Oblivious Rounding and the Integrality Gap*

Uriel Feige[†1], Michal Feldman[‡2], and Inbal Talgam-Cohen[§3]

1 **Department of Computer Science, Weizmann Institute of Science, Israel**
uriel.feige@weizmann.ac.il
2 **School of Computer Science, Tel-Aviv University, and Microsoft Research, Israel**
michal.feldman@cs.tau.ac.il
3 **School of Computer Science, Hebrew and Tel-Aviv Universities, Israel**
inbal.talgamcohen@mail.huji.ac.il

Abstract

The following paradigm is often used for handling NP-hard combinatorial optimization problems. One first formulates the problem as an integer program, then one relaxes it to a linear program (LP, or more generally, a convex program), then one solves the LP relaxation in polynomial time, and finally one rounds the optimal LP solution, obtaining a feasible solution to the original problem. Many of the commonly used rounding schemes (such as randomized rounding, threshold rounding and others) are *oblivious* in the sense that the rounding is performed based on the LP solution alone, disregarding the objective function. The goal of our work is to better understand in which cases oblivious rounding suffices in order to obtain approximation ratios that match the integrality gap of the underlying LP. Our study is information theoretic – the rounding is restricted to be oblivious but not restricted to run in polynomial time. In this information theoretic setting we characterize the approximation ratio achievable by oblivious rounding. It turns out to equal the integrality gap of the underlying LP on a problem that is the *closure* of the original combinatorial optimization problem. We apply our findings to the study of the approximation ratios obtainable by oblivious rounding for the maximum welfare problem, showing that when valuation functions are submodular oblivious rounding can match the integrality gap of the configuration LP (though we do not know what this integrality gap is), but when valuation functions are gross substitutes oblivious rounding cannot match the integrality gap (which is 1).

1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity

Keywords and phrases Welfare-maximization

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2016.8

1 Introduction

Rounding and Obliviousness

Consider a combinatorial maximization problem π, represented by a pair (V, X). The set V contains all possible problem *instances*, where an instance is the linear objective function to

* This work was partly done at Microsoft Research, Herzliya, Israel.
† The work of U. Feige was supported in part by the Israel Science Foundation (grant No. 621/12) and by the I-CORE Program of the Planning and Budgeting Committee and the Israel Science Foundation (grant no. 4/11).
‡ The work of M. Feldman was supported in part by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122.
§ The work of I. Talgam-Cohen was supported in part by the I-CORE Program of the Planning and Budgeting Committee and the Israel Science Foundation (grant no. 4/11), and by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122.

be maximized, represented as a vector in $\mathbb{R}^d_{\geq 0}$. The set X contains all feasible *solutions* to the problem, also represented as vectors in $\mathbb{R}^d_{\geq 0}$. The goal is, given an instance $v \in V$, to return a feasible solution $x \in X$ that maximizes the objective $v \cdot x$ among all feasible solutions. If the combinatorial problem is hard, the goal is to approximate rather than optimize the objective. As a concrete example, consider the problem of finding a max-cut in a complete weighted graph. In this case, V is the set of all possible edge weights, and X is the set of all valid cuts where each cut is represented by the set of edges in the cut. The objective value $v \cdot x$ that a cut x obtains for a weighted graph v is the total weight of edges in the cut.

A paradigmatic approach to solving combinatorial optimization problems is that of *relaxation and rounding*: The problem π is relaxed to a new problem $\pi' = (V, Y)$ where Y is such that $X \subset Y$, i.e., the new feasible solution set is a relaxation of the original one. Typically the new feasible domain is fractional while the original one is integral. To solve a given instance v, first a relaxed solution $y \in Y$ to the new problem π' is computed, and then it is (randomly) rounded to get a solution $x \in X$ to the original problem π. The algorithm designer aims to design a rounding process that does not lose too much in the objective value, i.e., for which the inner product $v \cdot x$ is not far in a multiplicative sense (and in expectation) from $v \cdot y$. If she succeeds we say that the rounding scheme guarantees a good *approximation ratio*. A rounding scheme is *oblivious* if y is rounded to x without knowledge of the objective function v.[1] In other words, v is used only to obtain a relaxed solution (e.g., to formulate and solve a linear program), and not to round it back to a feasible solution (e.g., a solution to a corresponding integer program).

Many rounding schemes in the optimization literature are oblivious, and many are not oblivious (see Section 5). This raises the following natural question: Is there a reason why for some problems oblivious rounding works well (achieves good approximation ratios to the optimal objective), while for others it fails miserably? For an algorithm designer it may be very useful to be able to predict in advance whether the relaxation she has formulated for the problem admits an oblivious rounding scheme with a good approximation ratio, or whether any good scheme will need to utilize the objective function to guide its rounding process. *The purpose of this paper is to initiate a systematic study of the power of oblivious rounding relative to its non-oblivious counterpart.* We study this question from an *information* perspective, imposing no polynomial time constraint on the rounding schemes. We remark that even non-polynomial time rounding schemes are of interest, for example, as a way of bounding the integrality gap of the underlying relaxation.

Advantages of Oblivious Rounding

There is also reason to try and aim specifically for a relaxation that admits good oblivious rounding, and/or to be able to prove the impossibility of getting a good approximation via oblivious rounding. The advantages of rounding that is oblivious are demonstrated nicely in the context of *welfare maximization in combinatorial auctions*, which will be the main domain in which we demonstrate the results of our study of oblivious rounding (see Section 4 for more details on welfare maximization). In this context, indivisible items are to be allocated among buyers, each with her own valuation function mapping bundles of items to values. The valuations are very large objects (exponential in the number of items), and there is extensive literature related to their communication complexity (see, e.g., [20]). Oblivious

[1] Our notion of oblivious rounding is not to be confused with the rounding technique of [27], which avoids solving a linear program – see the discussion of related work below.

rounding limits the algorithmic stage in which communication in required, and there is no need for communication after a relaxed solution is found. Also, as we show in Proposition 23, oblivious rounding gives the different buyers the same treatment in terms of the value they are guaranteed to obtain after the rounding, and so has a "built-in" fairness guarantee.

Recently in [7], oblivious rounding was studied in the context of incentive properties of allocation mechanisms. It turns out that when an algorithm is based on the relax-and-round paradigm, and the rounding is oblivious, there are price rules that can be added to the algorithm such that the worst equilibrium behavior (the price of anarchy) is determined by the relaxation and by the approximation ratio of the oblivious rounding. This is quite remarkable, as there is no a priori reason to believe that the consequences of strategic behavior would be determined by algorithmic properties of the rounding, and indeed this is not the case for non-oblivious rounding. Thus, an algorithmic mechanism designer may aim for a design based on obliviousness to get good strategic properties, and so it would be helpful to understand what a design based on oblivious rounding can hope to achieve.

Finally, in [12, 24], the issue of robustness of the welfare guarantees to noise in the objective function is studied. Ideally, an algorithm for approximating welfare will get a good approximation despite small perturbations in the buyers' valuations. In the common case that the welfare maximization problem is relaxed to a linear program which is then solved and rounded, it turns out that solving the LP is quite robust, and so if the rounding is oblivious this ensures the robustness of the entire algorithm.

Our Results

Consider a problem $\pi = (V, X)$ and a relaxed problem $\pi' = (V, Y)$. Our main result is to relate the approximation ratio achievable by oblivious rounding to the well-studied notion of *integrality gap*.

In the context of our work, we define the approximation ratio to be the worst case ratio, over all instances $v \in V$ and all relaxed solutions $y \in Y$, between the (expected) objective value $v \cdot x$ achieved by the (random) rounded solution $x \in X$, and between the objective $v \cdot y$ achieved by the relaxed solution to be rounded y. Note that there is another notion of approximation ratio, which compares $v \cdot x$ achieved by the rounding to $v \cdot x^*$ (rather than $v \cdot y$), where x^* is the optimal feasible solution to instance v. While different in general, in many cases the two notions coincide.

On the other hand, recall that the integrality gap is the worst case ratio, over all instances v and all relaxed solutions y, between the objective value $v \cdot x$ that can be achieved by the *best* feasible solution x, and between $v \cdot y$. As our starting point, we observe that no oblivious rounding can guarantee a better approximation factor than the integrality gap. Thus the question that we ask is: *For which problems does the approximation ratio achievable by oblivious rounding techniques match the integrality gap?* We stress that we do not require oblivious rounding to be polynomial time, but nevertheless the question is of interest due to the information-theoretic obliviousness requirement. This question also comes in another flavor, where one gets an *optimal* solution to the relaxed problem and needs to round it.

Our general results can be summarized informally by the following theorem. The *convex closure* of a problem π is obtained by taking the convex closure of its instance set. This may not actually change the problem, i.e., the instance set may be closed under convex combinations. For example, welfare maximization in a combinatorial auction setting with *submodular* valuations is an example of such a problem (because submodularity is preserved under convex combinations), but with *gross substitutes* (*GS*) valuations it is not (the average of two GS functions need not be GS – see Section 4). The convex closure of the class of GS valuations is the class *cone GS* (*CGS*) defined in [6].

▶ **Theorem** (Informal).

- *For optimization problems closed under convex combinations, the approximation ratio of the best oblivious rounding scheme equals the integrality gap.*
- *More generally, for optimization problems that are not closed under convex combinations, the approximation ratio of the best oblivious rounding scheme equals the integrality gap of the convex closure of the problem.*
- *If the relaxed solution to be rounded obliviously is guaranteed to be optimal, the approximation ratio of the best oblivious rounding scheme is at least the integrality gap of the convex closure of the problem, and may be strictly greater than it in some cases.*

See Section 3 for formal statements of these results.

We apply our general results to the welfare maximization problem for combinatorial auctions. In particular, we use the integrality gap of welfare maximization with *coverage* valuations – the convex closure of unit-demand valuations – to establish a bound on what the best oblivious rounding can achieve for unit-demand valuations.

▶ **Theorem.** *For the welfare maximization problem with unit-demand valuations and for its relaxation based on the* configuration linear program*, no oblivious rounding can get more than a* $5/6$*-approximation ratio for two buyers, and no oblivious rounding can get more than a* 0.782*-approximation ratio for* n *buyers. These bounds immediately extend to gross substitutes, for which the integrality gap is known to be 1.*

Another application of our general results to welfare maximization is the prediction that the above gap, which occurs between the integrality gap and the approximation ratio of the best oblivious rounding for unit-demand valuations, will *not* occur for classes of valuations like submodular valuations, which are closed under convex combinations. See Section 4 for formal statements of these results.

