European Journal of Operational Research 247 (2015) 548–559

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

Maximizing Nash product social welfare in allocating indivisible goods

Andreas Darmann[a,*], Joachim Schauer[b]

[a] University of Graz, Institute of Public Economics, Universitaetsstr. 15, A-8010 Graz, Austria
[b] University of Graz, Department of Statistics and Operations Research, Universitaetsstr. 15, A-8010 Graz, Austria

ARTICLE INFO

Article history:
Received 16 June 2014
Accepted 26 May 2015
Available online 3 June 2015

Keywords:
Fair division
Social welfare
Nash product
Computational complexity
Scoring rules

ABSTRACT

We consider the problem of allocating indivisible goods to agents who have preferences over the goods. In such a setting, a central task is to maximize social welfare. In this paper, we assume the preferences to be additive and measure social welfare by means of the Nash product. We focus on the computational complexity involved in maximizing Nash product social welfare when scores inherent in classical voting procedures such as approval or Borda voting are used to associate utilities with the agents' preferences. In particular, we show that the maximum Nash product social welfare can be computed efficiently when approval scores are used, while for Borda and lexicographic scores the corresponding decision problem becomes NP-complete.

1. Introduction

The allocation of goods (items, resources) to agents who have preferences over these goods (multiagent resource allocation) is a fundamental problem of economics, and, in particular, social choice theory. This problem has been tackled in various scenarios (see, e.g., Chevaleyre et al., 2006 for a survey), where, e.g., we distinguish between divisible and indivisible goods, and centralized and decentralized approaches. Here, we consider the case of m indivisible and non-shareable goods to be distributed among agents who report their preferences to a central authority. Typically, individual utilities of (bundles of) items are associated with the preferences over the items. In this work, this is done via numerical scores used in voting rules. Now, a major task is to find an allocation which maximizes the social welfare achieved. Different notions of social welfare have been introduced, the most important being utilitarian, egalitarian, and Nash product social welfare (cf. Brandt, Conitzer, & Endriss, 2013).

Utilitarian social welfare of an allocation is given by the sum of the agents' utilities resulting from the allocation. A more fine-grained approach is egalitarian social welfare, where the lowest of the agents' individual utilities in a given allocation is considered. In a certain sense, the Nash product social welfare links these two approaches: by measuring the product of the agents' utilities in an allocation, maximizing the Nash product social welfare targets at a "balanced" allocation (see also Nguyen, Nguyen, Roos, & Rothe, 2014).

The Nash product as a measure for social welfare satisfies several desirable properties (see Moulin, 2003 for an explicit treatment). For example, it satisfies the basic fairness criterion that it increases when inequality among two agents is reduced (given the respective change is mean-preserving; see also Ramezani & Endriss, 2010). Clearly, the Nash product also satisfies the monotonicity property that an increase of an agent's utility yields an increase in the Nash product. In addition, it is independent of both common utility scale and individual utility scale: the social welfare ordering, i.e., the ordering of the allocations according to their Nash product, remains unchanged both if (i) all agents rescale their utilities with the same factor, and (ii) each agent rescales her utility using a different factor.

A central question in maximizing social welfare is the computational complexity involved. We assume that the agents have additive preferences, i.e., for each agent, the utility of a set of goods is the sum of the utilities of the single goods it contains.

Clearly, maximizing utilitarian social welfare is an easy task – simply allocate each item to an agent who it yields the highest utility for (see also Brandt et al., 2013). In contrast, it is known that maximizing egalitarian social welfare and Nash product social welfare are NP-complete for additive utilities and general scoring functions (Roos & Rothe, 2010). Recently, Baumeister et al. (2013) have shown that maximizing egalitarian social welfare remains NP-complete for a number of prototypical scoring functions from voting theory: quasi-indifference, Borda, and lexicographic scoring. On the positive side, it is known that the maximum egalitarian social welfare can be computed in polynomial time for approval scores (Golovin, 2005). To the best of our knowledge, the computational complexity of maximizing Nash product social welfare under scoring functions such as approval, Borda, or lexicographic scoring has not been considered yet. In

* Corresponding author. Tel.: +43 3163807139; fax: +43 3163809530.
E-mail addresses: andreas.darmann@uni-graz.at (A. Darmann), joachim.schauer@uni-graz.at (J. Schauer).

http://dx.doi.org/10.1016/j.ejor.2015.05.071
0377-2217/

this paper, we investigate the computational complexity involved in maximizing Nash product social welfare under these classical scoring functions.

