of this polytope which is an optimal solution. Exploiting this, we have developed a simplex-like vertex-marching algorithm which runs in strongly polynomial time for many special cases.

We feel that the strongly polynomial algorithm by Orlin [11] is neither polytopal nor very intuitive. The algorithms, which are polytopal and simplex-like are generally easier to understand, simpler to implement using standard math libraries, and run faster in practice. Therefore, an obvious open problem is to give a strongly polynomial, simplex-like algorithm; even a polynomial bound will be interesting. Another open problem is to give a linear programming formulation that captures the equilibrium prices for the Fisher market. Therefore, it will be interesting to construct a linear cost function on our polytope so that optimum vertex gives the equilibrium prices.

References

1. Adsul, B., Babu, C.S., Garg, J., Mehta, R., Sohoni, M.: Nash equilibria in Fisher market. arXiv:1002.4832 (2010)
2. Birnbaum, B., Devanur, N.R., Xiao, L.: New convex programs and distributed algorithms for Fisher markets with linear and spending constraint utilities. Working Paper (2010)
3. Chakraborty, S., Devanur, N.R., Karande, C.: Market equilibrium with transaction costs. arXiv:1001.0393 (2010)
4. Deng, X., Papadimitriou, C.H., Safra, S.: On the complexity of equilibria. In: STOC 2002 (2002)
5. Devanur, N.R., Papadimitriou, C.H., Saberi, A., Vazirani, V.V.: Market equilibrium via a primal-dual type algorithm. JACM 55(5) (2008)
6. Eisenberg, E., Gale, D.: Consensus of subjective probabilities: The pari-mutuel method. Annals Math. Stat. 30, 165–168 (1959)
7. Garg, J.: Algorithms for market equilibrium. First APS Report (2008), `http://www.cse.iitb.ac.in/~jugal/aps1.pdf`
8. Garg, J.: Fisher markets with transportation costs. Technical Report (2009), `http://www.cse.iitb.ac.in/~jugal/techRep.pdf`
9. Jain, K.: A polynomial time algorithm for computing the Arrow-Debreu market equilibrium for linear utilities. In: FOCS 2004 (2004)
10. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press, Cambridge (2007)
11. Orlin, J.B.: Improved algorithms for computing Fisher's market clearing prices. In: STOC 2010 (2010)
12. Shmyrev, V.I.: An algorithm for finding equilibrium in the linear exchange model with fixed budgets. Journal of Applied and Industrial Mathematics 3(4), 505–518 (2009)

Nash Equilibria in Fisher Market

Bharat Adsul[1], Ch. Sobhan Babu[2], Jugal Garg[1],
Ruta Mehta[1], and Milind Sohoni[1]

[1] Indian Institute of Technology, Bombay
{adsul,jugal,ruta,sohoni}@cse.iitb.ac.in
[2] Indian Institute of Technology, Hyderabad
sobhan@iith.ac.in

Abstract. Much work has been done on the computation of market equilibria. However due to strategic play by buyers, it is not clear whether these are actually observed in the market. Motivated by the observation that a buyer may derive a better payoff by feigning a different utility function and thereby manipulating the Fisher market equilibrium, we formulate the *Fisher market game* in which buyers strategize by posing different utility functions. We show that existence of a *conflict-free allocation* is a necessary condition for the Nash equilibria (NE) and also sufficient for the symmetric NE in this game. There are many NE with very different payoffs, and the Fisher equilibrium payoff is captured at a symmetric NE. We provide a complete polyhedral characterization of all the NE for the two-buyer market game. Surprisingly, all the NE of this game turn out to be symmetric and the corresponding payoffs constitute a piecewise linear concave curve. We also study the correlated equilibria of this game and show that third-party mediation does not help to achieve a better payoff than NE payoffs.

1 Introduction

A fundamental market model was proposed by Walras in 1874 [21]. Independently, Fisher proposed a special case of this model in 1891 [3], where a market comprises of a set of buyers and divisible goods. The money possessed by buyers and the amount of each good is specified. The utility function of every buyer is also given. The market equilibrium problem is to compute prices and allocation such that every buyer gets the highest utility bundle subject to her budget constraint and that the market clears. Recently, much work has been done on the computation of market equilibrium prices and allocation for various utility functions, for example [6,7,11,15].

The payoff (*i.e.*, happiness) of a buyer depends on the equilibrium allocation and in turn on the utility functions and initial endowments of the buyers. A natural question to ask is, can a buyer achieve a better payoff by feigning a different utility function? It turns out that a buyer may indeed gain by feigning! This observation motivates us to analyze the strategic behavior of buyers in the Fisher market. We analyze here the linear utility case described below.

S. Kontogiannis, E. Koutsoupias, P.G. Spirakis (Eds.): SAGT 2010, LNCS 6386, pp. 30–41, 2010.

Let $\mathcal{B}$ be the set of buyers, and $\mathcal{G}$ be the set of goods, and $|\mathcal{B}| = m, |\mathcal{G}| = n$. Let m_i be the money possessed by buyer i, and q_j be the total quantity of good j in the market. The utility function of buyer i is represented by the non-negative utility tuple $\langle u_{i1}, \ldots, u_{in} \rangle$, where u_{ij} is the payoff, she derives from a unit amount of good j. Thus, if x_{ij} is the amount of good j allocated to buyer i, then the *payoff* she derives from her allocation is $\sum_{j \in \mathcal{G}} u_{ij} x_{ij}$. Market equilibrium or market clearing prices $(p_1, \ldots, p_n)$, where p_j is the price of good j, and equilibrium allocation $[x_{ij}]_{i \in \mathcal{B}, j \in \mathcal{G}}$ satisfy the following constraints:

- **Market Clearing:** The demand equals the supply of each good, *i.e.*, $\forall j \in \mathcal{G}$, $\sum_{i \in \mathcal{B}} x_{ij} = q_j$, and $\forall i \in \mathcal{B}$, $\sum_{j \in \mathcal{G}} p_j x_{ij} = m_i$.
- **Optimal Goods:** Every buyer buys only those goods, which give her the maximum utility per unit of money, *i.e.*, if $x_{ij} > 0$ then $\frac{u_{ij}}{p_j} = \max_{k \in \mathcal{G}} \frac{u_{ik}}{p_k}$.

