Sequential Bidding in the Bailey-Cavallo Mechanism[*]

Krzysztof R. Apt [*][†]   Evangelos Markakis [‡]

Abstract

We are interested in mechanisms that maximize social welfare. In [2] this problem was studied for multi-unit auctions with unit demand bidders and for the public project problem, and in each case social welfare undominated mechanisms in the class of feasible and incentive compatible 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. The resulting social welfare can be larger than the one obtained in the simultaneous setting. 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 equilibrium.

---

1   Introduction

1.1   Motivation

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, that is in determining the winner at a minimal cost.

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 [9]. It was proved that for the first domain the optimal-in-expectation linear (OEL) redistribution mechanisms, which were introduced in [12] and include the Bailey-Cavallo mechanism, are optimal, while for the second one the pivotal mechanism is optimal. We note here that other related aspects and objectives have also been recently studied in a series of works on redistribution and money-burning mechanisms, see among others [10, 13, 11, 14, 17, 8].

1.2   Sequentiality

We continue here 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 where it was shown that in the sequential pivotal mechanism natural strategies exist that allow the players to increase the social welfare that would be generated if they moved simultaneously. In this paper 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. [15, 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.

1.3   Results

We study here a sequential version of the Bailey-Cavallo mechanism of [6] and [7], as being a simplest, natural and most intuitive mechanism in the class of OEL mechanisms that were proven in [2] to be welfare undominated in the simultaneous setting. To this end, after recalling the next section the necessary background material, we introduce in Section 3 the natural and general framework of sequential pre-Bayesian games. Our main results start in Section 4, where we first show that in a large class of sequential Groves auctions that include the sequential Bailey-Cavallo mechanism no *dominant strategies* exist, except for the last player.

Therefore we settle on a weaker concept, that of an *optimal strategy*. An optimal strategy for a player i yields for all type vectors a best response to all joint strategies of the other players, under the assumption that the players that follow i are memoryless (i.e., their strategy does not depend on the types of the previous players). For example truth telling is a memoryless strategy and it is also optimal. Our motivation for focusing on optimal strategies is two-fold. First, we show that in sequential Groves auctions optimal strategies capture precisely the way a "prudent" player would play. Such a player would never try to win the item, when his actual utility is less than the currently maximum bid, and would bid his true valuation otherwise.

We proceed in Section 5 with proposing optimal strategies that differ from truth telling in the Bailey-Cavallo mechanism. While truth telling focuses only on player's own payoff, the proposed strategies are additionally beneficial for the society: we show that their vector yields a maximal social welfare among all possible vectors of optimal strategies. In addition, the same outcome is realized under these optimal strategies as under truth telling.

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 they do not form an ex-post equilibrium, a concept criticized in [5] 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

We collect here the necessary background material. Readers familiar with Groves mechanisms can safely move to Subsection 2.3.

2.1 Direct mechanisms

Assume that there is a finite set of possible outcomes or **decisions** D, a set $\{1,\ldots,n\}$ of players where $n \geq 2$, and for each player i a set of **types** Θ_i and an (**initial**) **utility function** $v_i : D \times \Theta_i \to \mathbb{R}$. Let $\Theta := \Theta_1 \times \cdots \times \Theta_n$. A **decision rule** is a function $f : \Theta \to D$ and **tax function** is a function $t : \Theta \to \mathbb{R}^n$, that we sometimes write as $(t_1,\ldots,t_n)$.

In a **direct mechanism**, in short a **mechanism**, each player reports a type θ_i and based on this, the mechanism selects a decision and computes the taxes to be paid/received by every player. Hence a mechanism is given by a pair of functions (f,t), where f is the decision rule and t is the tax function that are applied to the vector of reported types.

We assume that the (**final**) **utility function** for player i is a function $u_i : D \times \mathbb{R}^n \times \Theta_i \to \mathbb{R}$ defined by $u_i(d,t_1,\ldots,t_n,\theta_i) := v_i(d,\theta_i)+t_i$. For each vector θ of announced types, if $t_i(\theta) \geq 0$, player i **receives** $t_i(\theta)$, and if $t_i(\theta) < 0$, he **pays** $|t_i(\theta)|$. Thus when the true type of player i is θ_i and his announced type is θ_i', his final utility is

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

where θ_{-i} are the types announced by the other players.

We say that a mechanism (f,t) is

3

- **efficient** if for all $\theta \in \Theta$, $f(\theta) \in \mathrm{argmax}_{d \in D} \sum_{i=1}^{n} v_i(d, \theta_i)$, i.e., the taken decision maximizes the initial social welfare,

- **feasible** if for all θ, $\sum_{i=1}^{n} t_i(\theta) \leq 0$, i.e., the mechanism does not need to be funded by an external source,

- **incentive compatible** if for all θ, $i \in \{1, \ldots, n\}$ and θ_i',

$$u_i((f, t)(\theta_i, \theta_{-i}), \theta_i) \geq u_i((f, t)(\theta_i', \theta_{-i}), \theta_i),$$

  i.e., mispreporting one's type does not pay off.

2.2   Groves mechanisms

Each **Groves mechanism** is an efficient direct mechanism (f, t) such that for every i the tax function is of the form $t_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j) + h_i(\theta_{-i})$, where $h_i : \Theta_{-i} \to \mathbb{R}$ is an arbitrary function.

Recall now the following crucial result, see e.g., [16].

Groves Theorem Every Groves mechanism (f, t) is incentive compatible.

A special Groves mechanism, called the **pivotal** or **VCG mechanism** is obtained using $h_i(\theta_{-i}) := -\max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j)$. In this case, the tax, $t_i^p(\theta)$, is defined by

$$t_i^p(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j) - \max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j),$$

which shows that the pivotal mechanism is feasible.

Groves mechanisms for single item auctions

Given a sequence $a := (a_1, \ldots, a_j)$ of reals we denote the least l such that $a_l = \max_{k \in \{1, \ldots, j\}} a_k$ by $\mathrm{argsmax}_{k \in \{1, \ldots, j\}} a_k$ or simply by $\mathrm{argsmax}\, a$. A **single item sealed bid auction**, in short an **auction**, is modelled by choosing $D = \{1, \ldots, n\}$, each Θ_i to be the set of non-negative reals and $f(\theta) := \mathrm{argsmax}\, \theta$. If $\theta_i \in \Theta_i$ is player i's valuation then

$$v_i(d, \theta_i) := \begin{cases} \theta_i & \text{if } d = i \\ 0 & \text{otherwise} \end{cases}$$

Here decision $d \in D$ indicates which player is the winner. 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.[1]

By a **Groves auction** we mean a Groves mechanism for an auction setting. Below, given a sequence θ of reals we denote by θ^* its reordering in descending order. Then θ_k^* is the kth largest element in θ. For example, for $\theta = (1, 5, 0, 3, 2)$ we have $(\theta_{-2})_2^* = 2$ since $\theta_{-2} = (1, 0, 3, 2)$.

---

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

The **Vickrey auction** is the pivotal mechanism for an auction. In it the winner pays the second highest bid, that is,

$$t_i^p(\theta) := \begin{cases} -\theta_2^* & \text{if argsmax } \theta = i \\ 0 & \text{otherwise} \end{cases}$$

The Bailey-Cavallo mechanism

The **Bailey-Cavallo** mechanism, in short **BC auction**, is the mechanism originally proposed in [6] and [7]. (In fact, Bailey's mechanism is not always the same as Cavallo's mechanism, but it is in the setting in which we study it.) To define it note that each Groves mechanism is uniquely determined by a **redistribution function** $r := (r_1, \ldots, r_n)$, where each $r_i : \Theta_{-i} \to \mathbb{R}$ is an arbitrary function. Given a redistribution function r the tax for player i is defined by $t_i(\theta) := t_i^p(\theta) + r_i(\theta_{-i})$. So we can think of a feasible Groves auction as first running the pivotal mechanism and then redistributing the pivotal taxes.

