

Bayesian Sequential Auctions

VASILIS SYRGKANIS, Cornell University
ÉVA TARDOS, Cornell University

In many natural settings agents participate in multiple different auctions that are not simultaneous. In such auctions, future opportunities affect strategic considerations of the players. The goal of this paper is to develop a quantitative understanding of outcomes of such sequential auctions. In earlier work (Paes Leme et al. 2012) we initiated the study of the price of anarchy in sequential auctions. We considered sequential first price auctions in the full information model, where players are aware of all future opportunities, as well as the valuation of all players. In this paper, we study efficiency in sequential auctions in the Bayesian environment, relaxing the informational assumption on the players. We focus on two environments, both studied in the full information model in Paes Leme et al. 2012, matching markets and matroid auctions. In the full information environment, a sequential first price cut auction for matroid settings is efficient. In Bayesian environments this is no longer the case, as we show using a simple example with three players. Our main result is a bound of 3 on the price of anarchy in both matroid auctions and matching markets. To bound the price of anarchy we need to consider possible deviations at an equilibrium. In a sequential Bayesian environment the effect of deviations is more complex than in one-shot games; early bids allow others to infer information about the player's value. We create effective deviations despite the presence of this difficulty by introducing a bluffing technique of independent interest.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics; K.4.4 [Computers and Society]: Electronic Commerce; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: sequential auctions, first price auction, Bayesian games, price of anarchy, efficiency, unit-demand, matroid auctions

1. INTRODUCTION

Auctions typically used in practice are extremely simple, not truthful, and do not even run simultaneously. The Web provides an environment where running auctions becomes easy, an environment where simplicity of design and decreased need of coordination is more important than ever before.

The most well-known auction design is the truthful and efficient VCG auction, which requires large amount of coordination among sellers, having items available simultaneously, and requires the sellers to agree on how to divide the revenue. The classical field of mechanism design has a long and distinguished history. However, the mechanisms proposed are typically too complex, and require careful coordination between

Vasilis Syrgkanis was supported in part by ONR grant N00014-98-1-0589 and NSF grants CCF-0729006. Éva Tardos was supported in part by NSF grants CCF-0910940 and CCF-0729006, ONR grant N00014-08-1-0031, a Yahoo! Research Alliance Grant, and a Google Research Grant.

Author's addresses: V. Syrgkanis, Computer Science Department, Cornell University, vasilis@cs.cornell.edu; É. Tardos, Computer Science Department, Cornell University, eva@cs.cornell.edu

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

EC'12, June 4–8, 2012, Valencia, Spain.

Copyright 2012 ACM 978-1-4503-1415-2/12/06...\$10.00.

the sellers. Mechanism design for either multidimensional valuations or players with correlated types is proving to be rather challenging and leads to quite complicated mechanisms. In contrast, mechanisms typically used in practice tend to be simple, and do not satisfy the high standards of classical mechanism design even in simple environments. Bidders appear to prefer simple mechanisms. Given the prevalent use of simple mechanisms, it is important to design simple auctions with good performance and to understand properties of auction designs used in practice.

Several recent papers have studied properties of simple item-bidding auctions, such as using simultaneous second price auctions for each item [Christodoulou et al. 2008; Bhawalkar and Roughgarden 2011], or simultaneous first price auctions [Hassidim et al. 2011; Immorlica et al. 2005] and show that equilibria of these games have high social welfare. Other recent papers study simple auctions used in practice and show that high social welfare can be achieved even when players' valuations come from arbitrarily correlated distributions [Lucier and Paes Leme 2011; Caragiannis et al. 2012].

In this paper we will focus on sequential auctions. The simplest, most natural and most common way to auction items by different owners is to run individual single item auctions (e.g. sell each item separately on eBay). In [Paes Leme et al. 2012] we initiated the study of the quality of outcomes in sequential auctions in the full information setting. Full information is an extremely strong assumption: buyers have to be aware of all future opportunities, and the valuations of all bidders participating in future auctions. In contrast, in online auctions (see [Parkes 2007] for a survey) it is assumed that players make strategic decisions without having any information about the future. However, typical participants in auctions have some information about future events, and engage in strategic thinking about the upcoming auctions. The Bayesian environment considered in this paper allows us to model agents who are fully strategic, and have partial information about upcoming auctions. We focus on first price auctions as the price of anarchy for second price auctions can be arbitrarily bad in most environments already in the full information setting, including matching markets or even additive valuations.

The goal of this paper is to study sequential first price auctions in the Bayesian setting, relaxing the full information assumption of [Paes Leme et al. 2012]. Many of the results for simultaneous auctions, such as [Christodoulou et al. 2008; Bhawalkar and Roughgarden 2011], extend naturally, and sometimes without degradation of quality, to the Bayesian setting. However, such extension is more difficult in the sequential setting. To bound the price of anarchy we need to consider possible deviations at an equilibrium. The sequential and Bayesian environment presents additional difficulties that we need to overcome. Early bids allow others to infer information about the player's value, causing the player behavior in later auctions to become correlated even when valuations are drawn independently, and correlated bidding makes it harder to prove price of anarchy bounds - indicative of this is that Bayes-Nash price of anarchy proofs for simultaneous item bidding (e.g.[Bhawalkar and Roughgarden 2011; Hassidim et al. 2011]) would fail in the presence of correlated bidding.

Another difficulty that is more technical is that price of anarchy proofs typically consider possible deviations that the players have. However, in a sequential setting it is unclear how to predict player behavior outside the equilibrium path. To deal with this issue in [Paes Leme et al. 2012] we consider deviations where players remain on the equilibrium path until a particular item of interest arrives, the item assigned to them in the efficient allocation. In the Bayesian setting, it is harder to identify what can be such an item of interest, as the efficient solution varies with the valuations of other players. We use a random draw of the valuation to select the item in question,

and introduce a bluffing technique to de-correlate the behavior of players and relate the equilibrium outcome to the efficient solution in expectation.

We model strategic thinking in a sequential Bayesian environment by using the notion of Perfect Bayesian Equilibrium. Perfect Bayesian Equilibrium models players updating their beliefs about the values of other participants throughout the game based on the information inferred from previous auctions. This update of beliefs is greatly affected by the information structure of the game. It depends on what information is available to participants after each round of the auction. Updates are needed even if players only observe their own outcomes. The main technical problem with Bayesian sequential games is that the Bayesian update of beliefs is well defined only when players observe actions that are employed by some players type at equilibrium. Otherwise, the probability of observing such an action is zero and an update is not well defined. A Perfect Bayesian Equilibrium assumes that for such out of equilibrium histories of play, players can form any possible belief that is consistent with the history. A more restrictive notion is that of the Sequential Equilibrium, where each player uses a perturbed strategy, bidding uniformly at random with a small probability ϵ . Such perturbation makes all bids possible on the equilibrium path, and hence an equilibrium has well-defined beliefs for each possible history of play. A sequential equilibrium is the limit of such ϵ -perturbed equilibria as ϵ goes to 0. Our example of a matching auction shows that the outcome is inefficient even at a sequential equilibrium, but our bounds on the price of anarchy make no assumptions on the players' beliefs outside the equilibrium path. Hence, our upper bounds also hold for the concept of Perfect Bayesian Equilibrium.