Related work and our contribution. In the Santa Claus problem (Bansal & Sviridenko, 2006), the goal is to maximize egalitarian social welfare when indivisible items have to be allocated among agents. There, the agents associate an arbitrary numerical value with each item, i.e., the input is numerical, and the utility of an agent is the sum of the values of the items assigned to her. In contrast, as done in our work, in the setting presented in Brams, Edelman, and Fishburn (2003) the agents' preferences are expressed ordinally. Brams et al. (2003) introduce the notion of a Borda-optimal allocation, which is defined as an allocation maximizing egalitarian social welfare where an agent's utility is the sum of the Borda scores of the items the agent receives. We translate this concept to maximizing Nash product social welfare, and expand the perspective from Borda scores to approval (and, in particular, k-approval) and lexicographic scores. Informally speaking, in k-approval scoring a distinction between "good" (approved) and "bad" (disapproved) items is made: the top k items in the ranking of an agent receive a score of 1, while the remaining items get score 0. In Borda scoring, an agent's most preferred item gets a score of m, her second-ranked item a score of $m-1$, and so on; her least preferred item has a score of 1. In lexicographic scoring, the position of the item in the ranking is even more crucial: any item r yields a higher score than the total of all items ranked below r. Our goal is to analyze the computational complexity involved in maximizing Nash product social welfare with respect to these types of scores. In the context of maximizing social welfare in multiagent resource allocation, complexity results have been achieved with respect to different types of utility representation: the bundle form, k-additive form, or straight-line programs. For the bundle form representation, NP-completeness results for utilitarian (Chevaleyre, Endriss, Estivie, & Maudet, 2008), egalitarian (Roos & Rothe, 2010), and Nash product social welfare (Ramezani & Endriss, 2010; Roos & Rothe, 2010) are known. For straight-line programs, Dunne, Wooldridge, and Laurence (2005) show that maximizing utilitarian social welfare is NP-complete, while Nguyen et al. (2014) show that maximizing social welfare is NP-complete both for the egalitarian and Nash product approach. Both maximizing egalitarian social welfare and maximizing Nash product social welfare turn out to be NP-complete for 1-additive, i.e., additive utilities already (Lipton, Markakis, Mossel, & Saberi, 2004; Roos & Rothe, 2010). In these works, however, reductions from PARTITION are given, which do not imply the NP-completeness for any of the scoring functions considered in our work. Given additive utilities, Baumeister et al. (2013), besides many other results, have proven that maximizing egalitarian social welfare is NP-complete for Borda, lexicographic, and quasi-indifference scoring.

In this paper, we show that maximizing Nash product social welfare is NP-complete for Borda and lexicographic scores, whereas it is polynomially solvable for approval scores. The computational complexity involved when quasi-indifference scores are used is still open.

2. Formal framework

2.1. Preliminaries

Let $R = \{r_1, r_2, \dots, r_m\}$ be a set of m indivisible resources (items) and let $A = \{a_1, \dots, a_n\}$ be a set of n agents. An *allocation* is a mapping that assigns to each agent a subset of resources such that each resource is handed to exactly one agent. Formally, an allocation P is a mapping P: $A \to 2^R$ with $\bigcup_{a\in A} P(a) = R$ and $P(a_i) \cap P(a_j) = \emptyset$ whenever $i \neq j$.

