

Sequential Bidding in the Bailey-Cavallo Mechanism

Krzysztof R. Apt^{1,2} and Evangelos Markakis²

¹ CWI, Science Park 123, 1098 XG Amsterdam

² Institute of Logic, Language and Computation, University of Amsterdam
The Netherlands

³ Athens University of Economics and Business
Department of Informatics, Athens, Greece

Abstract. We are interested in mechanisms that maximize social welfare. In [2] this problem was studied for multi-unit auctions and for public project problems, and in each case social welfare undominated mechanisms were identified. One way to improve upon these optimality results is by allowing the players to move sequentially. With this in mind, we study here a sequential version of the Bailey-Cavallo mechanism, a natural mechanism that was proved to be welfare undominated in the simultaneous setting by [2]. Because of the absence of dominant strategies in the sequential setting, we focus on a weaker concept of an optimal strategy. We proceed by introducing natural optimal strategies and show that among all optimal strategies, the one we introduce generates maximal social welfare. Finally, we show that the proposed strategies form a safety level equilibrium and within the class of optimal strategies they also form a Pareto optimal ex-post equilibrium⁴.

1 Introduction

In many resource allocation problems a group of agents would like to determine who among them values a given object the most. A natural way to approach this problem is by viewing it as a single unit auction. Such an auction is traditionally used as a means of determining by a seller to which bidder and for which price the object is to be sold. The absence of a seller however changes the perspective and leads to different considerations since in our setting, the payments that the agents need to make flow out of the system (are "burned"). Instead of maximizing the revenue of the seller we are thus interested in maximizing the final social welfare.

This has led to the problem of finding mechanisms that are optimal in the sense that no other feasible, efficient and incentive compatible mechanism generates a larger social welfare. Recently, in [2] this problem was studied for two domains: multi-unit auctions with unit demand bidders and the public project problem of [8]. For the first domain a class of optimal mechanisms (which includes the Bailey-Cavallo mechanism) was identified, while for the second one

⁴ A full version of this work along with all the missing proofs is available at <http://pages.cs.aueb.gr/~markakis/research/pubs.html>.

it was proved that the pivotal mechanism is optimal. Other related aspects and objectives have also been recently studied in a series of works on redistribution and money-burning mechanisms, see among others [9, 12, 10, 13, 16, 7].

We continue this line of research by relaxing the assumption of simultaneity and allowing the players to move sequentially. This set up has been recently studied in [3] for the public project problem and here we consider such a modified setting for the case of single unit auctions. We call it *sequential bidding* as the concept of a “sequential auction” usually refers to a sequence of auctions, see, e.g. [14, chapter 15].

Hence we assume that there is a single object for sale and the players announce their bids sequentially in a fixed order. In contrast to the open cry auctions each player announces his bid *exactly once*. Once all bids have been announced, a mechanism is used to allocate the object to the highest bidder and determine the payments. Such a sequential setting can be very natural in many decision making or coordination problems without a central authority.

1.1 Results

We study here a sequential version of the Bailey-Cavallo mechanism of [5] and [6], as being a simplest, natural and most intuitive mechanism in the class of OEL mechanisms [11]. Our main results start in Section 4, where we first show that in a large class of sequential Groves auctions no *dominant strategies* exist. Therefore we settle on a weaker concept, that of an *optimal strategy*. An optimal strategy is a natural relaxation of the notion of dominant strategy, which also captures precisely the way a “prudent” player would play (see Lemma 1).

We proceed in Section 5 with proposing optimal strategies that differ from truth telling in the Bailey-Cavallo mechanism. We show that the proposed strategies yield maximal social welfare among all possible vectors of optimal strategies. Finally in Section 6 we further clarify the nature of the proposed strategies by studying what type of equilibrium they form. First we point that they do not form an ex-post equilibrium, a concept criticized in [4] and [1], where an alternative notion of a *safety-level equilibrium* was introduced for pre-Bayesian games. This concept captures the idea of an equilibrium in the case when each player is “prudent”. We prove that the proposed strategies form a safety-level equilibrium. We also show that our strategies form a Pareto optimal ex-post equilibrium within the class of optimal strategies.

