

Non-Redistributive Second Welfare Theorems

Bundit Laekhanukit¹, Guyslain Naves^{2**}, and Adrian Vetta³

¹ School of Computer Science, McGill University, Canada. Email: blaekh@cs.mcgill.ca

² Laboratoire d'Informatique Fondamentale de Marseille, Aix-Marseille Universités, France.

Email: guyslain.naves@lif.univ-mrs.fr

³ Department of Mathematics and Statistics and School of Computer Science, McGill University, Canada. Email: vetta@math.mcgill.ca

Abstract. The second welfare theorem tells us that social welfare in an economy can be maximized at an equilibrium given a suitable redistribution of the endowments. We examine welfare maximization without redistribution. Specifically, we examine whether the clustering of traders into k submarkets can improve welfare in a linear exchange economy. Such an economy always has a market clearing ε -approximate equilibrium. As $\varepsilon \rightarrow 0$, the limit of these approximate equilibria need not be an equilibrium but we show, using a more general price mechanism than the reals, that it is a “generalized equilibrium”. Exploiting this fact, we give a polynomial time algorithm that clusters the market to produce ε -approximate equilibria in these markets of near optimal social welfare, provided the number of goods and markets are constants. On the other hand, we show that it is NP-hard to find an optimal clustering in a linear exchange economy with a bounded number of goods and markets. The restriction to a bounded number of goods is necessary to obtain any reasonable approximation guarantee; with an unbounded number of goods, the problem is as hard as approximating the maximum independent set problem, even for the case of just two markets.

1 Introduction

The fundamental theorems of welfare economics are considered “the most remarkable achievements of modern microeconomic theory” [9] and are the “central set of propositions that economists have to offer the outside world - propositions that are in a real sense, the foundations of Western capitalism” [5]. Informally, they state (under certain conditions that we will discuss later)

First Fundamental Welfare Theorem. A competitive equilibrium is Pareto efficient.

Second Fundamental Welfare Theorem. Any Pareto efficient solution can be supported as a competitive equilibrium.

The First Welfare Theorem is widely viewed as “a mathematical statement of Adam Smith’s notion of the invisible hand leading to an efficient allocation” [12]. The Second Welfare Theorem implies that we can separate out issues of economic efficiency from issues

* This research was supported by Natural Sciences and Engineering Research Council of Canada (NSERC) grant no. 288334, 429598 and by Harold H Helm fellowship.

** The work was done while the second author was in Department of Mathematics and Statistics, McGill University, Canada.

of equity. Specifically, by redistributing the initial endowments (by lump-sum payments), a set of prices exists that can sustain any Pareto solution. This second theorem has “fundamental implications for how we think about economic organization” [13] and is “arguably the theoretical result that has had the most dramatic effect on economic thinking” [4]. Despite this, “much of public economics takes as its starting point the rejection of the practical value of the second theorem” [11]. Why this discrepancy? To understand this, note that lump-sum transfers are theoretically considered a very desirable form of taxation as they do not distort incentives within the pricing mechanism. However, this is essentially accomplished by a massive distortion of the initial market! Moreover, these are *personalized* liabilities which in turn can be viewed as an extremely unfair form of taxation in that they don’t depend upon the actions or behaviors of the agents, and are impractical for a myriad of implementational and political reasons (see, for example, [2], [3], [11] and [10]).

This observation motivates our work. Can the market mechanism be used to sustain Pareto allocations without redistribution? In particular, suppose that without redistribution a single market leads to low social welfare (or even has no competitive equilibrium at all). In these circumstances, can the market mechanism still be used to produce an allocation of high social welfare? We address this question under the classical model of exchange economy, and show that indeed this can often be achieved provided the single market can be clustered into submarkets.

1.1 The Exchange Economy.

We consider the classical model of an *exchange economy* – an economy without production. We have n traders $i \in \{1, 2, \dots, n\}$ and m goods $j \in \{1, 2, \dots, m\}$. (To avoid any ambiguity between traders and goods we will often refer to good j as good g_j). Each trader i has an initial endowment $\mathbf{e}_i \in \mathbb{R}_+^m$, where e_{ij} is the quantity of good g_j that she owns, and a utility function $u_i : \mathbb{R}_+^m \rightarrow \mathbb{R}$. The traders have no market power and so are *price-takers*. Given a set of prices $\mathbf{p} \in \mathbb{R}_+^m$, where p_j is the price of good g_j , trader i will demand the best bundle she can afford, that is, $\operatorname{argmax}_{\mathbf{x}_i} u_i(\mathbf{x}_i) \text{ s.t. } \mathbf{p} \cdot \mathbf{x}_i \leq \mathbf{p} \cdot \mathbf{e}_i$. These prices and demand bundles form a *Walrasian (competitive) equilibrium* if all markets clear: demand does not exceed supply for any good g_j . That is, $\sum_i x_{ij} \leq \sum_i e_{ij}$. In this paper, we focus on the basic case of linear utility functions – the *linear exchange model*. Here the function $u_i(\cdot)$ can be written as $u_i(\mathbf{x}_i) = \sum_{j=1}^m u_{ij} x_{ij}$ where $u_{ij} \geq 0$ is the utility per unit that trader i has for good g_j . (We denote by \mathbf{u}_i the vector of utility coefficients for trader i .)

1.2 The Fundamental Welfare Theorems.

An allocation is *Pareto efficient* if there is no feasible allocation in which some trader is strictly better off but no trader is worse off. The first welfare theorem states that any Walrasian equilibrium is Pareto efficient. It holds under very mild conditions, such as monotonic utilities or non-satiation. Clearly for this result to be of interest, though, we need this economy to possess Walrasian equilibria. In groundbreaking work, Arrow and Debreu [1] showed that this is indeed the case, under certain conditions such as concave utility functions and positive endowments.⁴ Interestingly, equilibria need not exist even in a

⁴ As well as the possibility of non-concave utility function, numerous other factors may affect the practicality of the welfare theorems: market power and the presence of price-makers; incomplete or asymmetric information; externalities; convergence issues for equilibria; the existence of multiple equilibria; economies of scale when production is added to the exchange economy, etc. Such issues are not our focus here.

linear exchange economy. However, there is a combinatorial characterization for existence due to Gale [6], and we discuss this characterization and other properties of the linear exchange economy in detail in Section 2.

