

Oblivious Rounding and the Integrality Gap*

Uriel Feige[†]

Michal Feldman[‡]

Inbal Talgam-Cohen[§]

May 4, 2017

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).

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 be maximized, represented as a vector in $\mathbb{R}_{\geq 0}^d$. The set X contains all feasible *solutions* to the problem, also represented as vectors in $\mathbb{R}_{\geq 0}^d$. 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.

*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.

[†]Weizmann Institute.

[‡]Tel-Aviv University.

[§]Hebrew University.

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 valuation functions in combinatorial auctions are often private knowledge of the buyers, and so a fundamental requirement of many allocation mechanisms is to incentivize *truthful* behavior on behalf of the buyers. One of the main known general methods for achieving truthfulness requires oblivious rounding [Lavi and Swamy, 2011]. Furthermore, 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., [Nisan and Segal, 2006]). Oblivious rounding limits the algorithmic stage in which communication is required, and there is no need for communication after a relaxed solution is found. Also (see Proposition 4.6), 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 [Dütting et al., 2015], 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.

¹Our notion of oblivious rounding is not to be confused with the rounding notion of Young [1995], which avoids solving a linear program – see the discussion of related work below.

Finally, oblivious rounding can lead to more robust welfare maximization algorithms in the following sense. Consider an approximation algorithm that guarantees an α -fraction of the optimal welfare provided that the buyers’ valuations belong to a “nice” class, e.g., the class of submodular valuations. A recent line of work (see, e.g., [Hassidim and Singer, 2016]) studies how robust such approximation algorithms are to noisy valuations, which belong to the nice class only approximately (for some appropriate notion of “approximately”). One natural notion is pointwise perturbation, where the value of every bundle of items is slightly perturbed, up to a $(1 + \epsilon)$ multiplicative factor, in comparison to a valuation that is exactly (say) submodular. Many welfare approximation algorithms are based on relax-and-round of a linear program called the “configuration LP”, which is solvable by “demand queries” to the valuations with no additional assumptions such as submodularity. If the rounding method is oblivious, its α -approximation guarantee for submodular valuations will extend to perturbed ones as well [Roughgarden et al., 2016, Proposition 6], ensuring robustness of the entire approximation algorithm, whereas if the rounding relies on the valuations and their submodularity then even small perturbations might lead to bad approximation in the worst-case (a related result appears in [Roughgarden et al., 2016, Proposition 8]).

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 $v \cdot y$. (Note that for maximization problems the integrality gap as defined here is at most 1, whereas for minimization problems it is at least 1.) 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 closed under convex combinations (because submodularity is preserved under convex combinations), but with *gross substitutes* valuations it is not (the average of two gross substitutes functions need not be gross substitutes itself – see Section 4).

Theorem (Main, 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 obviously 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 (Application, Informal). *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 The connection between various notions of rounding and algorithmic mechanism design has been studied in several works. Lavi and Swamy [2011] use the technique of randomized metarounding [Carr and Vempala, 2002] (see Section 5.1) 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). The work of Dütting et al. [2015] is directly related to the notion of oblivious rounding that we study (see also Proposition 4.6 below). Both works can be seen as strong motivation to systematically study oblivious rounding. Dughmi et al. [2011] require 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 [Feige and Jozeph, 2015a], 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. Theorem 1.8 in that work shows equivalence among the worst case approximation ratios 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.

Young [1995] introduces 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 and non-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 discuss 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 our set up of maximization problems, integrality gaps will be smaller than 1 (i.e., the larger and closer to 1, the more relevant to 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 solution of the original problem.

General Definitions Let dimension $d \in \mathbb{N}_{>0}$ 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 replaced 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 2.1. A problem π of dimension d is a pair (V, X) , where $V, X \subseteq \mathbb{R}_{\geq 0}^d$ 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

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 2.2. 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:

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

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

Definition 2.3. The problem $\text{cl}(\pi) = (C(V), X)$ is the closure of problem $\pi = (V, X)$.

²The 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.

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 $V' \subseteq \mathbb{R}_{\geq 0}^d \setminus \{0^d\}$ such that the instance set V is $\{v = cv' \mid c \in \mathbb{R}_{> 0} \text{ and } v' \in V'\}$. 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
- Recall that $V, X \subseteq \mathbb{R}_{\geq 0}^d$, and hence $v \cdot x \geq 0$ for every $v \in V$ and $x \in X$. 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 2.4 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 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 2.4. 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. First, in our formulation, the integrality gap $\rho_{\pi, \pi'}$ is always at most 1, and the larger and closer it is to 1, the better and tighter the relaxation.

³There 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.

Observation 2.5. *The integrality gap $\rho_{\pi, \pi'}$ is at most 1.*

Proof. 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. \square

Another observation is that taking the closure of a problem expands the instance set and so makes it “harder” to get a good relaxation:

Observation 2.6. *For every π and relaxation π' , $\rho_{\text{cl}(\pi), \text{cl}(\pi')} \leq \rho_{\pi, \pi'}$.*

Proof. 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 , $\rho_{\text{cl}(\pi), \text{cl}(\pi')}(y) \leq \rho_{\pi, \pi'}(y)$, and the observation follows. \square

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

Observation 2.7. *The overall integrality gap is not affected by the integrality gaps at relaxed solutions that lie strictly within the boundary: $\rho_{\pi, \pi'} = \min_{y \in \partial(Y)} \rho_{\pi, \pi'}(y)$.*

Proof. 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. \square

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 2.8. *A rounding scheme is a function $f : V \times Y \rightarrow C(X)$.*

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

Definition 2.9. *An oblivious rounding scheme is a function $f : Y \rightarrow C(X)$.*

Remark 2.10. *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 2.11. Consider an oblivious rounding scheme $f : Y \rightarrow C(X)$. 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.

Observation 2.12. For every $y \in Y \setminus \{0^d\}$, the approximation ratio of f at y is at most the integrality gap at y : $\alpha_{\pi, \pi'}(y) \leq \rho_{\pi, \pi'}(y)$. Therefore $\alpha_{\pi, \pi'} \leq \rho_{\pi, \pi'} \leq 1$.