In this market model, by scaling u_{ij}'s appropriately, we may assume that the quantity of every good is one unit, *i.e.*, $q_j = 1$, $\forall j \in \mathcal{G}$. Equilibrium prices are unique and the set of equilibrium allocations is a convex set [14]. The following example illustrates a small market.

Example 1. Consider a 2 buyers, 2 goods market with $m_1 = m_2 = 10$, $q_1 = q_2 = 1$, $\langle u_{11}, u_{12} \rangle = \langle 10, 3 \rangle$ and $\langle u_{21}, u_{22} \rangle = \langle 3, 10 \rangle$. The equilibrium prices of this market are $\langle p_1, p_2 \rangle = \langle 10, 10 \rangle$ and the unique equilibrium allocation is $\langle x_{11}, x_{12}, x_{21}, x_{22} \rangle = \langle 1, 0, 0, 1 \rangle$. The payoff of both the buyers is 10.

In the above market, does a buyer have a strategy to achieve a better payoff? Yes indeed, buyer 1 can force price change by posing a different utility tuple, and in turn gain. Suppose buyer 1 feigns her utility tuple as $\langle 5, 15 \rangle$ instead of $\langle 10, 3 \rangle$, then coincidentally, the equilibrium prices $\langle p_1, p_2 \rangle$ are also $\langle 5, 15 \rangle$. The unique equilibrium allocation $\langle x_{11}, x_{12}, x_{21}, x_{22} \rangle$ is $\langle 1, \frac{1}{3}, 0, \frac{2}{3} \rangle$. Now, the payoff of buyer 1 is $u_{11} * 1 + u_{12} * \frac{1}{3} = 11$, and that of buyer 2 is $u_{22} * \frac{2}{3} = \frac{20}{3}$. Note that the payoffs are still calculated w.r.t. the true utility tuples.

This clearly shows that a buyer could gain by feigning a different utility tuple, hence the Fisher market is susceptible to gaming by strategic buyers. Therefore, the equilibrium prices w.r.t. the true utility tuples may not be the actual operating point of the market. The natural questions to investigate are: What are the possible operating points of this market model under strategic behavior? Can they be computed? Is there a preferred one? This motivates us to study the Nash equilibria of the *Fisher market game*, where buyers are the players and strategies are the utility tuples that they may pose.

Related work. Shapley and Shubik [18] consider a market game for the exchange economy, where every good has a trading post, and the strategy of a buyer is to bid (money) at each trading post. For each strategy profile, the prices are determined naturally so that market clears and goods are allocated accordingly, however agents may not get the optimal bundles. Many variants [2,8] of this game have been extensively studied. Essentially, the goal is to design a mechanism to implement Walrasian equilibrium (WE), *i.e.*, to capture WE at a

NE of the game. The strategy space of this game is tied to the implementation of the market (in this case, trading posts). Our strategy space is the utility tuple itself, and is independent of the market implementation. It is not clear that bids of a buyer in the Shapley-Shubik game correspond to the feigned utility tuples.

In word auction markets as well, a similar study on strategic behavior of buyers (advertisers) has been done [4,9,19].

Our contributions. We formulate the Fisher market game, the strategy sets and the corresponding payoff function in Section 2. Every (pure) strategy profile defines a Fisher market, and therefore market equilibrium prices and a set of equilibrium allocations. The payoff of a buyer may not be same across all equilibrium allocations w.r.t. a strategy profile, as illustrated by Example 2 in Section 2. Furthermore, there may not exist an equilibrium allocation, which gives the maximum possible payoffs to all the buyers. This behavior causes a conflict of interest among buyers. A strategy profile is said to be *conflict-free*, if there is an equilibrium allocation which gives the maximum possible payoffs to all the buyers.

A strategy profile is called a Nash equilibrium strategy profile *(NESP)*, if no buyer can unilaterally deviate and get a better payoff. In Section 3, we show that all NESPs are conflict-free. Using the equilibrium prices, we associate a bipartite graph to a strategy profile and show that this graph must satisfy certain conditions when the corresponding strategy profile is a NESP.

Next, we define *symmetric* strategy profiles, where all buyers play the same strategy. We show that a symmetric strategy profile is a NESP iff it is conflict-free. It is interesting to note that a symmetric NESP can be constructed for a given market game, whose payoff is the same as the Fisher payoff, *i.e.*, payoff when all buyers play truthfully. Example 11 shows that all NESPs need not be symmetric and the payoff w.r.t. a NESP need not be Pareto optimal (*i.e.*, efficient). However, the Fisher payoff is always Pareto optimal (see First Theorem of Welfare Economics [20]).

Characterization of all the NESPs seems difficult; even for markets with only three buyers. We study two-buyer markets in Section 4 and the main results are:

- All NESPs are symmetric and they are a union of at most $2n$ convex sets.
- The set of NESP payoffs constitute a piecewise linear concave curve and all these payoffs are Pareto optimal. The strategizing on utilities has the same effect as differing initial endowments (see Second Theorem of Welfare Economics [20]).
- The third-party mediation does not help in this game.

Some interesting observations about two-buyer markets are:

- The buyer i gets the maximum payoff among all Nash equilibrium payoffs when she imitates the other, *i.e.*, when they play $(\boldsymbol{u}_{-i}, \boldsymbol{u}_{-i})$, where $\boldsymbol{u}_{-i}$ is the true utility tuple of the other buyer.
- There may exist NESPs, whose social welfare (*i.e.*, sum of the payoffs of both the buyers) is larger than that of the Fisher payoff (Example 17).