Observe that Pareto efficiency is not a particularly restrictive notion: an allocation is efficient unless there is an alternative that is *universally* agreed to be better (or at least as good). This requirement of unanimity has important implications. Allocations that may be viewed as societally better outcomes may be blocked by a single agent. For example, Pareto allocations can be extremely inequitable. The second welfare theorem attempts to address this concern: Any Pareto solution can be supported as a Walrasian equilibrium. Specifically, by redistributing the initial endowments via lump-sum payments, a set of prices exists that can sustain any Pareto solution. (The second theorem also requires concave utility functions.) Thus, the second welfare theorem implies that we can separate out issues of economic efficiency from issues of equity.

As stated, however, the second theorem is of limited practical value due to the infeasibility of direct transfer payments. Thus, our goal is to obtain non-redistributive second welfare theorems. Specifically, maximizing the social welfare, $\sum_{i=1}^n u_i(\mathbf{x}_i)$, is a fundamental question in economics; so, can we support at equilibrium an allocation with high social welfare? For example, in a linear exchange economy it is particularly easy to find an optimal social allocation. For each good g_j , simply give all of it to the trader i for whom it proffers the greatest utility per unit. However, even in this basic case, a Walrasian equilibrium may produce very low social welfare. Intuitively the reason is simple: a trader with a large utility coefficient for a good may not be able to afford many units of it. This may be because (a) the good is in high demand and thus has a high price and/or (b) the trader has a small budget because the goods she initially possesses are abundant and, thus, have a low price.

On the other hand, the second welfare theorem tells us that, with redistribution, it is possible to find prices that support an allocation of optimal welfare. Can any more practical, market-based mechanisms achieve this? To answer this, we consider a mechanism that is allowed to cluster the traders into trading groups.

1.3 Market Clustering.

Suppose we partition the traders into k separate markets, for some integer k . In each market t , trade then proceeds as normal with a distinct set of Walrasian prices \mathbf{p}^t and corresponding allocation generated. This partition and the collection of Walrasian equilibria (if they exist) form a *k-equilibrium*. The allocation induced by a *k-equilibrium* may be very different than one produced from a single market. So, first, can market clustering be used to improve social welfare? If so, second, can it be used to optimize social welfare?

The answer to the first question is **yes**. Trivially, the option to segment the market cannot hurt because we could simply place all the traders in the same market anyway. In fact, market clustering may dramatically improve social welfare; there is an example where the ratio between the social welfare with two markets and the social welfare with one market is unbounded.

The answer to the second question, however, is **no**. Not every Pareto solution can be supported by market clustering. In particular, there are cases where the optimal social solution cannot be obtained by clustering. Indeed, there is an example where the ratio between the optimal social welfare and the optimal welfare that can be generated by market clustering is also unbounded.

The main focus of this paper then becomes to efficiently obtain as large a welfare as possible under market clustering.

We remark that the basic idea underlying market clustering, i.e., the grouping and separation of traders, is a classical one in both economic theory and practice. In particular, it lies at the heart of the theory of trade. On the one hand, countries should trade together (grouping) to exploit the laws of comparative advantage; on the other hand, trade between countries may be restricted (separation) to protect the interests of certain subsections (e.g. specific industries or classes of worker). Interestingly, of course, whilst separation has a net negative effect on welfare in international trade models, our results show that it can have a large net positive effect in general equilibrium models. Other examples that can be viewed as market clustering arise in the regulation of oligopolies and in the issue of trading permits. A less obvious example concerns bandwidth auctions where participants are grouped into “large” (incumbent) and “small” (new-entrants). Trade, with the mechanism in the form of feasible bidding strategies, is then restricted depending upon the group.

1.4 Our Results.

Our main result, in Section 6, is a polynomial time algorithm that finds an ε -approximate k -equilibrium, of almost optimal social welfare, provided the number of goods and markets are constants. The key to this result is a limit theorem in Section 5 showing that, in a single market, ε -approximate equilibria converge to what we call a *generalized equilibrium*.

On the other hand, in Section 4 we show that it is NP-hard to find an optimal k -equilibrium in a linear exchange economy with a bounded number of goods and markets. The restriction to a bounded number of goods is necessary to obtain any reasonable approximation guarantees; for linear exchange economies with an unbounded number of goods, the problem is as hard as approximating the maximum independent set problem, even for the case of just 2 markets.

Due to space constraints many proofs of theorems and lemmas are omitted but can be found in the full paper.

2 Walrasian Equilibria in the Linear Exchange Model

Take an equilibrium with prices \mathbf{p} and allocations \mathbf{x}_i for the Walrasian model with linear utilities. Recall, we may assume that the followings hold:

Budget Constraints: Trader i cannot spend more than she receives: $\mathbf{p} \cdot \mathbf{x}_i \leq \mathbf{p} \cdot \mathbf{e}_i$ (1)

Optimality: Each trader i optimizes the bundle of goods she buys: (2)
 $\mathbf{u}_i \cdot \mathbf{x}_i$ is maximized subject to (1)

Market Clearing: Demand does not exceed supply, for any g_j : $\sum_i x_{ij} \leq \sum_i e_{ij}$ (3)

2.1 Properties of Equilibria

The following claims are well-known facts (see *e.g.* [8]).

Claim 1. *At equilibrium, the budget constraint (1) is tight for any trader.*

Claim 2. *At equilibrium, the market clearing condition (3) is tight for any g_j with $p_j > 0$.*

Claim 3. *At equilibrium, for any subset S of traders, there is a good g_j such that*

$$\sum_{i \in S} x_{ij} > \sum_{i \in S} e_{ij} \iff \text{there is a good } g_{j'} \text{ with } \sum_{i \notin S} x_{ij'} > \sum_{i \notin S} e_{ij'}$$

Claim 4. At equilibrium, for any i with $P_i := \mathbf{p} \cdot \mathbf{e}_i > 0$ and any good g_j with price $p_j > 0$

$$\frac{u_{ij}}{p_j} \leq \frac{\mathbf{u}_i \cdot \mathbf{x}_i}{P_i} \quad (4)$$

Moreover, the inequality is tight for any i, j with $x_{ij} > 0$.

2.2 The Existence of Equilibria in a Single Market

Gale [6] gave a characterization for when linear exchange economies possess equilibria. Observe that the price of every good will be strictly positive provided that each good is owned by at least one trader, and at least one trader desires it. We may assume this is the case as any good that does not satisfy this condition may be removed from the model; in this case, supply will exactly equal demand for each good. Gale also assumes that every trader is *non-altruistic* in that they each desire at least one good. (We say that a trader i is an *altruist* if $u_{ij} = 0$ for every good g_j .)