Proof. This observation follows from Definitions 2.4 (integrality gap) and 2.11 (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 $v \cdot x \geq v \cdot f(y)$. \square

Observation 2.12 upper-bounds the approximation ratio, and a natural class of interest is rounding schemes for which this bound is tight:

Definition 2.13. An oblivious rounding scheme $f : Y \rightarrow C(X)$ 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 of Reasonable Oblivious Rounding

The observations in this subsection are included to achieve a comprehensive picture. 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 scalar by which y needs to be multiplied to reach the boundary). As the following observation shows, assuming reasonability is without loss of generality:

Observation 2.14. 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 .

Proof. 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. \square

For reasonable rounding schemes, the approximation ratios matter only on the boundary.

Observation 2.15. 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: $\alpha_{\pi, \pi'} = \min_{y \in \partial(Y)} \alpha_{\pi, \pi'}(y)$.

Proof. 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. \square

By Observations 2.7 and 2.15, 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 2.13).

2.2.3 Approximation Ratio for Optimal Solutions

Recall from (1) that V_y^* denotes the set of all instances for which y is an optimal solution. In this section we consider oblivious rounding schemes when the relaxed solution $y \in Y$ is known to be an optimal solution to some instance of the relaxed problem (namely, only for those y for which V_y^* is nonempty) and the approximation ratio is measured only with respect to instances v for which y is indeed optimal (namely, $v \in V_y^*$).

Observation 2.16. *If V_y^* is nonempty then $y \in \partial(Y)$.*

Proof. 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. \square

The two definitions in this subsection are analogous to Definitions 2.11 (approximation ratio) and 2.13 (tightness) above:

Definition 2.17. *Consider an oblivious rounding scheme $f : Y \rightarrow C(X)$. 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 $\alpha_{\pi, \pi'}^ = \inf_{y \in Y: V_y^* \neq \emptyset} \{\alpha_{\pi, \pi'}^*(y)\}$.*

Notice that according to Definition 2.17, the approximation ratio for optimal solutions restricts both the solution space (to ones that are optimal for some instances), and also the instance space (to instances that are optimal for the given solution).

By definition, for every $y \in Y$ with nonempty V_y^* it holds that $\alpha_{\pi, \pi'}(y) \leq \alpha_{\pi, \pi'}^*(y)$, and so $\alpha_{\pi, \pi'} \leq \alpha_{\pi, \pi'}^*$. 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 3.1 below). This motivates the next definition:

Definition 2.18. *An oblivious rounding scheme $f : Y \rightarrow C(X)$ 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. Section 3.1 contains our main theorem and its proof. Section 3.2 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 $\text{cl}(\pi) = (C(V), X)$ (Definition 2.3). Our main general theorem relates the (pointwise) approximation ratio of oblivious rounding to the integrality gap of the problem's closure:

Theorem 3.1. *Given a problem $\pi = (V, X)$ and a relaxation $\pi' = (V, Y)$:*

1. *Upper bound: For every oblivious rounding scheme $f : Y \rightarrow C(X)$, at every point $y \in Y \setminus \{0^d\}$, the approximation ratio $\alpha_{\pi, \pi'}(y)$ is at most the integrality gap $\rho_{\text{cl}(\pi), \text{cl}(\pi')}(y)$ of the closure of problem π .*
2. *Tightness: There exists an oblivious rounding scheme $f : Y \rightarrow C(X)$ such that $\alpha_{\pi, \pi'}(y) = \rho_{\text{cl}(\pi), \text{cl}(\pi')}(y)$ 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 3.2. *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 $\text{cl}(\pi)$ such that $\frac{v \cdot x}{v \cdot y} = \rho_{\text{cl}(\pi), \text{cl}(\pi')}(y) = \alpha_{\pi, \pi'}(y)$, where the approximation ratio is that of the best oblivious rounding scheme at y .*

Proof. By Lemma 3.8. □

Two useful corollaries follow immediately from Theorem 3.1. First, we have already observed that $\alpha_{\pi, \pi'} \leq \rho_{\pi, \pi'}$ (Observation 2.12) and that $\rho_{\text{cl}(\pi), \text{cl}(\pi')} \leq \rho_{\pi, \pi'}$ (Observation 2.6). It follows from Theorem 3.1 that for the best oblivious rounding scheme in fact $\alpha_{\pi, \pi'} = \rho_{\text{cl}(\pi), \text{cl}(\pi')}$.

Corollary 3.3. *Given a problem $\pi = (V, X)$ and a relaxation $\pi' = (V, Y)$, there exists an oblivious rounding scheme $f : Y \rightarrow C(X)$ that achieves an approximation ratio of $\alpha_{\pi, \pi'} = \rho_{\text{cl}(\pi), \text{cl}(\pi')}$, and this is the best possible approximation ratio of any oblivious rounding scheme.*

Proof. By Definitions 2.4 (integrality gap) and 2.11 (approximation ratio), if for an oblivious rounding scheme f it holds that $\alpha_{\pi, \pi'}(y) = \rho_{\text{cl}(\pi), \text{cl}(\pi')}(y)$ for every $y \in Y \setminus \{0^d\}$, then $\alpha_{\pi, \pi'} = \inf_{y \in Y \setminus \{0^d\}} \alpha_{\pi, \pi'}(y) = \inf_{y \in Y \setminus \{0^d\}} \rho_{\text{cl}(\pi), \text{cl}(\pi')} = \rho_{\text{cl}(\pi), \text{cl}(\pi')}$. By Theorem 3.1 there exists such an oblivious rounding scheme. □

Corollary 3.4. *Given a problem $\pi = (V, X)$ whose instances form a convex set (i.e., $\pi = \text{cl}(\pi)$), for every relaxation $\pi' = (V, Y)$, there exists an oblivious rounding scheme $f : Y \rightarrow C(X)$ that is individually tight.*

Proof. By Theorem 3.1 there exists an oblivious rounding scheme f such that $\alpha_{\pi, \pi'}(y) = \rho_{\text{cl}(\pi), \text{cl}(\pi')}(y)$ for every $y \in Y \setminus \{0^d\}$, and by assumption, $\rho_{\text{cl}(\pi), \text{cl}(\pi')}(y) = \rho_{\pi, \pi'}(y)$. The proof follows from the definition of individual tightness (Definition 2.13). □

Unlike the statement in Corollary 3.3, Example 3.1 shows that the approximation ratio for optimal solutions $\alpha_{\pi, \pi'}^*$ may surpass the integrality gap of the closure $\rho_{\text{cl}(\pi), \text{cl}(\pi')}$. More details on oblivious rounding of optimal solutions appear in Section 3.2 below.