The BC auction is a Groves mechanism defined by using the following redistribution function $r := (r_1, \ldots, r_n)$ (to ensure that it is well-defined we need to assume that $n \geq 3$):

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

Given the sequence θ of submitted types, note that if player i is the first or the second highest bidder, then $(\theta_{-i})_2^* = \theta_3^*$. For the rest of the players, $(\theta_{-i})_2^* = \theta_2^*$.

Hence

$$\sum_{i=1}^n r_i(\theta_{-i}) = \sum_{i=1}^n \frac{(\theta_{-i})_2^*}{n} = \frac{n-2}{n}\theta_2^* + \frac{2}{n}\theta_3^*, \tag{1}$$

and

$$\sum_{i=1}^n t_i(\theta) = \frac{2}{n}(\theta_2^* - \theta_3^*). \tag{2}$$

So when the second highest type, θ_2^*, is strictly positive and the third highest type, θ_2^*, is non-negative, the BC auction yields a strictly higher social welfare than the pivotal mechanism. Note also that the aggregate tax is 0 when the second and third highest bids coincide.

In some situations it is useful to employ Groves auctions that maximize the final social welfare. This is for example the case when, as discussed in the introduction, the players want to determine at a minimal cost, using an incentive compatible mechanism, who among them values a given object most. For such applications by the above observation Vickrey auction is not an appropriate mechanism. In [2] it was proved that the BC auction is an appropriate mechanism in the sense that no Groves auction always generates a larger or equal social welfare than BC and sometimes a strictly larger social welfare. This explains our interest in BC auctions.

2.3 A useful lemma

In what follows the following observation will be helpful.

Lemma 2.1 *In each Groves auction for all $\theta \in \Theta$*

   (i) if $\theta_i > \max_{j \neq i} \theta_j$ and $argsmax\,(\theta_i', \theta_{-i}) \neq i$, then $\theta_i + t_i(\theta) > t_i(\theta_i', \theta_{-i})$,

   (ii) if $\theta_i' > \max_{j \neq i} \theta_j > \theta_i$, then $t_i(\theta) > \theta_i + t_i(\theta_i', \theta_{-i})$.

Part (i) states that if player i is a clear winner, given θ_{-i}, then it is strictly better for him to submit his true bid, θ_i, than a losing bid. In turn, part (ii) states that if player i is a clear loser, given θ_{-i}, it is strictly better for him to submit his true bid than a strictly winning bid.

Proof. Both properties clearly hold for Vickrey auction. In an arbitrary Groves auction the tax for player i is defined by $t_i(\theta) := t_i^p(\theta) + r_i(\theta_{-i})$, where $t_i^p(\theta)$ is his tax in Vickrey auction and $(r_1, \ldots, r_n)$ is a redistribution function (so r_i depends only on θ_{-i}). So both properties extend to an arbitrary Groves auction. $\qquad\square$

The above lemma will allow us to establish in Section 4 some general results about sequential Groves auctions, which we will later apply to sequential BC auctions.

3 Sequential pre-Bayesian games

Before we consider sequential mechanisms we introduce the corresponding general game-theoretic framework. To this end we introduce a modification of pre-Bayesian games (see, e.g., [5]). These games are distinguished by the fact that each player has a private type on which he can condition his strategy.

3.1 Pre-Bayesian games

Recall first[2] that a **pre-Bayesian game** for n players consists of

- a set A_i of **actions**,

- a set Θ_i of **types**,

- a **payoff function** $p_i : A_1 \times \ldots \times A_n \times \Theta_i \to \mathbb{R}$,

for each player i.

Let $A := A_1 \times \ldots \times A_n$ and $\Theta := \Theta_1 \times \ldots \times \Theta_n$. In a pre-Bayesian game Nature moves first and provides each player i with a type $\theta_i \in \Theta_i$. Each player knows only his type. Subsequently the players simultaneously select their actions. The payoff function of each player now depends on his type, so after each player selected his action, each player knows his payoff but does not know the payoffs of the other players.

A **strategy** for player i in a pre-Bayesian game is a function $s_i : \Theta_i \to A_i$. A strategy $s_i(\cdot)$ for player i is called **dominant** if for all $a \in A$ and $\theta_i \in \Theta_i$

$$p_i(s_i(\theta_i), a_{-i}, \theta_i) \geq p_i(a_i, a_{-i}, \theta_i).$$

---

[2] The definition given in [5] is considerably more general in that it allows for different states with different subsets of the players active in each state.

A pre-Bayesian game is of a ***revelation-type*** if $A_i = \Theta_i$ for all $i \in \{1, \ldots, n\}$. So in a revelation-type pre-Bayesian game the strategies of a player are the functions on his set of types.

As explained in [5] mechanism design can be viewed as an instance of the revelation-type pre-Bayesian games. Indeed, with each mechanism (f, t) we can associate a revelation-type pre-Bayesian game, by defining each payoff function p_i by

$$p_i((\theta_i', \theta_{-i}), \theta_i) := u_i((f, t)(\theta_i', \theta_{-i}), \theta_i).$$

It is then easy to check that a mechanism (f, t) is incentive compatible iff in the associated pre-Bayesian game for each player truth telling is a dominant strategy.

3.2   Sequential pre-Bayesian games

In this modification of pre-Bayesian games Nature moves again first and provides a type $\theta_i \in \Theta_i$ for each player i and an order, say (for simplicity) $1, \ldots, n$, in which the players sequentially select their actions. Each player plays only once according to the provided order. Each player i *observes* the actions announced by players $1, \ldots, i-1$ and uses this information to decide which action to select.

We call the resulting game a ***sequential pre-Bayesian game***. In this game a ***strategy*** of player i is now a function

$$s_i : A_1 \times \ldots \times A_{i-1} \times \Theta_i \to A_i.$$

Then if the vector of types that the players receive is θ and the vector of strategies that they decide to follow is $s(\cdot) := (s_1(\cdot), \ldots, s_n(\cdot))$, the resulting vector of the selected actions will be denoted by $[s(\cdot), \theta]$, where $[s(\cdot), \theta]$ is defined inductively by

$$[s(\cdot), \theta]_1 := s_1(\theta_1) \text{ and } [s(\cdot), \theta]_{i+1} := s_{i+1}([s(\cdot), \theta]_1, \ldots, [s(\cdot), \theta]_i, \theta_{i+1}).$$

We now modify in a natural way the standard notions that have been introduced for pre-Bayesian games to sequential pre-Bayesian games. Given $\theta \in \Theta$ and $i \in \{1, \ldots, n\}$ we denote the sequence $\theta_{i+1}, \ldots, \theta_n$ by $\theta_{>i}$ and the sequence $\Theta_{i+1}, \ldots, \Theta_n$ 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 $j \neq i$

$$p_i([(s_i(\cdot), s_{-i}(\cdot)), \theta], \theta_i) \geq p_i([(s_i'(\cdot), s_{-i}(\cdot)), \theta], \theta_i),$$

- ***safety-level dominant*** if for all $\theta_{\leq i} \in \Theta_{\leq i}$, all strategies $s_i'(\cdot)$ of player i and all vectors $s_{-i}(\cdot)$ of strategies of players $j \neq i$

$$\min_{\theta_{>i} \in \Theta_{>i}} p_i([(s_i(\cdot), s_{-i}(\cdot)), \theta], \theta_i) \geq \min_{\theta_{>i} \in \Theta_{>i}} p_i([(s_i'(\cdot), s_{-i}(\cdot)), \theta], \theta_i).$$

Intuitively, given the types $\theta_{\leq i} \in \Theta_{\leq i}$ of players $1, \ldots, i$ and the vector $s(\cdot)$ of strategies used by the players,

$$\min_{\theta_{>i} \in \Theta_{>i}} p_i([s(\cdot), \theta], \theta_i)$$

is the guaranteed payoff for player i. Clearly every dominant strategy is safety-level dominant.

