

Finding a Walrasian equilibrium is easy for a fixed number of agents **WORKING PAPER**

Federico Echenique, Division of the Humanities and Social Sciences, California Institute of Technology
Adam Wierman, Computing and Mathematical Sciences, California Institute of Technology

In this work, we study the complexity of finding a Walrasian equilibrium. Our main result gives an algorithm which can compute an approximate Walrasian equilibrium in an exchange economy with general, but well-behaved, utility functions in time that is polynomial in the number of goods when the number of agents is held constant. This result has applications to macroeconomics and finance, where applications of Walrasian equilibrium theory tend to deal with many goods but a fixed number of agents.

1. INTRODUCTION

The problem of computing Walrasian equilibria has been studied continually over many decades. The earliest studies include [Scarf 1977] and [Todd 1976]; but the literature has been reinvigorated by the recent efforts of the computer science community. Some recent surveys include [Vazirani 2007], [Codenotti and Varadarajan 2007], and [Codenotti et al. 2004].

The basic message that has emerged from the literature is negative: computing a Walrasian equilibrium tends to be “hard” in general settings. For example, [Chen et al. 2009a] prove that finding a Walrasian equilibrium is hard (PPAD-complete and hard to approximate) even in economies with piece-wise linear concave utilities that are separable by goods. Similarly, [Codenotti et al. 2006] prove that the problem is also hard in economies in which agents have Leontief preferences.¹

In contrast, for more limited settings, there exist positive results. In particular, there are studies presenting computationally tractable instances of Walrasian equilibrium, but they require very special assumptions on the markets in question. For example, the results assume that the economies satisfy the gross substitutes condition, or that all utilities are linear, or of the CES functional form (with certain values of the CES parameter) [Devanur et al. 2002; Codenotti et al. 2005; Jain 2004]. Additionally, [Segal 2007] presents positive results on the communication complexity of Walrasian equilibria. Other results focus on economies with a representative agent, through the device of a Fisher equilibrium [Jain and Vazirani 2007]. In general, positive results such as described above are deep and ingenious, but require assumptions that leave out most economic applications of general equilibrium theory.

¹We emphasize that [Codenotti et al. 2006] reduce arbitrary two-player games, in which one player has n strategies and the other has m strategies, into an economy with $n + m$ agents and goods. Thus they consider problems in which both goods and agents grow at the same rate.

This work is supported by the National Science Foundation, under grant CCF-1101470.

We thank Chris Umans, Jeremy Hurwitz, Daniel Golovin, and Ilya Segal for their advice and comments on this work.

Permission to make digital or hard copies of part or all 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 show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2012 ACM 0000-0000/2012/02-ART13 \$10.00

DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

As is evident from the above, there is a gap between the generality of the hardness results and the specificity of the instances for which finding an equilibrium is tractable. The goal of the current work is to find a middle ground where computing a Walrasian equilibrium is tractable, but which can still capture settings that include economic applications of the theory.

To that end, our focus is on the setting of many goods but a fixed number of agents. This restriction simplifies the computational problem dramatically, but still captures a large majority of the application of general equilibrium theory in economics. Specifically, most modern applications of general equilibrium theory are in the areas of macroeconomics and finance; in turn, most modern macro and finance deals with models of general equilibrium. In these settings the theory is most often applied to large economies, in the sense of having infinitely many goods, but the number of agents remains small and fixed. This is because in macroeconomic and financial applications the time horizon is usually infinite, which implies that there are infinitely many different goods. Additionally, there is often uncertainty, which gives rise to an infinite dimensional commodity space. On the other hand, the models typically assume finitely many (long-lived) agents or types of agents. Importantly, the complexity of finding an equilibrium in such settings is highly relevant, because researchers in macro and finance are heavy users of computational methods. They very often need (and use) algorithms to find equilibria.

Our main result (Theorem 2.1) exhibits an algorithm for finding an approximate Walrasian equilibrium that runs in polynomial-time in the number of goods when the number of agents is fixed. Importantly, our result applies to general, but well-behaved, utility functions (see Section 2) and applies to some settings where there are many agents but where the number of “types” of agents is fixed (see Section 5). Situations with this form of limited heterogeneity are of particular interest in economics. More specifically, our results apply directly to (i) “replica economies” where there are many agents, but each agent is a copy of some prototypical (small) set of types of agents; (ii) economies where the behavior of many agents can be aggregated into the behavior of a single “representative consumer”; and (iii) economies with “endowment classes,” where agents are fully heterogeneous in their preferences but not in their sources of income.

We obtain results on Fisher equilibria as a special case of our approach. Representative consumers will only exist under very stringent settings, but one (salient) such setting coincides with the model of Fisher equilibria. Fisher equilibria have received a lot of attention in the literature on the complexity of economic equilibria, and some of the most positive recent results on the complexity of economic models have been related to Fisher equilibria, e.g., [Jain and Vazirani 2007]. In Section 5, we highlight that our main result implies an efficient algorithm for computing approximate Fisher equilibria as well.

Note that we are not the first to consider the problem of finding Walrasian equilibrium when the number of agents is fixed. The closest papers in the literature to ours are [Deng et al. 2002] and [Devanur and Kannan 2008], who show that the problem of finding Walrasian equilibrium is easy when the number of agents is bounded and utilities are, respectively, linear or piecewise linear and separable by goods. Separability by goods means that marginal utilities can be treated independently; and represents a key simplification that is exploited by the algorithms in these papers. Our approach is different, it uses a combinatorial version of the Negishi approach for proving existence of Walrasian equilibria [Negishi 1960], see Section 3 for an overview. As a consequence, our result allows for general concave and non-separable utilities, under an additional technical condition that rules out boundary solutions to the consumers’ maximization problem (see the discussion in Section 2).

2. MAIN RESULT

In this paper, we study the standard model of an exchange economy with n agents and l goods. In such a model, agents are endowed with non-negative quantities of each good; they derive income from selling their endowments at the prevailing prices; and use the income to purchase a consumption bundle. In a model of an exchange economy, all economic activity reduces to pure barter, and there is no production of new goods. It should be noted, however, that our results extend naturally to economies with production. We would simply need to assume that firms' production technologies are convex.

Before discussing our result, we introduce the basic notation and definitions for the model.

2.1. Basic notation

We adopt the following notational conventions: \mathbf{R}_+^m denotes the positive orthant of the m -dimensional Euclidean space; \mathbf{R}_{++}^m is the set of vectors of \mathbf{R}^m that are strictly positive in all its components. We use the norm defined as $\|x\| = \sup_i |x_i|$, and understand the distance between two vectors x and y to be $\|x - y\|$.