Related Work

In recent years, the connection between various notions of rounding and algorithmic mechanism design has been studied in several works. [15] use the technique of randomized metarounding [3] to derive truthful-in-expectation mechanisms. They require their rounding-based approximation algorithms to satisfy a stronger property than obliviousness (the output expected allocation should be a scaled version of the input for a universal scaling factor). We have already mentioned the work of [7] above, which is directly related to the notion of oblivious rounding that we study (see also Proposition 23 below). Both works and the latter in particular can be seen as strong motivation to systematically study oblivious rounding. [6] requires a different property – convexity of the rounding – in order to derive truthful-in-expectation mechanisms.

In terms of techniques, our work is related to that of [9], which considers a class of *oblivious algorithms* for the max *directed cut* problem. These are algorithms in which each vertex independently decides at random on which side of the cut to place itself, based only on its own in-degree and its own out-degree. One of the results in that work (Theorem 1.8 in the journal version, Theorem 1.5 in the preliminary version) shows equivalence in the worst case approximation ratio of two different ways of using a finite set of oblivious algorithms, one called *mixed* (in which an algorithm is chosen at random), the other called *max* (in which the best algorithm is chosen). The proof of that theorem and the proof of our main theorem are based on similar principles.

[27] introduced a technique for developing approximation algorithms that avoid the bottleneck of first solving a linear program. This technique is also known as "oblivious rounding", but this notion is different than our definition of (objective-)oblivious rounding.

Examples of oblivious rounding techniques that appear in the literature are mentioned in Section 5.

Organization

In Section 2 we present our general framework. In Section 3 we formally state and prove our results for the general framework. Section 4 contains our results for the application of welfare maximization. In Section 5 we list known rounding techniques from the literature and how they fit into the framework.

2 Framework

In this section we present our framework. After several general definitions, in Section 2.1 we define an optimization problem and its relaxation, and recall the well-known notion of integrality gap – a measure of how "relevant" optimization of the relaxation is to optimization of the original problem. In Section 2.2 we introduce oblivious rounding and define the approximation ratio of such rounding schemes, according to how well they round a solution to the relaxed problem into a feasible solution of the original problem.

Let $d \in \mathbb{N}_{>0}$ be a positive integer. For every set $S \subseteq \mathbb{R}^d$ of d-dimensional vectors, let $C(S)$ denote its *convex hull*, i.e., $C(S) = \{\sum_{s\in S} \lambda_s s \mid \forall s \in S : \lambda_s \geq 0 \text{ and } \sum_{s\in S} \lambda_s = 1\}$. A set S is *compact* if it is *closed* (no infinite sequence of vectors converges to a vector outside the set), and *bounded* (there is some finite μ such that the norm of every vector in the set is at most μ). If S is convex and compact, let $\partial(S)$ denote its outer *boundary*, i.e., $\partial(S) = \{s \in S \mid \forall \text{ scalar } \delta \in \mathbb{R}, \delta > 1 : \delta s \notin S\}$.

For sets $S_1, S_2 \subseteq \mathbb{R}^d$, we use the notation $\min_{s_1\in S_1} \max_{s_2\in S_2}\{\cdot\}$ when we are optimizing by first choosing $s_1 \in S_1$, and then choosing $s_2 \in S_2$ based on knowledge of s_1; similarly, the notation $\max_{s_2\in S_2} \min_{s_1\in S_1}\{\cdot\}$ means that $s_2 \in S_2$ is chosen first and $s_1 \in S_1$ is chosen with prior knowledge of s_2. Here, min and max can be replaces by inf and sup where needed.

2.1 Problems, Relaxations, Closures

We consider optimization problems with linear objectives. We define a problem of dimension d as a collection of d-dimensional instances coupled with a feasible solution set. This means that in our formulation, problem instances of a certain dimension all share the same set of feasible solutions.

For concreteness our framework is developed for *maximization* problems (the results can be adapted also to minimization).

▶ **Definition 1.** A *problem* π of dimension d is a pair (V, X), where $V, X \subseteq \mathbb{R}^d_{\geq 0}$ are nonempty sets of d-dimensional vectors with non-negative entries. V contains the problem *instances* (also called *value functions* or *objectives*), and X is the set of feasible *solutions*. Given an instance $v \in V$, the *value* of solution $x \in X$ is the inner product $v \cdot x$, and x is *optimal* if it has maximum value among all feasible solutions.[2]

[2] The non-negativity in this definition of vectors in V, X can be replaced by a weaker condition of $v \cdot x \geq 0$ for every $v \in V, x \in X$, and our results will still hold.

For concreteness, recall the max-cut example mentioned in Section 1: The instances are modeled as weighted complete graphs over n nodes, all of which share the same set of possible cuts. An instance is thus simply a vector of $n(n-1)/2$ non-negative edge weights, and a feasible solution is a $\{0,1\}$-vector indicating the edges that participate in a cut.

We now define a problem relaxation, which is itself a problem achieved by expanding the original set of feasible solutions:

▶ **Definition 2.** A problem $\pi' = (V, Y)$ is a *relaxation* of problem $\pi = (V, X)$ if $X \subseteq Y$. The solutions in Y are referred to as *relaxed* solutions.

For every relaxed solution $y \in Y$, V_y^+ denotes all instances for which the value of y is strictly positive, and V_y^* denotes all instances for which y is optimal:

$$V_y^+ = \{v \in V \mid v \cdot y > 0\}; \quad V_y^* = \{v \in V \mid v \cdot y \geq v \cdot y' \ \forall y' \in Y\}. \tag{1}$$

Finally, we introduce the closure of a problem, achieved by convexifying the set of instances:

▶ **Definition 3.** The problem $\mathrm{cl}(\pi) = (C(V), X)$ is the *closure* of problem $\pi = (V, X)$.

2.1.1 Assumed Properties of Problems and Relaxations

All problems and relaxations we consider in this paper are assumed to have the natural properties of *compactness* and *positivity* unless stated otherwise, and all relaxations are assumed to be *convex*:

- A problem $\pi = (V, X)$ is *compact* if the feasible solution set X is compact, and there is a compact set $V' \subseteq \mathbb{R}^d_{\geq 0} \setminus \{0^d\}$ such that the instance set V is $\{v = cv' \mid c \in \mathbb{R}_{>0}$ and $v' \in V'\}$. Without loss of generality, the vectors in V' can also be assumed to be normalized (i.e., $\sum_k v'_k = 1$). This is a weaker assumption than assuming V is compact, since it allows unbounded instances as well as instances that approach, but do not reach, 0^d. Many common optimization problems, for example max-cut, are compact: Indeed, the solution set (cuts) is usually closed and bounded; the value functions that make up the instances (edge weights) usually exclude the zero function $v = 0^d$, and so can be normalized as above without loss of generality (without affecting multiplicative approximation factors). Thus V' can be taken to be the set of normalized instances, which is bounded and closed.[3]
- A problem $\pi = (V, X)$ is *positive* if for every $v \in V$ there is some $x \in X$ such that $v \cdot x > 0$ (in particular, V is not allowed to include 0^d), and for every $x \in X \setminus \{0^d\}$ there is some $v \in V$ such that $v \cdot x > 0$. In the max-cut example, the first positivity condition holds because $v \neq 0^d$ and so at least one edge must have nonzero weight. For the second positivity condition, a natural sufficient condition is that the graph has a spanning tree such that for every edge in the tree, there is an edge-weight function in V that assigns positive weight to that edge. For every cut x there is at least one edge of the spanning tree in the cut, and therefore at least one instance v such that $v \cdot x > 0$. Notice that by the positivity assumption applied to a relaxation π', V_y^+ is nonempty for every $y \in Y \setminus \{0^d\}$, ensuring that our definitions (such as Definition 4 below) are well-defined.
- A relaxation $\pi' = (V, Y)$ to problem $\pi = (V, X)$ is *convex* if the set Y of relaxed solutions is convex. For example, relaxations that result from formulating the problem as an integer

[3] There is also a version of our results that holds when V' is not closed, in which sup and inf replace max and min in the appropriate places.

program and relaxing it to a linear program are convex. If π' is convex then in particular Y includes the convex hull $C(X)$.

Observe that if a problem is compact and positive, then its closure is also compact and positive.

2.1.2 Integrality Gap

Given a problem $\pi = (V, X)$ and a relaxation $\pi' = (V, Y)$, an important measure of the quality of the relaxation is the integrality gap – the worst case (smallest) ratio, over all possible instances in V, between the value achievable for the instance by a feasible solution in X, and the value achievable for it by a relaxed solution in Y. Formally:

▶ **Definition 4.** Let $\pi = (V, X)$ and $\pi' = (V, Y)$ be a problem and its relaxation. For every relaxed solution $y \in Y \setminus \{0^d\}$ and instance $v \in V_y^+$, the *integrality gap at v, y* is

$$\rho_{\pi,\pi'}(v, y) = \max_{x \in X} \frac{v \cdot x}{v \cdot y}.$$

The integrality gap at solution y is then obtained by taking the worst case instance v, i.e., $\rho_{\pi,\pi'}(y) = \inf_{v \in V_y^+} \rho_{\pi,\pi'}(v, y)$. Similarly, the integrality gap at instance v is $\rho_{\pi,\pi'}(v) = \inf_{y : v \in V_y^+} \rho_{\pi,\pi'}(v, y)$. The (overall) integrality gap is $\rho_{\pi,\pi'} = \inf_{y \in Y \setminus \{0^d\}} \rho_{\pi,\pi'}(y)$.

We make several basic observations regarding the integrality gap. Short proofs appear for completeness in Appendix A.

▶ **Observation 5.** *The integrality gap $\rho_{\pi,\pi'}$ is ≤ 1.*

Informally, the closer $\rho_{\pi,\pi'}$ is to 1, the better the relaxation.