The updates of beliefs depend on what information is available to the players after each round. Updates are needed even if players only observe their own outcomes, but releasing more information to the players allows for more informative updates. In our result on matroid auctions we assume that the winner (and possibly the winning price) is public information after each round. In particular, players deviations from equilibrium are not observed by other players, as long as they do not effect the outcome. Our result on matching auctions does not require any assumption on the information structure.

To illustrate the issues discussed above we start in Section 2 by giving an example of a two-round sequential first price auction in the Bayesian setting. The example has two items on sale, and three bidders. Bidders 1 and 2 are interested in acquiring either of the two items and have equal value for them, while the 3rd bidder wants only the second item. In the full information setting, the only subgame perfect equilibrium of this game is to allocate the two items to the two bidders of larger value. We consider this game with valuations drawn from $[0, 1]$ uniformly at random. Note that in the 2nd auction the third bidder has an information advantage. In this auction two bidders compete for the second item, the newly arrived bidder, and the bidder who lost the first auction. The advantage comes from the fact that losing the first auction contains information about the valuation of this bidder. This information asymmetry, and the fact that first price auction is not truthful, naturally leads to inefficiency. While this example is stated as a matching market, the same example with 3 bidders can also be thought of as matroid auction, as we will show in Section 4.

1.1. Our Results

In section 3 we consider matching markets, where player i has valuation v_{ij} for item j , and is interested in acquiring only one item. We will assume that there is free disposal, so his value for a set of items J is $\max_{j \in J} v_{ij}$. We assume that the valuations of different players are independent, and the valuation of player i comes from a distribution \mathcal{F}_i . Note that the valuation of different items for the same player can be arbitrarily

correlated, for example, players can have the same value for each item of interest. We prove a bound of 3 on the price of anarchy of this game. This is an improvement even in the full information case, where in [Paes Leme et al. 2012] we showed a bound of 4 for the price of anarchy of mixed Nash equilibria for the full information version. However, it is still a bit larger than the bound of 2 for the case of pure equilibria in the full information case. It is an interesting open problem if this increase is necessary. Our proof for pure equilibria in the full information setting relies on evaluating deviating bids for each player on the item they are assigned in the efficient outcome. This simple deviation strategy doesn't work in the Bayesian setting, as the efficient outcome depends on all the players' valuations. Our technique of selecting a good deviation to consider leads to a small increase in the bound. The result also extends to versions of the problem where items are not auctioned one-by-one, but rather groups of items are auctioned simultaneously in each sequential step.

In section 4 we consider a sequential matroid auction. The matroid structure limits the possible set of people that can get served to be a basis of that matroid. We assume that the players' valuations are drawn from a joint distribution \mathcal{F} , that can be arbitrarily correlated. In each step of the auction, we run a first price auction for selecting an element of a cut that doesn't contain a previous winner. This auction structure corresponds to the standard greedy algorithm for selecting the basis of maximum value. In [Paes Leme et al. 2012] we showed that in the full information setting, this auction implements the VCG outcome (achieves the same allocation and prices). In contrast, as the example mentioned above shows, the Bayesian information structure leads to inefficiency. We show that the resulting inefficiency is not unlimited by giving a bound of 3 on the price of anarchy. An interesting example of the matroid auction is matchable subsets of a bipartite graph (a set of nodes in one side is independent, if there is a way to match them to the other side via a matching in the graph). We can think of this as a special case of the matching auction, when players have a single private value v_i but each can only be assigned to a subset of items S_i . For such a matroid setting a more natural auction is to run the sequential item auction that we studied in the matching markets case, instead of the sequential cut auction. We show that the sequential item auction for this special case of matching markets has price of anarchy at most 3 even when bidders valuations are correlated.

1.2. Related Work

The study of Sequential Auctions in the economics literature was initiated by the works of [Weber 1981] and [Milgrom and Weber 1982] that analyzed first and second price sequential auctions with unit-demand bidders in the incomplete information model. They considered a setting where items are identical and players have the same value for all the items available which is drawn independently from the same distribution. Moreover, they limit their study to symmetric equilibria which, in their setting, are always efficient. Much of the literature of sequential auctions used the Milgrom and Weber model to study sequential auctions [Ashenfelter 1989; McAfee and Vincent 1993]. Little is known for more complex valuation models in the incomplete information case (see e.g. [Krishna 2009] for a detailed exposition on results in sequential first price auctions), mainly due to the fact that for more complex valuation models or for asymmetric or correlated settings computing the equilibrium analytically is hard. However, our price of anarchy techniques allow us to expose meaningful properties for much more complex environments without the need to explicitly calculate the equilibrium strategies.

A few papers consider the complete information case for more complex settings. For example, [Gale and Stegeman 2001; Rodriguez 2009; Bae et al. 2009] study models of multi-unit demands but only for the case of two bidders. Most of this work depends

heavily on having just two bidders since then the subgame perfect equilibrium that survives elimination of weakly dominated strategies is unique. In [Paes Leme et al. 2012] we initiated the study of the price of anarchy in sequential auctions with multi-unit demands and more than two bidders. In this work we manage to give efficiency guarantees for sequential auctions not only among more than two bidders but also in the incomplete information case, even for asymmetric players and in some cases even for correlated bidders.

Recent work from the Algorithmic Game Theory community studied the outcomes of simple mechanisms for multi-item auctions. Christodoulou, Kovacs and Schapira [Christodoulou et al. 2008] and [Bhawalkar and Roughgarden 2011] study the case of running simultaneous second price item auctions for combinatorial auction settings. [Christodoulou et al. 2008] prove that for bidders with submodular valuations and incomplete information, the Bayes-Nash Price of Anarchy is 2. [Bhawalkar and Roughgarden 2011] study the more general case of bidders with subadditive valuations and show that under complete information, the Price of Anarchy of any Pure Nash Equilibrium is 2 and under incomplete information the Price of Anarchy of any Bayes-Nash Equilibrium is at most logarithmic in the number of items. [Hassidim et al. 2011] and [Immorlica et al. 2005] study the case of simultaneous first price auctions and show that the set of pure Nash equilibria of the game correspond to exactly the Walrasian equilibria. [Hassidim et al. 2011] also show that mixed Bayes-Nash equilibria have a price of anarchy of 2 for submodular bidders and logarithmic, in the number of items, for subadditive valuations.

