




Contents lists available at ScienceDirect

Artificial Intelligence

journal homepage: www.elsevier.com/locate/artint

Designing menus of contracts efficiently: The power of randomization ^{☆,☆☆}

Matteo Castiglioni ^{*}, Alberto Marchesi, Nicola Gatti

Politecnico di Milano, Piazza Leonardo da Vinci 32, Milan, Italy

ARTICLE INFO

Article history:

Received 13 August 2022

Received in revised form 6 February 2023

Accepted 8 February 2023

Available online 14 February 2023

Keywords:

Principal-agent problems

Contract theory

Bayesian games

ABSTRACT

We study *hidden-action principal-agent problems* in which a principal commits to an outcome-dependent payment scheme (called *contract*) so as to incentivize the agent to take a costly, unobservable action leading to favorable outcomes. In particular, we focus on *Bayesian* settings where the agent has *private information*. This is collectively encoded by the agent's *type*, which is unknown to the principal but randomly drawn according to a finitely-supported, commonly-known probability distribution. In our model, the agent's type determines both the probability distribution over outcomes and the cost associated with each agent's action. In Bayesian principal-agent problems, the principal may be better off by committing to a *menu of contracts* specifying a contract for each agent's type, rather than committing to a single contract. This induces a two-stage process that resembles interactions studied in classical *mechanism design*: after the principal has committed to a menu, the agent first *reports a type* to the principal, and, then, the latter puts in place the contract in the menu that corresponds to the reported type. Thus, the principal's computational problem boils down to designing a menu of contracts that incentivizes the agent to report their true type and maximizes expected utility.

Previous works showed that, in Bayesian principal-agent problems, computing an optimal menu of contracts or an optimal (single) contract is APX-hard, which is in sharp contrast from what happens in non-Bayesian settings, where an optimal contract can be computed efficiently. Crucially, previous works focus on menus of *deterministic* contracts. Surprisingly, in this paper, we show that, if one instead considers menus of *randomized* contracts defined as probability distributions over payment vectors, then an "almost-optimal" menu can be computed in polynomial time. Indeed, the problem of computing a principal-optimal menu of *randomized* contracts may *not* admit a maximum, but only a supremum. Nevertheless, we show how to design a polynomial-time algorithm that guarantees the principal with an expected utility arbitrarily close to the supremum. Besides this main result, we also close several gaps in the computational complexity analysis of the problem of computing menus of deterministic contracts. In particular, we prove that the problem cannot be approximated up to within any multiplicative factor and it does *not* admit an additive FPTAS unless $P = NP$, even in basic instances with a constant number of actions and only four outcomes. This considerably extends previously-known negative results. Then, we show that our hardness result is tight, by providing an additive PTAS that works in instances with a constant number of outcomes. We complete our analysis by showing that

[☆] This paper is a participant in the 2022 ACM Conference on Economics and Computation (EC) Forward-to-Journal Program.

^{☆☆} In the conference version of the paper [1], we incorrectly claim that the problem of finding an optimal menu of randomized contracts admits a maximum and that an optimal menu can be computed in polynomial time. In this version, we show that the problem does *not* admit a maximum but only a supremum, and we provide a polynomial-time algorithm that provides a solution whose value is arbitrarily close to that provided by the supremum.

^{*} Corresponding author.

E-mail addresses: matteo.castiglioni@polimi.it (M. Castiglioni), alberto.marchesi@polimi.it (A. Marchesi), nicola.gatti@polimi.it (N. Gatti).

an optimal menu of deterministic contracts can be computed in polynomial time when either there are only two outcomes or there is a constant number of types.

© 2023 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Principal-agent problems have recently received growing attention from the economics and computation community. These problems model the interaction between a *principal* and an *agent*, where the latter chooses an action determining some externalities on the former. In this work, we focus on *hidden-action* problems—also known as models with *moral hazard*—in which the principal cannot observe the action adopted by the agent, but only an outcome that is stochastically determined as an effect of such an action. The agent incurs a cost for performing the action, while the principal receives a reward associated with the realized outcome. Thus, the goal of the principal is to incentivize the agent to take an action resulting in favorable outcomes. This is accomplished by the principal by committing to a *contract*, which is an outcome-dependent payment scheme defining a payment from the principal to the agent for every possible outcome.

A classical example of the principal-agent problem is that of a salesperson (agent) working for a company (principal). The former has to decide on the level of effort (action) to put into selling some products on behalf of the company. The latter can observe the total number of products sold by the salesperson (outcome), but it has no information about the actual level of effort undertaken by the salesperson. This naturally fits hidden-action models. Moreover, the company usually incentivizes a high level of effort by paying the salesperson on the basis of the number of products that have been actually sold, as is the case in the classical definition of contracts. Besides this simple example, nowadays principal-agent problems are ubiquitous in digital economies, as they find application in several real-world scenarios, such as, *e.g.*, crowdsourcing platforms [2], blockchain-based smart contracts [3], and healthcare [4].

Most of the computational works on principal-agent problems have focused on the basic setting in which the principal knows everything about the agent, *i.e.*, they know both the probability distribution over outcomes and the cost associated with each agent's action. Very recently, three concurrent works [5–7] started the study of much more realistic *Bayesian* principal-agent problems—also known as models with *adversarial selection*—in which the agent has some *private information* that is unknown to the principal. As it is common in Bayesian models, these works assume that the agent's private information is collectively encoded by an unknown agent's *type*, and that the latter is drawn from a probability distribution over a set of possible types, which is known to the principal. For instance, in the company-salesperson example described above, the salesperson may have some private features (such as, *e.g.*, experience gained with past works and/or advanced training courses) that determine how effectively the undertaken level of effort converts into sales.

The addition of private information in principal-agent problems establishes an intimate connection with *mechanism design*. While the latter has received a lot of attention on its own by computational economics—thanks to its widespread application in auction settings—, only the very recent works by Guruganesh et al. [5] and Alon et al. [7] addressed the computational aspects of problems at the interface of the two fields. Indeed, in Bayesian principal-agent problems, it may be the case that the principal is better off committing to a *menu of contracts* specifying a contract for every possible agent's type, rather than committing to a single contract. This induces a two-stage process that resembles interactions usually studied in mechanism design. In particular, as a first stage after the principal's commitment, the agent *reports a type* to the principal, possibly different from their true type. Then, the interaction goes on as in a non-Bayesian principal-agent problem, with the principal selecting the contract in the menu that corresponds to the reported type. The principal's goal is to commit to menus of contracts that incentivize the agent to report their true type, choosing an expected-utility-maximizing menu among them. Notice that proposing menus of contracts is natural in many practical applications. For instance, in the company-salesperson example, one may imagine the company proposing a portfolio of different payment regimes to the agent, with the latter selecting the preferred one based on their private information. Intuitively, this could considerably boost the revenue of the company with respect to proposing a single contract.

1.1. Original contributions

In this paper, we investigate the computational complexity of finding an optimal menu of contracts for the principal, *i.e.*, the contract maximizing the principal's expected utility among those that incentivize the agent to truthfully report their type. In particular, we study general Bayesian settings where the agent's private information determines both the probability distributions and the costs of actions.

In Bayesian principal-agent settings, designing an optimal (single) contract is largely computationally intractable, with the exception of some specific cases [5,6]. Guruganesh et al. [5] unsuccessfully tried to circumvent this issue with menus of contracts, showing that the problem of computing an optimal menu is APX-hard even in instances with a constant number of actions. These results are in sharp contrast with what happens in non-Bayesian settings, where an optimal contract can be designed efficiently [8]. Crucially, the work by Guruganesh et al. [5] focuses on menus of *deterministic* contracts, in which no randomization is involved. The main result of our work is that, if one considers menus of *randomized* contracts, then an “almost-optimal” one can indeed be computed in polynomial time in arbitrary Bayesian principal-agent

problem instances. Indeed, the problem of computing a principal-optimal menu of *randomized* contracts may *not* admit a maximum, but only a supremum. Nevertheless, we show how to design a polynomial-time algorithm that guarantees the principal with an expected utility arbitrarily close to the supremum. This is surprising since randomized contracts generalize classical, deterministic ones by specifying probability distributions over payment vectors, so that, after the type-reporting stage, the distribution corresponding to the reported agent's type is employed by the principal to draw a contract that is communicated to the agent.

After the provision of the preliminaries in Section 2, in Section 3 we provide our main results on menus of randomized contracts. As a first step, we show that the problem of computing a principal-optimal menu of randomized contracts may *not* admit a maximum, but only a supremum. Then, we show how to design in polynomial time a menu of randomized contracts with the principal's expected utility greater than or equal to the value of the supremum minus ϵ , for any given $\epsilon < 0$. To do so, we first show that, for every $\epsilon > 0$, there always exists a menu of randomized contracts that achieves the principal's expected utility at most ϵ less than the supremum by using "small" payment values, which can be bounded above by $1/\epsilon$ and a suitably-defined exponential function of the instance size. This is crucial to show that, in order to find the desired menu, we can restrict the attention to randomized contracts placing positive probability on a specific finite set of deterministic contracts, whose size is exponential in the instance size. Given such a set, we can formulate the problem as a linear program with exponentially-many variables and polynomially-many constraints, whose dual can be solved in polynomial time by means of the *ellipsoid algorithm* provided that a suitable separation oracle can be implemented in polynomial time. Such an oracle can be formulated as an optimization problem over the finite set of deterministic contracts defined above, which we show that it can be solved in time polynomial in the instance size and in $\log(1/\epsilon)$, proving our main result and concluding the paper.

In Section 4, we strengthen our main positive findings on randomized contracts by providing a strongly negative result for the problem of computing an optimal menu of deterministic contracts. Such a result considerably generalizes the negative one by Guruganesh et al. [5], as we prove that the problem cannot be approximated up to within any constant multiplicative factor and it does *not* admit an additive FPTAS unless $P = NP$, even in instances with a constant number of actions and only four outcomes. We prove our result by resorting to a non-trivial and non-standard reduction from a suitably-defined promise problem related to finding maximal independent sets in undirected graphs with a bounded degree. Let us remark that our negative result is surprising since an optimal (single) contract can be computed in polynomial time in Bayesian settings with a constant number of outcomes [5,6]. Next, in Section 5, we close the gaps in the computational analysis of menus of deterministic contracts by showing that our hardness results are indeed tight. In particular, we provide an additive PTAS for the problem of computing an optimal menu in instances with a constant number of outcomes. Our approximation scheme works by finding an approximately-incentive-compatible menu of deterministic contracts (*i.e.*, one that does *not* perfectly incentivize the agent to report their true type, according to a suitable definition of approximation introduced for our purposes), which can be shown to provide a good additive approximation of the optimal principal's expected utility. We prove that such an approximate menu of deterministic contracts can be found in polynomial time by restricting the attention to menus that only employ a "small" number of different contracts. Finally, starting from the approximate menu, we show how to recover in polynomial time a menu of deterministic contracts that correctly incentivizes the agent to report their true type, only incurring a small additional loss in terms of the principal's expected utility. Finally, in Section 6 we provide two additional positive results that complete our computational analysis. In particular, we show that the problem of finding an optimal menu of deterministic contracts can be solved in polynomial time when either there are only two outcomes or there is a constant number of agent types (and outcomes and actions can be an arbitrary number).

All the proofs omitted from the main body of the paper are in the Appendix.

1.2. Related works

Hidden-action principal-agent problems have received considerable attention in the economic literature, as part of a broader subject called *contract theory* [9–12] (see the books by Mas-Colell et al. [13], Bolton et al. [14], and Laffont and Martimort [15] for a detailed treatment of the subject).

Strong interest in the computational aspects of contract theory has emerged only recently. In the following, we survey the major computational works on hidden-action principal-agent problems.

Works on non-Bayesian settings Most of the computational works on principal-agent problems focus on non-Bayesian settings in which the principal knows everything about the agent. Since in a classical non-Bayesian setting the principal's computational problem can be solved straightforwardly in polynomial time by means of linear programming, all the works on the topic introduced more complicated models. Babaioff et al. [16] study a model with multiple agents (see also its extended version [17] and its follow-ups [18,19]), focusing on how complex combinations of agents' actions influence the resulting outcome in presence of inter-agent externalities. Dütting et al. [20] and Dütting et al. [21] study other non-Bayesian principal-agent models, whose underlying structure is combinatorial. In particular, the former study the case in which the outcome space is defined implicitly through a suitably-defined succinct representation, while the latter addresses settings in which the agent can select a subset of actions (rather than a single one) out of a set of available actions. Other works worth citing are [22], which studies contract complexity in terms of the number of different payments specified by contracts, and [2], which proposes an online learning model and solves it by means of multi-armed bandit techniques. Another

important line of work is that initiated by Dütting et al. [8] and aimed at using the computational lens for the efficiency analysis (in terms of principal's expected utility) of *linear contracts* with respect to general ones, where the former are simple, pure-commission contracts that pay the agent a given fraction of the principal's reward associated with the obtained outcome. In particular, Dütting et al. [8] show that, in non-Bayesian principal-agent settings, linear contracts perform well under reasonable assumptions.

Works on Bayesian settings There are three works that are arguably the most related to ours, namely [5,7,6]. These works concurrently introduced similar Bayesian principal-agent models, in order to study their computational properties. In particular, Guruganesh et al. [5] focus on a model in which the unknown agent's type determines the probability distributions associated with the agent's actions. They analyze linear contracts by extending the work of Dütting et al. [8] from non-Bayesian to Bayesian settings, showing how their efficiency is affected by problem parameters. Guruganesh et al. [5] also investigate the computational complexity of computing an optimal (single) contract and an optimal menu of (deterministic) contracts, showing that both problems are APX-hard even in instances with a constant number of actions. Castiglioni et al. [6] take a more computational-oriented approach than that of [5], by analyzing the efficiency of linear contracts in Bayesian settings with respect to the more realistic benchmark defined as the best among tractable contracts, *i.e.*, those computable in polynomial time. Furthermore, Guruganesh et al. [5] only compare contracts in multiplicative terms, while Castiglioni et al. [6] investigate bi-approximation (*i.e.*, both multiplicative and additive) guarantees. Finally, Alon et al. [7] study a specific Bayesian principal-agent setting in which the agent's type is single-dimensional. In particular, they show that, in their setting, an optimal menu of (deterministic) contracts can be computed in polynomial time when the number of actions is constant. Moreover, in the following preprint [23], the same authors introduce menus of randomized contracts in their setting, showing an example in which randomization makes the principal better off by increasing their expected utility with respect to menus of deterministic contracts. The importance of randomization in principal-agent problem has also been investigated in [24]. Kadan et al. [24] provide several argument showing the potential benefits of randomization both theoretically and practically.

Works on the generalized principal-agent problem Our work is related to the generalized principal-agent problem introduced by Myerson [25]. Myerson [25] studies a generalization of principal-agent problems studied in contract theory, where the principal has an arbitrary, finite set of actions available rather than contracts. Our model is *not* a special case of the model in [25], since the principal's decision space (the set of possible contracts) is continuous and unbounded. Our model has different properties from the one studied in [25], such as the non-existence of an optimal principal's strategy. Gan et al. [26] generalize the model of Myerson [25] to general principal's action spaces.¹ Such a model is general enough to represent our problem of designing menus of randomized contracts. They provide an alternative algorithm to compute an optimal menu of contracts with a linear programming formulation, avoiding the practically inefficient ellipsoid method.

2. Preliminaries

In this section, we introduce all the elements needed for the rest of the paper. Section 2.1 formally describes standard Bayesian principal-agent problems, while Section 2.2 introduces our setting in which the principal proposes to the agent a menu of randomized contracts to choose from.

2.1. The Bayesian principal-agent problem

An instance of the *Bayesian principal-agent problem* is defined by a tuple (Θ, A, Ω) , where: Θ is a finite set of $\ell := |\Theta|$ agent's types; A is a finite set of $n := |A|$ agents' actions; and Ω is a finite set of $m := |\Omega|$ possible outcomes.² The agent's type is drawn according to a fixed probability distribution known to the principal. We let $\mu \in \Delta_\Theta$ be such a distribution, with μ_θ denoting the probability of type $\theta \in \Theta$ being selected.³ For every type $\theta \in \Theta$, we denote by $F_{\theta,a} \in \Delta_\Omega$ the probability distribution over outcomes Ω when an agent of type θ selects action $a \in A$, while $c_{\theta,a} \in [0, 1]$ is the agent's cost for that action.⁴ For the ease of notation, we let $F_{\theta,a,\omega}$ be the probability that $F_{\theta,a}$ assigns to outcome $\omega \in \Omega$, so that $\sum_{\omega \in \Omega} F_{\theta,a,\omega} = 1$. Each outcome $\omega \in \Omega$ has a reward $r_\omega \in [0, 1]$ for the principal. As a result, when an agent of type $\theta \in \Theta$ selects an action $a \in A$, then the principal achieves an expected reward of $\sum_{\omega \in \Omega} F_{\theta,a,\omega} r_\omega$.

In the standard model, the principal commits to a contract maximizing their expected utility. A *contract* specifies payments from the principal to the agent, which are contingent on the actual outcome achieved with the agent's action. We formally define a contract by a vector $p \in \mathbb{R}_+^m$, whose components p_ω represent payments associated with outcomes $\omega \in \Omega$. The assumption that payments are non-negative (*i.e.*, they can only be from the principal to the agent, and *not* the other

¹ The work of Gan et al. [26] is subsequent to the conference version of this work [27].

² For ease of presentation, we assume that all the agent types share the same action set. All the results in this paper can be easily extended to the case in which each agent's type $\theta \in \Theta$ has their own action set A_θ .

³ Given a finite set X , we denote with Δ_X the set of all the probability distributions defined over X .

⁴ In the rest of this work, we assume that rewards and costs are in $[0, 1]$. All the results in this paper can be easily generalized to the case of an arbitrary range of positive numbers, by applying a suitable normalization.

way around) is known as *limited liability* and it is common in contract theory [28]. When an agent of type $\theta \in \Theta$ selects an action $a \in A$, then the expected payment to the agent is $\sum_{\omega \in \Omega} F_{\theta,a,\omega} p_\omega$, while their utility is $\sum_{\omega \in \Omega} F_{\theta,a,\omega} p_\omega - c_{\theta,a}$. On the other hand, the principal's expected utility in that case is $\sum_{\omega \in \Omega} F_{\theta,a,\omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta,a,\omega} p_\omega$.

Given a contract $p \in \mathbb{R}_+^m$, an agent of type $\theta \in \Theta$ plays a *best response*, which is an action that is:

1. *incentive compatible* (IC), i.e., it maximizes their expected utility over actions in A ; and
2. *individually rational* (IR), i.e., it has non-negative expected utility (if there is no IR action, then the agent abstains from playing so as to preserve the *status quo*).

In the rest of this work, we make the following w.l.o.g. common assumption guaranteeing that IR is always enforced [8]. This allows us to focus on IC only. Intuitively, the following assumption ensures that each agent's type has always an action providing them with a non-negative utility, thus ensuring IR of any IC action.

Assumption 1. *There exists an action $a \in A$ such that $c_{\theta,a} = 0$ for all $\theta \in \Theta$.*

Formally, we denote by $\mathcal{B}_p^\theta := \arg \max_{a \in A} \{ \sum_{\omega \in \Omega} F_{\theta,a,\omega} p_\omega - c_{\theta,a} \}$ the set of *best responses* of an agent of type $\theta \in \Theta$ under a contract $p \in \mathbb{R}_+^m$, i.e., given Assumption 1, the set of all the actions that are IC for an agent of type θ under contract p . As it is common in the literature (see, e.g., [8]), we assume that the agent breaks ties in favor of the principal, selecting a best response that maximizes the principal's expected utility. In the following, we let $b^\theta : \mathbb{R}_+^m \rightarrow A$ be a function returning the best responses played by an agent of type $\theta \in \Theta$, where, for any contract $p \in \mathbb{R}_+^m$, we define $b^\theta(p) \in \arg \max_{a \in \mathcal{B}_p^\theta} \{ \sum_{\omega \in \Omega} F_{\theta,a,\omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta,a,\omega} p_\omega \}$.

2.2. Menus of randomized contracts

We study Bayesian principal-agent problems in which there is an additional *type-reporting* stage in which the principal proposes to the agent a menu of randomized contracts to choose from.

A *randomized contract* is a probability distribution γ over \mathbb{R}_+^m , i.e., over the set of vectors $p \in \mathbb{R}_+^m$ representing all the possible contracts. We use $p \sim \gamma$ to denote that the (random) contract p is distributed according to γ , and write $\mathbb{E}_{p \sim \gamma} [\cdot]$ to indicate the expectation taken with respect to the randomness of p . We denote by $\text{supp}(\gamma)$ the support of γ . When γ has a finite support, i.e., $|\text{supp}(\gamma)| < \infty$, we let γ_p be the probability that γ assigns to contract $p \in \mathbb{R}_+^m$.

A *menu of randomized contracts* is defined by a tuple $\Gamma = (\gamma^\theta)_{\theta \in \Theta}$ specifying a probability distribution γ^θ over \mathbb{R}_+^m for each agent's type $\theta \in \Theta$.

The interaction between the principal and an agent of type $\theta \in \Theta$ goes as follows:

- (i) the principal publicly commits to a menu $\Gamma = (\gamma^\theta)_{\theta \in \Theta}$ of randomized contracts;
- (ii) the agent reports a type $\hat{\theta} \in \Theta$ to the principal, possibly different from the true type θ ;
- (iii) the principal draws a contract $p \sim \gamma^{\hat{\theta}}$ and communicates it to the agent;
- (iv) the agent plays the best-response action $b^{\hat{\theta}}(p)$.

The goal of the principal is to commit to a utility-maximizing menu of randomized contracts, selecting among those that are *dominant-strategy incentive compatible* (DSIC).⁵ Formally, a menu $\Gamma = (\gamma^\theta)_{\theta \in \Theta}$ of randomized contracts is DSIC if the following holds:

$$\mathbb{E}_{p \sim \gamma^\theta} \left[\sum_{\omega \in \Omega} F_{\theta,b^\theta(p),\omega} p_\omega - c_{\theta,b^\theta(p)} \right] \geq \mathbb{E}_{p \sim \gamma^{\hat{\theta}}} \left[\sum_{\omega \in \Omega} F_{\theta,b^{\hat{\theta}}(p),\omega} p_\omega - c_{\theta,b^{\hat{\theta}}(p)} \right] \quad \forall \hat{\theta} \neq \theta \in \Theta. \tag{1}$$

Intuitively, the conditions above guarantee that the expected utility of an agent of type $\theta \in \Theta$ under the randomized contract γ^θ is greater than or equal to that obtained under $\gamma^{\hat{\theta}}$, for any $\hat{\theta} \neq \theta \in \Theta$. This ensures that the agent is always better off reporting their true type to the principal. Then, the principal's goal is to find a menu $\Gamma = (\gamma^\theta)_{\theta \in \Theta}$ that is optimal for the following problem:

$$\max_{\Gamma = (\gamma^\theta)_{\theta \in \Theta}} \sum_{\theta \in \Theta} \mu_\theta \mathbb{E}_{p \sim \gamma^\theta} \left[\sum_{\omega \in \Omega} F_{\theta,b^\theta(p),\omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta,b^\theta(p),\omega} p_\omega \right] \quad \text{s.t.} \tag{2}$$

Equation (1),

⁵ Notice that, by a revelation-principle-style argument (see the book by Shoham and Leyton-Brown [29] for some examples of these kinds of arguments), focusing on DSIC menus of contracts is w.l.o.g. when looking for a principal-optimal menu.

whose objective is the principal's expected utility for a DSIC menu of randomized contracts.

We also consider the case, already investigated by Guruganesh et al. [5], in which the menu of contracts is made by non-randomized contracts. Formally, we denote a *menu of deterministic contracts* by a tuple $P = (p^\theta)_{\theta \in \Theta}$, containing a contract $p^\theta \in \mathbb{R}_+^m$ for each agent's type $\theta \in \Theta$. In this case, the principal does *not* draw a contract in step (iii) of the interaction, but they simply adopt contract $p^{\hat{\theta}}$, where $\hat{\theta} \in \Theta$ is the type reported by the agent. Thus, all the definitions above can be specialized to the case of menus of deterministic contracts by dropping expectations.

3. How to find an “almost-optimal” menu of randomized contracts efficiently

We start by providing our main positive result. In particular, we show that randomization enables the polynomial-time computation of a DSIC menu of contracts providing the principal with an expected utility arbitrarily close to the best possible one.

As we show later in this section, the impossibility of designing an optimal menu of randomized contracts does *not* stem from computational challenges associated with the problem, but it is rather due to the fact that an optimal menu of randomized contracts may *not* exist. Indeed, generally speaking, the problem of finding a DSIC menu of randomized contracts maximizing the principal's expected utility only admits a supremum, and it may *not* admit a maximum. The supremum of the problem of finding a principal-utility-maximizing DSIC menu of randomized contracts can be computed by solving the following problem:

$$\sup_{\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}} \sum_{\theta \in \Theta} \mu_\theta \mathbb{E}_{p \sim \gamma^\theta} \left[\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega \right] \quad \text{s.t.} \tag{3a}$$

$$\mathbb{E}_{p \sim \gamma^\theta} \left[\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right] \geq \mathbb{E}_{p \sim \gamma^{\hat{\theta}}} \left[\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right] \quad \forall \theta \neq \hat{\theta} \in \Theta. \tag{3b}$$

As a first step, we show that the supremum defined in Problem (3) always assumes a finite value (i.e., it is never $+\infty$). In the following, for ease of notation, we denote by SUP the value of Problem (3). Then, we can formally prove:

Proposition 1. *Problem (3) has always value SUP $\in [0, 1]$.*

Proof. It is sufficient to observe that the maximum principal's expected utility in a menu is 1, since the rewards are in $[0, 1]$ and the payments can only provide a negative utility contribution to the principal. Let $p^* \in \mathbb{R}_+^m$ be a contract such that $p_\omega^* = 0$ for every $\omega \in \Omega$. Then, it is easy to check that $\text{SUP} \geq 0$, since the menu of randomized contracts that set $\gamma_{p^*}^\theta = 1$ for every $\theta \in \Theta$ provides the principal with an expected utility of at least 0. This proves the statement. \square

Next, we show that Problem (3) may *not* admit a maximum, i.e., there is no contract for which the (finite) value of the supremum is attained.

Before proving such a result, we introduce some additional notation. In particular, we let $\mathcal{P}^{\theta, a} := \{p \in \mathbb{R}_+^m \mid a \in \mathcal{B}_p^\theta\}$ be the set of contracts such that action $a \in A$ is IC for an agent of type $\theta \in \Theta$. Furthermore, we let $\hat{\mathcal{P}}^{\theta, a} := \{p \in \mathbb{R}_+^m \mid b^\theta(p) = a\}$ be the set of contracts in which an agent of type θ plays action a . Notice that the sets $\mathcal{P}^{\theta, a}$ are closed polyhedra (as they can be defined by a system of linear inequalities), while sets $\hat{\mathcal{P}}^{\theta, a}$ are not open nor closed polyhedra (as they can be defined by a system of linear inequalities, some of which are strict due to the tie-breaking rule).

We also need to prove a preliminary result (Lemma 1), which intuitively states that one can restrict the attention to randomized contracts by placing positive probability on a finite number of contracts. In particular, we prove that, for each type $\theta \in \Theta$, it is sufficient that the support of γ^θ contains at most one contract for each agent's action $a \in A$, with the latter being the action played by an agent of type θ in such a contract. Formally:

Lemma 1. *Given any DSIC menu of randomized contract $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$, there always exists a DSIC menu of randomized contracts $\bar{\Gamma} = \{\bar{\gamma}^\theta\}_{\theta \in \Theta}$ that provides the principal with at least the same expected utility as $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$ and such that, for every $\theta \in \Theta$, it holds $|\text{supp}(\bar{\gamma}^\theta) \cap \hat{\mathcal{P}}^{\theta, a}| \leq 1$ for all $a \in A$.*

Equipped with the result in Lemma 1, we can now show that Problem (3) does *not* admit a maximum in general. Formally:

Theorem 1. *There exist Bayesian principal-agent problem instances for which Problem (3) does not admit a maximum.*

Table 1

Description of the rewards, probability distributions of the actions, and action costs for the agent's types in the proof of Theorem 1.

			ω_1	ω_2	ω_3	ω_4	
$c_{\theta,a}$			1	$\frac{3}{4}$	0	0	
θ_1	a_1	0	$F_{\theta_1,a_1,\omega}$	1	0	0	0
	a_2	0	$F_{\theta_1,a_2,\omega}$	0	0	1	0
	a_3	0	$F_{\theta_1,a_3,\omega}$	0	0	1	0
θ_2	a_1	0	$F_{\theta_2,a_1,\omega}$	1	0	0	0
	a_2	0	$F_{\theta_2,a_2,\omega}$	0	1	0	0
	a_3	0	$F_{\theta_2,a_3,\omega}$	0	0	0	1
θ_3	a_1	$\frac{1}{4}$	$F_{\theta_3,a_1,\omega}$	0	1	0	0
	a_2	0	$F_{\theta_3,a_2,\omega}$	0	0	1	0
	a_3	0	$F_{\theta_3,a_3,\omega}$	0	0	1	0

Proof. Let us consider an instance defined as follows (see Table 1). There are three possible agent's types, namely $\Theta = \{\theta_1, \theta_2, \theta_3\}$, with $\mu_{\theta_1} = \mu_{\theta_2} = \mu_{\theta_3} = \frac{1}{3}$. The agent has three actions available, namely $A = \{a_1, a_2, a_3\}$, while the set of possible outcomes is $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$. Type θ_1 is such that $F_{\theta_1,a_1,\omega_1} = 1, c_{\theta_1,a_1} = 0, F_{\theta_1,a_2,\omega_3} = 1, c_{\theta_1,a_2} = 0, F_{\theta_1,a_3,\omega_3} = 1,$ and $c_{\theta_1,a_3} = 0$. Type θ_2 is such that $F_{\theta_2,a_1,\omega_1} = 1, c_{\theta_2,a_1} = 0, F_{\theta_2,a_2,\omega_2} = 1, c_{\theta_2,a_2} = 0, F_{\theta_2,a_3,\omega_4} = 1,$ and $c_{\theta_2,a_3} = 0$. Type θ_3 is such that $F_{\theta_3,a_1,\omega_2} = 1, c_{\theta_3,a_1} = \frac{1}{4}, F_{\theta_3,a_2,\omega_3} = 1, c_{\theta_3,a_2} = 0, F_{\theta_3,a_3,\omega_3} = 1,$ and $c_{\theta_3,a_3} = 0$. Finally, the principal's rewards are $r_{\omega_1} = 1, r_{\omega_2} = \frac{3}{4},$ and $r_{\omega_3} = r_{\omega_4} = 0$.

As a first step, we show that, for every $\epsilon > 0$, there exists a DSIC menu of randomized contracts with principal's expected utility at least $\frac{3}{4} - \epsilon$. Let $p^1 \in \mathbb{R}_+^m$ be a contract such that $p^1_{\omega} = 0$ for all $\omega \in \Omega, p^2 \in \mathbb{R}_+^m$ be such that $p^2_{\omega_4} = \frac{1}{12\epsilon}$ and $p^2_{\omega} = 0$ for all $\omega \neq \omega_4,$ and $p^3 \in \mathbb{R}_+^m$ be such that $p^3_{\omega_2} = \frac{1}{4}$ and $p^3_{\omega} = 0$ for all $\omega \neq \omega_2$. Let us consider the menu $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$ defined so that $\gamma_{p^1}^{\theta_1} = 1, \gamma_{p^2}^{\theta_2} = 3\epsilon, \gamma_{p^1}^{\theta_2} = 1 - 3\epsilon,$ and $\gamma_{p^3}^{\theta_3} = 1$. It is easy to check that such a menu is DSIC. Moreover, the principal's expected utility is at least of

$$\frac{1}{3} \left[\gamma_{p^1}^{\theta_1} + \gamma_{p^1}^{\theta_2} - \frac{1}{12\epsilon} \gamma_{p^2}^{\theta_2} + \left(\frac{3}{4} - \frac{1}{4} \right) \gamma_{p^3}^{\theta_3} \right] = \frac{1}{3} \left[1 + 1 - 3\epsilon - \frac{1}{4} + \frac{3}{4} - \frac{1}{4} \right] = \frac{3}{4} - \epsilon.$$

Hence, SUP is at least $\frac{3}{4}$ for the considered instance.

We conclude the proof by showing that any DSIC menu $\bar{\Gamma} = \{\bar{\gamma}^\theta\}_{\theta \in \Theta}$ results in a principal's expected utility that is strictly smaller than $\frac{3}{4}$. By contradiction, we assume that there exists a DSIC menu $\bar{\Gamma} = \{\bar{\gamma}^\theta\}_{\theta \in \Theta}$ of randomized contracts that provides the principal with expected utility greater than or equal to $\frac{3}{4}$.

First, by Lemma 1, we can focus on menus $\bar{\Gamma} = \{\bar{\gamma}^\theta\}_{\theta \in \Theta}$ such that for every $\theta \in \Theta,$ it holds $|\text{supp}(\bar{\gamma}^\theta) \cap \hat{\mathcal{P}}^{\theta,a}| \leq 1$ for all $a \in A$. If $|\text{supp}(\bar{\gamma}^{\theta_3}) \cap \hat{\mathcal{P}}^{\theta_3,a_1}| = 1,$ let $p^4 \in \mathbb{R}_+^m$ be the only contract in $\hat{\mathcal{P}}^{\theta_3,a_1}$ such that $\bar{\gamma}_{p^4}^{\theta_3} > 0$. Notice that $p^4_{\omega_2} \geq \frac{1}{4}$ by definition of $\hat{\mathcal{P}}^{\theta_3,a_1}$. Instead, if $|\text{supp}(\bar{\gamma}^{\theta_3}) \cap \hat{\mathcal{P}}^{\theta_3,a_1}| = 0,$ let $p^4 \in \mathbb{R}_+^m$ be such that $p^4_{\omega} = 0$ for all $\omega \in \Omega$. By the DSIC conditions for type $\theta_2,$ we have that type θ_2 gets an expected payment at least of $\frac{1}{4} \bar{\gamma}_{p^4}^{\theta_3}$. Moreover, if $|\text{supp}(\bar{\gamma}^{\theta_2}) \cap \hat{\mathcal{P}}^{\theta_2,a_1}| = 1,$ let $p^5 \in \mathbb{R}_+^m$ be the only contract in $\hat{\mathcal{P}}^{\theta_2,a_1}$ such that $\bar{\gamma}_{p^5}^{\theta_2} > 0$. Otherwise, let $p^5 \in \mathbb{R}_+^m$ be such that $p^5_{\omega} = 0$ for all $\omega \in \Omega$. Finally, if $|\text{supp}(\bar{\gamma}^{\theta_2}) \cap \hat{\mathcal{P}}^{\theta_2,a_2}| = 1,$ let $p^6 \in \mathbb{R}_+^m$ be the only contract in $\hat{\mathcal{P}}^{\theta_2,a_2}$ such that $\bar{\gamma}_{p^6}^{\theta_2} > 0;$ otherwise, let $p^6 \in \mathbb{R}_+^m$ be such that $p^6_{\omega} = 0$ for every $\omega \in \Omega$.

Now, suppose that action a_1 is incentivized for type θ_2 with probability strictly smaller than 1, i.e., $\bar{\gamma}_{p^5}^{\theta_2} < 1$. Then, the overall principal's expected utility is at most

$$\frac{1}{3} \left[1 + \bar{\gamma}_{p^5}^{\theta_2} + \frac{3}{4} \bar{\gamma}_{p^6}^{\theta_2} - \frac{1}{4} \bar{\gamma}_{p^4}^{\theta_3} + \left(\frac{3}{4} - \frac{1}{4} \right) \bar{\gamma}_{p^4}^{\theta_3} \right] = \frac{1}{3} \left[1 + \bar{\gamma}_{p^5}^{\theta_2} + \frac{3}{4} \bar{\gamma}_{p^6}^{\theta_2} + \frac{1}{4} \bar{\gamma}_{p^4}^{\theta_3} \right] < \frac{1}{3} \cdot \frac{9}{4} = \frac{3}{4},$$

which contradicts our initial assumption on $\bar{\Gamma} = \{\bar{\gamma}^\theta\}_{\theta \in \Theta}$.

As a result, it must hold $\bar{\gamma}_{p^5}^{\theta_2} = 1,$ and, by the DSIC conditions for type $\theta_2,$ it must be the case that $p^5_{\omega_1} \geq \frac{1}{4} \bar{\gamma}_{p^4}^{\theta_3}$. Then, by the DSIC conditions for type $\theta_1:$

$$\mathbb{E}_{p \sim \bar{\gamma}^{\theta_1}} \left[\sum_{\omega \in \Omega} F_{\theta,b^{\theta_1}(p),\omega} p_{\omega} \right] \geq \bar{\gamma}_{p^5}^{\theta_2} p^5_{\omega_1} \geq \frac{1}{4} \bar{\gamma}_{p^4}^{\theta_3}.$$

Thus, the overall principal's expected utility must be at most

$$\frac{1}{3} \left[1 - \frac{1}{4} \bar{\gamma}_{p^4}^{\theta_3} + 1 - \bar{\gamma}_{p^4}^{\theta_3} \frac{1}{4} + \bar{\gamma}_{p^4}^{\theta_3} \left(\frac{3}{4} - \frac{1}{4} \right) \right] = \frac{1}{3} [1 + 1] = \frac{2}{3},$$

which contradicts our initial assumption on $\bar{\Gamma}$, concluding the proof. \square

In the remaining part of this section, we show how to design an “almost-optimal” DSIC menu of randomized contracts, which is one providing the principal’s expected utility arbitrarily close to SUP. First, we show that bounded payments are sufficient to provide an almost-optimal solution. Formally:

Lemma 2. *For every $\epsilon > 0$, there always exists a DSIC menu of randomized contracts $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$ with principal’s expected utility at least $\text{SUP} - \epsilon$ such that, for every $\theta \in \Theta$, it holds (i) $|\text{supp}(\gamma^\theta) \cap \hat{\mathcal{P}}^{\theta, a}| \leq 1$ for all $a \in A$, and (ii) $p_\omega \leq C(I, \epsilon)$ for all $p \in \text{supp}(\gamma^\theta)$ and $\omega \in \Omega$, where $C(I, \epsilon) \in O(\frac{1}{\epsilon} \cdot 2^{\text{poly}(I)})$ and I denotes the size of the problem instance.*

Given a Bayesian principal-agent problem instance (Θ, A, Ω) and an $\epsilon > 0$, we let $C(I, \epsilon)$ be defined as in Lemma 2. Moreover, let $\mathcal{P}^\epsilon := [0, C(I, \epsilon)]^m$. For every $\mathbf{a} = (a_\theta)_{\theta \in \Theta} \in \times_{\theta \in \Theta} A$, where \mathbf{a} is a tuple specifying an action a_θ for each agent’s type $\theta \in \Theta$, we let $\mathcal{P}^{\mathbf{a}, \epsilon} := \mathcal{P}^\epsilon \cap \left(\bigcap_{\theta \in \Theta} \mathcal{P}^{\theta, a_\theta} \right)$ and, additionally, $\hat{\mathcal{P}}^{\mathbf{a}, \epsilon} := \mathcal{P}^\epsilon \cap \left(\bigcap_{\theta \in \Theta} \hat{\mathcal{P}}^{\theta, a_\theta} \right)$. We also define the set of all the vertexes of the closed polyhedra $\mathcal{P}^{\mathbf{a}, \epsilon}$ as $\mathcal{P}^{*, \epsilon} := \bigcup_{\mathbf{a} \in \times_{\theta \in \Theta} A} V(\mathcal{P}^{\mathbf{a}, \epsilon})$, where $V(\cdot)$ denotes the set of vertexes of the polyhedron given as input. Then, we can prove the following last lemma, which shows that for every $\epsilon > 0$, in order to obtain a DSIC menu of randomized contracts with principal’s expected utility at least $\text{SUP} - \epsilon$, it is sufficient to restrict the attention to randomized contracts placing positive probability only on contracts in the (finite) set $\mathcal{P}^{*, \epsilon}$.

Lemma 3. *For every $\epsilon > 0$, there always exists a DSIC menu of randomized contracts supported on $\mathcal{P}^{*, \epsilon}$ with principal’s expected utility at least $\text{SUP} - \epsilon$.*

Lemma 3 allows us to formulate the following LP which, given an $\epsilon > 0$, has an optimal solution of value at least $\text{SUP} - \epsilon$.

$$\max_{\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}} \sum_{\theta \in \Theta} \mu_\theta \sum_{p \in \mathcal{P}^{*, \epsilon}} \gamma_p^\theta \left[\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega \right] \quad \text{s.t.} \quad (4a)$$

$$\sum_{p \in \mathcal{P}^{*, \epsilon}} \gamma_p^\theta \left[\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right] \geq \sum_{p \in \mathcal{P}^{*, \epsilon}} \gamma_p^{\hat{\theta}} \left[\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right] \quad \forall \theta \neq \hat{\theta} \in \Theta \quad (4b)$$

$$\sum_{p \in \mathcal{P}^{*, \epsilon}} \gamma_p^\theta = 1 \quad \forall \theta \in \Theta. \quad (4c)$$

Notice that LP 4 has an exponential number of variables, since the probability distributions γ^θ are defined over contracts in $\mathcal{P}^{*, \epsilon}$, and these may be exponentially many in the size of the problem instance.

Nevertheless, the LP has polynomially many constraints, and, thus, as we show next, we can solve it in polynomial time by applying the *ellipsoid algorithm* to its dual program, which features polynomially-many variables and exponentially-many constraints.⁵

The dual of LP 4 reads as follows:

$$\min_{y \leq 0, t} \sum_{\theta \in \Theta} t_\theta \quad \text{s.t.} \quad (5a)$$

$$\sum_{\hat{\theta} \in \Theta: \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right) - \sum_{\hat{\theta} \in \Theta: \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\hat{\theta}, b^{\hat{\theta}}(p), \omega} p_\omega - c_{\hat{\theta}, b^{\hat{\theta}}(p)} \right) + t_\theta \geq \mu_\theta \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega \right) \quad \forall \theta \in \Theta, \forall p \in \mathcal{P}^{*, \epsilon}, \quad (5b)$$

⁵ An LP analogous to LP 4 also arises when dealing with Bayesian persuasion problems in which the receiver can be of multiple types (see [27]). However, in that case, the counterparts of the terms γ_p^θ and p_ω (namely, for a receiver’s type k , the probability γ_ξ^k of a given posterior ξ and the probability ξ_θ that the posterior assigns to a given state of nature θ) always appear in a product $\gamma_p^\theta p_\omega$, which can be replaced by a suitably-defined new variable (subject to some consistency constraints). This property of the Bayesian persuasion setting allows to formulate the problem as an LP with polynomially-many variables and constraints. However, in our principal-agent setting, the variables γ_p^θ also appear without being multiplied by p_ω , which prevents us from applying the same trick. As a result, in our setting, we had to resort to the ellipsoid algorithm.

where $y \in \mathbb{R}^{\ell(\ell-1)}$ is a vector of dual variables whose components $y_{\theta, \hat{\theta}}$ for $\theta, \hat{\theta} \in \Theta : \theta \neq \hat{\theta}$ are the dual variables corresponding to Constraints (4b), while $t \in \mathbb{R}^\ell$ is another vector of dual variables whose components t_θ for $\theta \in \Theta$ are the dual variables of Constraints (4c).

The dual LP 5 has polynomially-many variables and exponentially-many constraints, and it can be optimally solved in polynomial time by means of the ellipsoid algorithm provided that a suitable polynomial-time *separation oracle* is available. In particular, given a pair (y, t) assigning values to the variables of the dual LP, the separation oracle that we provide outputs a pair $\theta \in \Theta$, $p \in \mathcal{P}^{*,\epsilon}$ such that the corresponding inequality in Constraints (5b) is violated for (y, t) if any; otherwise, the oracle concludes that (y, t) is feasible for the dual LP. Formally, the separation oracle that we provide solves (in polynomial time) the following optimization problem for each $\theta \in \Theta$:

$$\begin{aligned} \max_{p \in \mathcal{P}^{*,\epsilon}} \left\{ \mu_\theta \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega \right) - \sum_{\hat{\theta} \in \Theta : \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right) \right. \\ \left. + \sum_{\hat{\theta} \in \Theta : \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\hat{\theta}, b^{\hat{\theta}}(p), \omega} p_\omega - c_{\hat{\theta}, b^{\hat{\theta}}(p)} \right) \right\}. \end{aligned} \tag{6}$$

Indeed, if the value of the maximization above is greater than t_θ for some $\theta \in \Theta$, then the separation oracle outputs the pair (θ, p) with $p \in \mathcal{P}^{*,\epsilon}$ being a contract for which the maximization is attained since the constraint corresponding to θ and p is violated. Instead, if t_θ is less than or equal to the value of the maximization above for every $\theta \in \Theta$, then the oracle concludes that (y, t) is feasible since no constraint is violated.

The following lemma shows that Problem (6) above can indeed be solved in polynomial time. Notice that ϵ appears only in the coefficients of Problem (4). Indeed, such coefficients are at most $C(I, \epsilon)$. Hence, our algorithm runs in time polynomial in the number of bits needed to represent the coefficients, and, thus, polynomial in $\log(1/\epsilon)$. Formally, we can state the following:

Lemma 4. *For every $\epsilon > 0$, there exists a separation oracle for LP 5, that runs in time $\text{poly}(I, \log(1/\epsilon))$, where I is the size of the problem instance.*

Proof. Given a pair $(y, t) \in \mathbb{R}^{\ell(\ell-1)} \times \mathbb{R}^\ell$ assigning values to the variables of LP 5, the separation oracle solves Problem (6) for every $\theta \in \Theta$. Then, if there exists an agent's type $\theta \in \Theta$ such that t_θ is less than the value of an optimal solution $p \in \mathcal{P}^{*,\epsilon}$ to Problem (6) for type θ , the oracle outputs the pair (θ, p) , which corresponds to a violated inequality in Constraints (5b). Instead, if t_θ is greater than or equal to the value of an optimal solution to Problem (6) for every $\theta \in \Theta$, the oracle concludes that (y, t) is feasible since all the inequalities in Constraints (5b) are satisfied.

Next, we show that, for every $\theta \in \Theta$, Problem (6) can be solved in polynomial time. In order to do that, we split the set $\mathcal{P}^{*,\epsilon}$ into the subsets $\hat{\mathcal{P}}^{\theta,a} \cap \mathcal{P}^{*,\epsilon}$, defined for every agents' action $a \in A$. Notice that, by the definitions of $\hat{\mathcal{P}}^{\theta,a}$ and $\mathcal{P}^{*,\epsilon}$, each $p \in \mathcal{P}^{*,\epsilon}$ belongs to exactly one subset $\hat{\mathcal{P}}^{\theta,a} \cap \mathcal{P}^{*,\epsilon}$. Thus, solving Problem (6) reduces to solving the following problem:

$$\begin{aligned} \max_{a \in A} \max_{p \in \hat{\mathcal{P}}^{\theta,a} \cap \mathcal{P}^{*,\epsilon}} \left\{ \mu_\theta \left(\sum_{\omega \in \Omega} F_{\theta, a, \omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta, a, \omega} p_\omega \right) - \sum_{\hat{\theta} \in \Theta : \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, a, \omega} p_\omega - c_{\theta, a} \right) \right. \\ \left. + \sum_{\hat{\theta} \in \Theta : \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\hat{\theta}, b^{\hat{\theta}}(p), \omega} p_\omega - c_{\hat{\theta}, b^{\hat{\theta}}(p)} \right) \right\}, \end{aligned} \tag{7}$$

where we used the fact that $b^\theta(p) = a$ for $p \in \hat{\mathcal{P}}^{\theta,a}$.

As a first step, we prove that Problem (7) is equivalent to the following one:

$$\begin{aligned} \max_{a \in A} \max_{p \in \mathcal{P}^{\theta,a} \cap \mathcal{P}^{*,\epsilon}} \left\{ \mu_\theta \left(\sum_{\omega \in \Omega} F_{\theta, a, \omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta, a, \omega} p_\omega \right) - \sum_{\hat{\theta} \in \Theta : \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, a, \omega} p_\omega - c_{\theta, a} \right) \right. \\ \left. + \sum_{\hat{\theta} \in \Theta : \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\hat{\theta}, b^{\hat{\theta}}(p), \omega} p_\omega - c_{\hat{\theta}, b^{\hat{\theta}}(p)} \right) \right\}, \end{aligned} \tag{8}$$

where we replace $\hat{\mathcal{P}}^{\theta,a}$ with $\mathcal{P}^{\theta,a}$ in the inner maximization. Indeed, by adding contracts $p \in (\mathcal{P}^{\theta,a} \setminus \hat{\mathcal{P}}^{\theta,a}) \cap \mathcal{P}^{*,\epsilon}$ to the domain of the inner maximizations for every $a \in A$ does *not* change the optimal value of the overall maximization problem. In order to see this, notice that, for every $p \in (\mathcal{P}^{\theta,a} \setminus \hat{\mathcal{P}}^{\theta,a}) \cap \mathcal{P}^{*,\epsilon}$, by definition of best response it

holds $\sum_{\omega \in \Omega} F_{\theta, b^{\hat{\theta}}(p), \omega} p_{\omega} - c_{\theta, b^{\hat{\theta}}(p)} = \sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega} - c_{\theta, a}$ and $\sum_{\omega \in \Omega} F_{\theta, b^{\hat{\theta}}(p), \omega} r_{\omega} - \sum_{\omega \in \Omega} F_{\theta, b^{\hat{\theta}}(p), \omega} p_{\omega} \geq \sum_{\omega \in \Omega} F_{\theta, a, \omega} r_{\omega} - \sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega}$ (tie-breaking rule). Thus, any pair (a, p) with $p \in (\mathcal{P}^{\theta, a} \setminus \hat{\mathcal{P}}^{\theta, a}) \cap \mathcal{P}^{*, \epsilon}$ would result in an objective value smaller than that of the pair $(b^{\hat{\theta}}(p), p)$. Recalling that $p \in \hat{\mathcal{P}}^{\theta, b^{\hat{\theta}}(p)}$ by definition, we can conclude that Problem (7) and Problem (8) are indeed equivalent.

We are left to prove that Problem (8) can be solved in polynomial time. First, notice that solving Problem (8) is equivalent to solving the following problem for every $a \in A$:

$$\begin{aligned} \max_{p \in \mathcal{P}^{\theta, a} \cap \mathcal{P}^{*, \epsilon}} & \left\{ \mu_{\theta} \left(\sum_{\omega \in \Omega} F_{\theta, a, \omega} r_{\omega} - \sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega} \right) - \sum_{\hat{\theta} \in \Theta: \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega} - c_{\theta, a} \right) \right. \\ & \left. + \sum_{\hat{\theta} \in \Theta: \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\hat{\theta}, b^{\hat{\theta}}(p), \omega} p_{\omega} - c_{\hat{\theta}, b^{\hat{\theta}}(p)} \right) \right\}. \end{aligned} \tag{9}$$

The first and second terms in the maximization above are linear in the payments p_{ω} for $\omega \in \Omega$. Moreover, it holds:

$$\sum_{\hat{\theta} \in \Theta: \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\hat{\theta}, b^{\hat{\theta}}(p), \omega} p_{\omega} - c_{\hat{\theta}, b^{\hat{\theta}}(p)} \right) = \sum_{\hat{\theta} \in \Theta: \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \max_{a' \in A} \left\{ \sum_{\omega \in \Omega} F_{\hat{\theta}, a', \omega} p_{\omega} - c_{\hat{\theta}, a'} \right\},$$

and the latter is a concave function of the payments p_{ω} since all the $y_{\theta, \hat{\theta}}$ are negative and the max is convex. These observations allow us to solve Problem (8) in polynomial time by first solving the following LP relaxation obtained by weakening the requirement $p \in \mathcal{P}^{\theta, a} \cap \mathcal{P}^{*, \epsilon}$ as $p \in \mathcal{P}^{\theta, a} \cap \mathcal{P}^{\epsilon}$:

$$\max_{p \in \mathcal{P}^{\epsilon}} \mu_{\theta} \left(\sum_{\omega \in \Omega} F_{\theta, a, \omega} r_{\omega} - \sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega} \right) - \sum_{\hat{\theta} \in \Theta: \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega} - c_{\theta, a} \right) + \sum_{\hat{\theta} \in \Theta: \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} z_{\hat{\theta}} \quad \text{s.t.} \tag{10a}$$

$$z_{\hat{\theta}} \geq \sum_{\omega \in \Omega} F_{\hat{\theta}, a', \omega} p_{\omega} - c_{\hat{\theta}, a'} \quad \forall \hat{\theta} \in \Theta: \hat{\theta} \neq \theta, \forall a' \in A \tag{10b}$$

$$\sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega} - c_{\theta, a} \geq \sum_{\omega \in \Omega} F_{\theta, a', \omega} p_{\omega} - c_{\theta, a'} \quad \forall a' \in A, \tag{10c}$$

where Constraints (10b) ensure that $z_{\hat{\theta}} = \max_{a' \in A} \left\{ \sum_{\omega \in \Omega} F_{\hat{\theta}, a', \omega} p_{\omega} - c_{\hat{\theta}, a'} \right\}$ for every $\hat{\theta} \in \Theta: \hat{\theta} \neq \theta$ (since in Objective (10a) all the coefficients of the variables $z_{\hat{\theta}}$ are negative), while Constraints (10c) are equivalent to $p \in \mathcal{P}^{\theta, a}$. Given an optimal solution to LP 10, it is possible to recover in polynomial time a contract $p \in \mathcal{P}^{\theta, a} \cap \mathcal{P}^{*, \epsilon}$ that is an optimal solution to Problem (9). Given a contract $p^* \in \mathcal{P}^{\theta, a} \cap \mathcal{P}^{\epsilon}$ that is an optimal solution to LP 10, let $\mathbf{a}^* = (a_{\hat{\theta}}^{p^*})$, where $a_{\hat{\theta}}^{p^*} = a$ and $a_{\hat{\theta}}^{p^*} = b^{\hat{\theta}}(p^*)$ for each $\hat{\theta} \in \Theta: \hat{\theta} \neq \theta$. Then, since $a_{\hat{\theta}}^{p^*} = a$, and for each contract $\hat{p} \in \mathcal{P}^{\mathbf{a}^*, \epsilon}$ and each $\hat{\theta} \in \Theta: \hat{\theta} \neq \theta$, $b^{\hat{\theta}}(\hat{p}) \in \mathcal{B}_{\hat{p}}^{\hat{\theta}}$, the following LP is equivalent to LP 10 restricted to the feasibility set $\mathcal{P}^{\mathbf{a}^*, \epsilon} \subseteq \mathcal{P}^{\theta, a} \cap \mathcal{P}^{\epsilon}$. Hence, the contract p^* defined above is also an optimal solution to the following LP:

$$\begin{aligned} \max_{p \in \mathcal{P}^{\epsilon}} & \mu_{\theta} \left(\sum_{\omega \in \Omega} F_{\theta, a, \omega} r_{\omega} - \sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega} \right) - \sum_{\hat{\theta} \in \Theta: \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega} - c_{\theta, a} \right) \\ & + \sum_{\hat{\theta} \in \Theta: \hat{\theta} \neq \theta} y_{\theta, \hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\hat{\theta}, b^{\hat{\theta}}(p^*), \omega} p_{\omega} - c_{\hat{\theta}, b^{\hat{\theta}}(p^*)} \right) \quad \text{s.t.} \end{aligned} \tag{11a}$$

$$\sum_{\omega \in \Omega} F_{\hat{\theta}, a_{\hat{\theta}}^{p^*}, \omega} p_{\omega} - c_{\hat{\theta}, a_{\hat{\theta}}^{p^*}} \geq \sum_{\omega \in \Omega} F_{\hat{\theta}, a', \omega} p_{\omega} - c_{\hat{\theta}, a'} \quad \forall \hat{\theta} \in \Theta, \forall a' \in A. \tag{11b}$$

Thus, there always exists an optimal solution to LP 11 that is a vertex of its feasibility polytope $\mathcal{P}^{\mathbf{a}^*, \epsilon}$, i.e., a contract in $V(\mathcal{P}^{\mathbf{a}^*, \epsilon}) \subseteq \mathcal{P}^{\theta, a} \cap \mathcal{P}^{*, \epsilon}$, and such optimal solution can be found in polynomial time [30]. \square

By the previous lemma, it immediately follows the main result of the paper.

Theorem 2. *In Bayesian principal-agent problems, there exists an algorithm that, given any $\epsilon > 0$, runs in time $\text{poly}(l, \log(1/\epsilon))$ and computes a DSIC menu of randomized contracts providing the principal with an expected utility at least of $SUP - \epsilon$, where l is the size of the problem instance.*

4. A striking negative result for menus of deterministic contracts

In this section, we provide a negative result for the principal's computational problem of finding an optimal menu of deterministic contracts. In particular, we prove that it is NP-hard to provide any constant multiplicative approximation to the optimal principal's expected utility, even if both the number of outcomes and that of the agent's actions are constants. This is in contrast to what happens with the problem of finding an optimal (single) contract in Bayesian instances, which can be solved in polynomial time when the number of outcomes is constant [5,6]. Moreover, our reduction also shows that the problem does *not* admit an additive FPTAS, unless $P = NP$. Let us remark that our reduction considerably strengthens an already-known hardness result for menus of deterministic contracts, which is that by Guruganesh et al. [5], who show that the problem is APX-hard even with a constant number of actions. Indeed, our result shows that the problem is *not* in APX, and that this holds even when both the number of outcomes and that of actions are constants.

Our reduction is from the following *promise problem*, which is related to the INDEPENDENT-SET problem on undirected graphs with bounded-degree nodes.

Definition 1 (*GAP-BOUNDED- $IS_{\alpha,k}$*). For every $\alpha \in [0, 1]$ and $k \in \mathbb{N}_+$, we define *GAP-BOUNDED- $IS_{\alpha,k}$* as the following promise problem:

- **Input** : An undirected graph $G = (V, E)$ in which each vertex has degree at most k and a constant $\eta \in [\frac{1}{k}, 1]$ such that either one of the following is true⁷:
 - there exists an independent set (i.e., a subset of vertices such that there is no edge connecting two of them) of size at least $\eta|V|$;
 - all the independent sets have size at most $\alpha\eta|V|$.
- **Output** : Determine which of the above two cases hold.

Notice that, for every constant $\alpha > 0$, there exists a constant $k = k(\alpha)$ depending on α such that *GAP-BOUNDED- $IS_{\alpha,k}$* is NP-hard [31,32]. This observation is exploited in order to prove the following theorem:

Theorem 3. *In Bayesian principal-agent problems, the following two results hold even when there are only four outcomes and a constant number of actions.*

- For every constant $\alpha > 0$, it is NP-hard to approximate the principal's expected utility in an optimal DSIC menu of deterministic contracts up to within an α multiplicative factor.
- There is no additive FPTAS for the problem of finding an DSIC optimal menu of deterministic contracts, unless $P = NP$.

Proof. We reduce from *GAP-BOUNDED- $IS_{\alpha,k}$* . In the rest of the proof, given an undirected graph $G = (V, E)$, for ease of presentation, we label vertices with natural numbers, so that we can write $V = \{1, \dots, s\}$, where $s := |V|$. Given an instance of *GAP-BOUNDED- $IS_{\alpha,k}$* , letting $l := \lceil \frac{k}{\alpha} \rceil$ and $\rho := s^{-3}$, we build an instance of Bayesian principal-agent problem such that:

- **Completeness**: If there is an independent set of G with size at least ηs , then there is a DSIC menu of deterministic contracts with principal's expected utility at least $\frac{1}{2}\eta\rho l 2^{-l}$.
- **Soundness**: If all the independent sets of G have size at most $\alpha\eta s$, then the principal's expected utility in any DSIC menu of deterministic contracts is at most $4\alpha\eta\rho l 2^{-l}$.

Since, for any $\alpha > 0$, there exists $k = k(\alpha)$ such that *GAP-BOUNDED- $IS_{\alpha,k}$* is NP-hard, the two points above immediately provide our hardness result for multiplicative approximations. Moreover, by noticing that the difference between the principal's expected utilities in the completeness and the soundness cases depends polynomially on the instance size, we also have that there is no additive FPTAS for the problem of finding an optimal menu of deterministic contracts unless $P = NP$.

Construction. There are four outcomes, namely $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, with $r_{\omega_3} = 1$ and all the other rewards equal to zero. There is an agent's type $\theta_v \in \Theta$ for each vertex $v \in V$ of G so that $l = s$. All types are equally probable, and, thus, $\mu_{\theta_v} = \frac{1}{s}$ for every $\theta_v \in \Theta$. Each agent's type $\theta_v \in \Theta$ has an action $a_v \in A_{\theta_v}$, and an additional action $a_{ui} \in A_{\theta_v}$ for every vertex $u \in V$ adjacent to v , i.e., $(v, u) \in E$, and index $i \in [l - 3]$, where $[x] := \{1, \dots, x\}$.⁸ For every $\theta_v \in \Theta$, we let

$$\bullet F_{\theta_v, a_v, \omega_1} = \frac{1}{4} \cos\left(\pi \frac{v}{2s}\right),$$

⁷ For $\eta < \frac{1}{k}$, the problem can be trivially solved since a maximal independent set has size at least $\frac{1}{k}$.

⁸ For the ease of presentation, in the instances that we define in the reduction proving Theorem 3 each agent's type $\theta \in \Theta$ has a different set of actions A_θ . It is easy to modify such instances by adding dummy actions so as to recover "equivalent" instances in which each type has the same set of actions available.

Table 2

Description of the rewards, probability distributions of the actions, and action costs for an agent of type θ_v , with $v \in V$, in the proof of Theorem 3. We denote by $(*)$ the probability with which the actions induce outcome ω_4 , which are such that the probability distributions $F_{\theta_v, a_v, \omega}$ and $F_{\theta_v, a_{ui}, \omega}$ are well defined.

	$c_{\theta_v, a}$	r_ω	ω_1	ω_2	ω_3	ω_4
			0	0	1	0
a_v	$\frac{1}{4} - \rho l 2^{-l}$	$F_{\theta_v, a_v, \omega}$	$\frac{1}{4} \cos(\pi \frac{v}{25})$	$\frac{1}{4} \sin(\pi \frac{v}{25})$	$\frac{1}{4}$	$(*)$
a_{ui}	$2^{-i-2} - \rho(l-i)2^{-l}$	$F_{\theta_v, a_{ui}, \omega}$	$\cos(\pi \frac{u}{25}) 2^{-i-2}$	$\sin(\pi \frac{u}{25}) 2^{-i-2}$	2^{-i-2}	$(*)$
\bar{a}	0	$F_{\theta_v, \bar{a}, \omega}$	0	0	0	1




Fig. 1. Left: Graphical representation of the probability vectors $F_{\theta_v, a}$ used in the proof of Theorem 3 in the ω_1 - ω_2 Cartesian coordinate system. In particular, node $v \in V$ has 2 adjacent nodes $u_1, u_2 \in V$, and $l = 6$. **Right:** Expected payments of contract p^{θ_v} for a vertex $v \in V^*$ in the completeness part of the proof of Theorem 3. The dashed line represents the direction of the vector encoding p^{θ_v} . The expected payments of actions a_v and $a_{u_2 1}$ are given by the projections of their probability vectors on p^{θ_v} , times $\|p^{\theta_v}\|$.

- $F_{\theta_v, a_v, \omega_2} = \frac{1}{4} \sin(\pi \frac{v}{25})$,
- $F_{\theta_v, a_v, \omega_3} = \frac{1}{4}$, and
- $F_{\theta_v, a_v, \omega_4} = 1 - F_{\theta_v, a_v, \omega_1} - F_{\theta_v, a_v, \omega_2} - F_{\theta_v, a_v, \omega_3}$.

Moreover, for every pair of adjacent vertices $v \in V, u \in V$ such that $(v, u) \in E$ and index $i \in [l-3]$, an agent of type θ_v has the following probabilities associated to action a_{ui} :

- $F_{\theta_v, a_{ui}, \omega_1} = \cos(\pi \frac{u}{25}) 2^{-i-2}$,
- $F_{\theta_v, a_{ui}, \omega_2} = \sin(\pi \frac{u}{25}) 2^{-i-2}$,
- $F_{\theta_v, a_{ui}, \omega_3} = 2^{-i-2}$, and
- $F_{\theta_v, a_{ui}, \omega_4} = 1 - F_{\theta_v, a_{ui}, \omega_1} - F_{\theta_v, a_{ui}, \omega_2} - F_{\theta_v, a_{ui}, \omega_3}$.

Finally, each agent's type $\theta_v \in \Theta$ has an additional action $\bar{a} \in A_{\theta_v}$ such that $F_{\theta_v, \bar{a}, \omega_4} = 1$ and $F_{\theta_v, \bar{a}, \omega} = 0$ for all $\omega \neq \omega_4 \in \Omega$. Notice that $|A_{\theta_v}| = O(k^2/\alpha)$ for every $\theta_v \in \Theta$. In order to conclude the construction, we need to define action costs. In particular, for every agent's type $\theta_v \in \Theta$, we set $c_{\theta_v, a_v} = \frac{1}{4} - \rho l 2^{-l}$ and $c_{\theta_v, a_{ui}} = 2^{-i-2} - \rho(l-i)2^{-l}$ for every vertex $u \in V$ such that $(v, u) \in E$ and index $i \in [l-3]$. Furthermore, the cost of action \bar{a} is $c_{\theta_v, \bar{a}} = 0$ for every agent's type $\theta_v \in \Theta$. Table 2 provides a summary of the main components of the instance.

Let us remark that we use the sin and cos trigonometric functions in order to easily represent the vectors $F_{\theta_v, a}$ in the ω_1 - ω_2 Cartesian coordinate system (see Fig. 1). One of the main properties of our construction is that, for a type θ_v , no vector representing an action a_v or a_{ui} is dominated by a convex combination of vectors representing actions with the same costs. This is necessary since an action represented by a dominated vector is never preferred over another action represented by an un-dominated vector, with their costs being the same. Notice that we use the functions sin and cos for ease of exposition. It is easy to change our instances so that our results continue to hold even if we replace such functions with sufficiently good approximations.⁹

Completeness. Suppose that there exists an independent set $V^* \subseteq V$ of G with size at least ηs . From that, we build a DSIC menu of deterministic contracts $P = (p^{\theta_v})_{\theta_v \in \Theta}$ with principal's expected utility at least $\frac{1}{2} \eta \rho l 2^{-l}$. For every agent's type $\theta_v \in \Theta$ such that $v \in V^*$, we define the contract $p^{\theta_v} \in \mathbb{R}_+^m$ so that $p_{\omega_1}^{\theta_v} = \cos(\pi \frac{v}{25}) (1 - \rho l 2^{-l+1})$, $p_{\omega_2}^{\theta_v} = \sin(\pi \frac{v}{25}) (1 - \rho l 2^{-l+1})$,

⁹ Indeed, it is sufficient to use approximations that can be represented in memory with a number of bits that is polynomial in the instance size, and, thus, that can be computed in polynomial time.

and $p_{\omega_3}^{\theta_v} = p_{\omega_4}^{\theta_v} = 0$. Moreover, for every agent's type $\theta_v \in \Theta$ such that $v \notin V^*$, we define $p^{\theta_v} = p^{\theta_u}$ for some $u \in V^*$ such that

$$\theta_u \in \operatorname{arg\,max}_{\theta_u \in \Theta: u \in V^*} \left\{ \max_{a \in A_{\theta_v}} \sum_{\omega \in \Omega} F_{\theta_v, a, \omega} p_{\omega}^{\theta_u} - c_{\theta_v, a} \right\},$$

where the term inside the $\operatorname{arg\,max}$ operator is the expected utility that a utility-maximizing action provides to an agent of type θ_v under a contract p^{θ_u} already defined above, since $u \in V^*$. Intuitively, contract p^{θ_v} represents a vector on the ω_1 - ω_2 plane. An action a_{ui} has a large expected payment if the angle between vector p^{θ_v} and the projection of vector $F_{\theta_v, a_{ui}}$ on the ω_1 - ω_2 plane is small. In particular, the expected payment is $\|p^{\theta_v}\|_2 \|(F_{\theta_v, a_{ui}, \omega_1}, F_{\theta_v, a_{ui}, \omega_2})\|_2 \cos(\beta)$, where β is the angle between the two vectors.

Next, we show that the menu of contracts defined above is DSIC. First, we prove that an agent of type $\theta_v \in \Theta$ with $v \in V^*$ gets an expected utility of $\rho \frac{1}{2} 2^{-l}$ under contract p^{θ_v} , since the best response of the agent under contract p^{θ_v} is a_v , namely $a_v = b^{\theta_v}(p^{\theta_v})$. Indeed, the agent's expected utility by selecting action a_v under contract p^{θ_v} is

$$\frac{1}{4} \cos^2\left(\pi \frac{v}{2s}\right) (1 - \rho l 2^{-l+1}) + \frac{1}{4} \sin^2\left(\pi \frac{v}{2s}\right) (1 - \rho l 2^{-l+1}) - \left(\frac{1}{4} - \rho l 2^{-l}\right) = \rho \frac{1}{2} 2^{-l},$$

where we used $\cos^2 x + \sin^2 x = 1$. All the other actions a_{ui} , for any vertex $u \in V$ such that $(v, u) \in E$ and index $i \in [l-3]$, provide the agent with an expected utility of

$$\begin{aligned} & \cos\left(\pi \frac{u}{2s}\right) 2^{-i-2} \cos\left(\pi \frac{v}{2s}\right) (1 - \rho l 2^{-l+1}) \\ & + \sin\left(\pi \frac{u}{2s}\right) 2^{-i-2} \sin\left(\pi \frac{v}{2s}\right) (1 - \rho l 2^{-l+1}) - (2^{-i-2} - \rho(l-i)2^{-l}) \\ & = \cos\left(\pi \frac{|v-u|}{2s}\right) 2^{-i-2} (1 - \rho l 2^{-l+1}) - (2^{-i-2} - \rho(l-i)2^{-l}) \\ & \leq \left(1 - \frac{1}{\pi} \left(\pi \frac{|v-u|}{2s}\right)^2\right) (1 - \rho l 2^{-l+1}) 2^{-i-2} - (2^{-i-2} - \rho(l-i)2^{-l}) \\ & \leq \left(1 - \left(\frac{|v-u|}{2s}\right)^2\right) (1 - \rho l 2^{-l+1}) 2^{-i-2} - (2^{-i-2} - \rho(l-i)2^{-l}) \\ & \leq 2^{-i-2} - \rho l 2^{-i-l-1} - \frac{1}{4s^2} 2^{-i-2} + \frac{1}{4s^2} \rho l 2^{-i-l-1} - (2^{-i-2} - \rho(l-i)2^{-l}) \\ & = -\rho l 2^{-i-l-1} - \frac{1}{4s^2} 2^{-i-2} + \frac{1}{4s^2} \rho l 2^{-i-l-1} + \rho(l-i)2^{-l} \\ & = 2^{-l} \left(-\rho l 2^{-i-1} - \frac{1}{4s^2} 2^{l-i-2} + \frac{1}{4s^2} \rho l 2^{-i-1} + \rho(l-i)\right), \end{aligned}$$

where in the first equality we used $\sin x \sin y + \cos x \cos y = \cos|x-y|$, in the first inequality the fact that $\cos x \leq 1 - \frac{x^2}{\pi}$ for $x \in [0, \frac{\pi}{2}]$, and in the third inequality $|v-u| \geq 1$. For s large enough, it holds that $\frac{1}{4s^2} 2^{l-i-2} \geq \frac{1}{4s^2} \rho l 2^{-i-l-1} + \rho(l-i)$ (recall that $\rho = s^{-3}$ and l is a constant depending only on k and α), and, thus, the agent's expected utility by playing any action a_{ui} is smaller than or equal to zero, proving that $a_v = b^{\theta_v}(p^{\theta_v})$.

In order to prove that the menu is DSIC, it remains to show that each agent's type $\theta_v \in \Theta$ is *not* better off by reporting a type different from their true type. Let us consider any contract p^{θ_u} with $\theta_u \in \Theta$ such that $u \in V^*$ (for contracts p^{θ_u} with $\theta_u \in \Theta$ such that $u \notin V^*$ the reasoning is analogous given how they are defined). Since V^* is an independent set of G , an agent of type $\theta_v \in \Theta$ has no action a_{ui} available, while, by selecting any action $a_{u'i}$ corresponding to a node $u' \neq u \in V$ and an index $i \in [l-3]$, under contract p^{θ_u} their expected utility is

$$\begin{aligned} & \cos\left(\pi \frac{u'}{2s}\right) 2^{-i-2} \cos\left(\pi \frac{u}{2s}\right) (1 - \rho l 2^{-l+1}) \\ & + \sin\left(\pi \frac{u'}{2s}\right) 2^{-i-2} \sin\left(\pi \frac{u}{2s}\right) (1 - \rho l 2^{-l+1}) - (2^{-i-2} - \rho(l-i)2^{-l}), \end{aligned}$$

which is less than or equal to zero by arguments similar to those employed above. Finally, by playing action a_v under contract p^{θ_u} , the expected utility of an agent of type θ_v is

$$\begin{aligned} & \frac{1}{4} \cos\left(\pi \frac{v}{2s}\right) \cos\left(\pi \frac{u}{2s}\right) (1 - \rho l 2^{-l+1}) + \frac{1}{4} \sin\left(\pi \frac{v}{2s}\right) \sin\left(\pi \frac{u}{2s}\right) (1 - \rho l 2^{-l+1}) - \left(\frac{1}{4} - \rho l 2^{-l}\right) \\ &= \frac{1}{4} \cos\left(\pi \frac{|v-u|}{2s}\right) (1 - \rho l 2^{-l+1}) - \left(\frac{1}{4} - \rho l 2^{-l}\right) \\ &\leq \frac{1}{4} \left(1 - \frac{1}{\pi} \left(\pi \frac{|v-u|}{2s}\right)^2\right) (1 - \rho l 2^{-l+1}) - \left(\frac{1}{4} - \rho l 2^{-l}\right) \\ &\leq 2^{-l} \left(-\frac{l}{2} \rho - \frac{1}{16s^2} 2^l + \frac{1}{8s^2} \rho l + \rho l\right), \end{aligned}$$

which, as it is easy to check, is smaller than or equal to zero for s large enough (recall that $\rho = s^{-3}$ and l is a constant depending only on k and α).

Now, we show that the principal's expected utility provided by the menu of deterministic contracts defined above is as desired. We do that by first proving that the contribution to the principal's expected utility due to each agent's type $\theta_v \in \Theta$ such that $v \in V^*$ is equal to $\frac{1}{2} \rho l 2^{-l}$. Since each agent reports their true type θ_v and plays action a_v , the principal's expected utility is

$$\begin{aligned} & F_{\theta_v, a_v, \omega_3} r_{\omega_3} - F_{\theta_v, a_v, \omega_1} p_{\omega_1}^{\theta_v} - F_{\theta_v, a_v, \omega_2} p_{\omega_2}^{\theta_v} \\ &= \frac{1}{4} - \frac{1}{4} \cos^2\left(\pi \frac{v}{2s}\right) (1 - \rho l 2^{-l+1}) - \frac{1}{4} \sin^2\left(\pi \frac{v}{2s}\right) (1 - \rho l 2^{-l+1}) \\ &= \frac{1}{4} - \frac{1}{4} (1 - \rho l 2^{-l+1}) = \frac{1}{4} \rho l 2^{-l+1} = \frac{1}{2} \rho l 2^{-l}. \end{aligned}$$

Moreover, we show that the principal's expected utility resulting from each agent's type $\theta_v \in \Theta$ such that $v \notin V^*$ is at least zero. For each such a type θ_v , let $u \in V^*$ be a vertex such that $p^{\theta_v} = p^{\theta_u}$. We need to consider four different cases. First, suppose that $b^{\theta_v} (p^{\theta_v}) = a_v$. Then, the principal's expected utility is equal to

$$\begin{aligned} & \frac{1}{4} - \frac{1}{4} \cos\left(\pi \frac{v}{2s}\right) \cos\left(\pi \frac{u}{2s}\right) (1 - \rho l 2^{-l+1}) \\ & \quad - \frac{1}{4} \sin\left(\pi \frac{v}{2s}\right) \sin\left(\pi \frac{u}{2s}\right) (1 - \rho l 2^{-l+1}) \geq \frac{1}{4} - \frac{1}{4} (1 - \rho l 2^{-l+1}) \geq 0, \end{aligned}$$

where in the inequality we use $\sin x \sin y + \cos x \cos y \leq 1$.

Second, assume that $b^{\theta_v} (p^{\theta_v}) = a_{ui}$ for some index $i \in [l-3]$. Then, the principal's utility is

$$\begin{aligned} & 2^{i-2} - \cos^2\left(\pi \frac{u}{2s}\right) 2^{-i-2} (1 - \rho l 2^{-l+1}) - \sin^2\left(\pi \frac{u}{2s}\right) 2^{-i-2} (1 - \rho l 2^{-l+1}) \\ & \geq 2^{i-2} - \cos^2\left(\pi \frac{u}{2s}\right) 2^{-i-2} - \sin^2\left(\pi \frac{u}{2s}\right) 2^{-i-2} = 0. \end{aligned}$$

Third, suppose that $b^{\theta_v} (p^{\theta_v}) = a_{u'i}$ for some vertex $u' \neq u \in V$ and index $i \in [l-3]$, then the principal's expected utility is

$$\begin{aligned} & 2^{i-2} - \cos\left(\pi \frac{u}{2s}\right) 2^{-i-2} \cos\left(\pi \frac{u'}{2s}\right) (1 - \rho l 2^{-l+1}) - \sin\left(\pi \frac{u}{2s}\right) 2^{-i-2} \sin\left(\pi \frac{u'}{2s}\right) (1 - \rho l 2^{-l+1}) \\ & \geq 2^{i-2} - \cos\left(\pi \frac{u}{2s}\right) 2^{-i-2} \cos\left(\pi \frac{u'}{2s}\right) - \sin\left(\pi \frac{u}{2s}\right) 2^{-i-2} \sin\left(\pi \frac{u'}{2s}\right) \\ & \geq 0. \end{aligned}$$

Finally, whenever $b^{\theta_v} (p^{\theta_v}) = \bar{a}$, expected rewards and payments are zero. Hence, we can conclude that the principal's expected utility when the agent's type is $\theta_v \in \Theta$ with $v \notin V^*$ is at least zero. Thus, the overall principal's expected utility is at least

$$\sum_{\theta_v \in \Theta: v \in V^*} \mu_{\theta_v} \left(\sum_{\omega \in \Omega} F_{\theta_v, b^{\theta_v}(p^{\theta_v}), \omega} r_{\omega} - \sum_{\omega \in \Omega} F_{\theta_v, b^{\theta_v}(p^{\theta_v}), \omega} p_{\omega}^{\theta_v} \right) \geq \frac{1}{2} \eta \rho l 2^{-l}.$$

Soundness. We prove that, if all the independent sets of G have size at most $\alpha \eta s$, then the principal's expected utility is at most $4\alpha \eta l \rho 2^{-l}$ in any DSIC menu of deterministic contracts $P = (p^\theta)_{\theta \in \Theta}$. First, we show that, when the agent plays an action a_{ui} , then the principal's expected utility is at most $3\rho 2^{-l}$. If an agent of type $\theta_v \in \Theta$ is incentivized to play an action $a_{ui} \in A_{\theta_v}$ with $i = l-3$ by contract p^{θ_v} , i.e., it holds $b^{\theta_v} (p^{\theta_v}) = a_{ui}$, then it is easy to see that the principal's expected utility

is at most $F_{\theta_v, a_{ui}, \omega_3} r_{\omega_3} - c_{\theta_v, a_{ui}} = 3\rho 2^{-l}$. Moreover, suppose that the agent plays an action a_{ui} with $i \in [l - 4]$. Since such an action is incentivized by contract p^{θ_v} over action a_{ui+1} , it is easy to check that by IC conditions it must be the case that

$$\begin{aligned} p_{\omega_1}^{\theta_v} \sin\left(\pi \frac{u}{2s}\right) 2^{-i-2} + p_{\omega_2}^{\theta_v} \cos\left(\pi \frac{u}{2s}\right) 2^{-i-2} - \left(2^{-i-2} - \rho(l-i)2^{-l}\right) \\ \geq p_{\omega_1}^{\theta_v} \sin\left(\pi \frac{u}{2s}\right) 2^{-i-3} + p_{\omega_2}^{\theta_v} \cos\left(\pi \frac{u}{2s}\right) 2^{-i-3} - \left(2^{-i-3} - \rho(l-i-1)2^{-l}\right), \end{aligned}$$

which implies that

$$\begin{aligned} \left(p_{\omega_1}^{\theta_v} \sin\left(\pi \frac{u}{2s}\right) + p_{\omega_2}^{\theta_v} \cos\left(\pi \frac{u}{2s}\right)\right) \left(2^{-i-2} - 2^{-i-3}\right) \\ \geq \left(2^{-i-2} - \rho(l-i)2^{-l}\right) - \left(2^{-i-3} - \rho(l-i-1)2^{-l}\right). \end{aligned}$$

Hence,

$$p_{\omega_1}^{\theta_v} \sin\left(\pi \frac{u}{2s}\right) + p_{\omega_2}^{\theta_v} \cos\left(\pi \frac{u}{2s}\right) \geq \left(2^{-i-3} - \rho 2^{-l}\right) 2^{i+3},$$

and the agent's expected payment is at least

$$\begin{aligned} p_{\omega_1}^{\theta_v} \sin\left(\pi \frac{u}{2s}\right) 2^{-i-2} + p_{\omega_2}^{\theta_v} \cos\left(\pi \frac{u}{2s}\right) 2^{-i-2} &\geq 2^{-i-2} \left(2^{-i-3} - \rho 2^{-l}\right) 2^{i+3} \\ &\geq 2^{-i-2} - \rho 2^{-l+1}. \end{aligned}$$

This implies that the principal's expected utility is at most

$$\begin{aligned} F_{\theta_v, a_{ui}, \omega_3} r_{\omega_3} - F_{\theta_v, a_{ui}, \omega_3} p_{\omega_1}^{\theta_v} - F_{\theta_v, a_{ui}, \omega_3} p_{\omega_2}^{\theta_v} &\leq 2^{-i-2} - \left(2^{-i-2} - \rho 2^{-l+1}\right) \\ &\leq \rho 2^{-l+1} \\ &\leq 3\rho 2^{-l}. \end{aligned}$$

Next, we prove that, given two agent's types $\theta_v \in \Theta$ and $\theta_u \in \Theta$ such that $(v, u) \in E$, then the principal's expected utility due to an agent of type θ_v or θ_u is at most $3\rho 2^{-l}$. As we have shown, when an agent of type θ_v plays an action different from a_v , the principal's expected utility is at most $3\rho 2^{-l}$. Hence, the principal's expected utility is strictly greater than $3\rho 2^{-l}$ for both the agent's types $\theta_v \in \Theta$ and $\theta_u \in \Theta$ only if they play a_v and a_u , respectively. However, in the following, we show that if $(v, u) \in E$ and the two agent types play the actions a_v and a_u , then the principal's expected utility is at most $3\rho 2^{-l}$ for at least one of the two types. Suppose that p^{θ_v} incentivizes action a_v and that p^{θ_u} incentivizes action a_u , namely $b^{\theta_v}(p^{\theta_v}) = a_v$ and $b^{\theta_u}(p^{\theta_u}) = a_u$. Then, since an agent of type θ_v prefers action a_v over action a_{ui} with $i = 1$, by IC we have

$$\begin{aligned} p_{\omega_1}^{\theta_v} \frac{1}{4} \sin\left(\pi \frac{v}{2s}\right) + p_{\omega_2}^{\theta_v} \frac{1}{4} \cos\left(\pi \frac{v}{2s}\right) - \left(\frac{1}{4} - \rho l 2^{-l}\right) \\ \geq p_{\omega_1}^{\theta_u} \frac{1}{8} \sin\left(\pi \frac{u}{2s}\right) + p_{\omega_2}^{\theta_u} \frac{1}{8} \cos\left(\pi \frac{u}{2s}\right) - \left(\frac{1}{8} - \rho(l-1)2^{-l}\right). \end{aligned}$$

Similarly, since an agent of type θ_u prefers action a_u over action a_{vi} with $i = 1$, we have

$$\begin{aligned} \frac{1}{4} p_{\omega_1}^{\theta_u} \sin\left(\pi \frac{u}{2s}\right) + p_{\omega_2}^{\theta_u} \frac{1}{4} \cos\left(\pi \frac{u}{2s}\right) - \left(\frac{1}{4} - \rho l 2^{-l}\right) \\ \geq p_{\omega_1}^{\theta_v} \frac{1}{8} \sin\left(\pi \frac{v}{2s}\right) + p_{\omega_2}^{\theta_v} \frac{1}{8} \cos\left(\pi \frac{v}{2s}\right) - \left(\frac{1}{8} - \rho(l-1)2^{-l}\right). \end{aligned}$$

Now, assume that

$$p_{\omega_1}^{\theta_v} \sin\left(\pi \frac{v}{2s}\right) + p_{\omega_2}^{\theta_v} \cos\left(\pi \frac{v}{2s}\right) \geq p_{\omega_1}^{\theta_u} \sin\left(\pi \frac{u}{2s}\right) + p_{\omega_2}^{\theta_u} \cos\left(\pi \frac{u}{2s}\right).$$

The other case is analogous. Then, by taking the difference between the two inequalities above, we get

$$\begin{aligned} p_{\omega_1}^{\theta_u} \frac{1}{8} \sin\left(\pi \frac{u}{2s}\right) + p_{\omega_2}^{\theta_u} \frac{1}{8} \cos\left(\pi \frac{u}{2s}\right) &\geq \left(\frac{1}{4} - \rho l 2^{-l}\right) - \left(\frac{1}{8} - \rho(l-1)2^{-l}\right) \\ &= \frac{1}{8} - \rho 2^{-l}. \end{aligned}$$

Then, the expected payment for an agent of type θ_u is

$$p_{\omega_1}^{\theta_u} \frac{1}{4} \sin\left(\pi \frac{u}{2s}\right) + p_{\omega_2}^{\theta_u} \frac{1}{4} \cos\left(\pi \frac{u}{2s}\right) \geq \frac{1}{4} - \rho 2^{-l+1},$$

and the principal's expected utility due to agent's type θ_u is at most $\rho 2^{-l+1}$.

As the last step, notice that the principal's expected utility is at most $\rho l 2^{-l}$ for every possible agent's type θ_v and menu of deterministic contracts $P = (p^\theta)_{\theta \in \Theta}$. In particular, we have shown that, by incentivizing an action different from a_v , the principal obtains at most $3\rho 2^{-l}$ expected utility. Moreover, by incentivizing action a_v , the principal gets at most $F_{\theta_v, a_v, \omega_3} r_{\omega_3} - c_{\theta_v, a_v} = \rho l 2^{-l}$.

To conclude the proof, for every possible menu of deterministic contracts $P = (p^\theta)_{\theta \in \Theta}$, the set of agent's types $\theta_v \in \Theta$ that result in a principal's expected utility strictly greater than $3\rho 2^{-l}$ defines an independent set $V^* \subseteq V$ of G . Moreover, for all the agent's types, the principal's expected utility is at most $\rho l 2^{-l}$. Hence, the overall principal's expected utility is:

$$\begin{aligned} \sum_{\theta_v \in \Theta} \mu_{\theta_v} \left(\sum_{\omega \in \Omega} F_{\theta, b^{\theta_v}(p^{\theta_v}), \omega} r_{\omega} - \sum_{\omega \in \Omega} F_{\theta, b^{\theta_v}(p^{\theta_v}), \omega} p_{\omega}^{\theta_v} \right) \\ \leq \frac{1}{s} \left(\sum_{\theta_v \in \Theta: v \in V^*} \rho l 2^{-l} + \sum_{\theta_v \in \Theta: v \notin V^*} 3\rho 2^{-l} \right) \\ \leq \alpha \eta \rho l 2^{-l} + 3\rho 2^{-l} \\ \leq 4\alpha \eta \rho l 2^{-l}, \end{aligned}$$

where the last inequality follows from $\alpha \eta l \geq \alpha \frac{1}{k} \frac{k}{\alpha} = 1$ (recall that $\eta \geq \frac{1}{k}$). \square

5. Menus of deterministic contracts: an additive PTAS with a constant number of outcomes

In the previous section, we show that the problem of computing a utility-maximizing DSIC menu of deterministic contracts does *not* admit any multiplicative approximation that can be computed in polynomial time even in Bayesian principal-agent instances with a constant number of outcomes (and actions). In this section, we study the domain of additive approximations, where we provide a PTAS for the problem that works in settings with a constant number of outcomes. This result is tight since by Theorem 3 the problem does *not* admit an additive FPTAS unless $P = NP$.

In order to define our PTAS, we first need to introduce some auxiliary results. In the following, we will rely on an abstract definition of *approximately-incentive-compatible* menu of deterministic contracts, which combines two levels of approximation, one related to the IC conditions on agent's actions, and the other to the DSIC constraints that the menu should satisfy. In particular, approximate menus do *not* only specify a contract for each agent's type, but also a tuple of action recommendations, one for each type, so that the action recommended to each agent's type is approximately IC under the contract corresponding to that type. Formally:

Definition 2 (ϵ -Approximate Menu of Deterministic Contracts). Given any Bayesian principal-agent instance (Θ, A, Ω) and $\epsilon > 0$, an ϵ -approximate menu of deterministic contracts is a tuple $(p^\theta, a^\theta)_{\theta \in \Theta}$ of contract-action pairs, with $p^\theta \in \mathbb{R}_+^m$ and $a^\theta \in A$ for all $\theta \in \Theta$, such that:

$$\sum_{\omega \in \Omega} F_{\theta, a^\theta, \omega} p_{\omega}^{\theta} - c_{\theta, a^\theta} \geq \sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p^{\hat{\theta}}), \omega} p_{\omega}^{\hat{\theta}} - c_{\theta, b^{\theta}(p^{\hat{\theta}})} - \epsilon \quad \forall \theta, \hat{\theta} \in \Theta. \tag{12}$$

Intuitively, an ϵ -approximate menu of deterministic contracts defined by $(p^\theta, a^\theta)_{\theta \in \Theta}$ satisfies the following two conditions: each action a^θ is ϵ -approximately IC for an agent of type $\theta \in \Theta$ under contract p^θ (see the case $\theta = \hat{\theta}$ in Equation (12)), and, additionally, the menu is ϵ -approximately DSIC assuming type θ plays a^θ when reporting their true type (see cases $\theta \neq \hat{\theta}$ in Equation (12)). In the following, when we refer to the principal's expected utility in an ϵ -approximate menu of deterministic contracts $(p^\theta, a^\theta)_{\theta \in \Theta}$, we mean the expected utility that the principal gets if each agent's type $\theta \in \Theta$ truthfully reports their type and plays a^θ under contract p^θ . Formally, the principal's expected utility in $(p^\theta, a^\theta)_{\theta \in \Theta}$ can be written as follows:

$$\sum_{\theta \in \Theta} \mu_{\theta} \left(\sum_{\omega \in \Omega} F_{\theta, a^\theta, \omega} p_{\omega}^{\theta} - c_{\theta, a^\theta} \right).$$

Now, we show that, when starting from an ϵ -approximate menu of deterministic contracts with the principal's expected utility equal to APX , we can recover in polynomial time a (non-approximate) DSIC menu of deterministic contracts providing the principal with an expected utility at least of $APX - 2\sqrt{\epsilon}$. This allows us to focus on the computation of ϵ -approximate menu of deterministic contracts. Formally, we prove the following lemma.

Lemma 5. Given a Bayesian principal-agent instance (Θ, A, Ω) and $\epsilon > 0$, let $(p^\theta, a^\theta)_{\theta \in \Theta}$ be an ϵ -approximate menu of deterministic contracts with principal's expected utility APX . Then, there exists a DSIC menu of deterministic contracts $\bar{P} = (\bar{p}^\theta)_{\theta \in \Theta}$ in which the principal's expected utility is at least $APX - 2\sqrt{\epsilon}$. Moreover, such a menu can be computed in polynomial time.

Next, we show that, in an optimal menu of deterministic contracts, large expected payments are only assigned to agent types occurring with small probability, as it will be useful in Lemma 7. Given a constant $L \geq 1$ and a menu of deterministic contracts $P = (p^\theta)_{\theta \in \Theta}$, we let $\Theta(L, P) \subseteq \Theta$ be the set of agent's types such that $\sum_{\omega \in \Omega} F_{\theta, b^\theta(p^\theta), \omega} p_\omega^\theta \geq L$. Formally, we prove the following:

Lemma 6. Given a Bayesian principal-agent instance (Θ, A, Ω) , let $P = (p^\theta)_{\theta \in \Theta}$ be an optimal menu of deterministic contracts. Then, for every constant $L > 1$, it holds $\sum_{\theta \in \Theta(L, P)} \mu_\theta \leq \frac{1}{L}$.

The next step is to show that there exists a menu of deterministic contracts that employs a small set of *different* contracts while providing almost optimal principal's expected utility. This allows us to represent a menu of deterministic contracts in a more succinct way. In particular, a menu that uses $k \in \mathbb{N}_+$ different contracts can be represented by a matrix $T \in \mathcal{T} := \mathbb{R}_+^{k \times m}$ and a function $f^T : \Theta \rightarrow \{1, \dots, k\}$. The matrix $T \in \mathcal{T}$ defines the k different contracts used by the menu on its rows, with $T_i \in \mathbb{R}_+^m$ (i.e., the i -th row of T) denoting the vector encoding the i -th contract and $T_{i, \omega}$ (i.e., the element indexed by i, ω) specifying the payment that the i -th contract associates to outcome $\omega \in \Omega$. Additionally, the function $f^T : \Theta \rightarrow \{1, \dots, k\}$ assigns contracts to agent's types, with $f^T(\theta) \in \{1, \dots, k\}$ defining the index of the contract proposed by menu $T \in \mathcal{T}$ when the agent reports type $\theta \in \Theta$ to the principal.

As a preliminary step, we prove the following crucial lemma.

Lemma 7. Given a Bayesian principal-agent instance (Θ, A, Ω) and a constant $\delta > 0$, there exists a DSIC menu of deterministic contracts $P = (p^\theta)_{\theta \in \Theta}$ that employs at most $O\left(\left(\frac{m}{\delta^3} \log \frac{m}{\delta}\right)^m\right)$ different contracts and provides the principal with an expected utility at least of $OPT - \delta$, where OPT is the principal's expected utility in an optimal DSIC menu of deterministic contracts.

The following theorem provides our PTAS. By Lemma 7, in order to find an optimal DSIC menu of deterministic contracts, we can optimize over the set of menus of contracts represented by the matrices in $\mathcal{T} = \mathbb{R}_+^{k \times m}$, for values of k that only depends on the desired constant approximation $\delta > 0$ and m . As a first step, we show that \mathcal{T} can be partitioned into polyhedra such that the principal's expected utility is linear in each of them, and, thus, an optimal solution is on a vertex of one of these polyhedra. Moreover, each polyhedron is defined by a subset of a common set of linear inequalities. Since the dimension of the space \mathcal{T} is km and, thus, it is constant, each vertex is defined by the intersection of km linearly independent hyperplanes. Finally, all the possible subsets of km inequalities can be enumerated in polynomial time since k and m are constants.

Theorem 4. There exists an additive PTAS for the problem of computing an optimal DSIC menu of deterministic contracts in Bayesian principal-agent instances with a constant number of outcomes.

Notice that our result cannot be extended to settings with an arbitrary number of outcomes. Indeed, with Bayesian principal-agent instances with an arbitrary number of outcomes, Guruganesh et al. [5] show that there exists an $\epsilon > 0$ such that it is NP-hard to approximate the principal's expected utility in an optimal menu of deterministic contracts up to within an ϵ additive factor unless $P = NP$.

6. Menus of deterministic contracts: two easy cases

In this section, we analyze two cases in which an optimal menu of deterministic contracts can be computed in polynomial time. In particular, Section 6.1 studies instances with two outcomes, complementing the hardness result for the setting with four outcomes, while Section 6.2 addresses the case with a constant number of types.

6.1. Case with two outcomes

Previously, we provide a PTAS for the problem of designing an optimal DSIC menu of deterministic contracts in Bayesian principal-agent instances with a constant number of outcomes. With four or more outcomes, this result is tight by Theorem 3. In this section, we show that the problem can be solved in polynomial time with only two outcomes. We leave as an open problem the analysis of the case with three outcomes. Our result generalizes that of Guruganesh et al. [5]. In particular, they prove a result that is analogous to ours but for the simpler case in which the action costs are the same for all the agent's types.

Our main result shows that using menus of deterministic contracts does *not* increase the principal's expected utility with respect to using a single contract.

Lemma 8. *In Bayesian principal-agent instances (Θ, A, Ω) with $|\Omega| = 2$, there exists a contract having the same principal's expected utility as an optimal DSIC menu of deterministic contracts.*

Let us remark that, in Bayesian principal-agent instances with two outcomes, Assumption 1 is *not* without loss of generality. Indeed, Alon et al. [33] prove that when Assumption 1 is not satisfied, then every optimal menu includes an arbitrarily large number of contracts.

By the previous lemma, in order to compute an optimal DSIC menu of deterministic contracts, it is sufficient to compute an optimal single contract. As shown by Castiglioni et al. [6] and Guruganesh et al. [5], an optimal single contract can be computed in polynomial time when the number of outcomes is constant. Hence, we obtain the following:

Theorem 5. *In Bayesian principal-agent instances (Θ, A, Ω) with $|\Omega| = 2$, there exists a polynomial-time algorithm that computes an optimal DSIC menu of deterministic contracts.*

6.2. Case with a constant number of types

As shown in Theorem 3, when there is an arbitrary number of agent's types, the principal's problem cannot be approximated efficiently, even with a constant number of outcomes and actions. We complement this result by showing that, when the number of types is constant, the problem can be solved in polynomial time. Formally:

Theorem 6. *In Bayesian principal-agent instances with a constant number of agent's types, there exists a polynomial-time algorithm that computes an optimal DSIC menu of deterministic contracts.*

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Nicola Gatti reports financial support was provided by Polytechnic of Milan Department of Electronics Information and Bioengineering. Nicola Gatti reports a relationship with Polytechnic of Milan Department of Electronics Information and Bioengineering that includes: employment and funding grants.

Data availability

No data was used for the research described in the article.

Acknowledgements

This paper is supported by FAIR (Future Artificial Intelligence Research) project, funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, Investment 1.3, Line on Artificial Intelligence).

Appendix A. Proofs omitted from Section 3

Lemma 1. *Given any DSIC menu of randomized contract $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$, there always exists a DSIC menu of randomized contracts $\tilde{\Gamma} = \{\tilde{\gamma}^\theta\}_{\theta \in \Theta}$ that provides the principal with at least the same expected utility as $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$ and such that, for every $\theta \in \Theta$, it holds $|\text{supp}(\gamma^\theta) \cap \hat{\mathcal{P}}^{\theta,a}| \leq 1$ for all $a \in A$.*

Proof. Let $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$ be a DSIC menu of randomized contracts, i.e., a feasible solution to Problem (3). We show that it is always possible to build a new DSIC menu of randomized contracts $\tilde{\Gamma} = \{\tilde{\gamma}^\theta\}_{\theta \in \Theta}$ that is still feasible to Problem (3), provides at least the same expected utility, and satisfies $|\text{supp}(\tilde{\gamma}^\theta) \cap \hat{\mathcal{P}}^{\theta,a}| \leq 1$ for all $\theta \in \Theta$ and $a \in A$. In particular, we let each probability distribution $\tilde{\gamma}^\theta$ be such that the support $\text{supp}(\tilde{\gamma}^\theta)$ contains all and only the contracts defined by $p^{\theta,a} := \mathbb{E}_{p \sim \gamma^\theta} [p \mid p \in \hat{\mathcal{P}}^{\theta,a}]$ for all the actions $a \in A$ such that $\text{supp}(\gamma^\theta) \cap \hat{\mathcal{P}}^{\theta,a} \neq \emptyset$. Moreover, we define $\tilde{\gamma}_{p^{\theta,a}}^\theta := \Pr_{p \sim \gamma^\theta} \{p \in \hat{\mathcal{P}}^{\theta,a}\}$ for these contracts. Clearly, the menu of randomized contracts $\tilde{\Gamma} = \{\tilde{\gamma}^\theta\}_{\theta \in \Theta}$ satisfies the required conditions on the supports.

It remains to show that $\tilde{\Gamma} = \{\tilde{\gamma}^\theta\}_{\theta \in \Theta}$ defines a solution to Problem (3). First, we prove that the principal's expected utility (the objective of Problem (3)) does *not* decrease with respect to that of the menu $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$. Formally:

$$\begin{aligned} & \sum_{\theta \in \Theta} \mu_\theta \mathbb{E}_{p \sim \gamma^\theta} \left[\sum_{\omega \in \Omega} F_{\theta,b^\theta(p),\omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta,b^\theta(p),\omega} p_\omega \right] \\ &= \sum_{\theta \in \Theta} \mu_\theta \sum_{a \in A} \Pr_{p \sim \gamma^\theta} \{p \in \hat{\mathcal{P}}^{\theta,a}\} \mathbb{E}_{p \sim \gamma^\theta} \left[\sum_{\omega \in \Omega} F_{\theta,b^\theta(p),\omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta,b^\theta(p),\omega} p_\omega \mid p \in \hat{\mathcal{P}}^{\theta,a} \right] \end{aligned} \tag{A.1a}$$

$$= \sum_{\theta \in \Theta} \mu_{\theta} \sum_{a \in A} \Pr_{p \sim \gamma^{\theta}} \left\{ p \in \hat{\mathcal{P}}^{\theta, a} \right\} \left(\sum_{\omega \in \Omega} F_{\theta, a, \omega} r_{\omega} - \sum_{\omega \in \Omega} F_{\theta, a, \omega} \mathbb{E}_{p \sim \gamma^{\theta}} [p_{\omega} \mid p \in \hat{\mathcal{P}}^{\theta, a}] \right) \tag{A.1b}$$

$$= \sum_{\theta \in \Theta} \mu_{\theta} \sum_{p \in \text{supp}(\tilde{\gamma}^{\theta})} \tilde{\gamma}_p^{\theta} \left(\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} r_{\omega} - \sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} p_{\omega} \right) \tag{A.1c}$$

$$= \sum_{\theta \in \Theta} \mu_{\theta} \mathbb{E}_{p \sim \tilde{\gamma}^{\theta}} \left[\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} r_{\omega} - \sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} p_{\omega} \right], \tag{A.1d}$$

where Equation (A.1a) holds since the sets $\hat{\mathcal{P}}_{\theta, a}$ define a partition of \mathbb{R}_+^m , Equation (A.1b) holds since for contracts $p \in \hat{\mathcal{P}}^{\theta, a}$ it is the case that $b^{\theta}(p) = a$, while Equation (A.1c) holds by definition of $\tilde{\Gamma} = \{\tilde{\gamma}^{\theta}\}_{\theta \in \Theta}$.

Finally, we need to prove that $\tilde{\Gamma} = \{\tilde{\gamma}^{\theta}\}_{\theta \in \Theta}$ satisfies Constraints (3b) of Problem (3). Formally, we have that for every $\theta \neq \hat{\theta} \in \Theta$ it holds:

$$\mathbb{E}_{p \sim \tilde{\gamma}^{\theta}} \left[\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} p_{\omega} - c_{\theta, b^{\theta}(p)} \right] = \sum_{p \in \text{supp}(\tilde{\gamma}^{\theta})} \tilde{\gamma}_p^{\theta} \left(\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} p_{\omega} - c_{\theta, b^{\theta}(p)} \right) \tag{A.2a}$$

$$= \sum_{a \in A} \Pr_{p \sim \gamma^{\theta}} \left\{ p \in \hat{\mathcal{P}}^{\theta, a} \right\} \left(\sum_{\omega \in \Omega} F_{\theta, a, \omega} \mathbb{E}_{p \sim \gamma^{\theta}} [p_{\omega} \mid p \in \hat{\mathcal{P}}^{\theta, a}] - c_{\theta, a} \right) \tag{A.2b}$$

$$= \sum_{a \in A} \Pr_{p \sim \gamma^{\theta}} \left\{ p \in \hat{\mathcal{P}}^{\theta, a} \right\} \mathbb{E}_{p \sim \gamma^{\theta}} \left[\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} p_{\omega} - c_{\theta, b^{\theta}(p)} \mid p \in \hat{\mathcal{P}}^{\theta, a} \right] \tag{A.2c}$$

$$= \mathbb{E}_{p \sim \gamma^{\theta}} \left[\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} p_{\omega} - c_{\theta, b^{\theta}(p)} \right] \tag{A.2d}$$

$$\geq \mathbb{E}_{p \sim \gamma^{\hat{\theta}}} \left[\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} p_{\omega} - c_{\theta, b^{\theta}(p)} \right] \tag{A.2e}$$

$$= \mathbb{E}_{p \sim \gamma^{\hat{\theta}}} \left[\max_{a \in A} \left\{ \sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega} - c_{\theta, a} \right\} \right] \tag{A.2f}$$

$$= \sum_{a' \in A} \Pr_{p \sim \gamma^{\hat{\theta}}} \left\{ p \in \hat{\mathcal{P}}^{\hat{\theta}, a'} \right\} \mathbb{E}_{p \sim \gamma^{\hat{\theta}}} \left[\max_{a \in A} \left\{ \sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega} - c_{\theta, a} \right\} \mid p \in \hat{\mathcal{P}}^{\hat{\theta}, a'} \right] \tag{A.2g}$$

$$\geq \sum_{a' \in A} \Pr_{p \sim \gamma^{\hat{\theta}}} \left\{ p \in \hat{\mathcal{P}}^{\hat{\theta}, a'} \right\} \max_{a \in A} \mathbb{E}_{p \sim \gamma^{\hat{\theta}}} \left[\sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega} - c_{\theta, a} \mid p \in \hat{\mathcal{P}}^{\hat{\theta}, a'} \right] \tag{A.2h}$$

$$= \sum_{a' \in A} \Pr_{p \sim \gamma^{\hat{\theta}}} \left\{ p \in \hat{\mathcal{P}}^{\hat{\theta}, a'} \right\} \max_{a \in A} \left\{ \sum_{\omega \in \Omega} F_{\theta, a, \omega} \mathbb{E}_{p \sim \gamma^{\hat{\theta}}} [p_{\omega} \mid p \in \hat{\mathcal{P}}^{\hat{\theta}, a'}] - c_{\theta, a} \right\} \tag{A.2i}$$

$$= \sum_{p \in \text{supp}(\tilde{\gamma}^{\hat{\theta}})} \tilde{\gamma}_p^{\hat{\theta}} \max_{a \in A} \left\{ \sum_{\omega \in \Omega} F_{\theta, a, \omega} p_{\omega} - c_{\theta, a} \right\} \tag{A.2j}$$

$$= \sum_{p \in \text{supp}(\tilde{\gamma}^{\hat{\theta}})} \tilde{\gamma}_p^{\hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} p_{\omega} - c_{\theta, b^{\theta}(p)} \right) \tag{A.2k}$$

$$= \mathbb{E}_{p \sim \tilde{\gamma}^{\hat{\theta}}} \left[\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} p_{\omega} - c_{\theta, b^{\theta}(p)} \right], \tag{A.2l}$$

where Equation (A.2b) holds by definition of $\tilde{\Gamma} = \{\tilde{\gamma}^{\theta}\}_{\theta \in \Theta}$, Equation (A.2c) holds by the fact that $b^{\theta}(p) = a$ for contracts $p \in \hat{\mathcal{P}}^{\theta, a}$, Equation (A.2d) holds since the sets $\hat{\mathcal{P}}_{\theta, a}$ define a partition of \mathbb{R}_+^m , Equation (A.2e) holds by the fact that $\Gamma = \{\gamma^{\theta}\}_{\theta \in \Theta}$ satisfies Constraints (3b) in Problem (3), Equation (A.2f) holds since action $a^{\theta}(p)$ is IC given contract p , Equation (A.2g) holds since sets $\hat{\mathcal{P}}_{\theta, a}$ define a partition of \mathbb{R}_+^m , Equation (A.2h) holds by Jensen's inequality and the fact that max is convex, Equation (A.2i) holds since $F_{\theta, a, \omega}$ and $c_{\theta, a}$ do not depend on p , Equation (A.2j) holds by definition of $\tilde{\Gamma} = \{\tilde{\gamma}^{\theta}\}_{\theta \in \Theta}$, while Equation (A.2k) holds by definition of $b^{\theta}(p)$. This concludes the proof. \square

Lemma 2. For every $\epsilon > 0$, there always exists a DSIC menu of randomized contracts $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$ with principal's expected utility at least $\text{SUP} - \epsilon$ such that, for every $\theta \in \Theta$, it holds (i) $|\text{supp}(\gamma^\theta) \cap \hat{\mathcal{P}}^{\theta,a}| \leq 1$ for all $a \in A$, and (ii) $p_\omega \leq C(I, \epsilon)$ for all $p \in \text{supp}(\gamma^\theta)$ and $\omega \in \Omega$, where $C(I, \epsilon) \in O(\frac{1}{\epsilon} \cdot 2^{\text{poly}(I)})$ and I denotes the size of the problem instance.

Proof. We show that for each $\epsilon > 0$, there exists a menu $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$ such that the required conditions are satisfied. By the definition of supremum for each $\eta > 0$, there exists a menu $\tilde{\Gamma} = \{\tilde{\gamma}^\theta\}_{\theta \in \Theta}$ that provides an expected utility to the principal of at least $1 - \eta$. Then, exploiting Lemma 1, we can derive a menu $\tilde{\Gamma} = \{\tilde{\gamma}^\theta\}_{\theta \in \Theta}$ with at least the same utility and such that for every $\theta \in \Theta$, it holds $|\text{supp}(\tilde{\gamma}^\theta) \cap \hat{\mathcal{P}}^{\theta,a}| \leq 1$ for all $a \in A$. Let $\tilde{\mathcal{P}}^\theta := \text{supp}(\tilde{\gamma}^\theta)$ for each $\theta \in \Theta$. Moreover, for each θ and a such that $\text{supp}(\tilde{\gamma}^\theta) \cap \hat{\mathcal{P}}^{\theta,a} \neq \emptyset$, let $\tilde{p}^{\theta,a} = \text{supp}(\tilde{\gamma}^\theta) \cap \hat{\mathcal{P}}^{\theta,a} \neq \emptyset$. Notice that, for each θ , $\tilde{\gamma}^\theta$ is supported on $\{\tilde{p}^{\theta,a}\}_{a \in A: \text{supp}(\tilde{\gamma}^\theta) \cap \hat{\mathcal{P}}^{\theta,a} \neq \emptyset}$.

Let $F := \min_{\theta \in \Theta, a \in A, \omega \in \Omega: F_{\theta,a,\omega} > 0} F_{\theta,a,\omega}$ be the minimum value of the probabilities $F_{\theta,a,\omega}$.¹⁰ Moreover, let $Y = \min_{\theta \in \Theta} \mu_\theta$.¹¹ Notice that $F \geq 2^{-\text{poly}(I)}$ and $Y \geq 2^{-\text{poly}(I)}$, where I is the instance size. Then, we show that for each $\theta \in \Theta$, $p \in \tilde{\mathcal{P}}^\theta$, and $\omega \in \Omega$, it holds $\tilde{\gamma}_p^\theta p_\omega \leq 4/FY$. By contradiction, assume that there exists $\tilde{\theta} \in \Theta$, $\tilde{p} \in \tilde{\mathcal{P}}^{\tilde{\theta}}$, and $\tilde{\omega} \in \Omega$ such that $\tilde{\gamma}_{\tilde{p}}^{\tilde{\theta}} \tilde{p}_{\tilde{\omega}} > \frac{4}{FY}$. Then, by the definition of F , there exists an agent's type θ' and an action a' , such that $F_{\theta',a',\tilde{\omega}} \geq F$. The agent of type θ' reporting type $\tilde{\theta}$ to the principal would incur an expected payment

$$\begin{aligned} \sum_{p \in \tilde{\mathcal{P}}^{\tilde{\theta}}} \tilde{\gamma}_p^{\tilde{\theta}} \left[\sum_{\omega \in \Omega} F_{\theta',b^{\theta'}(p),\omega} p_\omega \right] &\geq \sum_{p \in \tilde{\mathcal{P}}^{\tilde{\theta}}} \tilde{\gamma}_p^{\tilde{\theta}} \left[\sum_{\omega \in \Omega} F_{\theta',a',\omega} p_\omega - c_{\theta',a'} \right] + \sum_{p \in \tilde{\mathcal{P}}^{\tilde{\theta}}} \tilde{\gamma}_p^{\tilde{\theta}} c_{\theta',b^{\theta'}(p)} \\ &\geq \sum_{p \in \tilde{\mathcal{P}}^{\tilde{\theta}}} \tilde{\gamma}_p^{\tilde{\theta}} \sum_{\omega \in \Omega} F_{\theta',a',\omega} p_\omega - 1 \\ &\geq \tilde{\gamma}_{\tilde{p}}^{\tilde{\theta}} F_{\theta',a',\tilde{\omega}} \tilde{p}_{\tilde{\omega}} - 1 \\ &\geq \tilde{\gamma}_{\tilde{p}}^{\tilde{\theta}} \tilde{p}_{\tilde{\omega}} F - 1 > \frac{4}{FY} \cdot F - 1 \geq 3/Y. \end{aligned}$$

Thus, by the DSIC Constraints of type θ' , it must be the case that the expected payment to agent's type θ' , i.e., $\sum_{p \in \tilde{\mathcal{P}}^{\theta'}} \tilde{\gamma}_p^{\theta'} \left[\sum_{\omega \in \Omega} F_{\theta',b^{\theta'}(p),\omega} p_\omega \right]$, under menu $\tilde{\Gamma} = \{\tilde{\gamma}^\theta\}_{\theta \in \Theta}$ is at least $3/Y - 1 \geq 2/Y$, the constraints would be violated otherwise (recall that costs are in the range $[0, 1]$). Hence, the expected payment of the principal (over the agent's types) is at least

$$\sum_{\theta \in \Theta} \mu_\theta \sum_{p \in \tilde{\mathcal{P}}^\theta} \tilde{\gamma}_p^\theta \sum_{\omega \in \Omega} F_{\theta,b^\theta(p),\omega} p_\omega \geq \mu_{\theta'} \sum_{p \in \tilde{\mathcal{P}}^{\theta'}} \tilde{\gamma}_p^{\theta'} \left[\sum_{\omega \in \Omega} F_{\theta',b^{\theta'}(p),\omega} p_\omega \right] \geq 2.$$

As a result, since the principal's expected reward is at most 1, the principal's expected utility must be less than -1 , contradicting that the expected utility of the principal with menu $\tilde{\Gamma} = \{\tilde{\gamma}^\theta\}_{\theta \in \Theta}$ is at least $\text{SUP} - \eta \geq -\eta > -1$ for $\eta < 1$.

Let \mathcal{Z} the set of couples $(\theta, a) \in \Theta \times A$ such that there exists a $p \in \hat{\mathcal{P}}^{\theta,a}$ with $\tilde{\gamma}_p^\theta > 0$. Notice that $|\mathcal{Z}| \leq n\ell$. Then, for each couple $(\theta, a) \in \mathcal{Z}$ there exists a payment $p^{\theta,a}$ in $\hat{\mathcal{P}}^{\theta,a}$ such that $p^{\theta,a} \leq D$, where $D = 2^{\text{poly}(I)}$. This holds since $\hat{\mathcal{P}}^{\theta,a}$ is non-empty by the definition of \mathcal{Z} , while $p \leq D$ since $\hat{\mathcal{P}}^{\theta,a}$ is defined by linear inequalities such that the number of variables, the number of constraints, and the size of the binary representation of the coefficients can be bounded by a polynomial function of the instance size I [34].

Now we show how to build a menu $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$ with utility at least $\text{SUP} - \epsilon$ such that conditions i) and ii) hold for each θ . Consider the DSIC menu of contracts $\hat{\Gamma} = \{\hat{\gamma}^\theta\}_{\theta \in \Theta}$ such that for each $\theta' \in \Theta$ it holds $\hat{\gamma}_{p^{\theta',a}}^{\theta'} = 1/|\mathcal{Z}|$ for each $(\theta, a) \in \mathcal{Z}$. Then, consider the menu $\hat{\Gamma} = \{\hat{\gamma}^\theta\}_{\theta \in \Theta}$ such that for each $p \in \mathbb{R}_+^m$ it holds $\hat{\gamma}_p = \tilde{\gamma}_p(1-q) + \hat{\gamma}_p q$, where $q > 0$ will be defined in the following. By the linearity of the DSIC Constraints $\hat{\Gamma}$ is a DSIC menu of randomized contracts with principal's utility at least $(1-q)(\text{SUP} - \eta) - q \sum_{\theta \in \Theta} \mu_\theta \mathbb{E}_{p \sim \hat{\gamma}^\theta} \left[\sum_{\omega \in \Omega} F_{\theta,b^\theta(p),\omega} p_\omega - \sum_{\omega \in \Omega} F_{\theta,b^\theta(p),\omega} p_\omega \right] \geq \text{SUP} - \eta - q - q\eta - qD$, where the inequalities follow from $\mathbb{E}_{p \sim \hat{\gamma}^\theta} \left[\sum_{\omega \in \Omega} F_{\theta,b^\theta(p),\omega} p_\omega - \sum_{\omega \in \Omega} F_{\theta,b^\theta(p),\omega} p_\omega \right] \geq -D$ since distributions $\hat{\gamma}^\theta$ are supported on contracts with payments at most D . Hence, setting $\eta = q = \frac{\epsilon}{4D}$, we have that $\hat{\Gamma} = \{\hat{\gamma}^\theta\}_{\theta \in \Theta}$ provides principal's utility at least $\text{SUP} - \epsilon$.

We conclude the proof exploiting an argument equivalent to the one in Lemma 1 to build a menu $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$ that satisfies the required conditions. As shown in the proof of Lemma 1, given the menu $\hat{\Gamma}$, it is possible to build a menu of

¹⁰ We assume that each outcome $\omega \in \Omega$ occurs with positive probability at least for a type $\theta \in \Theta$ and an action $a \in A$. If it is not the case, we can safely remove outcome ω since it never occurs.

¹¹ We assume that for each type $\theta \in \Theta$ it holds $\mu_\theta > 0$. Otherwise, we can safely remove the type θ since it never occurs.

randomized contracts Γ with at least the same principal's utility such that for every $\theta \in \Theta$, it holds $|\text{supp}(\gamma^\theta) \cap \hat{\mathcal{P}}^{\theta,a}| \leq 1$ for all $a \in A$. We need to prove that also (ii) holds. For each $\theta \in \Theta$ and $a \in A$ such that $\text{supp}(\gamma^\theta) \cap \hat{\mathcal{P}}^{\theta,a} \neq \emptyset$, let $\bar{p}^{\theta,a} := \mathbb{E}_{p \sim \gamma^\theta} [p \mid p \in \hat{\mathcal{P}}^{\theta,a}]$. By the definition of Γ in Lemma 1 for each $\theta \in \Theta$ it holds $\text{supp}(\gamma^\theta) = \{\bar{p}^{\theta,a}\}_{a \in A: \text{supp}(\gamma^\theta) \cap \hat{\mathcal{P}}^{\theta,a} \neq \emptyset}$. To conclude the proof, it is sufficient to see that for each $\theta \in \Theta$ and $a \in A$ such that $\text{supp}(\gamma^\theta) \cap \hat{\mathcal{P}}^{\theta,a} \neq \emptyset$, it holds $(\theta, a) \in \mathcal{Z}$. Hence $\gamma_{\bar{p}^{\theta,a}}^\theta = \frac{1}{|\mathcal{Z}|}$ and $\gamma_{\bar{p}^{\theta,a}}^\theta \geq \frac{q}{|\mathcal{Z}|}$. For each θ and a , let $\mathcal{Z}^{\theta,a}$ be the set of couples $(\theta', a') \in \mathcal{Z}$ such that $p^{\theta',a'} \in \mathcal{P}^{\theta,a}$. Notice that $(\theta, a) \in \mathcal{Z}^{\theta,a}$. Thus, for each $\theta \in \Theta$ and $a \in A$ such that $\text{supp}(\gamma^\theta) \cap \hat{\mathcal{P}}^{\theta,a} \neq \emptyset$ and $\omega \in \Omega$, it holds

$$\begin{aligned} \bar{p}_\omega^{\theta,a} &= \mathbb{E}_{p \sim \gamma^\theta} [p \mid p \in \hat{\mathcal{P}}^{\theta,a}] = \frac{(1-q)\gamma_{\bar{p}^{\theta,a}}^\theta \bar{p}^{\theta,a} + \sum_{(\theta',a') \in \mathcal{Z}^{\theta,a}} \frac{q}{|\mathcal{Z}|} p_\omega^{\theta',a'}}{(1-q)\gamma_{\bar{p}^{\theta,a}}^\theta + \sum_{(\theta',a') \in \mathcal{Z}^{\theta,a}} \frac{q}{|\mathcal{Z}|}} \\ &\leq \frac{\frac{4}{FY} + qD}{\frac{q}{|\mathcal{Z}|}} \\ &\leq \frac{4n\ell D}{\epsilon} \left(\frac{4}{FY} + D \right) \in O\left(\frac{D^2 n \ell}{FY \epsilon} \right), \end{aligned}$$

where we assume $\bar{p}^{\theta,a}$ that is an arbitrary contract in $\hat{\mathcal{P}}^{\theta,a}$ when $\text{supp}(\gamma^\theta) \cap \hat{\mathcal{P}}^{\theta,a} = \emptyset$ and $\bar{p}^{\theta,a}$ is not been defined previously. Hence, for each $\theta \in \Theta$, $p \in \text{supp}(\gamma^\theta)$ and $\omega \in \Omega$, it holds $p_\omega = O(\frac{1}{\epsilon} \cdot 2^{\text{poly}(\ell)})$. \square

Lemma 3. For every $\epsilon > 0$, there always exists a DSIC menu of randomized contracts supported on $\mathcal{P}^{*,\epsilon}$ with principal's expected utility at least $\text{SUP} - \epsilon$.

Proof. By Lemma 2, there exists a menu of randomized contracts $\Gamma = \{\gamma^\theta\}_{\theta \in \Theta}$ with principal's utility at least $\text{SUP} - \epsilon$ such that, for every $\theta \in \Theta$, it holds $|\text{supp}(\gamma^\theta) \cap \hat{\mathcal{P}}^{\theta,a}| \leq 1$ for all $a \in A$ and $\text{supp}(\gamma^\theta) \subseteq \mathcal{P}^\epsilon$. We show that it is possible to build another optimal menu of randomized contracts $\tilde{\Gamma} = \{\tilde{\gamma}^\theta\}_{\theta \in \Theta}$ such that $\text{supp}(\tilde{\gamma}^\theta) \subseteq \mathcal{P}^{*,\epsilon}$ for every $\theta \in \Theta$.

In order to define $\tilde{\Gamma} = \{\tilde{\gamma}^\theta\}_{\theta \in \Theta}$, for every $\theta \in \Theta$ and contract $p \in \text{supp}(\gamma^\theta)$, let $\mathbf{a}^p := (a_{\theta'}^p)_{\theta' \in \Theta}$, where $a_{\theta'}^p := b^{\theta'}(p)$. Then, given a $\theta \in \Theta$ and a $p \in \text{supp}(\gamma^\theta)$, since $p \in \hat{\mathcal{P}}^{\theta,a^p}$, $\hat{\mathcal{P}}^{\theta,a^p} \in \mathcal{P}^{\mathbf{a}^p,\epsilon}$, and $\mathcal{P}^{\mathbf{a}^p,\epsilon}$ is a closed polytope, we can apply Carathéodory's theorem to conclude that there exists a probability distribution $\gamma^{\theta,p} \in \Delta_{\mathcal{P}^{*,\epsilon}}$ that is supported on the vertices of the polytope $\mathcal{P}^{\mathbf{a}^p,\epsilon}$, i.e., it holds $\text{supp}(\gamma^{\theta,p}) \subseteq V(\mathcal{P}^{\mathbf{a}^p,\epsilon})$, and such that $\mathbb{E}_{p' \sim \gamma^{\theta,p}} [p'] = \sum_{p' \in \text{supp}(\gamma^{\theta,p})} \gamma_{p'}^{\theta,p} p' = p$. Then, for every $\theta \in \Theta$, we define the distribution $\tilde{\gamma}^\theta \in \Delta_{\mathcal{P}^{*,\epsilon}}$ so that, for every $p \in \mathcal{P}^{*,\epsilon}$, it holds:

$$\tilde{\gamma}_p^\theta := \sum_{p' \in \text{supp}(\gamma^\theta)} \gamma_{p'}^\theta \gamma_{p'}^{\theta,p'}.$$

It remains to show that $\tilde{\Gamma} = \{\tilde{\gamma}^\theta\}_{\theta \in \Theta}$ is indeed a menu of randomized contracts with principal's utility at least $\text{SUP} - \epsilon$. First, it is immediate to see that the $\tilde{\gamma}^\theta$ are valid probability distributions over $\mathcal{P}^{*,\epsilon}$. Formally, for each $\theta \in \Theta$:

$$\sum_{p \in \mathcal{P}^{*,\epsilon}} \tilde{\gamma}_p^\theta = \sum_{p \in \mathcal{P}^{*,\epsilon}} \sum_{p' \in \text{supp}(\gamma^\theta)} \gamma_{p'}^\theta \gamma_{p'}^{\theta,p'} = \sum_{p' \in \text{supp}(\gamma^\theta)} \gamma_{p'}^\theta \sum_{p \in \mathcal{P}^{*,\epsilon}} \gamma_{p'}^{\theta,p'} = \sum_{p' \in \text{supp}(\gamma^\theta)} \gamma_{p'}^\theta = 1.$$

Next, we prove that $\tilde{\Gamma} = \{\tilde{\gamma}^\theta\}_{\theta \in \Theta}$ satisfies Constraints (3b) in Problem (3). First, notice that, for every $\theta \in \Theta$, $\hat{\theta} \in \Theta$, and $p \in \text{supp}(\gamma^\theta)$, it holds:

$$\sum_{p' \in \text{supp}(\gamma^{\theta,p})} \gamma_{p'}^{\theta,p} \left(\sum_{\omega \in \Omega} F_{\hat{\theta}, b^{\hat{\theta}}(p'), \omega} p'_\omega - c_{\hat{\theta}, b^{\hat{\theta}}(p')} \right) \tag{A.3a}$$

$$= \sum_{p' \in \text{supp}(\gamma^{\theta,p})} \gamma_{p'}^{\theta,p} \left(\sum_{\omega \in \Omega} F_{\hat{\theta}, a_\theta^p, \omega} p'_\omega - c_{\hat{\theta}, a_\theta^p} \right) \tag{A.3b}$$

$$= \sum_{\omega \in \Omega} F_{\hat{\theta}, a_\theta^p, \omega} \sum_{p' \in \text{supp}(\gamma^{\theta,p})} \gamma_{p'}^{\theta,p} p'_\omega - c_{\hat{\theta}, a_\theta^p} \tag{A.3c}$$

$$= \sum_{\omega \in \Omega} F_{\hat{\theta}, a_\theta^p, \omega} p_\omega - c_{\hat{\theta}, a_\theta^p}, \tag{A.3d}$$

where Equation (A.3b) holds since $a_\theta^p \in \mathcal{B}_{p'}^{\hat{\theta}}$ for contracts $p' \in \text{supp}(\gamma^{\theta,p}) \subseteq V(\mathcal{P}^{\mathbf{a}^p})$ and Equation (A.3d) holds by definition of $\gamma^{\theta,p}$. Moreover, for every $\theta \in \Theta$ and $\hat{\theta} \in \Theta$, it is the case that:

$$\sum_{p \in \mathcal{P}^{*,\epsilon}} \tilde{\gamma}_p^{\hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right) \quad (\text{A.4a})$$

$$= \sum_{p \in \mathcal{P}^{*,\epsilon}} \sum_{p' \in \text{supp}(\gamma^{\hat{\theta}})} \gamma_{p'}^{\hat{\theta}} \gamma_p^{\hat{\theta}, p'} \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right) \quad (\text{A.4b})$$

$$= \sum_{p' \in \text{supp}(\gamma^{\hat{\theta}})} \gamma_{p'}^{\hat{\theta}} \sum_{p \in \mathcal{P}^{*,\epsilon}} \gamma_p^{\hat{\theta}, p'} \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right) \quad (\text{A.4c})$$

$$= \sum_{p' \in \text{supp}(\gamma^{\hat{\theta}})} \gamma_{p'}^{\hat{\theta}} \sum_{p \in \text{supp}(\gamma^{\hat{\theta}, p'})} \gamma_p^{\hat{\theta}, p'} \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right) \quad (\text{A.4d})$$

$$= \sum_{p' \in \text{supp}(\gamma^{\hat{\theta}})} \gamma_{p'}^{\hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, a_{\theta}^{p'}, \omega} p'_\omega - c_{\theta, a_{\theta}^{p'}} \right) \quad (\text{A.4e})$$

$$= \sum_{p \in \text{supp}(\gamma^{\hat{\theta}})} \gamma_p^{\hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right) \quad (\text{A.4f})$$

$$= \mathbb{E}_{p \sim \gamma^{\hat{\theta}}} \left[\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right], \quad (\text{A.4g})$$

where Equation (A.4b) holds by definition of $\tilde{\gamma}^{\hat{\theta}}$, Equation (A.4d) holds since $\text{supp}(\gamma^{\hat{\theta}, p'}) \subseteq V(\mathcal{P}^{a^{p'}}) \subseteq \mathcal{P}^{*,\epsilon}$, Equation (A.4e) holds by Equation (A.3), while Equation (A.4f) holds since $a_{\theta}^{p'} = b^\theta(p)$. Thus, we can conclude that for every $\theta \in \Theta$ and $\hat{\theta} \in \Theta : \hat{\theta} \neq \theta$, it holds:

$$\begin{aligned} \sum_{p \in \mathcal{P}^{*,\epsilon}} \tilde{\gamma}_p^{\hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right) &= \mathbb{E}_{p \sim \gamma^{\hat{\theta}}} \left[\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right] \\ &\geq \mathbb{E}_{p \sim \gamma^{\hat{\theta}}} \left[\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right] \\ &= \sum_{p \in \mathcal{P}^{*,\epsilon}} \tilde{\gamma}_p^{\hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega - c_{\theta, b^\theta(p)} \right), \end{aligned}$$

where the equalities come from Equation (A.4) and the inequality from DSIC of Γ . Finally, we need to show that the principal's expected utility does *not* decrease. Formally, we have:

