# Approximating the Nash Social Welfare with Indivisible Items

RICHARD COLE
Courant Institute, New York University
and
VASILIS GKATZELIS
Computer Science, Stanford University

In this letter we briefly discuss our main result from [Cole and Gkatzelis 2015]. Given a set of indivisible items and a set of agents having additive valuations, our goal is to allocate the items to the agents in order to maximize the geometric mean of the agents' valuations, i.e., the Nash social welfare. This problem is known to be NP-hard, and our main result is the first efficient constant-factor approximation algorithm for this objective.

Categories and Subject Descriptors: F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms & Computations on discrete structures

General Terms: Algorithms, Theory, Economics

Additional Key Words and Phrases: Nash social welfare, approximation algorithms, fair division

## 1. INTRODUCTION

Imagine that we need to allocate a collection of $m$ indivisible items in a fair and efficient manner among a set of $n \leq m$ agents having additive valuations. That is, each agent $i$ has a non-negative value $v_{ij}$ for each item $j$, and her value for an allocation $x$ that assigns to her some bundle of items $B_i$, is $v_i(x) = \sum_{j \in B_i} v_{ij}$. Since the items are indivisible, an allocation $x$ assigns each item to a single agent, and different allocations lead to different distributions of value $v_i(x)$ among the agents. How should we allocate the items to ensure a distribution of value that balances fairness and efficiency?

Maximizing efficiency is commonly associated with maximizing the utilitarian social welfare objective, i.e., the *total value* across the agents: $\max_x \sum_i v_i(x)$. It is easy to verify though that this objective can often lead to some agents being allocated nothing, hence neglecting fairness considerations. On the other extreme, maximizing fairness is often associated with maximizing the egalitarian social welfare objective, i.e., the *minimum value* across all agents: $\max_x \min_i v_i(x)$. But, just as the utilitarian objective disregards fairness, the egalitarian objective can lead to really inefficient outcomes, allocating many items to hard to satisfy agents.

A well-studied objective that lies between these two extremes is the *Nash social welfare* (NSW), which maximizes the geometric mean of the agents' values, i.e.,

$$\max_x \left( \prod_i v_i(x) \right)^{1/n}.$$

Authors' addresses: `cole@cims.nyu.edu`, `gkatz@cs.stanford.edu`

The Nash social welfare objective goes back to the fifties [Nash 1950; Kaneko and Nakamura 1979], and it satisfies several appealing properties [Moulin 2003]. As we already mentioned, it strikes a natural balance between fairness and efficiency since maximizing the geometric mean favors more uniform value distributions, but without sacrificing the utilitarian social welfare too much. Furthermore, this objective is scale-free: its optimal allocation $x^*$ is independent of the scale of each agent's valuations, so choosing the desired allocation does not require interpersonal comparability of the individual's preferences. This property is particularly useful in settings where the agents are not paying for the items that they are allocated, in which case the scale in which their valuations are expressed may not have any real meaning (if they were paying, $v_{ij}$ could be interpreted as the amount that agent $i$ is willing to pay for item $j$). These are properties that neither the utilitarian nor the egalitarian objective satisfy.

## 2. PRELIMINARIES

Given a set $M$ of $m$ items and a set $N$ of $n$ agents ($n \leq m$) with additive valuations, our optimization problem can be expressed as the following integer program, IP:

$$\begin{aligned}
\text{maximize:} \quad & \left(\prod_{i \in N} u_i\right)^{1/n} \\
\text{subject to:} \quad & \sum_{j \in M} x_{ij} v_{ij} = u_i, \quad && \forall i \in N \\
& \sum_{i \in N} x_{ij} = 1, && \forall j \in M \\
& x_{ij} \in \{0, 1\}, && \forall i \in N, j \in M
\end{aligned}$$

Note that, for any fixed number of agents $n$, the objective of IP is equivalent to maximizing $\sum_{i \in N} \log u_i$. As a result, our optimization problem can be expressed as a convex integer program. In fact, the fractional relaxation of this convex program is an instance of the well studied Eisenberg-Gale program [Eisenberg and Gale 1959], whose solutions can be interpreted as the equilibrium allocation for the linear case of Fisher's market model [Nisan et al. 2007, Chapter 5]. In this model, each agent has a certain budget, and she is using this budget in order to buy fractions of the available items. Although the agents in our setting are *not* using money, this market-based interpretation of the optimal solution of IP's fractional relaxation provides some very useful intuition, which we use in order to design our approximation algorithm.

