MATHEMATICS OF OPERATIONS RESEARCH
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 0364-765X | EISSN 1526-5471 | 00 | 0000 | 0001

INFORMS
DOI 10.1287/xxxx.0000.0000

Algorithms as Mechanisms: The Price of Anarchy of Relax and Round*

Paul Dütting
Department of Mathematics, London School of Economics, Houghton Street, London WC2A 2AE, UK
p.d.duetting@lse.ac.uk

Thomas Kesselheim
Institute of Computer Science, University of Bonn, Endenicher Allee 19a, 53115 Bonn, Germany
thomas.kesselheim@uni-bonn.de

Éva Tardos
Department of Computer Science, Cornell University, 316 Gates Hall, Ithaca, NY 14853, USA
eva.tardos@cornell.edu

Many algorithms that are originally designed without explicitly considering incentive properties are later combined with simple pricing rules and used as mechanisms. A key question is therefore to understand which algorithms, or more generally which algorithm design principles, when combined with simple payment rules such as pay-your-bid, yield mechanisms with small price of anarchy.

Our main result concerns mechanisms that are based on the relax-and-round paradigm. It shows that oblivious rounding schemes approximately preserve price of anarchy guarantees provable via smoothness. By the virtue of being smoothness proofs, our price of anarchy bounds extend to Bayes-Nash equilibria and learning outcomes. In fact, they even apply out of equilibrium, requiring only that agents have no regret for deviations to half their value.

We demonstrate the broad applicability of our main result by instantiating it for a wide range of optimization problems ranging from sparse packing integer programs, over single-source unsplittable flow problems and combinatorial auctions with fractionally subadditive valuations, to a maximization variant of the traveling salesman problem.

Key words: Algorithmic mechanism design, non-truthful mechanisms, price of anarchy, smoothness framework, randomized rounding, oblivious rounding
MSC2000 subject classification: Primary: 91A10; secondary: 68W25
OR/MS subject classification: Primary: games/group decisions: noncooperative; secondary: analysis of algorithms: suboptimal algorithms
History: Received January 10, 2019; revised November 6, 2019; accepted February 19, 2020.

1. Introduction.

Mechanism design—or "reverse" game theory—is concerned with protocols, or mechanisms, through which potentially selfish agents interact with one another. The basic goal is to achieve a socially desirable outcome in strategic equilibrium, despite the fact that agents are selfish and may misreport the data that is optimized over.

* An extended abstract appeared in Proceedings of the 16th ACM Conference on Economics and Computation, EC 2015, Portland, OR, USA. Part of the work of P. Dütting was done while visiting Google Research. The work of T. Kesselheim was done while he was at Max Planck Institute for Informatics and Saarland University, supported in part by the DFG through Cluster of Excellence MMCI. É. Tardos was supported in part by NSF grant CCF-1408673, CCF-1563714, and ASOFR grant F5684A1.

A celebrated result of Vickrey, Clarke, and Groves [48, 12, 25] establishes that for the central goal of welfare maximization there is a dominant-strategy incentive-compatible (DSIC) mechanism that achieves this goal. This mechanism, however, requires the underlying optimization problem to be solved exactly, which for many optimization problems is an NP-hard task [41, 44]. Moreover, despite major efforts [31, 6, 17, 16, 2, 1], it is not known how to turn an arbitrary approximation algorithm into a DSIC mechanism with the same approximation guarantee.

An alternative, more recent approach, is to relax the incentive constraints, and instead of trying to design a possibly complicated dominant-strategy incentive-compatible mechanism, use a standard approximation algorithm, combine it with a simple payment rule, and take strategic behavior into account (e.g., [34, 11, 23, 35]). The standard way to evaluate the performance of such non-truthful mechanisms is to bound the price of anarchy [29], i.e., the worst-possible ratio between optimal welfare and welfare at equilibrium.

The central question in this area is: which algorithms or more generally which algorithm design principles when combined with simple payment rules lead to mechanisms with small price of anarchy. The "canonical" result in this context is a result of Lucier and Borodin [34] according to which every equilibrium of pay-your-bid mechanisms based on an α-approximate greedy algorithm is within $O(\alpha)$ of the optimal social welfare.

1.1. Our contribution. In this work we consider a different algorithm design paradigm, and show that it has a very desirable property in regard to price of anarchy guarantees.

Our results concern maximization problems and the algorithmic blueprint of *relaxation and rounding* (see, e.g., [47]). In this approach a problem Π is *relaxed* to a problem Π', with the purpose of rendering exact optimization computationally tractable. Having found the optimal relaxed solution x', another algorithm *rounds* x' to a solution x to the original problem.

Many rounding schemes in text books as well as highly sophisticated ones are *oblivious*. These rounding schemes do not require knowledge of the objective function, but instead ensure that the rounded solution provides at least a $1/\alpha$ fraction of the value of the relaxed solution for all possible objective functions. Up to this point, to the best of our knowledge, this property—though wide-spread—has never proven useful. In this paper, we show that oblivious rounding schemes approximately preserve bounds on the price of anarchy provable via smoothness.

1.1.1. Main result. The first ingredient to our main result is the smoothness framework of Roughgarden and Syrgkanis and Tardos [42, 43, 46], which is the main technique for proving price of anarchy guarantees. Guarantees proven through this technique extend to a broad range of equilibrium concepts and compose across problems.

At the heart of this framework is the notion of a (λ, μ)-*smooth mechanism*, where $\lambda, \mu \geq 0$. The main result is that a mechanism that is (λ, μ)-smooth achieves a price of anarchy of $\beta(\lambda, \mu) = \max(1, \mu)/\lambda$ with respect to a broad range of equilibrium concepts including learning outcomes. Furthermore, the simultaneous and sequential composition of (λ, μ)-smooth mechanisms is again (λ, μ)-smooth. Ideally, $\lambda = 1$ and $\mu \leq 1$ in which case this result tells us that all equilibria of the mechanism are socially optimal; otherwise, if $\lambda < 1$ or $\mu > 1$, then this result tells us which fraction of the optimal social welfare the mechanism is guaranteed to achieve at any equilibrium.

The other crucial ingredient to our main result is the notion of an α-*approximate oblivious rounding scheme*, where $\alpha \geq 1$. Such a rounding scheme is a (typically randomized) mapping from solutions x' to the relaxed problem Π' to solutions x to the original problem Π which guarantees that the expected value $\mathbf{E}[w(x)]$ of the rounded solution x is at least a $1/\alpha$ fraction of the value $w(x')$ of the relaxed solution x' for all possible objective functions w.

Clearly an α-approximate oblivious rounding scheme leads to an approximation ratio of α. We show that it also approximately preserves the price of anarchy of the relaxation. We focus on pay-your-bid mechanisms for concreteness. Our result actually applies to a broad range of mechanisms and can also be extended to include settings where the relaxation is not solved optimally; we discuss these extensions in Section 8.

THEOREM 1 **(Main theorem, informal).** *Consider problem Π and a relaxation Π'. Suppose the pay-your-bid mechanism M for Π is derived from the pay-your-bid mechanism M' for Π'. If M' is (λ, μ)-smooth, then M is $(\lambda/(2\alpha), \mu)$-smooth.*

COROLLARY 1. *The price of anarchy established via smoothness of mechanism M' of β translates into a smooth price of anarchy bound for mechanism M of $2\alpha\beta$ extending to both Bayes-Nash equilibria and learning outcomes.*

Smoothness generally only guarantees the existence of deviation strategies that yield high (parametrized in λ and μ) utility. A key ingredient in our proof is to show that for the type of mechanisms we consider it is approximately without loss of generality to consider deviations to half the value. On the one hand, this means that we can often save the factor two in our main theorem and its corollary by showing smoothness for deviations to half the value directly. On the other hand, it means that our guarantees apply already out-of-equilibrium requiring only that players have no regret for deviations to half their value.

1.1.2. Applications. We demonstrate the broad applicability of our main result by showing how it can be used to obtain price of anarchy guarantees for a wide range of optimization problems. For each of these problems we show the existence of a smooth relaxation and the existence of an oblivious rounding scheme.

Sparse packing integer programs. The first problem we consider are multi-unit auctions with n bidders and m items, where bidders have unconstrained valuations. The underlying optimization problem has a natural LP relaxation, which we show is $(1/2, 2)$-smooth. Using the 8-approximate oblivious rounding scheme of Bansal et al. [4], our framework yields a constant price of anarchy. This is quite remarkable as solving the integral optimization problem optimally leads to a price of anarchy that grows linearly in n and m.

We then consider the generalized assignment problem in which n bidders have unit-demand valuations for a certain amount of one of k services and allocations of services to bidders must respect the limited availability of each service. For this problem we also show $(1/2, 2)$-smoothness, and use the 8-approximate oblivious rounding scheme of [4] to obtain a constant price of anarchy.

Both these results are in fact special cases of a more general result regarding sparse packing integer programs (PIP) that we show. Namely, the pay-your-bid mechanism that solves the canonical relaxation of a PIP with column sparsity d is $(1/2, d+1)$-smooth. Multi-unit auctions and the generalized assignment problem have $d = 1$; combinatorial auctions in which each bidder is interested in at most d items simultaneously have $d \geq 1$. For general PIPs the rounding scheme of [4] is $O(d)$-approximate. We get a price of anarchy of $O(d^2)$.

To the best of our knowledge, these are the first price of anarchy guarantees for sparse PIPs or general multi-unit auctions. A few results give constant price of anarchy bounds in the context of multi-unit auctions [36, 46, 14]. The analyzed mechanisms ask each agent to submit a vector of bids expressing her marginal value for each additional item, and assign items depending on which agent's value increases most by an additional item. This approach inherently requires some form of subadditivity of the valuation functions, and the price-of-anarchy guarantees only apply in this case. In contrast to our results, they do not hold, for example, if a bidder has a high value for getting all items but no value for any strict subset.

Single source unsplittable flow. The second problem that we consider are multi-commodity flow problems with a single source. In these problems we are given a capacitated, directed network and a set of requests consisting of a target and a demand, corresponding to requests of, say different information, held at the source. We assume that each player corresponds to a target, and that she has a private value for routing a certain demand from the shared source to her position. The goal is to maximize the total value of the demand routed, subject to feasibility.

For this problem we show that the natural LP relaxation is $(1/2, 1)$-smooth. A $(1+\epsilon)$-approximate oblivious rounding scheme for high enough capacities is obtained through an adaptation of the "original" randomized rounding algorithm of Raghavan and Thompson [39, 40]. This yields a price of anarchy of $2(1+\epsilon)$.

An interesting feature of this result is that the LP can be solved greedily through a variant of Ford-Fulkerson which allows us to exploit the known connection to smoothness [34, 46]. Crucially, the reference to these results has to be on the fractional level, as a greedy procedure on the integral level achieves a significantly worse approximation guarantee.

It is again crucial to show smoothness of the relaxed problem, as solving the integral problem optimally leads to an unbounded price of anarchy.

Combinatorial auctions. We also consider the "canonical" mechanism design problem of combinatorial auctions. Our first result concerns fractionally subadditive, or XOS, valuations [33]. We show that the pay-your-bid mechanism for the canonical LP relaxation is $(1/2, 2)$-smooth. Using Feige's ingenious $e/(e-1)$-approximate oblivious rounding scheme [21], our main result implies an upper bound on the price of anarchy of $4e/(e-1) \approx 6.328$.

We then show how to extend this result to the recently proposed hierarchy of $\mathcal{MPH}$-k valuations [22]. Levels of the hierarchy correspond to the degree of complementarity in a given function. The lowest level $k = 1$ coincides with the class of XOS/fractionally subadditive valuations; the highest level $k = m$ can be shown to comprise all monotone valuation functions. We show that for $\mathcal{MPH}$-k valuations the LP relaxation is $(1/2, k+1)$ smooth. Together with the $O(k)$-approximate oblivious rounding scheme of [22] we obtain a price of anarchy of $O(k^2)$.

