/\epsilon)}{\epsilon^3}\right)$, with probability $1 - \epsilon$, for every i , $v_i(\widehat{L}_i) \geq (1 - \epsilon)\widehat{v}_i$.*

PROOF. In order to analyze the number of rounds needed, we use the following variants of the multiplicative Chernoff bound.

Let χ_1, \dots, χ_ℓ be ℓ independent random variables, such that $\chi_i \in [0, B]$ for some $B > 0$, and let $\mathbb{E}[\chi_i] = \mu$ for every $i \in \{1, \dots, \ell\}$. For every $\epsilon \in (0, 1]$, the following holds:

$$\Pr\left[\frac{1}{\ell} \sum_{i=1}^{\ell} \chi_i < (1 - \epsilon)\mu\right] < \exp\left(-\frac{\epsilon^2 \cdot \mu \cdot \ell}{2B}\right). \quad (6)$$

First, notice that $v_i(\widehat{L}_i) = \frac{\sum_{t=1}^{\ell} v_i(S_t^t)}{\ell}$. For every agent i , we define the following sets:

$$\text{Big}_i = \left\{S \subseteq M : v_i(S) \geq \frac{n}{\epsilon} \cdot \widehat{v}_i\right\} \text{ and } \text{Small}_i = \left\{S \subseteq M : v_i(S) < \frac{n}{\epsilon} \cdot \widehat{v}_i\right\}.$$

Let

$$v_i^{\text{Small}} = \sum_{S \in \text{Small}_i} \Pr_{S \sim \widehat{\mathcal{O}}^t} [S_i = S] \cdot v_i(S).$$

Let S^1, \dots, S^ℓ be the ℓ outputs of R during the run of ProduceLotteries. Let X_t be a random variable defined by

$$X_t = \begin{cases} v_i(S_t^t) & S_t^t \in \text{Small}_i \\ 0 & \text{otherwise} \end{cases},$$

let Y_t be a random variable defined by

$$Y_t = \begin{cases} 1 & S_t^t \in \text{Big}_i \\ 0 & \text{otherwise} \end{cases}.$$

Using the above definitions, we have that

$$v_i(\widehat{L}_i) = \frac{\sum_{t=1}^{\ell} (X_t + Y_t \cdot v_i(S_t^t))}{\ell}. \quad (7)$$

Note that $\mathbb{E}[X_t] = v_i^{\text{Small}}$ and that $X_t \in [0, \frac{n}{\epsilon} \cdot \widehat{v}_i]$. If $v_i^{\text{Small}} > (1 - \epsilon)\widehat{v}_i$, using Chernoff bound (6) and $\ell = \Theta\left(\frac{n \cdot \log(n/\epsilon)}{\epsilon^3}\right)$, we get that

$$\Pr\left[\frac{\sum_{t=1}^{\ell} X_t}{\ell} < (1 - \epsilon)v_i^{\text{Small}}\right] < \exp\left(-\frac{\epsilon^2 \cdot v_i^{\text{Small}} \cdot \ell}{\frac{2n}{\epsilon} \cdot \widehat{v}_i}\right) < \exp(-\log(n/\epsilon)) = \frac{\epsilon}{n},$$

where the second inequality comes from the definition of ℓ and from the fact that $v_i^{\text{Small}} > (1 - \epsilon)\alpha \cdot LP_i$ for some small ϵ .

If Big_i is non-empty, and by the monotonicity of v_i , $M \in \text{Big}_i$. Therefore, $\mathbb{E}[Y_t] \geq \Pr[S_i^t = M] \geq \frac{\epsilon}{n}$. Using Chernoff bound with $\ell = \Theta\left(\frac{n \log(n/\epsilon)}{\epsilon^3}\right)$, we get that

$$\Pr\left[\frac{\sum_{t=1}^{\ell} Y_t}{\ell} < (1 - \epsilon) \mathbb{E}[Y_t]\right] < \exp\left(-\frac{\epsilon^2 \cdot \mathbb{E}[Y_t] \cdot \ell}{2}\right) < \exp(-\log(n/\epsilon)) = \frac{\epsilon}{n},$$

where the second inequality comes from the definition of ℓ and from the fact that $\mathbb{E}[Y_j] \geq \frac{\epsilon}{n}$.

We inspect the following cases:

— $\text{Big}_i \neq \emptyset$: In this case, whenever $\frac{\sum_{t=1}^{\ell} Y_t}{\ell} \geq (1 - \epsilon) \mathbb{E}[Y_t] \geq (1 - \epsilon) \frac{\epsilon}{n}$, then

$$v_i(\hat{L}_i) \geq \frac{\sum_{t=1}^{\ell} Y_t \cdot v_i(S_i^t)}{\ell} \geq (1 - \epsilon) \cdot \frac{\epsilon}{n} \cdot \frac{n}{\epsilon} \cdot \hat{v}_i = (1 - \epsilon) \cdot \hat{v}_i$$

— $\text{Big}_i = \emptyset$: In this case, $v_i^{\text{Small}} = \mathbb{E}_{S \sim R}[v_i(S_i)] \geq (1 - \epsilon) \cdot \hat{v}_i$. Whenever $\frac{\sum_{t=1}^{\ell} X_t}{\ell} \geq (1 - \epsilon) \cdot v_i^{\text{Small}}$, then $v_i(\hat{L}_i) \geq \frac{\sum_{t=1}^{\ell} X_t}{\ell} \geq (1 - \epsilon)^2 \cdot \hat{v}_i > (1 - 2\epsilon) \cdot \hat{v}_i$.