A function $u : \mathbf{R}_+^m \rightarrow \mathbf{R}$ is C^1 if there is an open set $U \supseteq \mathbf{R}_+^m$ and a function $f : U \rightarrow \mathbf{R}$ that is differentiable and has a continuous derivative, such that f and u coincide on \mathbf{R}_+^m . The following definitions will be useful: Say that u is

- (1) *monotonic* if $x \leq y$ and $x \neq y$ imply that $u(x) < u(y)$;
- (2) *quasiconcave* if for any $x, y \in \mathbf{R}_+^m$, $x \neq y$ and $\alpha \in (0, 1)$, $u(x) \leq u(y) \Rightarrow u(x) \leq u(\alpha x + (1 - \alpha)y)$;
- (3) *concave* if for any $x, y \in \mathbf{R}_+^m$, $x \neq y$ and $\alpha \in (0, 1)$, $\alpha u(x) + (1 - \alpha)u(y) \leq u(\alpha x + (1 - \alpha)y)$;
- (4) *homothetic* if for any $x, y \in \mathbf{R}_+^l$, $u(y) = u(x)$ implies that $u(\alpha y) = u(\alpha x)$, for all scalar $\alpha \in \mathbf{R}_+$.

2.2. Exchange economies

Before presenting the result, we define exchange economies and describe the assumptions that we impose in our main result.

We denote an *exchange economy* as a tuple $(\omega_i, u_i)_{i=1}^n$ where $\omega_i \in \mathbf{R}_+^l$ and $u_i : \mathbf{R}_+^l \rightarrow \mathbf{R}$. The number l is the number of goods in the economy. The number of agents is n , and each one is characterized by two objects: a vector of *endowments* ω_i , and a *utility function* u_i . Further, an *allocation* in $(\omega_i, u_i)_{i=1}^n$ is a vector $x \in \mathbf{R}_+^{nl}$ for which $\sum_{i=1}^n x_i = \sum_{i=1}^n \omega_i$.

We denote the family of exchange economies that we study by E , and define it to satisfy the following. Each economy $(u_i, \omega_i)_{i=1}^n$ in E has the same number of agents, n . They may differ in the number of goods. We assume that all economies in E satisfy the following assumptions:

- (1) $\sum_{i=1}^n \omega_i \in \mathbf{R}_{++}^l$, i.e., existence;
- (2) u_i is C^1 , concave, and strictly monotonic, i.e., utilities are regular;
- (3) If $x \in \mathbf{R}_+^l$, $y \in \mathbf{R}_{++}^l$, and for some good s , $x_s = 0$, then $u(x) < u(y)$, i.e., an Inada condition;
- (4) for all $x \in \mathbf{R}_+^{nl}$ such that $\sum_{i=1}^n \omega_i = \sum_{i=1}^n x_i$, $u_i(x) \in [0, 1]$, i.e., a normalization condition.

Our assumptions on E deserve some discussion. The assumptions placed on utilities are ubiquitous in economic models. Many general equilibrium models assume that utilities are regular, and a boundary condition equivalent to our "Inada" condition: the role of this is assumption is to rule out that agents' optimal consumption choice has

zero consumption of some goods. The normalization condition puts a uniform bound on utilities evaluated within an allocation. Note that this bound is not global, it only holds over the bundles achieved in some allocation.

Importantly, concavity is stronger than the assumption of quasiconcavity in textbook treatments of general equilibrium theory, but is still commonly assumed in applications. Concavity is a requirement of our analysis in two ways: First, concavity is required by the Negishi approach for proving existence, which is the basis of our analysis. Second, we also use concavity to bound the degree of approximation in our algorithm away from the number of goods. It is fair to say that our assumptions coincide with the hypotheses of the Second Welfare Theorem (a crucial ingredient for us), with a strengthening of quasiconcavity to concavity of utilities.

2.3. Exact and approximate Walrasian equilibria

Given the notion of an exchange economy, we can now define a *Walrasian equilibrium* in $(\omega_i, u_i)_{i=1}^n$ as a pair (p, x) where

- (1) $p \in \mathbf{R}_{++}^l$, i.e., p a price vector,
- (2) $x = (x_i)_{i=1}^n \in \mathbf{R}_+^l$ is an allocation, i.e., supply equals demand,
- (3) for all i , $p \cdot \omega_i = p \cdot x_i$, and $u_i(y) > u_i(x_i) \Rightarrow p \cdot y > p \cdot x_i$, i.e., agents maximize utility.

In this paper, we are concerned with computing approximate Walrasian equilibrium. There are multiple versions of approximate equilibrium that have been studied in the literature, and the notion we define below is somewhat different than all of them, but is particularly natural for our proof approach. However, the particular notion of approximate equilibria considered is not crucial to our main result, which can be extended to other notions of approximate equilibria as discussed in Section 4.

We define an approximate equilibrium as follows. Informally, an approximate equilibrium consists of a price and an allocation in which agents are utility maximizing, supply equals demand, and agents' expenditures are approximately equal their incomes. Formally, a *Walrasian ε -equilibrium* is a pair (p, x) where $p \in \mathbf{R}_+^l$, x is an allocation, and for all i , $u_i(y) > u_i(x_i) \Rightarrow p \cdot y > p \cdot x_i$ and $|p \cdot \omega_i - p \cdot x_i| < \varepsilon$.

2.4. Main result

Given the definitions of the class of exchange economies E and the notion of approximate equilibrium above, we are now ready to state our main result. The main result of this paper is to exhibit an algorithm that computes an approximation of a Walrasian equilibrium in time that is polynomial in the number of goods when the number of agents is fixed.

THEOREM 2.1. *Let $\varepsilon > 0$. There is an algorithm that, for any economy in E , finds a Walrasian ε -equilibrium in time polynomial in l when n is fixed.*

The specifics of the algorithm are given in Section 3, where this result is proven. Following the proof we discuss extensions of the result to other notions of approximate equilibria (Section 4) and applications of the result to replica economies, endowment classes, representative consumers, and Fisher equilibria (Section 5).

3. A PROOF OF THEOREM 2.1

The key idea in our proof of Theorem 2.1 is to make use of the Negishi approach for proving existence of Walrasian equilibrium [Negishi 1960]. As a result, we start by giving a brief overview of Negishi's approach and then discuss the challenges for using it to define an efficient algorithm for finding an equilibrium. The algorithm is then stated in Section 3.4

3.1. An overview of Negishi's approach

Negishi's approach [Negishi 1960] for proving equilibrium existence consists of exploiting the second welfare theorem, which (informally) states that all Pareto efficient outcomes can be achieved via wealth redistributions, a.k.a., transfers, [MasColell et al. 1995]. These transfers are key to Negishi's approach.

In particular, let us start by defining a *Walrasian equilibrium with transfers* as a triple (p, x, T) where:

- (1) $p \in \mathbf{R}_+^l$, i.e., p is a price vector,
- (2) $T \in \mathbf{R}^n$ and $\sum_{i=1}^n T_i = 0$, i.e., a transfer vector,
- (3) x is an allocation, i.e., supply equals demand,
- (4) for all i , $u_i(y) > u_i(x_i) \Rightarrow p \cdot y > p \cdot \omega_i + T_i$, and $p \cdot x_i = p \cdot \omega_i + T_i$, i.e., agents maximize utility.

Note that a Walrasian equilibrium without transfers is just a Walrasian equilibrium with transfers (p, x, T) in which $T = 0$; and that a Walrasian ε -equilibrium is one where $\|T\| \leq \varepsilon$, where $\|x\| = \sum_i |x_i|$.