3.1 Proof of Theorem 3.1 via Minimax

Our goal in this section is to prove Theorem 3.1 via our main lemma (Lemma 3.8), 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 3.5 (von Neumann [1928]). *For every bipartite zero-sum game in which the players' pure strategy sets X and V are compact and the payoff function $g : V \times X \rightarrow \mathbb{R}$ is continuous, there exists a unique minimax value μ^* such that*

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

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 3.6. *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 $V = V'$.*

3.1.1 Proof Intuition

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 3.5), 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 2.11). 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 3.5, 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* $\text{cl}(\pi)$ 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 $\rho_{\text{cl}(\pi), \text{cl}(\pi')}(v, y)$ in Definition 2.4. 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 $\rho_{\text{cl}(\pi), \text{cl}(\pi')}(y)$. 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 3.1 should follow directly by invoking Theorem 3.5. 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 3.8 and its proof show how to circumvent this problem by defining an appropriate series of zero-sum games.

3.1.2 Main Lemma and Proof 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 3.7. *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, \dots, 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 $\epsilon > 0$. \square

Lemma 3.8. *Fix $y \in Y \setminus \{0^d\}$. There exists a value μ^* such that*

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

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 3.6), then Theorem 3.5 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 3.5, 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 3.7, $\max_{x \in C(X)} \min_{v \in V} \{v \cdot x\} \geq \min_{v \in C(V)} [x' \cdot v] \geq \epsilon$. 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

$$\begin{aligned} p_{\bar{\mu}} &= v_{\bar{\mu}} \cdot x_{\bar{\mu}} - \bar{\mu}(v_{\bar{\mu}} \cdot y) \\ &\leq v_\mu \cdot x_{\bar{\mu}} - \bar{\mu}(v_\mu \cdot y) \end{aligned} \tag{4}$$

$$\leq v_\mu \cdot x_\mu - \mu(v_\mu \cdot y), \tag{5}$$

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 $\mu^* = \frac{v_{\mu^*} \cdot x_{\mu^*}}{v_{\mu^*} \cdot y}$. This completes the proof. □

Proof of Theorem 3.1. Fix $y \in Y \setminus \{0^d\}$. On the one hand, for every oblivious rounding scheme f , recall from Definition 2.11 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

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

On the other hand, recall from Definition 2.4 that the integrality gap at y with respect to the closure is

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

So both parts of the theorem follow from Lemma 3.8, which states that (6) and (7) are equal. \square

3.2 Rounding Optimal Solutions

A corollary of Theorem 3.1 applies to the approximation guarantees of oblivious rounding for solutions known to be optimal. The corollary follows directly from the observation in Section 2.2.3 that for every $y \in Y$ with nonempty V_y^* , $\alpha_{\pi, \pi'}(y) \leq \alpha_{\pi, \pi'}^*(y)$.

Corollary 3.9. *Given a problem $\pi = (V, X)$ and a relaxation $\pi' = (V, Y)$, there exists an oblivious rounding scheme $f : Y \rightarrow C(X)$ that achieves an approximation ratio of $\alpha_{\pi, \pi'}^*(y) \geq \rho_{\text{cl}(\pi), \text{cl}(\pi')}(y)$ at every point y with nonempty V_y^* . The overall approximation ratio of f is $\alpha_{\pi, \pi'}^* \geq \rho_{\text{cl}(\pi), \text{cl}(\pi')}$.*

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

Proof. By Definitions 2.4 (integrality gap) and 2.17 (approximation ratio for optimal solutions), if for an oblivious rounding scheme f it holds that $\alpha_{\pi, \pi'}(y) = \rho_{\text{cl}(\pi), \text{cl}(\pi')}(y)$ for every $y \in Y \setminus \{0^d\}$, then $\alpha_{\pi, \pi'}^*(y) \geq \alpha_{\pi, \pi'}(y) = \rho_{\text{cl}(\pi), \text{cl}(\pi')}(y)$ for every y with nonempty V_y^* . By Theorem 3.1 there exists such an oblivious rounding scheme. It follows that $\alpha_{\pi, \pi'}^* \geq \rho_{\text{cl}(\pi), \text{cl}(\pi')}$. If $\rho_{\text{cl}(\pi), \text{cl}(\pi')}(y) = \rho_{\pi, \pi'}(y)$ 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 2.18). \square

The next example shows that, unlike the case in Corollary 3.3, there may be oblivious rounding schemes whose approximation ratio for optimal solutions $\alpha_{\pi, \pi'}^*$ surpasses the integrality gap of the closure $\rho_{\text{cl}(\pi), \text{cl}(\pi')}$. 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 3.1. *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 $\text{cl}(\pi), \text{cl}(\pi')$ 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 $\leq \frac{1/2}{3/4} = \frac{2}{3}$. 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 $\rho_{\text{cl}(\pi), \text{cl}(\pi')}(y)$ 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 $\rho_{\text{cl}(\pi), \text{cl}(\pi')}(y) = \frac{2}{3}$. This is equal to the approximation

ratio $\alpha_{\pi, \pi'}(y)$ of the best oblivious rounding scheme at y , as known from Theorem 3.1. It also follows that the overall integrality gap $\rho_{\text{cl}(\pi), \text{cl}(\pi')}$ 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 $1 = \alpha_{\pi, \pi'}^* > \rho_{\text{cl}(\pi), \text{cl}(\pi')} = \frac{2}{3}$.

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 present the configuration LP for the problem. In Section 4.3 we show a fairness property of oblivious rounding for welfare maximization. In Section 4.4 we bound the approximation ratio of oblivious rounding schemes for welfare maximization with unit-demand valuations. In Section 4.5 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 3.1). In Section 4.6 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 $\nu_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$ that assigns a real value to every subset of items $S \subseteq M$ (also called a *bundle*). Valuations are 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, \dots, 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 Section 4.2). 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 [Kelso and Crawford, 1982] and in discrete convex optimization [Murota, 2003]; for a recent algorithmic survey see [Paes Leme, 2014]. There are many equivalent definitions of gross substitutes valuations, one of which we give for completeness in Section 4.2.

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

Proposition 4.1 (Bikhchandani and Mamer [1997]). *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 [Murota, 1996a,b].