Taking the closure of a problem expands the instance set and so makes it "harder" to get a good relaxation:

▶ **Observation 6.** *For every π and relaxation π', $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')} \leq \rho_{\pi,\pi'}$.*

The next observation shows that to find the integrality gap, we may restrict attention to relaxed solutions that lie on the boundary. Recall that Y is compact, then:

▶ **Observation 7.** *The overall integrality gap is not affected by the integrality gaps at relaxed solutions that lie strictly within the boundary: $\rho_{\pi,\pi'} = \min_{y \in \partial(Y)} \rho_{\pi,\pi'}(y)$.*

2.2 Oblivious Rounding

For the definitions in this section, fix a problem $\pi = (V, X)$ and a relaxation $\pi' = (V, Y)$.

A (randomized) rounding scheme receives an instance $v \in V$ and a relaxed solution $y \in Y$, and returns a distribution over feasible solutions in X. Note that since our objective functions in V are linear, any distribution over feasible solutions in X can be summarized by its average, which lies in the convex hull $C(X)$. This leads to the following definition:

▶ **Definition 8.** A *rounding scheme* is a function $f : V \times Y \to C(X)$.

A rounding scheme is oblivious if it is not allowed to "see" the objective function when rounding a solution of the relaxed problem:

▶ **Definition 9.** An *oblivious* rounding scheme is a function $f : Y \to C(X)$.

▶ Remark. The rounding schemes we consider, whether oblivious or not, need not be computable in polynomial time.

2.2.1 Approximation Ratio of Oblivious Rounding

Our goal is to study the power of oblivious rounding schemes for approximation. For this we shall use the following definition – the approximation ratio of an oblivious rounding scheme is the worst case ratio, over all possible instances in V, between the value achieved for the instance by a rounded solution in X, and the value achievable for it by the corresponding relaxed solution in Y. Formally:

▶ **Definition 10.** Consider an oblivious rounding scheme $f : Y \to C(X)$. For every relaxed solution $y \in Y \setminus \{0^d\}$, the *approximation ratio of f at y* is

$$\alpha_{\pi,\pi'}(y) = \inf_{v \in V_y^+} \frac{v \cdot f(y)}{v \cdot y}.$$

The approximation ratio of f is $\alpha_{\pi,\pi'} = \inf_{y \in Y \setminus \{0^d\}} \alpha_{\pi,\pi'}(y)$.

A larger approximation ratio indicates better approximation by the rounding scheme. A basic observation regarding the approximation ratio is that it is upper-bounded by the integrality gap. A short proof appears in Appendix A for completeness.

▶ **Observation 11.** *For every $y \in Y \setminus \{0^d\}$, the approximation ratio of f at y is at most the integrality gap at y: $\alpha_{\pi,\pi'}(y) \le \rho_{\pi,\pi'}(y)$. Therefore $\alpha_{\pi,\pi'} \le \rho_{\pi,\pi'} \le 1$.*

Observation 11 upper-bounds the approximation ratio, and a natural class of interest is rounding schemes for which this bound is tight:

▶ **Definition 12.** An oblivious rounding scheme $f : Y \to C(X)$ is *tight* if $\alpha_{\pi,\pi'} = \rho_{\pi,\pi'}$, and *individually tight* if $\alpha_{\pi,\pi'}(y) = \rho_{\pi,\pi'}(y)$ for every relaxed solution $y \in Y \setminus \{0^d\}$.

By definition, individual tightness implies tightness.

2.2.2 Approximation Ratio for Optimal Solutions

We are also interested in the approximation guarantees of oblivious rounding schemes *only* for relaxed solutions $y \in Y$ which have the following promised property: they are known to be optimal solutions to some instance of the relaxed problem. Recall from (1) that V_y^* denotes the set of all instances for which y is an optimal solution.

▶ **Observation 13.** *If V_y^* is nonempty then $y \in \partial(Y)$.*

See Appendix A for a proof.

The two definitions in this subsection are analogous to Definitions 10 (approximation ratio) and 12 (tightness) above:

▶ **Definition 14.** Consider an oblivious rounding scheme $f : Y \to C(X)$. For every relaxed solution $y \in Y$ for which $V_y^* \neq \emptyset$, the approximation ratio *for optimal solutions* of f at y is

$$\alpha^*_{\pi,\pi'}(y) = \inf_{v \in V_y^*} \frac{v \cdot f(y)}{v \cdot y}.$$

The approximation ratio for optimal solutions of f is $\alpha^*_{\pi,\pi'} = \inf_{y \in Y : V_y^* \neq \emptyset} \{\alpha^*_{\pi,\pi'}(y)\}$.

By definition, for every $y \in Y$ with nonempty V_y^* it holds that $\alpha_{\pi,\pi'}(y) \le \alpha^*_{\pi,\pi'}(y)$, and so $\alpha_{\pi,\pi'} \le \alpha^*_{\pi,\pi'}$. Note that this inequality may be strict in some cases, and moreover it is not necessarily the case that the upper bound $\rho_{\pi,\pi'}$ on $\alpha_{\pi,\pi'}$ is also an upper bound on $\alpha^*_{\pi,\pi'}$ (see Example 35 below). This motivates the next definition:

▶ **Definition 15.** An oblivious rounding scheme $f : Y \to C(X)$ is tight *for optimal solutions* if $\alpha^*_{\pi,\pi'} \geq \rho_{\pi,\pi'}$, and individually tight *for optimal solutions* if $\alpha^*_{\pi,\pi'}(y) \geq \rho_{\pi,\pi'}(y)$ for every relaxed solution $y \in Y$ with nonempty V^*_y.

By definition, individual tightness for optimal solutions implies tightness for optimal solutions.

3 General Results

In this section we state our results for the general framework; some proofs are deferred to Appendix B.1. Appendix B.2 discusses implications for oblivious rounding of optimal solutions. Additional results that concern the applications of the framework to welfare maximization appear in Section 4.

Recall that the closure of problem $\pi = (V, X)$ is $\mathrm{cl}(\pi) = (C(V), X)$ (Definition 3). Our main general theorem relates the (pointwise) approximation ratio of oblivious rounding to the integrality gap of the problem's closure:

▶ **Theorem 16.** *Given a problem $\pi = (V, X)$ and a relaxation $\pi' = (V, Y)$:*
1. *Upper bound: For every oblivious rounding scheme $f : Y \to C(X)$, at every point $y \in Y \setminus \{0^d\}$, the approximation ratio $\alpha_{\pi,\pi'}(y)$ is at most the integrality gap $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y)$ of the closure of problem π.*
2. *Tightness: There exists an oblivious rounding scheme $f : Y \to C(X)$ such that $\alpha_{\pi,\pi'}(y) = \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y)$ for every $y \in Y \setminus \{0^d\}$.*

Moreover, our proof method yields the following proposition, by which the approximation ratio and integrality gap are achieved by the same instance and (random) feasible solution:

▶ **Proposition 17.** *Given a problem $\pi = (V, X)$, a relaxation $\pi' = (V, Y)$ and a relaxed solution $y \in Y \setminus \{0^d\}$, there exist an instance $v \in C(V)$ and a random feasible solution $x \in C(X)$ of the problem $\mathrm{cl}(\pi)$ such that $\frac{v \cdot x}{v \cdot y} = \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y) = \alpha_{\pi,\pi'}(y)$, where the approximation ratio is that of the best oblivious rounding scheme at y.*

Proof. By Lemma 33. ◀

Two useful corollaries follow immediately from Theorem 16. First, we have already observed that $\alpha_{\pi,\pi'} \leq \rho_{\pi,\pi'}$ (Observation 11) and that $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')} \leq \rho_{\pi,\pi'}$ (Observation 6). It follows from Theorem 16 that for the best oblivious rounding scheme in fact $\alpha_{\pi,\pi'} = \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}$.

▶ **Corollary 18.** *Given a problem $\pi = (V, X)$ and a relaxation $\pi' = (V, Y)$, there exists an oblivious rounding scheme $f : Y \to C(X)$ that achieves an approximation ratio of $\alpha_{\pi,\pi'} = \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}$, and this is the best possible approximation ratio of any oblivious rounding scheme.*

Proof. By Definitions 4 (integrality gap) and 10 (approximation ratio), if for an oblivious rounding scheme f it holds that $\alpha_{\pi,\pi'}(y) = \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y)$ for every $y \in Y \setminus \{0^d\}$, then $\alpha_{\pi,\pi'} = \inf_{y \in Y \setminus \{0^d\}} \alpha_{\pi,\pi'}(y) = \inf_{y \in Y \setminus \{0^d\}} \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')} = \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}$. By Theorem 16 there exists such an oblivious rounding scheme. ◀

▶ **Corollary 19.** *Given a problem $\pi = (V, X)$ whose instances form a convex set (i.e., $\pi = \mathrm{cl}(\pi)$), for every relaxation $\pi' = (V, Y)$, there exists an oblivious rounding scheme $f : Y \to C(X)$ that is individually tight.*

Proof. By Theorem 16 there exists an oblivious rounding scheme f such that $\alpha_{\pi,\pi'}(y) = \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y)$ for every $y \in Y \setminus \{0^d\}$, and by assumption, $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y) = \rho_{\pi,\pi'}(y)$. The proof follows from the definition of individual tightness (Definition 12). ◀

Unlike the statement in Corollary 18, Example 35 shows that the approximation ratio *for optimal solutions* $\alpha^*_{\pi,\pi'}$ may surpass the integrality gap of the closure $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}$. See Appendix B.2 for more details on oblivious rounding of optimal solutions.

3.1 Proof of Theorem 16 via Minimax

Our goal in this section is to prove Theorem 16 via our main lemma (Lemma 33), which is a version of von Neumann's minimax theorem. In the proof we shall use the classic minimax theorem for *non-finite* zero-sum games:

▶ **Theorem 20** ([26]). *For every bipartite zero-sum game in which the players' pure strategy sets X and V are compact and the payoff function $g : \mathcal{V} \times \mathcal{X} \to \mathbb{R}$ is continuous, there exists a unique minimax value μ^* such that*

$$\mu^* = \max_{x \in C(X)} \min_{v \in V} g(v, x) = \min_{v \in C(V)} \max_{x \in X} g(v, x). \tag{2}$$

Moreover, there are equilibrium strategies *$x^* \in C(X)$ and $v^* \in C(V)$ such that x^* maximizes $g(v^*, x)$, v^* minimizes $g(v, x^*)$, and $\mu^* = g(v^*, x^*)$.*

▶ Remark. Throughout this section we shall assume that every problem $\pi = (V, X)$ has a compact instance set V in which instances are normalized (i.e., $\sum_k v_k = 1$). This assumption is without loss of generality, as we assumed in Section 2.1.1 that $V = \{v = cv' \mid c \in \mathbb{R}_{>0} \text{ and } v' \in V'\}$ where V' is compact and normalized. Since an instance $v \in V$ appears exclusively within the expressions $\frac{v \cdot x}{v \cdot y}$ or $\frac{v \cdot f(y)}{v \cdot y}$, the multiplying constant c cancels out and we may as well assume that $V = V'$.