Now, in our model, we start with ordinal inputs, i.e., the agents rank resources, and map these ranks to numerical scores then. Note that we do not claim that these numerical scores are equivalent or at least close to the agents' actual utilities. However, starting with numerical inputs instead would have several drawbacks (see also Baumeister et al., 2013); e.g., often it is easier for agents to rank items instead of associating numerical values with each single item, especially in contexts where money is not a key factor. Next, as also pointed out in Baumeister et al. (2013), the use of numerical inputs has the severe disadvantage that it insinuates comparability of interpersonal preferences. Finally, note that our approach is very common in voting theory, as in fact it resembles the way that positional scoring rules proceed.[1]

In particular, we assume that agents have preferences over the single resources. The preferences are expressed by means of strict orders $\succ_{a_i}$ over R, which are summarized by the n-tuple $\pi = (\succ_{a_1}, \succ_{a_2}, \dots, \succ_{a_n})$ called *profile*. We denote by $rank_{a_i}(r)$ the rank of resource r in the ranking of agent a_i.

We adopt scores used in voting procedures to evaluate these preferences by means of utility functions $u_a : R \to \mathbb{Q},\ a \in A$. We assume that the utility functions are additive, i.e., for any subset $R' \subseteq R$ we have $u_a(R') = \sum_{r\in R'} u_a(R')$. For the sake of readability, we may write $u_a(P)$ instead of $u_a(P(a))$.

Given a profile π, we consider the following types of scores (where $r \in R$):

- *k-approval scores:* For each agent $a \in A$,
$$u_a(r) = \begin{cases} 1 & \text{if } rank_a(r) \leq k \\ 0 & \text{otherwise} \end{cases}$$
- *Borda scores:* For each agent $a \in A$, $u_a(r) = m + 1 - rank_a(r)$.
- *Lexicographic scores:* For each agent $a \in A$, $u_a(r) = 2^{m-rank_a(r)}$.

Given k-approval scores, for each $a \in A$, u_a partitions the set R into a set $S_a := \{r \in R :\ u_a(r) = 1\}$ (the set of resources agent a approves of) and a set $S_a^c := \{r \in R :\ u_a(r) = 0\}$ (the set of resources agent a disapproves of). Conversely, specifying the set S_a (of size k) for each agent a uniquely determines the corresponding k-approval scores. More generally (and slightly abusing notation), given a set $S(a) \subseteq R$ for each $a \in A$, *approval scores* are given by $u_a(r) = 1$ for $r \in S(a)$ and $u_a(r) = 0$ for $r \in R \setminus S(a)$.

Given an allocation P, the *Nash product social welfare* for P is given by $sw_N(P) = \prod_{1\leq i\leq n} u_{a_i}(P)$.

2.2. Problem definitions

In this paper, we consider the problem of maximizing the Nash product social welfare with respect to the above scores, i.e., utility functions. The corresponding decision problems are defined as follows.

Definition 2.1. (NASH PRODUCT SOCIAL WELFARE MAXIMIZATION-Approval)
GIVEN: Quadruple (R, A, S, k): R is a set of resources, A a set of agents, a collection $S = \{S_{a_1}, S_{a_2}, \dots, S_{a_n}\}$ of subsets $S_{a_i} \subseteq R$, and $k \in \mathbb{N}$.
QUESTION: Is there an allocation P such that $sw_N(P) \geq k$, where $u_{a_i}(r) = 1$ if $r \in S_{a_i}$ and $u_{a_i}(r) = 0$ otherwise?

Analogously, we defineNASH PRODUCT SOCIAL WELFARE MAXIMIZATION-*Borda*.

Definition 2.2. (NASH PRODUCT SOCIAL WELFARE MAXIMIZATION-Borda)
GIVEN: Quadruple (R, A, π, k): R is a set of resources, A a set of agents, π is a profile, and $k \in \mathbb{N}$.
QUESTION: Is there an allocation P such that $sw_N(P) \geq k$ for Borda scores?

[1] Obviously, with the clear difference that we are finally interested in allocations instead of winners of elections.