A subclass of gross substitutes valuations is the class of *unit-demand* valuations. A valuation ν is unit-demand if there exists a vector $(\nu_1, \dots, \nu_m) \in \mathbb{R}_{\geq 0}^m$ 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) $w : E \rightarrow \mathbb{R}_{\geq 0}$ 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., [Badanidiyuru et al., 2012; Du et al., 2014]).

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 A for a proof).

Lemma 4.2. *The class of coverage valuations is the convex hull of unit-demand valuations.*

4.2 Gross Substitutes and the Configuration LP

Definition 4.3. *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. Gross substitutes valuations are not closed under convex combinations; their convex closure is the class *cone GS* defined in [Dughmi et al., 2011].

Definition 4.4. *The integer programming (IP) formulation of the welfare maximization problem is the following:*

$$\begin{aligned} & \max \sum_{i,S} x_{i,S} v_{i,S} \\ \text{s. t.} & \\ & \sum_S x_{i,S} \leq 1 \quad \forall i \in N, & (8) \\ & \sum_{i,S:j \in S} x_{i,S} \leq 1 \quad \forall j \in M, & (9) \\ & x_{i,S} \in \{0, 1\} \quad \forall i \in N, S \subseteq M. \end{aligned}$$

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 4.5. *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:*

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

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.

4.3 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 4.6. *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. \square

4.4 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 4.7. *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 4.8. *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 4.1. 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 4.7. By Corollary 3.3 and Lemma 4.2, 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 [Feige and Vondrák, 2010] 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. We name the items $a_{11}, a_{12}, a_{21}, a_{22}$, and we think of them as being the cells of a two by two array. There are six elements $\{H_1, H_2, V_1, V_2, D_1, D_2\}$, where H stands for ‘‘horizontal’’, V stands for ‘‘vertical’’, and D stands for ‘‘diagonal’’. 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, $w^2(H_1) = w^2(H_2) = 2$, $w^2(V_1) = w^2(V_2) = 0$, 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. \square

Proof of Proposition 4.8. By Corollary 3.3 and Lemma 4.2, it is sufficient to show an instance with n coverage valuations and an integrality gap of ≈ 0.782 . We claim that the instance in [Feige and Vondrák, 2010] 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 i th 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 [Feige and Vondrák, 2010], the integrality gap of this instance is ≈ 0.782 . This establishes the assertion of the proposition. \square

4.5 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 2.4; 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 $\text{LP}(v, y)$.

Lemma 4.9. *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 4.6).

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 $\alpha^*(y) \sum_{ik} \lambda_{ik} \nu_{ik}(M) = \alpha^*(y) \sum_i \nu_i(M) = \alpha^*(y) \text{LP}(v, y)$. 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). \square

The following proposition follows directly from Lemma 4.9.

Proposition 4.10. *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 4.9, 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 2.17). \square

Corollary 4.11. *The impossibility results in Propositions 4.7 and 4.8 apply also to the approximation ratio of oblivious rounding of optimal solutions of the configuration LP.*

Proof. The proof is by applying Proposition 4.10 and verifying that the instances in the proofs of Propositions 4.7 and 4.8 admit an individually optimal fractional solution. In the instance used in the proof of Proposition 4.7, 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 4.8, 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 $1/n$. \square

4.6 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 4.7, 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 4.1. *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 4.7: $\bar{H}_1 = \{a_{11}, a_{12}\}$, $\bar{H}_2 = \{a_{21}, a_{22}\}$, $\bar{V}_1 = \{a_{11}, a_{21}\}$, $\bar{V}_2 = \{a_{12}, a_{22}\}$, $\bar{D}_1 = \{a_{11}, a_{22}\}$, 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 and Non-Oblivious Rounding in the Literature

In this section we describe several rounding schemes from the literature that are oblivious, as well as several schemes that are non-oblivious. We begin in Section 5.1 with a general oblivious rounding method

that originates in the paper “Randomized Metarounding” by Carr and Vempala [2002], which we refer to as *canonical* oblivious rounding. In Section 5.2 we include examples of rounding methods for various algorithmic applications. Note that we have selected only a small set of examples to demonstrate and highlight some aspects of our framework, so the following list is by no means exhaustive:

1. Examples of oblivious rounding schemes:

- Threshold rounding for vertex cover by Hochbaum [1982].
- Randomized rounding by Raghavan and Thompson [1987] for set cover.
- Random hyperplane rounding of a semidefinite program (SDP) for max-cut by Goemans and Williamson [1995].
- Welfare maximization for subadditive and fractionally subadditive (XOS) valuations by Feige [2009]; and for submodular valuations by Feige and Vondrák [2010].

2. Examples of non-oblivious rounding schemes:

- Rounding of SDPs for the constraint satisfaction problem (CSP) by Raghavendra and Steurer [2009].
- Welfare maximization for XOS valuations by Dobzinski and Schapira [2006]; and for capped-additive valuations by Srinivasan [2008]; Chakrabarty and Goel [2010].
- Rounding of an SDP with triangle constraints for max-cut in *bounded degree* graphs by Feige et al. [2002].

In the context of our framework, these examples may fall into several different categories depending on whether the problem is convex, whether the integrality gap is known, and whether the best known rounding scheme is oblivious. For instance, it is interesting to note that welfare maximization with capped-additive valuations is not closed under convex combinations, and the best known approximation is by non-oblivious rounding of the *assignment LP*, which matches the integrality gap (Section 5.2.7). We remark that the integrality gap of the *configuration LP* for capped-additive valuations is unknown, nor is a non-oblivious rounding scheme that matches it. In comparison, welfare maximization with submodular, XOS and subadditive valuations is closed under convex combinations, and the best-known approximations (which achieve the relevant integrality gaps for XOS and subadditive) are oblivious (Section 5.2.6).

For max-cut, in general the problem is convex, and the approximation ratio of Goemans and Williamson’s well-known oblivious rounding scheme matches the integrality gap (Section 5.2.3). However, in the bounded degree case, the closure of bounded degree graphs is the set of *all* graphs, and the rounding scheme that achieves an approximation ratio better than that for general graphs is non-oblivious (Section 5.2.4).

Another related observation is that for CSP, our framework predicts the existence of an oblivious rounding scheme with an approximation ratio which matches that of the existing non-oblivious one (which itself matches the integrality gap – see Section 5.2.8).