In both cases, $\Pr\left[v_i(\hat{L}_i) < (1 - 2\epsilon) \cdot \hat{L}_i\right] < \frac{\epsilon}{n}$. Using a union bound, we get the desired result. \square

6. CONCLUSIONS AND OPEN PROBLEMS

In this work we introduce *Lottery Pricing Equilibria* (LPE), extending the notion of Combinatorial Walrasian Equilibrium [Feldman et al. 2013] to settings with budgets by pricing lotteries over bundles. We use a generalization of the liquid welfare of [Dobzinski and Paes Leme 2014] to lotteries as our objective, in lieu of the social welfare. We present a weakly polynomial algorithm that turns an initial randomized allocation into a ϵ -LPE, losing only a constant fraction of the initial liquid welfare in the process. We also show how to efficiently compute an initial randomized allocation with high liquid welfare for subadditive valuations.

Our work raises several interesting questions. Can our loss of a factor of $\frac{3-\sqrt{5}}{2}$ be improved? Does there exist, and can we compute (in weakly or strongly polynomial time), an exact LPE with high liquid welfare? What about bounding the revenue which can be obtained by an LPE?

REFERENCES

- K. J. Arrow and G. Debreu. 1954. The Existence of an Equilibrium for a Competitive Economy. *Econometrica* 22, 3 (1954), 265–290.
- Robert J Aumann. 1975. Values of Markets with a Continuum of Traders. *Econometrica* 43, 4 (1975), 611–46.
- Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. 2008. Item pricing for revenue maximization. In *EC*. 50–59.
- Sushil Bikhchandani and John W. Mamer. 1997. Competitive Equilibrium in an Exchange Economy with Indivisibilities. *J. of Economic Theory* 74, 2 (1997), 385–413.
- Patrick Briest. 2006. Towards Hardness of Envy-Free Pricing. *ECCC* 13, 150 (2006).
- Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. 2010. Pricing Randomized Allocations. In *ACM-SIAM Symposium on Discrete Algorithms*. 585–597.

- Xi Chen, I. Diakonikolas, A. Orfanou, D. Pappas, Xiaorui Sun, and M. Yannakakis. 2015. On the Complexity of Optimal Lottery Pricing and Randomized Mechanisms. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*. 1464–1479. DOI: <http://dx.doi.org/10.1109/FOCS.2015.93>
- Maurice Cheung and Chaitanya Swamy. 2008. Approximation Algorithms for Single-minded Envy-free Profit-maximization Problems with Limited Supply. In *FOCS*. 35–44.
- Edith Cohen, Michal Feldman, Amos Fiat, Haim Kaplan, and Svetlana Olonetsky. 2012. Envy-Free Makespan Approximation. *SIAM J. Comput.* 41, 1 (2012), 12–25. DOI: <http://dx.doi.org/10.1137/100801597>
- Nikhil R. Devanur, Bach Q. Ha, and Jason D. Hartline. 2012. Prior-free Auctions for Budgeted Agents. *CoRR* abs/1212.5766 (2012).
- Shahar Dobzinski, Ron Lavi, and Noam Nisan. 2008. Multi-unit Auctions with Budget Limits. In *FOCS*. 260–269.
- Shahar Dobzinski and Renato Paes Leme. 2014. Efficiency Guarantees in Auctions with Budgets. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*. 392–404.
- Shahar Dobzinski and Michael Schapira. 2006. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. Society for Industrial and Applied Mathematics, 1064–1073.
- Uriel Feige. 2009. On Maximizing Welfare When Utility Functions Are Subadditive. *SIAM J. Comput.* 39, 1 (2009), 122–142. DOI: <http://dx.doi.org/10.1137/070680977>
- Michal Feldman, Amos Fiat, Stefano Leonardi, and Piotr Sankowski. 2012. Revenue maximizing envy-free multi-unit auctions with budgets. In *ACM Conference on Electronic Commerce, EC '12, Valencia, Spain, June 4-8, 2012*. 532–549. DOI: <http://dx.doi.org/10.1145/2229012.2229052>
- Michal Feldman, Nick Gravin, and Brendan Lucier. 2013. Combinatorial Walrasian Equilibrium. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing (STOC '13)*. ACM, New York, NY, USA, 61–70. DOI: <http://dx.doi.org/10.1145/2488608.2488617>
- Amos Fiat and Amiram Wingarten. 2009. Envy, Multi Envy, and Revenue Maximization. In *WINE*. 498–504.
- D. Foley. 1967. Resource Allocation and the Public Sector. *Yale Economics Essays* 7 (1967), 45–98.
- Gagan Goel, Vahab S. Mirrokni, and Renato Paes Leme. 2014. Clinching auctions beyond hard budget constraints. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*. 167–184.
- F. Gul and E. Stacchetti. 1999a. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory* 56 (1999), 95–124.
- Faruk Gul and Ennio Stacchetti. 1999b. Walrasian Equilibrium with Gross Substitutes. *J. of Economic Theory* 87, 1 (1999), 95–124.
- Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. 2005. On profit-maximizing envy-free pricing. In *SODA*. 1164–1173.
- Sergiu Hart and Noam Nisan. 2013. The Menu-size Complexity of Auctions. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC '13)*. 565–566. DOI: <http://dx.doi.org/10.1145/2482540.2482544>
- Jason D. Hartline and Qiqi Yan. 2011. Envy, truth, and profit. In *EC*. 243–252.
- Alexander S. Kelso, Jr. and Vincent P. Crawford. 1982. Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica* 50, 6 (1982), 1483–1504.