Theorem 5 ([6]). *An altruist-free linear exchange economy has a Walrasian equilibrium if and only if there is no super self-sufficient set of traders.*

Here, a subset S of traders is called *super self-sufficient* if

1. **Self-Sufficiency:** $\sum_{i \notin S} e_{ij} = 0$ for every good g_j such that $\sum_{i \in S} u_{ij} > 0$.
2. **Superfluity:** There is a good g_j such that $\sum_{i \in S} e_{ij} > 0$ and $\sum_{i \in S} u_{ij} = 0$.

It will be useful to reinterpret Gale's condition combinatorially using the market graph. The *market graph* G_M for a given market is a directed graph whose set of vertices is the set of goods in that market. There is an arc $j \rightarrow j'$ with label i if there is a trader i with $e_{ij} > 0$ and $u_{ij'} > 0$; thus, trader i has good g_j and, depending upon the prices, is willing to trade it for good $g_{j'}$. Let $h(S)$ denote the goods that are the heads of arcs with labels from traders in S , and let $t(\bar{S})$ denote the goods that are the tails of arcs corresponding to traders not in S . Then S is self-sufficient if $h(S) \cap t(\bar{S}) = \emptyset$. In this case, $h(S)$ induces a directed in-cut in the market graph. (Thus, a sufficient – but not necessary – condition for the existence of an equilibrium is the strong connectivity of the market graph. Moreover, any directed cut will correspond to a self-sufficient set.) If, in addition, $h(S)$ is a strict subset of $t(S)$, then S is super self-sufficient.

For example, the market graph shown in Figure 1 does not have an equilibrium. It represents a market with 6 traders and 5 goods g_1, \dots, g_5 : each arc $g \xrightarrow{i} h$ represents one trader i with $e_{ig} = 1$ and $u_{ih} = 1$; all other values are 0. Then traders $\{4, 5, 6\}$ form a super self-sufficient set, so this market does not have an equilibrium.

We can use the market graph to test Gale's condition efficiently. Furthermore, Jain [8] gave a polynomial time algorithm to find an equilibrium when one exists.

2.3 The Existence of Equilibria in a Market Clustering.

Recall that a trader i is an altruist if $u_{ij} = 0$ for every good g_j . An economy is *altruistic* if it is allowed to contain altruistic traders. It is important for us to understand the implications of altruism even for economies without altruists. This is because clustering may create de facto altruists in the submarkets. Moreover, such altruists are one of the factors that allow the equilibria to exist in a market clustering, even if the single market has no equilibrium. We can easily extend Gale's theorem to altruistic economies.

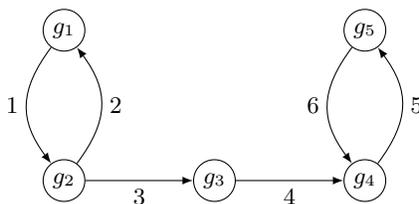


Fig. 1. A market graph with a super self-sufficient set and, therefore, has no equilibrium.

Theorem 6. *An altruistic, linear exchange economy has an equilibrium if and only if every super self-sufficient set of traders contains at least one altruist.*

So, altruistic economies need not have equilibria. However, they can always be clustered into markets with equilibria provided that the number of markets k is at least the number of goods m .

Theorem 7. *An altruistic, linear exchange economy with m goods has an m -equilibrium.*

Proof. We prove this by induction on the number of goods. An altruistic economy with one good g_j has a trivial equilibrium. Now take an altruistic economy with m goods. If it has no super self-sufficient set of traders consisting entirely of non-altruists then, by Theorem 6, it has an equilibrium. Otherwise, let S be a minimal super self-sufficient set of non-altruists. By minimality, the market induced by S contains an equilibrium as it has no super self-sufficient subset.

As they are not altruistic, each trader in S desires at least one good. By definition, however, traders in S desire no goods held by traders in \bar{S} . So, there is at least one good held by S that is not held by traders in \bar{S} . Thus, the market induced by the traders in \bar{S} contains at most $m - 1$ goods. By induction it can be partitioned into $m - 1$ clusters that each has an equilibrium. Together with the cluster S , we obtain an m -equilibrium. \square

For example, consider again the market in Figure 1. If we partition the traders into two, with trader 3 alone in the first market and traders $\{1, 2, 4, 5, 6\}$ in the second market, then both resulting markets have equilibria (with $x_{32} = 1$ in the first market, and $p_3 = 0$ in the second market).

3 Single Markets, Market Clustering and Welfare Redistribution.

In this section, we examine the potential benefits of market clustering and the limits of its power as a tool. First, we have seen that equilibria may not exist in the single market case (i.e., when market clustering is prohibited). In such instances, by Theorem 7, market clustering can always be applied to produce equilibria. Furthermore, even when equilibria do exist in the single market case, market clustering may lead to huge improvements in social welfare in comparison. On the other hand, market clustering is not as powerful as welfare redistribution; specifically, market clustering does not always support every Pareto allocation. To see this, we consider two measures regarding the social welfare function:

1. **The Clustering Ratio:** the ratio between the maximum social welfare under market clustering and the social welfare obtained in a single market.

2. **The Redistribution Ratio:** the ratio between the maximum achievable social welfare (with welfare redistribution) and the maximum welfare under market clustering.

Examples showing that both ratios can be unbounded can be found in the full paper.

4 The Hardness of Market Clustering

In this section, we consider the hardness of the k -market clustering problem. We show that the problem is NP-hard even if we only have a fixed number of goods and a fixed number of markets, that is, m and k are constant.

Theorem 8. *Given an instance of the 2-market clustering problem with five goods and linear utility functions, it is NP-hard to decide whether there is a clustering that yields a social welfare of value at least Z , for any $Z > 0$.*

The problem becomes much harder when the number of goods is unbounded.

Theorem 9. *For any constant $\delta > 0$ and maximum social welfare Z , unless $\text{NP} = \text{ZPP}$, it is hard to distinguish between the following two cases:*

- **Yes-Instance:** *There is a clustering that yields a social welfare of value at least $Z^{1-\delta}$.*
- **No-Instance:** *There is no clustering that yields a social welfare of value at least Z^δ .*

5 Approximate Walrasian Equilibria.

We are interested in finding an ε -approximate market equilibrium; that is, for each market, our algorithm outputs a price \mathbf{p} and an allocation \mathbf{x} satisfying the following conditions.