We begin with an intuitive (albeit imprecise) explanation of the connection between the minimax theorem and the approximation ratio of an oblivious rounding scheme. Fix a problem $\pi = (V, X)$, a relaxation $\pi' = (V, Y)$ and a relaxed solution y. We claim that an oblivious rounding scheme f, which maximizes the approximation ratio $\alpha_{\pi,\pi'}(y)$ at y, is equivalent to an optimal mixed strategy in the following zero-sum game (the games used in the actual proof are slightly different): Given y, the maximizing "rounding" player picks a mixed strategy $f(y) \in C(X)$ over feasible solutions in X, and the minimizing "instance" player picks an instance $v \in V$ as his pure strategy best-response to $f(y)$. The expected payoff of the rounding player is the ratio $\frac{v \cdot f(y)}{v \cdot y}$. By the minimax theorem (Theorem 20), the resulting zero-sum game has a minimax value achieved by the optimal mixed strategy $f(y)$ and the worst case v for $f(y)$. This value is thus precisely equal to the approximation ratio $\alpha_{\pi,\pi'}(y)$ of the optimal oblivious rounding scheme at y (recall Definition 10). Note that we require the rounding to be oblivious, hence the rounding player does not know the strategy v of the instance player when choosing her mixed strategy $f(y)$ given y.

Again by Theorem 20, the minimax value of the game $\alpha_{\pi,\pi'}(y)$ is alternatively achieved by first letting the instance player pick an optimal mixed strategy (a distribution $v \in C(V)$ over instances), and then allowing the rounding player to pick a best-response feasible solution $x \in X$. Notice that a mixed strategy v of the instance player is an instance of the *closure* $\mathrm{cl}(\pi)$ of the original problem π. Given y and v, the feasible solution x that maximizes the rounding player's expected payoff $\frac{v \cdot x}{v \cdot y}$ is precisely the same x that achieves the integrality gap $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(v, y)$ in Definition 4. Since the instance player is playing an optimal mixed strategy, we get that the value of the game $\alpha_{\pi,\pi'}(y)$ is equal to $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y)$. We conclude that the best approximation ratio at y and the integrality gap at y with respect to the closure coincide.

Given the above paragraphs, it may seem that the proof of Theorem 16 should follow directly by invoking Theorem 20. However, the classic minimax theorem is not immediately applicable in our setting due to a technical difficulty: While we can set the payoff function g in (2) to be $\frac{v \cdot x}{v \cdot y}$, the approximation ratio and integrality gap notions are defined with $\inf_{v \in V_y^+}$ instead of $\min_{v \in V}$ (to avoid division by zero). And while X and V_y^+ are bounded and X is also closed, V_y^+ may not be closed, and therefore may not be compact (unlike V). Lemma 33 and its proof show how to circumvent this problem by defining an appropriate series of zero-sum games. The lemma, its proof and the proof of Theorem 16 are deferred to Appendix B.1.

4 Application: Welfare Maximization

In this section we demonstrate our framework and results by applying them to the optimization problem of welfare maximization in combinatorial auctions. In Section 4.1 we state some preliminaries regarding the problem. In Section 4.2 we show a fairness property of oblivious rounding for welfare maximization. In Section 4.3 we bound the approximation ratio of oblivious rounding schemes for welfare maximization with unit demand valuations. In Section 4.4 we use the particular structure of the welfare maximization problem to extend our impossibility results to rounding of solutions that are guaranteed to be optimal (this is in contrast to the general case, see, e.g., Example 35). In Section 4.5 we give an explicit example of an instance that manages to "fool" oblivious rounding attempts.

4.1 Auction Preliminaries

A combinatorial auction involves a set $N = [n]$ of players and a set $M = [m]$ of indivisible items. Each player i has a valuation ν_i, which is a function $\nu_i : 2^M \to \mathbb{R}_{\geq 0}$ that assigns a real value to every subset of items $S \subseteq M$ (also called a *bundle*). Valuations are routinely assumed to be monotone (for every two bundles $S \subseteq T$, $\nu(S) \leq \nu(T)$), and bounded (assigning values up to some maximum value μ). An *allocation* $(S_1, \ldots, S_n)$ of the items is a (partial) partition of M into n bundles of items, one per player (some bundles may be empty). The *welfare* of a given allocation is the sum of the players' values for their allocated bundle, i.e., $\sum_{i=1}^{n} \nu_i(S_i)$. The goal of the welfare maximization problem is to find an allocation of the items that maximizes the welfare.

In the terminology of our framework, an instance of the welfare maximization problem is a vector v of dimension $n \cdot 2^m$ (indexed by pairs (i, S) of player and bundle) containing all the players' values for all the bundles, that is, $v_{i,S} = \nu_i(S)$. A feasible solution is a $\{0,1\}$-vector x of the same length, $n \cdot 2^m$, that indicates which player receives which bundle (up to one bundle per player), and does not over-allocate the items. Formally, $x_{i,S} \in \{0,1\}$, for every player i, $\sum_S x_{i,S} \leq 1$, and for every item j, $\sum_{i,S:j \in S} x_{i,S} \leq 1$.

The welfare maximization problem can be formulated as an integer program, and its standard relaxation is the associated linear program, called the *configuration LP* (see Appendix C.1). A relaxed solution is a vector y with $[0,1]$-entries, which can be thought of as an allocation of *fractional* rather than indivisible items, via an allocation of fractions of bundles. It must still hold that at most one of each item is allocated ($\sum_{i,S:j \in S} y_{i,S} \leq 1$ for every item j), and that each player receives at most one bundle ($\sum_S y_{i,S} \leq 1$ for every player i). In other words, a relaxed solution is any (fractional) feasible solution to the configuration LP.

A class of welfare maximization problems that has been extensively studied in the literature is welfare maximization with *gross substitutes* valuations. Such valuations play a crucial role in microeconomics [14] and in discrete convex optimization [19]; for a recent algorithmic survey see [21]. There are many equivalent definitions of gross substitutes valuations, one of which we give for completeness in Appendix C.1.

An important property of gross substitutes valuations is that the integrality gap of the configuration LP is 1.

▶ **Proposition 21** ([2]). *The integrality gap of the configuration LP for gross substitutes valuations is 1.*

Moreover, if all valuations are gross substitutes, then the welfare maximization problem can be solved optimally in polynomial time [17, 18].

A subclass of gross substitutes valuations is the class of *unit-demand* valuations. A valuation ν is unit-demand if there exists a vector $(\nu_1, \ldots, \nu_m) \in \mathbb{R}^m_{\geq 0}$ such that for every bundle S, $\nu(S) = \max_{j \in S}(\nu_j)$.

Also relevant to our study is the class of *coverage* valuations. A valuation ν is a coverage function if it can be described by a tuple $\nu = \langle E, w, \{E_j\}_j \rangle$, where: (1) E is a ground set of elements, (2) $w : E \to \mathbb{R}_{\geq 0}$ is a weight function that assigns a weight $w(e)$ for every element $e \in E$, and (3) for every item $j \in [m]$, $E_j \subseteq E$ is the subset of elements *covered* by item j; and for every bundle of items $S \subseteq M$, it holds that $\nu(S) = \sum_{e \in \bigcup_{j \in S} E_j} w(e)$. The class of coverage valuations is a strict superset of unit-demand valuations (and is incomparable with gross substitutes). Coverage valuations are well-studied, with a particular surge in attention in the context of social networks (see, e.g., [1, 5]).

The convex hull of the class of unit-demand valuations is strictly larger than the class itself. In particular, the following lemma asserts that the convex hull of unit-demand valuations is precisely the class of coverage valuations (see Appendix C.2 for a proof).

▶ **Lemma 22.** *The class of coverage valuations is the convex hull of unit-demand valuations.*

4.2 A Fairness Property

In the context of welfare maximization, oblivious rounding with good approximation guarantees also offers certain guarantees *per player*. The intuition is that a rounding scheme that is ignorant to the instance has no way of telling which player contributes what to the welfare, and so must approximately preserve the welfare contributions of *all* players from behind its veil of ignorance. This can be viewed as a fairness property of oblivious rounding.

▶ **Proposition 23.** *Consider an oblivious rounding scheme f for the welfare maximization problem and its configuration LP relaxation, which has approximation ratio α. Then for every instance v and fractional allocation y, $f(y)$ guarantees for each player i, in expectation, an α-fraction of the player's value $\sum_S v_{i,S} y_{i,S}$ in y.*

Proof. Assume for contradiction that there is a player i for which this is not the case. Then we can create a new instance v' in which only player i's valuation is non-zero, meaning that all welfare comes from this player (note that while we do not allow an all zero valuation, assigning zero valuations to all players other than player i is valid). Since f is oblivious, it should achieve the approximation ratio α for v', contradiction. ◀

4.3 Impossibility Results

In this section we prove two impossibility results on the approximation ratios of oblivious rounding schemes for unit-demand valuations. These bounds extend to gross substitutes valuations.

▶ **Proposition 24.** *The approximation ratio of any oblivious rounding scheme for welfare maximization with two unit-demand players and the configuration LP relaxation is at most* $5/6$.

▶ **Proposition 25.** *The approximation ratio of any oblivious rounding scheme for welfare maximization with* n *unit-demand players and the configuration LP relaxation is at most* ≈ 0.782.

These impossibility results are in stark contrast to Proposition 21. In particular, while the integrality gap of the configuration LP is 1 even for a strict superclass of unit demand (i.e., gross substitutes), oblivious rounding for unit-demand valuations is quite limited in its performance.

Proof of Proposition 24. By Corollary 18 and Lemma 22, it is sufficient to show an instance with two coverage valuations that has an integrality gap of $5/6$. We claim that the instance in [10] for two players with submodular valuations satisfies these conditions. Let us describe the example explicitly using our notation, and showing in the process that the players' valuations are coverage functions.

