Approximating Equilibrium under Constrained Piecewise Linear Concave Utilities with Applications to Matching Markets*

Jugal Garg† jugal@illinois.edu
Yixin Tao‡ Y.Tao16@lse.ac.uk
László A. Végh§ L.Vegh@lse.ac.uk

Abstract

We study the equilibrium computation problem in the Fisher market model with constrained piecewise linear concave (PLC) utilities. This general class captures many well-studied special cases, including markets with PLC utilities, markets with satiation, and matching markets. For the special case of PLC utilities, although the problem is PPAD-hard, Devanur and Kannan (FOCS 2008) gave a polynomial-time algorithm when the number of goods is constant. Our main result is a fixed parameter approximation scheme for computing an approximate equilibrium, where the parameters are the number of agents and the approximation accuracy. This provides an answer to an open question by Devanur and Kannan for PLC utilities, and gives a simpler and faster algorithm for matching markets as the one by Alaei, Jalaly and Tardos (EC 2017).

The main technical idea is to work with the stronger concept of thrifty equilibria, and approximating the input utility functions by 'robust' utilities that have favorable marginal properties. With some restrictions, the results also extend to the Arrow–Debreu exchange market model.

1 Introduction

Market equilibrium is one of the most fundamental solution concepts in economics, where prices and allocations are such that demand meets supply when each agent gets her most preferred and affordable bundle of goods. Due to the remarkable fairness and efficiency guarantees of equilibrium allocation, it is also one of the preferred solutions for fair division problems even though there may be no money involved in the latter case. A prominent example is competitive equilibrium with equal incomes (CEEI) [35], where a market is created by giving one dollar of virtual money to every agent.

In this paper, we focus on markets with divisible goods. Extensive work in theoretical computer science over the last two decades has led to a deep understanding of the computational complexity of equilibria for the classical models of Fisher and exchange markets, introduced by Fisher [4] and Walras [40] respectively in the late nineteenth century. In a Fisher market, agents have fixed budgets to spend on goods according to their preferences given by utility functions over bundles of goods. CEEI is a special case of this model, where each agent has a budget of one dollar. In the exchange (also known as Arrow–Debreu) market model, the goods are brought to the market by the agents, who can spend their revenue from selling their initial endowments.

Prevalent assumptions on the utility functions in the literature are *(a) monotonicity*, i.e., getting a bundle containing more of each good may not decrease the utility, and *(b) local non-satiation*, i.e., for every bundle of goods, an arbitrary neighborhood contains a bundle with strictly higher utility. A prominent example where these assumptions do not hold is the *one-sided matching market problem*, where each agent needs to be assigned *exactly* one unit of fractional goods in total. Hylland and Zeckhauser [26] introduced an elegant mechanism based on CEEI for the one-sided matching markets. However, more general allocation constraints remain largely unexplored.

In this paper, we consider the equilibrium computation problem when agents have *constrained piecewise linear concave (PLC)* utility functions, defined as follows.

*A full version of the paper can be accessed at https://arxiv.org/abs/2107.05700. The first author is supported by NSF Grant CCF–1942321 (CAREER). The second and third authors received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. ScaleOpt–757481).

†University of Illinois at Urbana-Champaign.
‡London School of Economics and Political Science.
§London School of Economics and Political Science.

DEFINITION 1.1. *The utility function* $u_i : \mathbb{R}^m_+ \to \mathbb{R} \cup \{-\infty\}$ *of agent* i *is* constrained PLC *if it is given in the following form. For some* $p_i, r_i \in \mathbb{N}$*, let* $A_i \in \mathbb{R}^{p_i \times m}$*,* $B_i \in \mathbb{R}^{p_i \times r_i}$*,* $q_i \in \mathbb{R}^m$*,* $s_i \in \mathbb{R}^{r_i}$*,*

$$u_i(x_i) = \begin{cases} \max_{t_i} q_i^\top x_i + s_i^\top t_i \ \textit{s.t.}\ A_i x_i + B_i t_i \le b_i & \textit{if } \exists t_i :\ A_i x_i + B_i t_i \le b_i\,, \\ -\infty & \textit{otherwise.} \end{cases}$$

This general model includes the following well-studied examples:

- Matching markets [1, 26, 38] in the form $u_i(x_i) = \sum_j a_{ij} x_{ij}$ if $\sum_j x_{ij} = 1$, and $-\infty$ otherwise.
- PLC utility functions (see e.g., [14, 21, 22]) $u_i(x_i) = \min_\ell\{\sum_j a^l_{ij} x_{ij} + b^l_i\}$ can be modeled as $u_i(x_i) = \max t$ s.t. $t \le \sum_j a^l_{ij} x_{ij} + b^l_i$, $\forall \ell$. This includes Leontief utilities as a special case.
- Markets with satiation, where an agent may either have the maximum utility limit [3, 11] or consumption constraints [30], which can be easily captured through constrained PLC functions in most cases.

For the special case of PLC utilities, although the problem is PPAD-hard [8][1], Devanur and Kannan [14] and Kakade, Kearns, and Ortiz [29] gave polynomial-time algorithms for computing exact and approximate equilibria, respectively, when the number of goods is a constant. However, the other significant case of constantly many agents turns out to be much more challenging. In [14] an algorithm is given for fixed number of agents with *separable* PLC utilities, but the case with non-separable PLC utilities remained open. Moreover, apart from theoretical interest, designing simpler and faster algorithms for these cases is crucial for their applications.

1.1 Our contributions Our main result is a fixed parameter approximation scheme for computing an approximate equilibrium in Fisher model under constrained PLC utilities, where the parameters are the number of agents and the approximation accuracy. The main technical ideas are to use the stronger concept of *thrifty* equilibria and to approximate the input utility functions by *robust* utilities that have favorable marginal properties.

Before reviewing our algorithm for fixed number of agents, let us start with an easier algorithm for fixed number of goods. In this case, a fairly simple grid search works over all possible price combinations with a small stepsize. This is applicable to the even more general class of *regular* concave utilities (Theorem 3.1). For each price combination, we compute the maximum utility of each agent at these prices, and check whether these utilities can be approximately attained by a feasible allocation also respecting the budget constraints. The existence of an equilibrium guarantees that we find a suitable solution for at least one price combination. This is similar to the grid search approaches used in [26, Appendix B] for matching markets, and for other markets in, e.g., [13, 29, 32].

The natural starting point for fixed number of agents is to perform a grid search over all possible combinations of utility values with a small stepsize. However, even after fixing the desired utility values for each agent, we need to find both allocations and prices, a significantly more challenging task. Our approach is to *(I)* first find an allocation of the goods that meet the utility requirements of each agent, and then *(II)* compute prices for which these allocations form a market equilibrium.

Consider an equilibrium with allocations $x^* = \{x^*_{ij}\}_{i,j}$, prices $p^* = \{p^*_j\}_j$ and utility values $u^* = \{u^*_i\}_i$. For such a two-stage grid search approach to work, a necessary requirement is that given approximate utility values $u^*_i - \delta \le \tilde{u}_i \le u^*_i$, (x, p^*) must form an approximate market equilibrium for every allocation x such that $u_i(x_i) \ge \tilde{u}_i$ for each agent i. This is not true for arbitrary utility functions: not only that x may be very far from x^*, but more importantly, the approximate utility value $\tilde{u}_i$ could be obtained by paying much less than $p^\top x^*_i$.

To address this problem, we make further assumptions both on the utility functions u_i as well as on the equilibrium (x^*, p^*). We require *robust* utility functions, where the change in the utility value is bounded by the change of the budget in a certain critical range of budgets. We then show that every constrained PLC utility function u_i can be approximated by a ξ-robust constrained-PLC utility u^ξ_i for any $\xi > 0$. We run the algorithm for u^ξ_i; the resulting approximate equilibrium will also be an approximate

[1] We note that even the subcases of PLC such as separable PLC is already PPAD-hard for both Fisher [8] and Arrow-Debreu models [6], and Leontief is PPAD-hard for the Arrow-Debreu model [10].

equilibrium for the original u_i with a slightly worse accuracy. The construction of the u_i^ξ's relies on using *perspective functions* of the u_i's.

Robustness on its own however does not suffice. A curious phenomenon for constrained utilities is that an agent may not need to spend their entire budget to obtain their most preferred bundle. For example, if the most favored good has price less than 1 in a matching market, the optimal choice of the agent is to purchase the full unit of this good. We will need to require that (x^*, p^*) is a *thrifty* equilibrium: the agents do not only get their most preferred bundle of goods, but purchase such a bundle at the cheapest possible costs. (In the matching market example, if there is a tie among most preferred goods, all priced less than 1, the agent is only allowed to purchase the cheapest one.) Fortunately, thriftiness can always be assumed (Theorem 2.1): we show that a thrifty equilibrium always exists in the Fisher model for regular concave utilities mentioned above. The proof uses Kakutani's fixed point theorem. To the extent of our knowledge, existence of (even a non-thrifty) equilibrium is not implied by previous results for constrained PLC utilities.[2]

Let us now describe the algorithm for robust utilities. In stage *(I)*, we can find allocations delivering the utility guesses by solving a linear program. In stage *(II)*, the goal is to find prices p that form an approximate equilibrium with x. The most challenging part is to ensure that the maximum utility profile available at p is close to the guess $\tilde{u}_i$ for each agent i. This is achieved by considering the dual of the utility maximizing linear program, and applying a variable transformation. After these reductions, suitable prices can be found by linear programming (Theorem 4.2).