- **Budget Constraints:** Trader i cannot spend more than she receives: $\mathbf{x}_i \cdot \mathbf{p} \leq \mathbf{e}_i \cdot \mathbf{p}$
- **Approximate Optimality:** Subject to the budget constraints, each trader i purchases a bundle \mathbf{x}_i whose utility is similar to that of the optimal bundle \mathbf{x}_i^* : $\mathbf{u}_i \cdot \mathbf{x}_i \geq (1-\varepsilon)\mathbf{u}_i \cdot \mathbf{x}_i^*$
- **Market Clearing:** Demand never exceeds supply: for any g_j , $\sum_i x_{ij} \leq \sum_i e_{ij}$

5.1 Existence of Approximate Walrasian Equilibria.

Compared to exact market equilibria, which do not always exist, there is always an approximate market equilibrium with arbitrary small approximation:

Theorem 10. *For $\varepsilon > 0$, every market has a market-clearing ε -approximate equilibrium.*

By *market-clearing ε -approximate equilibrium*, we mean an ε -approximate equilibrium for which the approximate market clearing inequality is tight. This theorem can be inferred from the algorithm of [7]. A direct proof of this fact can be found in the full paper.

5.2 Properties of Approximate Walrasian Equilibria.

We now discuss some properties of equilibria that will later be very useful to us in designing efficient algorithms. Given a market, we use the following definitions:

- $u_{\max} = \max_{i,j} u_{ij}$ is the maximum coefficient of any utility function.
- $u_{\min} = \min_{i,j:u_{ij}>0} u_{ij}$ is the minimum *non-zero* coefficient of any utility function.

- $p_{\max} = \max_j p_j$ is the maximum price of any good in the market.
- $p_{\min} = \min_{j:p_j>0} p_j$ is the minimum *non-zero* price of any good in the market.
- $e_{\min} = \min_{i,j:e_{ij}>0} e_{ij}$ is the minimum *non-zero* endowment for any good and trader.

Assume wlog that $\sum_i e_{ij} = 1$, for every good g_j . We can connect the above values via the market graph. Recall that n denotes the number of traders and m denotes the number of goods; then we obtain

Lemma 1. *If the market graph is strongly connected then, at a market equilibrium,*

$$\frac{p_{\max}}{p_{\min}} \leq e^{\frac{nm}{e}} \left(\frac{u_{\max}}{e_{\min} u_{\min}} \right)^{nm}$$

In particular, scaling so that $p_{\min} = 1$ gives $p_{\max} \leq e^{\frac{nm}{e}} \left(\frac{u_{\max}}{e_{\min} u_{\min}} \right)^{nm}$.

Proof. We may assume (solely for the duration of this proof) that each trader has a positive endowment for exactly one good and no two traders have positive endowments for the same good [8]. To do this, consider a trader i with endowment \mathbf{e}_i . For each good g_j such that $e_{i,j} > 0$, we create a new trader i_j with $e_{i_j j} = e_{ij}$, $\mathbf{u}_{i_j} = \mathbf{u}_i$ and $e_{i_j j'} = 0$, for all $j' \neq j$. So, each trader now has only one good. Furthermore, if two traders have the same good, then we simply give the good two different names (and replicate the utility functions of other traders accordingly). Now, each trader represents a unique good, i.e., a trader i has a positive endowment for the unique good g_i . So, we have at most nm traders/goods.

These transformations maintain the strong connectivity of the market graph. Moreover, all the copies of the same original good will have the same price in an equilibrium. After these transformations, the number of units of each good will in general be less than one. Thus, we scale all the initial endowments so that each trader i has one unit of a good g_i . In addition, we must scale the coefficients of utility functions; otherwise, the scaling would effect the social welfare. Specifically, for each good g_i , we divide the initial endowments of trader i by e_{ii} , and we multiply the utilities of every trader for this good by e_{ii} , so as to keep the prices unchanged.

We may assume that no good has a price of zero. By Equation (4), we have:

$$\frac{u'_{ij}}{p_j} \leq \frac{\sum_{\ell} u'_{i\ell} \cdot x_{i\ell}}{p_i}$$

For a pair i, j with $u'_{ij} > 0$, we get

$$\frac{p_i}{p_j} \leq \frac{\sum_{\ell} u'_{i\ell} \cdot x_{i\ell}}{u'_{ij}} \leq \frac{u'_{\max}}{u'_{\min}} \sum_{\ell} x_{i\ell}$$

Assume $p'_{\min} = p'_{i_0}$ and $p'_{\max} = p'_{i_s}$, where $s \leq nm$. Because the market is strongly connected, there is a sequence of traders with indices i_0, i_1, \dots, i_s , $s \leq nm$, such that $u'_{i_{j-1} i_j} > 0$ for all $i \in \{1, \dots, s\}$. Multiplying the previous inequalities for all consecutive terms of this sequence, we get

$$\frac{p_{\max}}{p_{\min}} = \prod_{j=0}^{s-1} \frac{p_{i_{j+1}}}{p_{i_j}} \leq \left(\frac{u'_{\max}}{u'_{\min}} \right)^s \cdot \prod_{j=0}^{s-1} \left(\sum_{\ell} x_{i_j \ell} \right) \leq \left(\frac{nm}{s} \right)^s \left(\frac{u'_{\max}}{u'_{\min}} \right)^{nm} \leq e^{\frac{nm}{e}} \left(\frac{u'_{\max}}{u'_{\min}} \right)^{nm}$$

Here the second inequality follows from the Arithmetic-Geometric Mean Inequality and the fact that $\sum_{j=0}^{s-1} \sum_{\ell} x_{i_j \ell} \leq \sum_{j=0}^{s-1} \sum_{\ell} e_{j\ell} \leq nm$ by the market clearing constraint (3). Now, observe that there is a pair i, j such that $u'_{\min} = e_{jj} u_{ij} \geq e_{\min} u_{\min}$, and there is a (different) pair i, j such that $u'_{\max} = e_{jj} u_{ij} \leq u_{\max}$. \square

The same reasoning applied to approximate equilibria gives:

Lemma 2. *If the market graph is strongly connected, at an ε -approximate market equilibrium, we have $\frac{p_{\max}}{p_{\min}} \leq e^{\frac{nm}{\varepsilon}} \left(\frac{u_{\max}}{(1-\varepsilon)u_{\min}e_{\min}} \right)^{nm}$.*

It is possible for a market that is not strongly connected to have an equilibrium: the owners of the goods reachable from a strongly connected set induce a self-sufficient set but not necessarily a super self-sufficient set. However, in this case, we cannot bound the ratio p_{\max}/p_{\min} as seen from the following lemma.