These results nicely complement recent work on the price of anarchy of simultaneous first- and second-price auctions [11, 5, 23, 18, 3], in which each item is sold in a separate single-item auction. In these mechanisms players cannot bid their valuation functions, which makes the overall mechanism *indirect*. Our mechanisms, in contrast, are *direct*, meaning that players can report entire valuation functions. This way, we address an open question due to Babaioff et al. [3] about the price of anarchy of direct mechanisms based on approximation algorithms in this setting.

While the price of anarchy guarantees that we obtain are slightly worse, figuring out how to bid in an indirect mechanism such as a simultaneous first- or second-price auction is a non-trivial task. Indeed a sequence of recent works has established that finding exact or approximate equilibria of these mechanisms is hard [7, 15, 13]. In our mechanisms, in contrast, players have simple fall-back strategies and the guarantees that we show apply whenever players have no regret with respect to these simple strategies.

Maximum traveling salesman. Our final application is the maximization variant of the classic traveling salesman problem. We think of the problem as a game where each edge has a value for being included, and the goal of the mechanism is to select a tour of maximum total value. The classic algorithm for this problem is a 2-approximation due to Fisher et al. [24]. It proceeds by computing a cycle cover, dropping an edge from each cycle, and connecting the resulting paths in an arbitrary manner to obtain a solution. We prove this can be thought off as a 2-approximate oblivious rounding scheme and show, through a novel combinatorial argument, that the relaxation is $(1/2, 3)$-smooth. We thus obtain a price of anarchy of 12.

The best approximation guarantee for max-TSP is a $3/2$-approximation due to Kaplan et al. [27]. The same approximation ratio is achieved by a (much simpler) algorithm of Paluch et al. [37]. We show that this algorithm—just as the basic algorithm—can be interpreted as a relax-and-round algorithm. Generalizing the arguments for the basic algorithm to the (different) relaxation used in this interpretation, we show that this algorithm achieves a price of anarchy that is by a factor $3/4$ better than the price of anarchy of the basic algorithm.

This application and these examples are interesting as they show how seemingly combinatorial algorithms can be re-stated within our framework. They also represent the first non-trivial price of anarchy bounds for this problem.

1.2. Further related work. Our work is closely related to the literature on so-called "back-box reductions", which has led to some of the most impressive results in algorithmic mechanism design (such as [31, 6, 16, 17, 2, 1]). This approach takes an algorithm, and aims to implement the algorithm's outcome via a game. To this end it typically modifies the algorithm and adds a sophisticated payment scheme. Our approach is different in that we consider an algorithm without any modification, introduce a simple payment rule, such as the "pay your bid" rule, and understand the expected outcomes of the resulting game.

Lavi and Swamy [31] use *randomized metarounding* [9] to turn LP-based approximation algorithms for packing domains into truthful-in-expectation mechanisms. Our result is similar in spirit as it demonstrates the implications of *obliviousness* for non-truthful mechanism design. The property that we need, however, is less stringent and shared by most rounding algorithms. Another important difference is that our approach is not limited to packing domains.

Briest et al. [6] show how pseudo-polynomial approximation algorithms for single-parameter problems can be turned into a truthful fully polynomial-time approximation schemes (FPTAS). Dughmi et al. [16] prove that every welfare-maximization problem that admits a FPTAS and can be encoded as a packing problem also admits a truthful-in-expectation randomized mechanism that is an FPTAS. Unlike our approach these approaches are limited to single-parameter problems, or to multi-parameter problems with packing structure.

Dughmi et al. [17] present a general framework that also looks at the fractional relaxation of the problem. They show that if the rounding procedure has a certain property, which they refer to as *convex rounding*, then the resulting algorithm is truthful. They instantiate this framework to design a truthful-in-expectation mechanism for combinatorial auctions with matroid-rank-sum valuations (which are strictly less general than submodular). The main difference to our work is that standard rounding procedures are often oblivious but typically not convex.

Babaioff et al. [2, 1] show how to transform a (cycle-)monotone algorithm into a truthful-in-expectation mechanism using a *single call* to the algorithm. The resulting mechanism coincides with the algorithm with high probability. This work differs from ours in that it only applies to monotone or cycle-monotone algorithms.

By insisting on truthfulness, or truthfulness-in-expectation, as a solution concept, all these approaches face certain natural barriers regarding the achievable approximation guarantees (see, e.g., [38, 10]). In addition, they typically do not lead to simple, practical mechanisms. For example, despite running times technically being polynomial, these mechanisms require far more computational effort than standard approximation algorithms for the underlying optimization problem. In some cases, for example when using randomized metarounding [31], the reduction yields mechanisms in which the approximation guarantee is tight on every single instance (not only in the worst case). That is, even when the optimization problem is trivial, the mechanism sacrifices the solution quality for incentives.

2. Preliminaries We begin by formally introducing the key concepts from algorithm design and mechanism design that we need in this work. There will be several sources of randomness. Whenever we write expectations without subscripts, the expectation is to be taken with respect to the randomness in the algorithm or mechanism.

Algorithm design basics. We consider maximization problems Π in which the goal is to determine a feasible outcome $x \in \Omega$ that maximizes total weight given by $w(x)$ for non-negative a weight function $w\colon \Omega \to \mathbb{R}_{\geq 0}$. A potentially randomized algorithm A receives the functions w as input and computes an output $A(w) \in \Omega$. The algorithm is an α-*approximation algorithm*, for $\alpha \geq 1$, if for all weights w, $\mathbf{E}\left[w(A(w))\right] \geq \frac{1}{\alpha} \cdot \max_{x \in \Omega} w(x)$.

We are interested in *relax-and-round algorithms*. These algorithms first relax the problem Π to Π' by extending the space of feasible outcomes to $\Omega' \supseteq \Omega$ and generalizing weight functions w to all $x \in \Omega'$. They compute an optimal solution $x' \in \Omega'$ to the relaxed problem. In the rounding step, a solution $x \in \Omega$ of the original problem is derived from $x' \in \Omega'$, typically in a randomized way.

A rounding algorithm is an α-*approximate oblivious rounding scheme* if, given some relaxed solution x', it computes a solution x such that for all w, $\mathbf{E}\left[w(x)\right] \geq \frac{1}{\alpha} w(x')$.

Clearly, a relax-and-round algorithm that first optimally solves the problem Π' and then applies an α-approximate oblivious rounding scheme is an α-approximation algorithm.

EXAMPLE 1. Consider the knapsack problem with a unit capacity knapsack, in which each item $i \in [n]$ has a size $s_i \geq 0$ and a value $v_i \geq 0$ for being included in the knapsack, and the sizes satisfy $s_i \leq 1/2$ for all i. Feasible solutions correspond to $x \in \{0,1\}^n$ such that $\sum_{i \in [n]} s_i x_i \leq 1$. The goal is to find a feasible solution x that maximizes $v(x) = \sum_{i \in [n]} v_i x_i$.

We relax the integrality constraint by allowing $x'_i \in [0,1]$, so that feasible solutions to the relaxation correspond to $x' \in [0,1]^n$ such that $\sum_{i \in [n]} s_i x'_i \leq 1$. We now compute an optimal solution x' to the relaxed problem for valuations v and turn it into a solution x of the original problem as follows: For each agent i, set $x''_i = 1$ with probability $\frac{x'_i}{2}$ and $x''_i = 0$ otherwise. If the resulting integral solution is feasible, i.e., $\sum_{i \in [n]} s_i x''_i \leq 1$, set $x = x''$. Otherwise, set $x = 0$.

Note that this rounding scheme ensures that $\sum_{i \in [n]} s_i x''_i \leq 1$ with probability at least $\frac{1}{2}$ by Markov's inequality. So $x_i = 1$ with probability at least $x'_i/4$. In particular, for any vector of valuations w (and not just v), we have $\mathbf{E}\left[w(x)\right] \geq \frac{1}{4} \cdot w(x')$. Therefore, the rounding scheme is a 4-approximate oblivious rounding scheme.

In the above example, setting $w = v$ allows us to conclude that the respective relax-and-round algorithm is a 4-approximation algorithm. We will see that the fact that the key inequality $\mathbf{E}\left[w(x)\right] \geq \frac{1}{4} \cdot w(x')$ actually holds for any w will be useful when reasoning about a strategic version of this problem and mechanisms for solving it.

Mechanism design basics. Our results apply to general multi-parameter mechanism design problems Π in which agents $N = \{1, \ldots, n\}$ interact to select an element from a set Ω of outcomes. Each agent has a valuation function $v_i\colon \Omega \to \mathbb{R}_{\geq 0}$. We use v for the valuation profile that specifies a valuation for each agent, and v_{-i} to denote the valuations of the agents other than i. The quality of an outcome $x \in \Omega$ is measured in terms of its social welfare $\sum_{i \in N} v_i(x)$.

We consider *direct mechanisms* M that ask the agents to report their valuations. We refer to the reported valuations as bids and denote them by b. The mechanism uses *outcome rule* f to compute an outcome $f(b) \in \Omega$ and *payment rule* p to compute payments $p(b) \in \mathbb{R}_{\geq 0}$. Both the computation of the outcome and the payments can be randomized. So in general, $f(b)$ and $p(b)$ are random variables. We are specifically interested in *pay-your-bid mechanisms*, in which agents are asked to pay what they have bid on the outcome they get. In other words, in a pay-your-bid mechanism $M = (f,p)$, $p_i(b) = b_i(f(b))$. We assume that the agents have quasi-linear utilities and that they are risk neutral. That is, we assume that agent i's expected utility given value v_i and bids b in mechanism $M = (f,p)$ is $u_i(b, v_i) = \mathbf{E}\left[v_i(f(b)) - p_i(b)\right]$.

For the game-theoretic analysis we distinguish two settings. In the *complete information* setting valuations v are fixed, and agents know each others' valuations. We write $\mathcal{B}_i$ for a distribution over bids b_i by agent i, and $\mathcal{B} = \prod_{i \in N} \mathcal{B}_i$ for the corresponding product distribution. We let $b_{-i} =$

$(b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$ and $\mathcal{B}_{-i} = \prod_{i' \neq i} \mathcal{B}_{i'}$. A distribution $\mathcal{B}$ over bids b is a *mixed Nash equilibrium* if for each agent i and alternative bid b'_i, $\mathbf{E}_{b\sim\mathcal{B}}[u_i(b_i, b_{-i}, v_i)] \geq \mathbf{E}_{b_{-i}\sim\mathcal{B}_{-i}}[u_i(b'_i, b_{-i}, v_i)]$. In the *incomplete information* setting valuations v are drawn from a product distribution $\mathcal{D} = \prod_{i\in N} \mathcal{D}_i$, and each agent $i \in N$ knows its own valuation v_i and the distributions $\mathcal{D}_{-i} = \prod_{i' \neq i} \mathcal{D}_{i'}$ from which the other agents' valuations are drawn. Instead of a single bid distribution for agent i, we know consider a collection of bid distributions $\mathcal{B}_i(v_i)$, one for each v_i in the support of $\mathcal{D}_i$. We let $\mathcal{B}(v) = \prod_{i\in N} \mathcal{B}_i(v_i)$ and $\mathcal{B}_{-i}(v) = \prod_{i' \neq i} \mathcal{B}_{i'}(v_{i'})$. A collection of distributions $\{\mathcal{B}(v)\}_v$ is a *Bayes-Nash equilibrium* if for each agent i, valuation v_i, and alternative bid b'_i, $\mathbf{E}_{v_{-i}\sim\mathcal{D}_{-i}, b(v)\sim\mathcal{B}(v)}[u_i(b_i, b_{-i}, v_i)] \geq \mathbf{E}_{v_{-i}\sim\mathcal{D}_{-i}, b_{-i}(v)\sim\mathcal{B}_{-i}(v)}[u_i((b'_i, b_{-i}), v_i)]$.