In the above algorithm, we assume that empty allocation is feasible, i.e., $u_i(0) = 0$. However, this is no longer true in matching markets where $u_i(0) = -\infty$, and hence the above algorithm does not directly apply here. We proceed with the natural approach by relaxing the matching constraints to $\sum_j x_{ij} \le 1$ for every i. For this relaxation to work, we need to add the requirement on both exact and approximate equilibria that the minimum price is 0. This can be ensured by exploiting a natural price transformation in the problem. We show that this approach works even for a more general model of PLC matching markets with $u_i(x_i) = \min_\ell\{\sum_j a_{ij}^l x_{ij} + b_i^l\}$ if $\sum_j x_{ij} = 1$ and $-\infty$ otherwise (Theorems 5.2 and 5.3).

The papers [14] and [1] give polynomial-time algorithms for computing exact equilibria for the special cases mentioned earlier using a *cell decomposition* technique. Note that in both PLC and matching markets, it is possible that all equilibria are irrational [18, 38]; exact equilibria in these works are represented as roots of polynomials. The cell decomposition arguments partition the parameter space by polynomial surfaces such that in each cell it is easy to decide whether a solution in the particular configuration exists; the number of cells can be bounded using results from algebraic geometry. While the number of cells is polynomial, the results for fixed number of agents (for separable PLC in [14] and for matching markets in [1]) require solving $m^{\mathrm{poly}(n)}$ subproblems and thus may not be practical. In contrast, our algorithm is a fixed parameter scheme in n and the accuracy ε; we need to solve $O((n/\varepsilon^2)^n)$ polynomial-size linear programs.[3] Hence, the complexity of finding an approximate equilibrium is much lower.

For matching markets, we also show that the set of equilibria is non-convex by a simple example of 3 agents and 3 goods with tri-valued utility values $a_{ij} \in \{0, 1, 2\}$. To the best of our knowledge, this is the first proof of the non-convexity of equilibria in matching markets.[4] Moreover, our example is the simplest one can hope for as for both the bi-valued utility values and two agents case, the set of equilibria is convex; see e.g., [23].

Finally, we show that our algorithms also extend to the more general case of the Arrow-Debreu model under PLC utilities. The additional challenge here is to handle budgets that now depend on the prices. For both cases of fixed number of agents and fixed number of goods, we approximate the utilities by robust utilities. This can be done in a simpler way using the special form of the utilities, and in particular it guarantees a lower bound on the minimum price. We refer the reader to the full version for the Arrow-Debreu model and its results.

[2] We note that [36] shows NP-hardness of checking equilibrium existence in Fisher model under separable PLC utilities, which seems to require two conditions: first, the sum of prices must equal the sum of budgets, and second, they implicitly assume that the agents are thrifty. Hence, there is no contradiction.

[3] We note that $\varepsilon \le 1/(cn)$ is needed for a meaningful approximate equilibrium.

[4] We note that [37] presents an example to show non-convexity of equilibria in matching markets. However, their latest version [38] does not contain that example, which seems to have only one equilibrium.

1.2 Related work Market equilibrium is an intensely studied concept with a variety of applications, so we briefly mention further relevant results. For the classical Fisher model, polynomial-time algorithms are obtained when agents have linear [15, 31, 39], weak gross substitutes [9], and homogeneous utility functions [19]. For separable PLC utilities, the problem is PPAD-hard [8].

For the constrained Fisher model, the most famous problem is the Hylland-Zeckhauser scheme for the one-sided matching markets, for which [26] shows the existence of equilibrium, which is recently simplified [5]. For matching markets, polynomial-time algorithms are obtained for special cases of constantly many agents (or goods) [1] and dichotomous utilities [38]. Settling its exact complexity is currently open.

Very recently, [28] considers Fisher markets with additional linear constraints, which includes matching markets but not the PLC utilities studied in this paper. It gives a simple fixed-point iterative scheme that converges to an equilibrium in numerical experiments, among other structural results. In particular, it provides a non-convexity example with additional linear utilities, which we note is not a matching market example.

For the classical Arrow-Debreu model, polynomial-time algorithms are obtained when agents have linear [16, 17, 24, 27, 41] and weak gross substitutes [2, 9, 20] utilities, and beyond that, the problem is essentially PPAD-hard [6, 7, 10, 22].

For the constrained Arrow-Debreu model, an exact equilibrium may not exist even in the case of matching markets [26]. For this, [23] gives the existence of an approximate equilibrium and a polynomial-time algorithm for computing it under dichotomous utilities.

Overview The rest of the paper is organized as follows. Section 2 defines all models and definitions. Section 3 presents an algorithm for computing an approximate Fisher equilibrium under regular concave utilities for a fixed number of goods. Section 4 gives an algorithm for computing an approximate Fisher equilibrium under constrained PLC utilities for a fixed number of agents. Section 5 extends algorithms to PLC matching markets and presents an example showing the non-convexity of equilibria.

2 Models and definitions

Consider a market with n agents and m divisible goods. We assume without loss of generality that there is a unit supply of each good. Each agent i has a concave utility function $u_i : \mathbb{R}^m_+ \to \mathbb{R} \cup \{-\infty\}$.

Definition 2.1. *We say that the utility function $u_i : \mathbb{R}^m_+ \to \mathbb{R} \cup \{-\infty\}$ is* regular, *if*

- *The function u_i is concave and the domain $K_i = \{x_i \in \mathbb{R}^m_+ : u_i(x_i) > -\infty\}$ is closed.*
- *u_i restricted to K_i is Lipschitz continuous, i.e. $|u_i(x_i) - u_i(y_i)| \le L\|x_i - y_i\|_2$ for $x_i, y_i \in K_i$.*
- *$u_i(0) = 0$.*

We assume that the Lipschitz constant L is the same for all utility functions and is known a priori. This will be relevant for the computational complexity of (approximately) solving convex programs with objective u_i. The main requirement in the assumption $u_i(0) = 0$ is that $0 \in K_i$, i.e., the empty allocation is feasible. If that holds, we can shift the utility function to $u_i(0) = 0$. Our main focus will be on the constrained PLC utilities defined in the introduction.

Lemma 2.1. *Every constrained PLC utility function u_i (Definition 1.1) with $u_i(0) = 0$ is regular. For the Lipschitz parameter L, $\log L$ is polynomially bounded in the bit-complexity of the input.*

Proof. The first property is immediate, and the Lipschitz bound follows by [34, Corollary 3.2a and Theorem 10.5]. □

For prices $p \in \mathbb{R}^m_+$ and a budget w_i, we define the optimal utility value

$$V_i(p, w_i) = \max_{x_i \in \mathbb{R}^m_+} \left\{ u_i(x_i) : p^\top x_i \le w_i \right\} \ ,$$

and the *demand correspondence* as the set of utility maximizing bundles that can be purchased at the given budget:

$$D_i(p, w_i) = \arg\max_{x_i \in \mathbb{R}^m_+} \left\{ u_i(x_i) : p^\top x_i \le w_i \right\} \ ,$$

Let

$$V_i^{\max} = \max_{x_i \in [0,1]^m} u_i(x_i)\,.$$

be the maximum utility value achievable by purchasing at most 1 unit from each good. Clearly, $V_i^{\max} \geq u_i(0) = 0$. Throughout, we make the following normalization assumption:

$$V_i^{\max} \leq 1 \quad \text{for each agent } i. \tag{2.1}$$

Let

$$C_i(p, w_i) = \min_{x_i \in \mathbb{R}_+^m} \left\{ p^\top x_i : x_i \in D_i(p, w_i) \right\}$$

be the minimum cost of an optimal bundle; we call this the *thrifty cost*. If the market satisfies non-satiation, then $C_i(p, w_i) = w_i$, but it can be strictly less otherwise. Note that if $C_i(p, w_i) < w_i$, then $V_i(p, w_i) = \max_{x_i \in \mathbb{R}_+^m} u_i(x_i)$. We define the *thrifty demand correspondence* as the set of cheapest optimal bundles.

$$D_i^t(p, w_i) = \arg\min_{x_i \in \mathbb{R}_+^m} \left\{ p^\top x_i : x_i \in D_i(p, w_i) \right\}\,.$$

Finally, we let

$$C_i^{\min}(p) = \min_{x_i \in \mathbb{R}_+^m} \left\{ p^\top x_i : x_i \in \mathbb{R}^m, u_i(x_i) \geq V_i^{\max} \right\}$$

denote the minimum cost to achieve $V_i^{\max}$ at prices p. Note that we also allow bundles here that are not in $[0,1]^m$, i.e., may use more than one unit of an good.

2.1 The Fisher market model In the Fisher market model, we are given n agents and m divisible goods of unit supply each. Each agent has a budget w_i and a regular utility function $u_i : \mathbb{R}_+^m \to \mathbb{R} \cup \{-\infty\}$. We assume

$$V_i^{\max} > 0 \quad \text{for each agent } i. \tag{2.2}$$

If $V_i^{\max} = 0$, then by concavity we must have $u_i(x_i) \leq 0$ for all $x_i \in \mathbb{R}_+^m$. We can remove such agents, as they can always be allocated $x_i = 0$ at equilibrium.