2 Preliminaries

Assume that there is a finite set of possible outcomes or *decisions* D , a set $\{1, \dots, n\}$ of players where $n \geq 2$, and for each player i a set of *types* Θ_i and an (*initial*) *utility function* $v_i : D \times \Theta_i \rightarrow \mathbb{R}$. Let $\Theta := \Theta_1 \times \dots \times \Theta_n$. A *decision rule* is a function $f : \Theta \rightarrow D$. A mechanism is given by a pair of functions (f, t) , where f is the decision rule and $t = (t_1, \dots, t_n)$ is the tax function that determines the players’ payments. We assume that the (*final*) *utility function* for player

i is a function u_i defined by $u_i(d, t_1, \dots, t_n, \theta_i) := v_i(d, \theta_i) + t_i$. Thus, when the true type of player i is θ_i and his announced type is θ'_i , his final utility under the mechanism (f, t) is:

$$u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i) = v_i(f(\theta'_i, \theta_{-i}), \theta_i) + t_i(\theta'_i, \theta_{-i}),$$

Given a sequence $a := (a_1, \dots, a_j)$ of reals we denote the least l such that $a_l = \max_{k \in \{1, \dots, j\}} a_k$ by $\operatorname{argmax} a$. A **single item sealed bid auction**, is modelled by choosing $D = \{1, \dots, n\}$, each Θ_i to be the set of non-negative reals and $f(\theta) := \operatorname{argmax} \theta$. Hence the object is sold to the highest bidder and in the case of a tie we allocate the object to the player with the lowest index.⁵

By a **Groves auction** we mean a Groves mechanism for an auction setting (for details on Groves mechanisms see [15]). Below, given a sequence θ of reals we denote by θ^* its reordering in descending order. Then θ_k^* is the k th largest element in θ . For example, for $\theta = (1, 5, 0, 3, 2)$ we have $(\theta_{-2})_2^* = 2$ since $\theta_{-2} = (1, 0, 3, 2)$. The **Vickrey auction** is the pivotal mechanism for an auction (also referred to as the VCG mechanism). In it the winner pays the second highest bid.

The **Bailey-Cavallo** mechanism, in short **BC auction**, was originally proposed in [5] and [6]. To define it note that each Groves mechanism is uniquely determined by its **redistribution function** $r := (r_1, \dots, r_n)$. Given the redistribution function r , the tax for player i is defined by $t_i(\theta) := t_i^p(\theta) + r_i(\theta_{-i})$, where t_i^p is the tax of player i in the Vickrey auction. So we can think of a Groves auction as first running the pivotal mechanism and then redistributing some amount of the pivotal taxes.

The BC auction is a Groves mechanism defined by using the following redistribution function $r := (r_1, \dots, r_n)$ (assuming that $n \geq 3$):

$$r_i(\theta_{-i}) := \frac{(\theta_{-i})_2^*}{n}$$

It can be seen that the BC auction always yields at least as high social welfare as the pivotal mechanism. Note also that the aggregate tax is 0 when the second and third highest bids coincide.

3 Sequential mechanisms

We are interested in sequential mechanisms, where players announce their types according to a fixed order, say, $1, 2, \dots, n$. Each player i *observes* the actions announced by players $1, \dots, i-1$ and uses this information to decide which action to select. Thus a **strategy** of player i is now a function $s_i : \Theta_1 \times \dots \times \Theta_{i-1} \times \Theta_i \rightarrow \Theta_i$. Then if the vector of types that the players have is θ and the vector of strategies that they decide to follow is $s(\cdot) := (s_1(\cdot), \dots, s_n(\cdot))$, the resulting vector of the selected actions will be denoted by $[s(\cdot), \theta]$, where $[s(\cdot), \theta]$ is defined inductively by $[s(\cdot), \theta]_1 := s_1(\theta_1)$ and $[s(\cdot), \theta]_{i+1} := s_{i+1}([s(\cdot), \theta]_1, \dots, [s(\cdot), \theta]_i, \theta_{i+1})$.