In the Fisher market corresponding to our problem the items are divisible and the valuation of agent $i$ who is receiving a fraction $x_{ij} \in [0, 1]$ of each item $j$ is $\sum_{j \in M} x_{ij} v_{ij}$. Each agent has the same budget of, say, \$1 to spend on items and each item $j$ has a price $p_j$. If agent $i$ spends $b_{ij}$ on an item whose price is $p_j$, then she receives a fraction $x_{ij} = b_{ij}/p_j$ of that item (there is one unit of each item, so $\sum_{i \in N} b_{ij} \leq p_j$). A vector of item prices $p = (p_1, \ldots, p_m)$ induces a market equilibrium if every agent is spending all of her budget on her "optimal" items given these prices, and the market clears, i.e., all of the items are allocated fully. The "optimal" items for agent $i$, given prices $p$, are the ones that maximize the ratio $v_{ij}/p_j$, also known as the *maximum bang per buck* (MBB) items.

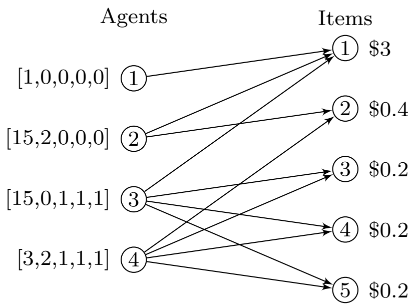


Fig. 1: The MBB graph.

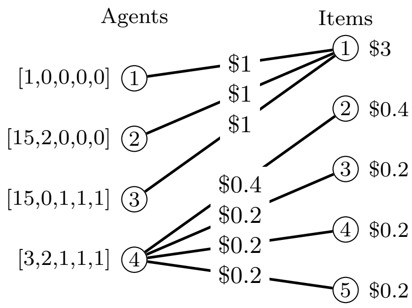


Fig. 2: The spending graph.

Example 2.1. *The graph of Figure 1 comprises 4 vertices on the left (the agents) and 5 vertices on the right (the items). The valuations of each agent appear to the left of her vertex, so the first agent values only the first item. The second agent values this item at 15 and the second item at 2, while she has no value for the other items, etc. The price to the right of each item's vertex corresponds to its price in the market equilibrium for this problem instance, and the directed edges point from each agent to its MBB items at the given prices. To verify that these are market clearing prices, Figure 2 shows how each agent could spend* all *of her budget on MBB items. The label of each edge $(i, j)$ denotes how much agent $i$ is spending on item $j$. Since the total money spent on each item is exactly equal to its price, all of the items are fully allocated, which implies that this equilibrium is the optimal solution of the fractional relaxation of* IP*: the first three agents each get a 1/3 fraction of the first item, and the fourth agent gets all of the other items.*

## 3. OUR MAIN RESULT

Let $x^*$ denote the integral allocation that maximizes the Nash social welfare. Our main result in [Cole and Gkatzelis 2015] is the *Spending-Restricted Rounding* (SRR) algorithm, an efficient algorithm which computes an integral allocation $\tilde{x}$ guaranteed to be within a constant factor of the optimal one. The best previously known approximation factor was $\Omega(m)$ [Nguyen and Rothe 2014].

Theorem 3.1. *The allocation $\tilde{x}$ computed by the SRR algorithm satisfies*

$$\left(\prod_{i\in N} v_i(x^*)\right)^{1/n} \leq 2.889 \left(\prod_{i\in N} v_i(\tilde{x})\right)^{1/n}.$$

Since we can solve the fractional relaxation of IP using the Eisenberg-Gale program, a standard technique for designing an approximation algorithm would be to take the fractional allocation and "round" it in an appropriate way to get an integral one. The hope would be that the fractional allocation provides some useful information regarding what a good integral allocation should look like, as well as an upper bound for the geometric mean of $x^*$. Unfortunately, we show that the *integrality gap* of IP, i.e., the ratio of the geometric mean of the fractional solution and that of $x^*$, is $\Omega(2^m)$, which implies that the fractional solution of IP cannot be used for proving a constant-factor approximation guarantee using standard techniques.