The most recent related work on Matroid Auctions is that of [Bikhchandani et al. 2011] who propose a centralized ascending auction for selling a base of a Matroid that results in the VCG outcome. We study a sequential version of this matroid base auction which potentially renders the auction more distributed. Several other recent works have focused on mechanism design questions in the case when the feasible set of allocations has a matroid structure [Chawla et al. 2010; Hartline and Roughgarden 2009].

2. AN ILLUSTRATIVE EXAMPLE AND PRELIMINARIES

We start our discussion with a simple illustrative example. Assume two items will be available for auction. First item 1 and then item 2. There are three buyers a, b, c . Two of the buyers a and b would like either of the two items, while bidder c is only interested in the second item. Suppose that all players valuations are drawn from the uniform distribution $U(0, 1)$ (and players a and b value either item the same). The following theorem, whose proof we defer to Appendix A describes the Sequential Equilibrium bidding strategy of the first stage auction of this game.

THEOREM 2.1. *There is a sequential equilibrium of this game, where in the first auction, players a and b bid using the bidding function $b(v) = 1 - \frac{\ln(1+v)}{v}$, and the players' beliefs about a player losing to a bid $b(v)$ is that the losing player has valuation uniformly distributed in the range $[0, v]$.*

Observe that players bid more conservatively in the first auction than in the standard first price auction without accounting for their expected utility from the second auction. In Figure 1 we show the equilibrium bidding function of the first auction as compared to the (myopic) equilibrium bidding function of $b(v) = v/2$ not taking into account the second round.

Our main point for presenting this example is not the particular form of the equilibrium, but rather the fact that it leads to an asymmetric auction in the second stage. Equilibria in such asymmetric auctions can be computed analytically (see e.g. [Krishna 2009]) and it is well known that asymmetric first price auctions lead to inefficient out-




Fig. 1. The left figure shows the symmetric equilibrium bidding function of the first stage auction compared to myopic bidding $b(v) = v/2$. The right figure shows the asymmetric bidding equilibrium of the second stage auction when player b wins the first auction and hence players a, c compete in the second auction. The value of player b is thus revealed and the belief on player a is updated to be uniform on $[0, v_b]$.

comes, i.e. there are regions of players values where the player with the smaller valuation wins. In contrast, in a full information game this simple auction has a unique equilibrium, that is socially optimal.

COROLLARY 2.2. *The above equilibrium is not guaranteed to be efficient.*

The loss of efficiency comes from a factor that is necessarily introduced by our incomplete information setting: Asymmetry. Even if players start with symmetric priors, and even if we keep the prices secret, the fact that players participate in different auctions causes asymmetry. Suppose that players started with a symmetric prior distribution F . Then in the second auction the loser of the first auction has a distribution that is F conditional on his value being smaller than the winner of the first auction. It is well-known that in such asymmetric settings the player with the weaker distribution is more aggressive ([Maskin and Riley 2000; Kirkegaard 2009; 2011]), which leads to inefficiency.

Next we formally define an equilibrium of a sequential game. In the example above, it is natural to assume that when player a loses to a bid w , and the bidders employ a monotone bidding strategy $b(x)$ at equilibrium, then the common belief on player a 's valuation is updated to condition a 's distribution on $v_a \leq b^{-1}(w)$, since he must have bid less than w if he lost. This is the natural Bayesian update of the information about the value, given that the players observed a losing the first auction. Note, however that in the equilibrium claimed above the range of bids in the first auction is $[0, 1 - \ln 2]$. To complete the definition of an equilibrium, we also have to define how beliefs are updated if players bid outside the range of the equilibrium, i.e. if they go off the equilibrium path.

2.1. Extensive-form Games of Incomplete Information

Defining an equilibrium in a sequential (or extensive form) game is more complex than in the full information setting. In a full information game the subgame perfect equilibrium is defined naturally via backwards induction. In each stage the player's strategies need to form an equilibrium, given the expected outcome defined in future games. In a Bayesian game we need to not only describe the strategies of the players, but also define how beliefs about player values are updated given the information available at each stage, and these strategies and belief updates need to be consistent.

A **Bayesian extensive form game** is defined in our setting as follows (see [Fudenberg and Tirole 1991] for a more comprehensive treatment). At any point in the game,

player's information correspond to pairs (v_i, h_a) of the player's valuation (vector) v_i and the history h_a of outcomes in auctions before auction a . In the games we consider, we will assume that the same information is available to all players, hence the information depends only on the stage of the auction, which simplifies notation. A strategy of a player is a bid for each information set. Thus a player's strategy consists of bids $b_i(v_i, h_a)$.

We can model a sequential game via a game tree, where the starting nodes correspond to different possible valuation profiles, and later nodes correspond to such valuation and a history of play so far. However, players can only base their updates on information available to them, so a player's information set consists of nodes for each possible valuation of the rest of the players v_{-i} . Player i doesn't have deterministic information as to which of these nodes is implemented, since the same history of play h_a could be implemented with different valuation profiles. However, the history of play signals incomplete information about the probability distribution of valuations. This is captured by the belief system of the sequential equilibrium. Specifically in a sequential equilibrium a player has to have beliefs over the nodes of an information set that are consistent with the strategies. This forward consistency is what creates the complexity of the incomplete information setting that we will briefly describe. Given the beliefs of the players at an information set, his bid has to maximize his expected utility. Moreover, the belief of a player at an information set is the Bayesian update of his initial beliefs given the history of play.

An additional difficulty arises when we consider examples, such as the three bidder example above, when the proposed bidding strategy has limited range. In order to fully describe an equilibrium, we need to also define how beliefs will be updated if players deviate by bidding above the range, or more generally use strategies that are not on an equilibrium path for some instance of the players valuations. For histories that have zero density in equilibrium, Bayesian update is not well defined. However, we need to define how beliefs will be updated to be able to evaluate if players would benefit from deviating to such strategies.

The notion of **Perfect Bayesian Equilibrium** allows for almost any updates in this case, as long as the system of beliefs is consistent with prior and the history. The **Sequential Equilibrium**, is a more refined notion. In this equilibrium the beliefs on the information sets that are never reached by some equilibrium path have to be limits of Bayesian updates for ϵ -perturbed strategies of the rest of the players, i.e. if the players were playing every strategy with at least some ϵ density, then the Bayesian update of the beliefs is well-defined for all information sets. Then for a belief to be consistent with a sequential equilibrium it has to be the limit of the beliefs of some sequence of such ϵ strategies.

Sequential Equilibrium has been mostly used in games with finite action and type spaces. One can think of a discretized version of our auction games, where types and bids are multiples of some ϵ and then take ϵ to zero. This is very natural in auction settings since, both types and valuations are multiples of pennies. However, we could also use a recent formal definition of Sequential Equilibrium for infinite type and action space games proposed by Myerson and Reny [Myerson and Reny 2011]. Our efficiency results are very robust and extend to all these different type of equilibrium definitions.