Price of anarchy. We evaluate the quality of mechanisms by their *price of anarchy*. The price of anarchy with respect to Nash equilibria (PoA) is the worst ratio between the optimal social welfare and the expected welfare in a mixed Nash equilibrium. Similarly, the price of anarchy with respect to Bayes-Nash equilibria (BPoA) is the worst ratio between the optimal expected social welfare and the expected welfare in a mixed Bayes-Nash equilibrium. Formally, define $\mathsf{NASH}(v)$ and $\mathsf{BNASH}(\mathcal{D})$ as the set of all mixed Nash and mixed Bayes Nash equilibria respectively. Then,

$$PoA = \max_{v} \max_{\mathcal{B}\in\mathsf{NASH}(v)} \frac{\max_{x\in\Omega} \sum_{i\in N} v_i(x)}{\mathbf{E}_{b\sim\mathcal{B}}[\sum_{i\in N} \mathbf{E}\left[v_i(f(b))\right]]}, \text{ and}$$

$$BPoA = \max_{\mathcal{D}} \max_{\{\mathcal{B}(v)\}_v\in\mathsf{BNASH}(\mathcal{D})} \frac{\mathbf{E}_{v\sim\mathcal{D}}[\sum_{i\in N} \max_{x\in\Omega} v_i(x)]}{\mathbf{E}_{v\sim\mathcal{D}, b\sim\mathcal{B}(v)}[\sum_{i\in N} \mathbf{E}\left[v_i(f(b))\right]]}.$$

The smoothness framework. An important ingredient in our result is the following notion of a smooth mechanism of Syrgkanis and Tardos [46]. A mechanism M = (f,p) is (λ, μ)-*smooth* for $\lambda, \mu \geq 0$ if for all valuation profiles v and all bid profiles b there exists an alternative bid b'_i for each agent i that may depend on the valuation profile v of all agents and the bid b_i of that agent such that

$$\sum_{i\in N} u_i((b'_i, b_{-i}), v_i) \geq \lambda \cdot \max_{x\in\Omega} \sum_{i\in N} v_i(x) - \mu \cdot \sum_{i\in N} \mathbf{E}\left[p_i(b)\right].$$

THEOREM 2 **(Syrgkanis and Tardos [46])**. *If a mechanism is (λ, μ)-smooth and agents have the possibility to withdraw from the mechanism, then the expected social welfare at any mixed Nash or mixed Bayes-Nash equilibrium is at least $\lambda / \max(\mu, 1)$ of the optimal social welfare.*

The smoothness definition in [46] also allows each agent i to deviate to a distribution $\mathcal{B}'_i$ over bids b'_i. Our meta-theorems in the following section continue to hold for this more permissive definition, but none of our smoothness proofs in the subsequent sections requires this additional power.

As shown in [46], (λ, μ)-smoothness also implies a bound of $\max(\mu, 1)/\lambda$ on the price of anarchy for correlated equilibria, which are the outcomes of learning dynamics. Furthermore, the simultaneous and sequential composition of multiple (λ, μ)-smooth mechanisms is again (λ, μ)-smooth.

In fact, our smoothness proofs will show an even slightly stronger property than smoothness, namely semi-smoothness as defined by [8]: the deviation strategy b'_i only depends on the respective agent's valuation v_i, but not on the agent's bid b_i or the other agents' valuations v_{-i}. Therefore, the same price of anarchy bounds also apply to coarse correlated equilibria and Bayes-Nash equilibria with correlated types.

3. Oblivious rounding and smooth relaxations. In this section, we show our main theorems. We consider mechanisms for a problem Π that are constructed as follows. First, one computes an optimal solution x' to a relaxed problem Π' that maximizes the declared welfare. That is, it maximizes $\sum_{i\in N} b_i(x')$. Afterwards, an α-approximate oblivious rounding scheme is applied to derive a feasible solution x to the original problem Π. Each bidder is charged $b_i(x)$, i.e., his declared value for this outcome.

THEOREM 3 **(Main result)**. *Consider problem Π and a relaxation Π'. Given a pay-your-bid mechanism $M' = (f', p')$ that is (λ, μ)-smooth, where f' is an exact declared welfare maximizer for the relaxation Π'. Then a pay-your-bid mechanism $M = (f, p)$ for the original problem Π that is obtained from the relaxation through an α-approximate oblivious rounding scheme is $(\lambda/(2\alpha), \mu)$-smooth.*

In many applications, smoothness is shown by the deviation strategy of reporting half one's true value. First we show that, while generally the deviation strategy b'_i can be arbitrary, it is sufficient to consider only this deviation $b'_i = \frac{1}{2}v_i$. We exploit the fact that f' performs exact optimization.

LEMMA 1. *Given a pay-your-bid mechanism $M = (f, p)$ that is (λ, μ)-smooth where f is an exact declared welfare maximizer. Then M is $(\lambda/2, \mu)$-smooth for deviations to half the value. That is, for all valuation profiles v, bid vectors b, and bids $b'_i = \frac{1}{2}v_i$ for all $i \in N$ it holds that $\sum_{i \in N} u_i((b'_i, b_{-i}), v_i) \geq \frac{\lambda}{2} OPT(v) - \mu \sum_{i \in N} p_i(b)$.*

Proof. We first use (λ, μ)-smoothness of M. For any valuations, there have to be deviation bids fulfilling the respective conditions. So, in particular, let us pretend that each bidder i has valuation $\frac{1}{2}v_i$. By smoothness, there are bids b''_i against b such that

$$\sum_{i \in N} u_i\left((b''_i, b_{-i}), \frac{1}{2}v_i\right) \geq \lambda OPT\left(\frac{v}{2}\right) - \mu \sum_{i \in N} p_i(b). \tag{1}$$

The next step is to relate the sum of utilities that agents with valuations v get in M when they unilaterally deviate from b to b'_i, i.e., $\sum_{i \in N} u_i((b'_i, b_{-i}), v_i) = \sum_{i \in N} \frac{1}{2} v_i(f(b'_i, b_{-i})) = \sum_{i \in N} b'_i(f(b'_i, b_{-i}))$, to the sum of utilities that they get in M with valuations $\frac{1}{2}v$ and unilateral deviations from b to b''_i, i.e., $\sum_{i \in N} u_i((b''_i, b_{-i}), \frac{1}{2}v_i)$.

The allocation function f optimizes exactly over its outcome space. Therefore, it can be used to implement a truthful mechanism $M^{\mathrm{VCG}} = (f, p^{\mathrm{VCG}})$ by applying VCG payments. As VCG payments are non-negative, we get

$$u_i((b'_i, b_{-i}), v_i) = \frac{1}{2} v_i(f(b'_i, b_{-i})) = b'_i(f(b'_i, b_{-i})) \geq b'_i(f(b'_i, b_{-i})) - p^{\mathrm{VCG}}(b'_i, b_{-i}).$$

Observe that the latter term is exactly the utility bidder i receives in M^{VCG} if his valuation and bid is b'_i. As M^{VCG} is truthful, this term is maximized by reporting the true valuation. In other words, it can only decrease, if bidder i changes his bid to b''_i (keeping the valuation b'_i). That is,

$$u_i((b'_i, b_{-i}), v_i) \geq b'_i(f(b'_i, b_{-i})) - p^{\mathrm{VCG}}(b'_i, b_{-i}) \geq b'_i(f(b''_i, b_{-i})) - p_i^{\mathrm{VCG}}(b''_i, b_{-i}).$$

Finally, we use that the VCG payment p^{VCG} is no larger than the pay-your-bid payment p because VCG payments never exceed bids, i.e., $p_i^{\mathrm{VCG}}(b''_i, b_{-i}) \leq b''_i(f(b''_i, b_{-i})) = p_i(b''_i, b_{-i})$. By furthermore changing b'_i back to $\frac{1}{2}v_i$, we get

$$u_i((b'_i, b_{-i}), v_i) \geq \frac{1}{2} v_i(f(b''_i, b_{-i})) - p_i(b''_i, b_{-i}) = u_i\left((b''_i, b_{-i}), \frac{1}{2}v_i\right).$$

Summing this inequality over all $i \in N$ and combining it with inequality (1), we get

$$\sum_{i \in N} u_i((b'_i, b_{-i}), v_i) \geq \lambda OPT\left(\frac{v}{2}\right) - \mu \sum_{i \in N} p_i(b) = \frac{\lambda}{2} OPT(v) - \mu \sum_{i \in N} p_i(b). \quad \square$$

It remains to show that smoothness of the relaxation for deviations to half the value, implies smoothness of the derived mechanism for the original problem. As it is often possible to directly show smoothness for deviations to half the value, we state the following stronger version of Theorem 3 for relaxations that are (λ, μ)-smooth for deviations to half the value.

Theorem 3 follows by first using Lemma 1 to argue that unconstrained (λ, μ)-smoothness of the relaxation implies $(\lambda/2, \mu)$-smoothness for deviations to half the value and then using Theorem 3′ to show that the derived mechanism is $(\lambda/(2\alpha), \mu)$-smooth.

THEOREM 3′ (**Stronger version of main theorem**). If the pay-your-bid mechanism $M' = (f', p')$ that solves the relaxation Π' optimally is (λ, μ)-smooth for deviations to $b'_i = \frac{1}{2}v_i$, then the pay-your-bid mechanism $M = (f, p)$ for Π that is obtained from the relaxation through an α-approximate oblivious rounding scheme is $(\lambda/\alpha, \mu)$-smooth.

Proof. For any valuation profile v and bid vector b, denote the utility of agent $i \in N$ under mechanism $M = (f, p)$ by $u_i(b, v_i) = \mathbf{E}[v_i(f(b)) - p_i(b)]$ and under mechanism $M' = (f', p')$ by $u'_i(b, v_i) = v_i(f'(b)) - p'_i(b)$.

For each bidder i, we consider the unilateral deviation to $b'_i = \frac{1}{2}v_i$. As M is a pay-your-bid mechanism, bidder i's utility when bidding b'_i against b_{-i} can be expressed by

$$u_i((b'_i, b_{-i}), v_i) = \mathbf{E}\left[v_i(f(b'_i, b_i)) - p_i(b'_i, b_{-i})\right] = \frac{1}{2}\mathbf{E}\left[v_i(f(b'_i, b_{-i}))\right].$$

Next we use that the outcome $f(b'_i, b_{-i})$ is derived from $f'(b'_i, b_{-i})$ by applying an α-approximate oblivious rounding scheme. Considering the weight function in which $w_i = v_i$ and $w_{i'} = 0$ for all $i' \neq i$, we conclude that $\mathbf{E}\left[v_i(f(b'_i, b_{-i}))\right] \geq \frac{1}{\alpha}v_i(f'(b'_i, b_{-i}))$. That is, for bidder i's utility, we get

$$u_i((b'_i, b_{-i}), v_i) \geq \frac{1}{2\alpha}v_i(f'(b'_i, b_{-i})) = \frac{1}{\alpha}u'_i((b'_i, b_{-i}), v_i),$$

where the last step uses that M' is a pay-your-bid mechanism as well.

Next, we apply the fact that M' is (λ, μ)-smooth for deviations to $b'_i = \frac{1}{2}v_i$. For the the sum of utilities in M we thus obtain that

$$\sum_{i \in N} u_i((b'_i, b_{-i}), v_i) \geq \frac{1}{\alpha}\sum_{i \in N} u'_i((b'_i, b_{-i}), v_i) \geq \frac{1}{\alpha}\left(\lambda OPT(v) - \mu \sum_{i \in N} p'_i(b)\right).$$

To bound the terms $p'_i(b)$, we use once more the fact that we are applying an α-approximate oblivious rounding scheme, this time to derive $f(b)$ from $f'(b)$ and considering the weight function in which $w_i = b_i$ and $w_{i'} = 0$ for all $i' \neq i$. This implies

$$p'_i(b) = b_i(f'(b)) \leq \alpha \mathbf{E}\left[b_i(f(b))\right] = \alpha \mathbf{E}\left[p_i(b)\right].$$

Overall, we get

$$\sum_{i \in N} u_i((b'_i, b_{-i}), v_i) \geq \frac{1}{\alpha}\lambda OPT(v) - \mu \sum_{i \in N} \mathbf{E}\left[p_i(b)\right],$$

as claimed. $\square$

We note that while we stated our main theorem for pay-your bid mechanisms and for exact optimization over the relaxed solution, both assumptions can be relaxed. We discuss these and further extensions in more detail in Section 8.