– For a particular payoff tuple, there is a convex set of NESPs and hence convex set of equilibrium prices. This motivates a seller to offer incentives to the buyers to choose a particular NESP from this convex set, which fetches the maximum price for her good. Example 18 illustrates this behavior.

Most qualitative features of these markets may carry over to oligopolies, which arise in numerous scenarios. For example, relationship between a few manufacturers of aircrafts or automobiles and many suppliers. Finally, we conclude in Section 5 that it is highly unlikely that buyers will act according to their true utility tuples in Fisher markets and discuss some directions for further research.

2 The Fisher Market Game

As defined in the previous section, a linear Fisher market is defined by the tuple $(\mathcal{B}, \mathcal{G}, (\boldsymbol{u_i})_{i\in\mathcal{B}}, \boldsymbol{m})$, where $\mathcal{B}$ is a set of buyers, $\mathcal{G}$ is a set of goods, $\boldsymbol{u_i} = (u_{ij})_{j\in\mathcal{G}}$ is the true utility tuple of buyer i, and $\boldsymbol{m} = (m_i)_{i\in\mathcal{B}}$ is the endowment vector. We assume that $|\mathcal{B}| = m, |\mathcal{G}| = n$ and the quantity of every good is one unit.

The *Fisher market game* is a one-shot non-cooperative game, where the buyers are the players, and the *strategy set* is all possible utility tuples that they may pose, *i.e.*, $\mathbb{S}_i = \{\langle s_{i1}, s_{i2}, \ldots, s_{in}\rangle \mid s_{ij} \geq 0, \sum_{j\in\mathcal{G}} s_{ij} \neq 0\}$, $\forall i \in \mathcal{B}$. Clearly, the set of all strategy profiles is $\mathbb{S} = \mathbb{S}_1 \times \cdots \times \mathbb{S}_m$. When a strategy profile $S = (\boldsymbol{s_1}, \ldots, \boldsymbol{s_m})$ is played, where $\boldsymbol{s_i} \in \mathbb{S}_i$, we treat $\boldsymbol{s_1}, \ldots, \boldsymbol{s_m}$ as utility tuples of buyers $1, \ldots, m$ respectively, and compute the equilibrium prices and a set of equilibrium allocations w.r.t. S and $\boldsymbol{m}$.

Further, using the equilibrium prices $(p_1, \ldots, p_n)$, we generate the corresponding *solution graph* G as follows: Let $V(G) = \mathcal{B} \cup \mathcal{G}$. Let b_i be the node corresponding to the buyer i, $\forall i \in \mathcal{B}$ and g_j be the node corresponding to the good j, $\forall j \in \mathcal{G}$ in G. We place an edge between b_i and g_j iff $\frac{s_{ij}}{p_j} = \max_{k\in\mathcal{G}} \frac{s_{ik}}{p_k}$, and call the edges of the solution graph as *tight edges*. Note that when the solution graph is a forest, there is exactly one equilibrium allocation, however this is not so, when it contains cycles. In the standard Fisher market (*i.e.*, strategy of every buyer is her true utility tuple), all equilibrium allocations give the same payoff to a buyer. However, this is not so when buyers strategize on their utility tuples: Different equilibrium allocations may not give the same payoff to a buyer. The following example illustrates this scenario.

Example 2. Consider the Fisher market of Example 1. Consider the strategy profile $S = (\langle 1, 19\rangle, \langle 1, 19\rangle)$. Then, the equilibrium prices $\langle p_1, p_2\rangle$ are $\langle 1, 19\rangle$ and the solution graph is a cycle. There are many equilibrium allocations and the allocations $[x_{11}, x_{12}, x_{13}, x_{14}]$ achieving the highest payoff for buyers 1 and 2 are $[1, \frac{9}{19}, 0, \frac{10}{19}]$ and $[0, \frac{10}{19}, 1, \frac{9}{19}]$ respectively. The payoffs corresponding to these allocations are $(11.42, 5.26)$ and $(1.58, 7.74)$ respectively. Note that there is no allocation, which gives the maximum possible payoff to both the buyers.

Let $\boldsymbol{p}(S) = (p_1, \ldots, p_n)$ be the equilibrium prices, $G(S)$ be the solution graph, and $\mathbb{X}(S)$ be the set of equilibrium allocations w.r.t. a strategy profile S.

The payoff w.r.t. $X \in \mathbb{X}(S)$ is defined as $(u_1(X), \ldots, u_m(X))$, where $u_i(X) = \sum_{j \in \mathcal{G}} u_{ij} x_{ij}$. Let $w_i(S) = \max_{X \in \mathbb{X}(S)} u_i(X), \forall i \in \mathcal{B}$.

Definition 3. *A strategy profile S is said to be* ***conflict-free*** *if $\exists X \in \mathbb{X}(S)$, s.t. $u_i(X) = w_i(S)$, $\forall i \in \mathcal{B}$. Such an X is called a* ***conflict-free allocation****.*

When a strategy profile $S = (\boldsymbol{s_1}, \ldots, \boldsymbol{s_m})$ is not conflict-free, there is a conflict of interest in selecting a particular allocation for the play. If a buyer, say k, does not get the same payoff from all the equilibrium allocations, *i.e.*, $\exists X \in \mathbb{X}(S)$, $u_k(X) < w_k(S)$, then we show that for every $\delta > 0$, there exists a strategy profile $S' = (\boldsymbol{s'_1}, \ldots, \boldsymbol{s'_m})$, where $\boldsymbol{s'_i} = \boldsymbol{s_i}$, $\forall i \neq k$, such that $u_k(X') > w_k(S) - \delta, \forall X' \in \mathbb{X}(S')$ (Section 3.1). The following example illustrates the same.