5.1 Canonical Oblivious Rounding (a.k.a. Randomized Meta-rounding)

Let $\pi = (V, X)$ be a maximization problem and let $\pi' = (V, Y)$ be its relaxation. Let $\beta \leq 1$ be largest so that for every $y \in Y$ it holds that $\beta y \in C(X)$. Then we call a randomized oblivious rounding technique *canonical* if for every relaxed solution y it returns a distribution σ_y over X (feasible solutions) such that its expectation dominates βy . Namely, $E_{x \leftarrow \sigma_y}[x] \geq \beta y$, where the inequality holds coordinate-wise. (For minimization problems $\beta \geq 1$, and we require $E_{x \leftarrow \sigma_y}[x] \leq \beta y$.) In many cases (some of which will be presented shortly) it turns out that $\beta = \rho_{\pi, \pi'}$, and then the approximation ratio of canonical oblivious rounding matches the integrality gap of the underlying relaxation. Canonical oblivious rounding is referred

to as *randomized meta-rounding* by Carr and Vempala [2002], where sufficient conditions are (implicitly) given for it to match the integrality gap. It is shown that in these settings any (non-oblivious) rounding scheme with approximation ratio β can be transformed in polynomial time into a canonical oblivious rounding scheme. For many problems (some of which will be presented shortly) the known oblivious rounding techniques are canonical (and do not require the complicated machinery of Carr and Vempala [2002]). For some other problems (we will see such examples in the context of the maximum welfare problem), the known oblivious rounding techniques are not canonical, and any attempt to make them canonical results in a substantial loss in the approximation ratio.

5.2 Examples of Rounding Schemes for Different Applications

5.2.1 Threshold Rounding for Minimum Cost Vertex Cover

Input A graph $G(U, E)$ with vertex costs $c : U \rightarrow \mathbb{R}_{\geq 0}$, where c_i denotes the cost of vertex i .

IP formulation and LP relaxation

$$\begin{aligned} & \min \sum_{i \in U} c_i x_i \\ & \text{s. t.} \\ & \quad x_i + x_j \geq 1 \quad \forall (i, j) \in E, \\ & \quad x_i \in \{0, 1\} \quad \forall i \in U. \end{aligned}$$

The LP relaxation replaces $x_i \in \{0, 1\}$ by $y_i \geq 0$.

Threshold rounding [Hochbaum, 1982] Round y_i to 1 if $y_i \geq \frac{1}{2}$, and to 0 otherwise.

Discussion This is a minimization problem rather than a maximization problem, but our discussion of oblivious rounding extends to minimization problems in a straightforward way. For minimization problems, approximation ratios will be larger than 1.

This rounding is oblivious (since it does not depend on c) and canonical. It has an approximation ratio of 2, which matches the integrality gap of the LP up to an additive term of $O(\frac{1}{|V|})$.

Remark 5.1. *Threshold rounding does not guarantee that the resulting vertex cover is minimal, namely, does not contain a strictly smaller vertex cover. By removing vertices until the vertex cover becomes minimal, a procedure that can be done obliviously, the approximation ratio improves by low order terms. However, here we shall not be concerned with these low order terms in the approximation ratio.*

The set of instances in the vertex cover problem is convex. In our framework, one may think of all instances as sharing the same graph G , and differing only by the cost vector c . The reason why we need all instances to share the same graph is because in our framework we assume that the set of feasible solutions is known to the rounding procedure, and this set depends on the input graph.

Remark 5.2. *Observe that our convention that all instances share the same graph implies that for the common case of unweighted vertex cover (in which all costs are 1), there is a better oblivious rounding procedure than threshold rounding in the following sense: One can compute the minimum cardinality vertex cover in G (this might take exponential time, but we do not limit the running time of the oblivious rounding procedure, as long as it does not access the cost vector), and output it. This is an example where the approximation ratio of the described rounding scheme according to our definition throughout the paper – the ratio between $c \cdot x$ achieved by the rounding and $c \cdot y$ achieved by the relaxed solution – differs from the ratio between $c \cdot x$ and the optimal solution, which is in this case 1.*

Remark 5.3. *It is an interesting question whether the fact that the graph is fixed and only the cost of the vertices differs between instances makes the minimum cost vertex cover problem easier to approximate (compared to the case that also the graph itself is part of the input). For a discussion of this issue in the context of the related problem of maximum weight independent set (which is the complement of minimum weight vertex cover), see [Feige and Jozeph, 2014].*

5.2.2 Randomized Rounding for Minimum Cost Set Cover

Input A collection \mathcal{S} of sets over a universe \mathcal{U} of items, with set costs $c : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$, where c_S denotes the cost of set S .

IP formulation and LP relaxation

$$\begin{aligned} & \min \sum_{S \in \mathcal{S}} c_S x_S \\ & \text{s. t.} \\ & \sum_{S \in \mathcal{S} | u \in S} x_S \geq 1 \quad \forall u \in \mathcal{U}, \\ & x_S \in \{0, 1\} \quad \forall S \in \mathcal{S}. \end{aligned}$$

The LP relaxation replaces $x_S \in \{0, 1\}$ by $y_S \geq 0$.

A scheme based on randomized rounding [Raghavan and Thompson, 1987] First apply the randomized rounding scheme of [Raghavan and Thompson, 1987]: independently, for each $S \in \mathcal{S}$, round y_S to 1 with probability $\min\{1, y_S \ln |U|\}$ (where $\ln = \log_e$ and $|U|$ is the number of items). After the randomized rounding step, if some items remain uncovered, pick for every uncovered item u a set $S \in \mathcal{S}$ that contains u , with probability proportional to y_S , and add it to the set cover.

Discussion This rounding is oblivious (does not depend on c) and canonical. Its approximation ratio is roughly $\ln |U|$, which matches the integrality gap of the LP up to low order additive terms. The set of instances is convex. In our framework, one may think of all instances as sharing the same set system, and differing only by the cost vector c . The second step of the rounding procedure certainly needs to know which items are contained in each set. Remarks 5.1, 5.2 and 5.3 apply here as well.

5.2.3 Random Hyperplane Rounding for Max-Cut

Input A complete graph $G(U, E)$ with edge weights $w : E \rightarrow \mathbb{R}_{\geq 0}$, where $w_{i,j}$ denotes the weight of edge (i, j) .

Integer quadratic program formulation and SDP relaxation

$$\begin{aligned} & \max \frac{1}{2} \sum_{i,j} w_{i,j} x_{i,j} \\ & \text{s. t.} \\ & z_{i,j} = 1 - x_{i,j} \quad \forall i, j, \\ & Z = (z_{i,j}) = \mathbf{z} \cdot \mathbf{z}^T \text{ where for every } i, z_i \in \{\pm 1\}. \end{aligned}$$

The SDP relaxation is as follows:

$$\begin{aligned} & \max \frac{1}{2} \sum_{i < j} w_{i,j} y_{i,j} \\ & \text{s. t.} \\ & z_{i,j} = 1 - y_{i,j} \quad \forall i, j, \\ & z_{i,i} = 1 \quad \forall i, \\ & Z = (z_{i,j}) \text{ is a symmetric and PSD matrix.} \end{aligned}$$

Hyperplane rounding [Goemans and Williamson, 1995] The rounding process has two stages: (1) Decompose Z into $B^T B$. (2) “Round” the vectors $\{b_i\}$ to ± 1 scalars $\{z_i\}$ by passing a random hyperplane through the origin of the unit sphere, and labeling vectors on one half by $+1$ and on the other half by -1 . Then recover $(z_{i,j})$ and $(x_{i,j})$ from the quadratic program formulation.

Discussion This rounding is oblivious (does not depend on w) and canonical. Its approximation ratio is ≈ 0.878 , which matches the integrality gap of the SDP [Feige and Schechtman, 2002]. The set of instances is convex.

5.2.4 Rounding with Misplacement Correction for Bounded Degree Max-Cut

Input A complete graph $G(U, E)$ with edge weights $w : E \rightarrow \{0, 1\}$, where $w_{i,j}$ indicates whether edge (i, j) appears. The graph has maximal degree Δ (the number of edges with non-zero weights adjacent to any vertex is at most Δ).

Integer quadratic program formulation and SDP relaxation The integer program is the same as in the unbounded degree case. The SDP relaxation is almost the same, with additional *triangle constraints* of the form $z_{i,j} + z_{i,k} + z_{j,k} \geq -1$ and $z_{i,j} - z_{i,k} - z_{j,k} \geq -1$ for every i, j, k .

Non-oblivious rounding [Feige et al., 2002] First run the oblivious rounding scheme of [Goemans and Williamson, 1995]. Then move “misplaced vertices” to the other side of the partition, where a vertex is misplaced if it and more than half of its neighbors (vertices with whom it shares a non-zero edge) lie on the same side of the partition.

Discussion This rounding is non-oblivious (depends on w). Its approximation ratio is $\alpha + \Omega(1/\Delta^4)$ where $\alpha \approx 0.878$ is the approximation ratio of the [Goemans and Williamson, 1995] algorithm. The set of instances is non-convex.

5.2.5 Rounding for Welfare Maximization with General Valuations

Input Recall from Section 4 that M is a set of m items and N is a set of n players. The input includes a monotone valuation function $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$ for each player $i \in N$. (The valuations are usually assumed to be given via oracle access to the set values.)

IP formulation and LP relaxation The IP formulation appears in Definition 4.4, and the LP relaxation – known as the configuration LP – appears in Definition 4.5. Recall that $y_{i,S}$ is the fractional allocation of item set S to player i .

Remark 5.4. *The configuration LP has exponentially many variables but can be solved optimally in polynomial time assuming access to a demand oracle, which returns the set of items a player would want to buy given a set of prices for the individual items.*

Canonical rounding Iteratively allocate bundles to players, where after each such allocation the allocated player and all bundles that intersect with the allocated bundle are removed. The next allocation is chosen according to one of two options, which is determined randomly once at the beginning of the rounding process: The first option is to pick a player i and a bundle S at random with probability proportional to $\{y_{i,S}\}$. The second option is to pick an item j at random, and then to pick at random $y_{i,S}$ such that $j \in S$.

Discussion This rounding is oblivious (does not depend on the valuations $\{v_i\}$) and canonical. Its approximation ratio is $\approx 1/\sqrt{m}$ (since the probability of allocating bundle S to player i is at least $\frac{1}{2}y_{i,S}(\frac{1}{1+(1-y_{i,S})|S|} + \frac{|S|}{m})$, and by case analysis over the size of S). This is tight up to lower order terms. The set of instances is convex.

5.2.6 Two-Step Rounding for Complement-Free Welfare Maximization

Input Here we consider welfare maximization when valuation functions are *subadditive*, i.e., for every i and item subsets S, T , $v_i(S \cup T) \leq v_i(S) + v_i(T)$. A strict subclass is *fractionally* subadditive valuations, also known as *XOS*, for which there is a set of additive valuations such that the value of a set S is the maximum value for S across all additive valuations. In turn, a strict subclass of XOS valuations is that of submodular valuations, in which for every i and item subsets S, T , $v_i(S \cup T) + v_i(S \cap T) \leq v_i(S) + v_i(T)$. All three sets of instances (i.e., subadditive, XOS and submodular) are convex.

IP formulation and LP relaxation The IP formulation appears in Definition 4.4, and the LP relaxation – known as the configuration LP – appears in Definition 4.5. Recall that $y_{i,S}$ is the fractional allocation of item set S to player i .

Two-step randomized rounding framework [Dobzinski et al., 2010] The first step is a tentative allocation step, where each player i chooses a tentative set of items S with probability $x_{i,S}$. The solution might be infeasible (i.e., tentative sets of different players might intersect). The second step is a *contention resolution* step, where for each item j , if it is allocated to several players under the tentative allocation, one of these players is chosen to receive item j . The rounding schemes described below all adopt this two-step framework but differ in their approach to the second step.

Greedy rounding for XOS valuations [Dobzinski and Schapira, 2006] In order to resolve contentions, access the valuations (via an “XOS oracle”) and allocate each item greedily to the player who values it the most, based on the supporting additive functions.

Fair rounding for XOS valuations [Feige, 2009] Contentions are resolved in a way that achieves the following “fairness” property, despite using no knowledge of the valuations: conditioned on being involved in a contention for an item, a player gets the item with probability at least $1 - 1/e$. Another version of this rounding scheme appears under the name “fair rounding” in [Feige and Vondrák, 2010].

Discussion The rounding scheme of [Dobzinski and Schapira, 2006] is non-oblivious (depends on the valuations $\{v_i\}$), and achieves an approximation ratio of $1 - 1/e$. The rounding scheme of [Feige, 2009] is oblivious, and also achieves an approximation ratio of $1 - 1/e$. This matches the integrality gap of the LP up to low order terms. Neither of the described rounding schemes are canonical (the probability that bundle S is allocated as a whole to a player i might be much smaller than $(1 - \frac{1}{e})y_{i,S}$, due to the contention resolution step).

Guiding graph rounding for subadditive valuations [Feige, 2009] Describing this rounding scheme is beyond the scope of this paper; it uses a different, yet oblivious contention resolution. Its approximation ratio of $1/2$ matches the integrality gap of the LP up to low order terms.

Fair and “butterfly” rounding for submodular valuations [Feige and Vondrák, 2010] Describing this rounding scheme is beyond the scope of this paper; it involves (among other ideas) choosing two tentative bundles per player and not just one, and uses an oblivious contention resolution method. This rounding has an approximation ratio of $1 - \frac{1}{e} + \epsilon$, where $\epsilon \geq 10^{-5}$. It is unknown (but believed to be unlikely) that this matches the integrality gap of the configuration LP for submodular valuations.

5.2.7 Rounding Methods for Capped-Additive Welfare Maximization

Input A monotone valuation function v_i for each player $i \in N$, where the valuations are *capped-additive* (also known as “budget-additive”); i.e., for every i there is an additive valuation ν_i with value $\nu_{i,j}$ for every item $j \in M$, and a cap c_i such that the value of a set S is $\min\{\nu_i(S), c_i\}$.

IP formulation and LP relaxation

$$\begin{aligned} & \max \sum_{i,j} x_{i,j} \nu_{i,j} \\ \text{s. t.} & \\ & \sum_j x_{i,j} \nu_{i,j} \leq c_i \quad \forall i \in N, \\ & \sum_i x_{i,j} \leq 1 \quad \forall j \in M, \\ & x_{i,j} \in \{0, 1\} \quad \forall i \in N, j \in M. \end{aligned}$$

The LP relaxation replaces $x_{i,j} \in \{0, 1\}$ by $y_{i,j} \geq 0$, and is known as the *assignment* LP. We remark that an alternative formulation leads to the configuration LP relaxation (Definition 4.5); the integrality gap of this relaxation for capped-additive valuations is unknown [Chakrabarty and Goel, 2010].

Iterative rounding [Chakrabarty and Goel, 2010] The structure of an extreme point solution to the LP guarantees that there is always at least one “leaf” player, i.e., a player such that all items she is fractionally allocated, except for at most one, are not fractionally allocated to any other player. The rounding assigns all such items to their leaf players. Then a residual LP is formulated (with carefully defined new budgets), and the process iterates.

Conversion ratio rounding [Srinivasan, 2008] A solution to the LP induces a bipartite graph, with players on one side, items on the other, and edge weights that correspond to the fractional allocation $\{y_{i,j}\}$. Weights can be shifted along cycles, using “conversion ratios” $\nu_{i,j}/\nu_{i,j'}$ as guidelines, until some $y_{i,j}$ becomes integer. Eventually the graph becomes a forest. Then weight shifts are continued in the same manner along maximal paths, where the direction of the shift is chosen randomly such that the expected shift is 0. This is repeated until all edges become integer.

Discussion Both rounding schemes are not oblivious (depend on the values $\{\nu_{i,j}\}$, and in the case of [Chakrabarty and Goel, 2010], also on the caps $\{c_i\}$), and both have an approximation ratio of 0.75, which matches the integrality gap of the assignment LP. The set of capped-additive instances is not convex.

5.2.8 Rounding SDPs for CSP

To avoid the introduction of heavy notation we consider here only Boolean k -CSPs, namely constraint satisfaction problems in which each constraint involves k Boolean variables. A comprehensive treatment that extends also to more general CSPs appears in [Raghavendra and Steurer, 2009, Section 3.1].

Input A variable set of size n , a finite set $\{P\}$ of allowable Boolean predicates, and a distribution over clauses, where each clause is composed of a predicate from $\{P\}$ and a tuple of k variables on which the predicate is applied. (For a given instance, the number of clauses is the support of the distribution is typically denoted by m .) The goal is to find a Boolean assignment to the variables that maximizes the sum of weights (according to the given distribution) of satisfied clauses.

Formulation and SDP relaxation A feasible solution includes n binary values which specify the assignment to the variables. This induces, for every predicate P over k variables, an indicator vector specifying the “local” assignment (out of the 2^k possibilities) to these variables.

A relaxed solution includes n vectors whose inner products form a symmetric positive semidefinite matrix, and for every predicate P , a distribution $\{\mu_{P,x}\}$ over local assignments x . The constraints of the relaxed formulation ensure that the inner products of the vectors corresponding to the variables of P match the distribution μ_P .

Generic rounding [Raghavendra and Steurer, 2009] At a high level, the rounding steps are: reduce the dimension of the SDP solution by random projection; discretize the projected vectors; solve the resulting CSP instance by brute force (using the distribution over clauses). In more detail, the rounding scheme is defined by a “folding” operation – given the SDP solution, the variables are clustered into a constant number of clusters, and by introducing a single new variable for every cluster, a constant-sized instance is formed. This can be solved by brute force and translated directly to a feasible solution of the original instance. The key is to fold without significantly changing the SDP value. This is achieved by random projection, so that one can get a constant number of clusters without distorting the inner products too much.

Discussion This rounding scheme is not oblivious, because of the brute force step. Its approximation ratio matches the integrality gap of the SDP up to low order terms. The set of CSP instances is convex.

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? Second, for which problems are there *polynomial-time* computable oblivious rounding schemes that are comparable to the integrality gap? Carr and Vempala [2002] have identified special cases for which the approximation ratio of a polynomial time, canonical oblivious rounding scheme is guaranteed to match the integrality gap. On the other hand, the work of [Feige and Jozeph, 2015b] demonstrates that in some cases matching the integrality gap in polynomial time can be hard. They do so by constructing integer programs whose relaxations have good exponential time rounding procedures, but no good polynomial time rounding procedures (under standard complexity assumptions). Finally, what else can we hope to learn about the most promising rounding techniques from properties of the combinatorial problem?