Also, as already noted in the introduction, the price of anarchy bounds implied by our main theorem and its strengthening apply even out-of-equilibrium assuming only that players have no regret for deviations to half their value. That is, the utility that they achieve is at least as high as the utility they would get when bidding half their value.

Furthermore, obliviousness is indeed necessary in the sense that it cannot be dropped without any replacement. The reason is that otherwise it would be feasible to use exact optimization over the original problem as rounding procedure, ignoring the relaxed solution entirely. We provide many examples throughout this paper where exact optimization over the original problem results in a mechanism with high price of anarchy whereas the one based on oblivious rounding performs much better.

4. Sparse packing integer programs. In a packing integer program (PIP) with r variables and s constraints we are given weights $w \in \mathbb{R}^r_{\geq 0}$, capacities $c \in \mathbb{R}^s_{\geq 0}$, and a constraint matrix $A \in \mathbb{R}^{s\times r}_{\geq 0}$ and we seek to solve

$$\max\{w^T x \mid x \in \{0,1\}^r, Ax \leq c\}.$$

The *column sparsity* d is the maximum number of non-zero entries in a single column of A. Formally, for each variable x_j, let S_j be the set of constraints in A with a non-zero coefficient, that is, $S_j = \{\ell \mid A_{\ell,j} \neq 0\}$. Now $d = \max_j |S_j|$.

We refer to such PIPs as d-sparse PIPs and consider algorithms that relax the above program by allowing fractional solutions $x \in [0,1]^s$, solve the resulting linear program optimally, and apply an α-approximate oblivious rounding scheme.

Letting $r = nK$, $w = v$, and grouping the variables into n disjoint sets $M_1, \ldots, M_n$ of size K, we obtain an interpretation as a mechanism-design problem in which player i can be served in K different ways and has additive preferences over the options, i.e., $v_i(x) = \sum_{j \in M_i} v_j x_j$.

This way we can encode knapsack auctions (with $d = 1$), multi-unit auctions with general valuations (with $d = 2$), the generalized assignment problem (with $d = 2$), and multi-minded combinatorial auctions in which each agent is interested in K bundles of items whose size is bounded by k (with $d = k+1$).

Theorem 4. *There is an oblivious rounding based, pay-your-bid mechanism for d-sparse packing integer programs that achieves a price of anarchy of $16d(d+1)$.*

To prove Theorem 4 we use the fact that an $8d$-approximate oblivious rounding scheme is available through [4]. In addition, we show in Lemma 2 that the canonical LP relaxation of a d-sparse packing integer program that this rounding scheme is based on is $(1/2, d+1)$-smooth for deviations to $b'_i = \frac{1}{2} v_i$. The claimed bound on the price of anarchy then follows from Theorem 3′.

Lemma 2. *The pay-your-bid mechanism that solves the canonical LP relaxation of a d-sparse packing integer program is $(1/2, d+1)$-smooth for deviations to $b'_i = \frac{1}{2} v_i$.*

To prove this lemma we show the following auxiliary lemma. It shows that the sum of externalities when unilaterally moving to a different fractional solution cannot be much higher than the optimal declared welfare. Given a bid vector b and a capacity vector c, we denote by $W^b(c)$ the value of the optimal LP solution.

Lemma 3. *Let $\bar{x}$ be an arbitrary fractional solution and let $\bar{x}^{(i)}$ denote the solution that is obtained from $\bar{x}$ by setting all variables not belonging to player i to 0. Then,*

$$\sum_{i \in N} \left(W^{b_{-i}}(c) - W^{b_{-i}}(c - A\bar{x}^{(i)}) \right) \leq (d+1) \cdot W^b(c).$$

Proof. Let $\hat{x}$ denote a fractional allocation that maximizes declared welfare for players N with bids b (we have $\hat{x}_j \in [0,1]$ for all j, $A\hat{x} \leq c$).

Now, define LP solution $\hat{x}^{-i}$ by setting $\hat{x}^{-i}_j = \left(1 - \delta^i_j\right) \hat{x}_j$, where $\delta^i_j = \max_{\ell \in S_j} \frac{(A\bar{x}^{(i)})_\ell}{c_\ell}$. Note that $\delta^i_j \leq 1$ for all j and observe that $A\hat{x}^{-i} \leq c - A\bar{x}^{(i)}$.

For all j, we have $\sum_i \delta_j^i \leq \sum_i \sum_{\ell \in S_j} \frac{(A\bar{x}^{(i)})_\ell}{c_\ell} = \sum_{\ell \in S_j} \sum_i \frac{(A\bar{x}^{(i)})_\ell}{c_\ell} \leq |S_j| \leq d$ and therefore $\sum_{i \neq j, i \in N}(1-\delta_j^i) \geq n-d-1$. This gives us

$$\begin{aligned}\sum_{i \in N} W^{b_{-i}}(c - A\bar{x}^{(i)}) &\geq \sum_{i \in N} \sum_{\substack{j \in N \\ j \neq i}} \sum_{k \in M_j} b_k \hat{x}_k^{-i} \\ &= \sum_{j \in N} \sum_{\substack{i \in N \\ i \neq j}} \sum_{k \in M_j} b_k \left(1 - \delta_k^i\right) \hat{x}_k \\ &\geq (n-d-1) \sum_{j \in N} \sum_{k \in M_j} b_k \hat{x}_k \\ &= (n-d-1) W^b(c),\end{aligned}$$

which gives the claimed bound as clearly $W^{b_{-i}}(c) \leq W^b(c)$. □

Proof of Lemma 2. Consider valuations v, bids b and deviations of each player $i \in N$ to $b_i' = 1/2 \cdot v_i$. Denote the optimal fractional allocation for bids (b_i', b_{-i}) by $\bar{x}_1(b_i', b_{-i}), \ldots, \bar{x}_n(b_i', b_{-i})$. Then, by the definition of b_i',

$$u_i((b_i', b_{-i}), v_i) = v_i(\bar{x}_i(b_i', b_{-i})) - b_i'(\bar{x}_i(b_i', b_{-i})) = b_i'(\bar{x}_i(b_i', b_{-i})).$$

Since $\bar{x}_1(b_i', b_{-i}), \ldots, \bar{x}_n(b_i', b_{-i})$ is the fractional allocation that maximizes declared welfare with respect to bids (b_i', b_{-i}),

$$\begin{aligned}b_i'(\bar{x}_i(b_i', b_{-i})) + W^{b_{-i}}(c) &\geq b_i'(\bar{x}_i(b_i', b_{-i})) + \sum_{j \neq i} b_j(\bar{x}_j(b_i', b_{-i})) \\ &\geq b_i'(\bar{x}_i(v)) + W^{b_{-i}}(c - A\bar{x}^{(i)}(v)),\end{aligned}$$

where again $\bar{x}^{(i)}$ denotes the solution that is obtained from $\bar{x}$ by setting all variables not belonging to player i to 0. Rearranging this gives

$$b_i'(\bar{x}_i(b_i', b_{-i})) \geq b_i'(\bar{x}_i(v)) - [W^{b_{-i}}(c) - W^{b_{-i}}(c - A\bar{x}^{(i)}(v))].$$

Summing over all players and applying Lemma 3, we obtain

$$\begin{aligned}\sum_{i \in N} u_i((b_i', b_{-i}), v_i) &= \sum_{i \in N} b_i'(\bar{x}_i(b_i', b_{-i})) \\ &\geq \sum_{i \in N} (b_i'(\bar{x}_i(v)) - [W^{b_{-i}}(c) - W^{b_{-i}}(c - A\bar{x}^{(i)}(v))].) \\ &\geq \sum_{i \in N} b_i'(\bar{x}_i(v)) - (d+1) \cdot \sum_{i \in N} b_i(\bar{x}_i(b)) \\ &= \frac{1}{2} \cdot \sum_{i \in N} v_i(\bar{x}_i(v)) - (d+1) \cdot \sum_{i \in N} b_i(\bar{x}_i(b)),\end{aligned}$$

which completes the proof. □

Our next proposition shows that it is crucial to take the detour via relaxation and rounding: The mechanism that solves the integral problem optimally has an unbounded price of anarchy even when $d = 1$.

PROPOSITION 1. *The pay-your-bid mechanism that maximizes the declared welfare over integral allocations in a multi-unit auction has a pure Nash equilibrium whose welfare is by a factor $(n-2)/2 = m/2$ smaller than the optimal welfare, where n is the number of bidders and m is the number of goods.*

Proof. Consider a setting with n bidders and $m = n-2$ units of an identical good. The valuations for bidders $i = 1, \ldots, n-2$ are $v_{i,0} = 0$ for not receiving any item and $v_{i,k} = 1$ for any set of $k \geq 1$ items. For bidders $i = n-1$ and $i = n$, we set $v_{i,k} = 0$ for $k < m$ and $v_{i,m} = 2$. It is socially optimal to allocate one item each to bidders $1, \ldots, m$, achieving welfare m.

On the other hand, consider bids $b_{i,k} = 0$ for $i = 1, \ldots, m$ and all k and $b_{i,k} = v_{i,k}$ for $i = m+1, m+2$ and all k. The social welfare is 2. It is a pure Nash equilibrium because bidders $1, \ldots, m$ get no items unless they bid at least 2, which is above their value. Bidders $m+1$ and $m+2$ both have zero utility but when lowering the bid they are outbid by the respective other bidder; increasing the bid will result in negative utility. □

We note that in all of the applications mentioned above with $d > 1$, the constraint matrix A has a unit-demand constraint of the form $\sum_{j \in M_i} x_j \leq 1$ for each player i. We remark that it is possible to treat these constraints separately, so that d is the maximum number of constraints excluding the unit-demand constraint that any given variable participates in, while still achieving a price of anarchy guarantee of $16d(d+1)$.

The analysis in this section can remain unchanged. In particular, Lemma 2 and Lemma 3 still apply with this modified definition of d. The rounding scheme in [4] has to be adapted slightly: Rather than setting each variable independently, the randomization has to make sure that at most one variable from each set M_i is set to 1. As the sets M_i are disjoint and $\sum_{j \in M_i} x'_j \leq 1$, this is no problem. We can interpret x'_j for $j \in M_i$ as a probability distribution, and choose $j \in M_i$ with probability $\frac{1}{4d} x'_j$. The rest of the algorithm and the analysis is unchanged, and we obtain a $8d$-approximate oblivious rounding scheme.

Finally, combining Lemma 2 for the new definition of d with the existence of a $8d$-approximate oblivious rounding scheme for the alternate problem formulation we obtain the claimed price of anarchy guarantee through Theorem 3′.

5. Single source unsplittable flow. We next consider the single source unsplittable multi-commodity flow problem, in which we are given a graph $G = (V, E)$ with edge capacities c_e for each edge $e \in E$. All bidders share a source node s and each bidder i has a sink node t_i. Bidder i asks for a path connecting s and t_i fulfilling his demand d_i. His value for this is $v_i \cdot d_i$, and he has no value for less flow than his demand. We assume that capacities and demands are rational numbers, so it is without loss of generality to assume they are integral. We define the capacity-to-demand ratio as $\rho = \min_{e \in E} c_e / \max_{i \in N} d_i$. We assume that the sink t_i and demand d_i for each player is common knowledge, so the player's bid b_i is a claimed value per unit of flow.

Let $\mathcal{P}_i$ be the paths connecting s and t_i. For each $P \in \mathcal{P}_i$, we have a variable $f_{i,P}$ denoting the amount of flow along path P. The problem requires single path routing, that is, all the d_i flow satisfying player i's demand must be carried by a single path. We use the canonical LP relaxation, which maximizes $\sum_{i \in N} \sum_{P \in \mathcal{P}_i} b_i f_{i,P}$ subject to $\sum_{i \in N} \sum_{P \in \mathcal{P}_i : e \in P} f_{i,P} \leq c_e$ for all $e \in E$ and $\sum_{P \in \mathcal{P}_i} f_{i,P} \leq d_i$ for all $i \in N$.

