Computers & Operations Research 75 (2016) 76–82

Contents lists available at ScienceDirect

Computers & Operations Research

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

A heuristic optimization of Bayesian incentive-compatible cake-cutting

Lê Nguyên Hoang [a], François Soumis [a], Georges Zaccour [b,*]

[a] GERAD, École Polytechnique, Montréal, Canada

[b] Chair in Game Theory and Management, GERAD, HEC Montréal, 3000 Côte-Sainte-Catherine, Montreal, Canada H3T 2A7

ARTICLE INFO

Article history:
Received 27 April 2015
Received in revised form
13 May 2016
Accepted 14 May 2016
Available online 20 May 2016

Keywords:
Cake-cutting problem
Fairness
Incentive compatibility
Return function

ABSTRACT

Cake-cutting is a metaphor for problems where a principal agent has to fairly allocate resources. Such problems cover various areas of operations research and management science, like, for instance, shift scheduling with employees' preferences. Recent work focuses on optimizing social efficiency while guaranteeing fairness, but ignore incentive-compatibility constraints, or vice versa. In this paper, we present a new approach to heuristic mechanism design with Bayesian incentive-compatibility. As opposed to other papers, we do not allow monetary transfer. Our approach relies on the revelation principle and the computation of Bayesian–Nash equilibria using the so-called return function. This computation consists in tracking a best-reply dynamics of return function, which are mappings of action to probability distribution on outcomes, instead of the more classical but harder-to-compute best-reply dynamics of strategies. In essence, our mechanism-design approach explores a parameterized class of revelation mechanisms, which we know by construction to be Bayesian incentive-compatible. We highlight the efficiency of this approach through numerical results on instances of respectively 2, 5 and 20 agents.

1. Introduction

In this paper, we propose a mechanism to divide a resource among a set of agents who have heterogenous preferences and may not want to reveal their preferences truthfully. We suppose that the mechanism is designed by a principal agent (or mechanism designer), whose aim is to achieve a specified objective, e.g., maximizing a weighted sum of collective efficiency and fairness indicators. Our contribution belongs to the vast literature dealing with the cake-cutting problem, which is considered a dynamic and challenging field of investigation in operations research and management science (see, e.g., [3,16,12,5]).

The motivation for this work comes from a shift scheduling problem with employees' preferences, which was submitted by a firm designing optimization tools for companies operating in different sectors. Shift scheduling is a cost-minimization problem involving a series of constraints, with some of them specified in a labor contract. Suppose that an employer is willing to take into account the preferences of her employees when deciding about the schedule. For instance, a parent may prefer to start later in the morning to avoid being stressed by traffic conditions when driving her kids to school, whereas another employee may prefer a night shift. Taking into account such preferences will undoubtedly lead to a higher cost than the optimal one, but it may be argued that happy employees are more productive, and a good working atmosphere has also an intangible value to the company. To keep it simple, suppose that we have n employees and n shifts to be allocated. Each allocation has a cost, and the objective of the employer is to min $cost(x)$, subject to $x = (x_1, \ldots, x_n) \in X$, where X is the set of all feasible shift allocations. One way for including employees' preferences is to define a set of feasible shift allocation of reasonable cost, e.g., $X(\alpha) = \{x \in X | cost(x) \le \alpha \min cost(X)\}$, where $\alpha \ge 1$. Denote by a_i the revealed preference of employee i (we shall refer to a_i as agent's type in a Bayesian framework) and let $u_i(a_i, x_i)$ be the utility that this employee obtains when she is allocated shift x_i. Suppose that the employer's objective is to maximize

$$\sum_{i=1}^{n} u_i(a_i, x_i) - \lambda \sum_{i=1}^{n} \left| u_i(a_i, x_i) - \frac{1}{n} \sum_{i=1}^{n} u_i(a_i, x_i) \right|,$$

subject to $x \in X(\alpha)$,

that is, the sum of employees' utilities minus the sum of deviations with respect to the average utility weighted by a parameter $\lambda > 0$. Then, we have a mechanism with employees' preferences as an input, that is, $(a_1, \ldots, a_n)$, and the allocation $(x_1, \ldots, x_n)$ as an output. Of course, the employer (mechanism designer) would want the employees to reveal their true preferences, which are private information, and not gaming the system. The approach proposed here will guarantee that in equilibrium, each employee (player) will indeed behave truthfully.

* Corresponding author.
E-mail addresses: len.hoang.lnh@gmail.com (L.N. Hoang), francois.soumis@gerad.ca (F. Soumis), georges.zaccour@gerad.ca (G. Zaccour).

http://dx.doi.org/10.1016/j.cor.2016.05.009
0305-0548/

The cake-cutting problem, which was first introduced in Steinhaus [17], consists in devising a method to fairly allocate a cake to a set of agents. Many procedures have been proposed over time to do so, including the well-known *last diminisher method* [17], the *divide-and-choose* approach, the *moving knife procedure* [3], and *the successive pairs algorithm* [16]. These mechanisms share the property of being weakly incentive-compatible, that is, an untruthful claimer may regret her untruthfulness at some point. However, if one requires a stronger concept of incentive-compatibility, e.g., dominant strategy incentive-compatibility (DSIC), meaning that agents always have incentives to be truthful, then the problem may end up having no conceptually satisfying solution. To illustrate, Mossel and Tamuz [12] proved that there exists no deterministic DSIC super-fair division to the cake-cutting problem. We recall that a super-fair division yields the exact (or proportional) division solution when all players have the same preferences, and does strictly better than exact division otherwise. Recently, Chen et al. [5] added the assumption that players have piecewise constant valuation functions, and provided a proportionally fair and envy-free deterministic DSIC mechanism. A randomized DSIC super-fair division is discussed in Mossel and Tamuz [12] and Chen et al. [5].