There are four items and two players. For reasons that will become apparent shortly, it will be convenient to name the items $a_{11}, a_{12}, a_{21}, a_{22}$. There are six elements $\{H_1, H_2, V_1, V_2, D_1, D_2\}$. In both valuation functions, the coverage of elements by items is identical, but they differ in the weights of the different elements. We first state the coverage structure. For every element e, we denote the set of items that cover element e by $\bar{e}$. Let $\bar{H_1} = \{a_{11}, a_{12}\}$, $\bar{H_2} = \{a_{21}, a_{22}\}$, $\bar{V_1} = \{a_{11}, a_{21}\}$, Similarly, let $\bar{V_2} = \{a_{12}, a_{22}\}$, $\bar{D_1} = \{a_{11}, a_{22}\}$, and $\bar{D_2} = \{a_{12}, a_{21}\}$.

We now state the weights of the elements according to ν_1 and ν_2. Let $w^i(e)$ denote the weight of element e according to ν_i. For player 1, $w^1(H_1) = w^1(H_2) = 0$, $w^1(V_1) = w^1(V_2) = 2$, and $w^1(D_1) = w^1(D_2) = 1$. For player 2, $w^2(H_1) = w^2(H_2) = 2$, $w^2(V_1) = w^1(V_2) = 0$, and $w^2(D_1) = w^2(D_2) = 1$.

For example, $\nu_1(\{a_{11}, a_{12}\}) = w^1(H_1) + w^1(V_1) + w^1(V_2) + w^1(D_1) + w^1(D_2) = 6$, and $\nu_2(\{a_{11}, a_{12}\}) = w^2(H_1) + w^2(V_1) + w^2(V_2) + w^2(D_1) + w^2(D_2) = 4$.

One may verify that the following fractional solution has welfare 12: player 1 receives a fraction $1/2$ of bundle $\bar{H_1}$ and a fraction $1/2$ of bundle $\bar{H_2}$. This gives player 1 value 6. Player 2 receives a fraction $1/2$ of bundle $\bar{V_1}$ and a fraction $1/2$ of bundle $\bar{V_2}$. This gives player 2 value 6. It can be verified that no integer assignment of items gives total welfare above 10, establishing that the integrality gap is no better than $5/6$. This establishes the assertion of the proposition. ◀

Proof of Proposition 25. By Corollary 18 and Lemma 22, it is sufficient to show an instance with n coverage valuations and an integrality gap of ≈ 0.782. We claim that the instance in [10] for n players with submodular valuations satisfies these conditions.

Let us recall the instance. There are n players and n^n items arranged in an n dimensional cube. A *line* in direction i is a set of n points whose projection on the ith coordinate gives all values from 0 to $n-1$. There are $n^{n(n-1)}$ lines in direction i. The valuation function ν_i is defined such that $\nu_i(S)$ equals the fraction of lines in direction i hit by set S. One can verify that the valuation of player i is the following coverage valuation: Associate an element with every line in direction i, and let each item cover the elements corresponding to lines that contain it. The weight of every element is $1/n^{n(n-1)}$. As shown in [10], the integrality gap of this instance is ≈ 0.782. This establishes the assertion of the proposition. ◀

4.4 Impossibility Results for Optimal Solutions

We now define a strong notion of per-player guarantee. Consider an instance of the welfare maximization problem. A relaxed solution $y = \{y_{i,S}\}$ (a fractional solution of the configuration LP) is said to be *individually optimal* for this instance if the fractional value of every player in the solution x is his maximum possible value. Assuming that all valuations are monotone, this means that $\sum_S y_{i,S}\nu_i(S) = \nu_i(M)$ for every player i.

The significance of individual optimality lies in the following lemma. Consider a class of valuations P. Let v be an instance with valuations $\nu_1, \dots, \nu_n \in C(P)$, and let $\rho(v)$ denote its integrality gap (as defined in Definition 4; we omit here the problem and relaxation from the notation). Let $y = \{y_{i,S}\}$ be an *optimal* relaxed solution for instance v, i.e., an optimal fractional solution to the configuration LP, whose welfare (LP objective value) we denote by $\mathrm{LP}(v, y)$.

▶ **Lemma 26.** *If y is individually optimal for v, then the approximation ratio for optimal solutions of any oblivious rounding scheme at y is at most $\rho(v)$.*

Proof. By definition, for every i, there exist valuations $\nu_{ik} \in P$ such that $\nu_i = \sum_k \lambda_{ik}\nu_{ik}$. For valuation ν_i and a fractional solution y, let $\nu_i(y) = \sum_S y_{i,S}\nu_i(S)$. Let $\alpha^*(y)$ denote the approximation ratio for optimal solutions of an oblivious rounding of y with respect to P, and let f be the oblivious rounding scheme achieving $\alpha^*(y)$. Consider random instances with valuations in P, where in every instance player i has valuation ν_{ik} with probability λ_{ik} (independently). For every random instance, the expected welfare obtained by f is at least $\alpha^*(y) \cdot \sum_i \nu_{ik}(y)$ (Proposition 23).

We claim that the individual optimality of y implies that y is also individually optimal for every random instance (i.e., $\nu_{ik}(y) = \nu_{ik}(M)$ for every i, k). Suppose otherwise, i.e., suppose there exist i, k such that $\nu_{ik}(y) < \nu_{ik}(M)$. Then, for that player i it follows (by monotonicity of the valuation) that $\sum_k \lambda_{ik}\nu_{ik}(y) < \sum_k \lambda_{ik}\nu_{ij}(M)$. On the other hand, $\sum_k \lambda_{ik}\nu_{ik}(y) = \nu_i(y) = \nu_i(M)$, so we get $\nu_i(M) > \sum_k \lambda_{ik}\nu_{ik}(M)$, contradiction.

Substituting $\nu_{ik}(y) = \nu_{ij}(M)$, and taking a weighted average over all instances, we get that the expected value obtained by f is at least $\alpha^*(y)\sum_{ik} \lambda_{ik}\nu_{ik}(M) = \alpha^*(y)\sum_i \nu_i(M) = \alpha^*(y)\,\mathrm{LP}(v, y)$. Now observe that f obtains the same ratio $\alpha^*(y)$ on the original instance v; therefore, $\alpha(y) \leq \rho(v)$ (otherwise, it contradicts the integrality gap of v). ◀

The following proposition follows directly from Lemma 26.

▶ **Proposition 27.** *Consider the problem of welfare maximization with valuations from P and its configuration LP relaxation. Let v be an instance attaining the integrality gap for $C(P)$, and let $y = \{y_{i,S}\}$ be an optimal solution of the configuration LP for instance v. If y is individually optimal, then the approximation ratio for optimal solutions α^* of any oblivious rounding scheme is at most the integrality gap of $C(P)$.*

Proof. Follows directly from Lemma 26, and from the definition of the approximation ratio for optimal solutions. Recall that this ratio is the infimum over the approximation ratios for optimal solutions of all $y \in Y$ for which V^*_y is nonempty (Definition 14). ◀

▶ **Corollary 28.** *The impossibility results in Propositions 24 and 25 apply also to the approximation ratio of oblivious rounding of optimal solutions of the configuration LP.*

Proof. The proof is by applying Proposition 27 and verifying that the instances in the proofs of Propositions 24 and 25 admit an individually optimal fractional solution. In the instance used in the proof of Proposition 24, the optimal fractional solution is individually optimal since this solution gives agent 1 a fraction 1/2 of each of $\bar{H}_1, \bar{H}_2$, and agent 2 a fraction 1/2 of each of $\bar{V}_1, \bar{V}_2$. In the instance used in the proof of Proposition 25, the optimal fractional solution is individually optimal since the solution gives a player i the n level sets with respect to coordinate i, each with probability $1/n$. ◀

4.5 How to Fool Oblivious Rounding

To gain intuition as to why oblivious rounding fails to round optimally, we now describe a two-player instance related to the instance in the proof of Proposition 24, and show why no oblivious rounding can succeed in rounding it with an approximation ratio better than 5/6. The instance is simple, including two players with unit-demand valuations and $\{0,1\}$ values.

▶ **Example 29.** There are four items and two players. The items are $a_{11}, a_{12}, a_{21}, a_{22}$. Recall that a unit-demand function ν_i can be expressed by $\{\nu_{ij}\}_{j\in M}$, where $\nu(S) = \max_{j\in S}\nu_{ij}$. In our example, $\nu_{ij} \in \{0,1\}$ for every i, j. We adopt the following notation used in the proof of Proposition 24: $\bar{H}_1 = \{a_{11}, a_{12}\}$, $\bar{H}_2 = \{a_{21}, a_{22}\}$, $\bar{V}_1 = \{a_{11}, a_{21}\}$, $\bar{V}_2 = \{a_{12}, a_{22}\}$, $\bar{D}_1 = \{a_{11}, a_{22}\}$, and $\bar{D}_2 = \{a_{12}, a_{21}\}$. We denote by S^i the items j such that $\nu_{ij} = 1$. The valuation functions are as follows: S^1 is $\bar{V}_1$ or $\bar{V}_2$, each with probability 1/3, and is $\bar{D}_1$ or $\bar{D}_2$, each with probability 1/6. S^2 is $\bar{H}_1$ or $\bar{H}_2$, each with probability 1/3, and is $\bar{D}_1$ or $\bar{D}_2$, each with probability 1/6.

Observe that for every realization of the valuations there exists an integral solution with social welfare 2. In addition, for every realization it holds that $\nu_1(\bar{H}_1) = \nu_1(\bar{H}_2) = 1$ and $\nu_2(\bar{V}_1) = \nu_2(\bar{V}_2) = 1$. Therefore, a fractional solution that assigns a fraction 1/2 of each of $\bar{H}_1$ or $\bar{H}_2$ to player 1, and a fraction 1/2 of each of $\bar{V}_1$ or $\bar{V}_2$ to player 2, obtains optimal welfare of 2.