Substituting $f_{i,P}$ by $d_i \bar{x}_{i,P}$, we would get an LP formulation in the spirit of Section 4. However, this LP is not necessarily sparse, as the column sparsity d corresponds to the maximum path length. Nevertheless we are able to establish the following theorem.

THEOREM 5. *There is a constant $c > 0$ such that for all $\epsilon > 0$ for which $\rho \geq c\epsilon^{-1} \log |E|$ there is an oblivious rounding based, pay-your-bid mechanism for the single source unsplittable flow problem with price of anarchy at most $2(1+\epsilon)$.*

For the setting considered here Raghavan and Thompson [39, 40] present a $(1+\epsilon)$-approximate oblivious rounding scheme. So it only remains to prove smoothness of the relaxed problem. The price of anarchy bound then follows by Theorem 3′.

LEMMA 4. *The pay-your-bid mechanism that solves the canonical LP relaxation is* $(1, 1/2)$-*smooth for deviations to* $b'_i = \frac{1}{2} v_i$.

We can represent solutions to the canonical LP relaxation by n-dimensional vectors $x = (x_1, \ldots, x_n)$, where $x_i \in \mathbb{R}_{\geq 0}$ is the amount of flow routed from s to t_i. A well known result by Federgruen and Groenevelt [20] establishes that the set of feasible solutions forms a polymatroid, i.e., there exists a submodular function g such that the set of feasible solutions is

$$Q_g = \left\{ x \in \mathbb{R}^n_{\geq 0} : \sum_{i \in N'} x_i \leq g(N') \ \forall N' \subseteq N \right\}.$$

A solution of maximum weight can thus be found using the standard greedy algorithm for finding a maximum weight basis of a polymatroid [45].

With integral capacities and demands, the submodular function g will in fact be integer valued so that it suffices to consider the corresponding integer polymatroid with $x \in \mathbb{N}^n$.

To prove Lemma 4 it therefore suffices to show the following proposition.

PROPOSITION 2. *Consider an integer polymatroid, in which each dimension corresponds to the allocation of a player and players have linear preferences. Then the pay-your-bid mechanism that maximizes declared welfare is* $(\frac{1}{2}, 1)$-*smooth for deviations to* $b'_i = \frac{1}{2} v_i$.

A key component in our proof of this proposition will be the following generalized Rota exchange property for polymatroids, which we prove in Appendix A.

LEMMA 5. *Let* $\mathcal{Q} = (N, Q)$ *be an integer polymatroid,* α, β *bases in* $\mathcal{Q}$, *and* $q \in \mathbb{N}$. *Let* $\alpha^1, \ldots, \alpha^n \in Q$ *be such that* $\alpha^1_i + \ldots + \alpha^n_i = k \cdot \alpha_i$ *for all* $i \in N$. *Then there are* $\beta^1, \ldots, \beta^n \in Q$ *such that* $\beta^1_i + \ldots + \beta^n_i = k \cdot \beta_i$ *for all* $i \in N$, *and for each* j, $\alpha^j + (\beta - \beta^j) \in Q$.

Proof of Proposition 2. Denote the welfare maximizing solution for valuations v by x^* and the declared welfare maximizing solution for bids b by x. Furthermore, define z as the point-wise minimum of x^* and x, that is, $z_i = \min\{x_i, x^*_i\}$. We now apply Lemma 5 to the polymatroid after contracting z, letting $\alpha = x^* - z$, $\alpha^i = (\alpha_i, 0)$, $\beta = x - z$. Define $y^i = \beta^i + (z_i, 0)$. By this definition, $(x^*_i, x_{-i} - y^i_{-i})$ is feasible and $y^i_i \geq z_i$ for all i. Furthermore $\sum_i y^i = \sum_i b^i + z = x^*$.

Let $S = \{i \mid x_i \geq x^*_i\}$. Consider an arbitrary player i and a deviation to $b'_i = v_i/2$. Denote the resulting solution by x'. If $i \in S$ and $\frac{v_i}{2} \geq b_i$, we have by monotonicity $x'_i \geq x_i \geq x^*_i$. Therefore $x'_i \cdot \frac{v_i}{2} \geq x^*_i \cdot \frac{v_i}{2} \geq x^*_i \frac{v_i}{2} - z_i b_i$. If $i \in S$ and $\frac{v_i}{2} < b_i$, we have $x'_i \cdot \frac{v_i}{2} \geq 0 = (x^*_i - z_i) \cdot \frac{v_i}{2} \geq x^*_i \frac{v_i}{2} - z_i b_i$. So, whenever $i \in S$, we have

$$x'_i \cdot \frac{v_i}{2} \geq 0 \geq x^*_i \frac{v_i}{2} - z_i b_i \geq x^*_i \frac{v_i}{2} - y^i_i b_i \geq x^*_i \frac{v_i}{2} - \sum_{i'} y^i_{i'} b_{i'}.$$

If $i \notin S$, then we use that x' maximizes declared welfare for bids $b' = (b'_i, b_{-i})$,

$$x'_i \cdot \frac{v_i}{2} + \sum_{i' \neq i} x'_{i'} \cdot b_{i'} \geq x^*_i \cdot \frac{v_i}{2} + \sum_{i' \neq i} (x - y^i)_{i'} \cdot b_{i'}. \tag{2}$$

By optimality of x on bids b,

$$\sum_i x_i \cdot b_i \geq \sum_i x'_i \cdot b_i \geq \sum_{i' \neq i} x'_{i'} \cdot b_{i'}. \tag{3}$$

Combining (2) and (3), and using that $i \notin S$, so $x_i - y^i_i = z_i - y^i_i \leq 0$, we get

$$\begin{aligned} x'_i \cdot \frac{v_i}{2} &\geq x^*_i \cdot \frac{v_i}{2} + \sum_{i' \neq i} (x - y^i)_{i'} \cdot b_{i'} - \sum_{i'} x_{i'} \cdot b_{i'} \\ &\geq x^*_i \cdot \frac{v_i}{2} + \sum_{i'} (x - y^i)_{i'} \cdot b_{i'} - \sum_{i'} x_{i'} \cdot b_{i'} = x^*_i \cdot \frac{v_i}{2} - \sum_{i'} y^i_{i'} \cdot b_{i'} \end{aligned}$$

Using that the utility of player i for bid $b'_i = v_i/2$ is $u_i(b'_i, b_{-i}) = x'_i \cdot \frac{v_i}{2}$ and summing over all players i we obtain,

$$\sum_i u_i((b'_i, b_{-i}), v_i) \geq \sum_i x^*_i \cdot \frac{v_i}{2} - \sum_i \sum_{i'} y^i_{i'} \cdot b_{i'} = \frac{1}{2} \sum_i x^*_i v_i - \sum_i x_i b_i$$

as claimed. $\square$

Importantly, the reference to greedy in the above proof is on the level of the relaxation with splittable flows, as the approximation ratio of the greedy algorithm for the original problem with unsplittable flows can be as bad as $\Omega(\sqrt{|E|})$ (see [28]). Also, as we show next, solving the original problem optimally again leads to an unbounded price of anarchy, even if there is a single source, a single sink and just one unit-capacity edge between the two.

PROPOSITION 3. *The pay-your-bid mechanism that solves the single source unsplittable flow problem optimally has a price of anarchy of at least $\frac{m}{2}$ even if there is a single source, a single sink and a unit-capacity edge connecting the two.*

Proof. Consider a network with two nodes, one of which is the source and the other is the sink, and a single edge with unit capacity between the two. There are $m+2$ players. Two big players with demand 1 and value $2 \cdot 1 = 2$, and m small players with demand $1/m$ and value $m \cdot 1/m = 1$.

Note that this setting is effectively the same as the one considered in the proof of Proposition 1. So, again, having the big players both bid 2 and the small players bid 0 is a pure Nash equilibrium with social welfare 2, where the optimal is m. $\square$

6. Combinatorial auctions. In this section, we consider combinatorial auctions. In a combinatorial auction, m items are sold to n bidders. Each item is allocated to at most one bidder and each bidder i has a valuation $v_i(S)$ for the subset $S \subseteq [m]$ of items he receives. The canonical relaxation as a *configuration LP* uses variables $x_{i,S} \in [0,1]$ representing the fraction that bidder i receives of set S. The goal is to maximize $\sum_{i \in N} \sum_{S \subseteq [m]} b_i(S) \bar{x}_{i,S}$ s.t. $\sum_{i \in N} \sum_{S : j \in S} x_{i,S} \leq 1$ for all $j \in [m]$ and $\sum_S x_{i,S} \leq 1$ for all $i \in N$.

For arbitrary valuation functions, only very poor approximation factors can be achieved for the optimization problem. Therefore, we focus on XOS or fractionally subadditive valuations. That is, each valuation function v_i has a representation of the following form. There are values $v^\ell_{i,j} \geq 0$ such that $v_i(S) = \max_\ell \sum_{j \in S} v^\ell_{i,j}$. Feige et al. [22] generalized the class of XOS functions to $\mathcal{MPH}$-k, where XOS is precisely the case $k=1$. A valuation function v_i belongs to class $\mathcal{MPH}$-k if there are values $v^\ell_{i,T} \geq 0$ such that $v_i(S) = \max_\ell \sum_{T \subseteq S, |T| \leq k} v^\ell_{i,T}$.

THEOREM 6. *There is a pay-your-bid mechanism for combinatorial auctions that is based on oblivious rounding and achieves a price of anarchy of $4\frac{e}{e-1}$ for XOS-valuations and of $O(k^2)$ for $\mathcal{MPH}$-k-valuations.*

For general $\mathcal{MPH}$-k-valuations Feige et al. [22] present a $O(k+1)$-approximate rounding scheme; a better constant of $\frac{e}{e-1}$ for the special case XOS can be achieved via the rounding scheme described in [21]. Both schemes are oblivious. Regarding smoothness we show below that the configuration LP that these rounding schemes are based on is $(1/2, k+1)$-smooth for deviations to $b'_i = \frac{1}{2} v_i$. The claimed price of anarchy bounds then follow from Theorem 3′.

As in the previous applications the key lemma in the smoothness proof is the following lemma that bounds the net loss of enforcing a feasible solution one player at a time. For a bid profile b and a vector of quantities q let $W^b(q)$ denote the optimal declared social welfare over all fractional allocations, constrained by capacity vector q.

LEMMA 6. *Let x be an arbitrary fractional solution to the configuration LP. Then,*

$$\sum_{i\in N}\left(W^{b_{-i}}(1)-W^{b_{-i}}(1-x_i)\right)\leq (k+1)\cdot W^b(1).$$

Proof. As removing a player can only decrease the achievable declared welfare, we have $\sum_{i\in N} W^{b_{-i}}(1)\leq \sum_{i\in N} W^b(1)$. It now remains to show that $\sum_{i\in N} W^{b_{-i}}(1-x_i)\geq (n-k-1)\sum_{i\in N} W^b(1)$. Subtracting this inequality from the first one implies then the claim.

Let $\hat{x}$ denote a fractional allocation that maximizes declared welfare for players N with bids b, so $\sum_S b_i(S)\hat{x}_{i,S}=W^b(1)$. Let $\hat{b}_{i,T}$ be the corresponding values such that

$$W^b(1)=\sum_S b_i(S)\hat{x}_{i,S}=\sum_S\left(\sum_{T\subseteq S,|T|\leq k}\hat{b}_{i,T}\right)\hat{x}_{i,S}\quad .$$

In order to bound $W^{b_{-i}}(1-x_i)$ for a fixed i, we will turn $\hat{x}$ into a feasible solution $\hat{x}^{-i}$ for the more restricted constraint capacities $1-x_i$. We derive $\hat{x}^{-i}$ as follows. If $x_{i,A}=1$ for a set $A\subseteq[m]$, then from every set allocated to any other player we remove the intersection with A. That is, the value of $\hat{x}_{i',U}$ is then redirected to the set $U\setminus A$. This procedure generalizes to arbitrary fractional allocation by taking the respective convex combination. For a formal definition, we simplify notation and assume that for every i we have $\sum_A x_{i,A}=1$. This is possible without loss of generality as we can increase $x_{i,\emptyset}$ without modifying the objective function or feasibility. The LP solution $\hat{x}^{-i}$ is now defined by setting $\hat{x}^{-i}_{i',S}=\sum_A x_{i,A}\sum_{U:S=U\setminus A}\hat{x}_{i',U}$.

