IEOR 3900: Undergraduate Research	May 12, 2025
Maximizing Two-Sided Nash Welfare	
Profs. Yuri Faenza, Jay Sethuraman	Shiv Kampani

1 Introduction

Let $\mathcal{F}$ be a set of firms with $|\mathcal{F}| = n$ and $\mathcal{W}$ be a set of workers with $|\mathcal{W}| = m$. We will study the problem of allocating workers to firms in a manner that is provably **efficient** and **fair**. We define the valuation matrix $v^{\mathcal{F}} \in \mathbb{R}_{\geq 0}^{n \times m}$ such that v_{ij} is the value to the firm $i \in \mathcal{F}$ from employing worker $j \in \mathcal{W}$. We will consider **additive valuations**.

From a utilitarian perspective, our priority should be to maximize the sum of utilities to each firm. While this may be an efficient allocation, it is often unfair as it prioritizes select firms that value workers highest. Maximizing the product (equivalently, geometric mean) of utilities, i.e., Nash welfare, results in a provably fair allocation. The problem of maximizing (one-sided) Nash welfare is well-studied. Most notably, Cole and Gkazetlis provide a polynomial-time 2-approximation by interpreting the problem as a market process.

In this project, we modify the model further by introducing another valuation matrix $v^{\mathcal{W}} \in \mathbb{R}_{\geq 0}^{n \times m}$ such that v_{ij} is the value to the worker $j \in \mathcal{W}$ from being employed by firm $i \in \mathcal{F}$. This makes the problem **two-sided**. Two-sided Nash welfare is then defined as the product of the geometric means of the utilities to firms and workers respectively.

We will not consider part-time workers. In other words, a firm i can employ any number of workers upto its capacity c_i; however, a worker j can only be employed by a single firm.

2 Primal program

Let the allocation be $\mu : \mathcal{F} \times \mathcal{W} \to \{0, 1\}$. $\mu(i, j) = x_{ij} = 1 \iff$ firm i hires worker j.

$$\max \quad \left(\prod_{i \in \mathcal{F}} u_i^{\mathcal{F}}\right)^{1/n} \cdot \left(\prod_{j \in \mathcal{W}} u_j^{\mathcal{W}}\right)^{1/m}$$

$$
\begin{aligned}
\text{subj. to} \quad & \forall i \in \mathcal{F} \quad \sum_{j \in \mathcal{W}} v_{ij}^{\mathcal{F}} \cdot x_{ij} = u_i^{\mathcal{F}} \\
& \forall j \in \mathcal{W} \quad \sum_{i \in \mathcal{F}} v_{ij}^{\mathcal{W}} \cdot x_{ij} = u_j^{\mathcal{W}} \\
& \forall i \in \mathcal{F} \quad \sum_{j \in \mathcal{W}} x_{ij} \leq c_i \\
& \forall j \in \mathcal{W} \quad \sum_{i \in \mathcal{F}} x_{ij} \leq 1 \\
& \forall i \in \mathcal{F}, j \in \mathcal{W} \quad x_{ij} \in \{0, 1\}
\end{aligned}
$$

$u^{\mathcal{F}}, u^{\mathcal{W}}$ represent the utilities to the firms and workers respectively. The constraints represent the definitions of utility, capacity constraints, and integrality constraints. We replace the objective function from maximizing a geometric mean of utilities to (equivalently) maximizing the arithmetic mean of logarithms of utilities. Below, we also relax the integrality constraint that comes from the "indivisibility" of workers (no part-time workers):

$$
\begin{aligned}
\max \quad & \frac{1}{n} \sum_{i \in \mathcal{F}} \log\left(u_i^{\mathcal{F}}\right) + \frac{1}{m} \sum_{j \in \mathcal{W}} \log\left(u_i^{\mathcal{W}}\right) \\
\text{subj. to} \quad & \forall i \in \mathcal{F} \quad \sum_{j \in \mathcal{W}} v_{ij}^{\mathcal{F}} \cdot x_{ij} = u_i^{\mathcal{F}} \\
& \forall j \in \mathcal{W} \quad \sum_{i \in \mathcal{F}} v_{ij}^{\mathcal{W}} \cdot x_{ij} = u_j^{\mathcal{W}} \\
& \forall i \in \mathcal{F} \quad \sum_{j \in \mathcal{W}} x_{ij} \leq c_i \\
& \forall j \in \mathcal{W} \quad \sum_{i \in \mathcal{F}} x_{ij} \leq 1 \\
& \forall i \in \mathcal{F}, j \in \mathcal{W} \quad x_{ij} \geq 0
\end{aligned}
$$

The relaxed primal is a **convex optimization problem** with linear constraints. This program is a generalization of the (relaxation of the) one-sided version of this problem studied by Cole and Gkatzelis. We can show this by setting $\forall i, j,\ v_{ij}^{\mathcal{W}} = c$, a constant, and by removing capacity constraints on firms. Note that the one-sided relaxed program is equivalent to the Eisenberg-Gale (EG) program, which is well-studied in the literature. It is worth pointing out at the outset that the two-sided program above has an exponential integrality gap (can be proved by a reduction to the one-sided case, which was shown by Cole and Gkatzelis).