⁵ If we make a different assumption on breaking ties, some of our proofs need to be adjusted, but similar results hold.

Given $\theta \in \Theta$ and $i \in \{1, \dots, n\}$ we denote the sequence $\theta_{i+1}, \dots, \theta_n$ by $\theta_{>i}$ and the sequence $\Theta_{i+1}, \dots, \Theta_n$ by $\Theta_{>i}$, and similarly with $\theta_{\leq i}$ and $\Theta_{\leq i}$.

A strategy $s_i(\cdot)$ of player i is called **dominant** if for all $\theta \in \Theta$, all strategies $s'_i(\cdot)$ of player i and all vectors $s_{-i}(\cdot)$ of strategies of players $j \neq i$

$$u_i((f, t)((s_i(\cdot), s_{-i}(\cdot)), \theta), \theta_i) \geq u_i((f, t)((s'_i(\cdot), s_{-i}(\cdot)), \theta), \theta_i),$$

We call a joint strategy $s(\cdot) = (s_1(\cdot), \dots, s_n(\cdot))$

- an **ex-post equilibrium** if for all $i \in \{1, \dots, n\}$, all strategies $s'_i(\cdot)$ of player i and all joint types $\theta \in \Theta$

$$u_i((f, t)((s_i(\cdot), s_{-i}(\cdot)), \theta), \theta_i) \geq u_i((f, t)((s'_i(\cdot), s_{-i}(\cdot)), \theta), \theta_i),$$

- a **safety-level equilibrium** if for all $i \in \{1, \dots, n\}$, all strategies $s'_i(\cdot)$ of player i and all $\theta_{\leq i} \in \Theta_{\leq i}$

$$\min_{\theta_{>i} \in \Theta_{>i}} u_i((f, t)((s_i(\cdot), s_{-i}(\cdot)), \theta), \theta_i) \geq \min_{\theta_{>i} \in \Theta_{>i}} u_i((f, t)((s'_i(\cdot), s_{-i}(\cdot)), \theta), \theta_i).$$

Intuitively, given the types $\theta_{\leq i} \in \Theta_{\leq i}$ of players $1, \dots, i$ and the vector $s(\cdot)$ of strategies used by the players, the quantity $\min_{\theta_{>i} \in \Theta_{>i}} u_i((f, t)((s(\cdot), \theta), \theta_i)$ is the minimum payoff that player i can guarantee to himself.

4 Sequential Groves auctions

In Groves auctions truth telling is a dominant strategy. In the case of sequential Groves auctions the situation changes as for a wide class, which includes sequential BC auctions no dominant strategies exist (except for the last player).

Theorem 1. *Consider a sequential Groves auction. Suppose that for player $i \in \{1, \dots, n-1\}$, the redistribution function r_i is such that there exists $z > 0$ such that $r_i(0, 0, \dots, z, 0, \dots, 0) \neq r_i(0, \dots, 0) + z$ (in the first term z is in the i th argument of r_i). Then no dominant strategy exists for player i .*

In light of this negative result, we would like to identify strategies that players could choose. We therefore focus on a concept that formalizes the idea that the players are “prudent” in the sense that they want to avoid the *winner’s curse* by winning the item at a too high price. Such a player i could argue as follows: if his actual type is no more than the currently highest bid among players $1, \dots, i-1$, then he can safely bid up to the currently highest bid. On the other hand, if his actual type is higher than the currently highest bid among players $1, \dots, i-1$, then he should bid truthfully (overbidding can result in a winner’s curse and underbidding can result in losing). Lemma 1 below shows that the above intuition is captured by the following definition.

Definition 1. *We call a strategy $s_i(\cdot)$ of player i **optimal** if for all $\theta \in \Theta$ and all $\theta'_i \in \Theta_i$*

$$u_i((f, t)(s_i(\theta_1, \dots, \theta_i), \theta_{-i}), \theta_i) \geq u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i).$$

By choosing truth telling as the strategies of players $j \neq i$ we see that each dominant strategy is optimal. For player n the concepts of dominant and optimal strategies coincide.