The first step is to show feasibility of this solution. For all $A\subseteq[m]$ and $j\in[m]$, we have

$$\sum_{i'\neq i}\sum_{S:j\in S}\sum_{U:S=U\setminus A}\hat{x}_{i',U}=\begin{cases}0 & \text{if } j\in A\\ \sum_{i'\neq i}\sum_{U:j\in U}\hat{x}_{i',U} & \text{if } j\notin A.\end{cases}$$

This implies that for all $j\in[m]$

$$\begin{aligned}\sum_{i'\neq i}\sum_{S:j\in S}\hat{x}^{-i}_{i',S}&=\sum_{i'\neq i}\sum_{S:j\in S}\sum_A x_{i,A}\sum_{U:S=U\setminus A}\hat{x}_{i',U}\\&=\sum_A x_{i,A}\sum_{i'\neq i}\sum_{S:j\in S}\sum_{U:S=U\setminus A}\hat{x}_{i',U}\\&=\sum_{A:j\notin A}x_{i,A}\sum_{i'\neq i}\sum_{U:j\in U}\hat{x}_{i',U}\quad .\end{aligned}$$

By feasibility of $\hat{x}$, we have $\sum_{i'\neq i}\sum_{S:j\in S}\sum_{U:S=U\setminus A}\hat{x}_{i',U}$. Furthermore, as we assumed $\sum_A x_{i,A}=1$, we also have $\sum_{A:j\notin A}x_{i,A}=1-\sum_{A:j\in A}x_{i,A}$. Therefore $\sum_{i'\neq i}\sum_{S:j\in S}\hat{x}^{-i}_{i',S}\leq 1-\sum_{A:j\in A}x_{i,A}$. That is, $\hat{x}^{-i}$ is a feasible solution with respect to the capacity vector $q=1-x_i$.

Next, we bound the value of this constructed solution $\hat{x}^{-i}$. Let us first consider the contribution to the declared welfare by player $i'\neq i$ in $\hat{x}^{-i}$. We get

$$\begin{aligned}\sum_S b_{i'}(S)\hat{x}^{-i}_{i',S}&\geq\sum_S\left(\sum_{T\subseteq S,|T|\leq k}\hat{b}_{i',T}\right)\hat{x}^{-i}_{i',S}\\&=\sum_S\left(\sum_{T\subseteq S,|T|\leq k}\hat{b}_{i',T}\right)\sum_A x_{i,A}\sum_{U:S=U\setminus A}\hat{x}_{i',U}\\&=\sum_A x_{i,A}\sum_U\left(\sum_{T\subseteq U\setminus A,|T|\leq k}\hat{b}_{i',T}\right)\hat{x}_{i',U}\quad .\end{aligned}$$

Taking the sum over all $i' \neq i$, this implies

$$W^{b_{-i}}(1-x_i) \geq \sum_{i' \neq i} \sum_A x_{i,A} \sum_U \left(\sum_{T \subseteq U \setminus A, |T| \leq k} \hat{b}_{i',T} \right) \hat{x}_{i',U} \quad .$$

In the remainder, we will bound the sum of all $W^{b_{-i}}(1-x_i)$ and bound it in terms of $W^b(1)$. Using the bound on $W^{b_{-i}}(1-x_i)$ obtained so far and reordering the sums, we get

$$\begin{aligned} \sum_i W^{b_{-i}}(1-x_i) &\geq \sum_{i'} \sum_{i \neq i'} \sum_A x_{i,A} \sum_U \left(\sum_{T \subseteq U \setminus A, |T| \leq k} \hat{b}_{i',T} \right) \hat{x}_{i',U} \\ &= \sum_{i'} \sum_U \left(\sum_{i \neq i'} \sum_A x_{i,A} \sum_{T \subseteq U \setminus A, |T| \leq k} \hat{b}_{i',T} \right) \hat{x}_{i',U} \quad . \end{aligned}$$

By reordering the sums further, we get

$$\begin{aligned} \sum_{i \neq i'} \sum_A x_{i,A} \sum_{T \subseteq U \setminus A, |T| \leq k} \hat{b}_{i',T} &= \sum_{T \subseteq U, |T| \leq k} \hat{b}_{i',T} \sum_{i \neq i'} \sum_{A: A \cap T = \emptyset} x_{i,A} \\ &= \sum_{T \subseteq U, |T| \leq k} \hat{b}_{i',T} \left(\sum_{i \neq i'} \sum_A x_{i,A} - \sum_{i \neq i'} \sum_{A: A \cap T \neq \emptyset} x_{i,A} \right) \quad . \end{aligned}$$

As we assumed $\sum_A x_{i,A} = 1$ for all i, we have

$$\sum_{i \neq i'} \sum_A x_{i,A} = n-1 \quad .$$

Furthermore, we use feasibility of x and the fact that $|T| \leq k$. This implies

$$\sum_{i \neq i'} \sum_{A: A \cap T \neq \emptyset} x_{i,A} \leq \sum_{j \in T} \sum_{i \neq i'} \sum_{A: j \in A} x_{i,A} \leq |T| \leq k \quad .$$

Overall, this implies

$$\sum_i W^{b_{-i}}(1-x_i) \geq (n-k-1) \sum_i \sum_S b_i(S) \hat{x}_{i,S} = (n-k-1) W^b(1) \quad .$$

As $W^b(1) \geq W^{b_{-i}}(1)$ for all i, this shows the claim. □

LEMMA 7. *The pay-your-bid mechanism that solves the configuration LP for $\mathcal{MPH}-k$ valuations exactly is $(1/2, d+1)$-smooth for deviations to $b'_i = \frac{1}{2} v_i$.*

Proof. Following the same steps as in the proof of Lemma 2 and using Lemma 6 instead of Lemma 3 completes the proof. □

7. Maximum traveling salesman. In the asymmetric maximization version of the traveling salesman problem, one is given a complete digraph $G = (V, E)$ with non-negative weights $(w_e)_{e \in E}$. Players are the edges with value w_e for being selected, and the mechanism aims to select a Hamiltonian cycle C that maximizes $\sum_{e \in C} w_e$. We show how existing combinatorial algorithms for this problem can be interpreted as relax-and-round algorithms, and derive the following theorem.

THEOREM 7. *There is a pay-your-bid mechanism for the maximum traveling salesman problem based on oblivious rounding that achieves a price of anarchy of* 9.

We present a proof based on the algorithm of Fisher et al. [24] that yields a slightly worse but easier to prove price of anarchy bound of 12; in Appendix B we show how to improve this bound to 9 using the algorithm of Paluch et al. [37] instead.

We will first argue how the algorithm of Fisher et al. can be interpreted as an oblivious, 2-approximate rounding scheme that relaxes the problem to the problem of finding a maximum-weight cycle cover (defined below). We will then argue that the pay-your-bid mechanism that finds a cycle cover is $(1/2, 3)$-smooth for deviations to half the value. Together with Theorem 3′ these two facts imply the claimed price of anarchy bound.

Fisher et al.'s algorithm uses cycle covers as relaxed solutions. A collection of cycles $C_1, \ldots, C_k$ in a (di-)graph is called a *cycle cover* if each vertex of the graph is contained in exactly one of the cycles. A maximum-weight cycle cover can be computed in polynomial time by computing a maximum-weight perfect matching in a suitably defined bipartite graph. In order to approximate the max-weight TSP tour, one first determines a max-weight cycle cover $C_1, \ldots, C_k$, and then from each of the obtained cycles the minimum-weight edge is dropped, resulting in a collection of vertex-disjoint paths $P_1, \ldots, P_k$. These paths are connected in an arbitrary way to obtain a Hamiltonian cycle C. Going from C_i to P_i, we lose at most half of the weight of this respective cycle. As all weights are non-negative, no weight is lost going from $P_1, \ldots, P_k$ to C. In combination, we have $\sum_{e \in C} w_e \geq \sum_{i \in [k]} \sum_{e \in P_i} w_e \geq \frac{1}{2} \sum_{i \in [k]} \sum_{e \in C_i} w_e$.

The final rounding step, which turns the cycle cover into a tour, can also be modified to work in an oblivious way without loss in the worst case by removing one edge uniformly at random from each cycle. This way, for each edge that was contained in the cycle cover, the individual probability to be also included in the output is at least $\frac{1}{2}$.

To be able to apply Theorem 3′ and obtain the price of anarchy bound it remains to show that the pay-your-bid mechanism that finds a cycle cover is $(1/2, 3)$-smooth for deviations to half the value.

Lemma 8. *The pay-your-bid mechanism for computing an optimal cycle cover is* $(1/2, 3)$*-smooth for deviations to* $b'_i = \frac{1}{2} v_i$.

Our proof of this lemma follows a similar pattern as our proof for sparse packing integer programs. The idea is again to bound the net loss in declared welfare for a given feasible allocation relative to the optimal declared welfare. Let $\mathcal{C}$ denote the set of all cycle covers. Given any bid vector b and any $\mathcal{C}' \subseteq \mathcal{C}$, we write $W^b(\mathcal{C}')$ for the maximum declared welfare of a cycle cover in $\mathcal{C}'$ with respect to b. Letting now $\mathcal{C}_e$ denote the set of all cycle covers that include edge $e \in E$, $W^{b_{-i}}(\mathcal{C}) - W^{b_{-i}}(\mathcal{C}_e)$ is the social cost of including $e \in E$ in the cycle cover. Specifically, we show the following auxiliary lemma.

Lemma 9. *Consider bids b. Let $C_1, \ldots, C_k$ be the cycle cover that maximizes reported welfare for bids b and use E_C to denote the set of edges used in this cycle cover. Consider any other cycle cover $C'_1, \ldots, C'_\ell$ with edge set $E_{C'}$. Then,*

$$\sum_{i \in N : e_i \in E_{C'}} \left(W^{b_{-i}}(\mathcal{C}) - W^{b_{-i}}(\mathcal{C}_{e_i}) \right) \leq 3 \cdot W^b(\mathcal{C}).$$

Proof. Let us first consider any fixed edge $e \in E$ and let us construct a cycle cover from the edge set E_C that contains e.

Figure 1 depicts the two possible cases. The thin edges are edges from E_C. The thick edge is e. Note that we can w.l.o.g. assume that v_1 and v_4 are distinct. (As otherwise e would already be contained in E_C and there would be nothing to show.)

The first case is when nodes v_1 and v_6 are distinct. In this case we can remove edges (v_3, v_1) and (v_6, v_4) from E_C and add edge (v_6, v_1). The resulting edge set is a valid cycle cover because this

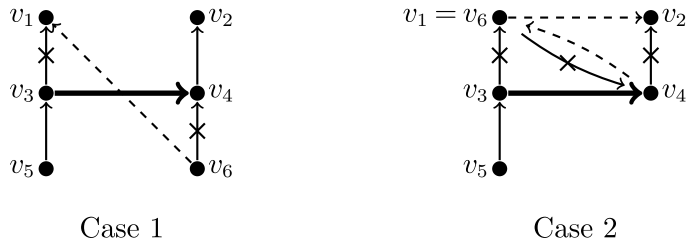

FIGURE 1. Smoothness of Cycle Cover LP

modification to E_C maintains the in-/outdegrees of all nodes and does not create self-loops as we assumed v_1 and v_6 to be distinct.

The second case is when nodes v_1 and v_6 are identical. In this case the modification just described would create a self-loop from/to node $v_1 = v_6$. We can avoid this by instead removing edges (v_3, v_1), (v_6, v_4), and (v_4, v_2) from E_C and adding edges (v_4, v_1) and (v_6, v_2). This leads to a valid cycle cover as it maintains in- and outdegrees and does not create self-loops as we assumed v_1 and v_6 to be identical and v_2 must be different from v_1 (as otherwise this node would have an indegree of 2, which would violate the cycle cover constraint).