- Herman B Leonard. 1983. Elicitation of Honest Preferences for the Assignment of Individuals to Positions. *J. of Political Economy* 91, 3 (1983), 461–79.
- Andreu Mas-Colell, Michael Dennis Whinston, Jerry R Green, and others. 1995. *Microeconomic theory*. Vol. 1. Oxford university press New York.
- Ahuva Mu’Alem. 2009. On Multi-dimensional Envy-Free Mechanisms. In *ADT*. 120–131.
- L. S. Shapley and M. Shubik. 1971. The assignment game I: The core. *International J. of Game Theory* 1 (1971), 111–130. Issue 1.
- John Thanassoulis. 2004. Hagglng over substitutes. *Journal of Economic Theory* 117, 2 (2004), 217 – 245. DOI: <http://dx.doi.org/10.1016/j.jet.2003.09.002>
- Hal R. Varian. 1974. Equity, envy, and efficiency. *J. of Economic Theory* 9, 1 (1974), 63–91.
- L. Walras. 1874. *Éléments d’économie politique pure; ou, Théorie de la richesse sociale*. Corbaz. <https://books.google.com/books?id=crUEAAAAMAAJ>
- Léon Walras. 2003. *Elements of Pure Economics: Or the Theory of Social Wealth*. Routledge.

A. A NOTE ON EXACT LPE

We give an algorithm that constructs an ϵ -LPE. In contrast, [Feldman et al. 2013] gives an algorithm that produces an exact CWE. In this section we explain why a straightforward adaptation of [Feldman et al. 2013] fails to give an exact LPE.

The fundamental difference between our procedure and the procedure used in [Feldman et al. 2013] is that in [Feldman et al. 2013] prices increase, whereas in our procedure probabilities decrease. Suppose that in [Feldman et al. 2013], say after the second agent picked her bundle, agents 1 and 2 are allocated bundles A and B , respectively. At this point, the prices of A and B simultaneously increase, until one of the agents 1,2 is indifferent between her own bundle and some unallocated bundle. Note that since utilities are linear in prices, by increasing the prices of the allocated bundles simultaneously, the demand structure within the set of allocated bundles remains intact. What changes is that now we have an additional demanded bundle outside the set of allocated bundles. In [Feldman et al. 2013] it is explained how this property is used to produce an exact CWE.

Crucially, we cannot mimic this process because we deal with reduced probabilities, not increased prices, and utility is not linear in probability. One would attempt to address this problem by decreasing probabilities at different rates. Unfortunately, this attempt fails to preserve the demand structure, as demonstrated in the following example.

Consider an instance with two unit-demand agents, 1 and 2, and two lotteries, A and B , with valuations $v_1(A) = 8$, $v_1(B) = 6$, $v_2(A) = 9$, $v_2(B) = 7$, and prices $p_A = 5$, $p_B = 3$. Let q_A and q_B be the scaling factors of lotteries A and B respectively. Simple calculation shows that for agent 1 to prefer A over B , it must hold that $q_A \geq \frac{3}{4}q_B + \frac{1}{4}$. Similarly, for agent 2 to prefer B over A , it must hold that $q_B \geq \frac{9}{7}q_A - \frac{2}{7}$. Combining, we get $q_A \geq 1$. Thus, to keep the desired demand structure, q_A must stay intact throughout the process. (This is in contrast to [Feldman et al. 2013], where prices can always increase simultaneously.) Consequently, we cannot achieve the desired property that one of the agents be indifferent between her lottery and an unallocated one. Suppose, for example, that lottery C is such that $v_1(C) = v_2(C) = \epsilon$. Agent 1 cannot be made indifferent between A and C , since q_A cannot be reduced. For agent 2 to be indifferent between B and C , q_B must drop, but then agent 2 would strictly prefer A over B , so the demand structure changes.