A Appendix: Coverage is the Closure of Unit-Demand

In this section we provide a proof of Lemma 4.2 for completeness (a related result appears in [Dughmi et al., 2011, Appendix A.1]).

Proof of Lemma 4.2. 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 A.1, 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 A.2, we show that the $C(\text{COV}) \subseteq \text{COV}$. The fact that $\text{COV} \subseteq C(\text{UD})$ is established in Lemma A.3, and this completes the proof. \square

Lemma A.1. *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 $v_1 \leq v_2 \leq \dots \leq v_m$, and let $\Delta_j = v_j - v_{j-1}$. Let D be the set of indices of distinct values; i.e., $D = \{j \in [m] \mid \Delta_j > 0\}$ (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

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

as desired. \square

Lemma A.2. *A convex combination of coverage functions is a coverage function.*

Proof. Let $v^1 = \langle E^1, w^1, \{E_j^1\}_j \rangle$ and $v^2 = \langle E^2, w^2, \{E_j^2\}_j \rangle$ 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 $E = E^1 \uplus E^2$, and let $w : E \rightarrow \mathbb{R}^{\geq 0}$ 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 $E_j = E_j^1 \uplus E_j^2$. Consider the coverage function $v = \langle E, w, \{E_j\}_j \rangle$. For every set $S \subseteq M$ it holds that

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

as desired. \square

Lemma A.3. *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 $S \subseteq M$,

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

References

- Ashwinkumar Badanidiyuru, Shahar Dobzinski, Hu Fu, Robert Kleinberg, Noam Nisan, and Tim Roughgarden. 2012. Sketching valuation functions. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*. 1025–1035.
- Sushil Bikhchandani and John W. Mamer. 1997. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory* 74, 2 (1997), 385–413.
- Robert Carr and Santosh Vempala. 2002. Randomized metarounding. *Random Structures and Algorithms* 20, 3 (2002), 343–352.
- Deeparnab Chakrabarty and Gagan Goel. 2010. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and GAP. *SIAM J. Comput.* 39, 6 (2010), 2189–2211.
- Shahar Dobzinski, Noam Nisan, and Michael Schapira. 2010. Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Oper. Res.* 35, 1 (2010), 1–13.
- Shahar Dobzinski and Michael Schapira. 2006. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*. 1064–1073.
- Nan Du, Yingyu Liang, Maria-Florina Balcan, and Le Song. 2014. Learning time-varying coverage functions. In *Proceedings of the 27th Neural Information Processing Systems Conference*. 3374–3382.
- Shaddin Dughmi, Tim Roughgarden, and Qiqi Yan. 2011. From convex optimization to randomized mechanisms: Toward optimal combinatorial auctions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*. 149–158.
- Paul Dütting, Thomas Kesselheim, and Éva Tardos. 2015. Algorithms as mechanisms: The price of anarchy of relax-and-round. In *Proceedings of the 16th ACM Conference on Economics and Computation*. 187–201.
- Uriel Feige. 2009. On maximizing welfare when utility functions are subadditive. *SIAM J. Comput.* 39, 1 (2009), 122–142.
- Uriel Feige and Shlomo Jozeph. 2014. Demand queries with preprocessing. In *Proceedings of the 41st International Colloquium on Automata, Languages and Programming*. 477–488.
- Uriel Feige and Shlomo Jozeph. 2015a. Oblivious algorithms for the maximum directed cut problem. *Algorithmica* 71, 2 (2015), 409–428.
- Uriel Feige and Shlomo Jozeph. 2015b. Separation between estimation and approximation. In *Proceedings of the 6th Conference on Innovations in Theoretical Computer Science*. 271–276.
- Uriel Feige, Marek Karpinski, and Michael Langberg. 2002. Improved approximation of Max-Cut on graphs of bounded degree. *J. Algorithms* 43, 2 (2002), 201–219.
- Uriel Feige and Gideon Schechtman. 2002. On the optimality of the random hyperplane rounding technique for MAX CUT. *Random Structures and Algorithms* 20, 3 (2002), 403–440.

- Uriel Feige and Jan Vondrák. 2010. The submodular welfare problem with demand queries. *Theory of Computing* 6, 1 (2010), 247–290.
- Michel X. Goemans and David P. Williamson. 1995. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM* 42, 6 (1995), 1115–1145.
- Avinatan Hassidim and Yaron Singer. 2016. Submodular optimization under noise. (2016). Working paper.
- Dorit S. Hochbaum. 1982. Approximation algorithms for the set covering and vertex cover problems. *SIAM J. Comput.* 11 (1982), 555–556.
- Alexander S. Kelso and Vincent P. Crawford. 1982. Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica* 50, 6 (1982), 1483–1504.
- Ron Lavi and Chaitanya Swamy. 2011. Truthful and near-optimal mechanism design via linear programming. *J. ACM* 58, 6 (2011). Article 25.
- Kazuo Murota. 1996a. Valuated matroid intersection I: Optimality criteria. *SIAM J. Discrete Math.* 9, 4 (1996), 545–561.
- Kazuo Murota. 1996b. Valuated matroid intersection II: Algorithms. *SIAM J. Discrete Math.* 9, 4 (1996), 562–576.
- Kazuo Murota. 2003. *Discrete Convex Analysis*. Society for Industrial and Applied Mathematics.
- Noam Nisan and Ilya Segal. 2006. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory* 129 (2006), 192–224.
- Renato Paes Leme. 2014. Gross substitutability: An algorithmic survey. (2014). Working paper.
- Prabhakar Raghavan and Clark D. Thompson. 1987. Randomized rounding: A technique for provably good algorithms and algorithmic proofs. *Combinatorica* 7, 4 (1987), 365–374.
- Prasad Raghavendra and David Steurer. 2009. How to round any CSP. In *Proceedings of the 50th Symposium on Foundations of Computer Science*. 586–594.
- Tim Roughgarden, Inbal Talgam-Cohen, and Jan Vondrák. 2016. When are welfare guarantees robust? (2016). Working paper available at <https://arxiv.org/abs/1608.02402>.
- Aravind Srinivasan. 2008. Budgeted allocations in the full-information setting. In *Proceedings of the 11th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*. 247–253.
- John von Neumann. 1928. Zur theorie der gesellschaftsspiele. *Math. Annalen*. 100 (1928), 295–320.
- Neal E. Young. 1995. Randomized rounding without solving the linear program. In *Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms*. 170–178.