In order to bound $\sum_{i \in N: e_i \in E_{C'}} (W^{b_{-i}}(\mathcal{C}) - W^{b_{-i}}(\mathcal{C}_{e_i}))$, we now assume that e was drawn uniformly at random from $E_{C'}$, the set of all edges in C'. As edge weights are non-negative, the loss in declared welfare by forcing e into C, i.e., $W^{b_{-i}}(\mathcal{C}) - W^{b_{-i}}(\mathcal{C}_e)$, is upper-bounded by the weight of all edges removed by our construction. For any edge e', the probability of being removed by our construction is at most $\frac{3}{|E_{C'}|}$. This is due to the fact that e' is only removed if it fulfills a certain role in relation to e, namely being the edge (v_3, v_1), (v_4, v_2), or (v_6, v_4) in Figure 1. As C' is a cycle cover, each edge e' can have each role only with respect to a single $e \in E_{C'}$. Overall, this gives

$$\mathbf{E}\left[W^{b_{-i}}(\mathcal{C}) - W^{b_{-i}}(\mathcal{C}_e)\right] \leq \sum_{e' \in E_C} w_{e'} \frac{3}{|E_{C'}|} = \frac{3}{|E_{C'}|} W^b(\mathcal{C}).$$

Using the definition of the expectation, and multiplying the previous inequality by $|E_{C'}|$ shows the claim. □

Proof of Lemma 8. Following the same steps as in the proof of Lemma 2 and using Lemma 9 instead of Lemma 3 completes the proof. □

8. Extensions. Throughout this paper, we focused on pay-your-bid rules. However, all of our results generalize to payment schemes that use arbitrary non-negative payments which are upper bounded by the respective bid. In this case, we resort to weak smoothness [46]. In our statements (λ, μ)-smoothness would be replaced by weak $(\lambda, 0, \mu)$-smoothness. Considering equilibria without overbidding, i.e., always $b_i(x) \leq v_i(x)$, this implies a price of anarchy bound of $(1+\mu)/\lambda$.

A second observation is that Theorem 3 also holds with a slightly weaker assumption on the rounding. It is sufficient if for all possible valuation profiles each agent is guaranteed to get, in expectation, a $1/\alpha$-fraction of the value that it would have had for the solution to the relaxed problem. That is, $w_i(f(w)) \geq \frac{1}{\alpha} w_i(f'(w))$ for all w.

Furthermore, Theorem 3 also holds if f' is not an exact declared welfare maximizer, but only allows implementation as a truthful mechanism. The interesting consequence is that it might make sense to only approximately solve the relaxation if this improves the smoothness guarantees. For example, a packing LP can be solved using the fractional-overselling mechanism in [26], which was originally introduced in [30]. The allocation rule is an $O(\log n + \log L)$-approximation for any packing LP with n bidders and L constraints between bidders. It allows implementation as a truthful mechanism but it is also a greedy algorithm in the sense of [34]. Therefore, the respective pay-your-bid mechanism is $(\frac{1}{O(\log n+\log L)}, 1)$-smooth. This means that combining this algorithm with any α-approximate oblivious rounding scheme for the respective packing LP, we get a pay-your-bid mechanism with price of anarchy at most $O(\alpha(\log n + \log L))$.

Finally, Carr and Vempala [9] introduced randomized metarounding, which is a technique to derive oblivious rounding schemes from non-oblivious ones. Lavi and Swamy [31] used this result to construct truthful mechanisms. However, they additionally need a packing structure. As in our case oblivious rounding is enough, any rounding scheme derived from the original version in [9] is enough for our considerations.

9. Conclusion and future work. In this paper we showed that oblivious rounding schemes approximately preserve price of anarchy bounds provable via smoothness, and used this result to derive new mechanisms for a broad range of applications.

An interesting direction for future work would be to identify additional algorithm design paradigms that yield mechanisms with low price of anarchy, or more generally to obtain a combinatorial characterization of algorithms with low price of anarchy. A first step towards this direction is [19], which provides such a result for settings where the private information held by each player is a single number and players can either win or lose.

Finally, in the tradition of [31], one could try to reduce mechanism design to algorithm design, by considering mechanisms that have only black-box access to an underlying approximation algorithm, but instead of aiming for a DSIC mechanism that approximately preserves the approximation guarantee of the underlying algorithm the goal would be to approximately preserve the approximation guarantee as a price of anarchy guarantee.

Acknowledgments. We would like to thank the anonymous reviewers of this and an earlier version of this work for their valuable feedback.

Appendix A: Proof of Lemma 5. To prove the lemma, we rely on the following lemma, which establishes a generalized Rota exchange property for matroids.

LEMMA 10 **(Lee et al. [32])**. *Let $\mathcal{M} = (E_{\mathcal{M}}, \mathcal{I})$ be a matroid and A, B bases in $\mathcal{M}$. Let $A_1, \dots, A_n$ be subsets of A such that each element of A appears in exactly k of them. Then there are subsets $B_1, \dots, B_n$ of B such that each element of B appears in exactly k of them, and for each j, $A_j \cup (B - B_j) \in \mathcal{I}$.*

To reduce our statement to matroids, we first define a vector m by setting $m_i = \min\{\alpha_i, \beta_i\}$. Clearly m is in Q, so we can contract m in $\mathcal{Q}$. Therefore, without loss of generality, we can assume that for each i we have $\alpha_i = 0$ or $\beta_i = 0$.

Now, let us consider the equivalent matroid $\mathcal{M} = (E_{\mathcal{M}}, \mathcal{I})$ (see, e.g., Chapter 44.6b of [45]): Form a ground set $E_{\mathcal{M}} = \cup_i E_i$ by introducing, for each polymatroid dimension $i \in N$, sufficiently many distinct elements E_i. Associate with each set of elements $I \subseteq E_{\mathcal{M}}$ a vector $s(I) \in \mathbb{N}_{\geq 0}^N$, where $s_i(I) = |I \cap E_i|$. A set $I \subseteq E_{\mathcal{M}}$ is independent if $s(I) \in Q$.

To apply Lemma 10, we let A and B be an arbitrary set of elements such that $|A \cap E_i| = \alpha_i$ and $|B \cap E_i| = \beta_i$ for all i. By equivalence of the matroid and polymatroid, A and B are bases in $\mathcal{M}$. Furthermore, let $A_1, \dots, A_n$ be any division of A into subsets such that each $e \in A$ is contained in exactly k of them.

By Lemma 10, there are subsets $B_1, \dots, B_n$ of B such that each element of B appears in exactly k of them, and for each j, $A_j \cup (B - B_j) \in \mathcal{I}$. Define $\beta_i^j = |B_j \cap E_i|$. By this construction $\sum_j \beta_i^j = k\beta_i$ because each element of B appears in exactly k sets $B_1, \dots, B_n$. It remains to show that $\alpha^j + (\beta - \beta^j) \in Q$ for all j. To this end, we use that $s_i(A_j \cup (B - B_j)) = |(A_j \cup (B - B_j)) \cap E_i| = (\alpha^j + (\beta - \beta^j))_i$ for all i. This is true for the following reason. If $\alpha_i = 0$, then also $\alpha^j = 0$ and $A_j = \emptyset$. So, the identity follows from the definitions of B and β^j. Otherwise, we have $\beta_i = 0$. So, also $B \cap E_i, B_j \cap E_i = \emptyset$ and the claim follows by the definition of α^j.

The proof is completed by the fact that $s(A_j \cup (B - B_j)) \in Q$ because the set is independent in $\mathcal{M}$.

Appendix B: Proof of Theorem 7. The approximation guarantee of the algorithm of Fisher et al. [24] could be improved to 3/2 if one could find a cycle cover that does *not* contain cycles of length two. Since finding such a cover is APX-hard, Paluch et al. [37] relax the problem even further by allowing so called half-edges.

More formally: From the original digraph $G = (V, E)$ with weights $w_e \in \mathbb{R}_+$ construct a new digraph $\tilde{G} = (\tilde{V}, \tilde{E})$ with weights $\tilde{w}_{\tilde{e}} \in \mathbb{R}_+$ as follows. First add all vertices $v_i \in V$ to $\tilde{V}$. Then, for each edge $(v_i, v_j) \in E$ add a vertex $v_{i,j}$ to $\tilde{V}$ and edges $(v_i, v_{i,j})$ and $(v_{i,j}, v_j)$ to $\tilde{E}$. Give each of these edges a weight of $w_{(v_i,v_j)}/2$. Now *a cycle cover without 2-cycles but with half-edges* is a collection of edges such that: (1) each vertex in the original graph $v_i \in V \subseteq \tilde{V}$ has in- and outdegree exactly one, (2) for each pair of edges in the original graph (v_i, v_j) and (v_j, v_i) either (a) none of the edges $(v_i, v_{i,j})$, $(v_{i,j}, v_j)$, $(v_j, v_{j,i})$, $(v_{j,i}, v_i)$ is used or (b) the cycle cover contains exactly two of the four edges, one incident to v_i and one incident to v_j. In other words, it is possible to take both heads or both tails of the edges between v_i and v_j.

Paluch et al. show that a cycle cover $\tilde{C}$ without 2-cycles but with half-edges can be computed in polynomial time. Furthermore, it is possible to derive from it three node disjoint paths P_1, P_2, P_3 in the original graph, whose weight is at least twice the weight of $\tilde{C}$. Since the optimal tour yields a cycle cover without 2-cycles but with half edges, choosing the path with maximum weight and extending it to a tour yields a 3/2-approximation. Since the derivation of the paths requires no knowledge of the weights, choosing one of the paths at random and extending it to a tour without looking at the weights yields an oblivious rounding scheme.

Our proof of the price of anarchy guarantee now follows a similar pattern as the proof for the pay-your-bid mechanism based on Fisher et al's algorithm. However, as each player now controls two edges, a direct translation of the argument would only show $(1/2, 6)$-smoothness for deviations to half the value and hence a price of anarchy of 18 via Theorem 3′. We therefore slightly deviate from our proof pattern by bounding only the net cost of enforcing a solution to the original problem, and showing that this suffices for our main theorem to go through.

LEMMA 11. *Consider any bid profile b. Let $\tilde{\mathcal{C}}$ be the set of all cycle covers without 2-cycles but with half-edges and for any edge e in the original graph let $\tilde{\mathcal{C}}_e$ be the set of all cycle covers without 2-cycles but with half-edges that contain edge e. Let T be any tour in the original graph with corresponding edge set E_T. Then,*

$$\sum_{i \in N : e_i \in E_T} \left(W^{b_{-i}}(\tilde{\mathcal{C}}) - W^{b_{-i}}(\tilde{\mathcal{C}}_{e_i}) \right) \leq 3 \cdot W^b(\tilde{\mathcal{C}}).$$

Proof. Let $\tilde{C}$ be the cycle cover without 2-cycles but with half edges that maximizes reported welfare for bids b and use $\tilde{E}_{\tilde{C}}$ to denote the set of edges used in this cycle cover. We begin by showing how to incorporate a given edge $e \in E_T$ from the tour T into the cycle cover $\tilde{C}$ by removing as few as possible edges from $\tilde{E}_{\tilde{C}}$.

Figure 2 depicts the possible configurations that can occur. The thick edge is e. One possible choice of the set of edges $\tilde{E}_{\tilde{C}}$ are the thin edges. Some of these edges can be replaced by alternative edges. The alternatives are drawn as dotted edges.

We will distinguish two cases. The case where v_1 and v_6 are distinct and there is no edge between v_1 and v_6 in the cycle cover, and the case where either v_1 and v_6 are identical or there is an edge between v_1 and v_6 in the cycle cover. Note that if we have $v_1 = v_4$ then v_6, which is distinct from

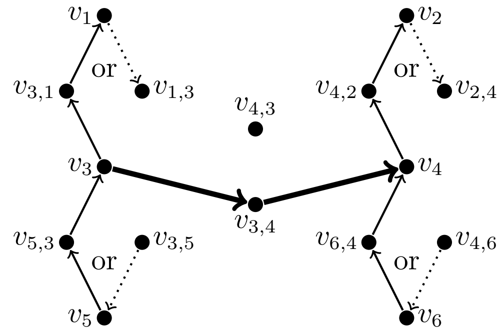

FIGURE 2. Possible Configurations of the Cycle Cover LP without 2-Cycles but with Half Edges

v_4, must be distinct from v_1. Hence the case $v_1 = v_4$ is covered by the first case, and in the second case we can assume that $v_1 \neq v_4$.