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

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

$$p_i([(s_i(\cdot), s_{-i}(\cdot)), \theta], \theta_i) \geq p_i([(s_i'(\cdot), s_{-i}(\cdot)), \theta], \theta_i),$$

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

$$\min_{\theta_{>i} \in \Theta_{>i}} p_i([(s_i(\cdot), s_{-i}(\cdot)), \theta], \theta_i) \geq \min_{\theta_{>i} \in \Theta_{>i}} p_i([(s_i'(\cdot), s_{-i}(\cdot)), \theta], \theta_i).$$

3.3 Sequential mechanisms

We are interested in sequential mechanisms, in particular sequential auction mechanisms, in which the players announce their types sequentially. Such mechanisms can be readily introduced as a special case of sequential pre-Bayesian games. The set-up is the same as in direct mechanisms, except that there is a fixed order of the players, which determines the order in which they announce their bid.

Then we focus on revelation-types sequential pre-Bayesian games where for a sequential mechanism (f, t), each payoff function p_i is defined by $p_i((\theta_i', \theta_{-i}), \theta_i) := u_i((f, t)(\theta_i', \theta_{-i}), \theta_i)$.

A ***strategy*** of player i in a sequential mechanism is a function $s_i : \Theta_1 \times \ldots \times \Theta_i \to \Theta_i$. In this context truth telling, as a strategy, is represented by the projection function $\pi_i(\cdot)$, defined by $\pi_i(\theta_1, \ldots, \theta_i) := \theta_i$.

4 Sequential Groves auctions

In Groves auctions truth telling is a dominant strategy. We now show that 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). The same holds for safety-level dominant strategies.

Theorem 4.1 *Consider a sequential Groves auction.*

- (i) *Suppose that for player $i \in \{1, \ldots, n-1\}$, the redistribution function r_i is such that there exists $z > 0$ such that $r_i(0, 0, \ldots, z, 0, \ldots, 0) \neq r_i(0, \ldots, 0) + z$ (in the first term z is in the ith argument of r_i). Then no dominant strategy exists for player i.*

- (ii) *Suppose that the redistribution function r satisfies that for every $i \in \{1, \ldots, n-2\}$, and for every $x \geq 0$, there exists $\epsilon > 0$ such that $r_i(0, 0, \ldots, x, x, \ldots, x) < r_i(0, 0, \ldots, x+\epsilon, \ldots, x+\epsilon)$. Then no safety-level dominant strategy exists for player $i \in \{1, \ldots, n-2\}$.*

Proof. See Appendix A. $\qquad\square$

In light of this negative result, we would like to identify strategies that players could choose. A natural possibility is to focus on **undominated strategies**, i.e., strategies that are not weakly dominated by another strategy. Unfortunately, as shown in the next section, in the sequential BC auctions *all* strategies for players $i \in \{1, \ldots, n-2\}$ are undominated.

Therefore we 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 should proceed as follows:

- If his actual type is no more than the currently highest bid among players $1, \ldots, i-1$, then he can safely bid up to the currently highest bid (recall that by the convention we have made, ties are in favor of the player with the lowest index).

- If his actual type is higher than the currently highest bid among players $1, \ldots, i-1$, then he should bid truthfully. (Overbidding can result in a winner's curse and underbidding can result in losing.)

Lemma 4.3 below shows that the above intuition is captured by the following definition.

Definition 4.2 We call a strategy $s_i(\cdot)$ of player i **optimal** if for all $\theta \in \Theta$ and all $\theta_i' \in \Theta_i$

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

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 4.2 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 $1, \ldots, j-1$. 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 $i+1, \ldots, n$ 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. We stipulate here and elsewhere that for $i = 1$ we have $\max_{j \in \{1, \ldots, i-1\}} \theta_j = -1$ so that for $i = 1$ we have $\theta_i > \max_{j \in \{1, \ldots, i-1\}} \theta_j$.

Lemma 4.3 *In each sequential Groves auction a strategy $s_i(\cdot)$ is optimal for player i if and only if the following holds for all $\theta_1, \ldots, \theta_i$:*

(i) *Suppose $\theta_i > \max_{j \in \{1, \ldots, i-1\}} \theta_j$ and $i < n$. Then $s_i(\theta_1, \ldots, \theta_i) = \theta_i$.*

(ii) *Suppose $\theta_i > \max_{j \in \{1, \ldots, i-1\}} \theta_j$ and $i = n$. Then $s_i(\theta_1, \ldots, \theta_i) > \max_{j \in \{1, \ldots, i-1\}} \theta_j$.*

(iii) *Suppose $\theta_i \leq \max_{j \in \{1, \ldots, i-1\}} \theta_j$ and $i < n$. Then $s_i(\theta_1, \ldots, \theta_i) \leq \max_{j \in \{1, \ldots, i-1\}} \theta_j$.*

(iv) *Suppose $\theta_i < \max_{j \in \{1, \ldots, i-1\}} \theta_j$ and $i = n$. Then $s_i(\theta_1, \ldots, \theta_i) \leq \max_{j \in \{1, \ldots, i-1\}} \theta_j$.*

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.

Proof. See Appendix A. □

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 4.4 *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 optimal in the sequential version of each Groves mechanism.

5 Sequential BC auctions

We will now focus on the sequential Bailey-Cavallo auction. As explained in the Introduction this 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 using appropriate optimal strategies that deviate from truth telling.

Theorem 4.1 applies for the BC auction, therefore no dominant strategies exist. Given a sequential pre-Bayesian game call a strategy $s_i(\cdot)$ of player i **weakly dominated** by a strategy $s_i'(\cdot)$ if for all $\theta \in \Theta$ and all vectors $s_{-i}(\cdot)$ of strategies of players $j \neq i$

$$p_i([(s_i'(\cdot), s_{-i}(\cdot)), \theta], \theta_i) \geq p_i([(s_i(\cdot), s_{-i}(\cdot)), \theta], \theta_i),$$

and for some $\theta \in \Theta$ and some vector $s_{-i}(\cdot)$ of strategies of players $j \neq i$

$$p_i([(s_i'(\cdot), s_{-i}(\cdot)), \theta], \theta_i) > p_i([(s_i(\cdot), s_{-i}(\cdot)), \theta], \theta_i).$$

In turn call a strategy **undominated** if it is not weakly dominated by any other strategy.

It is natural to assume that all players choose undominated strategies. Unfortunately, the following positive result renders this notion useless in the study of the sequential BC auction.

Theorem 5.1 *In the sequential BC auction for every player $i \in \{1, \dots, n-2\}$ every strategy is undominated.*

Proof. See Appendix B. □

We will therefore focus on the notion of an optimal strategy that is weaker than that of a dominant strategy. As implied by Lemma 4.4 many natural optimal strategies exist, e.g., strategies that given θ_{-i} produce the same winner as truth telling. In the sequel we will focus on the following optimal strategy which is tailored towards welfare maximization as we exhibit later on:

$$s_i(\theta_1, \dots, \theta_i) := \begin{cases} \theta_i & \text{if } \theta_i > \max_{j \in \{1,\dots,i-1\}} \theta_j \\ (\theta_1, \dots, \theta_{i-1})_1^* & \text{if } \theta_i \leq \max_{j \in \{1,\dots,i-1\}} \theta_j \\ & \text{and } i \leq n-1 \\ (\theta_1, \dots, \theta_{i-1})_2^* & \text{otherwise} \end{cases} \tag{3}$$

So according to strategy $s_i(\cdot)$ if player i cannot be a winner when bidding truthfully ($\theta_i \leq \max_{j \in \{1,\ldots,i-1\}} \theta_j$) 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 it does not change the decision that would be taken if player i were truthful, so Lemma 4.4 applies.