We next show that for every integral solution the expected social welfare is at most 5/3. Assigning $\bar{H}_1$ to player 1 and $\bar{H}_2$ to player 2, or vice versa, grants player 1 value 1 and player 2 an expected value of 2/3. An analogous argument holds for the assignment of $\bar{V}_1$ and $\bar{V}_2$; and the assignment of $\bar{D}_1$ and $\bar{D}_2$ grants every player an expected value of 5/6. Each of these assignments gives welfare 5/3. Finally, it is easy to see that assigning a single item to one player and a triplet to the other derives even less welfare (3/2).

We conclude that any oblivious rounding obtains welfare at most 5/3, which is 5/6 of the optimal solution.

5 Oblivious Rounding in the Literature

We list here several examples of rounding schemes which are oblivious, as well as schemes which are not oblivious. It is interesting to notice that for welfare maximization with budget additive valuations, which is not closed under convex combinations, the best known approximation is not oblivious, whereas for welfare maximization with submodular valuations, which is closed under convex combinations, the best-known approximation is oblivious. Additional examples appear in [7].

Examples of Oblivious Rounding

- Threshold rounding for vertex cover [13].
- Randomized rounding for set cover [22].
- Random hyperplane rounding for max cut [11].
- Welfare maximization for fractionally subadditive (XOS) and submodular valuations [8, 10].
- Randomized metarounding for congestion [3].

Examples of Non-Oblivious Rounding

- Rounding of semidefinite programs (SDPs) for the constraint satisfaction problem (CSP) [23].
- Welfare maximization with budget-additive valuations [25, 4].
- Facility location [16].

6 Conclusion and Open Questions

In this work we have systematically studied the notion of oblivious rounding and its approximation guarantees, with applications to the welfare maximization problem. We mention several directions for future research. First, are there optimization problems that are not closed under convex combinations, where the best known approximation is achieved by an oblivious rounding scheme, and can potentially be improved by considering non-oblivious rounding schemes? For which problems are there polynomial-time computable oblivious rounding schemes that are comparable to the integrality gap? Finally, what else can we hope to learn about the most promising rounding techniques from properties of the combinatorial problem?

References

1 Ashwinkumar Badanidiyuru, Shahar Dobzinski, Hu Fu, Robert Kleinberg, Noam Nisan, and Tim Roughgarden. Sketching valuation functions. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1025–1035, 2012.

2 Sushil Bikhchandani and John W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74(2):385–413, 1997.

3 Robert Carr and Santosh Vempala. Randomized metarounding. *Random Structures and Algorithms*, 20(3):343–352, 2002.

4 Deeparnab Chakrabarty and Gagan Goel. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and GAP. *SIAM J. Comput.*, 39(6):2189–2211, 2010.

5 Nan Du, Yingyu Liang, Maria-Florina Balcan, and Le Song. Learning time-varying coverage functions. In *Proceedings of the 27th Neural Information Processing Systems Conference*, pages 3374–3382, 2014.

6 Shaddin Dughmi, Tim Roughgarden, and Qiqi Yan. From convex optimization to randomized mechanisms: Toward optimal combinatorial auctions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, pages 149–158, 2011.

7 Paul Dütting, Thomas Kesselheim, and Éva Tardos. Algorithms as mechanisms: The price of anarchy of relax-and-round. In *Proceedings of the 16th ACM Conference on Economics and Computation*, pages 187–201, 2015.

8 Uriel Feige. On maximizing welfare when utility functions are subadditive. *SIAM J. Comput.*, 39(1):122–142, 2009.

9 Uriel Feige and Shlomo Jozeph. Oblivious algorithms for the maximum directed cut problem. *Algorithmica*, 71(2):409–428, 2015.

10 Uriel Feige and Jan Vondrák. The submodular welfare problem with demand queries. *Theory of Computing*, 6(1):247–290, 2010.

11 Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995.

12 Avinatan Hassidim and Yaron Singer. Submodular optimization under noise. Manuscript, 2016.

13 Dorit S. Hochbaum. Approximation algorithms for the set covering and vertex cover problems. *SIAM Journal on Computing*, 11:555—556, 1982.

14 Alexander S. Kelso and Vincent P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.

15 Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. *Journal of the ACM*, 58(6), 2011. Article 25.

16 Shi Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. *Information and Computation*, 222:45–58, 2013.

17 Kazuo Murota. Valuated matroid intersection I: Optimality criteria. *SIAM J. Discrete Math.*, 9(4):545–561, 1996.

18 Kazuo Murota. Valuated matroid intersection II: Algorithms. *SIAM J. Discrete Math.*, 9(4):562–576, 1996.

19 Kazuo Murota. *Discrete Convex Analysis.* Monographs on Discrete Mathematics and Applications. Society for Industrial and Applied Mathematics, 2003.

20 Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129:192–224, 2006.

21 Renato Paes Leme. Gross substitutability: An algorithmic survey. Working paper, 2014.

22 Prabhakar Raghavan and Clark D. Thompson. Randomized rounding: A technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7(4):365–374, 1987.

23 Prasad Raghavendra and David Steurer. How to round any CSP. In *Proceedings of the 50th Symposium on Foundations of Computer Science*, pages 586–594, 2009.

24 Tim Roughgarden, Inbal Talgam-Cohen, and Jan Vondràk. When are welfare guarantees robust? Working paper, 2016.

25 Aravind Srinivasan. Budgeted allocations in the full-information setting. In *Proceedings of the 11th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 247–253, 2008.

26 John von Neumann. Zur theorie der gesellschaftsspiele. *Math. Annalen.*, 100:295–320, 1928.

27 Neal E. Young. Randomized rounding without solving the linear program. In *Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 170–178, 1995.

A Appendix for Section 2

A.1 Missing Proofs

Proof of Observation 5. Since X is compact, there exist $x^* \in X$ and $v^* \in V$ such that $x^* = \arg\max_{x \in X} v^* \cdot x$ (and thus in particular $v^* \in V^+_{x^*}$). Set the relaxed solution y to be x^* (this is a valid choice since $x^* \in Y$). Then $\rho_{\pi,\pi'}(y) \leq 1$, since 1 can be achieved by choosing the instance v in the definition of $\rho_{\pi,\pi'}(y)$ to be $v = v^*$. The observation follows. ◀

Proof of Observation 6. For every $y \in Y \setminus \{0^d\}$, the set of instances v such that $v \cdot y > 0$ expands when we replace $v \in V$ by $v \in C(V)$. Thus for every y, $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y) \leq \rho_{\pi,\pi'}(y)$, and the observation follows. ◀

Proof of Observation 7. Let $y \in Y \setminus \partial(Y)$ be a point not on the boundary, and let $\delta > 1$ be a scalar such that $\delta y \in \partial(Y)$. For every $v \in V_y^+$, $v \cdot y \leq v \cdot \delta y$, and so $\rho_{\pi,\pi'}(v, y) \geq \rho_{\pi,\pi'}(v, \delta y)$. It follows that $\rho_{\pi,\pi'}(y) \geq \rho_{\pi,\pi'}(\delta y)$, proving the observation. ◀

Proof of Observation 11. This follows from Definitions 4 (integrality gap) and 10 (approximation ratio), and by noticing that even if $f(y) \in C(X) \setminus X$, for every $v \in V$ there must be some $x \in X$ with $v \cdot x \geq v \cdot f(y)$. ◀

Proof of Observation 30. Fix y and $v \in V_y^+$. Let f' be the oblivious rounding scheme that rounds y to $f(\delta y)$. So f' is reasonable by definition, and achieves $\frac{v \cdot f'(y)}{v \cdot y} = \frac{v \cdot f(\delta y)}{v \cdot y} \geq \frac{v \cdot f(\delta y)}{v \cdot \delta y}$. Taking the infimum over v maintains these relations, and so the observation holds. ◀

Proof of Observation 31. Let $y \in Y \setminus \partial(Y)$, $y \neq 0^d$ be a point not on the boundary, and let $\delta > 1$ be a scalar such that $\delta y \in \partial(Y)$. Since f is reasonable, $\alpha_{\pi,\pi'}(y) \geq \alpha_{\pi,\pi'}(\delta y)$, and since $\alpha_{\pi,\pi'}$ is achieved by taking the infimum over $Y \setminus \{0^d\}$, the observation follows. ◀

Proof of Observation 13. Let $y \in Y \setminus \partial(Y)$ be a point not on the boundary, and let $\delta > 1$ be a scalar such that $\delta y \in \partial(Y)$. Then for every $v \in V$ such that $v \cdot y > 0$, $v \cdot y < v \cdot \delta y$ and so $v \notin V_y^*$. If $v \cdot y = 0$, then by positivity of π', again $v \notin V_y^*$. We conclude that V_y^* is empty, completing the proof. ◀

A.1.1 Approximation Ratio of Reasonable Oblivious Rounding

We recall our assumption that π' is a convex relaxation. We say that an oblivious rounding scheme f is *reasonable* if it guarantees, for every relaxed solution $y \in Y \setminus \{0^d\}$, at least the approximation ratio $\alpha_{\pi,\pi'}(\delta y)$ that it achieves for $\delta y \in \partial(Y)$ (where $\delta \geq 1$ is the scaler by which y needs to be multiplied to reach the boundary). Assuming reasonability is without loss of generality as the following observation shows (see Appendix A for a proof):

▶ **Observation 30.** *For every oblivious rounding scheme f there is a reasonable oblivious rounding scheme f' such that for every $y \in Y' \setminus \{0^d\}$, the approximation ratio of f' at y is at least the approximation ratio of f at $\delta y \in \partial(Y)$, and so the overall approximation ratio of f' is at least that of f.*

For reasonable rounding schemes, the approximation ratios matter only on the boundary (see Appendix A for a proof):

▶ **Observation 31.** *The overall approximation ratio of any reasonable oblivious rounding scheme is not affected by the approximation ratios at relaxed solutions that lie strictly within the boundary: $\alpha_{\pi,\pi'} = \min_{y \in \partial(Y)} \alpha_{\pi,\pi'}(y)$.*

By Observations 7 and 31, if f is a reasonable oblivious rounding scheme and $\alpha_{\pi,\pi'}(y) = \rho_{\pi,\pi'}(y)$ for every $y \in \partial(Y)$, then f is tight (Definition 12).