3 Dual program

Let $\alpha_i, \beta_j, q_i, p_j$ represent the sets of dual variables associated with respective utility and capacity constraints in the relaxed primal problem. We will view the relaxed primal as a minimization problem by negating its objective. Consider the associated Lagrangian:

$$\begin{aligned}
\mathcal{L}(x, u^{\mathcal{F}}, u^{\mathcal{W}}, q, p) = &-\frac{1}{n}\sum_{i\in\mathcal{F}} \log\left(u_i^{\mathcal{F}}\right) - \frac{1}{m}\sum_{j\in\mathcal{W}} \log\left(u_i^{\mathcal{W}}\right) \\
&+ \sum_{i\in\mathcal{F}} \alpha_i \left(u_i^{\mathcal{F}} - \sum_{j\in\mathcal{W}} v_{ij}^{\mathcal{F}} x_{ij}\right) + \sum_{j\in\mathcal{W}} \beta_j \left(u_j^{\mathcal{W}} - \sum_{i\in\mathcal{F}} v_{ij}^{\mathcal{W}} x_{ij}\right) \\
&+ \sum_{i\in\mathcal{F}} q_i \left(\sum_{j\in\mathcal{W}} x_{ij} - c_i\right) + \sum_{j\in\mathcal{W}} p_j \left(\sum_{i\in\mathcal{F}} x_{ij} - 1\right)
\end{aligned}$$

Note: the domain of the Lagrangian has $\forall i, j\ x_{ij} \geq 0$. We define the Lagrange dual function:

$$\begin{aligned}
g(\alpha, \beta, q, p) &= \inf_{u^{\mathcal{F}}, u^{\mathcal{W}}, x} \{\mathcal{L}(x, p, q)\} \\
&= \inf_{u^{\mathcal{F}}} \{-\frac{1}{n}\sum_{i\in\mathcal{F}} \log\left(u_i^{\mathcal{F}}\right) + \sum_{i\in\mathcal{F}} \alpha_i u_i^{\mathcal{F}}\} \\
&+ \inf_{u^{\mathcal{W}}} \{-\frac{1}{m}\sum_{j\in\mathcal{W}} \log\left(u_i^{\mathcal{W}}\right) + \sum_{j\in\mathcal{W}} \beta_j u_j^{\mathcal{W}}\} \\
&+ \inf_{x_{ij}} \{\sum_{i,j} \left(-\alpha_i v_{ij}^{\mathcal{F}} - \beta_j v_{ij}^{\mathcal{W}} + q_i + p_j\right)\ x_{ij}\} - \sum_{i\in\mathcal{F}} q_i c_i + \sum_{j\in\mathcal{W}} p_j
\end{aligned}$$

Observe that $\inf_x\{-\frac{1}{n}\log(x) + cx\}$ has $x = \frac{1}{nc}$, so it is equal to $\frac{1}{n}(1 + \log(nc))$. Observe, also, that $g \to -\infty$ if $-\alpha_i v_{ij}^{\mathcal{F}} - \beta_j v_{ij}^{\mathcal{W}} + q_i + p_j < 0$ (since $x_{ij} \geq 0$). Thus, the dual program is:

$$\begin{aligned}
\min \quad & \sum_{i\in\mathcal{F}} q_i c_i + \sum_{j\in\mathcal{W}} p_j - \frac{1}{n}\sum_{i\in\mathcal{F}} \log\left(u_i^{\mathcal{F}}\right) - \frac{1}{m}\sum_{j\in\mathcal{W}} \log\left(u_i^{\mathcal{W}}\right) \\
\text{subj. to} \quad & \forall i \in \mathcal{F}, j \in \mathcal{W} \quad q_i + p_j \geq \alpha_i v_{ij}^{\mathcal{F}} + \beta_j v_{ij}^{\mathcal{W}} \\
& \forall i \in \mathcal{F} \quad q_i \geq 0 \\
& \forall j \in \mathcal{W} \quad p_j \geq 0
\end{aligned}$$

This is similar in form to the dual program for the one-sided case considered by Cole and Gkatzelis. As a sanity check, we see that the dual program has as many constraints as there are variables in the primal program.