We now show that within the universe of optimal strategies, if all players follow $s_i(\cdot)$, maximal social welfare is generated for all $\theta \in \Theta$. 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 5.2 *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 (3).

Proof. See Appendix B. □

This maximal final social welfare of the sequential BC auction under $s(\cdot)$ equals

$$SW(\theta, s(\cdot)) = \theta_i - \frac{2}{n}([s(\cdot), \theta]_2^* - [s(\cdot), \theta]_3^*)$$

where $i = \text{argsmax}\,\theta$. This is always greater than or equal to the final social welfare achieved in a BC auction when players bid truthfully, which is $\theta_i - 2/n(\theta_2^* - \theta_3^*)$. Additionally, for some inputs, for example when the last player is not the winner, it is strictly greater. It is also strictly greater when the last player is the winner but the last but one player does not have the highest type among the first $n - 1$ players.

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 problematic, since in pre-Bayesian game 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. By Groves theorem for all sequential Groves auctions the vector of truth telling strategies $\pi_i(\cdot)$ is an ex-post equilibrium. However, we do not know of any other ex-post equilibria in which players $i \neq n$ use a different strategy than $\pi_i(\cdot)$.

In particular, the vector of strategies $s_i(\cdot)$ defined in (3) 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 [5] and [1] and defined in Section 3. In the case of sequential mechanisms it captures a cautious approach by focusing for each player on his guaranteed payoff in view of his lack of any information about the types of the players who bid after him. This is similar in spirit to max-min strategies with the crucial difference that the possible malicious behaviour of the players who follow is replaced by the possibility of facing an unfortunate sequence of types of the players who follow. We have the following result.

Theorem 6.1 *The vector of strategies $s_i(\cdot)$ defined in (3) is a safety-level equilibrium in the sequential BC auction.*

Proof. See Appendix C. □

We illustrate next that not all vectors of optimal strategies form safety-level equilibria.

Example 6.2 Consider truth telling as the strategy for players $1, ..., n-2$. For any i, define $\hat{\theta}_i := \max_{j \in \{1,...,i-1\}} [s(\cdot), \theta]_j$. For player $n-1$ consider the strategy:

$$s_{n-1}(\theta_1, ..., \theta_{n-1}) = \begin{cases} \theta_{n-1} & \text{if } \theta_{n-1} > \hat{\theta}_{n-1}, \\ \max\{\hat{\theta}_{n-1} - 5, 0\} & \text{otherwise.} \end{cases}$$

Finally for player n define:

$$s_n(\theta_1, ..., \theta_n) = \begin{cases} \theta_n & \text{if } \theta_n > \hat{\theta}_n, \\ 0 & \text{if } \theta_n \leq \hat{\theta}_n \text{ and } \theta_{n-1} = \hat{\theta}_n, \\ \theta_{n-1} + \epsilon & \text{otherwise.} \end{cases}$$

where ϵ is some positive number in the interval $(0, \hat{\theta}_n - \theta_{n-1})$.

It is easy to check that all these are optimal strategies. Yet the vector $s(\cdot)$ does not form a safety-level equilibrium. To see this consider the type vector $\theta = (5, 10, 0, 0, ..., 0)$. Under $s(\cdot)$, the announced vector will be $(5, 10, 0, ..., 5, 5 + \epsilon)$ and the minimum (over θ_n) payoff to player $n-1$ is $(5 + \epsilon)/n$. On the other hand, if player $n-1$ switches to the strategy in which he bids $\max\{\hat{\theta}_{n-1} - 4, 0\}$, when his type is not the highest, his minimum payoff over all θ_n will be $(6 + \epsilon)/n$. □

One natural question is whether one can extend our Theorem 5.2 to show that our proposed vector of strategies in (3) generates maximal social welfare among all safety-level equilibria. The answer to this turns out to be negative as there exist type vectors at which other equilibria can generate higher welfare. This is illustrated by the next example:

Example 6.3 Consider truth telling as the strategy for players $1, ..., n-2$ and n. For player $n-1$ define the strategy:

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

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 $\theta = (0, 0, ...1, 15, 16)$. 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 (3), the sum of the taxes is $\frac{2}{n}(15 - 1)$. □

12

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 5.2 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 6.4 *If we allow only deviations to optimal strategies, then in the sequential BC auction, the vector of strategies $s_i(\cdot)$ defined in (3) is an ex-post equilibrium that is also Pareto optimal.*

Proof. See Appendix C. □

7 Final remarks

This paper and our previous recent work, [3], forms part of a larger research endevour in which we seek to improve the social welfare by considering sequential versions of commonly used incentive compatible mechanisms. The main conclusion of [3] was that in the sequential version of the public project problem 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 showed analogous results for the sequential Bailey-Cavallo mechanism. Additionally, we showed the absence of dominant and safety-level strategies and proved that the vector of the introduced strategies forms a safety-level equilibrium.

One could carry out a similar analysis for other sequential Groves auctions with relatively simple redistribution functions. However, we are not aware of a uniform approach to identify in sequential Groves auctions optimal strategies that maximize social welfare. In fact, not all Groves mechanisms behave the same with regard to these concepts. To illustrate this, we have included in Appendix D a similar analysis for the sequential Vickrey auction which highlights some some of the differences with the sequential BC auction.

A natural question is also whether analogous results can be established for other types of auctions. We believe that our results can be extended to multi-unit auctions with unit demand bidders. Finally, we would like to undertake similar study of the sequential version of the incentive compatible mechanism proposed in [18], concerned with purchasing a shortest path in a network.

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] K. R. Apt and E. Markakis. Optimal strategies in sequential bidding (short paper). In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1189–1190, 2009.

[5] 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.

[6] M. J. Bailey. The demand revealing process: To distribute the surplus. *Public Choice*, 91(2):107–126, 1997.

[7] 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.

[8] R. Cavallo. Efficiency and redistribution in dynamic mechanism design. In *EC 2008: ACM Conference on Electronic Commerce*, pages 220–229, 2008.

[9] E. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] J. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *STOC 2008: ACM Symposium on Theory of Computing*, pages 75–84, 2008.

[15] V. Krishna. *Auction Theory*. Academic Press, New York, third edition, 2002.

[16] A. Mas-Collel, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[17] H. Moulin. Efficient, strategy-proof and almost budget-balanced assignment, March 2007. Working Paper.

[18] N. Nisan and A. Ronen. Algorithmic mechanism design (extended abstract). In *STOC*, pages 129–140, 1999.

Appendix

A Proofs of Results of Section 4

Proof of Theorem 4.1.

(i) Suppose otherwise. Let $s_i(\cdot)$ be a dominant strategy of player i. Suppose that the redistribution function r_i satisfies that there exists $z > 0$ such that $r_i(0, 0, \ldots, z, 0, \ldots, 0) - r_i(0, \ldots, 0) < z$. Choose now $\theta \in \Theta$ such that $\theta_i = 2z$ and $\theta_j = 0$ for $j \neq i$. It is not difficult to prove that for a dominant strategy $s_i(\cdot)$, if θ_i is higher than all previously declared bids, then $s_i(\theta_1, \ldots, \theta_i) = \theta_i$ (the argument is the same as in Lemma 4.3(i)). Take now truth telling as strategy $s_j(\cdot)$ for players $j < i$ and the following strategy for players $j > i$:

$$s_j := \max\{\theta_{j-1} - z, 0\}$$

Then $u_i((f, t)([(s_i(\cdot), s_{-i}(\cdot)), \theta]), \theta_i) = 2z - z + r_i(0, 0, \ldots, z, 0, \ldots, 0)$, while for the strategy $s_i'(\cdot)$ such that $s_i'(\theta_1, \ldots, \theta_i) = z$ we have $u_i((f, t)([(s_i'(\cdot), s_{-i}(\cdot)), \theta]), \theta_i) = 2z + r_i(0, 0, \ldots, 0)$. By our assumptions on r_i, this shows that $s_i(\cdot)$ cannot be a dominant strategy. Consider now the case when for every $z > 0$, $r_i(0, 0, \ldots, z, 0, \ldots, 0) - r_i(0, \ldots, 0) \geq z$. Suppose there exists z, where this holds with strict inequality. Choose this z and consider the type vector θ where now $\theta_i = z$ and $\theta_j = 0$ for $j \neq i$. Again consider the same strategies $s_{-i}(\cdot)$ for the other players. If $s_i(\cdot)$ is dominant then if player i plays $s_i(\cdot)$, his payoff is $z + r_i(0, 0, \ldots, 0)$. On the other hand if he submits $2z$, then his payoff would be $r_i(0, 0, \ldots, z, 0 \ldots, 0)$, which is higher. Hence there can be no dominant strategy for player i unless for every $z > 0$, $r_i(0, 0, \ldots, z, 0, \ldots, 0) = r_i(0, \ldots, 0) + z$.