Now, using this definition, the second welfare theorem implies that for every vector of weights $(\lambda_1, \dots, \lambda_n)$, where λ_i is the weight on agent i , there exists a Walrasian equilibrium with transfers (p, x, T) in which x maximizes $\sum_i \lambda_i u_i(x_i)$. Note that the *welfare weights* λ are parameterizing the set of Pareto efficient allocations.

Without going into details, at this point one can define a mapping from λ into a vector of welfare weights that tries to “undo” the compensations introduced by the transfers T : these welfare weights seek to reward agents who have transfers $T_i < 0$ and punish agents with transfers $T_i > 0$. The details of this mapping are not crucial, but the point is that any fixed point of this mapping will necessarily have $T = 0$ since, any λ associated with positive or negative transfers would be “moved” by the mapping. Negishi exhibits such a mapping and proves that it has a fixed point, using Kakutani's fixed point theorem; thus establishing the existence of a Walrasian equilibrium with zero transfers, which is a standard Walrasian equilibrium.

From our perspective, the crucial feature of Negishi's approach is that the fixed point argument is done in a space of dimension n , not l , which is important since we are treating n as fixed. Note that this same feature has been crucial in the study of Walrasian equilibrium with infinite dimensional commodity spaces. See, for example, [Bewley 1991], [Magill 1981] or [Mas-Colell 1986].

3.2. Adapting Negishi's approach

In order to prove Theorem 2.1, we adopt the main structure of Negishi's approach in order to arrive at a fixed point problem in an n dimensional space rather than working in an l dimensional space, as is common. However, there are several important complications that we need to handle.

First, Kakutani's fixed point theorem is not constructive. Instead we base our algorithm on Sperner's lemma, the combinatorial underpinning of Kakutani's (and Brouwer's) theorem. In particular, we look directly for zero T , and do not explicitly use a fixed-point argument. We could have written a somewhat simpler version of our algorithm that relies on a fixed point theorem; however we have instead based the algorithm on Sperner's lemma because it is cleaner and more attractive computationally.

Second, and more importantly, as the number of goods changes, the mapping from welfare weights to transfers can change in ways that are difficult to control. Our algorithm needs to be robust to such changes: this is perhaps the main difficulty in applying Negishi's approach as a computational device. To handle this issue we exploit the concavity of utilities to obtain a Lipschitz bound on the mapping from welfare weights

to transfers. The bound allows us to approximate zero transfers in a way that is independent of the number of goods in the economy. The number of goods only enters the problem when one obtains a Walrasian equilibrium with transfers for given welfare weights, and we only perform this computation a fixed number of times.

We address each of these issues in detail in the following.

3.3. An overview of Sperner's lemma

As mentioned above, an important piece of our algorithm and proof is Sperner's lemma. To provide the terminology and background for what follows we include a brief overview here. The treatment is based on that of [Border 1989].

We say that a collection of vectors x_0, \dots, x_m in \mathbf{R}^n is *affinely independent* if $\sum_{i=0}^m \theta_i x_i = 0$ and $\sum_{i=0}^m \theta_i = 0$ implies that $\theta_0 = \theta_1 = \dots = \theta_m = 0$.

A *m-simplex* is the set of all strictly positive convex combinations of an affinely independent set of $m + 1$ vectors. A *closed m-simplex* is the convex hull of an affinely independent set of $m + 1$ vectors. Given the affinely independent vectors x_0, \dots, x_m , the simplex $\Delta(x_0, \dots, x_m)$ is the set

$$\Delta(x_0, \dots, x_m) = \left\{ \sum_{i=0}^m \theta_i x_i : \theta_i > 0, i = 0, \dots, m; \sum_{i=0}^m \theta_i = 1 \right\}.$$

Here, each x_i is a *vertex* of $\Delta(x_0, \dots, x_m)$ and each k -simplex $\Delta(x_{i_0}, \dots, x_{i_k})$ is a *face* of $\Delta(x_0, \dots, x_m)$. The *diameter* of a simplex is the largest distance between any two of its vertexes.

For each $y = \sum_{i=0}^m \theta_i x_i$ in the closure of $\Delta(x_0, \dots, x_m)$, let $\chi(y) = \{i : \theta_i > 0\}$. Note that if $\chi(y) = \{i_0, \dots, i_k\}$ then $y \in \Delta(x_{i_0}, \dots, x_{i_k})$.

Denote by e_i the vector in \mathbf{R}^n which has all its coordinates 0, except for a 1 in its i th coordinate. The *standard n-simplex* is the simplex $\Delta(e_1, \dots, e_n)$, denoted simply as Δ . Note that

$$\bar{\Delta} = \{x \in \mathbf{R}^n : x \geq 0 \wedge \sum_i x_i = 1\}.$$

A *simplicial subdivision* of $\bar{\Delta}$ is a collection A_1, \dots, A_J of simplexes such that $\bar{\Delta} = \cup_{j=1}^J \bar{A}_j$ and for each $j \neq h$ $\bar{A}_j \cap \bar{A}_h$ is either empty or the closure of a common face. The *mesh* of a simplicial subdivision is the largest diameter of any its simplexes.

Fix a simplicial subdivision A_1, \dots, A_J of Δ . Let V denote the collection of all the vertices of A_1, \dots, A_J . A function $f : V \rightarrow \{1, \dots, n\}$ for which $f(v) \in \chi(v)$ for all $v \in V$ is called a *proper labeling* of the simplicial subdivision.

LEMMA 3.1 (SPERNER'S LEMMA). *Let A_1, \dots, A_J be a simplicial subdivision of $\bar{\Delta}$, and f a proper labeling of this subdivision. Then there is (an odd number of) A_j such that f achieves all the values $\{1, \dots, n\}$ on the vertices of A_j .*

See [Border 1989] for a proof of Sperner's lemma. A simplex A_j for which f achieves all the values $\{1, \dots, n\}$ is called *completely labeled*.

3.4. An algorithm

We are now ready to state our algorithm for finding an approximate Walrasian equilibrium. Note that throughout the following we assume that for any concave maximization problem in \mathbf{R}^l there is an algorithm that finds an exact solution in time polynomial in l . Thus, we are in effect reducing the calculation of an approximate Walrasian equilibrium to a polynomial number of concave maximizations. Of course, in reality, these maximizations must be solved approximately; however we ignore this for the exposition in order to improve its clarity. However, it is interesting to observe that

even under under the assumption that these maximizations are solved exactly, only an approximate Walrasian equilibrium is found.

To begin the discussion of the algorithm, let $\varepsilon > 0$ and define the norm $\|x\| = \sup_i |x_i|$.

As in Negishi's approach, we work with welfare weights. Specifically, for each $\lambda \in \bar{\Delta}$, consider the problem

$$\begin{aligned} \Pi(\lambda) : \max & \sum_{i=1}^n \lambda_i u_i(x_i) \\ \text{s.t.} & \begin{cases} \sum_{i=1}^n x_i \leq \sum_{i=1}^n \omega_i, \\ x_i \geq 0. \end{cases} \end{aligned}$$

Let $x(\lambda)$ be a solution to the above problem.

Since $\lambda \in \bar{\Delta}$ there is at least one h with $\lambda_h > 0$. By the "Inada" (boundary) condition, we have that $x_h(\lambda) \in \mathbf{R}_{++}^l$. Let

$$p(\lambda) = \lambda_h D u_h(x_h(\lambda)).$$

Note that the first-order conditions of problem $\Pi(\lambda)$ imply that the definition of $p(\lambda)$ does not depend on the chosen consumer h with $\lambda_h > 0$.

Define, for all i ,

$$T_i(\lambda) = p(\lambda) \cdot (x_i(\lambda) - \omega_i).$$

and let $g : \bar{\Delta} \rightarrow \mathbf{R}^n$ be defined by $g(\lambda)_i = T_i(\lambda)$.

Observe that $(x(\lambda), p(\lambda), T(\lambda))$ is a Walrasian equilibrium with transfers, as $p(\lambda)$ is chosen to satisfy the first-order conditions of all agents with $\lambda_i > 0$, and agents with $\lambda_i = 0$ are maximizing utility trivially.

Now, the algorithm works as follows:

- (1) Construct a simplicial subdivision S of $\bar{\Delta}$ of mesh-size $\frac{\varepsilon}{(M-m)(n-1)^d}$, for parameters M and m that depend on the utilities and endowments (see equations (1) and (2)).
- (2) For each vertex λ of S calculate $g(\lambda)$. If there is a vertex for which $g(\lambda) \leq 0$, then we have found λ such that $T_i(\lambda) \leq 0$ and, thus, $\sum_{i=1}^n T_i(\lambda) = 0$ implies that $T = 0$. So, we have found a Walrasian equilibrium. If there is no such vertex, then for every vertex λ of the simplicial subdivision, there is an agent i with $g(\lambda)_i > 0$. In this case, define a labeling of the subsimplex as follows. Let the label of a vertex λ be the i for which the transfer to i in $g(\lambda)$ is largest; i.e. i is the largest component of $g(\lambda)$ (note it is strictly positive). If there is more than one, choose the smallest such i . We prove below that this yields a proper labeling.
- (3) Find a completely labeled subsimplex, say $\lambda^1, \dots, \lambda^n$, where λ^i is labeled i . One is guaranteed to exist by Sperner's lemma and it can be found either by exhaustive search or by some version of Scarf's algorithm. Let $\eta^i = g(\lambda^i)$ and report η^1 as a Walrasian ε -equilibrium.

As mentioned earlier, it is possible to "simplify" the algorithm by eliminating the application of Sperner's lemma through the use of the existence of a fixed point of some (smoothed) version of Negishi's mapping. With a fine enough mesh, some vertex of the simplicial subdivision will be close enough to the actual fixed point. Our proof (specifically Lemma 3.2) could then be used imply that we obtain a Walrasian ε -equilibrium. However, we have instead based the algorithm on Sperner's lemma because we think it is cleaner and more attractive computationally.

3.5. Three lemmas

We now prove that the algorithm defined above finds an approximate Walrasian equilibrium in time that is polynomial in l when n is fixed. The proof will follow from the application of three lemmas, one for each step of the algorithm.

To begin, we define two terms that are important in determining the appropriate mesh-size for the simplicial subdivision of Δ . Define

$$1 \geq M \geq \sup_i \left\{ u_i(x_i) : (x_1, \dots, x_n) \geq 0 \wedge \sum_j x_j \leq \sum_j \omega_j \right\}, \quad (1)$$

and

$$0 \leq m \leq \inf_i \left\{ u_i(x_i) : (x_1, \dots, x_n) \geq 0 \wedge \sum_j x_j \leq \sum_j \omega_j \right\}. \quad (2)$$

Note that M and m are well-defined because utility functions are continuous and the set of allocations is compact: they are bounded by 1 and 0 because of our normalization assumption.

The key challenge in defining the mesh size is to define it in a manner that is independent of l , the number of goods, since our algorithm needs to be robust in this dimension. To handle this issue we exploit the concavity of utilities to obtain a Lipschitz bound on the mapping from welfare weights to transfers. In particular, the lemma below implies that g is Lipschitz continuous with constant $(n-1)$, so it does not change with l .

LEMMA 3.2. *For $\lambda, \lambda' \in \Delta$, $\|g(\lambda) - g(\lambda')\| \leq (n-1)\|\lambda - \lambda'\|$.*

PROOF. The concavity of u_i implies that $Du_i(x)(\omega - x) \geq u(\omega) - u(x)$ ([Rockafellar 1970] Theorem 23.2), so

$$\begin{aligned} T_i(\lambda) &= p(\lambda) \cdot (x_i(\lambda) - \omega_i) \\ &= \lambda_i Du_i(x_i(\lambda))(x_i(\lambda) - \omega) \\ &\leq \lambda_i (u_i(x_i(\lambda)) - u_i(\omega_i)) \leq \lambda_i (M - m). \end{aligned}$$

On the other hand, $\sum_j T_j(\lambda) = 0$ so

$$T_i(\lambda) = - \sum_{j \neq i} T_j(\lambda) \geq - \sum_{j \neq i} (M - m) \lambda_j \geq -(n-1)(M - m) \|\lambda\|.$$

Thus, $|T_i(\lambda)| \leq (n-1)(M - m) \|\lambda\|$. \square

The key consequence of Lemma 3.2 is that it guarantees that the mesh size used in step (1) of the algorithm is small enough. In particular, the size of the mesh guarantees that if λ and λ' are vertexes in the same subsimplex, then $\|\lambda - \lambda'\| < \frac{\varepsilon}{(M-m)(n-1)^3}$. It then follows from Lemma 3.2 that

$$\|g(\lambda) - g(\lambda')\| < \frac{\varepsilon}{(n-1)^2}.$$

Next, we need to verify that the labeling performed in step (2) of the algorithm is a proper labeling, thus allowing the application of Sperner's lemma.

LEMMA 3.3. *If $\lambda_i = 0$ then $g(\lambda)_i \leq 0$.*

PROOF. If $x(\lambda)$ is a solution to $\Pi(\lambda)$ then $x_i(\lambda) = 0$, as utility functions are strictly monotonic. Then $T_i(\lambda) = -p(\lambda) \cdot \omega_i \leq 0$. \square

Lemma 3.3 ensures that the labeling performed in step (2) is indeed a proper labeling because if $i \notin \chi(\lambda)$ then $\lambda_i = 0$, so $g(\lambda)_i \leq 0$.

Finally, we need to ensure that the completely labeled subsimplex found in step (3) of the algorithm does indeed yield a Walrasian ε -equilibrium. Note that, by construction,

$$\|\eta^i - \eta^j\| \leq \frac{\varepsilon}{(n-1)^2}.$$

From this, fact, we can prove the following lemma, which completes the proof of correctness for the algorithm by guaranteeing that we have found a Walrasian ε -equilibrium.

LEMMA 3.4. $\|\eta^i\| < \varepsilon$.

PROOF. We shall prove that $\eta_i^i \leq \varepsilon/(n-1)$. This suffices to prove the lemma because if $\eta_j^i > 0$ then $\eta_j^i \leq \eta_i^i \leq \varepsilon$; and if $\eta_j^i < 0$ then $\sum_h \eta_h^i = 0$ implies that

$$\eta_j^i \geq - \sum_{h:\eta_h^i > 0} \eta_h^i \geq -(n-1)\eta_i^i \geq -\varepsilon,$$

as η_i^i is the largest value of a component of η^i ; so $|\eta_j^i| \leq \varepsilon$.

Suppose then, towards a contradiction, that there is i with $\eta_i^i > \varepsilon/(n-1)$. Since $\sum_j \eta_j^i = 0$, there is j with $\eta_j^i < -\varepsilon/(n-1)^2$. Then,

$$\|\eta^j - \eta^i\| \geq |\eta_j^j - \eta_j^i| \geq |\eta_j^i| > \frac{\varepsilon}{(n-1)^2}$$

where the first inequality is by definition of $\|\cdot\|$, the second because $\eta_j^j > 0$ and $\eta_j^i < 0$. But $\|\eta^j - \eta^i\| > \frac{\varepsilon}{(n-1)^2}$ contradicts the construction of the subsimplex and Lemma 3.2. \square

The last detail is to quantify the running time of the algorithm. As mentioned, our algorithm can be viewed as a reduction of the calculation of an approximate Walrasian equilibrium to a polynomial number of concave maximizations. In particular, the number of vertices in the mesh is bounded by $\theta((n^4/\varepsilon)^n)$. For each vertex one must solve a l -dimensional concave maximization. These can each be solved in time polynomial in l . Further, for each vertex there are $O((n^4/\varepsilon)^n)$ subsimplices containing that vertex. Therefore, a brute-force search in step (3) of the algorithm may require $\Theta((n^4/\varepsilon)^{2n})$ steps. Given that we are concerned with the situation where n is fixed, the running time of the algorithm is clearly polynomial in l . In fact, the algorithm is fixed parameter tractable with respect to n . *Further, the running time will still be polynomial in l when $n = O(\log l)$.*

4. EXTENSIONS TO OTHER NOTIONS OF APPROXIMATE EQUILIBRIA

Theorem 2.1 is proven for a particular notion of approximate Walrasian equilibrium, defined in Section 2.3. But, it is important to note that many varieties of approximate Walrasian equilibria have been studied before. The recent literature in computer science focuses, understandably, on approximate equilibria since exact equilibria can only have approximate discrete analogues, e.g., [Deng et al. 2002]. The older literature on general equilibrium theory in economics also looks at approximations of equilibria, for completely different reasons, e.g., [Starr 1969]. Our notion of Walrasian ε -equilibrium is somewhat different from the notions studied before and so it is important to relate it to those notions.

In this section, we show how, with a small modification, our algorithm finds approximate equilibria that are similar to the ones studied in the both the computer science literature, as well as in the general equilibrium literature from the 60s and 70s.

Specifically we consider two other notions of approximate equilibria, termed ε -approximate equilibrium and strong ε -approximate equilibrium.

In the case of ε -approximate equilibrium, the idea is to relax how consumers optimize. Specifically, the definition imposes ε -maximization of utility: An ε -approximate equilibrium in an exchange economy $(u_i, \omega_i)_{i=1}^n$ is a pair (p, x) where $p \in \mathbf{R}_+^l$, x is an allocation, and for all i

$$p \cdot y \leq p \cdot \omega_i \Rightarrow u_i(y) \leq u_i(x_i) + \varepsilon,$$

and $|p \cdot \omega_i - p \cdot x_i| < \varepsilon$. The notion of ε -approximate equilibrium is close to the one studied in [Deng et al. 2002]. They require that the consumers are ε -maximizing utility when consuming the bundles mandated by the equilibrium, but they assume that demand is only approximately equal to supply. In our definition, demand is exactly equal to supply. On the other hand, in our definition consumers are only approximately spending their incomes.²

To define strong ε -approximate equilibria, we instead relax the requirement that demand equals supply. This yields a notion that is similar in spirit (but stronger) than the one in [Starr 1969]. Specifically, a strong ε -approximate equilibrium in an exchange economy $(u_i, \omega_i)_{i=1}^n$ is a pair (p, x) where $p \in \mathbf{R}_+^l$, $x \in \mathbf{R}_+^{nl}$ with $\|\sum_i x_i - \sum_i \omega_i\| < \varepsilon$, and for all i

$$p \cdot y \leq p \cdot \omega_i \Rightarrow u_i(y) \leq u_i(x_i),$$

and $p \cdot \omega_i = p \cdot x_i$.

In order to extend Theorem 2.1 to provide to these alternative notions of approximate equilibria we need to impose an additional assumption on the economies we consider. In particular, define \tilde{E} to be the set of exchange economies that satisfy the conditions defining E and additionally satisfy the following. Suppose that there is a $\Theta > 0$ and a $\pi > 0$ such that, for all $(u_i, \omega_i)_{i=1}^n$ in \tilde{E} ,

$$\sup_{p \in \Delta} \left(p \cdot \sum_{i=1}^n \omega_i \right) \leq \Theta,$$

and if x is an allocation in $(u_i, \omega_i)_{i=1}^n$, then $D_s u_i(x_i) > \pi$.

Note that the first component of the additional assumption simply rules out arbitrarily large endowments. The number of different goods in the economies of \tilde{E} may grow, but the total “mass” in the economy must remain bounded. When endowments are bounded, it is easy to see that marginal utilities must be bounded below (using condition (3) in the definition of E). We require that the bound, π , be uniform across the economies in \tilde{E} . Additionally, the role of the bound Θ is to control how small the welfare weights in an economy may be (see Lemma 4.2 in the proof of Theorem 4.1). Using both bounds, Θ and π , we can also control how small prices can be in equilibrium.

Now, in the context of \tilde{E} we have the following extension of Theorem 2.1.

THEOREM 4.1. *Let $\varepsilon > 0$. There is an algorithm that, for any economy in \tilde{E} , finds an ε -approximate equilibrium, and a strong ε -approximate equilibrium, in time polynomial in l when n is fixed.*

In the remainder of this section we prove Theorem 4.1.

²Approximate budget exhaustion is assumed in part of the literature on core convergence, see for example [Anderson 1978].

Proof of Theorem 4.1

The following lemma provides the basis for proof of both statements in the theorem. It allows the algorithm presented in Section 3 to run on a subsimplex in which welfare weights are bounded away from zero. Specifically, we use the same algorithm as described in the Section 3, but modified to run on the simplex $\overline{\Delta}^\eta$, where $\Delta^\eta = \{\lambda \in \Delta : \lambda_i > \eta\}$.

LEMMA 4.2. *Suppose that E satisfies the assumption of bounded endowments. Let $\eta > 0$ be such that $-\Theta + (n-1)\eta < 0$. Then for all $\lambda \in \Delta$, $\lambda_i \leq \eta$ implies that $g(\lambda)_i < 0$.*

PROOF. Let $\lambda \in \Delta \setminus \Delta^\eta$. Let i be such that $\lambda_i \leq \eta$. Consider $\hat{\lambda}$, defined by

$$\hat{\lambda}_j = \begin{cases} \lambda_j + \lambda_i/(n-1) & \text{if } j \neq i \\ 0 & \text{if } j = i. \end{cases}$$

Then (by the same argument as in Lemma 3.3), $g(\hat{\lambda})_i = -p(\hat{\lambda})\omega_i$.