Definition 2.2. (Fisher equilibrium) *In a Fisher market with utilities $\{u_i\}_i$ and budgets $\{w_i\}_i$, the allocations and prices $(\{x_i\}_i, \{p_j\}_j)$ form a* market equilibrium *if*

- *$x_i \in D_i(p, w_i)$ for each agent i, i.e., each agent buys an optimal bundle subject to budget constraint;*
- *the market clears, i.e., $\sum_i x_{ij} \leq 1$, and $\sum_i x_{ij} = 1$ if $p_j > 0$ for every good j.*

Further, $(\{x_i\}_i, \{p_j\}_j)$ is a thrifty market equilibrium *if we require the stronger $x_i \in D_i^t(p, w_i)$ for each agent i.*

We prove the following theorem in the full version, which shows that regular utilities suffice for the existence of an equilibrium.

Theorem 2.1. *If all agents' utility functions are regular, then a thrifty market equilibrium always exists.*

Definition 2.3. (approximate Fisher equilibrium) *In a Fisher market with utilities $\{u_i\}_i$ and budgets $\{w_i\}_i$ that satisfies assumption* (2.1)*, the allocations and prices $(\{x_i\}_i, \{p_j\}_j)$ form a (σ, λ)-*approximate market equilibrium *if*

- $u_i(x_i) \geq V_i(p, w_i) - \lambda$;
- $p^\top x_i \leq w_i + \sigma \sum_i w_i$;
- $\sum_i x_{ij} \leq 1$, *and* $\sum_j p_j(1 - \sum_i x_{ij}) \leq \sigma \sum_i w_i$.

*Similarly, a (σ, λ)-*approximate thrifty market equilibrium *satisfies $p^\top x_i \leq C_i(p, w_i) + \sigma \sum_i w_i$ instead of the second constraint. A (σ, σ)-approximate (thrifty) market equilibrium will be also referred to as a σ-*approximate (thrifty) market equilibrium.

Note that, in order to get a meaningful approximate equilibrium solution, one needs to select $\sigma < 1/(cn)$ for some constant c, since the error term is $\sigma \sum_i w_i$.

3 Approximate Fisher equilibrium for fixed number of goods

As a warm-up, we give a simple algorithm for finding an ε-approximate thrifty market equilibrium in Fisher markets for a fixed number of goods. The algorithm amounts to approximately solving $O(n\left(\frac{m}{\varepsilon}\right)^m)$ convex programs. This is similar to the grid search approaches used, e.g., in [13, 26, 29, 32].

We assume that the utility functions $u_i(x_i)$ are represented by value oracles. For given prices p, the maximum utility $V_i(p, w_i)$ and the thrifty cost $C_i(p, w_i)$ can be obtained as the optimal solution to convex programs. Using a convex programming algorithm such as the ellipsoid method, we can compute a ε-approximate optimal solutions in oracle-polynomial time in n, m, the bit-complexity of the vector p, w_i, $\log L$, and $\log(1/\varepsilon)$ [25].

Further, we define the function $F : \mathbb{R}_+^m \to \mathbb{R}$ as the optimal solution to the following convex program in the variables $\{x_i\}_i$.

$$\begin{aligned} F(p) = \min \delta & \\ u_i(x_i) \geq V_i(p, w_i) - \delta, & \quad \forall i \\ \sum_i x_{ij} \leq 1, & \quad \forall j \\ \sum_j x_{ij} p_j \leq C_i(p, w_i) + \delta \sum_i w_i, & \quad \forall i \\ \sum_j p_j \left(1 - \sum_i x_{ij}\right) \leq \delta \sum_i w_i & \\ x, \delta \geq 0, & \end{aligned} \tag{3.3}$$

To compute a ε-approximate solution for given prices p, we first find $(\varepsilon/2)$-approximate values for $V_i(p, w_i)$ for all i and $(\varepsilon \sum_i w_i/2)$-approximate values for $C_i(p, w_i)$; then, we again use a convex programming algorithm to find a $(\varepsilon/2)$-approximate solution to the resulting program.

Lemma 3.1. *If $F(p) \leq \sigma$, then the prices p and allocations x_i give a σ-approximate thrifty market equilibrium. If (p^*, x^*) forms an exact thrifty market equilibrium, and prices $p \in \mathbb{R}^m$ satisfy $p_j^* \leq p_j \leq p_j^* + \frac{\sigma}{m} \sum_i w_i$, then $F(p) \leq \sigma$.*

Proof. The first claim is immediate by the definition of an approximate thrifty market equilibrium. For the second claim, $F(p^*) = 0$ with the optimal solution x^*. We show that (x^*, σ) is feasible to (3.3), showing that $F(p) \leq \sigma$. Since $p \geq p^*$, we have $V_i(p, w_i) \leq V_i(p^*, w_i)$, verifying the first constraint.

To verify the third constraint, we first show that $C_i(p, w_i) \geq C_i(p^*, w_i)$. This is immediate if $C_i(p, w_i) = w_i$. If $C_i(p, w_i) < w_i$, then $V_i(p, w_i) = \max_{x \in \mathbb{R}^m} V_i(x)$, the maximum utility without budget constraint; consequently, $V_i(p^*, w_i) = \max_{x \in \mathbb{R}^m} V_i(x)$. Purchasing such a maximum utility bundle cannot be cheaper at prices p, since $p \geq p^*$. Hence,

$$\sum_j p_j x_{ij}^* \leq \sum_j p_j^* x_{ij}^* + \left(\sum_j x_{ij}^*\right) \frac{\sigma}{m} \left(\sum_i w_i\right) \leq C_i(p^*, w_i) + \sigma \sum_i w_i \leq C_i(p, w_i) + \sigma \sum_i w_i \,.$$

The last constraint in (3.3) follows by $\sum_j p_j \sum_i x_{ij}^* \geq \sum_j p_j^* \sum_i x_{ij}^*$, and $\sum_j p_j \leq \sum_j p_j^* + \sigma \left(\sum_i w_i\right)$. ∎

Theorem 3.1. *Given a Fisher market with n agents, m goods, and regular concave utility functions u_i given by oracle access, we can compute a ε-approximate thrifty market equilibrium by approximately solving $O(n\left(\frac{m}{\varepsilon}\right)^m)$ convex programs, each in oracle-polynomial time in in n, m, the bit-complexity of the w_i's, $\log L$, and $\log(1/\varepsilon)$.*[5]

Proof. We enumerate all price vectors $p_j = k_j \frac{\varepsilon}{2m} \sum_i w_i$ for all integers k_i such that $0 \leq k_j \leq \frac{2m}{\varepsilon} + 1$. For each price vector p, we find a $(\varepsilon/2)$-approximate solution to (3.3). We output any price vector p for which a solution (x, δ) with $\delta \leq \varepsilon$ is found. Theorem 2.1 and Lemma 3.1 guarantee the existence of such a solution. ∎

[5] This is essentially $(0, \epsilon)$-approximate equilibrium if we can solve the convex program (3.3) exactly; otherwise, we can get (λ, ϵ)-approximate equilibrium for arbitrary λ at an additional $\log(1/\lambda)$ factor.

4 Approximate Fisher equilibrium for fixed number of agents

In this section, we design a polynomial-time algorithm to compute an approximate equilibrium with constant number of agents for constrained PLC utility functions. Recall that $V_i^{\max}$ is the maximum utility achievable on $[0,1]^m$, and $C_i^{\min}(p)$ is the minimum cost for achieving utility $V_i^{\max}$ at prices p. Also recall assumptions (2.1) and (2.2), that $V_i^{\max} \in (0,1]$ for all agents; we make this assumption throughout. The following class of utility functions plays a key role:

DEFINITION 4.1. (ξ-ROBUST UTILITY FUNCTION) *A regular utility function u_i is ξ-robust, if for any bundle $x_i \in \mathbb{R}_+^m$, prices $p \in \mathbb{R}_+^m$ and $\sigma > 0$ such that $p^\top x_i \le C_i^{\min}(p) - \sigma$, there exists a bundle $y \in \mathbb{R}_+^m$ such that $p^\top y \le p^\top x_i + \sigma$ and $u_i(y) \ge u_i(x_i) + \frac{\sigma\xi}{\sum_j p_j}$.*

For this class of utility functions, in Section 4.1, we show how to compute an approximate thrifty market equilibrium. However, not all regular utilities are ξ-robust for some $\xi > 0$. For example, $u_i(x_i) = 0$ is not ξ-robust for any positive ξ, while $u_i(x_i) = \varepsilon \sum_j x_{ij}$ is εm-robust for $\varepsilon > 0$. In Section 4.2, we show how to approximate any regular utilities by robust utilities.

In the overall algorithm, we approximate the regular utility functions by ξ-robust utility functions. Then, we can calculate an approximate thrifty market equilibrium for these ξ-robust utility functions. Finally, we show that the approximate market equilibrium we calculated will also be the approximate market equilibrium for the original utility functions (but not necessarily a thrifty one).

4.1 Approximate equilibria for ξ-robust utilities By Theorem 2.1, we know there exists a thrifty market equilibrium for ξ-robust utilities. Let (x^*, p^*) denote a thrifty market equilibrium and let u_i^* denote the utility achieved by agent i at the equilibrium.