B Appendix for Section 3

B.1 Proof of Theorem 16: Missing Details

We now formally state and prove our main lemma. We use $(C(V))_y^+$ to denote the set of instances $v \in C(V)$ such that $v \cdot y > 0$ (recall that V_y^+ is the set of such instances in V rather than in $C(V)$). We also use the following simple observation:

▶ **Observation 32.** *There exists $\epsilon > 0$ and $x' \in C(X)$ such that for every $v \in V$, $v \cdot x' \geq \epsilon$.*

Proof. Every feasible solution in X is a vector assigning nonnegative values to d variables $x_1, \ldots, x_d$. We may assume without loss of generality that for every coordinate $1 \leq j \leq d$, there is some solution $x^j \in X$ for which the variable x_j has strictly positive value. (Otherwise the variable x_j has value 0 in all feasible solutions and hence is redundant.) Consider a solution $x' = \frac{1}{d} \sum_{j=1}^{d} x^j \in C(X)$. All its coordinates are strictly positive. Recall that every $v \in V$ is nonnegative and not identically 0^d. Consequently $x' \cdot v > 0$ for every $v \in C(v)$. Moreover, our assumption that V is compact (see Section 2.1.1) together with the continuity of the inner product function implies that the function $f(v) = x' \cdot v$ attains a minimum over $v \in V$. Let $\epsilon = \min_{v \in V}[x' \cdot v]$ and note that $\epsilon > 0$. ◀

▶ **Lemma 33.** *Fix $y \in Y \setminus \{0^d\}$. There exists a value μ^* such that*

$$\mu^* = \max_{x \in C(X)} \inf_{v \in V_y^+} \frac{v \cdot x}{v \cdot y} = \inf_{v \in (C(V))_y^+} \max_{x \in X} \frac{v \cdot x}{v \cdot y}. \tag{3}$$

Moreover, there is a choice of $x^* \in C(X)$ and $v^* \in (C(V))_y^+$ such that x^* maximizes $\frac{v^* \cdot x}{v^* \cdot y}$, v^* minimizes $\frac{v \cdot x^*}{v \cdot y}$, and $\mu^* = \frac{v^* \cdot x^*}{v^* \cdot y}$.

Proof. Given $y \in Y \setminus \{0^d\}$, consider a series of two-player zero-sum games parameterized by $\mu \in \mathbb{R}_{\geq 0}$. In each such game, the *rounding* player has strategy set X, the *instance* player has strategy set V, and the payoff to the rounding player for choices $x \in X, v \in V$ is $v \cdot x - \mu(v \cdot y)$ (i.e., we use the difference as payoff instead of the ratio). Since X and V are both compact by assumption (see Section 2.1.1 and Remark 3.1), then Theorem 20 applies, and the unique minimax value p_μ of the game with parameter μ is

$$p_\mu = \max_{x \in C(X)} \min_{v \in V} \{v \cdot x - \mu(v \cdot y)\} = \min_{v \in C(V)} \max_{x \in X} \{v \cdot x - \mu(v \cdot y)\}.$$

Let $x_\mu \in C(X), v_\mu \in C(V)$ be equilibrium strategies that achieve the minimax value p_μ (by Theorem 20, such strategies are guaranteed to exist).

We now observe some properties of p_μ as a function of μ:

- p_μ is bounded: This is by the assumption that X and V are bounded.
- For sufficiently small μ, p_μ is positive: By Observation 32, $\max_{x \in c(X)} \min_{v \in V}\{v \cdot x\} \geq \min_{v \in C(V)}[x' \cdot v] \geq \epsilon$. Taking μ to be sufficiently small we can ensure that $\mu(v \cdot y) \leq \frac{\epsilon}{2}$ for every $v \in V$, because both y and V are bounded.
- For every μ such that $p_\mu \leq 0$ we have that $v_\mu \cdot y > 0$ for every equilibrium strategy v_μ (otherwise $p_\mu = v_\mu \cdot x_\mu$ and the rounding player can choose $x_\mu \in C(X)$ such that $v_\mu \cdot x_\mu > 0$). Hence $v_\mu \in (C(V))_y^+$.
- For large enough μ, p_μ is negative: Fix $v^+ \in C(V)_y^+$. Since $C(X)$ is bounded, we can set $\mu > (v^+ \cdot x)/(v^+ \cdot y)$ for every $x \in C(X)$. In particular, $v^+ \cdot x_\mu - \mu(v^+ \cdot y) < 0$, and so since v_μ is an equilibrium strategy, $p_\mu = v_\mu \cdot x_\mu - \mu(v_\mu \cdot y) \leq v^+ \cdot x_\mu - \mu(v^+ \cdot y) < 0$.
- p_μ is monotone (weakly) decreasing in μ: Let $\bar{\mu} > \mu$. Then
$$\begin{aligned} p_{\bar{\mu}} &= v_{\bar{\mu}} \cdot x_{\bar{\mu}} - \bar{\mu}(v_{\bar{\mu}} \cdot y) \\ &\leq v_\mu \cdot x_{\bar{\mu}} - \bar{\mu}(v_\mu \cdot y) \qquad (4) \\ &\leq v_\mu \cdot x_\mu - \mu(v_\mu \cdot y), \qquad (5) \end{aligned}$$
where (4) holds since $v_{\bar{\mu}}, x_{\bar{\mu}}$ are equilibrium strategies, and (5) holds since v_μ, x_μ are equilibrium strategies and $-\bar{\mu} < -\mu$.
- Let μ' be the smallest μ such that $p_{\mu'} \leq 0$, then p_μ is monotone *strictly* decreasing for $\mu \geq \mu'$: For every $\mu \geq \mu'$, by monotonicity $p_\mu \leq 0$, and so $v_\mu \cdot y > 0$. Thus for every $\mu \geq \mu'$, we can replace "$\leq$" by "$<$" in (5).
- p_μ is continuous when $\mu > 0$.

Given the above properties of p_μ, there is a unique $\mu^* > 0$ for which $p_{\mu^*} = 0$. We know that $v_{\mu^*} \cdot y > 0$, or equivalently, $v_{\mu^*} \in (C(V))^+_y$. The condition $v_{\mu^*} \cdot x_{\mu^*} - \mu^*(v_{\mu^*} \cdot y) = 0$ with positive $v_{\mu^*} \cdot y$ implies that $\mu^* = \frac{v^* \cdot x^*}{v^* \cdot y}$. This completes the proof. ◀

Proof of Theorem 16. Fix $y \in Y \setminus \{0^d\}$. On the one hand, for every oblivious rounding scheme f, recall from Definition 10 that the approximation ratio of f at y is $\alpha_{\pi,\pi'}(y) = \inf_{v \in V_y^+} \frac{v \cdot f(y)}{v \cdot y}$. Hence the oblivious rounding scheme with the optimal approximation ratio at y is the one that rounds y to $f(y) = \arg\max_{x \in C(X)} \inf_{v \in V_y^+} \frac{v \cdot x}{v \cdot y}$, achieving an approximation ratio of

$$\alpha_{\pi,\pi'}(y) = \max_{x \in C(X)} \inf_{v \in V_y^+} \frac{v \cdot x}{v \cdot y}. \tag{6}$$

On the other hand, recall from Definition 4 that the integrality gap at y with respect to the closure is

$$\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y) = \inf_{v \in (C(V))_y^+} \max_{x \in X} \frac{v \cdot x}{v \cdot y}. \tag{7}$$

So both parts of the theorem follow from Lemma 33, which states that (6) and (7) are equal. ◀

B.2 Rounding Optimal Solutions

A corollary of Theorem 16 applies to the approximation guarantees of oblivious rounding for solutions known to be optimal. The corollary follows directly from the observation in Section 2.2.2 that for every $y \in Y$ with nonempty V_y^*, $\alpha_{\pi,\pi'}(y) \leq \alpha^*_{\pi,\pi'}(y)$.

▶ **Corollary 34.** *Given a problem $\pi = (V, X)$ and a relaxation $\pi' = (V, Y)$, there exists an oblivious rounding scheme $f : Y \to C(X)$ that achieves an approximation ratio of $\alpha^*_{\pi,\pi'}(y) \geq \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y)$ at every point y with nonempty V_y^*. The overall approximation ratio of f is $\alpha^*_{\pi,\pi'} \geq \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}$.*

If $\pi = \mathrm{cl}(\pi)$ then there exists an oblivious rounding scheme $f : Y \to C(X)$ that is individually tight for optimal solutions.

Proof. By Definitions 4 (integrality gap) and 14 (approximation ratio for optimal solutions), if for an oblivious rounding scheme f it holds that $\alpha_{\pi,\pi'}(y) = \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y)$ for every $y \in Y \setminus \{0^d\}$, then $\alpha^*_{\pi,\pi'}(y) \geq \alpha_{\pi,\pi'}(y) = \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y)$ for every y with nonempty V_y^*. By Theorem 16 there exists such an oblivious rounding scheme. It follows that $\alpha^*_{\pi,\pi'} \geq \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}$. If $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y) = \rho_{\pi,\pi'}(y)$ for every y with nonempty V_y^*, then $\alpha^*_{\pi,\pi'}(y) \geq= \rho_{\pi,\pi'}(y)$, which by definition implies individual tightness for optimal solutions (Definition 15). ◀

The next example shows that, unlike the case in Corollary 18, there may be oblivious rounding schemes whose approximation ratio for optimal solutions $\alpha^*_{\pi,\pi'}$ surpasses the integrality gap of the closure $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}$. The reason for this difference is that $\alpha^*_{\pi,\pi'}$ only takes into account relaxed solutions that are guaranteed to be optimal for some instance of the relaxation.

▶ **Example 35.** Consider a problem $\pi = (V, X)$ of dimension 2, where the instances are $V = \{v_1, v_2\} = \{(1, 0), (0, 1)\}$ and the feasible solutions are $X = \{x_1, x_2, x_3\} = \{(0, 0), (1, 0), (0, 1)\}$. (For concreteness this example can be thought of as a welfare maximization problem with a single item and two buyers, where either: the first buyer has value 1 for the item and the other has value 0 – this is the first instance; or vice versa – this is the second instance. See Section 4 for more on welfare maximization.) Consider a relaxation $\pi' = (V, Y)$ where Y is a quadrilateral "kite" with vertices $\{(0,0),(1,0),(\frac{3}{4},\frac{3}{4}),(0,1)\}$. The closures $\mathrm{cl}(\pi), \mathrm{cl}(\pi')$ have an instance set $C(V)$ which is the set of vectors $\{(\lambda, 1-\lambda) \mid \lambda \in [0,1]\}$.