(ii) Let $s_i(\cdot)$ be a safety level dominant strategy of player i. Consider the type vector where $\theta_j = 0$ for $\jmath \in \{1, ..., i-1\}$ and take truth telling as the strategy of players $1, ..., i-1$. Suppose $s_i(0, 0, ..., \theta_i) = \theta_i'$. Then for $j = 1+1, ..., n$, consider the strategy $s_j(\cdot) = \theta_i$. If player i sticks to s_i, his minimum payoff over all $\theta_{>i}$ is $r_i(0, 0, ..., \theta_i', \theta_i', ..., \theta_i')$. On the other hand if he deviates to $s_i(0, 0, ..., \theta_i) = \theta_i' + \epsilon$, the resulting minimum payoff over $\theta_{>i}$ will be $r_i(0, 0, ..., \theta_i' + \epsilon, ..., \theta_i' + \epsilon)$. By the condition in the statement of the theorem this completes the proof. □

Proof of Lemma 4.3.

For the forward direction, let $s_i(\cdot)$ be an optimal strategy. We analyze each of the four cases outlined in the statement of the Lemma.

(i) Suppose otherwise. If $s_i(\theta_1, \ldots, \theta_i) < \theta_i$, then take $\epsilon > 0$ such that $s_i(\theta_1, \ldots, \theta_i) + \epsilon < \theta_i$ and set $\theta_{i+1} := \ldots := \theta_n := s_i(\theta_1, \ldots, \theta_i) + \epsilon$. Then $\theta_i > \max_{j \neq i} \theta_j > s_i(\theta_1, \ldots, \theta_i)$, so by Lemma 2.1($i$)

$$u_i((f, t)(\theta_i, \theta_{-i}), \theta_i) = \theta_i + t_i(\theta) > t_i(s_i(\theta_1, \ldots, \theta_i), \theta_{-i}) = u_i((f, t)(s_i(\theta_1, \ldots, \theta_i), \theta_{-i}), \theta_i).$$

15

This contradicts the optimality of $s_i(\cdot)$.

If $s_i(\theta_1, \ldots, \theta_i) > \theta_i$, then take $\epsilon > 0$ such that $\theta_i + \epsilon < s_i(\theta_1, \ldots, \theta_i)$ and set $\theta_{i+1} := \ldots :=$ $\theta_n := \theta_i + \epsilon$. Then $f(\theta) \neq i$, so $u_i((f,t)(\theta_i, \theta_{-i}), \theta_i) = r_i(\theta_{-i})$, where $t_i(\theta) := t_i^p(\theta) + r_i(\theta_{-i})$. On the other hand $f(s_i(\theta_1, \ldots, \theta_i), \theta_{-i}) = i$ and consequently

$$u_i((f,t)(s_i(\theta_1, \ldots, \theta_i), \theta_{-i}), \theta_i) = \theta_i - (\theta_i + \epsilon) + r_i(\theta_{-i}) < r_i(\theta_{-i}).$$

This again contradicts the optimality of $s_i(\cdot)$.

(ii) Suppose otherwise, that is $s_n(\theta_1, \ldots, \theta_n) \leq \max_{j \neq n} \theta_j$. Then $\mathrm{argsmax}(s_n(\theta_1, \ldots, \theta_n), \theta_{-n}) \neq n$, so by Lemma 2.1$(i)$

$$u_n((f,t)(\theta_n, \theta_{-n}), \theta_n) = \theta_n + t_n(\theta) > t_n(s_n(\theta_1, \ldots, \theta_n), \theta_{-n}) = u_n((f,t)(s_n(\theta_1, \ldots, \theta_n), \theta_{-n}), \theta_n).$$

This contradicts the optimality of $s_n(\cdot)$.

(iii) Suppose otherwise, i.e., $s_i(\theta_1, \ldots, \theta_i) > \max_{j \in \{1, \ldots, i-1\}} \theta_j$. Then take $\epsilon > 0$ such that $\max_{j \in \{1, \ldots, i-1\}} \theta_j + \epsilon < s_i(\theta_1, \ldots, \theta_i)$ and set

$$\theta_{i+1} := \ldots := \theta_n := \max_{j \in \{1, \ldots, i-1\}} \theta_j + \epsilon.$$

Then $s_i(\theta_1, \ldots, \theta_i) > \max_{j \neq i} \theta_j > \theta_i$, so by Lemma 2.1$(ii)$

$$u_i((f,t)(\theta_i, \theta_{-i}), \theta_i) = \theta_i > \theta_i + t_i(s_i(\theta_1, \ldots, \theta_i), \theta_{-i}) = u_i((f,t)(s_i(\theta_1, \ldots, \theta_i), \theta_{-i}), \theta_i).$$

This contradicts the optimality of $s_i(\cdot)$.

(iv) Suppose otherwise, that is $s_n(\theta_1, \ldots, \theta_n) > \max_{j \neq n} \theta_j$. Then $\mathrm{argsmax}(s_n(\theta_1, \ldots, \theta_n), \theta_{-n}) = n$, so by Lemma 2.1$(ii)$

$$u_n((f,t)(\theta_n, \theta_{-n}), \theta_n) = t_n(\theta) > \theta_n + t_n(s_n(\theta_1, \ldots, \theta_n), \theta_{-n}) = u_n((f,t)(s_n(\theta_1, \ldots, \theta_n), \theta_{-n}), \theta_n).$$

This contradicts the optimality of $s_n(\cdot)$.

For the reverse direction, suppose a strategy $s_i(\cdot)$ satisfies the conditions in the statement of the Lemma. Then one can easily check that for $i \in \{1, \ldots, n-1\}$ Lemma 4.4 applies and therefore $s_i(\cdot)$ is optimal. If $i = n$, then we will show that a strategy that satisfies the conditions for $i = n$ is dominant and hence optimal. By Groves theorem it suffices to prove that

$$u_n((f,t)((s_n(\theta_1, \ldots, \theta_n), \theta_{-n})), \theta_n) = u_n((f,t)(\theta), \theta_n). \tag{4}$$

Case 1. $f(s_n(\theta_1, \ldots, \theta_n), \theta_{-n}) = f(\theta)$. Then (4) holds since by the definition of the Groves mechanism for some $h_n : \Theta_{-n} \to \mathbb{R}$ we have for all θ_n'

$$u_n((f,t)(\theta_n', \theta_{-n}), \theta_n) = \sum_{j=1}^{n} v_j(f(\theta_n', \theta_{-n}), \theta_j) + h_n(\theta_{-n}).$$

Case 2. $f(s_n(\theta_1, \ldots, \theta_n), \theta_{-n}) \neq f(\theta)$. Then by the assumptions of $s_n(\cdot)$ we have $s_n(\theta_1, \ldots, \theta_n) > \theta_n = \max_{j \neq n} \theta_j$. Hence

$$u_n((f,t)((s_n(\theta_1, \ldots, \theta_n), \theta_{-n})), \theta_n) = \theta_n - \max_{j \neq n} \theta_j + r_n(\theta_{-n})$$
$$= r_n(\theta_{-n}) = u_n((f,t)(\theta), \theta_n),$$

where r is the redistribution function of (f,t). $\qquad \square$

In the proofs that follow we shall need the following helpful conclusion.