2.2. Efficiency in Auction Games

The core property satisfied by the above solution concepts that will be needed to prove the efficiency guarantees is that a player's expected utility conditional on his value is maximized at equilibrium. Specifically, consider the normal form representation of the extensive form game defined by our auction game. A strategy of a player i in this normal form representation is a function $b_i(t_i)$ that takes a player's type t as input and

outputs a whole contingency plan on what a player will bid at each auction for each possible history of play. If a strategy profile $b(t) = (b_i(t_i))_{i \in N}$ constitutes an equilibrium then it must satisfy that for any player i and type t_i :

$$u_i(t_i) = \mathbb{E}_{t_{-i}}[u_i(b(t))] \geq \mathbb{E}_{t_{-i}}[u_i(b'_i, b_{-i}(t_{-i}))] \quad (1)$$

In other words we need only that the equilibrium is a Bayes-Nash Equilibrium (BNE) of the normal form representation of our extensive form games. The Perfect Bayesian and Sequential Equilibrium concepts introduce further refinements on how the beliefs of the players should be updated according to the history of play observed and how the players behave when some non-equilibrium history is observed.

The social welfare $SW(b)$ of a bid profile b in our auction games is the sum of the utilities of the players and the auctioneer, which boils down to the total value of the players from the resulting allocation. Given a type profile t , we denote with $\text{OPT}(t)$ the allocation that maximizes social welfare.

We quantify the inefficiency of an equilibrium using the concept of the Bayes-Nash Price of Anarchy which is the fraction of the expected social welfare at the worst equilibrium concept that satisfies the above condition over the expected optimal social welfare.

$$\text{Bayes-PoA} = \sup_{b(\cdot) \text{ is BNE}} \frac{\mathbb{E}_t[\text{OPT}(t)]}{\mathbb{E}_t[SW(b(t))]} \quad (2)$$

3. SEQUENTIAL AUCTIONS FOR MATCHING MARKETS

In this section we consider the first example of a sequential auction. There are a set of m items on sale, each with a different seller. Items can go on sale simultaneously in groups, or individually. We assume that buyers are unit-demand, we use v_{ij} to denote player i 's value for item j and assume the value of player i for a set S of items is his maximum value item in the set: $v_i(S) = \max_{i \in S} v_{ij}$ (i.e. we assume free disposal). We will use $v_i \in \mathbb{R}^{[m]}$ to denote the vector of valuations of buyer i . We study such markets in incomplete information setting. We assume that each player's valuation vector is drawn independently from some prior distribution \mathcal{F}_i that is common knowledge. Note that a single user's valuations can be arbitrarily correlated, we only assume independence between the different buyers.

3.1. Bayes-Nash Price of Anarchy

In this section we prove that there is little inefficiency in the matching market setting we described. The proof is based on a "bluffing" technique. Bluffing here corresponds to a player with value vector v_i following a strategy as if he/she had a different valuation v'_i till a certain point in the auction. Such bluffing is useful, as it allows the player to use the equilibrium properties to predict prices so far, and possibly allows to take items cheaply from a different branch of the computation. Specifically, suppose that a player sets his mind on getting some item j . Then he can try to get that item by pretending to be some other type in all auctions previous to item j . This way he has the option to reach item j along several different equilibrium paths and for some of it the price on item j might be small enough that it would be better for player i to take it. However, equilibrium conditions state that none of these paths are preferable and in particular taking expectation over all such possible paths is not preferable. This bluffing is essential in the proof mainly because the bidding behavior of each player depends on the history of play and subsequently on the whole valuation vector. Thus the maximum other bid that a player faces at each auction is an implicit function of his own valuation. Such an effect didn't arise in simultaneous item auctions with independent bidders.

We first state the theorem in the simplest context of auctioning one item each step, then we show that the theorem extends to more general settings.

THEOREM 3.1. *The Bayes-Nash Price of Anarchy of a Sequential First Price Auction with unit-demand bidders is at most 3, that is, the expected social welfare in the Sequential First Price Auction is at least a 1/3 fraction of the social value of the efficient allocation.*

PROOF. Consider an instantiation of players values v . Let OPT_v be an efficient matching allocation for the instance v and $j_i^*(v)$ be the item that player i is matched to. So the allocation is maximizing $\sum_i v_{ij_i^*(v)}$. For simplicity of presentation, we will assume a pure equilibrium, that is, assume that play is a deterministic function of the type of the player and the history so far. The result naturally extends to mixed equilibria as well.

Also let $j_i(v)$ be the maximum value item that player i gets at the Perfect Bayesian Equilibrium PBE. We will also denote with U_{PBE} , SW_{PBE} and \mathcal{R}_{PBE} the expected player utility, social welfare and revenue respectively at PBE. Similarly U_{OPT} , SW_{OPT} , and \mathcal{R}_{OPT} denotes these quantities for the optimal outcome.

To bound the price of anarchy, we will consider deviations for each player i . If they have a deviation that creates high utility for them, they must have this much utility also in the equilibrium outcome. In the full information setting [Paes Leme et al. 2012] use a deviation where buyer i attempts to get $j_i^*(v)$. This is not an option in the Bayesian setting, since $j_i^*(v)$ depends on the valuation of all the players, and hence it is not known to i . We consider the following randomized version b'_i of the deviation for player i : he draws a random sample of a valuation profile w (including his own type), identifies $j_i^*(v_i, w_{-i})$, and the goal of the deviation will be to get this item cheaply. To do this, the player i will “bluff” and play as in the PBE had he been of type w_i up until item $j_i^*(v_i, w_{-i})$. Then he bids $(v_{ij_i^*(v_i, w_{-i})} - P_{ij_i^*(v_i, w_{-i})}^-(w_i, v_{-i}))/2$ at j , where $P_{ij}^-(w_i, v_{-i})$ is the price that player i has paid so far before item j when deviating. Let $f_j^{-i}(v) = \max_{k \neq i} b_{kj}(v_k, h_j(v))$, where $h_j(v)$ is the history of play at PBE up until item j when players have type v . Denote with $b'_i(w)$ the instantiation of the above mixed deviation for some random sample w .

There are two cases we need to consider. Either player i gets $j_i^*(v_i, w_{-i})$ and therefore has utility $u_i(b'_i(w), b_{-i}(v_{-i}))$ at least $(v_{ij_i^*(v_i, w_{-i})} - P_{ij_i^*(v_i, w_{-i})}^-(w_i, v_{-i}))/2$ or $f_{j_i^*(v_i, w_{-i})}^{-i}(w_i, v_{-i}) \geq (v_{ij_i^*(v_i, w_{-i})} - P_{ij_i^*(v_i, w_{-i})}^-(w_i, v_{-i}))/2$. Observe in the last formula that we used $f_{j_i^*(v_i, w_{-i})}^{-i}(w_i, v_{-i})$ instead of $f_{j_i^*(v_i, w_{-i})}^{-i}(v_i, v_{-i})$. This is because player i pretended to be of type w_i and hence the history that the rest of the players see up until item $j_i^*(v_i, w_{-i})$ is as if player i had type w_i . Moreover, $f_j^{-i}(t)$ depends on t_i only through the history up until item $j_i^*(v_i, w_{-i})$.