Oblivious rounding of the point $y = (\frac{3}{4},\frac{3}{4})$ gives a point $f(y)$ that belongs to $C(X)$, i.e., to the triangle with vertices $\{(0,0),(1,0),(0,1)\}$. For any such point $f(y)$, $\min\{v_1 \cdot f(y), v_2 \cdot f(y)\} \leq \frac{1}{2}$ whereas $v_1 \cdot y = v_2 \cdot y = \frac{3}{4}$, and so the approximation ratio $\alpha_{\pi,\pi'}(y)$ of any oblivious rounding scheme at y is $\leq \frac{1}{2}/\frac{3}{4} = \frac{2}{3}$. By rounding y to $(\frac{1}{2},\frac{1}{2})$ we get $\alpha_{\pi,\pi'}(y) = \frac{2}{3}$. It also follows that the overall approximation ratio of the best oblivious rounding scheme is $\leq \frac{2}{3}$.

Consider now the integrality gap $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y)$ at y with respect to the closure. For every $(\lambda, 1-\lambda) \in C(V)$, $\max\{(\frac{1}{2},\frac{1}{2}) \cdot x_1, (\frac{1}{2},\frac{1}{2}) \cdot x_2, (\frac{1}{2},\frac{1}{2}) \cdot x_3\} = \max\{\lambda, 1-\lambda\} \geq \frac{1}{2}$ whereas $(\lambda, 1-\lambda) \cdot y = \frac{3}{4}$, and so the integrality gap is $\leq \frac{2}{3}$. Since $(\frac{1}{2},\frac{1}{2}) \in C(V)$ we get $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}(y) = \frac{2}{3}$. This is equal to the approximation ratio $\alpha_{\pi,\pi'}(y)$ of the best oblivious rounding scheme at y, as known from Theorem 16. It also follows that the overall integrality gap $\rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')}$ is $\leq \frac{2}{3}$.

However, the point $y = (\frac{3}{4},\frac{3}{4})$ is not an optimal solution of the relaxation with respect to either of the instances in V. The set of optimal solutions $\{y \in Y \mid V_y^* \neq \emptyset\}$ includes only x_2 and x_3, and so the identity function is an oblivious rounding scheme with approximation ratio of 1 for optimal solutions. We conclude that $1 = \alpha^*_{\pi,\pi'} > \rho_{\mathrm{cl}(\pi),\mathrm{cl}(\pi')} = \frac{2}{3}$.

C Appendix for Section 4

C.1 Gross Substitutes and the Configuration LP

▶ **Definition 36.** A valuation ν is *gross substitutes* if the following holds. Consider any two item-price vectors $p, q \in \mathbb{R}^m$ such that $q \geq p$. Let S be a bundle such that $\nu(S) - \sum_{j \in S} p_j \geq \nu(T) - \sum_{j \in T} p_j$ for every bundle T. Let $S' = \{j \in S \mid q_j = p_j\}$. Then there exists a bundle U such that $S' \subseteq U$ and $\nu(U) - \sum_{j \in U} q_j \geq \nu(T) - \sum_{j \in T} q_j$ for every bundle T.

In words, a valuation is gross substitutes if for every bundle that maximizes the player's utility (value for the bundle minus the aggregate price of its items) given a price vector p, when prices of some of the items are raised, the items whose prices were not raised still participate in a bundle that maximizes the player's utility given the new price vector q. Intuitively, this monotonicity property facilitates the greedy approach in a similar way to matroid properties.

▶ **Definition 37.** The *integer programming* formulation of the welfare maximization problem is the following:

$$\max \textstyle\sum_{i,S} x_{i,S} v_{i,S}$$

s. t.

$$\textstyle\sum_{S} x_{i,S} \leq 1 \qquad \forall i \in N \tag{8}$$

$$\textstyle\sum_{i,S:j \in S} x_{i,S} \leq 1 \qquad \forall j \in M \tag{9}$$

$$x_{i,S} \in \{0,1\} \qquad \forall i \in N, S \subseteq M.$$

Constraint (8) corresponds to the requirement that no more than one bundle be allocated per player, and Constraint (9) corresponds to the requirement that no item is over-allocated. Note that the welfare maximization instance v appears only in the objective and does not affect X.

▶ **Definition 38.** The relaxed solution set Y of the *configuration LP relaxation* to the welfare maximization problem is the set of vectors y that are feasible solutions to the following LP:

$$\begin{aligned} & \max \textstyle\sum_{i,S} y_{i,S} v_{i,S} \\ \text{s. t.} & \\ & \textstyle\sum_S y_{i,S} \le 1 \qquad \forall i \in N \qquad (10) \\ & \textstyle\sum_{i,S:j\in S} y_{i,S} \le 1 \quad \forall j \in M \qquad (11) \\ & y_{i,S} \ge 0 \qquad \forall i \in N, S \subseteq M. \end{aligned}$$

Constraints (10) and (11) correspond to the same requirements as in the integer programming formulation above. However, the variables $y_{i,S}$ can now take any value in the interval $[0, 1]$, unlike the integral constraint in the IP problem.

C.2 Proof of Lemma 22: Coverage is the Closure of Unit-Demand

In this section we provide a proof of Lemma 22 for completeness (*cf.*, [6], Appendix A.1).

Proof of Lemma 22. Let UD and COV be the classes of unit-demand and coverage valuations, respectively. To prove the proposition we show that $C(\text{UD}) \subseteq \text{COV}$ and $\text{COV} \subseteq C(\text{UD})$. To show that $C(\text{UD}) \subseteq \text{COV}$, it is shown, in Lemma 39, that every unit-demand valuation is a coverage valuation (i.e., $\text{UD} \subseteq \text{COV}$ and thus $C(\text{UD}) \subseteq C(\text{COV})$), and, in Lemma 40, we show that the $C(\text{COV}) \subseteq \text{COV}$. The fact that $\text{COV} \subseteq C(\text{UD})$ is established in Lemma 41, and this completes the proof. ◀

▶ **Lemma 39.** *Every unit-demand valuation is a coverage valuation.*

Proof. Let v be a unit-demand valuation. We describe a coverage valuation v' satisfying $v'(S) = v(S)$ for every set $S \subseteq M$. Assume, by renaming, that $v_1 \le v_2 \le \cdots \le v_m$, and let $\Delta_j = v_j - v_{j-1}$. Let D be the set of indices of distinct values; i.e., $D = \{j \in [m] | \Delta_j > 0\}$ (with the convention that $v_0 = -1$). Associate an element with every distinct value v_j and set its weight to Δ_j. For every item j, E_j (i.e., the set of elements covered by j) is the set of elements corresponding to items up to item j. For example, if there are 4 items with values $v_1 = 1, v_2 = 1, v_3 = 3, v_4 = 8$, then there would be three elements, corresponding to items $1, 3, 4$ with weights $1, 2, 5$, respectively. For every $S \subseteq M$ it holds that

$$v'(S) = \sum_{e \in \bigcup_{j\in S} E_j} w(e) = \max_{j\in S} \sum_{e\in E_j} w(e) = \max_{j\in S} \sum_{k\in D \wedge k \le j} \Delta_k = \max_{j\in S} v_j = v(S),$$

as desired. ◀

▶ **Lemma 40.** *A convex combination of coverage functions is a coverage function.*

Proof. Let $v^1 = \langle E^1, w^1, \{E^1_j\}_j \rangle$ and $v^2 = \langle E^2, w^2, \{E^2_j\}_j \rangle$ be two coverage functions. It is sufficient to show that for every $\lambda \in [0, 1]$, $v(S) = \lambda v^1(S) + (1 - \lambda)v^2(S)$ is a coverage function, where S ranges over all subsets of M. Let $E = E^1 \biguplus E^2$, and let $w : E \to R^{\ge 0}$ be a weight function defined as $w(e) = \lambda w^1(e)$ for every $e \in E^1$, and $w(e) = (1 - \lambda)w^2(e)$ for every $e \in E^2$. Finally, for every item $j \in M$, let $E_j = E^1_j \biguplus E^2_j$. Consider the coverage function $v = \langle E, w, \{E_j\}_j \rangle$. For every set $S \subseteq M$ it holds that

$$\begin{aligned} v(S) &= \sum_{e \in \bigcup_{j \in S} E_j^1 \biguplus E_j^2} w(e) \\ &= \sum_{e \in \bigcup_{j \in S} E_j^1} w(e) + \sum_{e \in \bigcup_{j \in S} E_j^2} w(e) \\ &= \sum_{e \in \bigcup_{j \in S} E_j^1} \lambda w^1(e) + \sum_{e \in \bigcup_{j \in S} E_j^2} (1-\lambda) w^2(e) \\ &= \lambda v^1(S) + (1-\lambda) v^2(S), \end{aligned}$$

as desired. ◀

▶ **Lemma 41.** *Every coverage valuation can be expressed as a convex combination of unit-demand valuations.*

Proof. Let $v = \langle E, w, \{E_j\}_j \rangle$ be a coverage valuation, and let $k = |E|$ be the number of elements in E. We show that there exist k unit-demand valuations, whose average valuation for any set S equals $v(S)$. Associate a unit-demand function with every element as follows. For every element $e \in E$, let $S_e = \{j \in M : e \in E_j\}$ be the set of items that cover element e. The unit-demand valuation v^e associated with element e is defined by

$$v_j^e = \begin{cases} k \cdot w(e), & \text{if } j \in S_e \\ 0, & \text{otherwise.} \end{cases}$$

For every set of items $S \subseteq M$, let $E_S = \bigcup_{j \in S} E_j$, and let $\mathbb{1}\{e \in E_S\}$ be a binary function that returns 1 iff $e \in E_S$. We show that $v(S)$ can be written as a convex combination of the unit-demand functions described above. Indeed, for every set $S \subseteq M$,

$$\begin{aligned} \frac{1}{k} \sum_{e \in E} v^e(S) =& \frac{1}{k} \sum_{e \in E} \mathbb{1}\{e \in E_S\} k \cdot w(e) \\ =& \frac{1}{k} \sum_{e \in E_S} k \cdot w(e) \\ =& \sum_{e \in E_S} w(e) = v(S). \end{aligned}$$

◀