To circumvent the integrality gap, we introduce an interesting new constraint on the fractional solution. In particular, we restrict the total amount of money spent on any item to be at most \$1, i.e., at most the budget of a single agent. For any item $j$, the solution needs to satisfy $\sum_{i \in N} x_{ij} p_j \leq 1$; a constraint which combines both the primal ($x_{ij}$) and the dual ($p_j$) variables of the Eisenberg-Gale program.

Definition 3.2. *A* spending-restricted *(SR) outcome is a fractional allocation* $\overline{x}$ *and a price vector* $\overline{p}$ *such that every agent spends* all *of her budget on her MBB items at prices* $\overline{p}$*, and the total spending on each item is equal to* $\min\{1, \overline{p}_j\}$.

Example 3.3. *In the unrestricted equilibrium of the instance considered in Example 2.1, the price of the highly demanded first item was \$3, and three agents were spending all of their budgets on it. In a spending-restricted outcome this would not be acceptable, so the price of this item would need to be increased further until only the first agent, who has no other alternative, is spending her budget on it. The prices and the spending of the SR outcome can be seen in Figure 3. Note that, unlike the fractional solution of the unrestricted market equilibrium, this fractional solution reveals more information regarding the preferences of Agents 2 and 3.*

Our SRR algorithm begins by computing an SR allocation $\overline{x}$ and prices $\overline{p}$ and then appropriately allocates each item to one of the agents who is receiving some of it in $\overline{x}$. Each one of the items $j$ with a price $\overline{p}_j > 1/2$ is assigned to a distinct agent, while the rest of the items are allocated in a way that ensures a balanced distribution of value. To prove the approximation guarantee, we also show that the SR outcome implies an interesting upper bound for the optimal Nash social welfare.

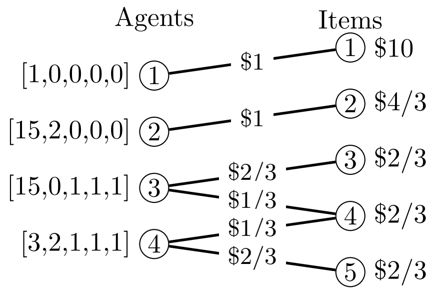


Fig. 3: The spending-restricted outcome.

## 4. CONCLUSION

Allocating indivisible items in a fair and efficient manner is a long standing problem that has received a lot of attention. Our algorithm provides a non-trivial approximation guarantee with respect to a well motivated objective that strikes a natural balance between fairness and efficiency. In designing this algorithm we use a fractional allocation that satisfies a novel constraint motivated by a market-theoretic interpretation of the fractional relaxation. In particular, this constraint restricts the amount of money that can be spent on any given item, thus forcing some agents to spend on less demanded items, thereby revealing useful information regarding the agents' preferences. This constraint combines both the primal and the dual variables of the Eisenberg-Gale convex program. How to compute the allocation that satisfies it is not obvious, but we show that it can, in fact, be computed in polynomial time.

## REFERENCES

COLE, R. AND GKATZELIS, V. 2015. Approximating the Nash social welfare with indivisible items. In *47th ACM Symposium on Theory of Computing, STOC*.

EISENBERG, E. AND GALE, D. 1959. Consensus of subjective probabilities: The pari-mutuel method. *Ann. Math. Stat. 30*, 165–168.

KANEKO, M. AND NAKAMURA, K. 1979. The nash social welfare function. *Econometrica 47,* 2, pp. 423–435.

MOULIN, H. 2003. *Fair Division and Collective Welfare.* The MIT Press.

NASH, J. 1950. The bargaining problem. *Econometrica 18,* 2 (April), 155–162.

NGUYEN, T. T. AND ROTHE, J. 2014. Minimizing envy and maximizing average nash social welfare in the allocation of indivisible goods. *Discrete Applied Mathematics 179,* 0, 54 – 68.

NISAN, N., ROUGHGARDEN, T., TARDOS, É., AND VAZIRANI, V. 2007. *Algorithmic Game Theory.* Cambridge University Press, New York, NY, USA.