Remark. (market-based interpretation) Like the one-sided case, we can also interpret the two-sided case as a market process, where dual variables q_i, p_j represent the prices of an employment contract between firm i and worker j, where q_i is the price to firm i and p_j is the price to firm j. In the style of Cole and Gkatzelis, we will try to take advantage of this market-based interpretation to design approximation algorithms. Note that we have still not specified equilibrium conditions for this market process.

4 KKT conditions

The KKT conditions are a set of optimality conditions for convex programs. For any optimization problem with convex and differentiable objective and constraint functions, any pair of primal and dual optimal points must satisfy the KKT conditions.

It is known that the KKT conditions of the relaxed one-sided program (EG) are equivalent to the market-clearing conditions for a linear Fisher market with unit-budgets. This forms the basis for Cole and Gkatzelis' approximation algorithm. In the two-sided case, it may be fruitful to apply the same approach of interpreting KKT conditions as market clearing conditions. This is explored in a subsequent section.

Before listing the KKT conditions, consider the gradient of the Lagrangian:

$$\begin{aligned}
\forall i \in \mathcal{F}, j \in \mathcal{W},\ x_{ij} > 0 \quad & \nabla_{x_{ij}} \mathcal{L} = q_i + p_j - \alpha_i v^{\mathcal{F}}_{ij} - \beta_j v^{\mathcal{W}}_{ij} \\
\forall i \in \mathcal{F} \quad & \nabla_{u^{\mathcal{F}}_i} \mathcal{L} = \alpha_i - \frac{1}{n\, u^{\mathcal{F}}_i} \\
\forall j \in \mathcal{W} \quad & \nabla_{u^{\mathcal{W}}_j} \mathcal{L} = \beta_j - \frac{1}{m\, u^{\mathcal{W}}_j}
\end{aligned}$$

Let $(x^*_{ij}, {u^{\mathcal{F}}}^*_i, {u^{\mathcal{W}}}^*_j)$ be any primal optimal point and $(\alpha^*_i, \beta^*_j, q^*_i, p^*_j)$ be any dual optimal point.

The KKT conditions for the two-sided program are:

1. **Primal Feasibility:** (see primal constraints)

$$\forall i \in \mathcal{F} \quad \sum_{j \in \mathcal{W}} v^{\mathcal{F}}_{ij} \cdot x^*_{ij} = {u^{\mathcal{F}}}^*_i \tag{1}$$

$$\forall j \in \mathcal{W} \quad \sum_{i \in \mathcal{F}} v^{\mathcal{W}}_{ij} \cdot x^*_{ij} = u^{\mathcal{W}*}_j \tag{2}$$

$$\forall i \in \mathcal{F} \quad \sum_{j \in \mathcal{W}} x^*_{ij} \leq c_i \tag{3}$$

$$\forall j \in \mathcal{W} \quad \sum_{i \in \mathcal{F}} x^*_{ij} \leq 1 \tag{4}$$

$$\forall i \in \mathcal{F}, j \in \mathcal{W} \quad x^*_{ij} \geq 0 \tag{5}$$

2. **Dual Feasibility:** (see dual constraints)

$$\forall i \in \mathcal{F}, j \in \mathcal{W} \quad q^*_i + p^*_j \geq \alpha^*_i v^{\mathcal{F}}_{ij} + \beta^*_j v^{\mathcal{W}}_{ij} \tag{6}$$

$$\forall i \in \mathcal{F} \quad q^*_i \geq 0 \tag{7}$$

$$\forall j \in \mathcal{W} \quad p^*_j \geq 0 \tag{8}$$

3. **Complementary Slackness:**

$$\forall i \in \mathcal{F} \quad q^*_i \left(\sum_{j \in \mathcal{W}} x^*_{ij} - c_i \right) = 0 \tag{9}$$

$$\forall j \in \mathcal{W} \quad p^*_j \left(\sum_{i \in \mathcal{F}} x^*_{ij} - 1 \right) = 0 \tag{10}$$

4. $\nabla \mathcal{L} = \mathbf{0}$: (see expression for $\nabla \mathcal{L}$)

$$\forall i \in \mathcal{F}, j \in \mathcal{W} \quad x_{ij} > 0 \implies q^*_i + p^*_j = \alpha^*_i v^{\mathcal{F}}_{ij} + \beta^*_j v^{\mathcal{W}}_{ij} \tag{11}$$

$$\forall i \in \mathcal{F} \quad \alpha^*_i = \frac{1}{n\, u^{\mathcal{F}*}_i} \tag{12}$$

$$\forall j \in \mathcal{W} \quad \beta^*_j = \frac{1}{m\, u^{\mathcal{W}*}_j} \tag{13}$$

We write the KKT conditions concisely, substituting the dual variables α^*_i, β^*_j (and assuming primal feasibility). The conditions listed below generalize the one-sided KKT conditions listed by Vazirani in "Algorithmic Game Theory."