Example 4. In Example 2, for $\delta = 0.1$, consider $S' = (\langle 1.1, 18.9\rangle, \langle 1, 19\rangle)$, *i.e.*, buyer 1 deviates slightly from S. Then, $\boldsymbol{p}(S') = \langle 1.1, 18.9\rangle$, and $G(S')$ is a tree; the cycle of Example 2 is broken. Hence there is a unique equilibrium allocation, and $w_1(S') = 11.41,\ w_2(S') = 5.29$.

Therefore, if a strategy profile S is not conflict-free, then for every choice of allocation $X \in \mathbb{X}(S)$ to decide the payoff, there is a buyer who may deviate and assure herself a better payoff. In other words, when S is not conflict-free, there is no way to choose an allocation X from $\mathbb{X}(S)$ acceptable to all the buyers. This suggests that only conflict-free strategies are interesting. Therefore, we may define the payoff function $\mathcal{P}_i : \mathbb{S} \to \mathbb{R}$ for each player $i \in \mathcal{B}$ as follows:

$$\forall S \in \mathbb{S},\ \mathcal{P}_i(S) = u_i(X),\ \text{where } X = \arg\max_{X' \in \mathbb{X}(S)} \prod_{i \in \mathcal{B}} u_i(X'). \tag{1}$$

Note that the payoff functions are well-defined and when S is conflict-free, $\mathcal{P}_i(S) = w_i(S),\ \forall i \in \mathcal{B}$.

3 Nash Equilibria: A Characterization

In this section, we prove some necessary conditions for a strategy profile to be a NESP of the Fisher market game defined in the previous section. Nash equilibrium [13] is a solution concept for games with two or more rational players. When a strategy profile is a NESP, no player benefits by changing her strategy unilaterally.

For technical convenience, we assume that $u_{ij} > 0$ and $s_{ij} > 0$, $\forall i \in \mathcal{B}, \forall j \in \mathcal{G}$. The boundary cases may be easily handled separately. Note that if $S = (\boldsymbol{s_1}, \ldots, \boldsymbol{s_m})$ is a NESP then $S' = (\alpha_1 \boldsymbol{s_1}, \ldots, \alpha_m \boldsymbol{s_m})$, where $\alpha_1, \ldots, \alpha_m > 0$, is also a NESP. Therefore, w.l.o.g. we consider only the normalized strategies $\boldsymbol{s_i} = \langle s_{i1}, \ldots, s_{in}\rangle$, where $\sum_{j \in \mathcal{G}} s_{ij} = 1$[1], $\forall i \in \mathcal{B}$. As mentioned in the previous section, the true utility tuple of buyer i is $\langle u_{i1}, \ldots, u_{in}\rangle$. For convenience, we may assume that $\sum_{j \in \mathcal{G}} u_{ij} = 1$ and $\sum_{i \in \mathcal{B}} m_i = 1$ (w.l.o.g.).

[1] For simplicity, we do use non-normalized strategy profiles in the examples.

We show that all NESPs are conflict-free. However, not all conflict-free strategies are NESPs. A symmetric strategy profile, where all players play the same strategy (*i.e.*, $\forall i, j \in \mathcal{B},\ \boldsymbol{s}_i = \boldsymbol{s}_j$), is a NESP iff it is conflict-free. If a strategy profile S is not conflict-free, then there is a buyer a such that $\mathcal{P}_a(S) < w_a(S)$. The ConflictRemoval procedure in the next section describes how she may deviate and assure herself payoff almost equal to $w_a(S)$.

3.1 Conflict Removal Procedure

Definition 5. *Let S be a strategy profile, $X \in \mathbb{X}(S)$ be an allocation, and $P = v_1, v_2, v_3, \ldots$ be a path in $G(S)$. P is called an **alternating path** w.r.t. X, if the allocation on the edges at odd positions is non-zero, i.e., $x_{v_{2i-1}v_{2i}} > 0, \forall i \geq 1$. The edges with non-zero allocation are called **non-zero edges**.*

Table 1. Conflict Removal Procedure

```
ConflictRemoval(S, b_a, δ)
    while b_a belongs to a cycle in G(S) do
        (p_1, ..., p_n) ← p(S);
        J ← {j ∈ 𝒢 | the edge (b_a, g_j) belongs to a cycle in G(S)};
        g_b ← arg max_{j∈J} u_{aj}/p_j;
        X ← an allocation in 𝕏(S) such that u_a(X) = w_a(S) and x_{ab} is maximum;
        S ← Perturbation(S, X, b_a, g_b, δ/n);
    endwhile
    return S;

Perturbation(S, X, b_a, g_b, γ)
    S' ← S;
    if (b_a, g_b) does not belong to a cycle in G(S) then
        return S';
    endif
    J_1 ← {v | there is an alternating path from b_a to v in G(S) \ (b_a, g_b) w.r.t. X};
    J_2 ← {v | there is an alternating path from g_b to v in G(S) \ (b_a, g_b) w.r.t. X};
    (p_1, ..., p_n) ← p(S);   l ← Σ_{g_j∈J_1} p_j;   r ← Σ_{g_j∈J_2} p_j;
    W.r.t. α, define prices of goods to be
        ∀g_j ∈ J_1 : (1 − α)p_j;   ∀g_j ∈ J_2 : (1 + lα/r)p_j;   ∀g_j ∈ 𝒢 \ (J_1 ∪ J_2) : p_j;
    Raise α infinitesimally starting from 0 such that none of the three events occur:
        Event 1: a new edge becomes tight;
        Event 2: a non-zero edge becomes zero;
        Event 3: payoff of buyer a becomes u_a(X) − γ;
    s'_{ab} ← s_{ab} (1 + lα/r)/(1 − α);   s'_a ← s'_a / Σ_{j∈𝒢} s'_{aj};
    return S';
```

The ConflictRemoval procedure in Table 1 takes a strategy profile S, a buyer a and a positive number δ, and outputs another strategy profile S', where $\boldsymbol{s'_i} = \boldsymbol{s_i},\ \forall i \neq a$ such that $\forall X' \in \mathbb{X}(S'),\ u_a(X') > w_a(S) - \delta$. The idea is that if a buyer, say a, does not belong to any cycle in the solution graph of a strategy profile S, then $u_a(X) = w_a(S)$, $\forall X \in \mathbb{X}(S)$. The procedure essentially breaks all the cycles containing b_a in $G(S)$ using the Perturbation procedure iteratively such that the payoff of buyer a does not decrease by more than δ.