Lemma 3. *Consider a market with equilibrium \mathbf{p}, \mathbf{x} . Let W be a proper subset of goods such that for any trader i , if there is some good $g_j \in W$ with $e_{ij} > 0$, then $u_{ik} = 0$ for all $k \notin W$ (i.e., W is the shore of a directed cut in the market graph). Then, for any $B > 1$, \mathbf{p}', \mathbf{x} is also an equilibrium where $p'_j = p_j$ if $g_j \notin W$ and $p'_j = Bp_j$ if $g_j \in W$.*

Proof. The lemma follows from the following two facts. First, $x_{ij} = 0$ for any good $g_j \notin W$ and any trader i with $\sum_{k \in W} e_{ik} > 0$ since then $u_{ij} = 0$ by the definition of W . Second, $x_{ij} = 0$ for any trader i with $\sum_{k \in W} e_{ik} = 0$ and $g_j \in W$. Consequently, scaling the prices of goods in W does not effect the equilibrium. \square

An implication of this lemma is that the strongly connected components have price allocations that are essentially independent of each other: for example one could decompose the problem, find local equilibria in each component, and then scale the prices accordingly to get a global equilibrium. Also, again by scaling the prices of W , we can assume that the minimum price in W is no more than u_{\max}/u_{\min} times the maximum price outside W , as it does not change the optimality of the allocations (it would be a problem if there was a trader with $\sum_{j \notin W} u_{ij} = 0$ and $\sum_{j \notin W} e_{ij} = 0$, but then taking this trader plus $\{i : \sum_{j \in W} e_{ij} > 0\}$ would give a super self-sufficient set). This gives the following strengthening:

Lemma 4. *Any market having an equilibrium has one such that $\frac{p_{\max}}{p_{\min}} \leq e^{\frac{nm}{\varepsilon}} \left(\frac{u_{\max}}{e_{\min}u_{\min}} \right)^{nm}$.*

5.3 Limits of Equilibria

By Theorem 10, for any $\varepsilon > 0$ there is a market-clearing ε -approximate equilibrium. When ε tends to 0, the prices of these approximate equilibria may diverge (if no exact equilibrium exists), but the allocations of goods to traders, as they are chosen from a compact set, admit at least one limit point, an allocation $\hat{\mathbf{x}}$. We call such an allocation a *limit allocation*. In particular, if the market admits an exact equilibrium, then $\hat{\mathbf{x}}$ is the allocation of an exact equilibrium (if $\hat{\mathbf{x}}$ could not be obtained as a limit of approximate equilibria with converging prices, one could exhibit a super self-sufficient set and this would be a contradiction). In any case, $\hat{\mathbf{x}}$ satisfies the market clearing constraints with equality.

The allocation $\hat{\mathbf{x}}$ may not be supported by a set of real prices. For example, there is obviously no set of prices supporting $\hat{\mathbf{x}}$ when the market does not have an exact equilibrium. We show that $\hat{\mathbf{x}}$ can be supported by taking prices from a set more general than the real numbers. Consider the set $Q = \mathbb{N} \times \mathbb{R}_+$, our new set of “prices”. We denote by π_1 and π_2 the first and second projection, i.e., $\pi_1(x, y) = \pi_2(y, x) = x$. We extend these projections to vectors (and abuse notation) by: $\pi_i((v_j)_j) = (\pi_i(v_j))_j$. We then redefine the notion of equilibrium in terms of Q . For $p \in Q^m$ and $\hat{\mathbf{x}} \in \mathbb{R}_+^{m \times n}$, let the *rank* r_i of i be the maximum a such that $\sum_j : \pi_1(p_j) = a \ e_{ij} > 0$, for all i . The pair \mathbf{p}, \mathbf{x} is a *generalized equilibrium* if

- **Budget Constraints:** For all $i \in \{1, \dots, n\}$, for all $a \geq r_i$,

$$\sum_{j : \pi_1(p_j)=a} \pi_2(p_j) \cdot \mathbf{x}_{ij} \leq \sum_{j : \pi_1(p_j)=a} \pi_2(p_j) \cdot \mathbf{e}_{ij}$$

- **Optimality:** For each trader, \mathbf{x}_i maximizes the utility $\mathbf{u}_i \cdot \mathbf{x}_i$ over all allocations satisfying the budget constraint.
- **Market Clearing:** No good is in deficit: $\sum_i x_{ij} \leq \sum_i e_{ij}$ for all goods j with $\pi_2(p_j) > 0$.

This is indeed a generalization. If we force the prices to be in $\{0\} \times \mathbb{R}_+$ then a generalized equilibrium would give a Walrasian equilibrium. An ε -approximate generalized equilibrium is defined by replacing the optimality condition by: $\mathbf{u}_i \cdot \mathbf{x}_i$ is at least $(1 - \varepsilon)$ times the utility of a best response of trader i , for all i .

Theorem 11. *For any market, each limit allocation $\hat{\mathbf{x}}$ gives a generalized equilibrium.*

A generalized equilibrium can be approximated by an approximate Walrasian equilibrium with almost as high welfare. The converse is not true. An approximate equilibrium may achieve a welfare arbitrarily high compared to a generalized equilibrium; consider the market with two traders and two goods where $e_{11} = e_{22} = 1$, $u_{12} = L$, $u_{22} = 1$ and all the other values are zero. In this market, the only generalized equilibrium has welfare 1, but there is an approximate equilibrium with welfare $\varepsilon \cdot L + (1 - \varepsilon)$, and this tends to $+\infty$ when L tends to $+\infty$.

Lemma 5. *Let $\hat{\mathbf{x}}, \hat{\mathbf{p}}$ be a generalized equilibrium. For any $\varepsilon > 0$, there is an approximate equilibrium with total welfare at least $1 - \varepsilon$ times the welfare of $\hat{\mathbf{x}}, \hat{\mathbf{p}}$.*

6 A Fully Polynomial Time Approximation Scheme

In this section, we exploit the structure we have now developed to obtain a polynomial time algorithm to find an ε -approximate equilibrium for the k -market clustering problem where the number of goods m and the number of markets k are constant. Moreover, this equilibrium has a very strong welfare guarantee: it gives social welfare of at least $1 - \varepsilon$ times the welfare of the optimal k -cluster generalized equilibrium.