Now, $\|\hat{\lambda} - \lambda\| = \lambda_i \leq \eta$, so Lemma 3.2 implies that $\|g(\hat{\lambda}) - g(\lambda)\| \leq (n-1)\eta$. Then $|g(\hat{\lambda})_i - g(\lambda)_i| \leq (n-1)\eta$; so $g(\hat{\lambda})_i = -p(\hat{\lambda})\omega_i$ implies that

$$g(\lambda)_i \leq -p(\hat{\lambda})\omega_i + (n-1)\eta \leq -\Theta + (n-1)\eta < 0$$

□

Note that the previous lemma guarantees that for any boundary λ , $\lambda \in \overline{\Delta}^\eta \setminus \Delta^\eta$, for all i with $\lambda_i = \eta$ we have that $g(\lambda)_i < 0$. Therefore, the labeling described in the algorithm is a proper labeling.

Case 1: ε -approximate equilibrium.

Let $\varepsilon' > 0$ be such that $\varepsilon' < \varepsilon\eta$. Our algorithm outputs a Walrasian ε' -equilibrium (x, p) . We shall prove that (x, p) is an ε -approximate equilibrium.

Observe that η , and therefore ε' , depends only on $(n-1)$ and Θ . It is therefore constant across the economies in E . Note that since $\lambda_i \geq \eta > 0$ for all i we have that $x_i \in \mathbf{R}_{++}^l$. Let x_i^* be a solution to

$$\begin{aligned} & \max u_i(\tilde{x}_i) \\ & \text{s.t. } p \cdot \tilde{x} \leq p \cdot \omega_i. \end{aligned}$$

The result now follows from proving that $u_i(x_i^*) - u_i(x_i) < \varepsilon$. First, if $p \cdot x \geq p \cdot \omega_i$ there is nothing to prove, as the desired conclusion follows from the definition of Walrasian ε -equilibrium. Let us then assume that $p \cdot x < p \cdot \omega_i$.

By the concavity of u_i ,

$$\begin{aligned} u_i(x_i^*) - u_i(x_i) & \leq Du_i(x_i) \cdot (x_i^* - x_i) \\ & = (1/\lambda_i)p \cdot (x_i^* - x_i) \\ & \leq (1/\eta)p \cdot \omega_i^* - p \cdot x_i \\ & \leq (1/\eta)\varepsilon' < \varepsilon. \end{aligned}$$

Case 2: strong ε -approximate equilibrium.

Let η be as above, and let $\varepsilon' > 0$ be such that

$$\varepsilon' < \frac{\varepsilon\eta\pi}{n}.$$

The algorithm on the $\overline{\Delta}^\eta$ simplex outputs a Walrasian ε' -equilibrium (x, p) . Observe that η and ε' , depend on $(n-1)$, π and Θ . It is therefore constant across the economies in E .

Let x_i^* be a solution to

$$\begin{aligned} & \max u_i(\tilde{x}_i) \\ & \text{s.t. } p \cdot \tilde{x} \leq p \cdot \omega_i. \end{aligned}$$

This defines $x^* \in \mathbf{R}_+^{nl}$. We shall prove that (x^*, p) is a strong ε -approximate equilibrium.

Define, for $1 \leq s \leq l$, $y^s \in \mathbf{R}_+^l$ by

$$y_h^s = \begin{cases} x_h & \text{if } h \neq s \\ x_h + \frac{p \cdot (x_i^* - x_i)}{p_h} & \text{if } h = s. \end{cases}$$

and $\theta_s = \frac{p_s \cdot (x_{si}^* - x_{si})}{p \cdot (x_i^* - x_i)}$ (here x_{si}^* denotes the amount of good s in bundle x^* , and similarly for x_{si}). So $\sum_s \theta_s = 1$. Then it is easy to verify that

$$x^* = \sum_{s=1}^l \theta_s y^s.$$

Note that $p = \lambda_i Du_i(x_i)$, so Lemma 4.2 implies that, for any good s , $p_s \geq \eta\pi$. Then, using the expression of x^* in terms of the vectors y^s we obtain:

$$\|x_i^* - x_i\| \leq \sup\left\{\frac{|p \cdot (x_i^* - x_i)|}{p_h}\right\} \leq \sup\left\{\frac{|p \cdot \omega_i - p \cdot x_i|}{\eta\pi}\right\} < \frac{\varepsilon'}{\eta\pi}$$

Then

$$\|x^* - x\| \leq \sum_i \|x_i^* - x_i\| < n \frac{\varepsilon'}{\eta\pi} < \varepsilon,$$

where the last inequality follows from the choice of ε' .

5. APPLICATIONS TO REPLICA ECONOMIES, REPRESENTATIVE CONSUMERS, FISHER EQUILIBRIA, AND ENDOWMENT CLASSES

We have presented our main result (Theorem 2.1) in a context with a fixed number of agents, but it also has applications more generally. As we emphasized in the introduction, many economic models assume that there are many goods, but a fixed number of agents. It is nonetheless very interesting to study economies with a large number of agents as well as goods. In this section we limit agent heterogeneity in ways that are a bit more subtle than assuming there is a fixed number of them. But, in all these cases, we show that when agent heterogeneity is limited, finding a Walrasian equilibrium is easy.

First, and most immediately, our result applies directly to *replica economies*. These are economies with many agents, where each agent is a copy of some prototypical (small) set of types of agents. Replica economies are one of the most important models of large economies in economics. They play a fundamental role in results on core convergence.³

In our second class of models, we discuss the existence of a representative consumer. We give a sufficient condition for an economy to admit a representative consumer, and observe that our result immediately gives an algorithm for approximating Walrasian equilibria. Clearly, representative consumers only exist under very stringent assumptions, but there are nevertheless many important models in economics that assume

³See Chapter 18 in [MasColell et al. 1995]. The theorem by Debreu and Scarf is one of the most important results in general equilibrium theory; see [Debreu and Scarf 1963].

their existence. Famously, many models in macroeconomics assume a representative consumer. For the literature on computing economic equilibria, the most important instance of representative consumers corresponds to Fisher equilibria. Some of the most positive recent results on the complexity of economic models have been related to Fisher equilibria ([Jain and Vazirani 2007]). We show that the results on Fisher equilibria are a special case of our result, the reason being that their models are essentially models of a representative consumer.

The third class of models we consider are those economies that have many agents but where each of them belongs to one of a small number of “endowment classes.” For example, some agents are laborers (endowed with labor), while others are endowed with land, and others with capital. If we can partition agents into a small number of classes, then—together with a strengthening of our assumption on preferences—we can in effect work with a model with a small number of agents, even if the actual number is large. The idea we are exploiting is that consumers may be heterogeneous in their preferences, but not in their sources of income. Under homotheticity, we can invoke some classical aggregation theorems to effectively work with a small set of agents.

5.1. Replica economies

In this section we present a direct application of our result to a model with an unbounded number of agents and replica economies. The main purpose of this section is to flesh out the idea that, in a world with limited heterogeneity on the part of agents, approximate Walrasian equilibria may be easy to find.