From the above reasoning we get:

$$\begin{aligned} u_i(b'_i(w), b_{-i}(v_{-i})) &\geq \frac{v_{ij_i^*(v_i, w_{-i})} - P_{ij_i^*(v_i, w_{-i})}^-(w_i, v_{-i})}{2} - f_{j_i^*(v_i, w_{-i})}^{-i}(w_i, v_{-i}) \\ &\geq \frac{v_{ij_i^*(v_i, w_{-i})} - P_i(w_i, v_{-i})}{2} - f_{j_i^*(v_i, w_{-i})}^{-i}(w_i, v_{-i}) \end{aligned}$$

where $P_i(v)$ is the total price that player i pays at PBE when the valuation profile is v .

Now taking expectation over all possible random samples, we can lower bound player i 's expected utility from deviating to b'_i :

$$\begin{aligned} u_i(b'_i, b_{-i}(v_{-i})) &= \mathbb{E}_w [u_i(b'_i(w), b_{-i}(v_{-i}))] \\ &\geq \mathbb{E}_w \left[\frac{v_{ij_i^*}(v_i, w_{-i}) - P_i(w_i, v_{-i})}{2} - f_{j_i^*}^{-i}(v_i, w_{-i})(w_i, v_{-i}) \right] \end{aligned}$$

By the equilibrium condition we have:

$$\begin{aligned} u_i(v_i) &\geq \mathbb{E}_{v_{-i}} [u_i(b'_i, b_{-i}(v_{-i}))] \\ &\geq \mathbb{E}_{v_{-i}} \mathbb{E}_w \left[\frac{v_{ij_i^*}(v_i, w_{-i}) - P_i(w_i, v_{-i})}{2} - f_{j_i^*}^{-i}(v_i, w_{-i})(w_i, v_{-i}) \right] \end{aligned}$$

We work separately for each part on the right hand side:

$$\mathbb{E}_{v_{-i}} \mathbb{E}_w \left[\frac{v_{ij_i^*}(v_i, w_{-i})}{2} \right] = \mathbb{E}_{v_{-i}} \mathbb{E}_{w_i} \mathbb{E}_{w_{-i}} \left[\frac{v_{ij_i^*}(v_i, w_{-i})}{2} \right] = \mathbb{E}_{w_{-i}} \left[\frac{v_{ij_i^*}(v_i, w_{-i})}{2} \right] = \mathbb{E}_{v_{-i}} \left[\frac{v_{ij_i^*}(v)}{2} \right]$$

where we used that w_{-i} is independent of w_i for the \mathbb{E}_{w_i} to be dropped, and the last equation on the first line is just a change of variables.

$$\begin{aligned} \mathbb{E}_{v_{-i}} \mathbb{E}_w \left[\frac{P_i(w_i, v_{-i})}{2} \right] &= \mathbb{E}_{v_{-i}} \mathbb{E}_{w_i} \mathbb{E}_{w_{-i}} \left[\frac{P_i(w_i, v_{-i})}{2} \right] \\ &= \mathbb{E}_{v_{-i}} \mathbb{E}_{w_i} \left[\frac{P_i(w_i, v_{-i})}{2} \right] = \mathbb{E}_v \left[\frac{P_i(v)}{2} \right] \end{aligned}$$

For the last term we will need a notation for the price of an item at some valuation profile: $p_j(v) = \max_k b_{kj}(v_k, h_j(v))$. Obviously for any instance v and for any player i : $f_j^{-i}(v) = \max_{k \neq i} b_{kj}(v_k, h_j(v)) \leq \max_k b_{kj}(v_k, h_j(v)) = p_j(v)$.

$$\begin{aligned} \mathbb{E}_{v_{-i}} \mathbb{E}_w \left[f_{j_i^*}^{-i}(v_i, w_{-i})(w_i, v_{-i}) \right] &= \mathbb{E}_{v_{-i}} \mathbb{E}_{w_i} \mathbb{E}_{w_{-i}} \left[f_{j_i^*}^{-i}(v_i, w_{-i})(w_i, v_{-i}) \right] \\ &= \mathbb{E}_{t_{-i}} \mathbb{E}_{t_i} \mathbb{E}_{v_{-i}} \left[f_{j_i^*}^{-i}(v_i, v_{-i})(t_i, t_{-i}) \right] \\ &= \mathbb{E}_{v_{-i}} \mathbb{E}_t \left[f_{j_i^*}^{-i}(v)(t) \right] \leq \mathbb{E}_{v_{-i}} \mathbb{E}_t \left[p_{j_i^*}(v)(t) \right] \end{aligned}$$

Where the first few lines are just convenient change of variable names. In the above calculations we heavily used independence so as to be able to consider (w_i, v_{-i}) as a random sample t of a valuation vector, and to be able to change the order of taking expectations at the last step. We point out that the bluffing technique is essential in this latter computation since otherwise we would end up with a term of the form: $\mathbb{E}_{v_{-i}} \mathbb{E}_{w_{-i}} \left[f_{j_i^*}^{-i}(v_i, w_{-i})(v_i, v_{-i}) \right]$. This would lead to a term $\mathbb{E}_{v_{-i}} \mathbb{E}_t \left[p_{j_i^*}(t_i, v_{-i})(t) \right]$ which when taking expectation over v_i and summing among all players would not correspond to the expected revenue at equilibrium. The fact that the players valuation affects both the history of play and the index of the item that he is aiming at creates a problem. The bluffing technique un-correlates this dependence since now the players value doesn't affect the history of play, but rather a random sample of his value does.

Combining the three latter equations we proved that:

$$u_i(v_i) \geq \mathbb{E}_{v_{-i}} \left[\frac{v_{ij_i^*}(v)}{2} \right] - \mathbb{E}_v \left[\frac{P_i(v)}{2} \right] - \mathbb{E}_{v_{-i}} \mathbb{E}_t \left[p_{j_i^*}(v)(t) \right] \quad (3)$$

Taking expectation over v_i for each player and adding for all players i we get:

$$\begin{aligned}
U_{\text{PBE}} &= \sum_i u_i \geq \sum_i \mathbb{E}_v \left[\frac{v_{ij_i^*(v)}}{2} \right] - \mathbb{E}_v \left[\frac{P_i(v)}{2} \right] - \mathbb{E}_v \mathbb{E}_t [p_{j_i^*(v)}(t)] \\
&= \frac{SW_{\text{OPT}}}{2} - \frac{\mathcal{R}_{\text{PBE}}}{2} - \mathbb{E}_v \mathbb{E}_t \left[\sum_{j \in [m]} p_j(t) \right] \\
&= \frac{SW_{\text{OPT}}}{2} - \frac{\mathcal{R}_{\text{PBE}}}{2} - \mathcal{R}_{\text{PBE}}
\end{aligned}$$

where the second line follows as the sum over elements $p_{j_i^*(v)}(t)$ is the sum over all elements $p_j(v)(t)$ as the optimal assignment only effects the order of summation. By rearranging, we get:

$$SW_{\text{PBE}} \geq \frac{SW_{\text{OPT}}}{2} - \frac{\mathcal{R}_{\text{PBE}}}{2} \geq \frac{SW_{\text{OPT}}}{2} - \frac{SW_{\text{PBE}}}{2} \implies SW_{\text{PBE}} \geq \frac{1}{3} SW_{\text{OPT}}$$

□

It is easy to generalize the latter proof to the case of mixed Bayes-Nash Equilibria, games where groups of items are sold simultaneously. Note that the bound of 3 of this result is an improvement over the bound of 4 for mixed Nash Equilibria of [Paes Leme et al. 2012] even in the full information setting.

COROLLARY 3.2. *The mixed Bayes-Nash Price of Anarchy of a Sequential First Price Auction where at each stage a set of items are on sale simultaneously and bidders are unit-demand is at most 3.*

Further, we can generalize the above theorems to approximately unit-demand environments in the following sense: given any instance of values v , the optimal matching allocation OPT_M is at least $1/\gamma$ the optimal allocation OPT . Then the latter two theorems carry over with a blow-up of γ .

4. MATROID AUCTIONS

In this section we consider a setting when N players each would like to get a service. However, due the constraints of the service only a subset of them can be selected. We will assume that the constraints have a matroid structure. Suppose there is a set of N bidders and there is a matroid $\mathcal{M} = (N, \mathcal{I})$, where the independent sets are sets that can be simultaneously served. Suppose each player i has a value v_i for getting the service. We assume that the valuation profile v of the bidders is drawn from a possibly correlated distribution F that is common knowledge. The goal of the mechanism is to find an independent set I of high value $\sum_{i \in I} v_i$ to serve.

The simple auction format will be based on the natural greedy algorithm for finding a large value independent set. A maximum value independent set is naturally a maximal independent set, a basis \mathcal{B} of the matroid. A cut of a matroid is a set $S \subset N$ that intersects all bases. For example, in the matroid of a graph, where an edge-set is independent if it contain no cycles, basis are the spanning trees, and cuts are the set of edges crossing a graph cut.

We consider the following class of simple sequential first-price auctions: At each time step the auctioneer picks a cut of the matroid that doesn't intersect previous winners and runs a first price sealed-bid auction among the bidders in the cut. At the end of each auction j the winner w_j and the price that he paid p_j are announced. Such an auction defines an extensive form game of incomplete information. If in each auctions bidders would announce their true value v_i as bids, the auction corresponds

to the standard greedy algorithm for matroids, and would result in picking the basis of maximum value. For many matroid structures cuts can be small, and hence the auction lends itself to a more distributed implementation, when only elements of a cut have to participate in one auction. The sequential structure of the auction allows participants to infer information about other players, and leads to strategies that possibly will not lead to an efficient outcome.

4.1. Inefficiency of the Incomplete Case

In [Paes Leme et al. 2012] we studied the complete information version of this matroid auction game and showed that it implements the VCG outcome in the sense that all equilibria select an efficient outcome, and make all winners pay their VCG price.

However, when we move to the incomplete information setting the auction no longer implements the VCG outcome, and inefficiency could arise at the sequential equilibria of the game.

As an example, consider the matroid on three players, where any two are independent, but not all three. This is the graphical matroid of a triangle graph. Suppose that all players valuations are drawn from the uniform distribution $U(0, 1)$. If we use the cut $\{a, b\}$ as the first cut, then in the next auction, the loser of the first auction is bidding against the 3rd player c . This is exactly the example of Lemma 2.2 cast as a matroid auction: players a and b participate in the first auction, and then c plays against the loser. As we showed the result is not guaranteed to be efficient.

4.2. Bayes-Nash Price of Anarchy

We first present a lemma from [Paes Leme et al. 2012] that will be very useful in our main theorem. We define the notion of the **participation graph** $\mathcal{P}(B)$ of a base B to be a bipartite graph between the elements in the base and the auctions that took place. An edge exists between an element of the base and an auction if that element participated in the auction.

LEMMA 4.1 ([PAES LEME ET AL. 2012]). *For any base B of the matroid \mathcal{M} , $\mathcal{P}(B)$ contains a perfect matching.*

THEOREM 4.2. *The Bayes-Nash Price of anarchy of the Sequential First Price Cut Auction is at most 3 even if valuations are arbitrarily correlated.*

PROOF. Given an instantiation of the value vector v , let $\text{OPT}(v)$ be the optimal base under value vector v and $\text{PBE}(v)$ the base produced at a perfect Bayesian equilibrium. From Lemma 4.1 we have that for any instantiation of values v there is a perfect matching between the auctions that take place and the players in $\text{OPT}(v)$. For a player $i \in \text{OPT}(v)$ let $A_v(i)$ be the index of the auction that player i is matched with in the latter matching. Also let $x_i(v) = \mathbf{1}_{i \in \text{PBE}(v)}$ and $y_i(v) = \mathbf{1}_{i \in \text{OPT}(v)}$.

Consider an instance of players' values v and a player $i \in \text{OPT}(v) - \text{PBE}(v)$. Suppose that i bids according to the following strategy: he bids $v_i/2$ in all the auctions until he either gets allocated or the auctions finish. Since $i \notin \text{PBE}(v)$ every time that he loses an auction j he doesn't affect the outcome (w_j, p_j) that is announced and therefore the change is transparent to the rest of the players. Hence, they continue using their equilibrium strategies in the subsequent auctions and the prices at subsequent auctions are the same as before the deviation. There are two cases: either player i gets allocated at some auction and therefore he gets utility $u_i(b'_i, b_{-i}(v_{-i})) \geq v_i/2$ or he is never allocated, which means that for all auctions j that he participated in $p_j(v) \geq v_i/2$ (consequently, $p_{A_v(i)}(v) \geq v_i/2$). Thus we get that:

$$u_i(b'_i, b_{-i}(v_{-i})) + p_{A_v(i)}(v)y_i(v) \geq \frac{v_i}{2}y_i(v) \quad (4)$$

If $i \in \text{OPT}(v) \cap \text{PBE}(v)$, i.e. $x_i(v) = 1$, then we cannot consider the latter deviation since i is winning at some auction at equilibrium and by switching he will definitely change either the winner or the price of that auction. Therefore the deviation is not transparent to the rest of the players. However, the following inequality, that is trivially true when $x_i(v) = 1$, will be sufficient to give us the desired bound.

$$\frac{v_i}{2}x_i(v) \geq \frac{v_i}{2}y_i(v) \quad (5)$$

Combining the last two inequalities we get that for all $i \in \text{OPT}(v)$:

$$u_i(b'(v_i, w_{-i}), b_{-i}(v_{-i})) + p_{A_v(i)}(v)y_i(v) + \frac{v_i}{2}x_i(v) \geq \frac{v_i}{2}y_i(v) \quad (6)$$

From the equilibrium condition we have:

$$\mathbb{E}_{v_{-i}|v_i}[u_i(b_i(v_i), b_{-i}(v_{-i}))] \geq \mathbb{E}_{v_{-i}|v_i}[\frac{v_i}{2}y_i(v) - p_{A_v(i)}(v)y_i(v) - \frac{v_i}{2}x_i(v)] \quad (7)$$

Taking expectation over v_i , adding for all i and using the fact that $A_v(i)$ is a matching between players in $\text{OPT}(v)$ and auctions that took place (and therefore winners of those auctions), we get:

$$\begin{aligned} \mathbb{E}_v[\sum_i u_i(b(v))] &\geq \mathbb{E}_v[\frac{1}{2} \sum_{i \in \text{OPT}(v)} v_i - \sum_{i \in \text{OPT}(v)} p_{A_v(i)}(v) - \frac{1}{2} \sum_{i \in \text{PBE}(v)} v_i] \\ &= \frac{1}{2}SW_{\text{OPT}} - \mathbb{E}_v[\sum_{i \in \text{PBE}(v)} p_i(v)] - \frac{1}{2}SW_{\text{PBE}} \end{aligned}$$

Rearranging we get:

$$\mathbb{E}_v[\sum_{i \in \text{PBE}(v)} u_i(b(v)) + p_i(v)] + \frac{1}{2}SW_{\text{PBE}} \geq \frac{1}{2}SW_{\text{OPT}} \implies SW_{\text{PBE}} + \frac{1}{2}SW_{\text{PBE}} \geq \frac{1}{2}SW_{\text{OPT}}$$

□

4.3. Single Value Matching Markets

In this part we consider a special case of our matching market setting where each player has a single private value v_i but participates in a subset of auctions S_i . The set of auctions where each player participates is common knowledge but the value is private. We call such a setting a Single Value Matching Market. For this special case of matching markets we show that the bound of 3 on the price of anarchy extends even when the values of the bidders are arbitrarily correlated and not independent.

In terms of an auction setting the latter is a special case of a matroid setting where the matroid is the matchable set of nodes on the one side of a bipartite graph: Let G be a bipartite graph with two sides A and B . This graph gives rise to a matroid on A , where a subset $S \subset A$ is independent if there is a matching of G that matches all nodes in A . We can think of the nodes of B as service stations and A are the players who want to connect to some station. This example can be thought of both as a matroid auction, but it is also a special case of a matching auction where the graph represents which items can be assigned to each player, and all players value all items that they can get identically. It is interesting to note the difference between the sequential cut auction when applied to this matroid and the sequential item auction. The matching auction, runs auctions selling off each element B individually, and greedily assigns the element auctioned to the winner. (With identical valuations winners will not bid on future items). In contrast the cut auction greedily commits to serve the winner w , but

does not commit to which node in B to assign it to. We can think of the mechanism as maintaining a tentative matching.

Therefore, if we want to claim a bound on the sequential item auction for this setting we cannot simply apply the bound of the sequential cut auction from the previous section. However, using a very similar proof with the same type of deviation (bid $v_i/2$ at all the item-auctions you participate in until you get allocated) leads to a proof for the item auction too. Moreover, such a proof holds for correlated bidders.

THEOREM 4.3. *The Bayes-Nash Price of anarchy of the Sequential First Price Item Auction in a single-value matching market is at most 3 even if the values of the different players are arbitrarily correlated.*

ACKNOWLEDGMENTS

We would like to thank William Sandholm for pointing to us the notes of Myerson and Reny on sequential equilibria for infinite type and action games. We would also like to thank the reviewers for the useful comments.

REFERENCES

- ASHENFELTER, O. 1989. How auctions work for wine and art. *Journal of Economic Perspectives* 3, 3, 23–36.
- BAE, J., BEIGMAN, E., BERRY, R., HONIG, M. L., AND VOHRA, R. 2009. On the efficiency of sequential auctions for spectrum sharing. In *Proceedings of the First ICST international conference on Game Theory for Networks*. GameNets'09. IEEE Press, Piscataway, NJ, USA, 199–205.
- BHAWALKAR, K. AND ROUGHGARDEN, T. 2011. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '11. SIAM, 700–709.
- BIKHCHANDANI, S., DE VRIES, S., SCHUMMER, J., AND VOHRA, R. V. 2011. An ascending vickrey auction for selling bases of a matroid. *Oper. Res.* 59, 2, 400–413.
- CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., KYROPOULOU, M., LUCIER, B., PAES LEME, R., AND TARDOS, É. 2012. On the efficiency of equilibria in generalized second price auctions. *CoRR abs/1201.6429*.
- CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM symposium on Theory of computing*. STOC '10. ACM, New York, NY, USA, 311–320.
- CHRISTODOULOU, G., KOVÁCS, A., AND SCHAPIRA, M. 2008. Bayesian combinatorial auctions. In *Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I*. ICALP '08. Springer-Verlag, Berlin, Heidelberg, 820–832.
- FUDENBERG, D. AND TIROLE, J. 1991. *Game Theory*. MIT Press.
- GALE, I. L. AND STEGEMAN, M. 2001. Sequential auctions of endogenously valued objects. *Games and Economic Behavior* 36, 1, 74–103.
- HARTLINE, J. D. AND ROUGHGARDEN, T. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*. EC '09. ACM, New York, NY, USA, 225–234.
- HASSIDIM, A., KAPLAN, H., MANSOUR, Y., AND NISAN, N. 2011. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM conference on Electronic commerce*. EC '11. ACM, New York, NY, USA, 295–296.
- IMMORLICA, N., KARGER, D., NIKOLOVA, E., AND SAMI, R. 2005. First-price path auctions. In *Proceedings of the 6th ACM conference on Electronic commerce*. EC '05. ACM, New York, NY, USA, 203–212.
- KIRKEGAARD, R. 2009. Asymmetric first price auctions. *Journal of Economic Theory* 144, 4, 1617–1635.
- KIRKEGAARD, R. 2011. Ranking asymmetric auctions using the dispersive order. Working Papers 1101, University of Guelph, Department of Economics.
- KRISHNA, V. 2009. *Auction theory*. Academic press.
- LUCIER, B. AND PAES LEME, R. 2011. Gsp auctions with correlated types. In *Proceedings of the 12th ACM conference on Electronic commerce*. EC '11. ACM, New York, NY, USA, 71–80.
- MASKIN, E. AND RILEY, J. 2000. Asymmetric auctions. *Review of Economic Studies* 67, 3, 413–438.
- MCAFEE, R. AND VINCENT, D. 1993. The declining price anomaly. *Journal of Economic Theory* 60, 1, 191–212.

- MILGROM, P. AND WEBER, R. 1982. A theory of auctions and competitive bidding II. Tech. Rep. 5, Northwestern University.
- MYERSON, R. AND RENY, P. J. 2011. Sequential equilibria of games with infinite sets of types and actions. Draft notes. May.
- PAES LEME, R., SYRGKANIS, V., AND TARDOS, E. 2012. Sequential auctions and externalities. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '12. SIAM, 869–886.
- PARKES, D. 2007. Online Mechanisms. *Algorithmic Game Theory* (eds, N. Nisan et al).
- RODRIGUEZ, G. E. 2009. Sequential auctions with multi-unit demands. *The B.E. Journal of Theoretical Economics* 9, 1, 45.
- WEBER, R. 1981. Multiple-object auctions. *Discussion Paper 496, Kellogg Graduate School of Management, Northwestern University*.

Appendix A: Proof of Theorem 2.1

In this section we present in detail the equilibrium computation of the sequential item auction example of two items 1 and 2 and three bidders a, b, c .

Consider the first auction. Since at that point both players face a symmetric setting, it is natural to look at equilibria where a, b bid symmetrically in the first auction. Thus suppose that at the first auction both a and b bid according to some strategy monotone $b(v)$. Based on such a symmetric play in the first auction, we will work out analytically the equilibrium and utilities in the second auction, and then solve the resulting differential equation that equilibrium in the first auction has to satisfy.

In the second auction, player c is playing with the loser of the first auction. Without loss of generality assume this is player a . From the fact that player a lost the first auction and the bidders use monotone strategies, it becomes common knowledge that a has a value less than $b^{-1}(p_1)$, where p_1 is the winning price of the first cut auction. Thus the auction for the second item is a first price auction with asymmetric bidders: player a is uniformly distributed in $[0, b^{-1}(p_1)]$ and player c is uniformly distributed in $[0, 1]$. The latter is a well known setting and the bidding strategies can be analytically computed (see e.g. [Krishna 2009]):

$$b_a(v_a) = \frac{1}{kv_a} (1 - \sqrt{1 - kv_a^2}) = \frac{v_a}{1 + \sqrt{1 - kv_a^2}}$$

$$b_c(v_c) = \frac{1}{kv_c} (\sqrt{1 + kv_c^2} - 1) = \frac{v_c}{1 + \sqrt{1 + kv_c^2}}$$

where $k = \frac{1}{(b^{-1}(p_1))^2} - 1 \geq 0$. Observe that if the "weak" player a has value v_a then for the strong player to win he must have value $v_c \geq b_c^{-1}(b_a(v_a)) = \frac{v_a}{\sqrt{1 - kv_a^2}} \geq v_a$. Hence, the allocation in this auction will not be efficient.

This bidding strategies allow us to explicitly compute player's a expected utility from the last auction given his value v_a , and given the price the first auction was won at. (Note that in the first auction player a has different utility for each possible price outcome even when he loses.) Let $u_a^2(v_a, p_1)$ denote player a 's expected utility from the last auction given his value v_a , and a price announcement of p_1 in the first auction won by player b .

$$\begin{aligned} u_a^2(v_a, p_1) &= F_c(b_c^{-1}(b_a(v_a)))(v_a - b_a(v_a)) = b_c^{-1}(b_a(v_a))(v_a - b_a(v_a)) \\ &= \frac{v_a}{\sqrt{1 - kv_a^2}}(v_a - b_a(v_a)) = \frac{v_a^2}{1 + \sqrt{1 - kv_a^2}} \end{aligned}$$

Since we assume a symmetric strategy in the first auction, the expected utility of player a in the first auction if he makes a bid as if he has value x is:

$$\begin{aligned} u_a^1(x) &= F_b(x)(v_a - b(x)) + (1 - F_b(x))\mathbb{E}_{v_b|v_b > x}[u_a^2(v_a, b(v_b))] \\ &= x(v_a - b(x)) + \int_x^1 u_a^2(v_a, b(t))dt \end{aligned}$$

The first part on the right hand side is the utility of the player from winning the auction and the second part is the utility of the player when he loses.

For $b(x)$ to be an equilibrium it has to be that the latter utility is maximized at v_a . Thus taking the first order conditions gives us:

$$\left. \frac{\partial u_a^1(x)}{\partial x} \right|_{x=v_a} = 0 \Leftrightarrow v_a - b(v_a) - v_a b'(v_a) - u_a^2(v_a, b(v_a)) = 0$$

We also impose the boundary condition that a player with zero value will bid zero $b(0) = 0$. It may be useful to observe that without the final part of $u_a^2(v_a, b(v_a))$ the last first order differential equation gives as a unique solution the well known first price equilibrium for uniformly distributed bidders of $b(v_a) = v_a/2$. In our case the latter part was introduced due to the externality introduced by the subsequent auction and the fact that the utility in the subsequent auction depends on price at the current one. The solution of the above differential equation gives us the equilibrium with the second auction.

Now, we analytically compute the equilibrium function at the first auction. First we observe that for $p_1 = b(v_a)$ we have: $k = \frac{1}{(b^{-1}(b(v_a)))^2} - 1 = \frac{1}{v_a^2} - 1$. Hence:

$$u_a^2(v_a, b(v_a)) = \frac{v_a^2}{1 + v_a}$$

Thus the differential equation for the first auction becomes:

$$v_a - b(v_a) - v_a b'(v_a) - \frac{v_a^2}{1 + v_a} = 0$$

This has a unique solution of $b(v) = 1 - \frac{\ln(1+v)}{v}$. So far we assumed a and b in the first auction bid in the range of $[0, b^{-1}(1)]$ possibly pretending to have a different value, but never bidding outside this range (above $b^{-1}(1)$), and that beliefs after he first auction about the loser are updated to a uniform distribution on $[0, b^{-1}(p)]$.

To complete the equilibrium computation we need to show that this equilibrium arises as a limit of equilibria where players bid uniformly random $U(0, 1)$ with a small probability ϵ and bid according to the equilibrium $b(x)$ with the remaining. For such strategies the Bayesian update of beliefs is well defined for any possible bid. In the equilibrium, the updated beliefs have to be limits of such Bayesian updates as ϵ goes to 0. For a price announcement in the range of $b(x)$ this updates are the classic updates we defined above. For a price announcement above the range of $b(x)$ there two cases: either the loser bid noisily and lost in which case there is no information revealed or he bid according to the equilibrium and lost in which case again no information is leaked about his value. Thus the updated belief for any ϵ will be $U(0, 1)$ for the loser. Thus the limit will also be $U(0, 1)$. So for price announcements that are not predicted by the equilibrium function the updated beliefs are as if there was a price announcement of $b(1)$. The latter detail completes the Sequential Equilibrium computation.