Definition 1 is a natural relaxation of the notion of dominant strategy as it calls for optimality w.r.t. a restricted subset of the other players' strategies. Call a strategy of player j *memoryless* if it does not depend on the types of players $1, \dots, j-1$. Then a strategy $s_i(\cdot)$ of player i is optimal if for all $\theta \in \Theta$ it yields a best response to all joint strategies of players $j \neq i$ in which the strategies of players $i+1, \dots, n$ are memoryless. In particular, an optimal strategy is a best response to the truth telling by players $j \neq i$.

The following lemma provides the announced characterization of optimal strategies. For any i , define $\bar{\theta}_i := \max_{j \in \{1, \dots, i-1\}} \theta_j$. We stipulate here and elsewhere that for $i=1$ we have $\bar{\theta}_1 = -1$ so that for $i=1$ we have $\theta_i > \bar{\theta}_i$.

Lemma 1. *In each sequential Groves auction a strategy $s_i(\cdot)$ is optimal for player i if and only if the following holds for all $\theta_1, \dots, \theta_i$:*

- (i) *Suppose $\theta_i > \bar{\theta}_i$ and $i < n$. Then $s_i(\theta_1, \dots, \theta_i) = \theta_i$.*
- (ii) *Suppose $\theta_i > \bar{\theta}_i$ and $i = n$. Then $s_i(\theta_1, \dots, \theta_i) > \bar{\theta}_i$.*
- (iii) *Suppose $\theta_i \leq \bar{\theta}_i$ and $i < n$. Then $s_i(\theta_1, \dots, \theta_i) \leq \bar{\theta}_i$.*
- (iv) *Suppose $\theta_i < \bar{\theta}_i$ and $i = n$. Then $s_i(\theta_1, \dots, \theta_i) \leq \bar{\theta}_i$.*

Note that no conclusion is drawn when $\theta_n = \max_{j \in \{1, \dots, n-1\}} \theta_j$. Player n can place then an arbitrary bid.

The following simple observation, see [3], provides us with a sufficient condition for checking whether a strategy is optimal in a sequential Groves mechanism.

Lemma 2. *Consider a Groves mechanism (f, t) . Suppose that $s_i(\cdot)$ is a strategy for player i such that for all $\theta \in \Theta$, $f(s_i(\theta_1, \dots, \theta_i), \theta_{-i}) = f(\theta)$. Then $s_i(\cdot)$ is optimal in the sequential version of (f, t) .*

In particular, truth telling is an optimal strategy.

5 Sequential BC auctions

As explained in the Introduction the BC mechanism cannot be improved upon in the simultaneous case, as shown in [2]. As we shall see here, the final social welfare can be improved in the sequential setting by appropriate optimal strategies that deviate from truth telling.

Theorem 1 applies for the BC auction, therefore no dominant strategies exist. We will thus focus on the notion of an optimal strategy. As implied by Lemma 2 many natural optimal strategies exist. In the sequel we will focus on the following optimal strategy which is tailored towards welfare maximization as we exhibit later on:

$$s_i(\theta_1, \dots, \theta_i) := \begin{cases} \theta_i & \text{if } \theta_i > \max_{j \in \{1, \dots, i-1\}} \theta_j \\ (\theta_1, \dots, \theta_{i-1})_1^* & \text{if } \theta_i \leq \max_{j \in \{1, \dots, i-1\}} \theta_j \\ & \text{and } i \leq n-1 \\ (\theta_1, \dots, \theta_{i-1})_2^* & \text{otherwise} \end{cases} \quad (1)$$

According to strategy $s_i(\cdot)$ if player i cannot be a winner when bidding truthfully he submits a bid that equals the highest current bid if $i < n$ or the second highest current bid if $i = n$. Note that $s_i(\cdot)$ is indeed optimal in the sequential BC auction, since Lemma 2 applies.