Consider an exchange economy $(\omega_i, u_i)_{i=1}^n$ under the assumptions we established in Section 2. The K th replica of $(\omega_i, u_i)_{i=1}^n$ is the exchange economy $(\omega_{i,k}, u_{i,k})_{i=1, \dots, n, k=1, \dots, K}$ where for all i, j and k we have that $\omega_{ik} = \omega_{jk}$ and $u_{ik} = u_{jk}$. In a replica economy, each agent is indexed by a pair of numbers i and k : i denotes the “type” of agent and k denotes the “serial number” of the agent, among those of his type.

It is trivial to verify that a Walrasian ε -equilibrium of $(\omega_i, u_i)_{i=1}^n$ is also a Walrasian ε -equilibrium of $(\omega_{i,k}, u_{i,k})_{i=1, \dots, n, k=1, \dots, K}$, for any k .⁴ Our algorithm gives an approximate Walrasian equilibrium of $(\omega_i, u_i)_{i=1}^n$. This is also an approximate Walrasian equilibrium of $(\omega_{i,k}, u_{i,k})_{i=1, \dots, n, k=1, \dots, K}$, for any K .

5.2. Representative consumers and Fisher equilibria

Our main result deals with economies with a fixed number of agents. One special case is, of course, when there is only one agent. Economists have long been interested in economies with many agents that behave as if there were only one. Without controlling how heterogeneous the agents can be, it is impossible that individual behavior will “aggregate” into the behavior of a “representative consumer.” Interestingly, there are plausible (but strong) restrictions on individual heterogeneity that guarantee the existence of a representative consumer.

Individual heterogeneity in our model has two sources: preferences and endowments. If endowments are not restricted, one needs preferences to be essentially identical.⁵ We shall instead restrict how endowments are allowed to vary and allow preferences to be heterogeneous. In the computer science literature, there is a precedent in this respect: many computer scientists have been interested in Fisher equilibria, which results from requiring that endowments are proportional. By focusing on aggregation,

⁴The approach is used to establish that the core of the replicated economies is, in the limit, equal to the set of Walrasian equilibrium allocation of the original economy $(\omega_i, u_i)_{i=1}^n$.

⁵This idea is captured by a result known as Antonelli’s Theorem.

and the existence of a representative consumer when we restrict the heterogeneity of endowments, we can generalize the ideas in computing Fisher equilibria.

Consider a collection of n agents. We shall assume that each one of them has a continuous, strictly monotone, and strictly quasiconcave utility function u_i . Let $d_i(p, m)$ denote the solution to the problem of maximizing u_i over $x_i \in \mathbf{R}_+^l$ such that $p \cdot x_i \leq m$; the function d_i is the *demand function* generated by u_i .

The following result determines the fraction of income (endowments) that goes to agent i under prices p . Informally, the above theorem says that when these fractions δ_i^* are chosen to optimize some social welfare function W , then one obtains a representative consumer.

THEOREM 5.1 (SAMUELSON'S AGGREGATION THEOREM). *Let $W : \mathbf{R}^n \rightarrow \mathbf{R}$ be strictly increasing. If, for every $p, \omega \in \mathbf{R}_{++}^l$, $\delta_1^*(p, \omega), \dots, \delta_n^*(p, \omega)$ solves the problem*

$$\max_{\delta \in \Delta} W(u_1(d_1(p, p \cdot \delta_1 \omega)), \dots, u_n(d_n(p, p \cdot \delta_n \omega)));$$

then there is a continuous, strictly monotonic and concave function u , generating a demand function d such that

$$d(p, p \cdot \omega) = \sum_{i=1}^n d_i(p, p \cdot \delta_i^*(p, \omega) \omega),$$

for all $p, \omega \in \mathbf{R}_{++}^l$. Further, u takes the form

$$u(x) = \max_{\text{s.t. } \sum_i x_i = x} W(u_1(x_1), \dots, u_n(x_n))$$

Importantly, the special case of Samuelson's theorem when δ_i^* are independent of prices and incomes coincides with the notion of Fisher equilibria. Specifically, given a set of n agents, with continuous, homothetic, monotone and strictly concave utility functions u_i , we can define a *Fisher equilibrium* as follows. Suppose there is 1 units of each of l goods, and that each agent i is endowed with α_i units of "money" (unit of account). Let $1 = \sum_i \alpha_i$. Then a *Fisher equilibrium* is a vector of prices $p \in \Delta$ such that

$$\sum_i d_i(p, \alpha_i) = (1, \dots, 1), \text{ i.e. supply equals demand.}$$

Now, a Fisher equilibrium is a special case of a Walrasian equilibrium in a model in which there is a representative consumer. The following Corollary, due to [Eisenberg 1961], is commonly used in the literature on Fisher equilibria and allows us to apply our algorithm in the context of Fisher equilibria.

COROLLARY 5.2 ([EISENBERG 1961]). *Fix $\alpha_1, \dots, \alpha_n \in \Delta$ and suppose that u_i is homothetic, in addition to the previously made assumptions. Then the utility function u defined by*

$$u(x) = \max_{\text{s.t. } \sum_i x_i = x} \prod (u_i(x_i))^{\alpha_i}$$

generates a demand function d such that

$$d(p, p \cdot \omega) = \sum_{i=1}^n d_i(p, p \cdot \delta_i^*(p, \omega) \omega),$$

for all $p, \omega \in \mathbf{R}_{++}^l$.

By Corollary 5.2, if $\omega = (1, \dots, 1)$ then a price vector p is a Fisher equilibrium if and only if it is a Walrasian equilibrium in the economy with a single (representative) consumer, with utility function u as defined in Corollary 5.2. Since this is an economy with a single consumer, Theorem 2.1 delivers an efficient algorithm for approximating Walrasian equilibria.

Note that the connection between Fisher equilibria and the Negishi approach is already remarked upon briefly in [Codenotti and Varadarajan 2007].

5.3. Fixed endowment classes

Consider a collection E of exchange economies $(\omega_i, u_i)_{i=1}^n$. Suppose now that both n and the number of goods can differ among the members of E . Assume that utilities and endowments satisfy all the assumptions of Section 2, and that in addition utilities are homothetic.

We limit the heterogeneity among agents in E by limiting how endowments differ among agents. For an exchange economy $(\omega_i, u_i)_{i=1}^n$, an *endowment class* is a set $P \subseteq \{1, \dots, n\}$, together with vectors $\omega \in \mathbf{R}_+^l$ and $(\alpha_i)_{i \in P}$ such that $\alpha_i \geq 0$, $\sum_{i \in P} \alpha_i = 1$, and $\omega_i = \alpha_i \omega$. Now, suppose that there is a fixed number K such that for every $(\omega_i, u_i)_{i=1}^n$ in E , there are at most K endowment classes that partition the set of agents $\{1, \dots, n\}$.

By Corollary 5.2, the homotheticity of utilities allow for the existence of a representative consumer for each of the K endowment classes. We can now find a Walrasian equilibrium for the economy populated by such representative consumers. From an equilibrium allocation and prices, one finds a final equilibrium allocation by solving the convex problem in Samuelson's Theorem. There is a fixed number of such problems to solve.