16

Corollary A.1 *In each sequential Groves auction for all $\theta \in \Theta$ and all vectors $s(\cdot)$ of optimal players' strategies*

(i) *for all $i \in \{1, \ldots, n-1\}$, $\max_{j \in \{1,\ldots,i\}} [s(\cdot), \theta]_j = \max_{j \in \{1,\ldots,i\}} \theta_j$,*

(ii) *for all $i \in \{1, \ldots, n-1\}$, $\text{argsmax}_{j \in \{1,\ldots,i\}} [s(\cdot), \theta]_j = \text{argsmax}_{j \in \{1,\ldots,i\}} \theta_j$,*

(iii) *for all $i \in \{1, \ldots, n-1\}$, if $\theta_i > \max_{j \in \{1,\ldots,i-1\}} [s(\cdot), \theta]_j$, then $\theta_i > \max_{j \in \{1,\ldots,i-1\}} \theta_j$ and for any vector of optimal players' strategies $s'(\cdot)$ we have $\theta_i > \max_{j \in \{1,\ldots,i-1\}} [s'(\cdot), \theta]_j$,*

(iv) *either $f([s(\cdot), \theta]) = f(\theta)$ or if not, $\theta_n = \max_{i \neq n} \theta_i$, $[s(\cdot), \theta]_n > \theta_n$ and $f([s(\cdot), \theta]) = n$.*

Informally, items $(i) - (iii)$ state that when each player follows an optimal strategy, the first $n-1$ entries of θ and $[s(\cdot), \theta]$ are very similar. In turn, item (iv) states that, except when $\theta_n = \max_{i \neq n} \theta_i$ and player n submits a larger bid, the same outcome is realized under an arbitrary vector of optimal strategies as under truth telling. This exception does not take place for the specific optimal strategies we consider in the sequel.

Proof.

(i) and (ii) follow by a straightforward induction using Lemma 4.3. (iii) is a direct consequence of (i). To prove (iv) note that if $\theta_n = \max_{i \neq n} \theta_i$ and $f([s(\cdot), \theta]) \neq f(\theta)$, then by (ii) $[s(\cdot), \theta]_n > \theta_n$ and hence $f([s(\cdot), \theta]) = n$. Otherwise $\theta_n \neq \max_{j \in \{1,\ldots,n-1\}} \theta_j$ and two cases arise.

Case 1 $\theta_n > \max_{j \in \{1,\ldots,n-1\}} \theta_j$.

By (i) we have $\theta_n > \max_{j \in \{1,\ldots,n-1\}} [s(\theta), \theta]_j$, so by Lemma 4.3$(ii)$ $[s(\cdot), \theta]_n > \max_{j \in \{1,\ldots,n-1\}} [s(\cdot), \theta]_j$. Hence by (i) and (ii) we get $\text{argsmax}\,[s(\cdot), \theta] = n$ and $\text{argsmax}\,\theta = n$, that is $f([s(\cdot), \theta]) = f(\theta)$.

Case 2 $\theta_n < \max_{j \in \{1,\ldots,n-1\}} \theta_j$.

By (i) $\theta_n < \max_{j \in \{1,\ldots,n-1\}} [s(\theta), \theta]_j$, so by Lemma 4.3$(iv)$ $[s(\cdot), \theta]_n \leq \max_{j \in \{1,\ldots,n-1\}} [s(\cdot), \theta]_j$. Hence $\text{argsmax}\,[s(\cdot), \theta] = \text{argsmax}_{j \in \{1,\ldots,n-1\}} [s(\cdot), \theta]_j$. Also $\text{argsmax}\,\theta = \text{argsmax}_{j \in \{1,\ldots,n-1\}} \theta_j$. So by (ii) we get $\text{argsmax}\,[s(\cdot), \theta] = \text{argsmax}\,\theta$, that is $f([s(\cdot), \theta]) = f(\theta)$. $\square$

B Proofs of Results of Section 5

Proof of Theorem 5.1.

Take a strategy $s'_i(\cdot)$ that differs from $s_i(\cdot)$. For some $\theta_1, \ldots, \theta_i$

$$s_i(\theta_1, \ldots, \theta_i) \neq s'_i(\theta_1, \ldots, \theta_i).$$

Let

$$\theta_{i+1} := \ldots = \theta_n = \max(\theta_1, \ldots, \theta_i, s_i(\theta_1, \ldots, \theta_i), s'_i(\theta_1, \ldots, \theta_i)) + 2.$$

and

$$\theta'_{i+1} := \ldots = \theta'_n = \theta_{i+1} - 1.$$

Take now for each $j \in \{1, \ldots, i-1\}$ the truth telling $\pi_j(\cdot)$ as the strategy $s_j(\cdot)$, and for $j \in \{i+1, \ldots, n\}$ a strategy $s_j(\cdot)$ such that

$$s_j(\theta_1, \ldots, \theta_{i-1}, s_i(\theta_1, \ldots, \theta_i), \theta_{i+1}, \ldots, \theta_n) = s_i(\theta_1, \ldots, \theta_i),$$
$$s_j(\theta_1, \ldots, \theta_{i-1}, s_i'(\theta_1, \ldots, \theta_i), \theta_{i+1}', \ldots, \theta_n') = s_i'(\theta_1, \ldots, \theta_i).$$

Then

$$u_i((f, t)([(s_i(\cdot), s_{-i}(\cdot)), \theta], \theta_i) = u_i((f, t)(s_i(\theta_1, \ldots, \theta_i), \theta_{-i}), \theta_i) = r_i(\theta_{-i}),$$
$$u_i((f, t)([(s_i'(\cdot), s_{-i}(\cdot)), \theta], \theta_i) = u_i((f, t)(\theta_1, \ldots, \theta_{i-1}, s_i'(\theta_1, \ldots, \theta_i), \theta_{i+1}', \ldots, \theta_n'), \theta_i)$$
$$= r_i(\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}', \ldots, \theta_n'),$$

where r is the redistribution function of (f, t). By the definition of the BC auction

$$r_i(\theta_{-i}) = \frac{\theta_{i+1}}{n} > \frac{\theta_{i+1} - 1}{n} = \frac{\theta_{i+1}'}{n} = r_i(\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}', \ldots, \theta_n'),$$

which shows that $s_i'(\cdot)$ does not weakly dominate $s_i(\cdot)$. □

Proof of Theorem 5.2.

Consider a type vector $\theta = (\theta_1, \ldots, \theta_n)$ and let $s'(\cdot)$ be an arbitrary vector of optimal strategies. By Corollary A.1(iv) the initial social welfare, i.e., the (initial) utility of the winner, is the same under $s(\cdot)$ and $s'(\cdot)$. Hence we only need to compare the aggregate tax paid by the players.

We now proceed in three steps. First suppose that θ is such that $\theta_n \leq \max_{j \in \{1, \ldots, n-1\}}[s(\cdot), \theta]_j$. By the definition of $s(\cdot)$ in (3) and the definition of the BC auction (see (2)), we know that player n submits the currently second highest bid under $s(\cdot)$ and as a result the taxes are 0. Hence for such type vectors maximal social welfare is generated when players follow $s(\cdot)$.

Thus we may assume $\theta_n > \max_{j \in \{1, \ldots, n-1\}}[s(\cdot), \theta]_j$. Suppose now $\theta_{n-1} \leq \max_{j \in \{1, \ldots, n-2\}}[s(\cdot), \theta]_j$. Then player $n-1$ submits the currently highest bid and since player n has the highest bid, the taxes sum up to 0 by (2).