The algorithm has two steps: first, guess each agent's utility at equilibrium and compute a feasible allocation giving each agent at least the guessed utility, and second, compute the prices that, together with the calculated allocation, give an approximate market equilibrium. In Theorem 4.1, we give a $(n\delta/\xi, 2\delta)$-approximate thrifty equilibrium for some parameter $\delta > 0$. This is achieved by solving $O(1/\delta^n)$ convex programs.

4.1.1 Guessing utilities and computing the allocation We first guess agents' utilities at equilibrium by enumerating all possible utilities of each agent, $\tilde{u}_i = k_i\delta$, for $0 \le k_i \le \lceil \frac{1}{\delta} \rceil + 1$. Then, we compute a feasible allocation $x = (x_1, \dots, x_n)$ giving $\tilde{u}_i$ utility to agent i using the following program:

$$\begin{aligned} & u_i(x_i) \ge \tilde{u}_i\,, \quad \forall i \\ & \sum_i x_{ij} \le 1\,, \quad \forall j \\ & x \ge 0 \end{aligned} \tag{4.4}$$

If (4.4) is infeasible, we move to the next utility profile. The following lemma shows that if we have a right guess on the utilities, then, with the equilibrium price p^*, the spending of agent i at x_i should be similar to that at x_i^*.

LEMMA 4.1. *Assume the utility functions are ξ-robust, $u_i^* - \delta < \tilde{u}_i \le u_i^*$, and $\{x_i\}_i$ is a feasible solution to* (4.4)*. Then,*

(i) $(p^*)^\top x_i^* - \dfrac{\delta \sum_i w_i}{\xi} \le (p^*)^\top x_i \le (p^*)^\top x_i^* + \dfrac{\delta n \sum_i w_i}{\xi}$, *and*

(ii) $\displaystyle\sum_j p_j^* \left(1 - \sum_i x_{ij}\right) \le \frac{n\delta \sum_i w_i}{\xi}$.

Proof. We first consider the lower bound in *(i)*. Let us denote $\sigma = \delta \sum_i w_i/\xi$. For a contradiction, assume $(p^*)^\top x_i^* - \sigma > (p^*)^\top x_i$. Since this is a thrifty equilibrium, $(p^*)^\top x_i^* = C_i(p^*, w_i) \le C_i^{\min}(p^*)$, since the optimal bundle x_i^* in $[0,1]^m$.

By the ξ-robustness property, there exists a bundle y such that $(p^*)^\top y \le (p^*)^\top x_i + \sigma < (p^*)^\top x_i^* \le w_i$ and

$$u_i(y) \ge u_i(x_i) + \frac{\sigma\xi}{\sum_j p_j^*} \ge u_i(x_i) + \frac{\delta \sum_i w_i}{\sum_j p_j^*} > u_i^*(x_i) - \delta + \delta = u_i^*(x_i)\,,$$

a contradiction since $u_i^*(x_i)$ is the maximum utility at budget w_i. The third inequality uses that all goods with $p_j^* > 0$ are fully sold, and therefore the sum of the prices is at most the sum of the budgets.

For the upper bound, note that $p_j^* = 0$ if $\sum_i x_{ij}^* < 1$. Therefore, $\sum_j p_j^* \left(\sum_i x_{ij}^*\right) \geq \sum_j p_j^* \left(\sum_i x_{ij}\right)$. Additionally, the lower bound give $\sum_j p_j^* x_{ij} \geq \sum_j p_j^* x_{ij}^* - \frac{\delta \sum_i w_i}{\xi}$ for each i, which completes the proof. Part *(ii)* is immediate by summing up the lower bounds in part *(i)* for all i. ∎

4.1.2 Computing the prices Assume that for the guesses $\{\tilde{u}_i\}_i$, we found an allocation x that satisfies (4.4), i.e., and allocation of the goods that provides at least $\tilde{u}_i$ amount of utility to each i. In what follows, our goal is to find prices p that form an approximate equilibrium with x. This is the most challenging part of the algorithm. The prices have to satisfy the following three conditions.

First condition: utility upper bound If the guesses $\{\tilde{u}_i\}_i$ were approximately correct, then the maximum utility $V_i(p, w_i)$ achievable at prices p should be close to $\tilde{u}_i$ for each agent i.

Since the utility functions are constrained PLC, we can compute $V_i(p, w_i)$ as

$$\begin{aligned}
\max \quad & q_i^\top z_i + s_i^\top t_i \\
\text{s.t.} \quad & A_i z_i + B_i t_i \leq b_i \\
& p^\top z_i \leq w_i \\
& z_i \geq 0
\end{aligned}$$

The dual of this program is as follows, using variables γ_i and β_i for the first two constraints, respectively.

$$\begin{aligned}
\min \quad & b_i^\top \gamma_i + w_i \beta_i \\
\text{s.t.} \quad & A_i^\top \gamma_i + \beta_i p \geq q_i \\
& B_i^\top \gamma_i = s_i \\
& \gamma_i, \beta_i \geq 0
\end{aligned} \tag{4.5}$$

For every feasible dual solution, the objective value provides an upper bound on the optimal utility agent i can get. Therefore, $V_i(p, w_i) \leq \tilde{u}_i + \delta$ if and only if there exists a feasible solution (γ_i, β_i) to (4.5) such that $\gamma_i^\top b_i + \beta_i w_i \leq \tilde{u}_i + \delta$.

However, if we also consider the prices p as variables, the program is not linear anymore. For this reason, we use a variable substitution, by letting

$$\frac{1}{\beta_i} \triangleq \overline{\beta_i} \quad \text{and} \quad \frac{\gamma_i}{\beta_i} \triangleq \overline{\gamma_i}$$

be the variables and we set a lower bound on $\frac{1}{\beta_i}$ such that the optimal solution of (4.5) doesn't change much.

Lemma 4.2. *For $\delta \in (0,1)$, consider a feasible solution $(\overline{\gamma_i}, \overline{\beta_i}, p)$ to following program,*

$$\begin{aligned}
& b_i^\top \overline{\gamma_i} + w_i \leq \overline{\beta_i}(\tilde{u}_i + 2\delta) \\
& {A_i}^\top \overline{\gamma_i} + p \geq \overline{\beta_i} q_i \\
& {B_i}^\top \overline{\gamma_i} = \overline{\beta_i} s_i \\
& w_i \leq \overline{\beta_i}
\end{aligned} \tag{4.6}$$

Then, the optimal utility for agent i to achieve with price p is at most $\tilde{u}_i + 2\delta$. Additionally, if $u_i^* - \delta < \tilde{u}_i \leq u_i^*$, then there exist $\overline{\gamma_i}$ and $\overline{\beta_i}$ such that $(\overline{\gamma_i}, \overline{\beta_i}, p^*)$ is a solution to this program.

Proof. First, if $(\overline{\gamma_i}, \overline{\beta_i}, p)$ is a feasible solution to (4.6), then $(\beta_i = \frac{1}{\overline{\beta_i}}, \gamma_i = \frac{\overline{\gamma_i}}{\overline{\beta_i}}, p)$ is a feasible solution to (4.5). Therefore, the optimal utility one can get at price p is at most $b_i^\top \gamma_i + \beta_i w_i = \frac{1}{\overline{\beta_i}}(b_i^\top \overline{\gamma_i} + w_i) \leq \tilde{u}_i + 2\delta$.

For the second part, consider the optimal solution (γ_i, β_i) to (4.5) with the price p^*. Then, clearly, $b_i^\top \gamma_i + \beta_i w_i = u_i(x_i^*)$ and (γ_i, β_i) is also a feasible solution with price p^* and $w_i = 0$, which implies $b_i^\top \gamma_i \geq u_i(0) = 0$. Combining with the fact that $b_i^\top \gamma_i + \beta_i w_i = u_i(x_i^*) \leq 1$, we get $\beta_i \leq \frac{1}{w_i}$. Therefore, if we consider the solution $\overline{\gamma_i} = \frac{\gamma_i}{\max\{\beta_i, \frac{\delta}{w_i}\}}$, $\overline{\beta_i} = \frac{1}{\max\{\beta_i, \frac{\delta}{w_i}\}}$, then it satisfies all conditions in (4.6) as $b_i^\top \gamma_i + \beta_i w_i = u_i(x_i^*) < \tilde{u}_i + \delta$. ∎

We note that this is the only part in the algorithm where we rely on the particular form of constrained PLC utilities; all other arguments work more generally, for regular utilities.

Second condition: budget constraint The cost of the allocation at prices p must not violate the budget constraints by much:

$$p^\top x_i \le w_i + \frac{n\delta \sum_i w_i}{\xi}, \quad \forall i\,. \tag{4.7}$$

Third condition: market clearing The market needs to approximately clear:

$$\sum_j p_j \left(1 - \sum_i x_{ij}\right) \le \frac{n\delta \sum_i w_i}{\xi}. \tag{4.8}$$

Lemma 4.1 implies the following:

LEMMA 4.3. *Assume that the utility functions are ξ-robust. If for all agents i, $u_i^* - \delta < \tilde{u}_i \le u_i^*$, then for any allocation x for which* (4.4) *holds, the optimal prices p^* satisfy* (4.7) *and* (4.8)*.*

Note that (4.6), (4.7), and (4.8) are linear in p, $\overline{\beta}$, and $\overline{\gamma}$. Combining Lemmas 4.2 and 4.3, we have the following:

THEOREM 4.1. *Suppose the utility functions are ξ-robust. For any $\tilde{u}_i$, x and p such that* (4.4), (4.6), (4.7)*, and* (4.8) *holds, (x, p) is a $(n\delta/\xi, 2\delta)$-approximate market equilibrium. Additionally, let (x^*, p^*) be any thrifty market equilibrium. If $u_i(x_i^*) - \delta < \tilde{u}_i \le u_i(x_i^*)$ for all i, then for any x such that* (4.4) *holds, p^* is a solution to* (4.6), (4.7)*, and* (4.8)*.*

4.2 Approximating regular utilities by ξ-robust utilities We introduce an approach to approximate regular utilities by ξ-robust utilities. The construction works for the general class of regular functions; for constrained PLC utilities, we show that this operation yields a constrained PLC utility.

Consider a regular utility function $u_i : \mathbb{R}_+^m \to \mathbb{R} \cup \{-\infty\}$, and define the *perspective function* with domain $\mathbb{R}_+^{m+1}$; see [33, Chapter 5].

$$\hat{u}_i(x, \alpha) = \begin{cases} \alpha u_i\left(\frac{x}{\alpha}\right)\,, & \text{if } \alpha > 0\,, \\ \lim_{\alpha\to 0} \alpha u_i\left(\frac{x}{\alpha}\right)\,, & \text{if } \alpha = 0\,. \end{cases}$$

If u_i is concave and upper semicontinuous (that hold for regular utilities), then so is $\hat{u}_i$ [12, Proposition 2.3(ii)]. Also note that u_i is positively homogeneous. For given $\xi > 0$, we define $u_i^\xi : \mathbb{R}_+^m \to \mathbb{R} \cup \{-\infty\}$, where

$$\begin{aligned} u_i^\xi(x_i) = \max \quad & \hat{u}_i(x', \alpha) + \hat{u}_i(x'', 1-\alpha) + (1-\alpha)\xi \\ \text{s.t.} \quad & \hat{u}_i(x'', 1-\alpha) \ge (1-\alpha)V_i^{\max} \\ & x' + x'' = x_i \\ & 0 \le \alpha \le 1 \\ & x', x'' \ge 0\ . \end{aligned}$$

Note that the maximum exists as the objective is convex over a compact domain.

LEMMA 4.4. *For every regular utility function u_i and $\xi > 0$, the following hold for u_i^ξ:*

(i) $u_i(x_i) \le u_i^\xi(x_i) \le u_i(x_i) + \xi$.

(ii) We have $\max_{x\in[0,1]^m} u_i^\xi(x) = V^{\max} + \xi$*. For any price vector p, the minimum cost of achieving utility $V_i^{\max} + \xi$ for u_i^ξ is the same as the minimum cost $C_i^{\min}(p)$ of achieving utility $V_i^{\max}$ for u_i.*

Proof. **Part *(i)*:** The lower bound $u_i(x_i) \le u_i^\xi(x_i)$ holds since $x' = x_i$, $x'' = 0$, $\alpha = 1$ is a feasible solution in the definition of $u_i^\xi(x_i)$. The upper bound follows by the concavity of $\hat{u}_i$. For any $x' + x'' = x_i$ and $\alpha \in [0, 1]$, we have

$$\hat{u}_i(x', \alpha) + \hat{u}_i(x'', 1-\alpha) + (1-\alpha)\xi \le \hat{u}_i\left(\frac{x_i}{2}, \frac{1}{2}\right) + (1-\alpha)\xi \le u_i(x_i) + \xi\,.$$

Part *(ii)*: By the previous part, $\max_{x\in[0,1]^m} u_i^\xi(x) \le V^{\max} + \xi$. Equality is achieved for any $x_i \in [0,1]^m$ for which $u_i(x_i) = V_i^{\max}$ with the choice $x' = 0$, $x'' = x$, $\alpha = 0$. For the second part, consider any price vector p. Note that the set $\{x_i \mid u_i^\xi(x_i) \ge V_i^* + \xi\}$ is the same as $\{x_i \mid u_i(x_i) \ge V_i^*\}$, which completes the proof. ∎

The next lemma asserts the key property for ξ-robustness:

LEMMA 4.5. *Given prices $p \in \mathbb{R}_+^m$ and $\sigma > 0$, let $x_i \in \mathbb{R}_+^m$ be a bundle such that $p^\top x_i \le C_i^{\min}(p) - \sigma$. Then, there exists a bundle $y \in \mathbb{R}_+^m$ such that $p^\top y \le p^\top x_i + \sigma$, and $u_i^\xi(y) \ge u_i^\xi(x_i) + \frac{\sigma\xi}{\sum_j p_j}$.*

Proof. Let us use the notation $C = C_i^{\min}(p)$; by Lemma 4.4(ii), this is the minimum cost of a bundle of utility $V_i^{\max} + \xi$. Let us use the combination $x_i = x_i' + x_i''$ and $\alpha \in [0,1]$ that gives the value of $u_i^\xi(x_i)$, that is,

$$u_i^\xi(x_i) = \hat{u}_i(x', \alpha) + \hat{u}_i(x'', 1-\alpha) + (1-\alpha)\xi,$$

such that $\hat{u}_i(x'', 1-\alpha) \ge (1-\alpha)V_i^{\max}$. By definition, $p^\top x'' \ge C$. Thus, $\alpha p^\top x' + (1-\alpha)C \le p^\top x_i \le C - \sigma$, implying

$$\alpha \ge \frac{\sigma}{C} \ge \frac{\sigma}{\sum_j p_j},$$

where the last inequality uses C is the cost of a bundle in $[0,1]^m$. Let $z \in [0,1]^m$ be a bundle such that $u_i(z) = V_i^{\max}$ and $p^\top z = C$; such a bundle exists since $[0,1]^m$ is a compact domain. Let

$$\beta = \frac{\sigma}{\sum_j p_j}, \quad y' = \frac{\alpha - \beta}{\alpha} \cdot x', \quad y'' = x'' + \beta z, \quad \text{and} \quad y = y' + y''.$$

Note that

$$p^\top y < p^\top x + \beta p^\top z = p^\top x + \frac{\sigma C}{\sum_j p_j} \le p^\top x + \sigma,$$

satisfying the required bound on the cost. The rest of the proof amounts to showing $u_i^\xi(y) \ge u_i^\xi(x_i) + \frac{\sigma\xi}{\sum_j p_j}$.

CLAIM 4.1. $\hat{u}_i(y'', 1-\alpha+\beta) \ge \hat{u}_i(x'', 1-\alpha) + \beta V_i^{\max} \ge (1-\alpha+\beta)V_i^{\max}$.

Proof. By the homogeneity and concavity of $\hat{u}_i$,

$$\hat{u}_i(y'', 1-\alpha+\beta) = 2\hat{u}_i\left(\frac{y''}{2}, \frac{1-\alpha+\beta}{2}\right) \ge \hat{u}_i(x'', 1-\alpha) + \hat{u}_i(\beta z, \beta) \ge \hat{u}_i(x'', 1-\alpha) + \beta V_i^{\max}.$$

The last inequality in the claim follows by noting that also $\hat{u}_i(x'', 1-\alpha) \ge (1-\alpha)V_i^{\max}$. ∎

CLAIM 4.2. $\hat{u}_i(y', \alpha-\beta) = \frac{\alpha-\beta}{\alpha}\hat{u}_i(x', \alpha) > \hat{u}_i(x', \alpha) - \beta V_i^{\max}$.

Proof. The first inequality is by definition of $\hat{u}_i$. The second inequality is equivalent to $\hat{u}_i(x', \alpha) < \alpha V_i^{\max}$. Assume for a contradiction $\hat{u}_i(x', \alpha) \ge \alpha V_i^{\max}$. Then, replacing x'' by $x' + x''$ and α by 0 results in a better combination using the concavity of $\hat{u}_i$ as in the previous claim. ∎

By Claim 4.1, $y = y' + y''$ is a feasible decomposition in the definition of u_i^ξ with coefficient $\alpha - \beta$. Further, note that $\hat{u}_i(y', \alpha-\beta) = \frac{\alpha-\beta}{\alpha}\hat{u}_i(x', \alpha)$ We get

$$\begin{aligned} u_i^\xi(y) &\ge \hat{u}_i(y', \alpha-\beta) + \hat{u}_i(y'', 1-\alpha+\beta) + (1-\alpha+\beta)\xi \\ &\ge \hat{u}_i(y', \alpha-\beta) + \hat{u}_i(x'', 1-\alpha) + \beta V_i^{\max} + (1-\alpha+\beta)\xi \\ &\ge \hat{u}_i(x', \alpha) + \hat{u}_i(x'', 1-\alpha) + (1-\alpha+\beta)\xi \\ &\ge u_i^\xi(x_i) + \beta\xi = u_i^\xi(x_i) + \frac{\sigma\xi}{\sum_j p_j}, \end{aligned}$$

as required. The second inequality used Claim 4.1 and the third inequality used Claim 4.2. ∎

Let us now turn to constrained PLC utilities.