6. CONCLUDING REMARKS

In this paper, we present a new algorithm for efficiently computing approximate Walrasian equilibria in a setting where the number of agents is fixed and small, and the number of goods is large.

This work is motivated by the question of understanding what leads to hardness in the computation of Walrasian equilibria. As mentioned, there are a number of recent results showing that computing an approximate Walrasian equilibrium is “hard” in general settings. For example, [Chen et al. 2009a; Codenotti et al. 2006]. These existing hardness results have a relation to the celebrated results of [Sonnenschein 1972; Mantel 1974; Debreu 1974], which state that any continuous function that satisfies Walras's Law is the aggregate excess demand function of some well-behaved economy *with at least as many agents as there are goods*. The hardness of computing fixed points can be translated into the hardness of computing equilibria using this result, but note that this source of hardness requires the number of agents to be at least as large as the number of goods. Similarly, [Codenotti et al. 2006] reduces arbitrary two-player games, in which one player has n strategies and the other has m strategies, into an economy with $n + m$ agents and goods. Thus they consider problems in which both goods and agents grow at the same rate.

In contrast, our work proves that when the number of agents n is fixed, Walrasian equilibrium can be found efficiently. This highlights an interesting contrast between Walrasian equilibria and Nash equilibria. In the case of Nash equilibria, hardness emerges already for 2 agents, in both the context of exact [Chen et al. 2009b] and approximate equilibria [Chen et al. 2006]; whereas in the case of Walrasian equilibria we have exhibited an efficient algorithm for finding an approximate equilibrium given any fixed number of agents. In fact, as noted in Section 3, our proof actually guarantees that when the number of agents $n = O(\log l)$ our algorithm finds an approximate

Walrasian equilibrium in time polynomial in the number of goods l . An interesting open question is to understand how large the number of agents can be before hardness emerges.

Finally, we would like to highlight that our focus on a fixed number of agents is motivated by *models* in macroeconomics and finance, which typically have many goods and a fixed, small number of agents. We emphasize *models* because it is hard to doubt the importance of the computational complexity of a model; and researchers in these fields are heavy users of computational methods to find equilibria in their models. However, it is important to note that the relevance of complexity for actual economies, and for positive economics, is not guaranteed, as we explain elsewhere [Echenique et al. 2011].

REFERENCES

- ANDERSON, R. M. 1978. An elementary core equivalence theorem. *Econometrica* 46, 6, pp. 1483–1487.
- BEWLEY, T. 1991. A theorem on the existence of competitive equilibria in a market with a finite number of agents and whose commodity space is l^∞ . In *Equilibrium Theory in Infinite Dimensional Spaces*, M. Khan and N. Yannellis, Eds. Springer-Verlag.
- BORDER, K. C. 1989. *Fixed point theorems with applications to economics and game theory*. Cambridge Univ Pr.
- CHEN, X., DAI, D., DU, Y., AND TENG, S.-H. 2009a. Settling the complexity of arrow-debreu equilibria in markets with additively separable utilities. In *Proceedings of FOCS*. 273–282.
- CHEN, X., DENG, X., AND TENG, S.-H. 2006. Computing nash equilibria: Approximation and smoothed complexity. In *Proceedings of FOCS*. 603–612.
- CHEN, X., DENG, X., AND TENG, S.-H. 2009b. Settling the complexity of computing two-player Nash equilibria. *J. ACM* 56, 3, 1–57.
- CODENOTTI, B., PEMMARAJU, S., AND VARADARAJAN, K. 2004. The computation of market equilibria. *ACM SIGACT News* 35, 4, 23–37.
- CODENOTTI, B., PEMMARAJU, S., AND VARADARAJAN, K. 2005. On the polynomial time computation of equilibria for certain exchange economies. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 72–81.
- CODENOTTI, B., SABERI, A., VARADARAJAN, K., AND YE, Y. 2006. Leontief economies encode nonzero sum two-player games. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. ACM, 659–667.
- CODENOTTI, B. AND VARADARAJAN, K. 2007. Computation of market equilibria by convex programming. In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, Eds. Cambridge University Press.
- DEBREU, G. 1974. Excess demand functions. *Journal of Mathematical Economics* 1, 1, 15–21.
- DEBREU, G. AND SCARF, H. 1963. A limit theorem on the core of an economy. *International Economic Review* 4, 3, pp. 235–246.
- DENG, X., PAPADIMITRIOU, C., AND SAFRA, S. 2002. On the complexity of equilibria. In *Proceedings of the thirty-fourth annual ACM symposium on the theory of computing (STOC)*. ACM, 67–71.
- DEVANUR, N., PAPADIMITRIOU, C., SABERI, A., AND VAZIRANI, V. 2002. Market equilibrium via a primal-dual-type algorithm. In *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*. IEEE, 389–395.
- DEVANUR, N. R. AND KANNAN, R. 2008. Market equilibria in polynomial time for

- fixed number of goods or agents. In *IEEE Symposium on Foundations of Computer Science (FOCS)*. 45–53.
- ECHENIQUE, F., GOLOVIN, D., AND WIERMAN, A. 2011. A revealed preference approach to computational complexity in economics. In *Proceedings of the ACM conference on Electronic Commerce*.
- EISENBERG, E. 1961. Aggregation of utility functions. *Management Science* 7, 4, pp. 337–350.
- JAIN, K. 2004. A polynomial time algorithm for computing the arrow-debreu market equilibrium for linear utilities. *Foundations of Computer Science, Annual IEEE Symposium on 0*, 286–294.
- JAIN, K. AND VAZIRANI, V. 2007. Eisenberg-gale markets: Algorithms and structural properties. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*. ACM, 364–373.
- MAGILL, M. 1981. An equilibrium existence theorem. *Journal of Mathematical Analysis and Applications* 84, 1, 162–169.
- MANTEL, R. 1974. On the characterization of aggregate excess demand. *Journal of Economic Theory* 7, 3, 348–353.
- MAS-COLELL, A. 1986. The price equilibrium existence problem in topological vector lattices. *Econometrica* 54, 5, pp. 1039–1053.
- MASCOLELL, A., WHINSTON, M., AND GREEN, J. 1995. *Microeconomic theory*. Vol. 1. Oxford university press New York.
- NEGISHI, T. 1960. Welfare economics and existence of an equilibrium for a competitive economy. *Metroeconomica* 12, 2-3, 92–97.
- ROCKAFELLAR, R. T. 1970. *Convex Analysis*. Princeton Univ Press.
- SCARF, H. 1977. *The computation of equilibrium prices: an exposition*. Cowles Foundation for Research in Economics at Yale University.
- SEGAL, I. R. 2007. The communication requirements of social choice rules and supporting budget sets. *Journal of Economic Theory* 136, 1, 341 – 378.
- SONNENSCHNEIN, H. 1972. Market excess demand functions. *Econometrica*, 549–563.
- STARR, R. M. 1969. Quasi-equilibria in markets with non-convex preferences. *Econometrica* 37, 1, pp. 25–38.
- TODD, M. 1976. *The computation of fixed points and applications*. Vol. 1. Springer-Verlag Berlin.
- VAZIRANI, V. V. 2007. Combinatorial algorithms for market equilibria. In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, Eds. Cambridge University Press.