Hence the remaining case is when $\theta_n > \theta_{n-1} > \max_{j \in \{1, \ldots, n-2\}}[s(\cdot), \theta]_j$. In this case we know that when the players follow $s(\cdot)$, player n is the winner and player $n-1$ submits θ_{n-1}. We claim that the same happens under any other optimal strategy.

Lemma 4.3 and Corollary A.1(iii) imply that when players follow $s'(\cdot)$, player n is the winner and player $n-1$ submits θ_{n-1}. Thus the aggregate tax under $s(\cdot)$ is:

$$\sum_i t_i(\theta) = \frac{2}{n}(\theta_{n-1} - [s(\cdot), \theta]_3^*)$$

and under $s'(\cdot)$:

$$\sum_i t_i(\theta) = \frac{2}{n}(\theta_{n-1} - [s'(\cdot), \theta]_3^*).$$

It suffices now to compare the third highest bid in $s(\cdot)$ and $s'(\cdot)$. For this we need the following more general claim.

Claim 1 *Let $\theta \in \Theta$, $z = ([s(\cdot), \theta]_1, \ldots, [s(\cdot), \theta]_{n-1})$ and $z' = ([s'(\cdot), \theta]_1, \ldots, [s'(\cdot), \theta]_{n-1})$, where $s'(\cdot)$ is a vector of optimal strategies. Then for all $i \in \{1, \ldots, n-1\}$, $z_i \geq z_i'$.*

18

This claim also helps in understanding the intuition behind strategy $s_i(\cdot)$. In particular, under the joint strategy $s(\cdot)$, any bidder, except the last one, whose type is no more than the currently highest bid, is bidding the maximal possible value among the set of his optimal strategies.

Proof. The proof is by induction using Lemma 4.3. If $i = 1$, then by Lemma 4.3, $z_1 = z_1' = \theta_1$. For the induction step, suppose the Lemma holds for all $i < k$, where $k \geq 2$. We argue about z_k and z_k'. By the induction hypothesis $\max_{j \in \{1,\ldots,k-1\}} z_j \geq \max_{j \in \{1,\ldots,k-1\}} z_j'$.

Case 1 $\theta_k > \max_{j \in \{1,\ldots,k-1\}} z_j \geq \max_{j \in \{1,\ldots,k-1\}} z_j'$. Then by Lemma 4.3 $z_k = z_k' = \theta_k$.

Case 2 $\max_{j \in \{1,\ldots,k-1\}} z_j \geq \max_{j \in \{1,\ldots,k-1\}} z_j' \geq \theta_k$. Then Lemma 4.3 implies $z_k = \max_{j \in \{1,\ldots,k-1\}} z_j$. On the other hand:

$$z_k' \leq \max_{j \in \{1,\ldots,k-1\}} z_j' \leq \max_{j \in \{1,\ldots,k-1\}} z_j = z_k.$$

Case 3 $\max_{j \in \{1,\ldots,k-1\}} z_j \geq \theta_k > \max_{j \in \{1,\ldots,k-1\}} z_j'$. Then $z_k = \max_{j \in \{1,\ldots,k-1\}} z_j$ and $z_k' = \theta_k$.
□

The above claim does not hold if we also take the last player into account. The reason is that when the last player's type equals the currently highest bid, then he is indifferent between winning the item or not and the set of his optimal strategies is the whole type space Θ.

To complete the proof now, note that the above claim implies that the second highest bid among the first $n - 1$ bids is at least as big in $s(\cdot)$ as in $s'(\cdot)$. But since we are in the case that $\theta_n > \max_{j \in \{1,\ldots,n-1\}} [s(\cdot), \theta]_j$, this means that $[s(\cdot), \theta]_3^* \geq [s'(\cdot), \theta]_3^*$. □

C   Proofs of Results of Section 6

Proof of Theorem 6.1.   For the last player since the notions of optimal and dominant strategies coincide, there is nothing to prove. Consider now a player i with $2 \leq i \leq n - 2$. Let $\hat{\theta} := \max_{j \in \{1,\ldots,i-1\}} \theta_j$.

Case 1 $\theta_i > \hat{\theta}$.

Then $s_i(\cdot) = \theta_i$. To compute the minimum utility over $\theta_{>i}$, there are two events we need to consider. If $\theta_{>i}$ is such that player i wins, then since all players play according to $s(\cdot)$, all remaining players after i declare θ_i. Hence both the highest and second highest bids other than i equal θ_i and the total utility of i will be $\frac{\theta_i}{n}$. If $\theta_{>i}$ is such that i loses, then the second highest bid other than i will be at least θ_i. Hence:

$$min_{\theta_{>i}} u_i((f, t)([s(\cdot), \theta]), \theta_i) = \frac{\theta_i}{n}$$

To compare this against the minimum utility of any other strategy $s_i'(\cdot)$ of i, we will consider 3 cases for $s_i'(\cdot)$. Let $s_i'(\cdot) = \theta_i'$.

Subcase 1a $\theta_i' > \theta_i$. For vectors $\theta_{>i}$ such that i wins, the remaining bidders will all bid θ_i'. Hence the utility of i will be $\theta_i - \theta_i' + \frac{\theta_i'}{n}$. The minimum over all $\theta_{>i}$ will be at most this quantity, which is no more than $\frac{\theta_i}{n}$.

Subcase 1b $\theta'_i < \theta_i$. In this case if $\theta'_i \leq \hat{\theta}$, i never wins, whereas if $\theta'_i > \hat{\theta}$ there exist vectors $\theta_{>i}$ such that i does not win the item. If i does not win, the second highest bid will be in worst case either $\hat{\theta}$ (if the winner is the last player and the remaining players submit $\hat{\theta}$) or θ'_i in the case that $\theta'_i > \hat{\theta}$. But both $\frac{\theta'_i}{n}$ and $\frac{\hat{\theta}}{n}$ are less than $\frac{\theta_i}{n}$.

Case 2 $\theta_i \leq \hat{\theta}$.

Then according to $s(\cdot)$, $s_i(\cdot) = \hat{\theta}$. Player i does not win the item and in worst case the second highest bid can be as low as $\hat{\theta}$. Thus

$$min_{\theta_{>i}} u_i((f,t)([s(\cdot), \theta]), \theta_i) = \frac{\hat{\theta}}{n}$$

Again let player i deviate to some other strategy s'_i and let $\theta'_i = s'_i(\cdot)$.

Subcase 2a $\theta'_i > \hat{\theta}$. For vectors $\theta_{>i}$ such that i wins, the highest bid other than i will be θ'_i and similarly for the second highest bid. Hence the utility will be $\theta_i - \theta'_i + \frac{\theta'_i}{n}$. But then one can easily check that $\frac{\hat{\theta}}{n} \geq \theta_i - \theta'_i + \frac{\theta'_i}{n}$.

Subcase 2b $\theta'_i \leq \hat{\theta}$. Then i never wins and the minimum possible utility is $\frac{\hat{\theta}}{n}$, which is the same as under $s_i(\cdot)$.

This completes the argument for $2 \leq i \leq n - 2$. Suppose now that $i = n - 1$. We will do a similar but simpler case analysis.

Case 1 $\theta_i > \hat{\theta}$.

Then $s_i(\cdot) = \theta_i$. Note now that $\theta_{>i}$ is simply the bid of player n. To compute the minimum utility that i gets, if i ends up winning the item, it means that the last player has submitted the second highest bid, which is $\hat{\theta}$. Therefore, the utility of i will be $\theta_i - \hat{\theta} + \frac{\hat{\theta}}{n}$. On the other hand if i does not win, it means that the last player submitted something higher than θ_i and hence the utility of i is $\frac{\hat{\theta}}{n}$. Among these two alternatives the minimum is $\frac{\hat{\theta}}{n}$. We compare now $\frac{\hat{\theta}}{n}$ against all other strategies of player i. If $s'_i(\cdot) > \hat{\theta}$ then in the scenario that i loses, the last player submits something higher and therefore the second highest bid not including i's bid is $\hat{\theta}$. Hence the minimum possible utility is $\frac{\hat{\theta}}{n}$. On the other hand, if $s'_i(\cdot) \leq \hat{\theta}$, then i never wins and in worst case the redistribution will be at most $\frac{\hat{\theta}}{n}$.