LEMMA 4.6. *Let u_i be a constrained PLC utility with $u_i(0) = 0$ and $V_i^{\max} > 0$, and $\xi > 0$. Then, u_i^ξ is also a constrained PLC utility, and is ξ-robust. The bit-length of the LP description of u_i^ξ is polynomial in the LP description of u_i and of $\log \xi$.*

Proof. Recall the form of the constrained PLC utility as

$$u_i(x_i) = \max_{t_i} q_i^\top x_i + s_i^\top t_i \text{ s.t. } A_i x_i + B_i t_i \le b_i \,,$$

where the value is $-\infty$ if the problem is infeasible. It is easy to verify that the following linear program gives an equivalent description of u_i^ξ:

$$\begin{aligned} u_i^\xi(x) &= q_i^\top x_i + s_i^\top t_i + (1-\alpha)\xi \\ & A_i x_i' + B_i t_i' \le \alpha b_i \\ & A_i x_i'' + B_i t_i'' \le (1-\alpha) b_i \\ & q_i^\top x_i'' + s_i^\top t_i'' \ge (1-\alpha) V_i^{\max} \\ & x_i' + x_i'' = x_i \\ & t_i' + t_i'' = t_i \\ & x_i', x_i'' \ge 0 \\ & 1 \ge \alpha \ge 0 \end{aligned} \tag{4.9}$$

Hence, u_i^ξ is also constrained PLC. If it is regular, then it is ξ-robust by Lemma 4.5. By Lemma 2.1, regularity only requires $u_i^\xi(0) = 0$ in this case. This follows by Lemma 4.4(i) and the assumptions $u_i(0) = 0 < V_i^{\max}$. Finally, the statement on bit-complexity follows since $V_i^{\max}$ is the optimum value of a linear program formed by the LP defining u_i and a box constraint. Therefore, $V_i^{\max}$ is polynomially bounded in the input. ∎

THEOREM 4.2. *In a Fisher market with n agents, m goods and regular constrained PLC utility functions and $\sigma < 1$, we can find a σ-approximate market equilibrium by solving $O\left(\left(\frac{n}{\sigma^2}\right)^n\right)$ linear programs, each polynomially bounded in the input size.*

Proof. Let us set $\delta = \sigma^2/(2n)$ and $\xi = \sigma/2$. We first replace the utility functions u_i by u_i^ξ as in Lemma 4.6. Then, we guess all combinations $\tilde{u}_i = k_i\delta$, for $0 \le k_i \le \lceil \frac{1}{\delta} \rceil + 1$. We calculate the allocations x as in (4.4); if no such allocation exists, we proceed to the next guess. If x is feasible to (4.4), then we check if prices p satisfying (4.6), (4.7), and (4.8) exist. Theorem 4.1 guarantees the existence both x and p for at least one choice of the $\tilde{u}_i$'s. This gives a $(n\delta/\xi, 2\delta)$-approximate equilibrium for the utilities u_i^ξ, and by Lemma 4.4(i), a $(n\delta/\xi, 2\delta + \xi)$-approximate equilibrium for the original utilities u_i. By the choice of δ and ξ, this is a σ-approximate equilibrium. ∎

5 PLC Matching Markets

In the Hylland-Zeckhauser matching market equilibrium [26], agents have unit budgets and linear utilities with the additional restriction that every agent has to purchase exactly one unit of good. We now consider the following generalization with PLC utilities for nonnegative values $a_{ij}^l, b_i^l \ge 0$.

$$u_i(x_i) = \begin{cases} \min_l \left\{ \sum_j a_{ij}^l x_{ij} + b_i^l \right\} \,, & \text{if } \sum_j x_{ij} = 1 \,, \\ -\infty & \text{otherwise.} \end{cases} \tag{5.10}$$

Throughout this section, we assume $w_i = 1$ for all agents, as standard in the matching market model. We also assume $n \le m$, i.e., there are at least as many goods as agents that is necessary for feasibility. [6] We refer to this problem as the *PLC matching market* problem. Let $\tilde{V}_i = \max_{x_i \in \mathbb{R}_+^m} u_i(x_i)$ be the maximum achievable utility of agent i; the matching constraint guarantees this is finite. Similarly to (2.1), without loss of generality we can apply affine transformations to the utilities so that

$$\tilde{V}_i \le 1 \quad \text{and} \quad \min_l b_i^l = 0 \quad \text{for each agent } i. \tag{5.11}$$

[6] Here we assume each good has exactly one copy. Note that our method can be generalized to the case that goods have multiple copies. In this case, we assume $n \le \sum_j s_j$, where s_j is the number of copies of good j.

Note that, even though this utility function is constrained PLC, it is *not regular*: $u_i(0) = -\infty$. For this reason, we cannot directly apply the results in Sections 3 and 4. The existence of an equilibrium is also not covered by Theorem 2.1. A key tool to tackle this model is the following price transformation with strong invariance properties that enables us to restrict our attention to (approximate) equilibria where $\min_j p_j = 0$.

LEMMA 5.1. ([38]) *For a PLC matching market model, let $p \in \mathbb{R}^m_+$, and $r > 0$ such that $p'_j = 1 + r(p_j - 1) \geq 0$ for all goods j. Then, $D_i(p,1) = D_i(p',1)$ for every agent i. Consequently, if there exists a market equilibrium $(\{x^*_i\}_i, \{p^*_j\}_j)$, then there exists one with* $\min_j p_j = 0$.

Proof. Both $D_i(p,1)$ and $D_i(p',1)$ only contain bundles x_i with $\sum_j x_{ij} = 1$. Since $1-(p')^\top x_i = r(1-p^\top x_i)$ for such a bundle, the price of a bundle satisfies $p^\top x_i \leq 1$ if and only if it satisfies $(p')^\top x_i \leq 1$. This implies $D_i(p,1) = D_i(p',1)$. For the second part, consider any market equilibrium $(\{x^*_i\}_i, \{p^*_j\}_j)$. If there exists a good at price $p^*_j < 1$, then we can select the largest r value such that this transformation gives $\min_j p'_j = 0$. The first part guarantees that $(\{x^*_i\}_i, \{p'_j\}_j)$ is also a market equilibrium. Otherwise, $p^*_j = 1$ for all j. In this case, setting $p_j = 0$ for all j will also be a matching market equilibrium. ∎

In light of this transformation, we note that the (σ, λ)-approximate (thrifty) equilibrium concept as in Definition 2.3 is unsatisfactory. Assume $n = m$, i.e., the number of goods is the same as the number of agents. Let $(\{x_i\}_i, \{p_j\}_j)$ be a (σ, λ)-approximate equilibrium. Then, for any choice of $0 < \sigma' \leq \sigma$, we can select $r > 0$ such that $(\{x_i\}_i, \{p'_j\}_j)$ will be a (σ', λ)-approximate (thrifty) equilibrium. This is because $(p')^\top x_i$ becomes arbitrarily close to 1, and the third constraint is satisfied since $\sum_i x_{ij} = 1$ for all j follows if $n = m$.

In accordance with Lemma 5.1, we will look for approximate (thrifty) equilibria with the additional requirement that $\min_j p_j = 0$. In Section 5.1, we show that approximate equilibrium results can be obtained by reducing to an associated partial matching market. In Section 5.2, we give a simple counterexample showing that the set of equilibria is non-convex already for the standard matching market model with three agents and three goods.

5.1 From partial to perfect matchings Both for showing the existence of equilibria, as well as for the algorithms, we relax the perfect matching requirement $\sum_j x_{ij} = 1$ to the *partial* matching constraint $\sum_j x_{ij} \leq 1$. That is, for the same parameters a^l_{ij}, b^l_i, we let

$$u'_i(x_i) = \begin{cases} \min_l \left\{ \sum_j a^l_{ij} x_{ij} + b^l_i \right\} , & \text{if } \sum_j x_{ij} \leq 1 \,, \\ -\infty & \text{otherwise.} \end{cases} \tag{5.12}$$

Using the assumption (5.11), $u'_i(0) = 0$, and therefore the u'_i's are regular utilities. For a PLC matching market with utilities u_i as in (5.10), we will refer to the market that replaces the u_i's by the u'_i's as the associated *PLC partial matching market*.

The next two lemmas show the close relationship between equilibria in these markets. In the proofs, we use $V_i(p,1)$ for the optimal utility for u_i and $C_i(p,1)$ the minimum price of an optimal bundle; we let $V'_i(p,1)$ and $C'_i(p,1)$ denote the same for u'_i. Clearly, $V'_i(p,1) \geq V_i(p,1)$.

LEMMA 5.2. *Let $(\{x_i\}_i, \{p_j\}_j)$ be a thrifty PLC matching market equilibrium with $\min_j p_j = 0$. Then, $(\{x_i\}_i, \{p_j\}_j)$ is also a thrifty market equilibrium in the associated PLC partial matching market.*

Proof. Using that $p_k = 0$ for some good k, for every x'_i with $\sum_j x'_{ij} \leq 1$ there exists a bundle $\tilde{x}_i \geq x'_i$ with $\sum_j \tilde{x}_{ij} = 1$ that has the same cost and $u_i(\tilde{x}_i) \geq u'_i(x'_i)$. Consequently, $V'_i(p,1) = V_i(p,1)$, and by the same token, $C_i(p,1) = C'_i(p,1)$. The statement follows. ∎

In the other direction, we show that approximate (thrifty) equilibria in the associated PLC partial matching market have $\min_j p_j = 0$, then this can be extended to the original PLC matching market. This also applies to exact equilibria with $\sigma = \lambda = 0$.

LEMMA 5.3. *For a PLC matching market, consider a (σ, λ)-approximate (thrifty) equilibrium $(\{x'_i\}_i, \{p'_j\}_j)$ in the associated PLC partial matching market, and assume $\min_j p'_j = 0$. Then, in $O(m)$ time we can construct a $(2\sigma, \lambda)$-approximate (thrifty) matching equilibrium $(\{x_i\}_i, \{p'_j\}_j)$ in the original market.*

Proof. Given $(\{x'_i\}_i, \{p'_j\}_j)$, we arbitrarily assign those goods which are not fully allocated to those agents such that $\sum_j x'_{ij} < 1$; this can be easily done in $O(m)$ time (recall $m \geq n$). Let $\{x_i\}_i$ denote the resulting allocations with $\sum_j x_{ij} = 1$.

Recalling that all $w_i = 1$, the approximate equilibrium means $p^\top x_i \leq 1 + n\sigma$ and $\sum_j p'_j(1 - \sum_i x'_{ij}) \leq \sigma \sum_i w_i = n\sigma$. Hence, the spending for each agent (after assignment) can be at most $2n\sigma + 1$. As in the previous proof, $\min_j p_j = 0$ guarantees that $V_i(p, 1) = V'_i(p, 1)$ and $C_i(p, w_i) = C'_i(p, w_i)$. The utility requirement follows since $V_i(p, 1) - \lambda = V'_i(p, 1) - \lambda \leq u'_i(x'_i) \leq u_i(x_i)$. Further, if $(\{x'_i\}_i, \{p'_j\}_j)$ was an approximate thrifty market equilibrium, then thriftiness for $(\{x_i\}_i, \{p_j\}_j)$ follows since the spending can only be increased by $n\sigma$. ∎

We can now derive the existence of an equilibrium, as well as algorithms for approximate equilibria, by making use of the results for PLC partial matchings that are regular utilities.

THEOREM 5.1. *In every PLC matching market, there exists a thrifty market equilibrium* $(\{x_i\}_i, \{p_j\}_j)$ *with* $\min_j p_j = 0$.

Proof. For the u'_i utilities in the associated PLC partial matching market, Theorem 2.1 guarantees the existence of an equilibrium $(\{x'_i\}_i, \{p'_j\}_j)$. If $\min_j p'_j = 0$, then Lemma 5.3 for $\sigma = \lambda = 0$ gives an equilibrium in the PLC matching market with the u_i's. If $\min_j p'_j > 0$, then all goods must be fully sold, hence $\sum_{i,j} x'_{ij} = m \geq n$. This cannot happen if $m > n$; and if $m = n$ this implies that all agents are getting one unit in x', i.e., $\sum_j x'_{ij} = 1$ for all i. Consequently, $(\{x'_i\}_i, \{p'_j\}_j)$ is already an equilibrium in the PLC matching market. By Lemma 5.1, this can be transformed to one with $\min_j p_j = 0$. ∎

For fixed number of goods, we can thus use the algorithm in Section 3 for u'_i, and transform it using Lemma 5.3 for u_i. In order to find an approximate thrifty equilibrium for the u'_i's with $\min_j p_j = 0$; we only enumerate over price combinations where one of the prices is 0. Theorem 5.1 and Lemma 5.2 guarantee the existence of such a solution.

THEOREM 5.2. (THRIFTY PLC MATCHING MARKET EQUILIBRIUM WITH FIXED NUMBER OF GOODS) *Given a PLC matching market with n agents, m goods, and PLC utilities $\{u_i\}_i$, we can compute an ε-approximate thrifty PLC matching market equilibrium by solving $O(n\left(\frac{m}{\varepsilon}\right)^m)$ linear programs, each in polynomial time in the input size.*

Similarly, for fixed number of agents, we can use the results in Section 4 for u'_i in conjunction with Lemma 5.3 to compute an approximate PLC matching market equilibrium (but not necessarily a thrifty one). The only modification needed is that we fix the price of some good to $p_j = 0$; this results in an additional factor m in the running time.

THEOREM 5.3. (PLC MATCHING MARKET EQUILIBRIUM WITH FIXED NUMBER OF AGENTS) *Given a PLC matching market with n agents, m goods, and PLC utilities $\{u_i\}_i$, we can compute a σ-approximate PLC matching market equilibrium by solving $O\left(m\left(\frac{n}{\sigma^2}\right)^n\right)$ linear programs, each in polynomial time in the input size.*

Finally, for the original Hylland-Zeckhauser model with linear utilities, we show that the stronger concept of an approximate *thrifty* equilibrium can also be computed, by exploiting the simpler structure in this case.

THEOREM 5.4. (THRIFTY MATCHING MARKET EQUILIBRIUM WITH FIXED NUMBER OF AGENTS) *Given a matching market with n agents, m goods, and linear utility function*

$$u_i(x_i) = \begin{cases} \sum_j a_{ij}x_{ij}\,, & \textit{if } \sum_j x_{ij} = 1\,, \\ -\infty & \textit{otherwise.} \end{cases}$$

we can compute a σ-approximate thrifty market equilibrium by solving $O\left(m\left(\frac{n}{\sigma^2}\right)^n\right)$ linear programs, each in polynomial time in the input size.

Proof. Similar to the PLC case, we first calculate a thrifty approximate equilibrium for the associated partial matching market such that $\min_j p_j = 0$ and then transform it into a thrifty approximate matching market equilibrium. The transformed ξ-robust utility $u_i'^{\xi}(x_i)$ used in the algorithm (see (4.9)) can be written in the following simpler form. Let $J = \arg\max_j a_{ij}$, and

$$u_i'^{\xi}(x_i) = \begin{cases} \sum_{j \notin J} a_{ij}x_{ij} + \sum_{j \in J}(a_{ij} + \xi)x_{ij}\,, & \text{if } \sum_j x_{ij} \le 1\,, \\ -\infty & \text{otherwise.} \end{cases} \tag{5.13}$$

Let us calculate an approximate thrifty market equilibrium $(\{x_i\}_i, \{p_j\}_j)$ for $u_i'^{\xi}$ as in Section 4.1 with two slight modifications.

We first enumerate all possible $\tilde{u}_i$ for $\tilde{u}_i = \delta k_i$ for $0 \le k_i \le \lceil \frac{1+\xi}{\delta} \rceil + 1$ and one j such that $p_j = 0$; then, we calculate $\{x_i\}_i$ by (4.4); and finally, we calculate the price $\{p_j\}_j$. When calculating the price, in addition to (4.6), (4.7), and (4.8), we add constraints $p_j = 0$ and

$$p^\top x_i \le p_{j'} + \frac{n^2\delta}{\xi} \qquad \forall j' \in J\,. \tag{5.14}$$

Recall that $\frac{n^2\delta}{\xi} = \frac{n\delta}{\xi}\sum_i w_i$ by the assumption that all budgets are 1. This additional last inequality makes the difference compared to the general PLC algorithm. We exploit this in the following claim.

Claim 5.1. *If $p^\top x_i \le p_{j'} + \frac{n\delta}{\xi}\sum_i w_i$ for $j' \in J$; and $p^\top x_i \le 1 + \frac{n^2\delta}{\xi}$, then $p^\top x_i \le C_i(p,1) + \frac{n^2\delta}{\xi}$. Additionally, let $(\{x_i^*\}_i, \{p_j^*\}_j)$ be any thrifty market equilibrium. If $u_i^* - \delta < \tilde{u}_i \le u_i^*$ for all i, then for any $\{x_i\}_i$ such that (4.4) holds, (5.14) is also valid for p^*.*

Proof. The first part follows as $C_i(p,1) = \min\{1, \min_{j\in J} p_j\}$. The second part is true because, for $j' \in J$,

$$\begin{aligned} {p^*}^\top x_i &\le {p^*}^\top x_i^* + \frac{n^2\delta}{\xi} && \text{(by Lemma 4.1)} \\ &\le p_{j'}^* + \frac{n^2\delta}{\xi}. && (\text{as } {p^*}^\top x_i^* = C_i(p^*,1) \le p_j') \end{aligned}$$

□

Combining this observation with Theorem 4.1, Lemma 5.2, and Theorem 5.1, this procedure will output a $(\delta n/\xi, \xi + 2\delta)$-approximate thrifty equilibrium for the associated partial matching market.

Finally, by Lemma 5.3, we construct an approximate matching market equilibrium from the approximate Fisher market equilibrium. The theorem follows by choosing $\delta = \sigma^2/(4n)$ and $\xi = \sigma/2$. □

5.2 Non-convexity example In this section, we give a simple example which shows that the sets of allocations and prices are non-convex. The example consists of three agents, three goods and the utilities are linear for these agents: $u_i(x_i) = \sum_j a_{ij}x_{ij}$. Each agent has a budget of 1 dollar.

	good 1	good 2	good 3
agent 1	1	1	2
agent 2	0	1	2
agent 3	1	1	2

Table 1: Utility matrix (a_{ij})

Given the utility functions, the following prices and allocations are two of the equilibria of the matching market.

The following two lemmas show that neither the set of allocations nor the set of prices is convex.

Lemma 5.4. $\frac{p^{(1)}+p^{(2)}}{2}$ *is not an equilibrium price.*

Proof. Note that $\frac{p^{(1)}+p^{(2)}}{2} = (0, 0.5, 2.5)$. In this case, both agent 1 and agent 3 will not be interested in good 2. This implies agent 2 will get good 2 fully. However, given the price, agent 2 will buy some of good 3, which provides a contradiction. □

	good 1	good 2	good 3
agent 1	0.5	0	0.5
agent 2	0	1	0
agent 3	0.5	0	0.5
price	0	1	2

Table 2: Price 1 ($p^{(1)}$) and Allocation 1 ($x^{(1)}$)

	good 1	good 2	good 3
agent 1	2/3	0	1/3
agent 2	0	2/3	1/3
agent 3	1/3	1/3	1/3
price	0	0	3

Table 3: Price 2 ($p^{(2)}$) and Allocation 2 ($x^{(2)}$)

LEMMA 5.5. $\frac{x^{(1)}+x^{(2)}}{2}$ *is not an equilibrium allocation.*

Proof. Note that in any equilibrium, the price of good 3 should be strictly larger than 1. This implies all agents will spend out all their budgets. Let the price of good 3 be $1+\alpha$ for some $\alpha > 0$. Since agent 1 get 7/12 of good 1 and 5/12 of good 3, the price of good 1 is $1 - \frac{5}{7}\alpha$. Similarly, since agent 2 get 5/6 of good 2 and 1/6 of good 3, the price of good 2 is $1 - \frac{1}{5}\alpha$. Since $\alpha > 0$, given the price $(1 - \frac{5}{7}\alpha, 1 - \frac{1}{5}\alpha, 1 + \alpha)$, agent 3 will not buy good 2, which contradicts allocation $\frac{x^{(1)}+x^{(2)}}{2}$. ∎

References

[1] S. Alaei, P. J. Khalilabadi, and É. Tardos. Computing equilibrium in matching markets. In *Proc. 18th Conf. Economics and Computation (EC)*, pages 245–261, 2017.

[2] X. Bei, J. Garg, and M. Hoefer. Ascending-price algorithms for unknown markets. *ACM Trans. Algorithms*, 15(3):37:1–37:33, 2019.

[3] X. Bei, J. Garg, M. Hoefer, and K. Mehlhorn. Earning and utility limits in Fisher markets. *ACM Trans. Econom. Comput.*, 7(2):10:1–10:35, 2019.

[4] W. Brainard and H. Scarf. How to compute equilibrium prices in 1891. *Cowles Foundation Discussion Paper*, 1270, 2000.

[5] M. Braverman. Optimization-friendly generic mechanisms without money. *CoRR*, abs/2106.07752, 2021.

[6] X. Chen, D. Dai, Y. Du, and S. Teng. Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities. In *Proc. 50th Symp. Foundations of Computer Science (FOCS)*, pages 273–282, 2009.

[7] X. Chen, D. Paparas, and M. Yannakakis. The complexity of non-monotone markets. *J. ACM*, 64(3):20:1–20:56, 2017.

[8] X. Chen and S. Teng. Spending is not easier than trading: On the computational equivalence of Fisher and Arrow-Debreu equilibria. In *Proc. 20th Intl. Symp. Algorithms and Computation (ISAAC)*, pages 647–656, 2009.

[9] B. Codenotti, S. Pemmaraju, and K. Varadarajan. On the polynomial time computation of equilibria for certain exchange economies. In *Proc. 16th Symp. Discrete Algorithms (SODA)*, pages 72–81, 2005.

[10] B. Codenotti, A. Saberi, K. Varadarajan, and Y. Ye. Leontief economies encode two-player zero-sum games. In *Proc. 17th Symp. Discrete Algorithms (SODA)*, pages 659–667, 2006.

[11] R. Cole, N. Devanur, V. Gkatzelis, K. Jain, T. Mai, V. Vazirani, and S. Yazdanbod. Convex program duality, Fisher markets, and Nash social welfare. In *Proc. 18th Conf. Economics and Computation (EC)*, 2017.

[12] P. L. Combettes. Perspective functions: Properties, constructions, and examples. *Set-Valued and Variational Analysis*, 26(2):247–264, 2018.

[13] X. Deng, C. Papadimitriou, and S. Safra. On the complexity of equilibria. In *Proceedings of the 34th annual ACM Symposium on Theory of Computing (STOC)*, pages 67–71, 2002.

[14] N. Devanur and R. Kannan. Market equilibria in polynomial time for fixed number of goods or agents. In *Proc. 49th Symp. Foundations of Computer Science (FOCS)*, pages 45–53, 2008.

[15] N. Devanur, C. Papadimitriou, A. Saberi, and V. Vazirani. Market equilibrium via a primal–dual algorithm for a convex program. *J. ACM*, 55(5), 2008.

[16] R. Duan, J. Garg, and K. Mehlhorn. An improved combinatorial polynomial algorithm for the linear Arrow-Debreu market. In *Proc. 27th Symp. Discrete Algorithms (SODA)*, pages 90–106, 2016.

[17] R. Duan and K. Mehlhorn. A combinatorial polynomial algorithm for the linear Arrow-Debreu market. *Inf. Comput.*, 243:112–132, 2015.

[18] B. C. Eaves. A finite algorithm for the linear exchange model. *J. Math. Econom.*, 3:197–203, 1976.

[19] E. Eisenberg. Aggregation of utility functions. *Management Sci.*, 7(4):337–350, 1961.

[20] J. Garg, E. Husić, and L. A. Végh. Auction algorithms for market equilibrium with weak gross substitute demands and their applications. In *Proc. 38th Symp. Theoret. Aspects of Computer Science (STACS)*, pages 33:1–33:19, 2021.

[21] J. Garg, R. Mehta, and V. V. Vazirani. Dichotomies in equilibrium computation and membership of PLC markets in FIXP. *Theory of Computing*, 12(1):1–25, 2016.

[22] J. Garg, R. Mehta, V. V. Vazirani, and S. Yazdanbod. Settling the complexity of Leontief and PLC exchange markets under exact and approximate equilibria. In *Proc. 49th Symp. Theory of Computing (STOC)*, pages 890–901, 2017.

[23] J. Garg, T. Tröbst, and V. V. Vazirani. An Arrow-Debreu extension of the Hylland-Zeckhauser scheme: Equilibrium existence and algorithms. *CoRR*, abs/2009.10320, 2020.

[24] J. Garg and L. A. Végh. A strongly polynomial algorithm for linear exchange markets. In *Proc. 51st Symp. Theory of Computing (STOC)*, pages 54–65, 2019.

[25] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2. Springer Science & Business Media, 2012.

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

[27] K. Jain. A polynomial time algorithm for computing the Arrow-Debreu market equilibrium for linear utilities. *SIAM J. Comput.*, 37(1):306–318, 2007.

[28] D. Jalota, M. Pavone, Q. Qi, and Y. Ye. Fisher markets with linear constraints: Equilibrium properties and efficient distributed algorithms. *CoRR*, abs/2106.10412, 2021.

[29] S. M. Kakade, M. J. Kearns, and L. E. Ortiz. Graphical economics. In *Learning Theory, 17th Annual Conference on Learning Theory, COLT*, pages 17–32, 2004.

[30] A. Mas-Colell. Equilibrium theory with possibly satiated preferences. In M. Majumdar, editor, *Equilibrium and Dynamics: Essays in Honor of David Gale*. Macmillan Press, 1982.

[31] J. Orlin. Improved algorithms for computing Fisher's market clearing prices. In *Proc. 42nd Symp. Theory of Computing (STOC)*, pages 291–300, 2010.

[32] C. H. Papadimitriou and M. Yannakakis. On the approximability of trade-offs and optimal access of web sources. In *Proceedings 41st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 86–92. IEEE, 2000.

[33] R. T. Rockafellar. *Convex analysis.* Princeton University Press, 1970, 2015.

[34] A. Schrijver. *Theory of linear and integer programming.* John Wiley & Sons, 1998.

[35] H. Varian. Equity, envy and efficiency. *J. Econom. Theory*, 29(2):217–244, 1974.

[36] V. Vazirani and M. Yannakakis. Market equilibrium under separable, piecewise-linear, concave utilities. *J. ACM*, 58(3):10, 2011.

[37] V. V. Vazirani and M. Yannakakis. Computational complexity of the Hylland-Zeckhauser scheme for one-sided matching markets. *CoRR*, abs/2004.01348, 2020.

[38] V. V. Vazirani and M. Yannakakis. Computational complexity of the Hylland-Zeckhauser scheme for one-sided matching markets. In *Proc. 12th Symp. Innovations in Theoret. Computer Science (ITCS)*, 2021. `https://www.ics.uci.edu/~vazirani/VY.pdf`.

[39] L. A. Végh. A strongly polynomial algorithm for a class of minimum-cost flow problems with separable convex objectives. *SIAM J. Comput.*, 45(5):1729–1761, 2016.

[40] L. Walras. *Éléments d'économie politique pure, ou théorie de la richesse sociale (Elements of Pure Economics, or the theory of social wealth)*. Lausanne, Paris, 1874. (1899, 4th ed.; 1926, rev ed., 1954, Engl. transl.).

[41] Y. Ye. A path to the Arrow-Debreu competitive market equilibrium. *Math. Prog.*, 111(1-2):315–348, 2008.