In the above-cited studies, little attention has been given to social efficiency, which could be a legitimate objective. Some recent work focuses on optimizing social efficiency while guaranteeing fairness, but ignore incentive-compatibility constraints (see, e.g., [4,2]). In Cohler et al. [6], the authors provided a tractable, nearly optimal envy-free mechanism when the agents truthfully report their valuations of the cake. They also gave optimal envy-free mechanisms for certain specific structures of the agents' preferences. Note that the difficulty in determining a socially efficient solution is due to the assumptions that (i) the cake is divisible into an infinite number of portions, and (ii) the agents' utility functions are complex mathematical objects involving probability measures. Here, we assume that the cake is made of homogeneous portions, with the agents having constant valuations over each of these portions. In particular, this will enable us to write the set of admissible allocations as a polytope.

Our approach has its roots in mechanism-design theory, where one retains Bayesian incentive-compatibility (BIC) instead of DSIC. It means that we require truthfulness to be a Bayesian–Nash equilibrium. In other words, assuming others truthful, it is in every agent's best interest to reveal her preferences truthfully. One popular approach for BIC design introduced by Myerson [14] for the context of auctions relies on so-called agents' virtual valuations. This approach has gained recent interests through developments of *ironing* methods in multi-dimensional auctions (see, e.g., [15,1,9,10,8]). However, such approaches strongly rely on monetary transfers and agents' risk neutrality.

The main contribution of this paper is a general approach to BIC mechanism design, which requires none of these assumptions: we do not allow monetary transfers nor assume risk neutrality. Instead, we will rely on the revelation principle introduced in Gibbard [7] and Myerson [13,14]. Applying this revelation principle requires the computation of a Bayesian–Nash equilibrium. To do so, we implement the algorithm proposed in Hoang et al. [11], which uses the concept of return functions. We wish to mention from the outset that there is a huge difference between this paper and Hoang et al. [11], namely, the focus here is on selecting or designing a mechanism, whereas in Hoang et al. [11] no such issue is involved as the mechanism is given. A return function is a mapping of an agent's action to the induced probability distribution on her outcomes. The authors showed that any strategy profile in a Bayesian game generates a return-function profile that captures all the information required to describe the best-reply dynamics. Consequently, any best-reply dynamics of strategies is naturally mapped to a best-reply dynamics of return functions. It is significant that return functions and best-replies to return functions are much easier to compute than strategies and best-replies to strategies. This advantage is particularly valuable when the beliefs or the Bayesian game cannot be described analytically, as will be the case in our cake-cutting problem. In a second step, we apply the revelation principle to obtain a BIC mechanism. Finally, we compute the value of the mechanism designer's objective function. This is done iteratively, where at each iteration, the principal draws types of players according to beliefs.

We shall apply this general approach to a cake-cutting problem that aims at a balance of social efficiency and fairness. For this problem, we present the principal agent's (mechanism designer's) optimal mechanism in the case where she has knowledge of players' preferences, which we call the ideal mechanism. Then, we show that this ideal mechanism performs poorly when players play a Bayesian–Nash equilibrium through numerical simulations. Third, we will provide a family of quality parameterized cake-cutting mechanisms to which we shall apply our general approach to BIC mechanism design. The computed results quantify the principal's objective values for the revelations of our family of parameterized mechanisms, which turn out to be significantly improving on the revelation of the ideal mechanism.

The rest of the paper is organized as follows: in Section 2, we provide the theoretical foundations to our approach by describing a general setting of mechanism design and by introducing the return function. Present an algorithm for computing Bayesian–Nash equilibria, implementing the revelation principle and optimizing the mechanism designer's objective. In Section 3, we provide an illustrative example of a cake-cutting problem and report the computational results. Finally, we conclude in Section 4.

2. General model

Denote by $N = \{1, \ldots, n\}$ the set of players participating in the cake-sharing problem. The mechanism designer (or principal) in charge of dividing the cake seeks an allocation that optimizes a given criterion specified below. Denote by x_i the share of player i, with $x_i \in X_i$, and by θ_i the type of player i, with $\theta_i \in \Theta_i$. Player i's preferences are described by a utility function that depends only on these data, namely, $u_i(\theta_i, x_i) \in \mathbb{R}$. We assume that the principal and all players but i have the same incomplete information about player i's type θ_i. Player i's type is drawn from some probability distribution, called belief, $\tilde{\theta}_i \in \Delta(\Theta_i)$, known to the principal and the players. We suppose that the beliefs about the players' types are independent.

2.1. Direct mechanisms

The principal would ideally choose (possibly stochastically) an allocation $x = (x_1, \ldots, x_n) \in X$ with the knowledge of the type profile $\theta = (\theta_1, \ldots, \theta_n) \in \Theta$. Such a choice is known as a direct mechanism.

Definition 1. A direct mechanism $\mathcal{D}$ is a mapping of type profiles $\theta \in \Theta$ into probability distributions on outcomes $\mathcal{D}(\theta) \in \Delta(X)$, i.e., $\mathcal{D}: \Theta \to \Delta(X)$. We denote by $\mathbb{D}$ the set of direct mechanisms.

The principal's payoff is defined by the function $\mathcal{P}: \mathbb{D} \to \mathbb{R}$, and her optimization (or mechanism design) problem consists in maximizing $\mathcal{P}$ over a set of incentive-compatible direct mechanisms. The objective function may depend on the belief $\tilde{\theta} = (\tilde{\theta}_1, \ldots, \tilde{\theta}_n)$. For instance, the objective function could combine social efficiency and fairness as follows: