

Simultaneous Bayesian Auctions and Computational Complexity

YANG CAI, UC Berkeley
CHRISTOS PAPADIMITRIOU, UC Berkeley

Bayesian equilibria of simultaneous auctions for individual items have been explored recently [Christodoulou et al. 2008; Bhawalkar and Roughgarden 2011; Hassidim et al. 2011; Feldman et al. 2013] as an alternative to the well-known complexity issues plaguing combinatorial auctions with incomplete information, and some strong positive results have been shown about their performance. We point out some very serious complexity obstacles to this approach: Computing a Bayesian equilibrium in such auctions is hard for **PP** — a complexity class between the polynomial hierarchy and **PSPACE** — and even finding an approximate such equilibrium is as hard as **NP**, for some small approximation ratio (additive or multiplicative); therefore, the assumption that such equilibria will be arrived at by rational agents is quite problematic. In fact, even recognizing a Bayesian Nash equilibrium is intractable. Furthermore, these results hold even if bidder valuations are quite benign: Only one bidder valuation in our construction is unit demand or monotone submodular, while all others are additive. We also explore the possibility of favorable price of anarchy results for no-regret dynamics of the Bayesian simultaneous auctions game, and identify complexity obstacles there as well.

1. INTRODUCTION

Auction design has proven an extremely fertile arena for productive interplay and dialogue between ideas from economics and computational concepts. Even though novel algorithmic techniques have occasionally helped advance the state of the art in important fronts [Cai et al. 2012, 2013; Papadimitriou and Pierrakos 2011], more often computational complexity considerations have shown that the constraint of truthfulness mixes poorly with algorithmic efficiency [Papadimitriou et al. 2008; Dobzinski and Vondrak 2012; Cai et al. 2013]. Roughly speaking, we now know that the auctions of Vickrey and Myerson are isolated areas of light in a sea of dark, while the new computationally efficient auctions discovered by computer scientists generally lack the compelling simplicity of those archetypes. This revealed complexity of auction design has prompted researchers to design auctions that are *simple* and only *approximately optimal* [Hartline and Roughgarden 2009; Dhangwatnotai et al. 2010; Daskalakis and Pierrakos 2011].

In recent years, one fruitful and much traveled such direction has been to approximate combinatorial auctions through simultaneous independent auctions of single items [Christodoulou et al. 2008; Bhawalkar and Roughgarden 2011; Hassidim et al. 2011; Feldman et al. 2013]. In a *combinatorial auction* bidders are interested in *bundles* of items, and bidder valuations are no longer single numbers but functions from the powerset of all items to the reals (presented either through an oracle, or by some succinct representation). Because of this immediate complexity, combinatorial auctions have played in algorithmic mechanism design a role akin to that enjoyed by the traveling salesman problem in algorithmic combinatorial optimization: a paradigmatic hard nut on which all new ideas must be tried. This new approach seeks to approximate

This work is supported by the National Science Foundation NSF Award CCF-0964033.

Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'14, June 8–12, 2014, Stanford, CA, USA. Copyright © 2014 ACM 978-1-4503-2565-3/14/06 ...\$15.00.
<http://dx.doi.org/10.1145/2600057.2602877>

the objective (revenue or welfare) by inviting the bidders to participate in independent auctions for single items. By bidding to these different auctions for each of the items, the bidders will try to assemble a good bundle. This defines a game — a Bayesian game if, as it is common in auctions, priors of the bidder valuations are known — and it has been shown that, once this simple process has converged to a Bayesian Nash equilibrium, the objective is often approximated satisfactorily. This approach was first used by Christodoulou et al. [Christodoulou et al. 2008], who showed that conducting simultaneous Vickrey auctions for the items results in Bayesian Nash equilibria that approximate the optimum welfare within a factor of two in the special case in which the bidder valuations are *fractionally subadditive* (a special case of subadditive also known as XOS valuations) and under the additional assumption of “no overbidding” (needed to rule out some unnatural equilibria, see below) [Christodoulou et al. 2008]. Soon after that, Bhawalkar and Roughgarden showed that this ratio is optimal for Bayesian Nash equilibria of simultaneous auctions, and becomes worse for more general valuations [Bhawalkar and Roughgarden 2011], while Hassidim et al. [Hassidim et al. 2011] showed a similar result for first-price simultaneous auctions with combinatorial valuations. Finally, in STOC 2013 Feldman et al. [Feldman et al. 2013] showed approximation ratios of four and two for all subadditive valuations, in the case of both Vickrey and first-price auctions, respectively.

In view of these impressive positive results, it is natural to ask: *But is it reasonable to assume that bidders will attain a Bayesian Nash equilibrium*, for example in simultaneous Vickrey auctions for which the bidders have combinatorial valuations? This question is even more natural because Bayesian Nash equilibria are known to be crawling with computational difficulty: Even telling whether a finite two-person game has a *pure* Bayesian Nash equilibrium is NP-complete [Conitzer and Sandholm 2008] (a task that is of course trivial in ordinary games). Finding a mixed Bayesian equilibrium in a Bayesian game is of course PPAD-hard [Daskalakis et al. 2009], but appears to be even harder.

This is an important question for reasons that go beyond a complexity theorist’s reflex. Here we have a situation in which an auction designer instead of a true auction sets up a Bayesian game, and then waits for the players to eventually reach a Bayesian Nash equilibrium — the inevitable solution concept for this kind of game, and for which good things have been proved regarding the designer’s objectives. If finding such an equilibrium were to be shown intractable, this approach would be cast in doubt.

In this paper we show some very strong intractability results for precisely this situation. We prove that finding a Bayesian Nash equilibrium in a simultaneous Vickrey auction game in which the bidders have combinatorial valuations of a very simple form, is hard for the class **PP** (the decision version of $\#\mathbf{P}$, far above **NP**, and in fact hovering near PSPACE and over the polynomial hierarchy by Toda’s Theorem [Toda 1991]). Not just that, but even if we were hit in the head by a Bayesian Nash equilibrium of such a game, we would not be able to tell, because certifying best responses is itself **PP**-hard. What is more, the valuations we need to establish all this are remarkably simple: All bidders have additive valuations, with the only exception of one bidder who has unit demand valuation (arguably just the next step in complexity beyond the additive valuations).

What is even more scary is, these are only the lower bounds on the complexity of the problem that we have been able to establish. We only have *triplly exponential* upper bounds for an approximation of the problem, while we do not know of a proof that exact Bayesian Nash equilibria for this game even *exist* (see the discussion in the end of Section 2).

This strong form of intractability prevails when we require to find an ϵ -approximate Bayesian Nash equilibrium (the players do not change their behavior for less than ϵ

additional utility) for an exponentially small ϵ . But even if we are willing to tolerate a *constant* approximation in our equilibria, we still have a negative result: By a different reduction we establish that it is **NP-hard** to find such an equilibrium — or to recognize one. The bidder valuations needed are different here, but they still have a rather simple structure: one monotone submodular valuation among additive ones. The upper bound on the problem’s complexity is now not very far: **NP^{BPP}**.

Finally, in Section 4 we discuss the possibility of employing more relaxed solution concepts for the simultaneous Bayesian auctions game, such as coarse correlated equilibria, as a possible way around the complexity obstacles proven here, and find reasons for serious complexity concern there as well.

Definitions

Recall that in a *game* we have a finite set of *players*, and for each player i a finite set A_i of *actions or strategies*, as well as a *utility*, a function mapping $A = \prod_i A_i$ to the reals. A *pure Nash equilibrium* is a tuple $a \in A$ such that for all players i changing the i th component of a does not increase i ’s utility. An ϵ -*Nash equilibrium* is a distribution on A_i for each player i such that the expected utility of i in the resulting distribution over A does not improve by more than $\epsilon \geq 0$ by changing i ’s distribution. A *mixed Nash equilibrium* is a 0-Nash equilibrium; all games have one (this was shown by Nash in 1950).

Now in a *Bayesian game* every player has a set T_i of *types*, and each type of each player has a different utility. There is also a known (possibly correlated) prior distribution over $T = \prod_i T_i$. A *Bayesian (Nash) equilibrium* is a distribution on A_i for each *type* of each player i such that no *type* of a player can improve its expected utility (expected over all other types and their mixed strategies) by changing its own distribution. A Bayesian game can be transformed into an ordinary game in which the action space of player i is now all functions $T_i \mapsto A_i$. Thus, a Bayesian equilibrium always exists in finite games.

The definition of a game can be extended to the case in which A_i is infinite, or even a compact set such as $[0, 1]$. All concepts carry over, but *not* Nash’s theorem. The games we shall consider are of this form.

In a *combinatorial auction* we have a set of n *items* and a set of m *bidders*. Each bidder i has a *valuation* v_i mapping $2^{[n]}$ to the reals — without loss of generality to $[0, 1]$ (in our proofs and constructions we shall feel free to use larger numbers as convenient). In all combinatorial auctions constructed in our proofs, the valuations of the bidders will be succinct and easy to describe. In a *Bayesian auction*, each bidder i has a set of *types* T_i , and a distribution on those (we assume that the types of the bidders are distributed independently).

There are certain classes of valuations worth mentioning. An *additive* valuation is of the form $v(S) = \sum_{i \in S} v(i)$. A valuation is *subadditive* if for any subsets $S_1, S_2 \subseteq [n]$, $v(S_1) + v(S_2) \geq v(S_1 \cup S_2)$. Such valuations are considered well behaved; valuations that are not subadditive are valuations with *complements* (consider items like left and right boots, or cart and horse, etc.), and render the auction problem harder. An even more benign class is the *submodular* valuations, for which $S_1 \subseteq S_2$ and $a \notin S_2$ implies $v(S_1 \cup \{a\}) - v(S_1) \geq v(S_2 \cup \{a\}) - v(S_2)$.

Given a Bayesian auction, it has been proposed in several recent works to construct a Bayesian game (*not* a finite one), in which the players are the bidder types, the action sets are $A_i = [0, 1]^n$, and the utilities are calculated as follows: Given an action profile $a \in A = [0, 1]^{mn}$, we determine who among the bidders is the highest bidder for each item, and then the utility of each bidder is the bidder’s valuation for the bundle

consisting of all items for which the bidder bid highest, minus the sum of the *second-highest bids* on these items.

As is common and necessary in the analysis of Bayesian auctions, we restrict the action space by disallowing *overbidding*, that is, by requiring that each type bid for each item no more than its value for the item. This has the effect of removing spurious equilibria in which certain types bid ridiculously high but have no incentive to lower their bid because the second-price rule shields them from its consequences. The above requirement is weaker than the requirement used in [Christodoulou et al. 2008; Bhawalkar and Roughgarden 2011], and incomparable but far simpler and more natural than the in-expectation form of [Feldman et al. 2013]. We believe it is the most natural such restriction — in any event, it makes little difference with our simple valuations.

There is, of course, the complication of ties in the bids. There are two equally unsatisfactory (and equally inconsequential) ways to handle this: Flip a coin among winners, or break ties in favor of lower-indexed bidders.

Either way, we are aware of no proof that a Bayesian equilibrium exists in these games (see Theorem 2 of [Jackson and Swinkels 2005] for some very intricate results showing existence in far simpler contexts than the present one). Our reductions establish that such equilibria, *when they exist* are hard to find. Of course we can find an *discretized Bayesian equilibrium* by embedding in the action space $[0, 1]^n$ a discrete net of M^n points. Now we have a finite game, and Nash’s theorem is restored. We do not know, however, whether this equilibrium is in any sense an ϵ -Bayesian Nash equilibrium of the original game.

2. THE COMPLEXITY OF EXACT EQUILIBRIA

In this section we show that finding an exact Bayesian equilibrium for simultaneous Vickrey auctions is **PP**-hard, even when the bidders’ valuations are additive or unit demand (in fact, additive functions with only one unit demand type). Since we are not sure that a Bayesian equilibrium exists in each game, we use the surrogate of an ϵ -equilibrium for an exponentially small $\epsilon > 0$.

We first describe two easy and useful facts about the Vickrey auction game corresponding to a combinatorial auction. Recall that, in a game, a *dominant strategy* for a player (and type) is a strategy that is best response no matter what everybody else (players and their types) is playing.

LEMMA 1. *For all additive types, being truthful (that is, bidding the true valuation of all items) is a dominant strategy.*

PROOF. It is well known [Vickrey 1961] that in a single-item Vickrey auction, being truthful is a dominant strategy. For an additive type A , her utility in an n -item simultaneous Vickrey auction is always the same as the sum of her utility in these n independent single-item Vickrey auctions, one for each item¹. The lemma follows. \square

LEMMA 2. *Let A and B be two additive types of two different bidders i_1 and i_2 respectively, and assume that, for some item j , $A(\{j\}) = x > B(\{j\}) = y$. If i_1 and i_2 are the only two bidders who are interested in j (that is, if $v_i^t(\{j\}) = 0$ for all $i \notin \{i_1, i_2\}$ and all $t \in T_i$), then at any Bayesian equilibrium, type A will never bid lower than y .*

PROOF. By contradiction. Assume A bid $b_{i_1j} = z < y$ with non-zero probability. Since being truthful is a dominant strategy for any additive type (Lemma 1), then B ’s strategy in any Bayesian equilibrium should be at least as good as bidding truthfully.

¹Note that this does not hold for non-additive types.

That means B should be getting item j always. However, bidding z for item j is no longer a best response, because A can strictly increase her utility by changing her bid for j to x . Hence, in any Bayesian equilibrium A should bid no less than y . \square

We now state the main results of the section:

THEOREM 3. *Finding a Bayesian equilibrium in the simultaneous Vickrey combinatorial auction when there is a unit-demand type and all other types are additive, and with a tie breaking rule that favors the unit demand bidder is **PP-hard**.*

We postpone the proof to Appendix. This result requires the auction to have a particular tie breaking rule, we show next a result in which there is no tie-breaking rule, but the unit demand type is replaced with a slightly more complicated valuation.

THEOREM 4. *Finding a Bayesian equilibrium in the simultaneous Vickrey combinatorial auction with monotone submodular bidders is **PP-hard**.*

PROOF. We start from the standard **PP**-complete problem.

MAJSAT: Given a Boolean formula Φ is the number of satisfying truth assignments greater than the number of non-satisfying truth assignments?

Given an instance Φ of MAJSAT with n variables, we construct a combinatorial auction with $n + 1$ items and $n + 2$ bidders. Next, we describe the bidder valuations.

- For $i = 1, 2, \dots, n$, Bidder i is only interested in the i -th item and has uniform valuation distribution for it over $\{1, 3\}$.
- Define $v(S) = 2|S| - (\frac{|S|}{n+1})^2$. Bidder $n+1$ is only interested in the $n+1$ -st item, and has value $D - (2^n - 1)\delta$, where $D = \frac{1}{2^n} \sum_{S \in 2^{[n]}} (v(S \cup \{n+1\}) - v(S))$ and $\delta = \frac{1}{2^{n+2}(n+1)^2}$.
- Finally, Bidder 0 has three types, each with probability $1/3$. The first two types are additive. In the first type, her value for each of the first n items is $3 - \delta$ and $D - (2^n - 1/2)\delta$ for item $n + 1$. In the second type, she has the same value for item $n + 1$, but has value $1 - \delta$ for items $1, 2, \dots, n$.
- The third type of Bidder 0, denoted v^Φ , encodes the instance Φ of MAJSAT. To define $v^\Phi(S)$ for a set $S \subseteq \{1, \dots, n+1\}$, first define x_S as the truth assignment (x_1, \dots, x_n) such that x_i is true if $i \in S$ and false otherwise. Now, if $n+1 \in S$ and x_S satisfies Φ , then $v^\Phi(S) = v(S) - 2^{n+1} \cdot \delta$, otherwise $v^\Phi(S) = v(S)$.

LEMMA 5. $v^\Phi(\cdot)$ is a monotone submodular function.

PROOF. We first prove it is monotone. Let S and T be two subsets of $[n+1]$, where $T \subset S$. Then

$$\begin{aligned} v^\Phi(S) - v^\Phi(T) &\geq 2(|S| + 1) - \left(\frac{|S| + 1}{n+1}\right)^2 - 2^{n+1} \cdot \delta - (2|T| - \left(\frac{|T|}{n+1}\right)^2) \\ &\geq 2(|S| + 1) - \left(\frac{|S| + 1}{n+1}\right)^2 - 2^{n+1} \cdot \delta - (2|S| - \left(\frac{|S|}{n+1}\right)^2) \\ &\geq 2 - 2^{n+1} \cdot \delta - \frac{(2|S| + 1)}{(n+1)^2} \geq 0. \end{aligned}$$

Next, we prove it is submodular. Take any $i \in [n+1]$, without loss of generality we can assume $i \notin S$. Then

$$v^\Phi(S \cup \{i\}) - v^\Phi(S) \leq 2 + 2^{n+1} \cdot \delta - \frac{(2|S| + 1)}{(n+1)^2},$$

and

$$v^\Phi(T \cup \{i\}) - v^\Phi(T) \geq 2 - 2^{n+1} \cdot \delta - \frac{(2|T| + 1)}{(n + 1)^2}.$$

Since $1/(n + 1)^2 \geq 2^{n+1} \cdot \delta$,

$$v^\Phi(S \cup \{i\}) - v^\Phi(S) \leq v^\Phi(T \cup \{i\}) - v^\Phi(T).$$

□

We now turn to the auction. Since all bidders except 0 are additive, we expect that their strategies are “close” to truthful in any Bayesian equilibrium. The following Proposition proves that this is exactly the case.

PROPOSITION 6. *In any Bayesian equilibrium of the above auction, for all $i \in \{1, \dots, n\}$, bidder i will bid 0 on every item except item i , and will bid between $3 - \delta$ and 3 with probability $1/2$ and bid between $1 - \delta$ and 1 with probability $1/2$ for item i . For bidder $n + 1$, she will always bid between $D - (2^n - 1/2)\delta$ and $D - (2^n - 1)\delta$.*

PROOF. It is clear that everyone will bid 0 on anything they are not interested, since we assume no overbidding. Now look at item i ; only bidder 0 and bidder i have types that are interested in it. From Lemma 2, we know when bidder i values the i -th item 3, she will bid some number in $[3 - \delta, 3]$, and when she values item 1, she will bid some number in $[1 - \delta, 1]$. Similarly, we can argue bidder n always bid between $D - (2^n - 1/2)\delta$ and $D - (2^n - 1)\delta$. □

How about Bidder 0? We are only interested in her submodular type.

LEMMA 7. *Type v^Φ should never buy an item in $\{1, 2, \dots, n\}$ with price $\geq 3 - 2\delta$, or fail to buy an item with price ≤ 1 .*

PROOF. For any bidder i in $\{1, \dots, n\}$, we say her bid is *high* if it is no less than $3 - \delta$ and *low* if it is no greater than 1. In Proposition 6 we show that in any Bayesian equilibrium any bidder $i \in \{1, \dots, n\}$ should bid high on item i if her value is 3 and low if her value is 1. This Lemma claims that any best response bid vector for type v^Φ should buy item i if and only if its price is low.

Assume b_0 is a best response for v^Φ , but bidder 0 wins item i with price higher than $3 - 2\delta$. Now consider changing bidder 0's bid for item i to 2 and call this new bid vector b'_0 . When the bidder i 's bid on item i is low, b_0 and b'_0 give bidder 0 the same utility. But when bidder i 's bid on item i is high, which happens with probability $1/2$, b'_0 gives bidder 0 strictly higher utility. Let S be the set of items bidder 0 wins when bidding b'_0 , then $S \cup \{i\}$ is the set of items she wins when bidding b_0 . The difference in value for the two sets is no greater than $2 + 2^{n+1} \cdot \delta$, but she pays at least $3 - \delta$ more when bidding b_0 . Hence, bidding b'_0 is strictly better. Similarly, we can argue that bidder 0 should never lose an item when the price is low. □

LEMMA 8. *In any Bayesian equilibrium, type v^Φ wins item $n + 1$ iff the number of satisfying assignments for Φ is less than 2^{n-1} .*

PROOF. Let b_{n+1} be bidder $n + 1$'s bid for item $n + 1$, and $\#\Phi$ be the number of satisfying assignments for Φ . Using Lemma 7, we can compute U which is the difference of v^Φ 's utility between winning and losing item n .

$$a = \frac{1}{2^n} \sum_{S \subseteq 2^{[n]}} (v^\Phi(S \cup \{n + 1\}) - v^\Phi(S)) - b_n = D - \#\Phi \cdot 2\delta - b_n.$$

From Proposition 6, we know b_{n+1} is between $D - (2^n - 1/2)\delta$ and $d - (2^n - 1)\delta$. Therefore, $a \leq -\delta/2$ if $\#\Phi$ is at least 2^{n-1} and $a \geq \delta$ if $\#\Phi$ is less than 2^{n-1} . \square

Note that this concludes the proof of the Theorem: One can decide the given instance of MAJSAT by simply looking at the “gross outcome” of any Bayesian equilibrium of the auction (which type gets which items).

An interesting consequence of our proof is the following:

COROLLARY 9. *In the simultaneous auction with combinatorial valuations, finding a best response for a single bidder, or even verifying that a response is best, given the strategies of other players and types, is **PP-hard**.*

PROOF. In our construction, let every type except v^Φ be truthful. Then a best response by type v^Φ gives a Bayesian equilibrium, and whether a response is best depends on the outcome of the MAJSAT instance. \square

This is ominous for the prospects of computationally efficient regret minimization [Blum et al. 2008] in Bayesian simultaneous auctions, see Section 4 for further discussion.

Upper bounds. Is the problem in **PP**? We seriously doubt it. If one insists on equilibria of the continuous game, even existence of equilibria is unclear. It is known that if the action set is compact and the utilities are continuous, then Nash equilibria exist; however, utilities here have discontinuities at bidding ties. There is classical work by Simon and Zame [Simon and Zame 1990] showing that, under conditions, equilibria exist *if* appropriate “sharing rules” (tie-breaking rules at discontinuities) are adopted; however, the proof of [Simon and Zame 1990] is complex and measure-theoretic, and gives few clues as to the nature and complexity of these sharing rules or equilibria. [Jackson and Swinkels 2005] show existence of equilibria in auctions (see e.g. their Theorem 2), albeit in much more restricted contexts than the present. We only know how to find ϵ -equilibria of the continuous game with exponentially small ϵ , albeit through quantifier elimination in the theory of reals — a multiply exponential affair, as it applies a doubly exponential algorithm to an exponentially long input. If we resort to approximating the cube by an exponentially dense net to obtain an approximate discrete equilibrium (and we do not know in what sense it is “approximate”), the complexity of computing a Nash equilibrium becomes **EXPPAD** (the exponential variant of **PPAD**).

3. APPROXIMATING A BAYESIAN EQUILIBRIUM

The complexity result of the previous section holds for exponentially small ϵ . In this section, we show that for some *constant* ϵ , finding an ϵ -Bayesian equilibrium in simultaneous Vickrey combinatorial auctions with a single monotone submodular bidding type and all other bidders additive is **NP-hard**. Our hardness result is obtained by a reduction from SET COVER.

We should be clear about what type of approximation we mean: In an ϵ -Bayesian equilibrium, the expected utility of any bidder is within factor of $(1 - \epsilon)$ of the utility for her best response strategy. We will first prove our hardness result for this approximation. There is another type of approximate Bayesian equilibrium, the additive ϵ -Bayesian equilibrium. For this, we first normalize valuations to $[0, 1]$, and we call a strategy profile an additive ϵ -Bayesian equilibrium if deviating will increase their utility by at most ϵ . Finding such an equilibrium is also **NP-hard** (see the Appendix).

Let \mathcal{C} be an instance of the Set Cover problem. The universe is $U = [m]$ and $B = \{X_1, \dots, X_n\}$ is a family of n sets whose union is U . Let S^* be the optimal solution for \mathcal{C} . S^* , that is, the smallest subset of B , such that the sets in it cover U .

Now from any such instance \mathcal{C} , we shall create an auction. The items of the auction are the elements of B , and we define a monotone submodular function $v^{\mathcal{C}}$ over the items — the *coverage function*. The value of $v^{\mathcal{C}}(S)$ for a set of items S is c times the number of elements in the union of the sets in S , where $c > 1$ is some absolute constant:

$$v^{\mathcal{C}}(S) = c \cdot \left| \bigcup_{i \in S} X_i \right|.$$

Now we are ready to define the bidders in the auction. There are $2n + 1$ bidders, labeled through 0 to $2n$. Bidder 0 has two types, with equal probabilities. The first type values any set of items as 0. The second type's valuation is $v^{\mathcal{C}}$. All bidders in $[2n]$ have only one type. Bidder $2i - 1$ and $2i$ are only interested in item i and both have value 1 for it.

We start with a simple fact about the strategies bidder in $[2n]$ will use in a Bayesian equilibrium.

LEMMA 10. *Let σ and δ be two small positive constants. For a $\frac{\sigma \cdot \delta}{5}$ -Bayesian equilibrium, with probability at most δ , the highest bid between bidder $2i - 1$ and $2i$ for item i is less than $1 - \sigma$.*

PROOF. By contradiction. Suppose that in a $\frac{\sigma \cdot \delta}{5}$ -Bayesian equilibrium the highest bid for item i is less than $1 - \sigma$ with probability larger than δ . Without loss of generality, assume that when bidder 0 has the first type and the highest bid is less than $1 - \sigma$ for item i , bidder $2i$ wins the item with probability less than $1/2$. Now consider switching bidder $2i$'s bid to 1 for item i (because of the no overbidding assumption, her bids are 0 always on other items). From Lemma 1, we know that this will not decrease $2i$'s utility under any bid profiles of the others. Moreover, when bidder 0 has the first type and the highest bid is less than $1 - \sigma$ for item i , bidder $2i$ always win item i now. That means, being truthful increases bidder $2i$'s utility by at least $1/2 \cdot \delta/2 \cdot (1 - (1 - \sigma)) = \frac{\sigma \cdot \delta}{4} > \frac{\sigma \cdot \delta}{5}$. Contradiction. \square

Given the structure of the strategies of bidders other than 0, we are now ready to show the following lemma (proof in the Appendix) analyzing the strategies of bidder 0 in any approximate Bayesian equilibrium.

LEMMA 11. *Given any ϵ -Bayesian equilibrium of the simultaneous auction, there is a polynomial time randomized algorithm which, with probability $1 - O(\exp(-n))$, produces a solution to the instance of SET COVER whose size is at most $\frac{1}{1-\epsilon} ((1 - \epsilon)|S^*| + 2\epsilon cm) + 2\delta n$, where σ, δ are two constants in $[0, 1]$ and $\sigma \cdot \delta = 5\epsilon$.*

From this lemma we can show the following:

THEOREM 12. *Unless $\mathbf{BPP} = \mathbf{NP}$, there is no polynomial time algorithm for finding an ϵ -Bayesian equilibrium in simultaneous Vickrey combinatorial auction with monotone submodular bidders for some absolute constant ϵ .*

PROOF. If we can find a restricted version of SET COVER such that n and m is linear in $|S^*|$ and it is \mathbf{NP} -hard to approximate within some factor $a > 1$, then Lemma 11 implies our theorem. More specifically, when n and m are linear in $|S^*|$, Lemma 11 shows that for any ϵ -Bayesian equilibrium, with probability at least $1 - O(\exp(-n))$, we can convert it into a solution for \mathcal{C} whose size is at most $(1 + f(\epsilon))|S^*|$,² where $f(\cdot)$ is a monotone function with $f(0) = 0$. Now let $\epsilon = f^{-1}(a - 1)$. If there is a polynomial time algorithm for finding an ϵ -Bayesian equilibrium, then $\mathbf{BPP} = \mathbf{NP}$.

²We can get such a function by fixing σ and δ . For example, let $\sigma = \sqrt{\epsilon}$.

In fact such a restricted version exists, and it is called Δ -VERTEX COVER. The problem is to find the minimum vertex cover in a graph with bounded degree Δ . Now n is the number of vertices and m is the number of edges. Since every vertex has at most Δ incident edges, the minimum vertex cover has size at least m/Δ . Clearly, $\Delta|S^*| \geq m$ and $(\Delta + 1)|S^*| \geq |S^*| + m \geq n$. Let A_Δ be a number such that it is **NP**-hard to approximate Δ -VERTEX COVER to within a factor of A_Δ . Berman and Karpinski show that $A_4 \in [1.0128, 1.2857]$ and $A_5 \in [1.0138, 1.625]$ [Berman and Karpinski 1999]. If we let $c = 2$, $\sigma = \sqrt{5(\Delta + 1)\epsilon}$ and $\delta = \sqrt{5/(\Delta + 1)\epsilon}$. We can show that for any ϵ if there exists some Δ such that $(4\Delta - 1)\epsilon + 4\sqrt{5(\Delta + 1)\epsilon} \leq A_\Delta - 1$, then finding an ϵ -Bayesian equilibrium is **NP**-hard. Plugging in the results of [Berman and Karpinski 1999], we prove that finding an ϵ -Bayesian equilibrium is **NP**-hard for some $\epsilon \in [7.44 \times 10^{-7}, 7 \times 10^{-4}]$. \square

To extend this result to additive approximate equilibria one has to overcome some technical difficulties, explained in the Appendix.

Incidentally, the complexity upper bound for the two versions of approximate equilibria is quite close to the lower bound just proved: **NP**^{BPP}. To see this, one can embed an ϵ -grid into the action space, guess an approximate Bayesian equilibrium of the approximate game,³ and then make sure that it is an approximate Bayesian equilibrium by sampling repeatedly the action/type space and relying on the union bound to verify that it is indeed an approximate Bayesian equilibrium.

4. ON BAYESIAN NO-REGRET DYNAMICS

The complexity obstacles to Bayesian equilibria shown in the previous sections cast a shadow on the positive price-of-anarchy results in the literature. It is natural to ask whether price of anarchy guarantees can be shown regarding solution concepts for simultaneous Bayesian auction games that are more relaxed — and less complex — than the Bayesian equilibrium. Hope in this regard comes from the work of Roughgarden [Roughgarden 2009] who showed that certain price of anarchy results (encompassing, in fact, most such results in the literature), under certain conditions, and if established in a particular style called the “smoothness framework,” can be shown to also hold for correlated equilibria and no-regret dynamics. Recently, [Syrgkanis and Tardos 2013] extended this framework to accommodate certain Bayesian settings. In this section we discuss briefly this approach.

One encounters three difficulties when looking for good price of anarchy results through no-regret dynamics and the smoothness framework in Bayesian simultaneous auctions. First, it is not at all clear what no-regret learning means in a Bayesian game: there is no mention of Bayesian games in the *Learning in Games* book [Fudenberg and Levine 1998], and even the concept of correlated equilibria of Bayesian games is controversial [Bergemann and Morris 2013]. Second, once we have decided on a notion of “no-regret dynamics” in the context of Bayesian games, we must show that Roughgarden’s smoothness framework extends to the Bayesian case. Third, to the extent that the methodology does generalize, one has to apply it to prove price-of-anarchy bounds for Bayesian simultaneous auctions.

There are two fundamentally different kinds of learning dynamics one can define for Bayesian games, see [Dekel et al. 2004] for a detailed discussion albeit from a different point of view: (a) dynamics in which types are re-drawn from scratch at every stage of the game (we call these *strongly Bayesian dynamics*), and (b) dynamics in which types are drawn once and for all in the beginning of the process and retained throughout (these are *weakly Bayesian dynamics*). In either case, one can define a *coarse correlated*

³The results of [Lipton et al. 2003] guarantee that an approximate equilibrium with small supports (polynomial in n and m) always exists.

equilibrium, or *no-regret dynamics* to be a polynomially computable sequence of n -tuples of mixed strategies (one for each of the players) with the property that the overall expected utility of each player is no smaller than the expected utility of the best single pure strategy play by the same player (or smaller only by a vanishingly small quantity). Note that the sequence of play by a player in a coarse correlated equilibrium need not be best response.

Weakly Bayesian dynamics seem intuitively to be a significant departure from — some would say a betrayal of — the incomplete information nature of Bayesian games. This intuition is supported by the following observation: It turns out that the problem of weakly Bayesian no-regret learning *can be reduced* to full-information no-regret learning. To see how, let us describe the sequence of play in the reduction: Once types are drawn at the first stage, the players start by announcing through their strategies their type to the other players. This may take a few moves, and cause a (quadratic at worst) loss in utility to each player, which will however be amortized over the remaining (polynomial) play. The game needs to have some minimal complexity for this coding to be achieved, but it is easy to see that it is always possible in auction-like games with finitely many types. Once the types are common knowledge, the players proceed with the assumed no-regret dynamics for the full-information game, completing the reduction.

Given the reduction sketched above, weak Bayesian dynamics is essentially a linear combination of (distribution over) exponentially many full-information game dynamics. Therefore, price of anarchy results obtained for full-information simultaneous auction games through the smoothness framework carry over to the weak Bayesian case, and these include the results of [Bhawalkar and Roughgarden 2011] and [Dütting et al. 2013], but also [Christodoulou et al. 2008] (since it is noted in [Bhawalkar and Roughgarden 2011] that their proof can be rendered within the smoothness framework). We conclude that *weakly no-regret learning in Bayesian simultaneous auctions for subadditive valuations has bounded price of anarchy with respect to welfare*.

Strongly Bayesian dynamics are much more faithful to the Bayesian nature of the auction, but are also far more problematic. We can show, by modifying the proof of Theorem 12, that it is NP-hard to find a strongly Bayesian approximate coarse correlated equilibrium in the simultaneous Bayesian auction game. The proof is given in the Appendix, but the basic intuition is this: Since types are re-drawn from the type distribution at each stage, it is impossible to learn the types. And since our intractability proofs show that it is difficult for a bidder to play “well” even if the behavior of all other types is fixed and known, one cannot learn the behavior of the types either — simply because it is already known. Thus, there is nothing to learn, and each play of the game is essentially a one-shot game — that is to say, hard to play with no substantial regret. The constant inapproximability result of Theorem 12 is needed, because a lesser gap could be absorbed in the allowed deviation from optimality.

5. OPEN PROBLEMS

Our work suggests many open problems:

- Show that Bayesian Nash equilibria in *first-price* simultaneous auctions are hard to find. We are confident this is true, but there are still gaps in our proof.
- Prove that, perhaps under assumptions, the simultaneous auction game has Bayesian Nash equilibria.
- Pin down more precisely the complexity of finding such equilibria and approximating them. Improve the approximation lower bound.
- Understand the computational implications of Simon-Zame theory of endogenous tie-breaking rules.

—What is the precise complexity of finding Bayesian equilibria in finite games?

Acknowledgment:. Many thanks to Costis Daskalakis for suggesting this problem to us, and for his advice about the no-regret dynamics section.

REFERENCES

- BERGEMANN, D. AND MORRIS, S. 2013. Bayes correlated equilibrium and the comparison of information structure. Working papers. Feb.
- BERMAN, P. AND KARPINSKI, M. 1999. On some tighter inapproximability results. In *the 26th International Colloquium in Automata, Languages and Programming (ICALP)*.
- BHAWALKAR, K. AND ROUGHGARDEN, T. 2011. Welfare guarantees for combinatorial auctions with item bidding. In *SODA*. 700–709.
- BLUM, A., HAJIAGHAYI, M., LIGETT, K., AND ROTH, A. 2008. Regret minimization and the price of total anarchy. In *STOC*. 373–382.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2012. Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. In *the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2013. Understanding Incentives: Mechanism Design becomes Algorithm Design. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- CHRISTODOULOU, G., KOVÁCS, A., AND SCHAPIRA, M. 2008. Bayesian combinatorial auctions. In *ICALP (1)*. 820–832.
- CONITZER, V. AND SANDHOLM, T. 2008. New complexity results about nash equilibria. *Games and Economic Behavior* 63, 2, 621–641.
- DASKALAKIS, C., GOLDBERG, P. W., AND PAPADIMITRIOU, C. H. 2009. The Complexity of Computing a Nash Equilibrium. *SIAM Journal on Computing* 39, 1, 195–259.
- DASKALAKIS, C. AND PIERRAKOS, G. 2011. Simple, optimal and efficient auctions. In *WINE*. 109–121.
- DEKEL, E., FUDENBERG, D., AND LEVINE, D. K. 2004. Learning to play bayesian games. *Games and Economic Behavior* 46, 2, 282–303.
- DHANGWATNOTAI, P., ROUGHGARDEN, T., AND YAN, Q. 2010. Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*. 129–138.
- DOBZINSKI, S. AND VONDRAK, J. 2012. The Computational Complexity of Truthfulness in Combinatorial Auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*.
- DÜTTING, P., HENZINGER, M., AND STARNBERGER, M. 2013. Valuation compressions in vcg-based combinatorial auctions. In *WINE*. 146–159.
- FELDMAN, M., FU, H., GRAVIN, N., AND LUCIER, B. 2013. Simultaneous auctions are (almost) efficient. In *STOC*. 201–210.
- FUDENBERG, D. AND LEVINE, D. K. 1998. *The Theory of Learning in Games*. MIT Press Books Series, vol. 1. The MIT Press.
- HARTLINE, J. D. AND ROUGHGARDEN, T. 2009. Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*. 225–234.
- HASSIDIM, A., KAPLAN, H., MANSOUR, Y., AND NISAN, N. 2011. Non-price equilibria in markets of discrete goods. In *ACM Conference on Electronic Commerce*. 295–296.
- JACKSON, M. O. AND SWINKELS, J. M. 2005. Existence of equilibrium in single and double private value auctions. *Econometrica* 73, 1, 93–139.
- LIPTON, R. J., MARKAKIS, E., AND MEHTA, A. 2003. Playing Large Games Using Simple Strategies. In *the 4th ACM Conference on Electronic Commerce (EC)*.
- PAPADIMITRIOU, C. H. AND PIERRAKOS, G. 2011. On Optimal Single-Item Auctions. In *the 43rd ACM Symposium on Theory of Computing (STOC)*.

- PAPADIMITRIOU, C. H., SCHAPIRA, M., AND SINGER, Y. 2008. On the hardness of being truthful. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- ROUGHGARDEN, T. 2009. Intrinsic robustness of the price of anarchy. In *STOC*.
- SIMON, L. K. AND ZAME, W. R. 1990. Discontinuous games and endogenous sharing rules. *Econometrica* 58, 4, 861–72.
- SYRGKANIS, V. AND TARDOS, É. 2013. Composable and efficient mechanisms. In *STOC*. 211–220.
- TODA, S. 1991. Pp is as hard as the polynomial-time hierarchy. *SIAM J. Comput.* 20, 5, 865–877.
- VICKREY, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16, 1, 8–37.

Appendix I: Proof of Theorem 3

PROOF. We start from the standard PP-complete problem, MAJSAT: Given a Boolean formula Φ , is the number of satisfying truth assignments greater than $\frac{1}{2} \cdot 2^n$? Given an instance Φ of MAJSAT with n variables and m clauses, we construct a combinatorial auction with m items and $n + 1$ bidders, as follows:

- The items C_1, \dots, C_m correspond to the clauses of the formula.
- The first n bidders correspond to the variables of the formula. Bidder i has two equiprobable types denoted by the two literals associated with the i th variable: x_i and \bar{x}_i . Each type/literal λ has an additive valuation, where her valuation for clause C_j is 1 if the λ appears in the clause, and zero otherwise.
- The last bidder $n + 1$ also has two equiprobable types. The *active* type has valuation $1 - \delta$ for all items. δ is a positive constant to be specified later. Finally, the *unit demand* type has valuation v for all items and unit demand. Here v is a parameter greater than one, to be specified soon.

LEMMA 13. *At any Bayesian Nash equilibrium, each literal type bids at least $1 - \delta$ for all clauses containing it, and 0 on all other clauses.*

PROOF. Since the literal type has value 0 on all clauses that do not contain it, she must bid 0 on all of them. For any clause that contains the literal type, Lemma 2 says that bidding truthfully strictly dominates any strategy that bids less than $1 - \delta$ on that clause. \square

How about the unit demand type of the $n + 1$ st bidder? We can show that only two kinds of bids make sense for this type:

LEMMA 14. *In any Bayesian Nash equilibrium the unit demand type:*

- *either chooses one clause and bids v for it (or anything between 1 and v — this is called the one high bid;*
- *or, otherwise, he bids something strictly between 0 and 1 on all clauses (many low bids), and in such case his utility is the same as bidding strictly less than $1 - \delta$ on all clauses.*

PROOF. We first argue that bidding high on more than one clause is a strictly dominated strategy. Since if he bids 1 on only one of those clauses, he still wins a clause with probability 1 and his value will be the same, but the price is strictly lower.

Next, we consider the case that he uses many low bids. For each clause C , whenever one of the additive bidders has value 1 on it, the unit-demand type loses. Otherwise, the additive bidder can bid 1 and strictly increase her utility. That means the unit-demand type wins a clause only if all the additive bidders have value 0 on it, which is the same outcome as bidding strictly between 0 and $1 - \delta$. \square

Which of these two strategies should the unit demand type choose? With the one high bid strategy he gets a utility of v with probability $\frac{1}{8}$, and utility between $v - 1$ and $v - (1 - \delta)$ with probability $\frac{7}{8}$, and therefore an expected utility of $v - \frac{7}{8}$. On the other hand, with many low bids he gets utility v with probability $1 - s_\Phi 2^{-n}$, and utility zero with probability $s_\Phi 2^{-n}$, where s_Φ is the number of satisfying truth assignments of Φ . Setting $v = \frac{7}{4} - \frac{1}{2^n}$ and $\delta = \frac{1}{7 \cdot 2^{n-3}}$ concludes the proof.

Appendix II: Proof of Lemma 11

Lemma 11: Given any ϵ -Bayesian equilibrium of the simultaneous auction, there is a polynomial time randomized algorithm which, with probability $1 -$

$O(\exp(-n))$, produces a solution to the instance of SET COVER whose size is at most $\frac{1}{1-\sigma}((1-\epsilon)|S^*| + 2\epsilon cm) + 2\delta n$, where σ, δ are two constants in $[0, 1]$ and $\sigma \cdot \delta = 5\epsilon$.

PROOF. First, consider bidders other than zero. Because no bidder is overbidding, the highest bid for any item is at most 1. If type v^c bids c on any item in S^* and 0 on everywhere else, her utility is at least $cm - |S^*|$. Since type v^c occurs with probability $1/2$, bidder 0's utility is at least $\frac{cm - |S^*|}{2}$. For any ϵ -Bayesian equilibrium, type v^c should expect utility at least $(1-\epsilon)(cm - |S^*|)$, as the first type always has utility 0.

Take the given ϵ -Bayesian equilibrium and fix bidder 0's type to be v^c . Sample k bid profiles from the ϵ -Bayesian equilibrium. We say a profile is good if the number of items, whose highest bid between bidder $2i$ and $2i-1$ is less than $1-\sigma$, is no more than $2\delta n$, and type v^c 's utility is no less than $(1-\epsilon)(cm - |S^*|) - \epsilon cm$. As the highest bid between bidder $2i$ and $2i-1$ is less than $1-\sigma$ with probability at most δ (Lemma 10), by a Chernoff bound we can argue that for every bid profile with probability at most $\exp(-2n\delta^2)$ the number of items whose highest bid between bidder $2i$ and $2i-1$ is more than $2\delta n$. Similarly, for any sampled bid profile, the utility for type v^c is less than $(1-\epsilon)(cm - |S^*|) - \epsilon cm$ with probability at most $\exp(-2\epsilon^2)$. For any constant ϵ and δ , if we take n to be sufficiently large, $\exp(-2n\delta^2) + \exp(-2\epsilon^2) < p$, for some absolute constant $p < 1$. By union bound, a sampled profile is good with probability at least $1-p$. Let $k = O(n)$, then with probability at least $1 - O(\exp(-n))$ there is a good profile among the k samples.

Let S be the set of items won by v^c in the good profile. If the items (sets) in S do not cover U , repeat the following until done. Add an item to S , if it covers any new element. Since getting one more item costs us at most 1, but covering one more element increases the value by c , in the end the total utility increases. As there are at most n items, this will take at most n rounds. Now let S be the set we got after applying this procedure, and $S' \subseteq S$ be the set of items whose prices are least $1-\sigma$. The following must hold.

$$\begin{aligned} cm - (1-\sigma)|S'| &\geq \text{utility of } v^c \geq (1-\epsilon)(cm - |S^*|) - \epsilon cm \\ \implies (1-\epsilon)|S^*| + 2\epsilon cm &\geq (1-\sigma)|S'| \\ \implies \frac{1}{1-\sigma}((1-\epsilon)|S^*| + 2\epsilon cm) &\geq |S'| \\ \implies \frac{1}{1-\sigma}((1-\epsilon)|S^*| + 2\epsilon cm) + 2\delta n &\geq |S| \end{aligned}$$

□

Appendix III: Additive ϵ -Bayesian Nash Equilibrium

Next we consider additive ϵ -Bayesian equilibria, the main reason that we cannot directly apply our previous result is that after normalization bidders in $[2i]$ will have values $O(1/n)$; to get an additive ϵ -Bayesian equilibrium, they can use any strategy. To fix this, we will essentially just merge all odd-numbered bidders into bidder 1 and all even number bidders except bidder 0 into bidder 2. Now both bidder 1 and 2 have only one additive type, which values each item 1. Let \mathcal{C} be an instance of Δ -VERTEX COVER. With the exception of the first type of bidder 0, every type has largest value that is $\Theta(n)$. To find an additive ϵ -Bayesian equilibrium for the normalized version of this auction is the same as finding an additive ϵn -Bayesian equilibrium in the original auction.

The proof will be essentially the same as for ϵ -Bayesian equilibrium, except that we need to establish some structure for the bids of bidder 1 and 2 similar to those in Lemma 10.

LEMMA 15. For an additive $\frac{\sigma \cdot \delta \cdot \lambda n}{5}$ -Bayesian equilibrium, with probability at most δ , there are at least λn items whose highest bid between bidder 1 and 2 is less than $1 - \sigma$.

PROOF. By contradiction. Let $\epsilon = \frac{\sigma \cdot \delta \cdot \lambda}{5}$, E denote the event that at least λn items whose highest bid between bidder 1 and 2 is less than $1 - \sigma$. Imagine there is an additive ϵn -Bayesian equilibrium such that the probability for E to happen is bigger than δ . Without loss of generality, we can assume that when bidder 0 has her first type and E happened, bidder 1 wins less of the low price items (whose highest bid is less than $1 - \sigma$) than bidder 2. Now switch bidder 1's bids to the all 1 bid. From Lemma 1, we know that this will not decrease 1's utility under any bid profiles of the others. Moreover, when bidder 0 has her first type and E happens, bidder 1's expected utility actually increases by at least $1/2 \cdot \delta \lambda n / 2 \cdot \sigma = 5\epsilon n / 4$. Contradiction. \square

With Lemma 15, we can use an almost identical proof to show that for sufficiently small ϵ , an additive ϵ -Bayesian equilibrium for the normalized auction can be efficiently converted to a factor $1 + g(\epsilon)$ approximate solution for \mathcal{C} with probability almost 1, if \mathcal{C} is an instance of Δ -VERTEX COVER (or other restricted version of the set cover problem where both n and m are linear in $|S^*|$). Then similarly as in Theorem 12, we can prove the following theorem.

THEOREM 16. Unless $\mathbf{BPP} = \mathbf{NP}$, there is no polynomial time algorithm for finding an additive ϵ -Bayesian equilibrium in simultaneous Vickrey combinatorial auction with monotone submodular bidders for some absolute constant ϵ .

Appendix III: Strongly Bayesian No-Regret Dynamics

We first define strongly Bayesian no-regret dynamics. Let \mathcal{D} be the distribution of players' types, $p_i(s|\theta)$ be the expected payoff for player i , if the strategy profile is s and the type profile is θ . Now, let $\{s^t|\theta^t\}_{t \in [T]}$ be a sequence of strategy profiles that players will use from day 1 to day T , when the corresponding types are θ . We say $\{s^t|\theta^t\}_{t \in [T]}$ is *small-regret* if for each player i , and for any pure strategy a of i ,

$$\mathbb{E}_{t \sim \mathcal{D}} \left[\sum_{t=1}^T p_i(a, s_{-i}^t | \theta^t) - \sum_{t=1}^T p_i(s^t | \theta^t) \right] = o(T).$$

We call this difference, maximized over all players and all pure strategies, the *regret* of the sequence.

We will show that, for simultaneous Vickrey auctions, finding a polynomial sequence of strategies that has $O(T)$ regret is already \mathbf{NP} -hard. We use the same hard instance as in Section 3. The crux of the proof is to establish the following informal claim:

Claim.: For most $t \in T$, and for most $i \in [n]$, the highest bid between bidders $2i - 1$ and $2i$ is close to 1 with high probability.

Once this is proved, we know that the types in $[2n]$ have small total payoff, and therefore the VERTEX COVER type should use a nearly best response strategy with respect to the others' strategies. We can then argue that, if such a strategy sequence exists that is polynomially computable, we can recover from the play of the VERTEX COVER player an approximate solution of the VERTEX COVER instance. This is done in a manner essentially identical to Lemma 11 and Theorem 12, and we do not repeat it here.

We now turn to the claim. In the sequel, when we say the *high bid* of item i , we mean the higher bid between $2i - 1$ and $2i$. We shall also call the elements of $[T]$ "days."

LEMMA 17. *For any sequence with regret at most $\frac{\sigma^4 T}{2}$, there are at most σT days, such that more than σn items have expected high bids less than $1 - \sigma^2$.*

PROOF. For any $i \in [n]$, let b_i^t be the expected high bid for item i at day t . Let $b_i = \sum_{t=1}^T b_i^t$. Consider switching bidder $2i - 1$'s strategy to bidding 1 in all days, clearly this will only increase the utility of bidder $2i - 1$, we call the increment g_{2i-1} . Similarly, we define g_{2i} . An interesting property is that $g_{2i-1} + g_{2i} = 1 - b_i$, since when the high bid is b , increasing the high bidder's bid to 1 will not change her utility, but increasing the low bidder's bid to 1 will increase her utility by $1 - b$. Because, the sequence has regret $\frac{\sigma^4 T}{2}$, $g_{2i-1}, g_{2i} \leq \frac{\sigma^4 T}{2}$. Hence, $1 - b_i \leq \sigma^4 T$.

If there are more than σT days, at which more than σn items have expected high bids less than $1 - \sigma^2$, then $\sum_{i \in [n]} 1 - b_i > \sigma T \cdot \sigma n \cdot \sigma^2 = \sigma^4 T n$. Contradiction. \square

To complete the proof of the claim, we only need to notice that if the expected high bid is close to 1, then the high bid is close to 1 with high probability.