We now exhibit that within the universe of optimal strategies, if all players follow $s_i(\cdot)$, maximal social welfare is generated. Given θ and a vector of strategies $s(\cdot)$, define the final social welfare of a sequential mechanism (f, t) as:

$$SW(\theta, s(\cdot)) = \sum_{i=1}^n u_i((f, t)([s(\cdot), \theta]), \theta_i) = \sum_{i=1}^n v_i(f([s(\cdot), \theta]), \theta_i) + \sum_{i=1}^n t_i([s(\cdot), \theta]).$$

Theorem 2. *In the sequential BC auction for all $\theta \in \Theta$ and all vectors $s'(\cdot)$ of optimal players' strategies,*

$$SW(\theta, s(\cdot)) \geq SW(\theta, s'(\cdot))$$

where $s(\cdot)$ is the vector of strategies $s_i(\cdot)$ defined in (1).

The maximal final social welfare of the sequential BC auction under $s(\cdot)$ is always greater than or equal to the final social welfare achieved in a BC auction when players bid truthfully.

6 Implementation in Safety-level equilibrium

In this section we clarify the status of the strategies studied in Section 5 by analyzing what type of equilibrium they form. The notion of an ex-post equilibrium is somewhat problematic, since in pre-Bayesian games (the games we study here are a special class of such games) it has a different status than Nash equilibrium in strategic games. Indeed, as explained in [1], there exist pre-Bayesian games with finite sets of types and actions in which *no* ex-post equilibrium in mixed strategies exists.

The vector of strategies $s_i(\cdot)$ defined in (1) is *not* an ex-post equilibrium in the sequential BC auction. Indeed, take three players and $\theta = (1, 2, 5)$. Then for player 1 it is advantageous to deviate from $s_1(\cdot)$ strategy and submit, say 4. This way player 2 submits 4 and player's 1 final utility becomes $4/3$ instead of $2/3$.

We believe that an appropriate equilibrium concept for the (sequential) pre-Bayesian games is the safety-level equilibrium introduced by [4] and [1] and defined in Section 3. In the case of sequential mechanisms it captures a cautious approach by focusing on each player's guaranteed payoff in view of his lack of any information about the types of the players who bid after him. We have the following result.

Theorem 3. *The vector of strategies $s_i(\cdot)$ defined in (1) is a safety-level equilibrium in the sequential BC auction.*

One natural question is whether one can extend our Theorem 2 to show that our proposed vector of strategies in (1) generates maximal social welfare among all safety-level equilibria. The answer to this turns out to be negative as illustrated by the next example:

Example 1. Consider truth telling as the strategy for players $1, \dots, n-2$ and n . For any i , define $\hat{\theta}_i := \max_{j \in \{1, \dots, i-1\}} [s(\cdot), \theta]_j$. For player $n-1$ define the strategy:

$$s'_{n-1}(\theta_1, \dots, \theta_{n-1}) = \begin{cases} \hat{\theta}_{n-1} + \epsilon & \text{if } \theta_{n-1} > \hat{\theta}_{n-1}, \\ \theta_{n-1} & \text{otherwise.} \end{cases}$$

where ϵ is a positive number in the interval $(\hat{\theta}_{n-1}, \theta_{n-1})$. This vector of strategies forms a safety-level equilibrium (we omit the proof here due to lack of space). Consider now the vector $\theta = (0, 0, \dots, 1, 15, 16)$. The sum of taxes under the set of strategies we have defined will be $\frac{2\epsilon}{n}$. On the other hand, under the vector $s(\cdot)$ defined in (1), the sum of the taxes is $\frac{2}{n}(15-1)$. \square

The set of safety-level equilibria is quite large. The above example illustrates that we can construct many other safety-level equilibria, by slight deviations from the truth telling strategy. In fact, there are even equilibria in which some players overbid and yet for some type vectors they generate higher social welfare than our proposed strategies. These equilibria, however, may be unlikely to form by prudent players and Theorem 2 guarantees that among equilibria where all players are prudent our proposed strategies generate maximal welfare.

