

Near-Optimal and Robust Mechanism Design for Covering Problems with Correlated Players*

Hadi Minooei[†]

Chaitanya Swamy[†]

Abstract

We consider the problem of designing incentive-compatible, ex-post individually rational (IR) mechanisms for covering problems in the Bayesian setting, where players' types are drawn from an underlying distribution and may be correlated, and the goal is to minimize the expected total payment made by the mechanism. We formulate a notion of incentive compatibility (IC) that we call *support-based IC* that is substantially more robust than Bayesian IC, and develop black-box reductions from support-based-IC mechanism design to algorithm design. For single-dimensional settings, this black-box reduction applies even when we only have an LP-relative *approximation algorithm* for the algorithmic problem. Thus, we obtain near-optimal mechanisms for various covering settings including single-dimensional covering problems, multi-item procurement auctions, and multidimensional facility location.

1 Introduction

In a *covering mechanism-design problem*, there are players who provide covering objects and a buyer who needs to obtain a suitable collection of objects so as to satisfy certain covering constraints (e.g., covering a ground set); each player incurs a certain private cost, which we refer to as his *type*, for providing his objects, and the mechanism must therefore pay the players from whom objects are procured. We consider the problem of designing incentive-compatible, ex-post individually rational (IR) mechanisms for covering problems (also called *procurement auctions*) in the *Bayesian setting*, where players' types are drawn from an underlying distribution and may be *correlated*, and the goal is to *minimize the expected total payment made by the mechanism*. Consider the simplest such setting of a *single-item procurement auction*, where a buyer wants to buy an item from any one of n sellers. Each seller's private type is the cost he incurs for supplying the item and the sellers must therefore be incentivized via a suitable payment scheme. Myerson's seminal result [19] solves this problem (and other single-dimensional problems) when players' private types are *independent*. However, no such result (or characterization) is known when players' types are correlated. This is the question that motivates our work.

Whereas the analogous revenue-maximization problem for packing domains, such as combinatorial auctions (CAs), has been extensively studied in the algorithmic mechanism design (AMD) literature, both in the case of independent and correlated (even interdependent) player-types (see, e.g., [3, 4, 2, 1, 5, 9, 8, 21, 2, 25] and the references therein), surprisingly, there are *almost no results* on the payment-minimization problem for covering settings in the AMD literature (see however the discussion in "Related work" for some exceptions). The economics literature does contain various general results that apply to both covering and packing problems. However much of this work

*A preliminary version appeared as [18].

[†]{hminooei, cswamy}@uwaterloo.ca. Dept. of Combinatorics and Optimization, Univ. Waterloo, Waterloo, ON N2L 3G1. Research supported partially by NSERC grant 327620-09 and the second author's Discovery Accelerator Supplement Award, and Ontario Early Researcher Award.

focuses on characterizing special cases; see, e.g., [14, 13]. An exception is the work of Crémer and McLean [6, 7], which shows that under certain conditions, one can devise a Bayesian-incentive-compatible (BIC) mechanism whose expected total payment exactly equal to the expected cost incurred by the players, albeit one where players may incur negative utility under certain type-profile realizations.

Our contributions. We initiate a study of payment-minimization (PayM) problems from the AMD perspective of designing computationally efficient, near-optimal mechanisms. We develop black-box reductions from mechanism design to algorithm design whose application yields a variety of optimal and near-optimal mechanisms. As we elaborate below, covering problems turn out to behave quite differently in certain respects from packing problems, which necessitates new approaches (and solution concepts).

Formally, we consider the setting of *correlated players* in the explicit model, that is, where we have an explicitly-specified *arbitrary discrete* joint distribution of players’ types. A commonly-used solution concept in Bayesian settings is Bayesian incentive compatibility (BIC) and interim individual rationality (interim IR), wherein at the *interim* stage when a player knows his type but is oblivious of the random choice of other players’ types, truthful participation in the mechanism by all players forms a Bayes-Nash equilibrium. Two serious drawbacks of this solution concept (which are exploited strikingly and elegantly in [6, 7]) are that: (i) a player may regret his decision of participating and/or truth-telling *ex post*, that is, after observing the realization of other players’ types; and (ii) it is overly-reliant on having precise knowledge of the true underlying distribution making this a rather *non-robust* concept: if the true distribution differs, possibly even slightly, from the mechanism designer’s and/or players’ beliefs or information about it, then the mechanism could lose its IC and IR properties.

On the other hand, the solution concept of dominant-strategy incentive compatibility (DSIC) and ex-post IR ensures that truthful participation is the best choice, and a no-regret choice, for every player *regardless* of the other players’ reported types. This is the most robust notion of IC and IR, since it is a completely distribution-independent.

We consider the problem of designing near-optimal mechanisms for payment-minimization problems with robustness being an important consideration. This makes (DSIC, ex-post IR), which we abbreviate to (DSIC, IR), as the natural ideal. However, certain difficulties arise in achieving this goal for covering problems (see “Differences . . . packing problems” below). We formulate a notion of incentive compatibility that we call *support-based IC*¹ that, while somewhat weaker than DSIC, is still substantially more robust than BIC and at the same time is flexible enough that it allows one to obtain various polytime near-optimal mechanisms satisfying this notion. A *support-based-(IC, IR)* mechanism (see Section 2) ensures that truthful participation in the mechanism is in the best interest of every player (i.e. a “no-regret” choice) for every type profile in a certain *subset* of the type space: for every player i , we impose the IC and IR conditions for all type profiles of the form (\cdot, c_{-i}) for every profile c_{-i} of other players’ types coming from the support of the underlying distribution. In particular, a support-based-(IC, IR) mechanism ensures that truthful participation is the best choice for every player *even at the ex-post stage* when the other players’ (randomly-chosen) types are revealed to him. Such a mechanism is significantly more robust than a (BIC, interim-IR) mechanism since it retains its IC and IR properties for a large class of distributions that contains (in particular) every distribution whose support is a subset of the support of the actual distribution. In other words, in keeping with Wilson’s doctrine of detail-free mechanisms, the mechanism functions robustly even under fairly limited information about the type-distribution.

¹The conference version [18] of this paper referred to this as “robust Bayesian IC.”

We show that for a variety of settings, one can reduce the support-based-(IC, IR) payment-minimization (PayM) mechanism-design problem to the algorithmic cost-minimization (CM) problem of finding an outcome that minimizes the total cost incurred. Moreover, this black-box reduction applies to: (a) single-dimensional settings even when we only have an LP-relative *approximation algorithm* for the CM problem (that is required to work only with nonnegative costs) (Theorem 4.2); and (b) *multidimensional problems* with additive types (Corollary 3.3).

Our reduction yields near-optimal support-based (IC-in-expectation, IR) mechanisms for a variety of covering settings such as (a) various single-dimensional covering problems including single-item procurement auctions (Table 1); (b) multi-item procurement auctions (Theorem 5.1); and (c) multidimensional facility location (Theorem 5.3). (Support-based IC-in-expectation means that the support-based-IC guarantee holds for the *expected utility* of a player, where the expectation is over the random coin tosses of the mechanism.) In Section 6, we consider some extensions involving both weaker (but more robust than (BIC, interim IR)) solution concepts and the stronger (DSIC, IR) solution concept. We obtain the same guarantees under the various weaker solution concepts, and adapt our techniques to obtain (DSIC-in-expectation, IR) mechanisms with the same guarantees for single-dimensional problems in time exponential in the number of players (Section 6.2). These are the first results for the PayM mechanism-design problem for covering settings with correlated players under a notion stronger than (BIC, interim IR). To our knowledge, our results are new even for the simplest covering setting of single-item procurement auctions.

Our techniques. The starting point for our construction is a linear-programming (LP) relaxation (P) for the problem of computing an optimal support-based-(IC, IR)-in-expectation mechanism. This was also the starting point in the work of [8], which considers the revenue-maximization problem for CAs, but the covering nature of the problem makes it difficult to apply certain techniques utilized successfully in the context of packing problems (as described below).

We show that an optimal solution to (P) can be computed given an optimal algorithm \mathcal{A} for the CM problem since \mathcal{A} can be used to obtain a separation oracle for the dual LP. Next, we prove that a feasible solution to (P) can be extended to a support-based-(IC-in-expectation, IR) mechanism with no larger objective value.

For single-dimensional problems, we show that even LP-relative ρ -approximation algorithms for the CM problem can be utilized, as follows. We move to a relaxation of (P), where we replace the set of allocations with the feasible region of the CM-LP. This can be solved efficiently, since the separation oracle for the dual can be obtained by optimizing over the feasible region of CM-LP, which can be done efficiently! But now we need to work harder to “round” an optimal solution (x, p) to the relaxation of (P) and obtain a support-based-(IC-in-expectation, IR) mechanism. Here, we exploit the Lavi-Swamy [12] convex-decomposition procedure, using which we can show (roughly speaking) that we can decompose ρx into a convex combination of allocations. This allows us to obtain a support-based-(IC-in-expectation, IR) mechanism while blowing up the payment by a ρ -factor.

In comparison with the reduction in [8], which is the work most closely-related to ours, our reduction from support-based-IC mechanism design to the algorithmic CM problem is stronger in the following sense. For single-dimensional settings, it applies even with LP-relative *approximation algorithms*, and the approximation algorithm is required to work only for “proper inputs” with nonnegative costs. (Note that whereas for packing problems, allowing negative-value inputs can be benign, this can change the character of a covering problem considerably.) In contrast, Dobzinski et al. [8] require an exact algorithm for the analogous social-welfare-maximization (SWM) problem.

Differences with respect to packing problems. At a high level, our method of writing an LP for the underlying mechanism-design problem and solving it given an algorithm for an associated algorithmic problem is similar to the procedure in Dobzinski et al. [8]. However, we encounter three distinct sources of difficulty when dealing with covering problems vis-a-vis packing problems.

First, as noted in [8], the LP can only encode the IC and IR conditions for a finite set of type profiles, whereas, with an infinite type space, both support-based-(IC, IR) and (DSIC, IR) require the mechanism to satisfy the IC and IR conditions for an infinite set of type profiles. Therefore, to translate the LP solution to a suitable mechanism, we need to solve the “extension problem” of extending an allocation and pricing rule defined on a (finite) subset of the type space to the entire type space while preserving its IC and IR properties. This turns out to be a much more difficult task for covering problems than for packing domains. The key difference is that for a packing setting such as combinatorial auctions, one can show that *any* LP solution—in particular, the optimal LP solution—can be converted into a (DSIC-in-expectation, IR) mechanism without any loss in expected revenue (see Section C.1). (Consequently, [8] obtain (DSIC-in-expectation, IR) mechanisms.) Intuitively, this works because one can focus on a single player by allocating no items to the other players. Clearly, one cannot mimic this for covering problems: dropping players may render the problem infeasible, and it is not clear how to extend an LP-solution to a (DSIC-in-expectation, IR) mechanism for covering problems. We suspect that not every LP solution or support-based-(IC, IR) mechanism can be extended to a (DSIC, IR) mechanism, and that there is a gap between the optimal expected total payments of support-based-(IC-in-expectation, IR) and (DSIC, IR) mechanisms. We leave these as open problems.²

Due to this complication, we sacrifice a modicum of the IC, IR properties in favor of obtaining polytime near-optimal mechanisms and settle for the weaker, but still quite robust notion of support-based (IC-in-expectation, IR). We consider this to be a reasonable starting point for exploring mechanism-design solutions for covering problems, which leads to various interesting research directions.³

A second difficulty, which we have alluded to above, arises due to the fact that solving the LP requires one to solve the CM problem with negative-valued inputs. This is also true of packing problems [8] (where one needs to solve the SWM problem), even in the single-item setting [19] (where reserve prices arise due to negative virtual valuations). While this is not a problem if we have an optimal algorithm for the CM problem, it creates serious issues, even in the single-dimensional setting, if we only have an approximation algorithm at hand; in particular, the standard notion of approximation becomes meaningless since the optimum could be negative. In contrast, for packing problems with single-dimensional types (or additive types [2, 3, 4]), these issues are more benign since one may always discard players (or options of players) with negative value. In particular, an approximation algorithm can be used to obtain an approximate separation oracle for the dual LP, and thus obtain a near-optimal solution to the primal LP via a well-known technique in approximation algorithms. (We sketch this extension of a result of [8] in Appendix C.1.)

Finally, a stunning aspect where covering and packing problems diverge can be seen when one considers the idea of a k -lookahead auction [22, 8]. This was used by [8] to convert their results in

²We show in Section 6.2 that if we expand our LP to include IC and IR constraints for a much larger (exponential-sized) set of type-profiles then, for single-dimensional settings, it is possible to extend an LP-solution to a (DSIC-in-expectation, IR) mechanism.

³For covering problems, even formulating the LP in a way that its solution can be extended to a support-based-(IC, IR) mechanism is somewhat tricky due to the following complication. We need to argue (see Lemma 2.3) that there is an optimal support-based-(IC-in-expectation, IR) mechanism such that for every player i , and every profile c_{-i} coming from the support of the underlying distribution, there is a type profile (m_i, c_{-i}) under which the mechanism never procures objects from i , and include this condition in the LP (otherwise, the IC and IR conditions would force the LP-extension to make arbitrarily large payments).

the explicit model to the oracle model introduced by [22]. This however fails spectacularly in the covering setting. One can show that even for single-item procurement auctions, dropping even a single player can lead to an arbitrarily large payment compared to the optimum (see Appendix B).

Other related work. In the economics literature, the classical results of Crémer and McLean [6, 7] and McAfee and Reny [16], also apply to covering problems, and show that one can devise a (BIC, interim IR) mechanism with correlated players whose expected total payment is at most the expected total cost incurred provided the underlying type-distribution satisfies a certain full-rank assumption. These mechanisms may however cause a player to have negative utility under certain realizations of the random type profile.

The AMD literature has concentrated mostly on the independent-players setting [3, 4, 2, 1, 5, 9]. There has been some, mostly recent, work that also considers correlated players [22, 23, 8, 21, 2, 25]. Much of this work pertains to the revenue-maximization setting; an exception is [23], which is discussed below. Ronen [22] considers the single-item auction setting in the *oracle model*, where one samples from the distribution conditioned on some players' values. He proposes the (1-) lookahead auction and shows that it achieves a $\frac{1}{2}$ -approximation. Papadimitriou and Pierrakos [21] show that the optimal (DSIC, IR) mechanism for the single-item auction can be computed efficiently with at most 2 players, and is *NP*-hard otherwise. Cai et al. [2] give a characterization of the optimal auction under certain settings. Roughgarden and Talgam-Cohen [25] consider interdependent types, which generalizes the correlated type-distribution setting, and develop an analog of Myerson's theory for certain such settings.

Ronen and Lehmann [23] consider the PayM problem in the setting where a buyer wants to buy an item from sellers who can supply the item in one of many configurations and incur private costs for supplying the item. However, this procurement problem is in fact a *packing problem*: one can view a solution to be feasible if it selects at most one configuration for procurement; in particular, the buyer has the flexibility of *not procuring* the item. As noted earlier, this flexibility drastically alters the character of the mechanism-design problem. Not surprisingly, the results therein, which are based on lookahead auctions, do not apply in the covering setting (as noted above).

Various reductions from revenue-maximization to SWM are given in [2, 3, 4]. The reductions in [2, 4] also apply to covering problems and the PayM objective, but they are incomparable to our results. These works focus on the (BIC, interim-IR) solution concept, which is a rather weak/liberal notion for correlated distributions. Most (but not all) of these consider independent players and additive valuations, and often require that the SWM-algorithm also work with negative values, which is a benign requirement for downward-closed environments such as CAs but is quite problematic for covering problems when only has an approximation algorithm. Cai et al. [2] consider correlated players and obtain mechanisms having running time polynomial in the maximum support-size of the marginal distribution of a player, which could be substantially smaller than the support-size of the entire distribution. This savings can be traced to the use of the (BIC, interim-IR) notion which allows [2] to work with a compact description of the mechanism. It is unclear if these ideas are applicable when one considers robust-(BIC, IR) mechanisms. A very interesting open question is whether one can design robust-(BIC-in-expectation, IR) mechanisms having running time polynomial in the support-sizes of the marginal player distributions (as in [2, 8]).

2 Preliminaries

Covering mechanism-design problems. We adopt the formulation in [17] to describe general covering mechanism-design problems. There are some items that need to be covered, and n players who provide covering objects. Let $[k]$ denote the set $\{1, \dots, k\}$. Each player i provides a set \mathcal{T}_i

of covering objects. All this information is public knowledge. Player i has a *private cost* or *type* vector $c_i = \{c_{i,v}\}_{v \in \mathcal{T}_i}$, where $c_{i,v} \geq 0$ is the cost he incurs for providing object $v \in \mathcal{T}_i$; for $T \subseteq \mathcal{T}_i$, we use $c_i(T)$ to denote $\sum_{v \in T} c_{i,v}$. A feasible solution or allocation selects a subset $T_i \subseteq \mathcal{T}_i$ for each agent i , denoting that i provides the objects in T_i , such that $\bigcup_i T_i$ covers all the items. Given this solution, each agent i incurs the private cost $c_i(T_i)$, and the mechanism designer incurs a publicly known cost $\text{pub}(T_1, \dots, T_n) \geq 0$, which may be used to encode any feasibility constraints in the covering problem.

Let C_i denote the set of all possible types of agent i , and $C = \prod_{i=1}^n C_i$. We assume (for notational simplicity) that $C_i = \mathbb{R}_+^{|\mathcal{T}_i|}$. Let $\Omega := \{(T_1, \dots, T_n) : \text{pub}(T_1, \dots, T_n) < \infty\}$ be the (finite) set of all feasible allocations. For a tuple $x = (x_1, \dots, x_n)$, we use x_{-i} to denote $(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$. Similarly, let $C_{-i} = \prod_{j \neq i} C_j$. For an allocation $\omega = (T_1, \dots, T_n)$, we sometimes use ω_i to denote T_i , $c_i(\omega)$ to denote $c_i(\omega_i) = c_i(T_i)$, and $\text{pub}(\omega)$ to denote $\text{pub}(T_1, \dots, T_n)$. We make the mild assumption that $\text{pub}(\omega') \leq \text{pub}(\omega)$ if $\omega_i \subseteq \omega'_i$ for all i ; so in particular, if ω is feasible, then adding covering objects to the ω_i s preserves feasibility.

A (direct revelation) *mechanism* $M = (\mathcal{A}, p_1, \dots, p_n)$ for a covering problem consists of an allocation algorithm $\mathcal{A} : C \mapsto \Omega$ and a payment function $p_i : C \mapsto \mathbb{R}$ for each agent i . Each agent i reports a cost function c_i (that might be different from his true cost function). The mechanism computes the allocation $\mathcal{A}(c) = (T_1, \dots, T_n) = \omega \in \Omega$, and pays $p_i(c)$ to each agent i . The *utility* $u_i(c_i, c_{-i}; \bar{c}_i)$ that player i derives when he reports c_i and the others report c_{-i} is $p_i(c) - \bar{c}_i(\omega_i)$ where \bar{c}_i is his true cost function, and each agent i aims to maximize his own utility. We refer to $\max_i |\mathcal{T}_i|$ as the dimension of a covering problem. Thus, for a *single-dimensional* problem, each player i 's cost can be specified as $c_i(\omega) = c_i \alpha_{i,\omega}$, where $c_i \in \mathbb{R}_+$ is his private type and $\alpha_{i,\omega} = 1$ if $\omega_i \neq \emptyset$ and 0 otherwise.

The above setup yields a *multidimensional covering mechanism-design problem* with *additive types*, where by additivity we mean that the private cost that a player i incurs for providing a set $T \subseteq \mathcal{T}_i$ of objects is additive across the objects in T (i.e., it is $\sum_{v \in T} c_{i,v}$). Notice that if $c_i, c'_i \in C_i$, then the type $c_i + c'_i$ defined by $\{c_{i,v} + c'_{i,v}\}_{v \in \mathcal{T}_i}$ is also in C_i and satisfies $(c_i + c'_i)(\omega) = c_i(\omega) + c'_i(\omega)$ for all $\omega \in \Omega$. It is possible to define more general multidimensional settings, but additive types is a reasonable starting point to explore the multidimensional covering mechanism-design setting. (As noted earlier, there has been almost no work on designing polytime, near-optimal mechanisms for covering problems.)

The Bayesian setting. We consider *Bayesian* settings where there is an underlying publicly-known *discrete* and possibly *correlated* joint type-distribution on C from which the players' types are drawn. We consider the so-called explicit model, where the players' type distribution is explicitly specified. We use $\mathcal{D} \subseteq C$ to denote the support of the type distribution, and $\Pr_{\mathcal{D}}(c)$ to denote the probability of realization of $c \in C$. Also, we define $\mathcal{D}_i := \{c_i \in C_i : \exists c_{-i} \text{ s.t. } (c_i, c_{-i}) \in \mathcal{D}\}$, and \mathcal{D}_{-i} to be $\{c_{-i} : \exists c_i \text{ s.t. } (c_i, c_{-i}) \in \mathcal{D}\}$.

Solution concepts. A mechanism sets up a game between players, and the solution concept dictates certain desirable properties that this game should satisfy, so that one can reason about the outcome when rational players are presented with a mechanism satisfying the solution concept. The two chief properties that one seeks to capture relate to incentive compatibility (IC), which (roughly speaking) means that every agent's best interest is to reveal his type truthfully, and individual rationality (IR), which is the notion that no agent is harmed by participating in the mechanism. Differences and subtleties arise in Bayesian settings depending on the stage at which we impose these properties and how robust we would like these properties to be with respect to the underlying type distribution.

Definition 2.1 A mechanism $M = (\mathcal{A}, \{p_i\})$ is Bayesian incentive compatible (BIC) and interim IR if for every player i and every $\bar{c}_i \in \mathcal{D}_i, c_i \in C_i$, we have $\mathbb{E}_{c_{-i}}[u_i(\bar{c}_i, c_{-i}; \bar{c}_i)|\bar{c}_i] \geq \mathbb{E}_{c_{-i}}[u_i(c_i, c_{-i}; \bar{c}_i)|\bar{c}_i]$ (BIC) and $\mathbb{E}_{c_{-i}}[u_i(\bar{c}_i, c_{-i}; \bar{c}_i)|\bar{c}_i] \geq 0$ (interim IR), where $\mathbb{E}_{c_{-i}}[\cdot|\bar{c}_i]$ denotes the expectation over the other players' types conditioned on i 's type being \bar{c}_i .

As mentioned earlier, the (BIC, interim-IR) solution concept may yet lead to ex-post “regret”, and is quite non-robust in the sense that the mechanism’s IC and IR properties rely on having detailed knowledge of the distribution; thus, in order to be confident that a BIC mechanism achieves its intended functionality, one must be confident about the “correctness” of the underlying distribution, and learning this information might entail significant cost. To remedy these weaknesses, we propose and investigate the following stronger IC and IR notions.

Definition 2.2 A mechanism $M = (\mathcal{A}, \{p_i\})$ is support-based IC and support-based IR, if for every player i , every $\bar{c}_i, c_i \in C_i$, and every $c_{-i} \in \mathcal{D}_{-i}$, we have $u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq u_i(c_i, c_{-i}; \bar{c}_i)$ (support-based IC) and $u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq 0$ (support-based IR).

Support-based (IC, IR) ensures that participating truthfully in the mechanism is in the best interest of every player even at the ex-post stage when he knows the realized types of all players. To ensure that support-based IC and support-based IR are compatible, we focus on *monopoly-free* settings: for every player i , there is some $\omega \in \Omega$ with $\omega_i = \emptyset$.

Notice that support-based (IC, IR) is subtly weaker than the notion of (dominant-strategy IC (DSIC), IR), wherein the IC and IR conditions of Definition 2.2 must hold for all $c_{-i} \in C_{-i}$, ensuring that truth-telling and participation are no-regret choices for a player even if the other players’ reports are outside the support of the underlying type-distribution. We focus on support-based IC because it forms a suitable middle-ground between BIC and DSIC: it inherits the desirable robustness properties of DSIC, making it much more robust than BIC (and closer to a worst-case notion), and yet is flexible enough that one can devise polytime mechanisms satisfying this solution concept.

It might seem strange that in the definition of support-based (IC, IR), we consider player i ’s incentives for types outside of i ’s support and for type-profiles that are inconsistent with the underlying distribution. Keeping robustness in mind, our goal here is to approach the ideal of (DSIC, IR) and we have therefore formulated the most-robust notion that permits us to devise polytime near-optimal mechanisms satisfying this notion. In Section 6, we consider various alternate solution concepts that, while weaker than support-based (IC, IR), still retain its robustness properties to a large extent, and show that our results extend easily to these solution concepts.

The above definitions are stated for a deterministic mechanism, but they have analogous extensions to a randomized mechanism M ; the only change is that the $u_i(\cdot)$ and $p_i(\cdot)$ terms are now replaced by the expected utility $\mathbb{E}_M[u_i(\cdot)]$ and expected price $\mathbb{E}_M[p_i(\cdot)]$ respectively, where the expectation is over the random coin tosses of M . We denote the analogous solution concept for a randomized mechanism by appending “in expectation” to the solution concept, e.g., a (BIC, interim IR)-in-expectation mechanism denotes a randomized mechanism whose expected utility satisfies the BIC and interim-IR requirements stated in Definition 2.1.

A support-based-(IC, IR)-in-expectation mechanism $M = (\mathcal{A}, \{p_i\})$ can be easily modified so that the IR condition holds *with probability 1* (with respect to M ’s coin tosses) while the expected payment to a player (again over M ’s coin tosses) is unchanged: on input c , if $\mathcal{A}(c) = \omega \in \Omega$ with probability q , the new mechanism returns, with probability q , the allocation ω , and payment $c_i(\omega) \cdot \frac{\mathbb{E}_M[p_i(c)]}{\mathbb{E}_M[c_i(\omega)]}$ to each player i (where we take 0/0 to be 0, so if $c_i(\omega) = 0$, the payment to i is 0). Thus, we obtain a mechanism whose expected utility satisfies the support-based-IC condition, and

IR holds with probability 1 for all $\bar{c}_i \in C_i, c_{-i} \in \mathcal{D}_{-i}$. A similar transformation can be applied to a (DSIC, IR)-in-expectation mechanism.

Optimization problems. Our main consideration is to minimize the expected total payment of the mechanism. It is natural to also incorporate the mechanism-designer’s cost into the objective. Define the *disutility* of a mechanism $M = (\mathcal{A}, \{p_i\})$ under input c to be $\sum_i p_i(c) + \kappa \cdot \text{pub}(\mathcal{A}(c))$, where $\kappa \geq 0$ is a given scaling factor. Our objective is to devise a polynomial-time support-based-(IC (in-expectation), IR)-mechanism with minimum expected disutility. Since most problems we consider have $\text{pub}(\omega) = 0$ for all feasible allocations, in which case disutility equals the total payment, abusing terminology slightly, we refer to the above mechanism-design problem as the *payment-minimization* (PayM) problem. (An exception is metric *uncapacitated facility location* (UFL), where players provide facilities and the underlying metric is public knowledge; here, $\text{pub}(\omega)$ is the total client-assignment cost of the solution ω .) We always use O^* to denote the expected disutility of an optimal mechanism for the PayM problem under consideration.

We define the *cost minimization* (CM) problem to be the *algorithmic* problem of finding $\omega \in \Omega$ that minimizes the total cost $\sum_i c_i(\omega) + \text{pub}(\omega)$ incurred.

The following technical lemma, whose proof we defer to Appendix A, will prove quite useful since it allows us to restrict the domain to a bounded set, which is essential to achieve IR with finite prices. (For example, in the single-dimensional setting, the payment is equal to the integral of a certain quantity from 0 to ∞ , and a bounded domain ensures that this is well defined.) Note that such complications do not arise for packing problems. Let $\mathbf{1}_{\mathcal{T}_i}$ be the $|\mathcal{T}_i|$ -dimensional all 1s vector. Let \mathcal{I} denote the input size.

Lemma 2.3 *We can efficiently compute an estimate $m_i > \max_{c_i \in \mathcal{D}_i, v \in \mathcal{T}_i} c_{i,v}$ with $\log m_i = \text{poly}(\mathcal{I})$ for all i such that there is an optimal support-based-(IC-in-expectation, IR) mechanism $M^* = (\mathcal{A}^*, \{p_i^*\})$ where $\mathcal{A}^*(m_i \mathbf{1}_{\mathcal{T}_i}, c_{-i})_i = \emptyset$ with probability 1 (over the random choices of M^*) for all i and all $c_{-i} \in \mathcal{D}_{-i}$.*

It is easy to obtain the stated estimates if we consider only *deterministic* mechanisms, but it turns out to be tricky to obtain this when one allows randomized mechanisms due to the artifact that a randomized mechanism may choose arbitrarily high-cost solutions as long as they are chosen with small enough probability. In the sequel, we set $\bar{\mathcal{D}}_i := \mathcal{D}_i \cup \{m_i \mathbf{1}_{\mathcal{T}_i}\}$ for all $i \in [n]$, and $\bar{\mathcal{D}} := \bigcup_i (\bar{\mathcal{D}}_i \times \mathcal{D}_{-i})$. Note that $|\bar{\mathcal{D}}| = O(n|\mathcal{D}|^2)$.

3 LP-relaxations for the payment-minimization problem

The starting point for our results is the LP (P) that essentially encodes the payment-minimization problem. Throughout, we use i to index players, c to index type-profiles in $\bar{\mathcal{D}}$, and ω to index Ω . We use variables $x_{c,\omega}$ to denote the probability of choosing ω , and $p_{i,c}$ to denote the expected payment to player i , for input c . For $c \in \bar{\mathcal{D}}$, let $\Omega(c) = \Omega$ if $c \in \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$, and otherwise if $c = (m_i \mathbf{1}_{\mathcal{T}_i}, c_{-i})$, let $\Omega(c) = \{\omega \in \Omega : \omega_i = \emptyset\}$ (which is non-empty since we are in a monopoly-free setting).

$$\min \quad \sum_{c \in \mathcal{D}} \Pr_{\mathcal{D}}(c) \left(\sum_i p_{i,c} + \kappa \sum_{\omega} x_{c,\omega} \text{pub}(\omega) \right) \quad (\text{P})$$

$$\text{s.t.} \quad \sum_{\omega} x_{c,\omega} = 1 \quad \forall c \in \bar{\mathcal{D}} \quad (1)$$

$$p_{i,(c_i,c_{-i})} - \sum_{\omega} c_i(\omega) x_{(c_i,c_{-i}),\omega} \geq p_{i,(c'_i,c_{-i})} - \sum_{\omega} c_i(\omega) x_{(c'_i,c_{-i}),\omega} \quad \forall i, c_i, c'_i \in \bar{\mathcal{D}}_i, c_{-i} \in \mathcal{D}_{-i} \quad (2)$$

$$p_{i,(c_i,c_{-i})} - \sum_{\omega} c_i(\omega) x_{(c_i,c_{-i}),\omega} \geq 0 \quad \forall i, c_i \in \bar{\mathcal{D}}_i, c_{-i} \in \mathcal{D}_{-i} \quad (3)$$

$$p, x \geq 0, \quad x_{c,\omega} = 0 \quad \forall c, \omega \notin \Omega(c). \quad (4)$$

(1) encodes that an allocation is chosen for every $c \in \bar{\mathcal{D}}$, and (2) and (3) encode the support-based-IC and support-based-IR conditions respectively. Lemma 2.3 ensures that (P) correctly encodes PayM, so that $OPT := OPT_{\text{P}}$ is a lower bound on the expected disutility of an optimal mechanism.

Our results are obtained by computing an optimal solution to (P), or a further relaxation of it, and translating this to a near-optimal support-based-(IC-in-expectation, IR) mechanism. Both steps come with their own challenges. Except in very simple settings (such as single-item procurement auctions), $|\Omega|$ is typically exponential in the input size, and therefore it is not clear how to solve (P) efficiently. We therefore consider the dual LP (D), which has variables γ_c , $y_{i,(c_i,c_{-i}),c'_i}$ and $\beta_{i,(c_i,c_{-i})}$ corresponding to (1), (2) and (3) respectively.

$$\max \quad \sum_c \gamma_c \quad (\text{D})$$

$$\text{s.t.} \quad \sum_{i:c \in \bar{\mathcal{D}}_i \times \mathcal{D}_{-i}} \left(\sum_{c'_i \in \bar{\mathcal{D}}_i} (c_i(\omega) y_{i,(c_i,c_{-i}),c'_i} - c'_i(\omega) y_{i,(c'_i,c_{-i}),c_i}) + c_i(\omega) \beta_{i,c} \right) + \kappa \cdot \Pr_{\mathcal{D}}(c) \text{pub}(\omega) \geq \gamma_c \quad \forall c \in \bar{\mathcal{D}}, \omega \in \Omega(c) \quad (5)$$

$$\sum_{c'_i \in \bar{\mathcal{D}}_i} (y_{i,(c_i,c_{-i}),c'_i} - y_{i,(c'_i,c_{-i}),c_i}) + \beta_{i,c_i,c_{-i}} \leq \Pr_{\mathcal{D}}(c) \quad \forall i, c_i \in \bar{\mathcal{D}}_i, c_{-i} \in \mathcal{D}_{-i} \quad (6)$$

$$y, \beta \geq 0. \quad (7)$$

With additive types, one can encode the LHS of (5) as $\sum_i \tilde{c}_i(\omega)$ for a suitably-defined additive type (depending on c) $\tilde{c}_i = (\tilde{c}_{i,v})_{v \in \mathcal{T}_i}$. Thus, the separation problem for constraints (5) amounts to determining if the optimal value of the CM problem defined by a certain additive type profile, with possibly negative values, is at least γ_c . Hence, an optimal algorithm for the CM problem can be used to solve (D), and hence, (P), efficiently.

Theorem 3.1 *With additive types, one can efficiently solve (P) given an optimal algorithm for the CM problem.*

Proof: Let \mathcal{A} be an optimal algorithm for the CM problem. Note that \mathcal{A} is only required to work with nonnegative inputs. We first observe that we can use \mathcal{A} to find a solution that minimizes $\sum_i c_i(\omega) + \kappa \cdot \text{pub}(\omega)$ for any $\kappa \geq 0$, even for an input $c = \{c_{i,v}\}_{i,v \in \mathcal{T}_i}$ where some of the $c_{i,v}$ s are negative. Let $A_i = \{v \in \mathcal{T}_i : c_{i,v} < 0\}$. Clearly, if ω^* is an optimal solution, then $A_i \subseteq \omega_i^*$ (since $\text{pub}(\cdot)$ does not increase upon adding covering objects). Define $c_{i,v}^+ := \max(0, c_{i,v})$ and $c_i^+ := \{c_{i,v}^+\}_{v \in \mathcal{T}_i}$.

Let $\Gamma = \frac{1}{\kappa}$ if $\kappa > 0$; otherwise let $\Gamma = NU$, where U is a strict upper bound on $\max_{\omega \in \Omega} \text{pub}(\omega)$ and N is an integer such that all the $c_{i,v}^+$ s are integer multiples of $\frac{1}{N}$. Note that for any $\omega, \omega' \in \Omega$,

if $\sum_i c_i^+(\omega) - \sum_i c_i^+(\omega')$ is non-zero, then its absolute value is at least $\frac{1}{N}$. Also, U and N may be efficiently computed (for rational data) and $\log(NU)$ is polynomially bounded. Let (S_1, \dots, S_n) be the solution returned by \mathcal{A} for the CM problem on the input where all the $c_{i,v}^+$ s are scaled by Γ . The choice of Γ ensures that

$$\sum_i c_i^+(S_i) + \kappa \cdot \text{pub}(S_1, \dots, S_n) \leq \sum_i c_i^+(\omega_i^*) + \kappa \cdot \text{pub}(\omega^*) = \sum_i (c_i(\omega_i^*) - c_i(A_i)) + \kappa \cdot \text{pub}(\omega^*).$$

(The first inequality clearly holds if $\kappa > 0$. If $\kappa = 0$ and $\sum_i c_i^+(S_i) > \sum_i c_i^+(\omega_i^*)$, then we have that $\Gamma \sum_i c_i^+(S_i) \geq \Gamma \sum_i c_i^+(\omega_i^*) + \frac{\Gamma}{N} > \Gamma \sum_i c_i^+(\omega_i^*) + \text{pub}(\omega^*)$, which contradicts the optimality of (S_1, \dots, S_n) for the input $\{\Gamma c_{i,v}^+\}_{i,v \in \mathcal{T}_i}$.) So setting $\omega'_i = A_i \cup S_i$ for every i yields a feasible solution such that $\sum_i c_i(\omega') + \kappa \cdot \text{pub}(\omega') \leq \sum_i c_i(\omega^*) + \kappa \cdot \text{pub}(\omega^*)$; hence ω' is an optimal solution.

Given a dual solution (y, β, γ) , we can easily check if (6), (7) hold. Fix $c \in \bar{\mathcal{D}}$ and player i . Since we have additive types, if we define $\theta_{i,v}^c = \sum_{c'_i \in \bar{\mathcal{D}}_i} (c_{i,v} y_{i,(c_i, c_{-i}), c'_i} - c'_{i,v} y_{i,(c'_i, c_{-i}), c_i}) + c_{i,v} \beta_{i,c}$, then for every $\omega \in \Omega$, we can equate $\sum_{c'_i \in \bar{\mathcal{D}}_i} (c_i(\omega) y_{i,(c_i, c_{-i}), c'_i} - c'_i(\omega) y_{i,(c'_i, c_{-i}), c_i}) + c_i(\omega) \beta_{i,c}$ with $\theta_{i,v}^c(\omega) := \sum_{v \in \omega_i} \theta_{i,v}^c$.

Let $I = \{i : c \in \bar{\mathcal{D}}_i \times \mathcal{D}_{-i}\}$. Constraints (5) for c can then be written as $\min_{\omega \in \Omega(c)} (\sum_{i \in I} \theta^c(\omega) + \kappa \Pr_{\mathcal{D}}(c) \text{pub}(\omega)) \geq \gamma_c$. Define \tilde{c} as follows:

$$\tilde{c}_{i,v} = \begin{cases} \gamma_c + 1 & \text{if } c_i = m_i \mathbf{1}_{\mathcal{T}_i}, \\ \theta_{i,v}^c & \text{if } i \in I, c_i \in \mathcal{D}_i, \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to see that (5) holds for c iff $\min_{\omega \in \Omega} (\sum_i \tilde{c}_i(\omega) + \kappa \cdot \Pr_{\mathcal{D}}(c) \text{pub}(\omega))$ —which can be computed using \mathcal{A} —is at least γ_c . Thus, we can use the ellipsoid method to solve (D). This also yields a compact dual consisting of constraints (6), (7) and the polynomially-many (5) constraints that were returned by the separation oracle during the execution of the ellipsoid method, whose optimal value is $OPT_{\mathcal{D}}$. The dual of this compact dual is an LP of the same form as (P) but with polynomially many $x_{c,\omega}$ -variables; solving this yields an optimal solution to (P). ■

Complementing Theorem 3.1, we argue that a feasible solution (x, p) to (P) can be extended to a support-based-(IC-in-expectation, IR) mechanism having expected disutility at most the value of (x, p) (Theorem 3.2). Combining this with Theorem 3.1 yields the corollary that an optimal algorithm for the CM problem can be used to obtain an optimal mechanism for the PayM problem (Corollary 3.3).

Theorem 3.2 *We can extend a feasible solution (x, p) to (P) to a support-based-(IC-in-expectation, IR) mechanism with expected disutility $\sum_c \Pr_{\mathcal{D}}(c) (\sum_i p_{i,c} + \kappa \sum_{\omega} x_{c,\omega} \text{pub}(\omega))$.*

Proof: Let $\Omega' = \{\omega : x_{c,\omega} > 0 \text{ for some } c \in \bar{\mathcal{D}}\}$. We use x_c to denote the vector $\{x_{c,\omega}\}_{\omega \in \Omega'}$. Consider a player i , $c_{-i} \in \mathcal{D}_{-i}$, and $c_i, c'_i \in \bar{\mathcal{D}}_i$. Note that (2) implies that if $x_{(c_i, c_{-i})} = x_{(c'_i, c_{-i})}$, then $p_{i,(c_i, c_{-i})} = p_{i,(c'_i, c_{-i})}$. For $c_{-i} \in \mathcal{D}_{-i}$, define $R(i, c_{-i}) = \{x_{(c_i, c_{-i})} : (c_i, c_{-i}) \in \bar{\mathcal{D}}\}$, and for $y = x_{(c_i, c_{-i})} \in R(i, c_{-i})$ define $p_{i,y}$ to be $p_{i,(c_i, c_{-i})}$ (which is well defined by the above argument).

We now define the randomized mechanism $M = (\mathcal{A}, \{q_i\})$, where $\mathcal{A}(c)$ and $q_i(c)$ denote respectively the probability distribution over allocations and the expected payment to player i , on input c . We sometimes view $\mathcal{A}(c)$ equivalently as the random variable specifying the allocation chosen for input c . Fix an allocation $\omega_0 \in \Omega$. Consider an input c . If $c \in \bar{\mathcal{D}}$, we set $\mathcal{A}(c) = x_c$, and $q_i(c) = p_{i,c}$ for all i . So consider $c \notin \bar{\mathcal{D}}$. If there is no i such that $c_{-i} \in \mathcal{D}_{-i}$, we simply set $\mathcal{A}(c) = \omega_0$, $q_i(c) = c_i(\omega_0)$

for all i ; such a c does not figure in the support-based (IC, IR) conditions. Otherwise there is a unique i such that $c_{-i} \in \mathcal{D}_{-i}$, $c_i \in C_i \setminus \overline{\mathcal{D}}_i$. Set $\mathcal{A}(c) = \arg \max_{y \in R(i, c_{-i})} (p_{i,y} - \sum_{\omega \in \Omega'} c_i(\omega) y_\omega)$ and $q_j(c) = p_{j, \mathcal{A}(c)}$ for all players j . Note that (c_i, c_{-i}) figures in (2) *only* for player i . Crucially, note that since $y = x_{(m_i, c_{-i})} \in R(i, c_{-i})$ and $\sum_{\omega \in \Omega} c_i(\omega) y_\omega = 0$ by definition, we always have $q_i(c) - \mathbb{E}_{\mathcal{A}}[c_i(\mathcal{A}(c))] \geq 0$. Thus, by definition, and by (2), we have ensured that M is support-based (IC, IR)-in-expectation and its expected disutility is exactly the value of (x, p) . This can be modified so that IR holds with probability 1.

The above procedure is efficient if $\sum_{\omega \in \Omega'} c_i(\omega) x_{c, \omega}$ can be calculated efficiently. This is clearly true if $|\Omega'|$ is polynomially bounded, but it could hold under weaker conditions as well. ■

Corollary 3.3 *Given an optimal algorithm for the CM problem, we can efficiently obtain an optimal support-based-(IC-in-expectation, IR) mechanism for the PayM problem in multidimensional settings with additive types.*

Using approximation algorithms. The CM problem is however often *NP*-hard (e.g., for vertex cover), and we would like to be able to exploit approximation algorithms for the CM problem to obtain near-optimal mechanisms. The usual approach is to use an approximation algorithm to “approximately” separate over constraints (5).⁴ However, this does not work here since the CM problem that one needs to solve in the separation problem involves negative costs, which renders the usual notion of approximation meaningless. Instead, if the CM problem admits a certain type of LP-relaxation (C-P), then we argue that one can solve a relaxation of (P) where the allocation-set is the set of extreme points of (C-P) (Theorem 3.4). For single-dimensional problems (Section 4), we leverage this to obtain strong and far-reaching results. We show that a ρ -approximation algorithm relative to (C-P) can be used to “round” the optimal solution to this relaxation to a support-based-(IC-in-expectation, IR)-mechanism losing a ρ -factor in the disutility (Theorem 4.2). Thus, we obtain near-optimal mechanisms for a variety of single-dimensional problems.

Suppose that the CM problem admits an LP-relaxation of the following form, where $c = \{c_{i,v}\}_{i \in [n], v \in \mathcal{T}_i}$ is the input type-profile.

$$\min \quad c^T x + d^T z \quad \text{s.t.} \quad Ax + Bz \geq b, \quad x, z \geq 0. \quad (\text{C-P})$$

Intuitively x encodes the allocation chosen, and $d^T z$ encodes $\text{pub}(\cdot)$. For $x \geq 0$, define $z(x) := \arg \min\{d^T z : (x, z) \text{ is feasible to (C-P)}\}$; if there is no z such that (x, z) is feasible to (C-P), set $z(x) := \perp$. Define $\Omega_{\text{LP}} := \{x : z(x) \neq \perp, \quad 0 \leq x_{i,v} \leq 1 \quad \forall i, v \in \mathcal{T}_i\}$. We require that: (a) a $\{0, 1\}$ -vector x is in Ω_{LP} iff it is the characteristic vector of an allocation $\omega \in \Omega$, and in this case, we have $d^T z(x) = \text{pub}(\omega)$; (b) $A \geq 0$; (c) for any input $c \geq 0$ to the covering problem, (C-P) is not unbounded, and if it has an optimal solution, it has one where $x \in \Omega_{\text{LP}}$; (d) for any c , we can efficiently find an optimal solution to (C-P) or detect that it is unbounded or infeasible.

We extend the type c_i of each player i and pub to assign values also to points in Ω_{LP} : define $c_i(x) = \sum_{v \in \mathcal{T}_i} c_{i,v} x_{i,v}$ and $\text{pub}(x) = d^T z(x)$ for $x \in \Omega_{\text{LP}}$. Let Ω_{ext} denote the finite set of extreme points of Ω_{LP} . Condition (a) ensures that Ω_{ext} contains the characteristic vectors of all feasible allocations. Let (P') denote the relaxation of (P), where we replace the set of feasible allocations Ω with Ω_{ext} (so ω indexes Ω_{ext} now), and for $c \in \overline{\mathcal{D}}$ with $c_i = m_i \mathbf{1}(\mathcal{T}_i)$, we now define $\Omega(c) :=$

⁴For revenue-maximization problems in packing domains, this simple approach *does* indeed work in single-dimensional settings and settings with additive types. This is because the dual separation problem is now an SWM problem, and it is easy to use a ρ -approximation algorithm for the SWM problem for nonnegative inputs to obtain a ρ -approximate solution with *arbitrary*, positive or negative, inputs. This yields a simple extension of some of the results in [8]; see Appendix C.1.

$\{\omega \in \Omega_{\text{ext}} : \omega_{i,v} = 0 \forall v \in \mathcal{T}_i\}$. Since one can optimize efficiently over Ω_{LP} , and hence Ω_{ext} , even for negative type-profiles, we have the following.

Theorem 3.4 *We can efficiently compute an optimal solution to (P').*

4 Single-dimensional problems

Corollary 3.3 immediately yields results for certain single-dimensional problems (see Table 1), most notably, *single-item procurement auctions*. We now substantially expand the scope of PayM problems for which one can obtain near-optimal mechanisms by showing how to leverage ‘‘LP-relative’’ approximation algorithms for the CM problem. (As noted earlier, and sketched in Appendix C.1, a simpler approach can be used to leverage approximation algorithms for revenue-maximization in packing domains.) Suppose that the CM problem can be encoded as (C-P). An *LP-relative ρ -approximation algorithm* for the CM problem is a polytime algorithm that for any input $c \geq 0$ to the covering problem, returns a $\{0, 1\}$ -vector $x \in \Omega_{\text{LP}}$ such that $c^T x + d^T z(x) \leq \rho \text{OPT}_{\text{C-P}}$. Using the convex-decomposition procedure in [12] (see Section 5.1 of [12]), one can show the following; the proof appears at the end of this section.

Lemma 4.1 *Let $x \in \Omega_{\text{LP}}$. Given an LP-relative ρ -approximation algorithm \mathcal{A} for the CM problem, one can efficiently obtain $(\lambda^{(1)}, x^{(1)}), \dots, (\lambda^{(k)}, x^{(k)})$, where $\sum_{\ell} \lambda^{(\ell)} = 1, \lambda \geq 0$, and $x^{(\ell)}$ is a $\{0, 1\}$ -vector in Ω_{LP} for all ℓ , such that $\sum_{\ell} \lambda^{(\ell)} x_{i,v}^{(\ell)} = \min(\rho x_{i,v}, 1)$ for all $i, v \in \mathcal{T}_i$, and $\sum_{\ell} \lambda^{(\ell)} d^T z(x^{(\ell)}) \leq \rho d^T z(x)$.*

Theorem 4.2 *Given an LP-relative ρ -approximation algorithm for the CM problem, one can obtain a polytime ρ -approximation support-based-(IC-in-expectation, IR) mechanism for the PayM problem.*

Proof: We solve (P') to obtain an optimal solution (x, p) . Since $|\mathcal{T}_i| = 1$ for all i , it will be convenient to view $\omega \in \Omega_{\text{LP}}$ as a vector $\{\omega_i\}_{i \in [n]} \in [0, 1]^n$, where $\omega_i \equiv \omega_{i,v}$ for the single covering object $v \in \mathcal{T}_i$. Fix $c \in \overline{\mathcal{D}}$. Define $y_c = \sum_{\omega \in \Omega_{\text{ext}}} x_{c,\omega} \omega$ (which can be efficiently computed since x has polynomial support). Then, $\sum_{\omega \in \Omega_{\text{ext}}} c_i(\omega) x_{c,\omega} = c_i y_{c,i}$ and $d^T z(y) \leq \sum_{\omega \in \Omega_{\text{ext}}} \text{pub}(\omega) x_{c,\omega}$. By Lemma 4.1, we can efficiently find a point $\tilde{y}_c = \sum_{\omega \in \Omega} \tilde{x}_{c,\omega} \omega$, where $\tilde{x}_c \geq 0, \sum_{\omega \in \Omega} \tilde{x}_{c,\omega} = 1$, in the convex hull of the $\{0, 1\}$ -vectors in Ω_{LP} such that $\tilde{y}_{c,i} = \min(\rho y_{c,i}, 1)$ for all i , and $\sum_{\omega \in \Omega} \tilde{x}_{c,\omega} \text{pub}(\omega) \leq \rho d^T z(y)$.

We now argue that one can obtain payments $\{q_{i,c}\}$ such that (\tilde{x}, q) is feasible to (P) and $q_{i,c} \leq \rho p_{i,c}$ for all $i, c \in \overline{\mathcal{D}}$. Thus, the value of (\tilde{x}, q) is at most ρ times the value of (x, p) . Applying Theorem 3.2 to (\tilde{x}, q) yields the desired result.

Fix i and $c_{-i} \in \mathcal{D}_{-i}$. Constraints (4) and (2) ensure that $y_{(m_i, c_{-i}), i} = 0$, and $y_{(c_i, c_{-i}), i} \geq y_{(c'_i, c_{-i}), i}$ for all $c_i, c'_i \in \overline{\mathcal{D}}_i$ s.t. $c_i < c'_i$. Hence, $\tilde{y}_{(m_i, c_{-i}), i} = 0$, $\tilde{y}_{(c_i, c_{-i}), i} \geq \tilde{y}_{(c'_i, c_{-i}), i}$ for $c_i, c'_i \in \overline{\mathcal{D}}_i, c_i > c'_i$. Define $q_{i, (m_i, c_{-i})} = 0$. Let $0 \leq c_i^1 < c_i^2 < \dots < c_i^{k_i}$ be the values in \mathcal{D}_i . For $c_i = c_i^\ell$, define

$$q_{i, (c_i, c_{-i})} = c_i \tilde{y}_{(c_i, c_{-i}), i} + \sum_{t=\ell+1}^{k_i} (c_i^t - c_i^{t-1}) \tilde{y}_{(c_i^t, c_{-i}), i}.$$

Since $\sum_{\omega \in \Omega} c_i(\omega) \tilde{x}_{(c_i, c_{-i}), \omega} = c_i \tilde{y}_{(c_i, c_{-i}), i}$, (3) holds. By construction, for consecutive values $c_i = c_i^\ell, c_i = c_i^{\ell+1}$, we have $q_{i, (c_i, c_{-i})} - q_{i, (c'_i, c_{-i})} = c_i (\tilde{y}_{(c_i, c_{-i}), i} - \tilde{y}_{(c'_i, c_{-i}), i})$, which is at most

$$\rho \cdot c_i (y_{(c_i, c_{-i}), i} - y_{(c'_i, c_{-i}), i}) \leq \rho (p_{i, (c_i, c_{-i})} - p_{i, (c'_i, c_{-i})}).$$

Since $q_{i,(m_i,c_{-i})} = 0 \leq \rho p_{i,(m_i,c_{-i})}$, this implies that $q_{i,(c_i,c_{-i})} \leq \rho p_{i,(c_i,c_{-i})}$. Finally, it is easy to verify that for any $c_i, c'_i \in \mathcal{D}_i$, we have $q_{i,(c_i,c_{-i})} - q_{i,(c'_i,c_{-i})} \geq c_i(\tilde{y}_{(c_i,c_{-i}),i} - \tilde{y}_{(c'_i,c_{-i}),i})$, so (\tilde{x}, q) satisfies (2). \blacksquare

Corollary 3.3 and Theorem 4.2 yield polytime near-optimal mechanisms for a host of single-dimensional PayM problems. Table 1 summarizes a few applications. Even for single-item procurement auctions, these are the *first* results for PayM problems with correlated players satisfying a notion stronger than (BIC, interim IR).

Problem	Approximation	Due to
<i>Single-item procurement auction</i> : buy one item provided by n players	1	Corollary 3.3
<i>Metric UFL</i> : players are facilities, output should be a UFL solution	1.488 using [15]	Theorem 4.2
<i>Vertex cover</i> : players are nodes, output should be a vertex cover	2	Theorem 4.2
<i>Set cover</i> : players are sets, output should be a set cover	$O(\log n)$	Theorem 4.2
<i>Steiner forest</i> : players are edges, output should be a Steiner forest	2	Theorem 4.2
<i>Multicut</i> (a), <i>Multicut</i> (b): players are edges, output should be a multicut in (a), or a multicut in (b)	2 for (a) $O(\log n)$ for (b)	Theorem 4.2

Table 1: Results for some representative single-dimensional PayM problems.

Proof of Lemma 4.1: It suffices to show that the LP (Q) can be solved in polytime and its optimal value is 1. Throughout, we use ℓ to index $\{0, 1\}$ vectors in Ω_{LP} . (Recall that these correspond to feasible allocations.)

$$\begin{array}{l|l}
\max & \sum_{\ell} \lambda^{(\ell)} & \text{(Q)} \\
\text{s.t.} & \sum_{\ell} \lambda^{(\ell)} x_{i,v}^{(\ell)} = \min(\rho x_{i,v}, 1) \quad \forall i, v \in \mathcal{T}_i & (8) \\
& \sum_{\ell} \lambda^{(\ell)} d^T(z(x^{(\ell)})) \leq \rho d^T z(x) & (9) \\
& \sum_{\ell} \lambda^{(\ell)} \leq 1 & (10) \\
& \lambda \geq 0. &
\end{array}
\quad \left| \quad \begin{array}{l}
\min & \sum_{i,v \in \mathcal{T}_i} \min(\rho x_{i,v}, 1) \alpha_{i,v} + \rho d^T z(x) \cdot \beta + \theta & \text{(R)} \\
\text{s.t.} & \sum_{i,v \in \mathcal{T}_i} x_{i,v}^{(\ell)} \alpha_{i,v} + d^T(z(x^{(\ell)})) \beta + \theta \geq 1 \quad \forall \ell & (11) \\
& \beta, \theta \geq 0. &
\end{array}$$

Here the $\alpha_{\ell s}$, β and θ are the dual variables corresponding to constraints (8), (9), and (10) respectively. Clearly, $OPT_{\text{(R)}} \leq 1$ since $\theta = 1$, $\alpha_{i,v} = 0 = \beta$ for all i, v is a feasible dual solution.

Suppose $(\hat{\alpha}, \hat{\beta}, \hat{\theta})$ is a feasible dual solution of value less than 1. Set $\tilde{\alpha}_{i,v} = \hat{\alpha}_{i,v}$ if $\hat{\alpha}_{i,v} \geq 0$ and $\rho x_{i,v} \leq 1$, and $\tilde{\alpha}_{i,v} = 0$ otherwise. Let $\Gamma = \frac{1}{\hat{\beta}}$ if $\hat{\beta} > 0$ and equal to $2N d^T z$ otherwise, where N is such that for all $\{0, 1\}$ -vectors $x^{(\ell)} \in \Omega_{\text{LP}}$, we have that $c^T x^{(\ell)} > c^T x$ implies $c^T x^{(\ell)} \geq c^T x + \frac{1}{N}$. Note that we can choose N so that its size is $\text{poly}(\mathcal{I}, \text{size of } x)$. Consider the CM problem defined by the input $\Gamma \hat{\alpha}$. Running \mathcal{A} on this input, we obtain a $\{0, 1\}$ -vector $x^{(\ell)} \in \Omega_{\text{LP}}$ whose total cost is at most ρ times the cost of the fractional solution $(x, z(x))$. This translates to

$$\sum_{i,v} x_{i,v}^{(\ell)} \tilde{\alpha}_{i,v} + d^T(z(x^{(\ell)})) \hat{\beta} \leq \rho \left(\sum_{i,v} x_{i,v} \tilde{\alpha}_{i,v} + d^T z(x) \cdot \hat{\beta} \right). \quad (12)$$

Now augment $x^{(\ell)}$ to the following $\{0,1\}$ -vector \hat{x} : set $\hat{x}_{i,v} = 1$ if $\rho x_{i,v} > 1$ or $\hat{\alpha}_{i,v} < 0$, and $x_{i,v}^{(\ell)}$ otherwise. Then \hat{x} is the characteristic vector of a feasible allocation, since we have only added covering objects to the allocation corresponding to $x^{(\ell)}$; hence $\hat{x} \in \Omega_{\text{LP}}$. We have $d^T z(\hat{x}) = \text{pub}(\hat{x}) \leq \text{pub}(x^{(\ell)}) = d^T(z(x^{(\ell)}))$ and

$$\sum_{i,v} \hat{x}_{i,v} \hat{\alpha}_{i,v} = \sum_{\substack{i,v:\rho x_{i,v}>1 \\ \text{or } \hat{\alpha}_{i,v}<0}} \hat{\alpha}_{i,v} + \sum_{i,v} x_{i,v}^{(\ell)} \tilde{\alpha}_{i,v} \leq \sum_{\substack{i,v:\rho x_{i,v}>1 \\ \text{or } \hat{\alpha}_{i,v}<0}} \min(\rho x_{i,v}, 1) \hat{\alpha}_{i,v} + \sum_{i,v} x_{i,v}^{(\ell)} \tilde{\alpha}_{i,v}.$$

Combined with (12), this shows that

$$\begin{aligned} \sum_{i,v} \hat{x}_{i,v} \hat{\alpha}_{i,v} + d^T z(\hat{x}) \hat{\beta} &\leq \sum_{\substack{i,v:\rho x_{i,v}>1 \\ \text{or } \hat{\alpha}_{i,v}<0}} \min(\rho x_{i,v}, 1) \hat{\alpha}_{i,v} + \sum_{i,v:\tilde{\alpha}_{i,v}>0} \rho x_{i,v} \tilde{\alpha}_{i,v} + \rho d^T z(x) \cdot \hat{\beta} \\ &= \sum_{i,v} \min(\rho x_{i,v}, 1) \hat{\alpha}_{i,v} + \rho d^T z(x) \cdot \hat{\beta} < 1 - \hat{\theta} \end{aligned}$$

which contradicts that $(\hat{\alpha}, \hat{\beta}, \hat{\theta})$ is feasible to (R). Hence, $OPT_{(\text{Q})} = OPT_{(\text{R})} = 1$.

Thus, we can add the constraint $\sum_{i,v \in \mathcal{T}_i} \min(\rho x_{i,v}, 1) \alpha_{i,v} + \rho d^T z(x) \cdot \beta + \theta \leq 1$ to (R) without altering anything. If we solve the resulting LP using the ellipsoid method, and take the inequalities corresponding to the violated inequalities (11) found by \mathcal{A} during the ellipsoid method, then we obtain a compact LP with only a polynomial number of constraints that is equivalent to (R). The dual of this compact LP yields an LP equivalent to (Q) with a polynomial number of $\lambda^{(\ell)}$ variables which we can solve to obtain the desired convex decomposition. \blacksquare

5 Multidimensional problems

We obtain results for multidimensional PayM problems via two distinct approaches. One is by directly applying Corollary 3.3 (e.g., Theorem 5.1). The other approach is based on again moving to an LP-relaxation of the CM problem and utilizing Theorem 3.4 in conjunction with a *stronger* LP-rounding approach. This yields results for multidimensional (metric) UFL and its variants (Theorem 5.3).

Multi-item procurement auctions. Here, we have n sellers and k (heterogeneous) items. Each seller i has a *supply vector* $s_i \in \mathbb{Z}_+^k$ denoting his supply for the various items, and the buyer has a *demand vector* $d \in \mathbb{Z}_+^k$ specifying his demand for the various items. This is public knowledge. Each seller i has a *private cost-vector* $c_i \in \mathbb{R}_+^k$, where $c_{i,\ell}$ is the cost he incurs for supplying *one unit* of item ℓ . A feasible solution is an allocation specifying how many units of each item each seller supplies to the buyer such that for each item ℓ , each seller i provides at most $s_{i,\ell}$ units of ℓ and the buyer obtains d_ℓ total units of ℓ . The corresponding CM problem is a min-cost flow problem (in a bipartite graph), which can be efficiently solved optimally, thus we obtain a polytime optimal mechanism.

Theorem 5.1 *There is a polytime optimal support-based-(IC-in-expectation, IR) mechanism for multi-unit procurement auctions.*

Multidimensional budgeted (metric) uncapacitated facility location (UFL). Here, we have a set \mathcal{E} of clients that need to be serviced by facilities, and a set \mathcal{F} of locations where facilities may be opened. Each player i may provide facilities at the locations in $\mathcal{T}_i \subseteq \mathcal{F}$. We may assume that the \mathcal{T}_i s are disjoint. For each facility $\ell \in \mathcal{T}_i$ that is opened, i incurs a private opening cost $f_\ell \equiv f_{i,\ell}$, and assigning client j to an open facility ℓ incurs a publicly-known assignment cost $d_{\ell j}$, where the $d_{\ell j}$ s form a metric. We are also given a public assignment-cost budget B . The goal in **Budget-UFL** is to open a subset $F \subseteq \mathcal{F}$ of facilities and assign each client j to an open facility $\sigma(j) \in F$ so as to minimize $\sum_{\ell \in F} f_\ell + \sum_{j \in \mathcal{E}} d_{\sigma(j)j}$ subject to $\sum_{j \in \mathcal{E}} d_{\sigma(j)j} \leq B$; **UFL** is the special case where $B = \infty$. We can define $\text{pub}(\mathcal{T}_1, \dots, \mathcal{T}_n)$ to be the total assignment cost if this is at most B , and ∞ otherwise.

Let O^* denote the expected disutility of an optimal mechanism for **Budget-UFL**. We obtain a mechanism with expected disutility at most $2O^*$ that always returns a solution with expected assignment cost at most $2B$. Consider the following LP-relaxation for **Budget-UFL**.

$$\begin{aligned} \min \quad & \sum_{\ell \in \mathcal{F}} f_\ell x_\ell + \sum_{j \in \mathcal{E}, \ell \in \mathcal{F}} d_{\ell j} z_{\ell j} \quad \text{s.t.} & \text{(BFL-P)} \\ & \sum_{j \in \mathcal{E}, \ell \in \mathcal{F}} d_{\ell j} z_{\ell j} \leq B, \quad \sum_{\ell \in \mathcal{F}} z_{\ell j} \geq 1 \quad \forall j \in \mathcal{E}, \quad 0 \leq z_{\ell j} \leq x_\ell \quad \forall \ell \in \mathcal{F}, j \in \mathcal{E}. \end{aligned}$$

Let (FL-P) denote (BFL-P) with $B = \infty$, and $OPT_{\text{FL-P}}$ denote its optimal value. We say that an algorithm \mathcal{A} is a *Lagrangian multiplier preserving* (LMP) ρ -approximation algorithm for **UFL** if for every instance, it returns a solution (F, σ) such that $\rho \sum_{\ell \in F} f_\ell + \sum_{j \in \mathcal{E}} d_{\sigma(j)j} \leq \rho \cdot OPT_{\text{FL-P}}$. In [17], it is shown that given such an algorithm \mathcal{A} , one can take any solution (x, z) to (FL-P) and obtain a convex combination of *UFL solutions* $(\lambda^{(1)}; F^{(1)}, \sigma^{(1)}), \dots, (\lambda^{(k)}; F^{(k)}, \sigma^{(k)})$ —so $\lambda \geq 0$, $\sum_r \lambda^{(r)} = 1$ —such that $\sum_{r: \ell \in F^{(r)}} \lambda^{(r)} = x_\ell$ for all ℓ and $\sum_r \lambda^{(r)} (\sum_j d_{\sigma^{(r)}(j)j}) \leq \rho \sum_{j, \ell} d_{\ell j} z_{\ell j}$. An LMP 2-approximation algorithm for **UFL** is known [10].

Lemma 5.2 *Given an LMP ρ -approximation algorithm for UFL, one can design a polytime support-based-(IC-in-expectation, IR) mechanism for Budget-UFL whose expected disutility is at most ρO^* while violating the budget by at most a ρ -factor.*

Proof: The LP-relaxation (BFL-P) for the CM problem is of the form (C-P) and satisfies the required properties. Recall that for $x \geq 0$, $z(x)$ denotes the min-cost completion of x to a feasible solution to (BFL-P) if one exists, and is \perp if there is no such completion of x . Let $\Omega_{\text{LP}} := \{x : z(x) \neq \perp, 0 \leq x_\ell \leq 1 \forall \ell\}$. For integral $\omega \in \Omega_{\text{LP}}$, $z(\omega)$ specifies the assignment where each client j is assigned to the nearest open facility. By Theorem 3.4, one can efficiently compute an optimal solution (X, p) to the relaxation of (P) where the set of feasible allocations is the set Ω_{ext} of extreme points of Ω_{LP} .

We round (X, p) to a feasible solution to (P) by proceeding as in the proof of Theorem 4.2. Let Ω_{UFL} be the set of characteristic vectors of open facilities of all integral **UFL** solutions. We use ℓ to index facilities in \mathcal{F} and j to index clients in \mathcal{E} . Fix $c \in \overline{\mathcal{D}}$. Define $y_c = \sum_{\omega \in \Omega_{\text{ext}}} X_{c,\omega} \omega$, so $\sum_{\omega \in \Omega_{\text{ext}}} c_i(\omega) X_{c,\omega} = \sum_{\ell \in \mathcal{T}_i} f_\ell y_{c,\ell}$. Let $z_c = \sum_{\omega \in \Omega_{\text{ext}}} X_{c,\omega} z(\omega)$, so $\sum_{j, \ell} z_{c,\ell j} d_{\ell j} \leq B$. We use the LMP ρ -approximation algorithm to express y_c as a convex combination $\sum_{\omega \in \Omega_{\text{UFL}}} \tilde{x}_{c,\omega} \omega$ of (integral) **UFL**-solutions such that the expected assignment cost $\sum_{\omega \in \Omega_{\text{UFL}}} \tilde{x}_{c,\omega} \sum_{j, \ell} z(\omega)_{\ell j} d_{\ell j}$ is at most $\rho \sum_{j, \ell} d_{\ell j} z_{c,\ell j} \leq \rho B$. Hence, (\tilde{x}, p) is a feasible solution to (P). Theorem 3.2 now yields the desired mechanism. \blacksquare

Theorem 5.3 *There is a polytime support-based-(IC-in-expectation, IR) mechanism for Budget-UFL with expected disutility at most $2O^*$, which always returns a solution with expected assignment cost at most $2B$.*

6 Extensions: alternative solution concepts

We now investigate the PayM problem under various alternative solution concepts. In Section 6.1, we consider solution concepts weaker than support-based (IC, IR), but yet robust enough to ensure that truthful participation is an ex-post no-regret choice for every player at every type profile in the support of the underlying distribution. We show that *all* our guarantees extend readily to these solution concepts. (Note that a weaker solution concept does not necessarily mean that the corresponding PayM mechanism-design problem is a simpler problem; a weaker solution concept enlarges the space of allowed mechanisms, which could make it more- or less- difficult to search for an optimal mechanism.) In Section 6.2, we consider the stronger solution concept of (DSIC (in-expectation, IR)), and obtain results for single-dimensional settings but at the expense of increasing the running time to exponential in the number of players.

6.1 Solution concepts weaker than support-based (IC, IR)

Consider the following weakenings of support-based (IC, IR) (Definition 2.2).

For every player i ,

$$u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq u_i(c_i, c_{-i}; \bar{c}_i), \quad u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq 0, \quad \text{for all } (\bar{c}_i, c_{-i}) \in \mathcal{D}, \quad c_i : (c_i, c_{-i}) \in \mathcal{D} \quad (13)$$

$$u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq u_i(c_i, c_{-i}; \bar{c}_i), \quad u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq 0, \quad \text{for all } (\bar{c}_i, c_{-i}) \in \mathcal{D}, \quad c_i \in C_i \quad (14)$$

$$u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq u_i(c_i, c_{-i}; \bar{c}_i), \quad u_i(\bar{c}_i, c_{-i}; \bar{c}_i) \geq 0, \quad \text{for all } \bar{c}_i, c_i \in \mathcal{D}_i, \quad c_{-i} \in \mathcal{D}_{-i} \quad (15)$$

All three solution concepts, (13)–(15), ensure that truthful participation is in the best interest of every player i at every type-profile in \mathcal{D} even at the ex-post stage when he knows the realized types of all players, but for varying choices of lies: in (14), the lie could be anything, so a mechanism satisfying (14) is (BIC, interim IR) for every distribution whose support is a subset of \mathcal{D} ; in (15), the “best interest” is among lies consistent with i ’s support; and in (13), the “best interest” is among lies consistent with the support of the distribution.

We now argue that our results extend to these notions. For notions (13) and (15), one can simply incorporate all the IC and IR constraints in the LP. Note that there are $O(n|\mathcal{D}|^2)$ such constraints under (13), and $O(n|\mathcal{D}|^3)$ constraints under (15), so the size of the resulting LP is $\text{poly}(n, |\mathcal{D}|)$. Theorem 3.1 continues to hold for the resulting LP, due to the same arguments. For both notions, an LP solution immediately yields a randomized mechanism satisfying that notion, except that utility is replaced by expected utility; for type profiles not included in the LP, we may output any outcome ω_0 (and any prices). The refinements for single-dimensional settings and multidimensional FL work in the same fashion as before: the appropriate LP (e.g., (P’)) is modified to include the appropriate set of IC and IR constraints and solved as before. The rounding of an LP solution to obtain a suitable mechanism proceeds as before. Consequently, all of our results extend to these two notions with minimal effort.

For notion (14), we incorporate constraints (14) but restrict c_i to lie in $\bar{\mathcal{D}}_i$. Recall that $\bar{\mathcal{D}}_i = \mathcal{D}_i \cup \{m_i \mathbf{1}_{\mathcal{T}_i}\}$. We also include constraints (4) as before. Again, the resulting LP can be solved given an optimal algorithm for the CM problem. The LP solution can be extended to a mechanism *exactly* as in Theorem 3.2, and it is easy to see that this extension satisfies (14). Therefore, all our results extend to this notion as well.

6.2 Dominant-strategy IC mechanisms

We can strengthen our results from Section 4 to obtain (near-) optimal *dominant-strategy incentive compatible* (DSIC) mechanisms for single-dimensional problems in time exponential in n . Thus, we obtain polytime mechanisms for any constant number of players.

The key change is in the LP (P) (or (P')), where we now enforce (1)–(4) for every player i and every type profile in $\prod_i \bar{\mathcal{D}}_i$. (Note that, as before, we can only enforce IC and IR constraints for a finite set of type profiles.) Theorem 3.1, as also the rounding procedure and arguments in Theorem 4.2 proceed essentially identically to yield a near-optimal solution to this LP. We prove that in single-dimensional settings, enforcing the IC, IR constraints for the set $\prod_i \bar{\mathcal{D}}_i$ of type profiles enables one to extend the LP solution to a (DSIC-in-expectation, IR) mechanism without increasing the expected disutility. Thus, we obtain the same guarantees as in Table 1, but under the stronger solution concept of (DSIC-in-expectation, IR).

We focus on single-dimensional settings here because at various places, our arguments rely on the well-known equivalence between monotonic allocation rules and DSIC-implementable allocation rules. We do not know if a similar result holds for multidimensional settings. This, and various other unanswered questions emerge from our result; we mention a few of these below, before delving into our construction for single-dimensional settings.

- (a) In multidimensional settings, what finite subsets $C' \subseteq C$ of the type space have the property that enforcing the IC and IR constraints for every player i and every type profile in C' suffice to extend an LP solution to a (DSIC-in-expectation, IR mechanism)?
- (b) Does $C' = \prod_i \bar{\mathcal{D}}_i$ have this extension property (as is the case in single-dimensional settings)?
- (c) Is there some C' of size $\text{poly}(n, |D|)$ with this extension property?

We now describe briefly the changes required to obtain (DSIC-in-expectation, IR) mechanisms. Analogous to Lemma 2.3, we can obtain estimates m_i such that there is an optimal mechanism M^* such that on any input $c \in \prod_i (\mathcal{D}_i \cup \{m_i\})$ where $c_i < m_i$ for at least one i , M^* only buys the item with non-zero probability from a player i with $c_i < m_i$ (the same proof approach works). Let $\tilde{\mathcal{D}} := \prod_i \bar{\mathcal{D}}_i$ and $\tilde{\mathcal{D}}_{-i} := \prod_{j \neq i} \bar{\mathcal{D}}_j$; also, let $\tilde{\mathcal{D}}_i := \bar{\mathcal{D}}_i := \mathcal{D}_i \cup \{m_i\}$ for uniformity of notation. For $c \in \tilde{\mathcal{D}}$, define $\Omega(c) = \{\omega \in \Omega : \omega_i = \emptyset \text{ for all } i \text{ s.t. } c_i = m_i\}$, if there is some i such that $c_i < m_i$, and Ω otherwise. In our LP (P), or its relaxation (P') (where we move to an LP-relaxation of the CM problem and consider the allocation-set Ω_{ext}), we now enforce (1)–(4) for all i , all $c_i, c'_i \in \tilde{\mathcal{D}}_i$ and all $c_{-i} \in \tilde{\mathcal{D}}_{-i}$.

Let (K-P), and (K-P') (with allocation-set Ω_{LP}) denote these new LPs. When n is a constant, both LPs have a polynomial number of constraints. So again by considering the dual, we can efficiently compute: (i) an optimal solution to (K-P) given an optimal algorithm for the CM problem; and (ii) an optimal solution to (K-P'). If the CM problem can be encoded via (C-P) and we have an LP-relative approximation algorithm for the CM problem, then one can use the rounding procedure described in Theorem 4.2 to round the optimal solution to (K-P') to a near-optimal solution to (K-P); the arguments are essentially identical.

So suppose that we have a near-optimal solution (x, p) to (K-P). We extend (x, p) to a (DSIC-in-expectation, IR) mechanism $M = (\mathcal{A}, \{q_i\})$ without increasing the expected disutility. Here, $\mathcal{A}(c)$ and $q_i(c)$ denote as before the allocation-distribution and expected payment to i , on input c .

Define $y_c = \sum_{\omega} x_{c,\omega} \omega$, where we treat ω as a vector in $\{0, 1\}^n$ with $\omega_i \equiv \omega_{i,v}$ for the single covering object $v \in \mathcal{T}_i$. Let $0 \leq c_i^1 < c_i^2 < \dots < c_i^{k_i} = c_i^{\max}$ be the values in \mathcal{D}_i , and set $c_i^{k_i+1} := m_i$. Define the mapping $H : C \rightarrow \tilde{\mathcal{D}}$ as follows: set $H(c) := (H_i(c_i))_{i=1,\dots,n}$, where $H_i(c_i)$ is c_i^{r+1} if $c_i \in (c_i^r, c_i^{r+1}]$, $r \leq k_i$, and m_i if $c_i \geq m_i$. Define $H_{-i}(c_{-i}) := (H_j(c_j))_{j \neq i}$.

Consider $c \in C$. If $c_i \leq c_i^{\max}$ for at least one i , we set $\mathcal{A}(c) = y_{H(c)}$. If $c_i > c_i^{\max}$ for all i , we set $\mathcal{A}(c)$ as in the VCG mechanism. Since we are in the single-dimensional setting, if we show that for all i , $c_{-i} \in C_{-i}$, $\mathcal{A}(c)_i$ is non-increasing in c_i and hits 0 at some point, then we know that setting $q_i(c) = c_i \mathcal{A}(c)_i + \int_{c_i}^{\infty} \mathcal{A}(t, c_{-i})_i dt$ ensures such that $M = (\mathcal{A}, \{q_i\})$ is (DSIC, IR)-in-expectation.

Consider some i , $c_{-i} \in C_{-i}$. If $c_j \leq c_j^{\max}$ for some $j \neq i$, then $\mathcal{A}(c) = y_{H(c)}$. Since H_i is non-decreasing in c_i and $y_{c,i}$ is non-increasing in c_i (which is easily verified), it follows that $\mathcal{A}(c)_i$ is non-increasing in c_i . Also, if $c_{-i} \in \tilde{D}_{-i}$, then one can argue as in the proof of Theorem 4.2 that $q_i(c) \leq p_{i,c}$. Hence, M has expected total payment at most $\sum_{c,i} \Pr_{\mathcal{D}}(c) p_{i,c}$. Suppose $c_j > c_j^{\max}$ for all $j \neq i$. Then, $H_j(c_j) = m_j$ for all $j \neq i$. So $\mathcal{A}(c) = y_{H(c)}$ for $c_i \leq c_i^{\max}$, and is the VCG allocation for $c_i > c_i^{\max}$. Therefore, $\mathcal{A}(c)_i = 1$ for $c_i \leq c_i^{\max}$, and the VCG allocation for $c_i > c_i^{\max}$, which is clearly non-increasing in c_i .

Theorem 6.1 *For single-dimensional problems with a constant number of players, we obtain the same guarantees as in Table 1, but under the stronger solution concept of DSIC-in-expectation and IR.*

Acknowledgments

We thank the anonymous reviewers and referees of the conference and journal submissions for various useful comments. Section 6.1 was prompted by questions raised by an anonymous referee of the journal submission regarding the solution concepts (13)-(15).

References

- [1] S. Alaei, H. Fu, N. Haghpanah, J. Hartline, and A. Malekian. Bayesian optimal auctions via multi- to single-agent reduction. In *Proc., EC*, 17, 2012.
- [2] Y. Cai, C. Daskalakis, and S. Weinberg. Optimal multi-dimensional mechanism design: reducing revenue to welfare maximization. In *Proc., FOCS*, 130–139, 2012.
- [3] Y. Cai, C. Daskalakis, and S. Weinberg. Reducing revenue to welfare maximization: approximation algorithms and other generalizations. In *Proc., SODA*, 2013.
- [4] Y. Cai, C. Daskalakis, and S. M. Weinberg. Understanding incentives: mechanism design becomes algorithm design. In *Proc., FOCS*, 2013.
- [5] S. Chawla, J. Hartline, D. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proc., STOC*, 311–320, 2010.
- [6] J. Crémer, and R. McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53, 2, 345–361, 1985.
- [7] J. Crémer, and R. McLean. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* 56, 6, 1247–1257, 1988.
- [8] S. Dobzinski, H. Fu, and R. Kleinberg. Optimal auctions with correlated bidders are easy. In *Proc. STOC*, 129–138, 2011.
- [9] S. Hart and N. Nisan. Approximate revenue maximization with multiple items. In *Proc., EC*, 656, 2012.

- [10] K. Jain, M. Mahdian, E. Markakis, A. Saberi, and V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *J. ACM*, 50:795–824, 2003.
- [11] K. Jain, M. Mahdian, and M. Salavatipour. Packing Steiner trees. In *Proc. SODA*, pages 266–274, 2003.
- [12] R. Lavi and C. Swamy. Truthful and near-optimal mechanism design via linear programming. *Journal of the ACM*, 58(6): 25, 2011.
- [13] P. Klemperer. Auction theory: A guide to the literature. *Journal of Economic Surveys* 13, 227–286, 1999.
- [14] V. Krishna. *Auction Theory*. Academic Press, 2010.
- [15] S. Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. *Inf. Comput.*, 222:45–58, 2013.
- [16] R. McAfee and P. Reny. Correlated information and mechanism design. *Econometrica* 60, 2, 395–421, 1992.
- [17] H. Minooei and C. Swamy. Truthful mechanism design for multidimensional covering problems. In *Proc., WINE*, 448–461, 2012. Detailed version at <http://www.math.uwaterloo.ca/~cswamy/papers/coveringmech-arxiv.pdf>.
- [18] H. Minooei and C. Swamy. Near-optimal and robust mechanism design for covering problems with correlated players. In *Proc. WINE*, 2013.
- [19] R. Myerson. Optimal auction design. *Math. of Operations Research*, 6:58–73, 1981.
- [20] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, September 2007.
- [21] C. Papadimitriou and G. Pierrakos. On optimal single-item auctions. In *Proc., STOC*, 119–128, 2011.
- [22] A. Ronen. On approximating optimal auctions. In *Proc. EC*, 11–17, 2001.
- [23] A. Ronen and D. Lehmann. Nearly optimal multi attribute auctions. In *Proc. EC*, pages 279–285, 2005.
- [24] A. Ronen and A. Saberi. On the hardness of optimal auctions. In *Proc., FOCS*, 396–405, 2002.
- [25] T. Roughgarden and I. Talgam-Cohen. Optimal and near-optimal mechanism design with interdependent values. In *Proc., EC*, 767–784, 2013.

A Proof of Lemma 2.3

Consider the following LP, which is the same as (P) except that we only consider $c \in \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$.

$$\min \quad \sum_{c \in \mathcal{D}} \Pr_{\mathcal{D}}(c) \left(\sum_i q_{i,c} + \kappa \sum_{\omega} x_{c,\omega} \text{pub}(\omega) \right) \quad (\text{LP})$$

$$\text{s.t.} \quad \sum_{\omega} x_{c,\omega} = 1 \quad \forall c \in \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i}) \quad (16)$$

$$q_{i,(c_i,c_{-i})} - \sum_{\omega} c_i(\omega) x_{(c_i,c_{-i}),\omega} \geq q_{i,(c'_i,c_{-i})} - \sum_{\omega} c_i(\omega) x_{(c'_i,c_{-i}),\omega} \quad \forall i, c_i, c'_i \in \mathcal{D}_i, c_{-i} \in \mathcal{D}_{-i} \quad (17)$$

$$q_{i,(c_i,c_{-i})} - \sum_{\omega} c_i(\omega) x_{(c_i,c_{-i}),\omega} \geq 0 \quad \forall i, c_i \in \overline{\mathcal{D}}_i, c_{-i} \in \mathcal{D}_{-i} \quad (18)$$

$$q, x \geq 0. \quad (19)$$

Let $M = (\mathcal{A}, \{p_i\})$ be an optimal mechanism. Recall that O^* is the expected disutility of M . Then, M naturally yields a feasible solution (x, q) to (LP) of objective value O^* , where $x_{c,\omega} = \Pr_M[\mathcal{A}(c) = \omega]$ and $q_{i,c} = \mathbb{E}_M[p_i(c)]$. Let (\hat{x}, \hat{q}) be an *optimal basic solution* to (LP). Then, for some N such that $\log N$ is polynomially bounded in the input size \mathcal{I} , we can say that the values of all variables are integer multiples of $\frac{1}{N}$, and $\log(N\hat{x}_{c,\omega}), \log(N\hat{q}_{i,c}) = \text{poly}(\mathcal{I})$ for all $i, c \in \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i}), \omega$.

First, we claim that we may assume that for every $i, c_i \in \mathcal{D}_i, c_{-i} \in \mathcal{D}_{-i}$, if whenever $\hat{x}_{c,\omega} > 0$ we have $\omega_i = \emptyset$ (where $c = (c_i, c_{-i})$), then $\hat{q}_{i,c} = 0$. If not, then (17) implies that $\hat{q}_{i,(\tilde{c}_i,c_{-i})} - \sum_{\omega} \tilde{c}_i(\omega) x_{(\tilde{c}_i,c_{-i}),\omega} \geq \hat{q}_{i,c}$ for all $\tilde{c}_i \in \mathcal{D}_i$ and decreasing $\hat{q}_{i,(\tilde{c}_i,c_{-i})}$ by $\hat{q}_{i,c}$ for all $\tilde{c}_i \in \mathcal{D}_i$ continues to satisfy (17)–(19).

Set $m_i := \max(2 \sum_{i,v \in \mathcal{T}_i} \max_{c_i \in \mathcal{D}_i} c_{i,v}, N \sum_{i,c} \hat{q}_{i,c})$ for all i . So $\log m_i = \text{poly}(\mathcal{I})$. Recall that $\overline{\mathcal{D}}_i := \mathcal{D}_i \cup \{m_i \mathbf{1}_{\mathcal{T}_i}\}$ for all $i \in [n]$, and $\overline{\mathcal{D}} := \bigcup_i (\overline{\mathcal{D}}_i \times \mathcal{D}_{-i})$.

Now we extend (\hat{x}, \hat{q}) to (\tilde{x}, \tilde{q}) that assigns values also to type-profiles in $\overline{\mathcal{D}} \setminus \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$ so that constraints (16)–(19) hold for all $i, c_i, c'_i \in \overline{\mathcal{D}}_i, c_{-i} \in \mathcal{D}_{-i}$. First set $\tilde{x}_{c,\omega} = \hat{x}_{c,\omega}, \tilde{q}_{i,c} = \hat{q}_{i,c}$ for all $i, \omega, c \in \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$. Consider $c \in \overline{\mathcal{D}} \setminus \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$, and let i be such that $c_i = m_i \mathbf{1}_{\mathcal{T}_i}$ (there is exactly one such i). We “run” VCG on c considering only the cost incurred by the players. That is, we set $\tilde{x}_{c,\omega} = 1$ for $\omega = \omega(c) := \arg \min_{\omega \in \Omega} \sum_i c_i(\omega)$ and pay $\tilde{q}_{i,c} = \min_{\omega \in \Omega: \omega_i = \emptyset} \sum_j c_j(\omega) - \sum_{j \neq i} c_j(\omega(c))$ to each player i . Note that the choice of m_i ensures that $\omega(c)_i = \emptyset$ and hence, $\tilde{q}_{i,c} = 0$.

We claim that this extension satisfies (16)–(19) for all $i, c_i, c'_i \in \overline{\mathcal{D}}_i, c_{-i} \in \mathcal{D}_{-i}$. Fix $i, c_i, c'_i \in \overline{\mathcal{D}}_i, c_{-i} \in \mathcal{D}_{-i}$. It is clear that (16), (19) hold. If $c_i \in \mathcal{D}_i$, then (18) clearly holds; if $c_i = m_i \mathbf{1}_{\mathcal{T}_i}$, then it again holds since $\tilde{x}_{c,\omega} = 1$ for $\omega = \omega(c)$ and $\omega(c)_i = \emptyset$. To verify (17), we consider four cases. If $c_i, c'_i \in \mathcal{D}_i$, then (17) holds since (\tilde{x}, \tilde{q}) extends (\hat{x}, \hat{q}) . If $c_i = c'_i = m_i \mathbf{1}_{\mathcal{T}_i}$, then (17) trivially holds. If $c_i \in \mathcal{D}_i, c'_i = m_i \mathbf{1}_{\mathcal{T}_i}$, then (17) holds since the RHS of (17) is 0 (as $\tilde{x}_{(c'_i,c_{-i}),\omega} = 1$ and $\omega(c'_i, c_{-i})_i = \emptyset$). We are left with the case $c_i = m_i \mathbf{1}_{\mathcal{T}_i}$ and $c'_i \in \mathcal{D}_i$. If whenever $\tilde{x}_{(c'_i,c_{-i}),\omega} = \hat{x}_{(c'_i,c_{-i}),\omega} > 0$ we have $\omega_i = \emptyset$, then we also have $\tilde{q}_{i,(c'_i,c_{-i})} = \hat{q}_{i,(c'_i,c_{-i})} = 0$ by our earlier claim, so the RHS of (17) is 0, and (17) holds. Otherwise, we have $\sum_{\omega} c_i(\omega) \tilde{x}_{(c'_i,c_{-i}),\omega} \geq \frac{m_i}{N} \geq \hat{q}_{i,(c'_i,c_{-i})}$, so the RHS of (17) is at most 0, and (17) holds.

Thus, we have shown that (\tilde{x}, \tilde{q}) is a feasible solution to (P). Now we can apply Theorem 3.2 to extend (\tilde{x}, \tilde{q}) and obtain a support-based-(IC-in-expectation, IR) mechanism M^* whose expected disutility is at most $\sum_{c,i} \Pr_{\mathcal{D}}(c) (\tilde{q}_{i,c} + \sum_{\omega} \tilde{x}_{c,\omega} \text{pub}(\omega)) \leq O^*$. Since $\tilde{x}_{(m_i \mathbf{1}_{\mathcal{T}_i}, c_{-i}),\omega} > 0$ implies that $\omega_i = \emptyset$ for all i , M^* satisfies the required conditions. ■

B Inferiority of k -lookahead procurement auctions

The following k -lookahead auction was proposed by [8] for the single-item revenue-maximization problem generalizing the 1-lookahead auction considered by [22, 24]: on input $v = (v_1, \dots, v_n)$, pick the set I of k players with highest values, and run the revenue-maximizing (DSIC, IR) mechanism for player-set I where the distribution we use for I is the conditional distribution of the values for I given the values $(v_i)_{i \notin I}$ for the other players. Dobzinski et al. [8] show that the k -lookahead auction achieves a constant-fraction of the revenue of the optimal (DSIC, IR) mechanism.

For any $k \geq 2$, we can consider an analogous definition of k -lookahead auction for the single-item procurement problem: on input c , we pick the set I of k players with smallest costs, and run the payment-minimizing support-based-(IC-in-expectation, IR) mechanism for I for the conditional distribution of I 's costs given $(c_i)_{i \notin I}$. We call this the *k -lookahead procurement auction*. The following example shows that the expected total payment of the k -lookahead procurement auction can be arbitrarily large, even when $k = n - 1$ (that is, we drop only 1 player), and compared to the optimal expected total payment of even a deterministic (DSIC, IR) mechanism.

Let $t = K + \epsilon$ where $\epsilon > 0$, and $\delta > 0$. The distribution \mathcal{D} consists of n points: each c in $\{c : c_n = t, \exists i \in [n-1] \text{ s.t. } c_i = 0, c_j = K \forall j \neq i, n\}$ has probability $\Pr_{\mathcal{D}}(c) = \frac{1-\delta}{n-1}$, and the type-profile c where $c_i = K \forall i \neq n, c_n = t$ has probability $\Pr_{\mathcal{D}}(c) = \delta$.

Let $k = n - 1$. The k -lookahead procurement auction will always select the player-set $I = \{1, \dots, n - 1\}$, and the conditional distribution of values of players in I is simply \mathcal{D} . Let M' be the support-based-(IC-in-expectation, IR) mechanism for the players in I under this conditional distribution \mathcal{D} . Suppose that on input (K, K, \dots, K, t) , the k -lookahead auction (which runs M') buys the item from player $i \in I$ with probability x_i . Clearly, $\sum_{i \in I} x_i = 1$. Then, on the input \tilde{c} where $\tilde{c}_i = 0, \tilde{c}_j = K$ for all $j \neq i, n, \tilde{c}_n = t$, the mechanism must also buy the item from player i with probability at least x_i since M' is support-based IC. So since M' is support-based (IC-in-expectation, IR), the payment to player i under input \tilde{c} is at least K , and therefore the expected total payment of the k -lookahead auction is at least $K(\sum_{i \in I} x_i) \cdot \frac{1-\delta}{n-1} = \frac{K(1-\delta)}{n-1}$.

Now consider the following mechanism $M = (\mathcal{A}, \{p_i\})$. Consider input c . If some player $i < n$ has $c_i = 0$, M buys the item from such a player i (breaking ties in some fixed way). Otherwise, if $c_n \leq t$, M buys from player n ; else, M buys from the player i with smallest c_i . It is easy to verify that for every i and c_{-i} , this allocation rule \mathcal{A} is monotonically decreasing in c_i . Let $p_i(c) = 0$ if M does not buy the item from i on input c , and $\max\{z : M \text{ buys the item from } i \text{ on input } (z, c_{-i})\}$ otherwise. By a well-known fact, (see, e.g., Theorem 9.39 in [20]), M is DSIC and IR. Then, the total payment under any $c \in \mathcal{D}$ for which $c_i = 0$ for some i , is 0, and the total payment under the input where $c_i = K$ for all $i \neq n, c_n = t$ is t . So the expected total payment of M is $t\delta$.

Thus the ratio of the expected total payments of M' and M is at least $\frac{K(1-\delta)}{t\delta(n-1)}$, which can be made arbitrarily large by choosing δ and $\epsilon = t - K$ small enough.

C Insights for revenue-maximization in packing domains

Our study of PayM problems also leads to some interesting insights into the revenue-maximization problem in packing settings with correlated players. We state two results that are obtained via relatively-simple observations, but we believe are nevertheless of interest.

In Section C.1, we justify our comment in the Introduction that the problem of extending an LP solution to a suitable mechanism becomes much easier in a packing setting such as combinatorial auctions (CAs). We show that any solution to an LP-relaxation similar to (P) for the revenue-maximization problem in CAs can be extended to a (DSIC-in-expectation, IR) mechanism without

any loss in revenue. In Section C.2, we obtain a noteworthy extension of a result in [8]. We show that in single-dimensional packing settings, a ρ -approximation algorithm for the SWM problem can be used to obtain a ρ -approximation (DSIC-in-expectation, IR)-mechanism for the revenue-maximization problem. We obtain this by noting that a ρ -approximation for the SWM problem can be used to obtain a ρ -approximate separation oracle for the dual of the revenue-maximization LP, despite the fact that the separation problem is an SWM problem possibly involving *negative-valued* inputs⁵, which then yields, in a fairly-standard way, a ρ -approximate solution to the revenue-maximization LP.

C.1 Extending LP solutions to (DSIC-in-expectation, IR) mechanisms

We consider the prototypical problem of combinatorial auctions; similar arguments can be made for other packing domains. In CAs, a feasible allocation ω is one that allots a disjoint set ω_i of items (which could be empty) to each player i , and player i 's value under allocation ω is $v_i(\omega_i)$, where $v_i : 2^{[m]} \mapsto \mathbb{R}_+$ is player i 's private valuation function. We use $v_i(\omega)$ to denote $v_i(\omega_i)$. Let V_i denote the set of all private types of player i , and $V_{-i} = \prod_{j \neq i} V_j$. As before, let Ω be the set of all feasible solutions.

We consider the following LP along the lines of (P). Since we are in a packing setting, we do not need the m_i estimates. We may assume that each \mathcal{D}_i contains the valuation 0_i , where $0_i(\omega) = 0$ for all ω , since if not, we can just add this to \mathcal{D}_i , and set $\Pr_{\mathcal{D}}(0_i, v_{-i}) = 0$. Let $\overline{\mathcal{D}}' := \bigcup_i (\mathcal{D}_i \times \mathcal{D}_{-i})$.

$$\begin{aligned}
\max \quad & \sum_{v \in \mathcal{D}} \Pr_{\mathcal{D}}(v) \sum_i p_{i,v} & \text{(R-P)} \\
\text{s.t.} \quad & \sum_{\omega \in \Omega} x_{v,\omega} \leq 1 & \forall v \in \overline{\mathcal{D}}' & (20) \\
& \sum_{\omega \in \Omega} v_i(\omega) x_{(v_i, v_{-i}), \omega} - p_{i,(v_i, v_{-i})} \geq \sum_{\omega \in \Omega} v_i(\omega) x_{(v'_i, v_{-i}), \omega} - p_{i,(v'_i, v_{-i})} & \forall i, v_i, v'_i \in \mathcal{D}_i, v_{-i} \in \mathcal{D}_{-i} \\
& \sum_{\omega \in \Omega} v_i(\omega) x_{(v_i, v_{-i}), \omega} - p_{i,(v_i, v_{-i})} \geq 0 & \forall i, v \in \overline{\mathcal{D}}' \\
& p, x \geq 0.
\end{aligned}$$

The above LP-relaxation is similar to the LP in [8] for single-dimensional packing problems. For CAs, the LP in [8] is subtly different from (R-P): their allocation variables encode the probability that a player i receives a set S of items. If we let the allocation space in (R-P) be the set Ω_{ext} of extreme points of the standard LP-relaxation for the CA problem, then our formulations coincide since a feasible solution to the standard LP specifies the extent to which each player receives each set. (The convex-decomposition technique in [12] directly implies that an integrality-gap verifying ρ -approximation algorithm for the SWM problem can be used to decompose the fractional allocation specified by the LP in [8] scaled by ρ into a distribution over Ω , and thereby obtain a solution to (R-P).)

Let (\tilde{x}, \tilde{p}) be a solution to (R-P). We convert (\tilde{x}, \tilde{p}) to a (DSIC-in-expectation, IR) mechanism $M = (\mathcal{A}, \{q_i\})$ with no smaller expected total revenue. Here, $\mathcal{A}(v)$ and $q_i(v)$ are the allocation-distribution and expected price of player i on input v . Since any support-based-(IC, IR)-in-expectation mechanism yields a feasible solution to (R-P), this also shows that *any support-based-(IC, IR)-in-expectation mechanism for CAs can be extended to a (DSIC-in-expectation, IR) mechanism without any loss in revenue.*

⁵Recall that negative costs are quite problematic for the CM problem, and hence, we had to resort to the method outlined in Section 4.

Our argument is similar to that in the proof of Theorem 3.2, but the packing nature of the problem simplifies things significantly. We may assume that all constraints (20) are tight; otherwise if there is some $v \in \overline{\mathcal{D}}'$ with $\sum_{\omega \in \Omega} \tilde{x}_{v,\omega} < 1$, then letting ω^0 be allocation where $\omega_i^0 = \emptyset$ for all i , we can increase \tilde{x}_{v,ω^0} to make this sum 1 without affecting feasibility.

First, we set $\mathcal{A}(v) = \tilde{x}_v$, $q_i(v) = \tilde{p}_{i,v}$ for all $v \in \overline{\mathcal{D}}'$ and all i , so it is clear that the expected total revenue of M is the value of (\tilde{x}, \tilde{p}) .

If $|\{i : v_i \notin \mathcal{D}_i\}| \geq 2$, then we give everyone the empty-set and charge everyone 0. Otherwise, suppose $v_i \notin \mathcal{D}_i$, $v_{-i} \in \mathcal{D}_{-i}$. Let $\bar{v}^{(i)} = \arg \max_{\bar{v}_i \in \mathcal{D}_i} (\sum_{\omega} v_i(\omega) \tilde{x}_{(\bar{v}_i, v_{-i}), \omega} - \tilde{p}_{i,(\bar{v}_i, v_{-i})})$ and $\bar{y}^{(i)} = \tilde{x}_{(\bar{v}^{(i)}, v_{-i})}$. For $\omega \in \Omega$, let $\text{proj}_i(\omega)$ denote the allocation where player i receives $\omega_i \subseteq [m]$, and the other players receive \emptyset . Viewing $\mathcal{A}(v)$ as the random variable specifying the allocation selected, we set $\mathcal{A}(v) = \text{proj}_i(\omega)$ with probability $\bar{y}_{\omega}^{(i)}$. We set $q_i(v) = \tilde{p}_{i,(\bar{v}^{(i)}, v_{-i})}$. Since $0_i \in \mathcal{D}_i$, we have $\mathbb{E}_{\mathcal{A}}[v_i(\mathcal{A}(v))] - q_i(v) \geq \sum_{\omega} 0_i(\omega) \tilde{x}_{(0_i, v_{-i}), \omega} - \tilde{p}_{i,(0_i, v_{-i})} \geq 0$, so M is IR-in-expectation.

To see that M is DSIC in expectation, consider some i , $v_i, v'_i \in V_i$, $v_{-i} \in V_{-i}$. If $v_{-i} \notin \mathcal{D}_{-i}$, then player i always receives the empty set and pays 0. Otherwise, we have ensured by definition that player i does not benefit by lying.

C.2 Utilizing approximation algorithms for the SWM problem

We briefly sketch how to utilize a ρ -approximation algorithm for the SWM problem to obtain a ρ -approximation (DSIC-in-expectation, IR)-mechanism for the revenue-maximization problem in single-dimensional settings. Similar arguments apply to packing settings with additive types.

Consider again the LP (R-P). For all i , ω , we now have $v_i(\omega) = v_i \alpha_{i,\omega}$, where $v_i \in \mathbb{R}$ is i 's private type and $\alpha_{i,\omega} \geq 0$ is public knowledge. This is identical to the LP in [8] for single-parameter revenue-maximization problems. It will be convenient to view $\omega \in \Omega$ as the vector $\{\alpha_{i,\omega}\}_{i \in [n]} \in \mathbb{R}_+^n$. Since we are in a packing setting, Ω is downward closed, so $\omega \in \Omega$ and $\omega' \leq \omega$ implies that $\omega' \in \Omega$. The dual of (R-P) is:

$$\min \quad \sum_v \gamma_v \tag{R-D}$$

$$\text{s.t.} \quad \sum_{i: v \in \mathcal{D}_i \times \mathcal{D}_{-i}} \left(\sum_{v'_i \in \mathcal{D}_i} (v_i(\omega) y_{i,(v_i, v_{-i}), v'_i} - v'_i(\omega) y_{i,(v'_i, v_{-i}), v_i}) + v_i(\omega) \beta_{i,v} \right) \leq \gamma_v \quad \forall v \in \overline{\mathcal{D}}', \omega \in \Omega \tag{21}$$

$$\sum_{v'_i \in \mathcal{D}_i} (y_{i,(v_i, v_{-i}), v'_i} - y_{i,(v'_i, v_{-i}), v_i}) + \beta_{i,v_i, v_{-i}} \geq \Pr_D(v) \quad \forall i, v \in \overline{\mathcal{D}}' \tag{22}$$

$$y, \beta, \gamma \geq 0. \tag{23}$$

Let optr be the common optimal value of (R-P) and (R-D). Define $\theta_i^v = \sum_{v'_i \in \mathcal{D}_i} (v_i y_{i,(v_i, v_{-i}), v'_i} - v'_i y_{i,(v'_i, v_{-i}), v_i}) + v_i \beta_{i,v}$ if $v \in \mathcal{D}_i \times \mathcal{D}_{-i}$, and 0 otherwise. Then, the separation problem for (R-D) amounts to determining if $\max_{\omega \in \Omega} \sum_i \theta_i^v \alpha_{i,\omega} \leq \gamma_v$ for every $v \in \overline{\mathcal{D}}'$; that is, solving an SWM problem over the allocation space Ω under the input $\theta^v = \{\theta_i^v\}_{i \in [n]}$ for every v . Given a ρ -approximation algorithm \mathcal{A} for the SWM problem (where $\rho \geq 1$) that only works with nonnegative inputs, we can obtain a ρ -approximate solution for the input θ^v as follows. Set $(\theta^+)_i^v := \max(0, \theta_i^v)$, use \mathcal{A} with the input $(\theta^+)^v = \{(\theta^+)_i^v\}_{i \in [n]}$ to obtain a solution ω' , and let $\omega_i'' = \omega_i'$ if $\theta_i \geq 0$ and 0 otherwise. Thus, \mathcal{A} can be used to approximately separate over constraints (21). We now argue that this implies that we can obtain a solution to (R-P) of value at least optr/ρ . This follows a routine argument in the approximation-algorithms literature (see, e.g., [11]).

Define $\mathcal{P}(\nu) := \{(y, \beta, \gamma) : (21)-(23), \sum_v \gamma_v \leq \nu\}$. Note that optr is the smallest ν such that $\mathcal{P}(\nu) \neq \emptyset$. Given ν, y, β, γ , we either show that $(y, \beta, \gamma \rho) \in \mathcal{P}(\nu \rho)$, or we exhibit a hyperplane

separating (y, β, γ) from $\mathcal{P}(\nu)$. To do this, we first check if $\sum_v \gamma_v \leq \nu$, (22), (23) hold, and if not use the appropriate inequality as the separating hyperplane. Next, for every $v \in \overline{\mathcal{D}}'$, we use \mathcal{A} as specified above to obtain some $\omega'' \in \Omega$. If in this process, the LHS of (21) exceeds γ_v for some v , then we return the corresponding inequality as the separating hyperplane. Otherwise, for all $v \in \overline{\mathcal{D}}'$ and all $\omega \in \Omega$, the LHS of (21) is at most $\gamma\rho$, and so $(y, \beta, \gamma\rho) \in \mathcal{P}(\nu\rho)$.

Thus, for a fixed ν , in polynomial time, the ellipsoid method either certifies that $\mathcal{P}(\nu) = \emptyset$, or returns a point (y, β, γ) with $(y, \beta, \gamma\rho) \in \mathcal{P}(\nu\rho)$. We find the smallest value ν^* (via binary search) such that the ellipsoid method run for ν^* (with the above separation oracle) returns a solution (y^*, β^*, γ^*) with $(y^*, \beta^*, \gamma^*\rho) \in \mathcal{P}(\nu^*\rho)$; hence, $\text{optr} \leq \nu^*\rho$. For any $\epsilon > 0$, running the ellipsoid method for $\nu^* - \epsilon$ yields a polynomial-size certificate for the emptiness of $\mathcal{P}(\nu^* - \epsilon)$. This consists of the polynomially many violated inequalities returned by the separation oracle during the execution of the ellipsoid method and the inequality $\sum_v \gamma_v \leq \nu^* - \epsilon$. By duality (or Farkas' lemma), this means that here is a polynomial-size solution (\tilde{x}, \tilde{p}) to (R-P) whose value is at least $\nu^* - \epsilon$. Taking ϵ to be $1/\exp(\text{input size})$ (so $\ln(\frac{1}{\epsilon})$ is polynomially bounded), this also implies that (\tilde{x}, \tilde{p}) has value at least $\nu^* \geq \text{optr}/\rho$.

We can now convert (\tilde{x}, \tilde{p}) to a (DSIC-in-expectation, IR) mechanism with revenue at least optr/ρ via the procedure described in Section C.1.