The Perturbation procedure takes a strategy profile S, a buyer a, a good b, an allocation $X \in \mathbb{X}(S)$, where x_{ab} is maximum among all allocations in $\mathbb{X}(S)$ and a positive number γ, and outputs another strategy profile S' such that $\boldsymbol{s'_i} = \boldsymbol{s_i}$, $\forall i \neq a$ and $w_a(S') > u_a(X) - \gamma$. It essentially breaks all the cycles containing the edge (b_a, g_b) in $G(S)$.

A detailed explanation of both the procedures is given in [1]. In the next theorem, we use the ConflictRemoval procedure to show that all the NESPs in the Fisher market game are conflict-free.

Theorem 6. *If S is a NESP, then*

(i) $\exists X \in \mathbb{X}(S)$ *such that* $u_i(X) = w_i(S), \forall i \in \mathcal{B}$, *i.e., S is conflict-free.*
(ii) the degree of every good in $G(S)$ is at least 2.
(iii) for every buyer $i \in \mathcal{B}$, $\exists k_i \in K_i$ *s.t.* $x_{ik_i} > 0$, *where* $K_i = \{j \in \mathcal{G} \mid \frac{u_{ij}}{p_j} = \max_{k\in\mathcal{G}} \frac{u_{ik}}{p_k}\}$, $(p_1, \ldots, p_n) = \boldsymbol{p}(S)$ *and* $[x_{ij}]$ *is a conflict-free allocation.*

Proof. Suppose there does not exist an allocation $X \in \mathbb{X}(S)$ such that $u_i(X) = w_i(S)$, $\forall i \in \mathcal{B}$, then there is a buyer $k \in \mathcal{B}$, such that $\mathcal{P}_k(S) < w_k(S)$. Clearly, buyer k has a deviating strategy (apply ConflictRemoval on the input tuple (S, k, δ), where $0 < \delta < (w_k(S) - \mathcal{P}_k(S))$), which is a contradiction.

For part (ii), if a good b is connected to exactly one buyer, say a, in $G(S)$, then buyer a may gain by reducing s_{ab}, so that price of good b decreases and prices of all other goods increase by the same factor.

For part (iii), if there exists a buyer i such that $x_{ik_i} = 0$, $\forall k_i \in K_i$, then she may gain by increasing the utility for a good in K_i. ⊓⊔

The following example shows that the above conditions are not sufficient.

Example 7. Consider a market with 3 buyers and 2 goods, where $\boldsymbol{m} = \langle 50, 100, 50\rangle$, $\boldsymbol{u_1} = \langle 2, 0.1\rangle, \boldsymbol{u_2} = \langle 4, 9\rangle$, and $\boldsymbol{u_3} = \langle 0.1, 2\rangle$. Consider the strategy profile $S = (\boldsymbol{u_1}, \boldsymbol{u_2}, \boldsymbol{u_3})$ given by the true utility tuples. The payoff tuple w.r.t. S is $(1.63, 6.5, 0.72)$. It satisfies all the necessary conditions in the above theorem, however S is not a NESP because buyer 2 has a deviating strategy $\boldsymbol{s'_2} = \langle 2, 3\rangle$ and the payoff w.r.t. strategy profile $(\boldsymbol{s_1}, \boldsymbol{s'_2}, \boldsymbol{s_3})$ is $(1.25, 6.75, 0.83)$.

3.2 Symmetric and Asymmetric NESPs

Recall that a strategy profile $S = (\boldsymbol{s_1}, \ldots, \boldsymbol{s_m})$ is said to be a *symmetric* strategy profile if $\boldsymbol{s_1} = \cdots = \boldsymbol{s_m}$, *i.e.*, all buyers play the same strategy.

Proposition 8. *A symmetric strategy profile S is a NESP iff it is conflict-free.*

Proof. ($\Rightarrow$) is easy (Theorem 6). For ($\Leftarrow$), suppose a buyer i may deviate and gain, then the prices have to be changed. In that case, all buyers except buyer i will be connected to only those goods, whose prices are decreased. This leads to a contradiction (refer to [1] for details). ⊓⊔

Let $S^f = [s_{ij}]$ be a strategy profile, where $s_{ij} = u_{ij}, \forall i \in \mathcal{B}, \forall j \in \mathcal{G}$, *i.e.*, true utility functions. All allocations in $\mathbb{X}(S^f)$ give the same payoff to the buyers (*i.e.*, $\forall i \in \mathcal{B}, u_i(X) = w_i(S^f)$, $\forall X \in \mathbb{X}(S^f)$), and we define *Fisher payoff* $(u_1^f, \ldots, u_m^f)$ to be the payoff derived when all buyers play truthfully.

Corollary 9. *A symmetric NESP can be constructed, whose payoff is the same as the Fisher payoff.*

Proof. Let $S = (\boldsymbol{s}, \ldots, \boldsymbol{s})$ be a strategy profile, where $\boldsymbol{s} = \boldsymbol{p}(S^f)$. Clearly S is a symmetric NESP, whose payoff is the same as the Fisher payoff (refer to [1] for details). □

Remark 10. The payoff w.r.t. a symmetric NESP is always Pareto optimal. For a Fisher market game, there is exactly one symmetric NESP iff the degree of every good in $G(S^f)$ is at least two [10].