Theorem 12. *For any $\varepsilon > 0$ and for fixed $k, n \in \mathbb{N}$, there is an algorithm that, given a market M with m goods, computes within time polynomial in $\frac{1}{\varepsilon}$ and the size of the M , an ε -approximate generalized k -equilibrium for M with welfare at least $1 - \varepsilon$ the optimal welfare of a generalized k -equilibrium.*

6.1 The Dynamic Program

Our dynamic program takes as an input a set of generalized prices for every good in each market. This follows as we may try all possible prices selected from a finite set of prices in $\{1, \dots, m\} \times P$ where $P = \{1, 1 + 1/b, (1 + 1/b)^2, \dots, (1 + 1/b)^{\sigma-2}\}$. Here $b \in \mathbb{N}_+$ is a parameter to be set later, and σ is such that $(1 + 1/b)^{\sigma-3} \leq e^{nm/e} \left(\frac{m \cdot u_{max}}{e_{min} u_{min}} \right)^{nm} \leq (1 + 1/b)^{\sigma-2}$. A generalized price (a, p) encodes a real price $v(a, p)$ given by $L^a \cdot p$, where L is an arbitrarily large constant. So, we have $m \cdot \sigma$ possible prices for each good in each market. Thus, the number of combinations of prices is $(m\sigma)^{km}$, where k is the number of markets and m is the number of goods.

Given the estimated prices, the dynamic program runs over each market to compute an approximate equilibrium that maximizes the total social welfare. We denote the estimated prices in market t by $\mathbf{p}^t \in (\{1, \dots, m\} \times P)^m$, *i.e.*, p_j^t is the price of a good g_j in a market t . We denote an initial endowment and a final allocation in a clustered market by \mathbf{e}_i^t and \mathbf{x}_i^t . (Thus, $\sum_{t=1}^k \mathbf{e}_i^t = \mathbf{e}_i$.)

The algorithm considers each trader iteratively. At each iteration, it assigns the trader to a market and gives a near-optimal bundle to this trader, according to her utility function and the prices in that market. Once the i th trader is assigned, the algorithm only remembers the deficit (or surplus) of each good in every market – this will be a key in obtaining an efficient algorithm. Once every trader is assigned, it selects the best possible solution that satisfies the approximate market clearing constraint, *i.e.*, the deficits must be small. Thus, we encode the state of each market t by a vector \mathbf{y}^t , where y_j^t denotes the surplus (or deficit) of the good g_j . Let I_t denote the set of traders already in the market t , and $\mathbf{x}_i, i \in I_t$ the bundles given to these traders. Ideally, we would like to have $y_j^t = \sum_{i \in I_t} e_{ij} - \sum_{i \in I_t} x_{ij}$. Hence, the value of y_j^t could be any real value between -1 and 1 . However, as we cannot afford to store all possible values for y^t , we round these values into a set \widetilde{W} of cardinality $\alpha \cdot 4n$, where α will be set later. To define \widetilde{W} , we first define a coarser set W .

$$W = \left\{ \left(\frac{b}{b+1} \right)^\alpha, \left(\frac{b}{b+1} \right)^{\alpha-1}, \dots, \frac{b}{b+1}, 1 \right\}.$$

We then choose α to be minimal such that

$$\left(\frac{b}{b+1} \right)^\alpha \leq \frac{e_{\min}}{2n(b+1)} \quad \text{and} \quad \left(\frac{b}{b+1} \right)^\alpha < \frac{u_{\min} \cdot e_{\min}}{m \cdot u_{\max} \cdot p_{\max}} \cdot \frac{1}{(b+1)^2}$$

Observe that W induces a set of intervals $[(b/(b+1))^\ell, (b/(b+1))^{\ell-1}]$, for $\ell = 1, \dots, \alpha$. We can now create the set \widetilde{W} by dividing each interval of W and its negation into subintervals. Specifically, for each interval $[(b/(b+1))^\ell, (b/(b+1))^{\ell-1}]$ (resp., for each interval $[-(b/(b+1))^{\ell-1}, -(b/(b+1))^\ell]$), we divide W equally into $2n(b+1)$ subintervals and put the boundary points in \widetilde{W} . Thus,

$$\begin{aligned} \widetilde{W} = & \bigcup_{\ell=1}^{\alpha-1} \left\{ \left(\frac{b}{b+1} \right)^\ell \left(1 + \frac{q}{2nb(b+1)} \right) : q \in \{0, 1, \dots, 2n(b+1)\} \right\} \cup \{0\} \cup \\ & \bigcup_{\ell=1}^{\alpha-1} \left\{ \left(-\frac{b}{b+1} \right)^\ell \left(1 + \frac{q}{2nb(b+1)} \right) : q \in \{0, 1, \dots, 2n(b+1)\} \right\} \end{aligned}$$

We insist that the algorithm selects allocations of goods that take values from W , and we then round down the surplus (or deficits) to values in \widetilde{W} . Formally, $\mathbf{x}^t \subseteq W^{n \times m}$ and $\mathbf{y}^t \subseteq \widetilde{W}^m$ for any market t . Towards this goal, let $\lfloor a \rfloor_{\widetilde{W}}$ denote the value of a rounded down to the closest value in \widetilde{W} , *i.e.*, $a' = \lfloor a \rfloor_{\widetilde{W}}$ is the largest value in \widetilde{W} such that $a' \leq a$.

We now need to ensure that these allocations correspond to an approximate generalized equilibrium (which in turn will correspond to an approximate equilibrium). To do this, for a market t and a trader i , we say that an allocation x_i of goods is *compatible* with i and t if, for $r_i = \max_{j: e_{ij} > 0} \pi_1(p_j)$:

- $u_{ij} > 0 \implies \pi_1(p_j^t) \geq r_i$,
- $x_{ij} > 0 \implies \pi_1(p_j^t) \leq r_i$,
- $\sum_{j: \pi_1(p_j^t) \geq r_i} \pi_2(p_j^t) \cdot x_{ij} \leq \sum_{j: \pi_1(p_j^t) \geq r_i} \pi_2(p_j^t) \cdot e_{ij}$ (Budget constraint).

An allocation \mathbf{x}_i compatible with i and t is *nearly-optimal* if $\mathbf{u}_i \cdot \mathbf{x}_i \geq (1 - \varepsilon) \max_{\mathbf{z}} \mathbf{u}_i \cdot \mathbf{z}$ where \mathbf{z} ranges over all the allocations compatible with i and t . By this definition, assuming L is large enough, a nearly-optimal allocation is an approximate best response for trader i in market t . The recurrence relation of our dynamic programming algorithm is then

$$f(0, \mathbf{y}^1, \mathbf{y}^2, \dots, \mathbf{y}^k) = \begin{cases} 0 & \text{if } \mathbf{y}^1 = \mathbf{y}^2 = \dots = \mathbf{y}^k = \mathbf{0} \\ -\infty & \text{otherwise} \end{cases}$$

$$f(i, \mathbf{y}^1, \mathbf{y}^2, \dots, \mathbf{y}^k) = \max_{t \in \{1, 2, \dots, k\}, \mathbf{x}_i} f(i-1, \mathbf{y}^1, \dots, \lfloor \mathbf{y}^t - \mathbf{e}_i + \mathbf{x}_i \rfloor_{\widetilde{W}}, \dots, \mathbf{y}^k) + \mathbf{u}_i \mathbf{x}_i$$

where \mathbf{x}_i ranges over the nearly-optimal allocations compatible with i and t .