Case 2 $\theta_i \leq \hat{\theta}$.

According to $s_i(\cdot)$, i submits $\hat{\theta}$ and the minimum utility of i will be $\frac{\hat{\theta}}{n}$. Suppose now i deviates to some strategy that bids $\theta'_i > \hat{\theta}$. Then if player n bids something higher, i loses and the utility of i will be $\frac{\hat{\theta}}{n}$, hence he cannot have a higher minimum utility. If he deviates to a strategy that bids $\theta'_i \leq \hat{\theta}$, then i will never win and his utility will never be higher than $\frac{\hat{\theta}}{n}$.

Finally we also need to handle player 1. For player 1 we know that $s_1(\theta_1) = \theta_1$. For $\theta_{>1}$ such that player 1 wins, all other players submit θ_1 under $s(\cdot)$. Hence the utility would be $\theta_1 - \theta_1 + \frac{\theta_1}{n} = \frac{\theta_1}{n}$. For vectors $\theta_{>1}$ such that 1 loses, the second highest bid will always be at least θ_1 if all play $s(\cdot)$. Hence again the minimum utility would be no less than $\frac{\theta_1}{n}$. Consider now a strategy $s'_1(\cdot)$ such that $s'_1(\cdot) > \theta_1$. For vectors $\theta_{>1}$ such that player 1 wins, all remaining bids equal θ'_1. Hence the minimum utility would be at most $\theta_1 - \theta'_1 + \frac{\theta'_1}{n}$, which is less than $\frac{\theta_1}{n}$. Consider now strategies for which the submitted bid θ'_1 is less than θ_1. Then for type vectors

such that i does not win, the second highest bid will be in worst case equal to θ_1' (if the highest bid is the last one and all remaining ones have type less or equal to θ_1'). Thus the minimum utility in this case would be less than $\frac{\theta_1}{n}$. $\qquad \square$

Proof of Theorem 6.4. For the first player there is no other optimal strategy other than $s_1(\theta_1)$. For the last player, we have seen that $s_n(\cdot)$ is dominant. Hence we need only consider a player i such that $2 \leq i \leq n-1$. Suppose now that player i switches from $s_i(\cdot)$ to some other optimal strategy $s_i'(\cdot)$. We need to compare the payoffs that he receives in $[s(\cdot), \theta]$ and $[(s_i'(\cdot), s_{-i}(\cdot)), \theta]$. If i is the winner in $[s(\cdot), \theta]$ then clearly his type is bigger than $\max_{j \in \{1,\dots,i-1\}}[s(\cdot), \theta]_j$. By Lemma 4.3 this means that there is no other optimal strategy for player i.

Therefore the only interesting case to consider is when i is not the winner in $[s(\cdot), \theta]$. By definition of $s(\cdot)$ this means that $[s(\cdot), \theta]_i = \max_{j \in \{1,\dots,i-1\}}[s(\cdot), \theta]_j$ and by Lemma 4.3 the bid of player i in $s_i'(\cdot)$ is strictly less than $[s(\cdot), \theta]_i$. By the definition of the BC auction, the payoff of player i in $[s(\cdot), \theta]$ is

$$u_i([s(\cdot), \theta], \theta_i) = \frac{([s(\cdot), \theta]_{-i})_2^*}{n}$$

and when i plays $s_i'(\cdot)$, it is:

$$u_i([(s_i'(\cdot), s_{-i}(\cdot)), \theta]), \theta_i) = \frac{([(s_i'(\cdot), s_{-i}(\cdot)), \theta]_{-i})_2^*}{n}$$

Now we claim that the first $n-1$ players apart from i announce the same bid in $[s(\cdot), \theta]$ and in $[(s_i'(\cdot), s_{-i}(\cdot)), \theta]$. To see this notice that the way we have defined $s(\cdot)$ in (3), the strategy of the first $n-1$ players only looks at the highest other bid that has been announced so far (no other information from the past is used). But whether player i played according to $s_i(\cdot)$ or according to $s_i'(\cdot)$, he announces a bid which is at most the highest bid declared so far. Therefore for all the subsequent players, the highest bid preceding them is the same both in $[s(\cdot), \theta]$ and in $[(s_i'(\cdot), s_{-i}(\cdot)), \theta]$. The only player that may differ is player n. If n is not the winner, then he declares the second highest bid preceding him. But this may only be smaller in $[(s_i'(\cdot), s_{-i}(\cdot)), \theta]$ in the case that the second highest bid was achieved by player i, otherwise it is also the same with the bid of player n in $[s(\cdot), \theta]$. In summary, the bids in $[s(\cdot), \theta]_{-i}$ and $[(s_i'(\cdot), s_{-i}(\cdot)), \theta]_{-i}$ are the same except possibly for the last player, who may have a lower bid in $[(s_i'(\cdot), s_{-i}(\cdot)), \theta]_{-i}$. Hence the second highest bid in $[s(\cdot), \theta]_{-i}$ is higher and player i has no incentive to switch to $s_i'(\cdot)$.

Finally, Pareto optimality follows directly from Theorem 5.2. $\qquad \square$

D Sequential Vickrey auctions

We now illustrate some similarities and differences in trying to answer the same questions for different sequential Groves mechanisms. We provide here some answers for the sequential Vickrey auction.

The following result shows that optimal strategies are payoff equivalent, unlike the sequential BC auctions where this does not hold.

Theorem D.1 *Consider a sequential Vickrey auction. For all $\theta \in \Theta$, all optimal strategies $s_i(\cdot)$ and $s_i'(\cdot)$ of player i and all vectors $s_{-i}(\cdot)$ of strategies of players $j \neq i$*

$$u_i((f,t)([s(\cdot),\theta]),\theta_i) = u_i((f,t)([(s_i'(\cdot),s_{-i}(\cdot)),\theta]),\theta_i).$$

However, optimal strategies may differ when the players take into account the utility of other players, in particular, the social welfare. The following natural strategy for player i is an example of an optimal strategy that deviates from truth telling:

$$s_i(\theta_1,\ldots,\theta_i) := \begin{cases} \theta_i & \text{if } \theta_i > \max_{j\in\{1,\ldots,i-1\}} \theta_j, \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

Note that strategy $s_i(\cdot)$ is indeed optimal, since it does not change the decision that would be taken if player i is truthful, and hence Lemma 4.4 can be applied.

In analogy to Theorem 5.2, we can show that when each player i follows strategy $s_i(\cdot)$ maximal social welfare is generated.

Theorem D.2 *In the sequential Vickrey auction for all $\theta \in \Theta$ and 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 (5).

This maximal final social welfare under $s(\cdot)$ equals

$$SW(\theta,s(\cdot)) = \theta_i - \max_{i\in\{1,\ldots,i-1\}} \theta_i,$$

where $i = \text{argsmax } \theta$. This is always greater than or equal to the final social welfare achieved in a Vickrey auction when players bid truthfully, which is $\theta_i - \max_{i\neq j} \theta_i$. Additionally, for some inputs, for instance those of the form θ^*, with the first three entries different, it is strictly greater.

Finally we have also studied safety-level equilibria for the Vickrey auction. As it turns out, all optimal strategies can form safety-level equilibria. In contrast, this is not true for sequential BC auctions as shown in Section 6.

Theorem D.3 *Any vector of optimal strategies in the sequential Vickrey auction is a safety-level equilibrium.*