The characterization of all the NESPs for the general market game seems hard; even for markets with only three buyers. The following example illustrates an asymmetric NESP, whose payoff is not Pareto optimal.

Example 11. Consider a market with 3 buyers and 2 goods, where $\boldsymbol{m} = \langle 50, 100, 50\rangle$, $\boldsymbol{u_1} = \langle 2, 3\rangle, \boldsymbol{u_2} = \langle 4, 9\rangle$, and $\boldsymbol{u_3} = \langle 2, 3\rangle$. Consider the two strategy profiles given by $S_1 = (\boldsymbol{s_1}, \boldsymbol{s_2}, \boldsymbol{s_3})$ and $S_2 = (\boldsymbol{s}, \boldsymbol{s}, \boldsymbol{s})$, where $\boldsymbol{s_1} = \langle 2, 0.1\rangle, \boldsymbol{s_2} = \langle 2, 3\rangle, \boldsymbol{s_3} = \langle 0.1, 3\rangle$, and $\boldsymbol{s} = \langle 2, 3\rangle$. The payoff tuples w.r.t. S_1 and S_2 are $(1.25, 6.75, 1.25)$ and $(1.25, 7.5, 1.25)$ respectively. Note that both S_1 and S_2 are NESPs for the above market (refer to [1] for details).

4 The Two-Buyer Markets

A two-buyer market consists of two buyers and a number of goods. These markets arise in numerous scenarios. The two firms in a duopoly may be considered as the two buyers with a similar requirements to fulfill from a large number of suppliers, for example, relationship between two big automotive companies with their suppliers.

In this section, we study two-buyer market game and provide a complete polyhedral characterization of NESPs, all of which turn out to be symmetric. Next, we study how the payoffs of the two buyers change with varying NESPs and show that these payoffs constitute a piecewise linear concave curve. For a particular payoff tuple on this curve, there is a convex set of NESPs, hence a convex set of equilibrium prices, which leads to a different class of non-market behavior such as incentives. Finally, we study the correlated equilibria of this game and show that third-party mediation does not help to achieve better payoffs than any of the NESPs.

Lemma 12. *All NESPs for a two-buyer market game are symmetric.*

Proof. If a NESP $S = (\boldsymbol{s_1}, \boldsymbol{s_2})$ is not symmetric, then $G(S)$ is not a complete bipartite graph. Therefore there is a good, which is exclusively bought by a buyer, which is a contradiction (Theorem 6, part (*ii*)). □

4.1 Polyhedral Characterization of NESPs

In this section, we compute all the NESPs of a Fisher market game with two buyers. Henceforth we assume that the goods are so ordered that $\frac{u_{1j}}{u_{2j}} \geq \frac{u_{1(j+1)}}{u_{2(j+1)}}$, for $j = 1, \ldots, n-1$. Chakrabarty et al. [5] also use such an ordering to design an algorithm for the linear Fisher market with two agents. Let $S = (\boldsymbol{s}, \boldsymbol{s})$ be a NESP, where $\boldsymbol{s} = (s_1, \ldots, s_n)$ and $(p_1, \ldots, p_n) = \boldsymbol{p}(S)$. The graph $G(S)$ is a complete bipartite graph. Since $m_1 + m_2 = 1$ and $\sum_{j=1}^{n} s_j = 1$, we have $p_j = s_j, \forall j \in \mathcal{G}$. In a conflict-free allocation $X \in \mathbb{X}(S)$, if $x_{1i} > 0$ and $x_{2j} > 0$, then clearly $\frac{u_{1i}}{p_i} \geq \frac{u_{1j}}{p_j}$ and $\frac{u_{2i}}{p_i} \leq \frac{u_{2j}}{p_j}$.

Definition 13. *An allocation $X = [x_{ij}]$ is said to be a* ***nice allocation****, if it satisfies the property: $x_{1i} > 0$ and $x_{2j} > 0 \;\Rightarrow\; i \leq j$.*

The main property of a *nice allocation* is that if we consider the goods in order, then from left to right, goods get allocated first to buyer 1 and then to buyer 2 exclusively, however they may share at most one good in between. Note that a symmetric strategy profile has a unique nice allocation.

Lemma 14. *Every NESP has a unique conflict-free nice allocation.*

Proof. The idea is to convert a conflict-free allocation into a nice allocation through an exchange s.t. payoff remains same (refer to [1] for details). □

The non-zero edges in a nice allocation either form a tree or a forest containing two trees. We use the properties of nice allocations and NESPs to give the polyhedral characterization of all the NESPs. The convex sets B_k for all $1 \leq k \leq n$, as given in Table 2, correspond to all possible conflict-free nice allocations, where non-zero edges form a tree, and the convex sets B'_k for all $1 \leq k \leq n-1$, as given in Table 3, correspond to all possible conflict-free nice allocations, where non-zero edges form a forest[2]. Let $\mathbb{B} = \cup_{k=1}^{n} B_k \cup_{k=1}^{n-1} B'_k$ and $S^{NE} = \{(\boldsymbol{\alpha}, \boldsymbol{\alpha}) \mid \boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n) \in \mathbb{B}\}$. Note that S^{NE} is a connected set.