Finally, an allocation is *valid* if $0 \leq \mathbf{y}^t \leq 1 - (b/(b+1))^3$ for all t , i.e., there is no market having a positive deficit, or a surplus greater than $1 - (b/(b+1))^3$. Notice that some of the surplus may have been lost in the rounding steps; we will show later that these additional losses amount to a fraction at most $1/(b+1)$ of each good. Hence, if $f(n, \mathbf{y}^1, \dots, \mathbf{y}^t)$ is finite for some valid allocation, and provided that b is sufficiently large, it gives an approximate *generalized equilibrium*: the budget constraints and approximate optimality constraints are guaranteed by the restriction on the choice of \mathbf{x}_i at each step, and market clearing is guaranteed by the validity of the allocation.

This completes the description of the dynamic program. It remains to compare the total social welfare to that of the optimal clustering and to analyze the running time of the algorithm.

6.2 Quality Analysis

Because of the rounding step, our dynamic programming algorithm loses some fraction of each good g_j . We have to bound the number of units of the good g_j that we lose. By the scaling on \widetilde{W} , each time we round y_j^t , we lose at most $(2n(b+1))^{-1} \sum_{i \in I_t} e_{ij}$ units when y_j^t is rounded to a positive value, and we lose at most $(b/(b+1))^\alpha \leq e_{\min} (2n(b+1))^{-1}$ units when y_j^t is rounded to zero. By the definition of e_{\min} , we have $e_{\min} \leq \sum_{i \in I_t} e_{ij}$ for all goods g_j . Furthermore, we round down y_j^t at most n times, once for each trader. Thus, summing them up, we lose at most $1/(2(b+1)) \cdot \sum_{i \in I_t} e_{ij} < 1/(2b+2)$ units of each good g_j .

Now, we compare the social welfare of the approximate generalized equilibrium obtained by our algorithm with that of some exact generalized equilibrium.

Lemma 6. *For any market $M = (n, m, \mathbf{u}, \mathbf{e})$ and any $\varepsilon > 0$, there is an exact generalized equilibrium such that*

$$\frac{\max_j \pi_2(p_j)}{\min_{j: \pi_2(p_j) \neq 0} \pi_2(p_j)} \leq e^{nm/e} \left(\frac{p_{\max} \cdot m}{e_{\min} \cdot p_{\min}} \right)^{nm}$$

Proof. Take an exact generalized equilibrium \mathbf{x}, \mathbf{p} that maximizes the total welfare. Let $G_r := \{g_j : \pi_1(p_j) = r\}$, $I_r := \{i : r_i = r\}$ and $A_r := \{i \in I_r : \forall j \in G_r, u_{ij} = 0\}$. For each possible rank r , let M_r be the market obtained by restricting M to the set of goods G_r . Define $x_{ij}^r = x_{ij}$ and $p_j^r = \pi_2(p_j)$ for any trader i and good g_j in M_r . Fix some rank r .

Claim 13. *We may assume that $\sum_{j \in G_r} p_j^r x_{ij} = \sum_{j \in G_r} p_j^r e_{ij}$ for any i .*

Proof. By market-clearing for \mathbf{x}, \mathbf{p} , we have $\sum_i x_{ij} = \sum_i e_{ij}$ for any $j \in G_r$. By optimality, we also have $\sum_{j \in G_r} p_j^r e_{ij} = \sum_{j \in G_r} p_j^r x_{ij}$ for any $i \in I_r - A_r$. Subtracting the two equalities,

we get $\sum_{j \in G_r} (\sum_{i \in A_r} p_j^r e_{ij} + \sum_{i: r_i > r} p_j^r e_{ij}) = \sum_{j \in G_r} (\sum_{i \in A_r} p_j^r x_{ij} + \sum_{i: r_i > r} p_j^r x_{ij})$. Because $u_{ij}^r = 0$ for any i with $r_i > r$ or any $i \in A_r$, we can redistribute the goods of G_r allocated to these traders to fulfill the condition of the claim. \square

Then $\mathbf{x}^r, \mathbf{p}^r$ is an equilibrium in M_r because of the previous claim (the optimality constraints and market-clearing constraints follow from the definition of generalized equilibrium). Let r be some rank and consider M_r . We want to compute an approximate equilibrium for M_r such that we can bound the prices. If there is no altruist in M_r , then we apply Lemma 4. Otherwise, for any altruist i , let j_i be such that $p_{j_i} x_{ij_i}^r$ is maximized. Define the utility vector \mathbf{u}^r

$$u_{ij}^r = \begin{cases} u_{ij} & \text{if } i \text{ is not an altruist in } M_r, \\ 1 & \text{if } i \text{ is altruist in } M_r \text{ and } j = j_i, \\ 0 & \text{otherwise.} \end{cases}$$

Then the market defined by \mathbf{e} and $\mathbf{u}^r, \mathbf{x}^r, \mathbf{p}^r$ is an approximate solution, where only the altruists of M_r do not follow a best response, and the approximation ratio is at most $(m-1)/m$ by the choice of \mathbf{u}_i^r for altruist trader i (and because there are at most m goods in M_r). Hence, by applying Lemma 2, we get $\frac{p_{\max}}{p_{\min}} \leq e^{nm/e} \left(\frac{mu_{\max}}{u_{\min} e_{\min}} \right)^{nm}$. \square

Fix any optimal clustering and consider any market t . Take a generalized equilibrium $\mathbf{p}^*, \mathbf{x}^*$ as in Lemma 6. We show that the dynamic program outputs an approximate generalized equilibria with total welfare $(1 - \varepsilon)$ times the welfare of $\mathbf{p}^*, \mathbf{x}^*$. This is done by building from $\mathbf{p}^*, \mathbf{x}^*$ a set of generalized prices \mathbf{p}' and allocations \mathbf{x}' computable by the dynamic program.