Finally, if we assume that players select only optimal strategies, then we could consider an ex-post equilibrium in the universe of optimal strategies. We have then the following positive result.

Theorem 4. *If we allow only deviations to optimal strategies, then in the sequential BC auction, the vector of strategies $s_i(\cdot)$ defined in (1) is an ex-post equilibrium that is also Pareto optimal.*

7 Final remarks

This paper and our previous recent work, [3], forms part of a larger research endeavour in which we seek to improve the social welfare by considering sequential versions of commonly used incentive compatible mechanisms. The main conclusion of [3] and of this work is that in the sequential version of single-item auctions and public project problems there exist optimal strategies that deviate from truth telling and can increase the social welfare. Further, the vector of these strategies generates the maximal social welfare among the vectors of all optimal strategies. Here, we also showed that the vector of the introduced strategies forms a safety-level equilibrium.

We would like to undertake a similar study of the sequential version of the incentive compatible mechanism proposed in [17], concerned with purchasing a shortest path in a network.

Acknowledgements

We thank Hervé Moulin and Moshe Tennenholtz for helpful discussions on the subject of this paper.

References

1. M. Aghassi and D. Bertsimas. Robust game theory. *Mathematical Programming*, 107:231–273, 2006.
2. K. R. Apt, V. Conitzer, M. Guo, and V. Markakis. Welfare undominated Groves mechanisms. In *Proc. 4th International Workshop on Internet and Network Economics (WINE 2008)*, Lecture Notes in Computer Science 5385, pages 426–437. Springer, 2008.
3. K. R. Apt and A. Estévez-Fernández. Sequential pivotal mechanisms for public project problems, 2008. Computing Research Repository (CoRR), <http://arxiv.org/abs/0810.1383>. To appear in *SAGT 2009: 2nd Symposium on Algorithmic Game Theory*.
4. I. Ashlagi, D. Monderer, and M. Tennenholtz. Resource selection games with unknown number of players. In *AAMAS '06: Proceedings 5th Int. Joint Conf. on Autonomous Agents and Multiagent Systems*, pages 819–825. ACM Press, 2006.
5. M. J. Bailey. The demand revealing process: To distribute the surplus. *Public Choice*, 91(2):107–126, 1997.
6. R. Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *AAMAS '06: Proceedings of the 5th Int. Joint Conf. on Autonomous Agents and Multiagent Systems*, pages 882–889. ACM Press, 2006.
7. R. Cavallo. Efficiency and redistribution in dynamic mechanism design. In *EC 2008: ACM Conference on Electronic Commerce*, pages 220–229, 2008.
8. E. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
9. M. Guo and V. Conitzer. Worst-case optimal redistribution of VCG payments. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce*, pages 30–39, New York, NY, USA, 2007. ACM Press.
10. M. Guo and V. Conitzer. Better redistribution with inefficient allocation in multi-unit auctions with unit demand. In *EC 2008: Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 210–219, 2008.
11. M. Guo and V. Conitzer. Optimal-in-expectation redistribution mechanisms. In *AAMAS 2008: Proceedings of the 7th Int. Joint Conf. on Autonomous Agents and Multi Agent Systems*, pages 1047–1054, 2008.
12. M. Guo and V. Conitzer. Undominated VCG redistribution mechanisms. In *AAMAS 2008: Proc. of 7th Int. Conf. on Autonomous Agents and Multi Agent Systems*, 2008.
13. J. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *STOC 2008: ACM Symposium on Theory of Computing*, pages 75–84, 2008.
14. V. Krishna. *Auction Theory*. Academic Press, New York, third edition, 2002.
15. A. Mas-Collel, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
16. H. Moulin. Efficient, strategy-proof and almost budget-balanced assignment, March 2007. Working Paper.
17. N. Nisan and A. Ronen. Algorithmic mechanism design (extended abstract). In *STOC*, pages 129–140, 1999.