Table 2. B_k

$\sum_{i=1}^{k-1} \alpha_i < m_1$	
$\sum_{i=k+1}^{n} \alpha_i < m_2$	
$\sum_{i=1}^{n} \alpha_i = m_1 + m_2$	
$u_{1j}\alpha_i - u_{1i}\alpha_j \leq 0$	$\forall i \leq k, \forall j \geq k$
$u_{2i}\alpha_j - u_{2j}\alpha_i \leq 0$	$\forall i \leq k, \forall j \geq k$
$\alpha_i \geq 0$	$\forall i \in \mathcal{G}$

Table 3. B'_k

$\sum_{i=1}^{k} \alpha_i = m_1$	
$\sum_{i=k+1}^{n} \alpha_i = m_2$	
$u_{1j}\alpha_i - u_{1i}\alpha_j \leq 0$	$\forall i \leq k, \forall j \geq k+1$
$u_{2i}\alpha_j - u_{2j}\alpha_i \leq 0$	$\forall i \leq k, \forall j \geq k+1$
$\alpha_i \geq 0$	$\forall i \in \mathcal{G}$

Lemma 15. *A strategy profile S is a NESP iff $S \in S^{NE}$.*

Proof. ($\Leftarrow$) is easy by the construction and Proposition 8. For the other direction, we know that every NESP has a conflict-free nice allocation (Lemma 14), and $\mathbb{B}$ corresponds to all possible conflict-free nice allocations. □

[2] In both the tables α_i's may be treated as price variables.

4.2 The Payoff Curve

In this section, we consider the payoffs obtained by both the players at various NESPs. Recall that whenever a strategy profile S is a NESP, $\mathcal{P}_i(S) = w_i(S)$, $\forall i \in \mathcal{B}$. Henceforth, we use $w_i(S)$ as the payoff of buyer i for the NESP S. Let $\mathbb{F} = \{(w_1(S), w_2(S)) \mid S \in S^{NE}\}$ be the set of all possible NESP payoff tuples.

Let $\mathcal{X}$ be the set of all nice allocations, and $\mathbb{H} = \{(u_1(X), u_2(X)) \mid X \in \mathcal{X}\}$. For $\alpha \in [0,1]$, let $t(\alpha) = (\langle s_1, \ldots, s_n\rangle, \langle s_1, \ldots, s_n\rangle)$, where $s_i = u_{1i} + \alpha(u_{2i} - u_{1i})$, and $\mathbb{G} = \{(w_1(S), w_2(S)) \mid S = t(\alpha), \alpha \in [0,1]\}$.

Proposition 16. $\mathbb{F}$ *is a piecewise linear concave (PLC) curve.*

Proof. The proof is based on the following steps (refer to [1] for details).

1. $\mathbb{H}$ is a PLC curve with $(0,1)$ and $(1,0)$ as the end points.
2. $\forall \alpha \in [0,1]$, $t(\alpha) \in S^{NE}$, then clearly $\mathbb{G} \subset \mathbb{H}$. Since the nice allocation w.r.t. $t(\alpha)$ changes continuously as α moves from 0 to 1, so we may conclude that $\mathbb{G}$ is a PLC curve with the end points $(w_1(S^1), w_2(S^1))$ and $(w_1(S^2), w_2(S^2))$, where $S^1 = t(0)$ and $S^2 = t(1)$.
3. $\mathbb{F} = \mathbb{G}$. □

The next example demonstrates the payoff curve for a small market game.

Example 17. Consider a market with 3 goods and 2 buyers, where $\boldsymbol{m} = \langle 7, 3\rangle$, $\boldsymbol{u_1} = \langle 6, 2, 2\rangle$, and $\boldsymbol{u_2} = \langle 0.5, 2.5, 7\rangle$. The payoff curve for this game is shown in the figure. The first and the second line segment of the curve correspond to the sharing of good 2 and 3 respectively. The payoffs corresponding to the boundary NESPs $S^1 = t(0)$ and $S^2 = t(1)$ are $(7, 8.25)$ and $(9.14, 3)$ respectively.

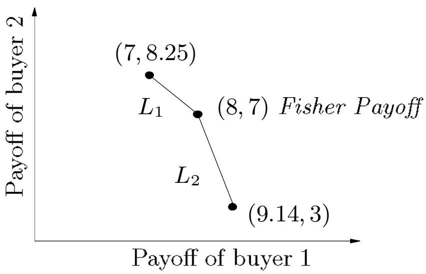

Furthermore, the *Fisher payoff* $(8, 7)$ may be achieved by a NESP $t(0.2)$. Note that in this example the social welfare (*i.e.*, sum of the payoffs of both the buyers) from the Fisher payoff (15) is lower than that of the NESP S^1 (15.25).

4.3 Incentives

For a fixed payoff tuple on the curve $\mathbb{F}$, there is a convex set of NESPs and hence a convex set of prices, giving the same payoffs to the buyers, and these may be computed using the similar inequalities as defined in Tables 2 and 3. This leads to a different class of behavior, *i.e.*, motivation for a seller to offer incentives to the buyers to choose a particular NESP from this convex set, which fetches the maximum price for her good. The following example illustrates this possibility.

Example 18. Consider a market with 2 buyers and 4 goods, where $\boldsymbol{m} = \langle 10, 10\rangle$, $\boldsymbol{u_1} = \langle 4, 3, 2, 1\rangle$, and $\boldsymbol{u_2} = \langle 1, 2, 3, 4\rangle$. Consider the two NESPs given by $S_1 = (\boldsymbol{s_1}, \boldsymbol{s_1})$ and $S_2 = (\boldsymbol{s_2}, \boldsymbol{s_2})$, where $\boldsymbol{s_1} = \langle \frac{20}{3}, \frac{20}{3}, \frac{10}{3}, \frac{10}{3}\rangle$ and $\boldsymbol{s_2} = \langle \frac{20}{3}, \frac{20}{3}, \frac{9}{3}, \frac{11}{3}\rangle$. Both S_1 and S_2 gives the payoff $(5.5, 8)$, however the prices are different, *i.e.*, $\boldsymbol{p}(S_1) = \langle \frac{20}{3}, \frac{20}{3}, \frac{10}{3}, \frac{10}{3}\rangle$ and $\boldsymbol{p}(S_2) = \langle \frac{20}{3}, \frac{20}{3}, \frac{9}{3}, \frac{11}{3}\rangle$. Clearly in S_2, good 3 is penalized and good 4 is rewarded (compared to S_1).