(i) $\forall i \in \mathcal{F} \quad q^*_i \geq 0$
$\forall j \in \mathcal{W} \quad p^*_j \geq 0$

(ii) $\forall i \in \mathcal{F} \quad q^*_i > 0 \implies \sum_{j \in \mathcal{W}} x^*_{ij} = c_i$

$$\forall j \in \mathcal{W} \quad p_j^* > 0 \implies \textstyle\sum_{i \in \mathcal{F}} x_{ij}^* = 1$$

(iii) $\forall i \in \mathcal{F}, j \in \mathcal{W} \quad q_i^* + p_j^* \geq \frac{v_{ij}^{\mathcal{F}}}{n\, u^{\mathcal{F}*}_i} + \frac{v_{ij}^{\mathcal{W}}}{m\, u^{\mathcal{W}*}_j}$

(iv) $\forall i \in \mathcal{F}, j \in \mathcal{W} \quad x_{ij} > 0 \implies q_i^* + p_j^* = \frac{v_{ij}^{\mathcal{F}}}{n\, u^{\mathcal{F}*}_i} + \frac{v_{ij}^{\mathcal{W}}}{m\, u^{\mathcal{W}*}_j}$

5 Explorations with `cvxpy`

We know that the integrality gap of the two-sided program considered is exponential. This can be proved by a reduction from the one-sided program, whose integrality gap was proved to be exponential by Cole and Gkatzelis. Since rounding doesn't work, it is not immediately clear what approach towards an approximation algorithm will work best.

It is worth numerically solving the two-sided program and understanding what the optimal allocations x_{ij} look like. This can be accomplished using the `cvxpy` package.

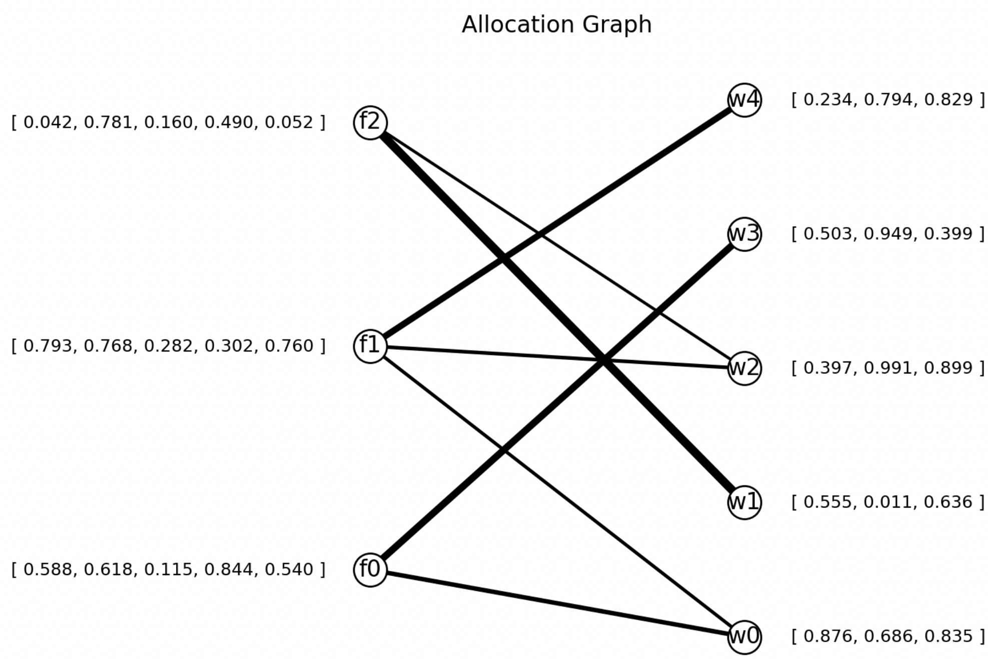

Figure 1: *Allocation graph from Python script for 3 firms, 5 workers, and random valuations.*

It is important to note that only a subset of all possible bipartite edges (i, j) are actually present in the optimal allocation. Given prices q, p, is it possible to know which subset of edges (i, j) the optimal allocation always consists of? This is a problem for future work.