We assume wlog that $p_{\min}^* := \min_{j: p_j^* > 0} \pi_2(p_j^*) = 1$. Denote $p_{\max} := \max_{j: p_j^* > 0} \pi_2(p_j^*)$. \mathbf{p}' is obtained by rounding down the second components of prices to values in P ; hence, $b/(b+1) \cdot \pi_2(\mathbf{p}^*) \leq \pi_2(\mathbf{p}') \leq \pi_2(\mathbf{p}^*)$. Consequently, we have for each trader i and $a \geq r_i$:

$$\frac{b}{b+1} \sum_{j: \pi_1(p'_j) = a} \pi_2(p'_j) x_{ij}^* \leq \frac{b}{b+1} \sum_{j: \pi_1(p'_j) = a} \pi_2(p_j^*) e_{ij} \leq \sum_{j: \pi_1(p'_j) = a} \pi_2(p'_j) e_{ij}$$

Hence, $b/(b+1) \cdot \mathbf{x}^*$ satisfies the budget constraint for the prices \mathbf{p}' .

Next, we have to modify $b/(b+1) \cdot \mathbf{x}^*$ further so that it satisfies the condition in our dynamic programming algorithm. Namely, we round down the coefficients of $b/(b+1) \cdot \mathbf{x}^*$ to W . This gives an allocation \mathbf{x}' with the properties:

$$\left(\frac{b}{b+1} \right)^2 x_{ij}^* \leq x'_{ij} \leq x_{ij}^* \quad \text{if } x_{ij}^* \geq \left(\frac{b}{b+1} \right)^\alpha, \text{ and } x'_{ij} = 0 \text{ otherwise.}$$

Clearly, \mathbf{x}' also satisfies the budget constraint inequalities. It remains to show that \mathbf{x}'_i is an almost optimal choice for trader i . To see this, consider any good g_j . We have

$$\left(\frac{b}{b+1} \right)^3 \pi_2(p_j^*) \cdot x_{ij}^* \leq \frac{b}{b+1} \pi_2(p_j^*) \cdot x'_{ij} \leq \pi_2(p'_j) \cdot x'_{ij} \leq \pi_2(p_j^*) \cdot x_{ij}^*$$

when $x_{ij}^* \geq \left(\frac{b}{b+1} \right)$ or $x_{ij}^* = 0$, otherwise $u_{ij} x'_{ij} = 0$ and $u_{ij} x_{ij}^* \leq \left(\frac{b}{b+1} \right)^\alpha u_{ij}$. We have to handle the latter special case, when $x'_{ij} = 0 < x_{ij}^*$. For this purpose, notice first that the welfare of a trader i with $\sum_{j: \pi_1(p_j^*) = r_i} u_{ij} > 0$ is lower bounded by $\frac{e_{\min} \cdot u_{\min}}{p_{\max}}$

as e_{\min} is the minimum possible budget for a trader (other traders have welfare 0). The maximum ratio utility per unit of price achievable is $\frac{u_{\max}}{p_{\min}}$. Thus, any allocation $x_{ij}^* \leq (b/(b+1))^\alpha$ contributes to the total welfare at most $\left(\frac{b}{b+1}\right)^\alpha \frac{u_{\max} p_{\max}}{u_{\min} e_{\min}}$, which is less than $(1/m)(1/(b+1))^2 \leq (1/m)(1 - (b/(b+1))^3)$ by the choice of α . Hence, over all goods, the fraction of welfare lost in rounding down \mathbf{x}^* is at most $1 - 2(b/(b+1))^3$. This bounds our approximation ratio to $1 - \varepsilon \leq 2\left(\frac{b}{b+1}\right)^3$, which is true for $b = \lceil 3/\varepsilon \rceil$.

With this choice, \mathbf{x}' and \mathbf{p}' satisfy the approximate market equilibrium constraints, thus the dynamic algorithm will find a solution with welfare at least $\mathbf{u} \cdot \mathbf{x}' \geq (b/(b+1))^2 \mathbf{u} \cdot \mathbf{x}^*$.

6.3 Running Time Analysis.

Now, consider the complexity of the dynamic program. It can be seen that the running of our algorithm for one set of prices is $O(nk\alpha^m(\alpha nb)^{km})$, and the running time for all possible price allocation is $(m\sigma)^{km}$. This implies the total running time of $O(n^{km+1}k\alpha^{(k+1)m}(mb\sigma)^{km})$. Thus, we have a $(1 - \varepsilon)$ -approximation algorithm with a running time of

$$O\left(n^{(3k+1)m+1}m^{(2k+1)m}k\left(1 + \frac{3}{\varepsilon}\right)^{3km+m}\left(\frac{1}{e} + \log\frac{m \cdot u_{\max}}{e_{\min}u_{\min}}\right)^{(k+1)m}\left(\log\frac{m \cdot u_{\max}}{u_{\min}e_{\min}}\right)^{km}\right)$$

The input of the k -market clustering problem in a standard binary representation has a size of $\Omega(n(\log(1/e_{\min}) + \log(u_{\max}/u_{\min})))$. Thus, the running time of our algorithm is polynomial in the size of the input when m and k are constant.

Acknowledgements. We thank Neil Olver for interesting discussions and the referees for helpful comments and suggestions.

References

- [1] K. Arrow and G. Debreu, "The existence of an equilibrium for a competitive economy", *Econometrica*, **22**, 265-290, 1954.
- [2] M. Blaug, "The fundamental theorems of modern welfare economics, historically contemplated", *History of Political Economy*, **39(2)**, pp185-207, 2007.
- [3] W. Baumol, *The Free Market Innovation Machine: Analyzing the Growth Miracle of Capitalism*, Princeton University Press, 2002.
- [4] M. Fafchamps, "Microeconomic theory", *Lecture Notes*, Oxford University, 2007.
- [5] F. Fisher, "Disequilibrium and stability", in *General Equilibrium: Problems and Prospects*, F. Petri and F. Hahn (eds.), Routledge, 2003.
- [6] D. Gale, "The linear exchange model", *Journal of Mathematical Economics*, **3(2)**, pp205-209, 1976.
- [7] K. Jain, M. Mahdian and A. Saberi, "Approximating market equilibria", *APPROX*, 2003.
- [8] K. Jain, "A Polynomial Time Algorithm for Computing an Arrow-Debreu Market Equilibrium for Linear Utilities", *SIAM Journal on Computing*, **37(1)**, pp303-318, 2007.
- [9] D. Luenberger, "Optimality and the theory of value", *Journal of Economic Theory*, **63(2)**, pp147-169, 1994.
- [10] J. Mirrlees, "The theory of optimal taxation", in *Handbook of Mathematical Economics*, K. Arrow and M. Intriligator (eds.), North-Holland, 1986.
- [11] G. Myles, *Public Economics*, Cambridge University Press, 1995.
- [12] R. Starr, *General Equilibrium Theory: An Introduction*, Cambridge University Press, 1997.
- [13] J. Stiglitz, "The invisible hand and modern welfare economics", *NBER Working Papers Series*, National Bureau of Economic Research, 1991.