4.4 Correlated Equilibria

We have seen in Section 4.2 that the two-buyer market game has a continuum of Nash Equilibria, with very different and conflicting payoffs. This makes it difficult to predict how a particular game will actually play out in practice, and if there is a different solution concept which may yield an outcome liked by both the players.

We examine the correlated equilibria framework as a possibility. Recall that according to the correlated equilibria, the mediator decides and declares a probability distribution π on all possible pure strategy profiles $(\boldsymbol{s_1}, \boldsymbol{s_2}) \in \mathbb{S}_1 \times \mathbb{S}_2$ beforehand. During the play, she suggests what strategy to play to each player privately, and no player benefits by deviating from the advised strategy. The question we ask: Is there a correlated equilibrium π such that the payoff w.r.t. π lies above the curve $\mathbb{H}$? We continue with our assumption that $\frac{u_{1j}}{u_{2j}} \geq \frac{u_{1(j+1)}}{u_{2(j+1)}}, \forall j < n$.

Lemma 19. *For any strategy profile $S = (\boldsymbol{s_1}, \boldsymbol{s_2})$, for every allocation $X \in \mathbb{X}(S)$, there exists a point (x_1, x_2) on $\mathbb{H}$ such that $x_1 \geq u_1(X)$ and $x_2 \geq u_2(X)$.*

Proof. Any allocation X may be converted to a nice allocation through an exchange such that no buyer worse off (refer to [1] for details). □

Corollary 20. *The correlated equilibrium does not give better payoff than any NE payoff to all the buyers.*

Remark 21. [10] extends this result for the general Fisher market game.

5 Conclusion

The main conclusion of the paper is that Fisher markets in practice will rarely be played with true utility functions. In fact, the utilities employed will usually be a mixture of a player's own utilities and her conjecture on the other player's true utilities. Moreover, there seems to be no third-party mediation which will induce players to play according to their true utilities so that the true Fisher market equilibrium may be observed. Further, any notion of market equilibrium should examine this aspect of players strategizing on their utilities. This poses two questions: (i) is there a mechanism which will induce players into revealing their true utilities? and (ii) how does this mechanism reconcile with the "invisible hand" of the market? The strategic behavior of agents and the question whether true preferences may ever be revealed, has been of intense study in economics [12,17,20]. The main point of departure for this paper is that buyers strategize directly on utilities rather than market implementation specifics, like trading posts and bundles. Hopefully, some of these analysis will lead us to a more effective computational model for markets.

On the technical side, the obvious next question is to completely characterize the NESPs for the general Fisher market game. We assumed the utility functions of the buyers to be linear, however Fisher market is gameable for the other class of utility functions as well. It will be interesting to do a similar analysis for more general utility functions.

References

1. Adsul, B., Babu, C.S., Garg, J., Mehta, R., Sohoni, M.: Nash equilibria in Fisher market. arXiv:1002.4832 (2010)
2. Amir, R., Sahi, S., Shubik, M., Yao, S.: A strategic market game with complete markets. Journal of Economic Theory 51, 126–143 (1990)
3. Brainard, W.C., Scarf, H.E.: How to compute equilibrium prices in 1891. Cowles Foundation, Discussion Paper-1272 (2000)
4. Bu, T., Deng, X., Qi, Q.: Forward looking Nash equilibrium for keyword auction. Inf. Process. Lett. 105(2), 41–46 (2008)
5. Chakrabarty, D., Devanur, N.R., Vazirani, V.V.: New results on rationality and strongly polynomial solvability in Eisenberg-Gale markets. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) WINE 2006. LNCS, vol. 4286, pp. 239–250. Springer, Heidelberg (2006)
6. Codenotti, B., Pemmaraju, S., Varadarajan, K.: On the polynomial time computation of equilibria for certain exchange economies. In: SODA 2005 (2005)
7. Devanur, N.R., Papadimitriou, C.H., Saberi, A., Vazirani, V.V.: Market equilibrium via a primal-dual type algorithm. JACM 55(5) (2008)
8. Dubey, P., Geanakoplos, J.: From Nash to Walras via Shapley-Shubik. Journal of Mathematical Economics 39, 391–400 (2003)
9. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. The American Economic Review 97(1), 242–259 (2007)
10. Garg, J.: Nash equilibria in Fisher market. Working Manuscript (2010)
11. Jain, K.: A polynomial time algorithm for computing the Arrow-Debreu market equilibrium for linear utilities. In: FOCS 2004 (2004)
12. Mas-Colell, A., Whinston, M.D., Green, J.R.: Microeconomic Theory. Oxford University Press, Oxford (1995)
13. Nash, J.F.: Equilibrium points in n-person games. Proc. of the National Academy of Sciences of the United States of America 36(1), 48–49 (1950)
14. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press, Cambridge (2007)
15. Orlin, J.B.: Improved algorithms for computing Fisher's market clearing prices. In: STOC 2010 (2010)
16. Samuelson, P.A.: A note on the pure theory of consumers' behaviour. Economica 5, 61–71 (1938)
17. Samuelson, P.A.: Foundations of Economic Analysis. Harward University Press (1947)
18. Shapley, L., Shubik, M.: Trade using one commodity as a means of payment. Journal of Political Economy 85(5), 937–968 (1977)
19. Varian, H.: Position auctions. International Journal of Industrial Organization 25, 1163–1178 (2007)
20. Varian, H.: Microeconomic Analysis, 3rd edn (1992)
21. Walras, L.: Elements of Pure Economics, Translated by Jaffé, Allen & Urwin. London (1954)