In the first case we remove the edges between v_3 and v_1 and between v_6 and v_4 from $\tilde{E}_{\tilde{C}}$, and add an edge between v_1 and v_6. We direct the half-edges in this edge so that they fit the in-/outgoing edges at v_1 and v_6. The resulting edge set is a valid cycle cover without 2-cycles but with half-edges because this modification to E_C maintains the in-/outdegrees of all nodes, does not create self-loops (as in this case we assumed that $v_1 \neq v_6$), and does not create a 2-cycle (as in this case we assumed that there was no edge between v_1 and v_6).

In the second case we remove edges between v_3 and v_1, between v_6 and v_4, and between v_4 and v_2 from $\tilde{E}_{\tilde{C}}$, and add an edge between v_4 and v_1 and an edge between v_6 and v_2. We direct the half-edges so that they fit with the in- and outgoing edges at v_1, v_2, and v_6. This leads to a valid cycle cover without 2-cycles but with half edges as it maintains in- and outdegrees and does not create self-loops or 2-cycles. To see that it does not create self-loops note that in this case we can assume that $v_1 \neq v_4$, and $v_2 = v_6$ would imply that there was a 2-cycle in the cycle cover. To see that it does not create 2-cycles first observe that we need not be worried about an edge between v_1 and v_4 because in the case where $v_1 = v_6$ we removed this edge, and in the case where $v_1 \neq v_6$ we assumed that there is an edge between v_1 and v_6 and so if there was an edge between v_1 and v_4 in the cycle cover then v_1 would have had three incident edges. By a similar argument we also need not be worried about an edge between v_2 and v_6. Namely, in the case where $v_1 = v_6$ this would imply $v_2 \neq v_3$ because otherwise there would have been a 2-cycle involving $v_1 = v_6$ and $v_2 = v_3$, but then v_6 would have had three incident edges. Otherwise, $v_1 \neq v_6$ and there is an edge between v_1 and v_6. Then either $v_2 \neq v_1$ in which case v_6 would have had three incident edges, or $v_2 = v_1$ in which case v_2 would have had three incident edges.

We conclude that in order to add any edge $e \in E_T$ from the tour we need to remove at most three edges from the cycle cover. Since each removed edge plays each role in the above construction at most once, we obtain

$$\sum_{i \in N : e_i \in E_T} \left(W^{b_{-i}}(\tilde{\mathcal{C}}) - W^{b_{-i}}(\tilde{\mathcal{C}}_{e_i}) \right) \leq 3 \cdot W^b(\tilde{\mathcal{C}}). \qquad \square$$

LEMMA 12. *Consider valuation profile v and bid profile b. Denote the welfare achieved by the welfare maximizing tour by $OPT_T(v)$. Then for the pay-your-bid mechanism that computes an optimal cycle cover without 2-cycles but with half-edges and bids $b'_i = \frac{1}{2} v_i$ for all $i \in N$,*

$$\sum_{i \in N} u_i((b'_i, b_{-i}), v_i) \geq \frac{1}{2} OPT_T(v) - 3 \cdot \sum_{i \in N} p_i(b).$$

Proof. The proof is analogous to the proof of Lemma 2. The only two differences are that (1) instead of switching to the optimal relaxed solution for v we switch to the optimal original solution T for v and (2) we then apply Lemma 11 instead of Lemma 3 to bound the social cost $\sum_{i \in N: e_i \in E_T} (W^{b_{-i}}(\tilde{\mathcal{C}}) - W^{b_{-i}}(\tilde{\mathcal{C}}_{e_i})$ of enforcing the optimal solution T. $\square$

To establish the price of anarchy guarantee we can now follow the same steps as in proof of Theorem 3′. The only exception is that instead of invoking smoothness for deviations to half the value for the pay-your-bid mechanism for the relaxation and lower bounding the optimal fractional solution with the optimal original solution, we directly apply Lemma 12.

References

[1] Babaioff M, Kleinberg R, Slivkins A (2013) Multi-parameter mechanisms with implicit payment computation. *Proceedings of the 13th Conference on Economics and Computation*, 35–52.

[2] Babaioff M, Kleinberg RD, Slivkins A (2015) Truthful mechanisms with implicit payment computation. *Journal of the ACM* 62(2):10:1–10:37.

[3] Babaioff M, Lucier B, Nisan N, Leme RP (2014) On the efficiency of the walrasian mechanism. *Proceedings of the 15th Conference on Economics and Computation*, 783–800.

[4] Bansal N, Korula N, Nagarajan V, Srinivasan A (2010) On k-column sparse packing programs. *Proceedings of the 14th Conference on Integer Programming and Combinatorial Optimization*, 369–382.

[5] Bhawalkar K, Roughgarden T (2011) Welfare guarantees for combinatorial auctions with item bidding. *Proceedings of the 22nd Symposium on Discrete Algorithms*, 700–709.

[6] Briest P, Krysta P, Vöcking B (2011) Approximation techniques for utilitarian mechanism design. *SIAM Journal on Computing* 40(6):1587–1622.

[7] Cai Y, Papadimitriou CH (2014) Simultaneous bayesian auctions and computational complexity. *Proceedings of the 15th Conference on Electronic Commerce*, 895–910.

[8] Caragiannis I, Kaklamanis C, Kanellopoulos P, Kyropoulou M, Lucier B, Leme RP, Tardos É (2015) On the efficiency of equilibria in generalized second price auctions. *Journal of Economic Theory* 343–388.

[9] Carr RD, Vempala S (2002) Randomized metarounding. *Random Structures and Algorithms* 20(3):343–352.

[10] Chawla S, Immorlica N, Lucier B (2012) On the limits of black-box reductions in mechanism design. *Proceedings of the 44th Symposium on Theory of Computing*, 435–448.

[11] Christodoulou G, Kovács A, Schapira M (2016) Bayesian combinatorial auctions. *Journal of the ACM* 63(2):11:1–11:19.

[12] Clarke E (1971) Multipart pricing of public goods. *Public Choice* 11:17–33.

[13] Daskalakis C, Syrgkanis V (2016) Learning in auctions: Regret is hard, envy is easy. *Proceedings of the 57th Symposium on Foundations of Computer Science*, 219–228.

[14] de Keijzer B, Markakis E, Schäfer G, Telelis O (2013) Inefficiency of standard multi-unit auctions. *Proceedings of the 21st Annual European Symposium on Algorithms*, 385–396.

[15] Dobzinski S, Fu H, Kleinberg RD (2015) On the complexity of computing an equilibrium in combinatorial auctions. *Proceedings of the 26th Symposium on Discrete Algorithms*, 110–122.

[16] Dughmi S, Roughgarden T (2014) Black-box randomized reductions in algorithmic mechanism design. *SIAM Journal on Computing* 43(1):312–336.

[17] Dughmi S, Roughgarden T, Yan Q (2016) Optimal mechanisms for combinatorial auctions and combinatorial public projects via convex rounding. *Journal of the ACM* 63(4):30:1–30:33.

[18] Dütting P, Henzinger M, Starnberger M (2018) Valuation compressions in VCG-based combinatorial auctions. *ACM Transactions on Economics and Computation* 6:5:1–5:19.

[19] Dütting P, Kesselheim T (2015) Algorithms against anarchy: Understanding non-truthful mechanisms. *Proceedings of the 16th Conference on Economics and Computation*, 239–255.

[20] Federgruen A, Groenevelt H (1986) Optimal flows in networks with multiple sources an sinks, with applications to oil and gas lease investment programs. *Operations Research* 34.

[21] Feige U (2009) On maximizing welfare when utility functions are subadditive. *SIAM Jornal on Computing* 39(1):122–142.

[22] Feige U, Feldman M, Immorlica N, Izsak R, Lucier B, Syrgkanis V (2015) A unifying hierarchy of valuations with complements and substitutes. *Proceedings of the 29th Conference on Artificial Intelligence*, 872–878.

[23] Feldman M, Fu H, Gravin N, Lucier B (2018) Simultaneous auctions without complemets are (almost) efficient. *Games and Economic Behavior* In Press.

[24] Fisher ML, Nemhauser GL, Wolsey LA (1979) An analysis of approximations for finding a maximum weight hamiltonian circuit. *Operations Research* 27(4):799–809.

[25] Groves T (1973) Incentives in teams. *Econometrica* 41:617–631.

[26] Hoefer M, Kesselheim T, Vöcking B (2016) Truthfulness and stochastic dominance with monetary transfers. *ACM Transactions on Economics and Computation* 4(2):11:1–11:18.

[27] Kaplan H, Lewenstein M, Shafrir N, Sviridenko M (2005) Approximation algorithms for asymmetric TSP by decomposing directed regular multigraphs. *Journal of the ACM* 52(4):602–626.

[28] Kleinberg J (1996) *Approximation Algorithms for Disjoint Paths Problems* (Ph.D. Thesis, Department of EECS, MIT).

[29] Koutsoupias E, Papadimitriou CH (1999) Worst-case equilibria. *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science*, 404–413.

[30] Krysta P, Vöcking B (2012) Online mechanism design (randomized rounding on the fly). *Proceedings of the 39th International Colloquium on Automata, Languages and Programming*, 636–647.

[31] Lavi R, Swamy C (2011) Truthful and near-optimal mechanism design via linear programming. *Journals of the ACM* 58(6):25:1–25:24.

[32] Lee J, Sviridenko M, Vondrák J (2010) Submodular maximization over multiple matroids via generalized exchange properties. *Mathematics of Operations Research* 35(4):795–806.

[33] Lehmann B, Lehmann DJ, Nisan N (2006) Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55(2):270–296.

[34] Lucier B, Borodin A (2017) Equilibria of greedy combinatorial auctions. *SIAM Journal on Computing* 46(2):620–660.

[35] Lucier B, Syrgkanis V (2015) Greedy algorithms make efficient mechanisms. *Proceedings of the 16th Conference on Economics and Computation*, 221–238.

[36] Markakis E, Telelis O (2015) Uniform price auctions: Equilibria and efficiency. *Theory of Computing Systems* 57(3):549–575.

[37] Paluch KE, Elbassioni KM, van Zuylen A (2012) Simpler approximation of the maximum asymmetric traveling salesman problem. *Proceedings of the 29th International Symposium on Theoretical Aspects of Computer Science*, 501–506.

[38] Papadimitriou CH, Schapira M, Singer Y (2008) On the hardness of being truthful. *Proceedings of the 49th Symposium on Foundations of Computer Science.*

[39] Raghavan P (1988) Probabilistic construction of deterministic algorithms: Approximating packing integer programs. *Journal on Computer Systems Sciences* 37(2):130–143.

[40] Raghavan P, Thompson C (1987) Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica* 7:365–374.

[41] Rothkopf M, Pekec A, Harstad R (1998) Computationally manageable combinatorial auctions. *Management Science* 44:1131–1147.

[42] Roughgarden T (2015) Intrinsic robustness of the price of anarchy. *Journal of the ACM* 62(5):32:1–32:42.

[43] Roughgarden T (2015) The price of anarchy in games of incomplete information. *ACM Transactions on Economics and Computation* 3(1):6:1–6:20.

[44] Sandholm T (2002) Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence* 135(1):1 – 54.

[45] Schrijver A (2003) *Combinatorial Optimization - Polyhedra and Efficiency* (Springer-Verlag Berlin Heidelberg).

[46] Syrgkanis V, Tardos É (2013) Composable and efficient mechanisms. *Proceedings of the 45th Symposium on Theory of Computing*, 211–220.

[47] Vazirani VV (2001) *Approximation Algorithms* (Springer-Verlag Berlin Heidelberg).

[48] Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16:8–37.