$$\begin{aligned} \sum_{\theta \in \Theta} \mu_\theta \sum_{p \in \mathcal{P}^{*,\epsilon}} \tilde{\gamma}_p^{\hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega \right) \\ = \sum_{\theta \in \Theta} \mu_\theta \sum_{p \in \mathcal{P}^{*,\epsilon}} \sum_{p' \in \text{supp}(\gamma^{\hat{\theta}})} \gamma_{p'}^{\hat{\theta}} \gamma_p^{\hat{\theta}, p'} \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega \right) \end{aligned} \quad (\text{A.5a})$$

$$= \sum_{\theta \in \Theta} \mu_\theta \sum_{p' \in \text{supp}(\gamma^{\hat{\theta}})} \gamma_{p'}^{\hat{\theta}} \sum_{p \in \text{supp}(\gamma^{\hat{\theta}, p'})} \gamma_p^{\hat{\theta}, p'} \left(\sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta, b^\theta(p), \omega} p_\omega \right) \quad (\text{A.5b})$$

$$\geq \sum_{\theta \in \Theta} \mu_\theta \sum_{p' \in \text{supp}(\gamma^{\hat{\theta}})} \gamma_{p'}^{\hat{\theta}} \sum_{p \in \text{supp}(\gamma^{\hat{\theta}, p'})} \gamma_p^{\hat{\theta}, p'} \left(\sum_{\omega \in \Omega} F_{\theta, a_{\theta}^{p'}, \omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta, a_{\theta}^{p'}, \omega} p_\omega \right) \quad (\text{A.5c})$$

$$= \sum_{\theta \in \Theta} \mu_\theta \sum_{p' \in \text{supp}(\gamma^{\hat{\theta}})} \gamma_{p'}^{\hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, a_{\theta}^{p'}, \omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta, a_{\theta}^{p'}, \omega} \sum_{p \in \text{supp}(\gamma^{\hat{\theta}, p'})} \gamma_p^{\hat{\theta}, p'} p_\omega \right) \quad (\text{A.5d})$$

$$= \sum_{\theta \in \Theta} \mu_\theta \sum_{p' \in \text{supp}(\gamma^{\hat{\theta}})} \gamma_{p'}^{\hat{\theta}} \left(\sum_{\omega \in \Omega} F_{\theta, a_{\theta}^{p'}, \omega} r_\omega - \sum_{\omega \in \Omega} F_{\theta, a_{\theta}^{p'}, \omega} p'_\omega \right) \quad (\text{A.5e})$$

$$= \sum_{\theta \in \Theta} \mu_{\theta} \sum_{p \in \text{supp}(\gamma^{\theta})} \gamma_p^{\theta} \left(\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} r_{\omega} - \sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} p_{\omega} \right) \tag{A.5f}$$

$$= \sum_{\theta \in \Theta} \mu_{\theta} \mathbb{E}_{p \sim \gamma^{\theta}} \left[\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} r_{\omega} - \sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p), \omega} p_{\omega} \right], \tag{A.5g}$$

where Equation (A.5a) holds by definition of $\tilde{\gamma}^{\theta}$, Equation (A.5b) holds by re-arranging the summations and by $\text{supp}(\gamma^{\theta, p'}) \subseteq V(\mathcal{P}^{a^{p'}}) \subseteq \mathcal{P}^{*, \epsilon}$, Equation (A.5c) holds since $a_{\theta}^{p'} \in \mathcal{B}_p^{\theta}$ for each $p \in \text{supp}(\gamma^{\theta, p'})$ and $b^{\theta}(p)$ breaks ties in favor of the principal, Equation (A.5e) holds by definition of $\gamma^{\theta, p'}$, while Equation (A.5f) holds since $a_{\theta}^{p'} := b^{\theta}(p)$. This concludes the proof. \square

Appendix B. Proofs omitted from Section 5

Lemma 5. *Given a Bayesian principal-agent instance (Θ, A, Ω) and $\epsilon > 0$, let $(p^{\theta}, a^{\theta})_{\theta \in \Theta}$ be an ϵ -approximate menu of deterministic contracts with principal's expected utility APX . Then, there exists a DSIC menu of deterministic contracts $\bar{P} = (\bar{p}^{\theta})_{\theta \in \Theta}$ in which the principal's expected utility is at least $APX - 2\sqrt{\epsilon}$. Moreover, such a menu can be computed in polynomial time.*

Proof. Given $(p^{\theta}, a^{\theta})_{\theta \in \Theta}$ as in Definition 2, we let $\hat{P} = (\hat{p}^{\theta})_{\theta \in \Theta}$ be an auxiliary menu of deterministic contracts such that $\hat{p}_{\omega}^{\theta} = (1 - \sqrt{\epsilon}) p_{\omega}^{\theta} + \sqrt{\epsilon} r_{\omega}$ for every $\theta \in \Theta$ and $\omega \in \Omega$. Intuitively, this modification helps to align the agent's utility with the principal's reward. Then, we define the menu of contracts $\bar{P} = (\bar{p}^{\theta})_{\theta \in \Theta}$ as follows. For every agent's type $\theta \in \Theta$, we let $\bar{p}^{\theta} = \hat{p}^{\theta'}$, where θ' is an optimal report of θ to \hat{P} . Formally, $\theta' \in \arg \max_{\theta' \in \Theta} \max_{a \in A} \left\{ \sum_{\omega \in \Omega} F_{\theta, a, \omega} \hat{p}_{\omega}^{\theta'} - c_{\theta, a} \right\}$. \hat{P} is essentially the DSIC version of \bar{P} obtained through the revelation principle.

Notice that $\bar{P} = (\bar{p}^{\theta})_{\theta \in \Theta}$ is DSIC by definition. Next, we show that, for each agent's type $\theta \in \Theta$, the principal's expected utility contribution due to that type under contract \bar{p}^{θ} decreases by at most $2\sqrt{\epsilon}$ compared with that obtained when the agent plays a^{θ} under p^{θ} , proving the lemma.

For every agent's type $\theta \in \Theta$, let $\theta' \in \Theta$ be the agent's type such that $\bar{p}^{\theta} = \hat{p}^{\theta'}$. Notice that it may be the case that $\theta = \theta'$ and/or $b^{\theta}(\bar{p}^{\theta}) = a^{\theta}$. First, we prove the following:

$$\begin{aligned} \sum_{\omega \in \Omega} F_{\theta, b^{\theta}(\bar{p}^{\theta}), \omega} \left((1 - \sqrt{\epsilon}) p_{\omega}^{\theta'} + \sqrt{\epsilon} r_{\omega} \right) - c_{\theta, b^{\theta}(\bar{p}^{\theta})} &= \sum_{\omega \in \Omega} F_{\theta, b^{\theta}(\bar{p}^{\theta}), \omega} \bar{p}_{\omega}^{\theta} - c_{\theta, b^{\theta}(\bar{p}^{\theta})} \\ &\geq \sum_{\omega \in \Omega} F_{\theta, a^{\theta}, \omega} \hat{p}_{\omega}^{\theta} - c_{\theta, a^{\theta}} \\ &= \sum_{\omega \in \Omega} F_{\theta, a^{\theta}, \omega} \left((1 - \sqrt{\epsilon}) p_{\omega}^{\theta} + \sqrt{\epsilon} r_{\omega} \right) - c_{\theta, a^{\theta}}, \end{aligned}$$

where the inequality follows from the fact that playing action $b^{\theta}(\bar{p}^{\theta})$ under contract \bar{p}^{θ} is preferred by the agent over playing action a^{θ} under contract \hat{p}^{θ} , by definition of contract \bar{p}^{θ} . Moreover, by definition of the ϵ -approximate menu of deterministic contracts $(p^{\theta}, a^{\theta})_{\theta \in \Theta}$, it holds

$$\sum_{\omega \in \Omega} F_{\theta, a^{\theta}, \omega} p_{\omega}^{\theta} - c_{\theta, a^{\theta}} \geq \sum_{\omega \in \Omega} F_{\theta, b^{\theta}(p^{\theta}), \omega} p_{\omega}^{\theta} - c_{\theta, b^{\theta}(p^{\theta})} - \epsilon \geq \sum_{\omega \in \Omega} F_{\theta, b^{\theta}(\bar{p}^{\theta}), \omega} p_{\omega}^{\theta} - c_{\theta, b^{\theta}(\bar{p}^{\theta})} - \epsilon.$$

By summing the two relations obtained above, we get

$$\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(\bar{p}^{\theta}), \omega} \left(-\sqrt{\epsilon} p_{\omega}^{\theta'} + \sqrt{\epsilon} r_{\omega} \right) \geq \sum_{\omega \in \Omega} F_{\theta, a^{\theta}, \omega} \left(-\sqrt{\epsilon} p_{\omega}^{\theta} + \sqrt{\epsilon} r_{\omega} \right) - \epsilon,$$

which implies that

$$\sum_{\omega \in \Omega} F_{\theta, b^{\theta}(\bar{p}^{\theta}), \omega} \left(r_{\omega} - p_{\omega}^{\theta'} \right) \geq \sum_{\omega \in \Omega} F_{\theta, a^{\theta}, \omega} \left(r_{\omega} - p_{\omega}^{\theta} \right) - \sqrt{\epsilon}.$$

Then, the principal's expected utility when the agent is of type is $\theta \in \Theta$ is equal to

$$\begin{aligned} \sum_{\omega \in \Omega} F_{\theta, b^{\theta}(\bar{p}^{\theta}), \omega} \left(r_{\omega} - \bar{p}_{\omega}^{\theta} \right) &= \sum_{\omega \in \Omega} F_{\theta, b^{\theta}(\bar{p}^{\theta}), \omega} \left(r_{\omega} - (1 - \sqrt{\epsilon}) p_{\omega}^{\theta'} - \sqrt{\epsilon} r_{\omega} \right) \\ &= (1 - \sqrt{\epsilon}) \sum_{\omega \in \Omega} F_{\theta, b^{\theta}(\bar{p}^{\theta}), \omega} \left(r_{\omega} - p_{\omega}^{\theta'} \right) \end{aligned}$$

$$\begin{aligned} &\geq (1 - \sqrt{\epsilon}) \left(\sum_{\omega \in \Omega} F_{\theta, a^\theta, \omega} (r_\omega - p_\omega^\theta) - \sqrt{\epsilon} \right) \\ &\geq \sum_{\omega \in \Omega} F_{\theta, a^\theta, \omega} (r_\omega - p_\omega^\theta) - 2\sqrt{\epsilon}. \end{aligned}$$

This concludes the proof. \square

Lemma 6. Given a Bayesian principal-agent instance (Θ, A, Ω) , let $P = (p^\theta)_{\theta \in \Theta}$ be an optimal menu of deterministic contracts. Then, for every constant $L > 1$, it holds $\sum_{\theta \in \Theta(L, P)} \mu_\theta \leq \frac{1}{L}$.

Proof. By contradiction, we show that, if $\sum_{\theta \in \Theta(L, P)} \mu_\theta > \frac{1}{L}$, then the menu $P = (p^\theta)_{\theta \in \Theta}$ provides a negative expected utility to the principal. In particular, if $\sum_{\theta \in \Theta(L, P)} \mu_\theta > \frac{1}{L}$, then the expected payment due to agent's types in $\Theta(L, P)$ is greater than 1, while the expected reward to the principal is at most 1. Hence, since a menu of all-zero contract vectors provides the principal with an expected utility at least of zero, the menu P cannot be optimal, reaching a contradiction. \square

Lemma 7. Given a Bayesian principal-agent instance (Θ, A, Ω) and a constant $\delta > 0$, there exists a DSIC menu of deterministic contracts $P = (p^\theta)_{\theta \in \Theta}$ that employs at most $O\left(\left(\frac{m}{\delta^3} \log \frac{m}{\delta}\right)^m\right)$ different contracts and provides the principal with an expected utility at least of $OPT - \delta$, where OPT is the principal's expected utility in an optimal DSIC menu of deterministic contracts.

Proof. The proof is organized in two steps. As a first step, we show that, once given an optimal menu of deterministic contracts $P = (p^\theta)_{\theta \in \Theta}$, we can build a $\frac{\delta^2}{16}$ -approximate menu of deterministic contracts $(\bar{p}^\theta, \bar{a}^\theta)_{\theta \in \Theta}$ that provides the principal with an expected utility of $OPT - \frac{\delta}{2}$ and such that it adopts a limited number of different payment values. In particular, all payment values \bar{p}_ω^θ , for every agent's type $\theta \in \Theta$ and outcome $\omega \in \Omega$, belong to a small, finite set. This implies that also the set of possible different contracts is small. Then, as a second step, we employ Lemma 5 to build a DSIC menu of deterministic contracts as in the statement of the lemma, starting from $(\bar{p}^\theta, \bar{a}^\theta)_{\theta \in \Theta}$, and only incurring in a small loss in the principal's expected utility.

Let $M(\omega) := \max_{\theta \in \Theta} p_\omega^\theta$ for every $\omega \in \Omega$, and let $I := \left\{0, 1, \dots, \left\lceil \frac{\log \eta}{\log(1-\eta)} \right\rceil\right\}$ with $\eta := \frac{\delta^3}{64m}$. We split Θ into two sets. Letting $L := \frac{4}{\delta}$, the first set is $\Theta_1 := \Theta(L, P)$, while the second one is $\Theta_2 := \Theta \setminus \Theta_1$, which includes all the agent's types that are not in Θ_1 . Moreover, we let $\hat{\Theta} \subseteq \Theta$ be the subset of agent's types $\theta \in \Theta$ such that there exists an action $a^\theta \in A$ and an outcome $\omega^\theta \in \Omega$ that simultaneously satisfy $p_{\omega^\theta}^\theta < \eta M(\omega^\theta)$ and $F_{\theta, a^\theta, \omega^\theta} p_{\omega^\theta}^\theta \geq \frac{\delta^2}{8m}$.

First, we prove that $\hat{\Theta} \subseteq \Theta_1$. For every $\theta \in \hat{\Theta}$, we have $F_{\theta, a^\theta, \omega^\theta} \geq \frac{1}{p_{\omega^\theta}^\theta} \frac{\delta^2}{8m} > \frac{\delta^2}{8m} \frac{1}{\eta M(\omega^\theta)} = \frac{8}{\delta M(\omega^\theta)}$. Then, by IC conditions and the definition of $M(\omega^\theta)$, it holds

$$\sum_{\omega \in \Omega} F_{\theta, b^\theta(p^\theta), \omega} p_\omega^\theta - c_{\theta, b^\theta(p^\theta)} \geq F_{\theta, a^\theta, \omega^\theta} M(\omega^\theta) - c_{\theta, a^\theta},$$

which implies that

$$\sum_{\omega \in \Omega} F_{\theta, b^\theta(p^\theta), \omega} p_\omega^\theta \geq F_{\theta, a^\theta, \omega^\theta} M(\omega^\theta) - 1 \geq \frac{8}{\delta} - 1 \geq \frac{4}{\delta}.$$

Next, we build a $\frac{\delta^2}{16}$ -approximate menu of deterministic contracts $(\bar{p}^\theta, \bar{a}^\theta)_{\theta \in \Theta}$, as follows. For each agent's type $\theta \in \Theta_2$ and outcome $\omega \in \Omega$, we set $\bar{p}_\omega^\theta = 0$ if $p_\omega^\theta < (1 - \eta)^{\left\lceil \frac{\log \eta}{\log(1-\eta)} \right\rceil} M(\omega) \leq \eta M(\omega)$, while we set $\bar{p}_\omega^\theta = (1 - \eta)^{i^*} M(\omega)$ otherwise, where i^* is the smallest integer such that $(1 - \eta)^{i^*} M(\omega) \leq p_\omega^\theta$. Notice that i^* is at most $\left\lceil \frac{\log \eta}{\log(1-\eta)} \right\rceil$, and, thus, $i^* \in I$. Intuitively, contracts \bar{p}^θ are defined by lowering the payments of the corresponding contracts p^θ to the nearest value $(1 - \eta)^i$ with $i \in I$, so that, in the second case above, it holds that $\bar{p}_\omega^\theta \geq (1 - \eta)p_\omega^\theta$. Moreover, we set $\bar{a}^\theta = b^\theta(p^\theta)$ for all $\theta \in \Theta_2$. Finally, for each agent's type $\theta \in \Theta_1$, we set $\bar{p}^\theta = p^{\hat{\theta}}$ and $\bar{a}^\theta = b^\theta(\bar{p}^\theta)$, where, for the ease of notation, we let $g(\theta) \in \arg \max_{\hat{\theta} \in \Theta_2} \left\{ \sum_{\omega \in \Omega} F_{\theta, b^\theta(\bar{p}^\theta), \omega} \bar{p}_\omega^\theta - c_{\theta, b^\theta(\bar{p}^\theta)} \right\}$.

First, we show that the menu $(\bar{p}^\theta, \bar{a}^\theta)_{\theta \in \Theta}$ is $\frac{\delta^2}{16}$ -approximate. This holds by construction for all the agent's types in Θ_1 . For any type $\theta \in \Theta_2$, as a first step, we prove that

$$F_{\theta, \bar{a}^\theta, \omega} \bar{p}_\omega^\theta \geq F_{\theta, \bar{a}^\theta, \omega} p_\omega^\theta - \frac{\delta^2}{16m} \quad \forall \omega \in \Omega. \tag{B.1}$$

We consider two cases. If $F_{\theta, \bar{a}^\theta, \omega} P_\omega^\theta < \frac{\delta^2}{16m}$, then $F_{\theta, \bar{a}^\theta, \omega} P_\omega^\theta - \frac{\delta^2}{16m} < 0$, and, thus, the inequality above holds since $F_{\theta, \bar{a}^\theta, \omega} \bar{p}_\omega^\theta \geq 0$. Otherwise, since $\theta \notin \hat{\Theta} \subseteq \Theta_1$, it holds $p_\omega^\theta \geq \eta M(\omega)$ and

$$F_{\theta, \bar{a}^\theta, \omega} \bar{p}_\omega^\theta \geq F_{\theta, \bar{a}^\theta, \omega} (1 - \eta) P_\omega^\theta \geq F_{\theta, \bar{a}^\theta, \omega} p_\omega^\theta - \eta F_{\theta, \bar{a}^\theta, \omega} P_\omega^\theta \geq F_{\theta, \bar{a}^\theta, \omega} p_\omega^\theta - \eta \frac{4}{\delta} = F_{\theta, \bar{a}^\theta, \omega} p_\omega^\theta - \frac{\delta^2}{16m},$$

where the last inequality follows from $F_{\theta, \bar{a}^\theta, \omega} p_\omega^\theta \leq \sum_\omega F_{\theta, \bar{a}^\theta, \omega} p_\omega^\theta \leq \frac{4}{\delta}$ since $\theta \notin \Theta_1$. Then, we prove that $(\bar{p}^\theta, \bar{a}^\theta)_{\theta \in \Theta}$ satisfies Equation (12) for every agent's type $\theta \in \Theta_2$. Notice that, since for each contract \bar{p}^θ with $\theta \in \Theta_1$ there exists a contract $\bar{p}^{g(\theta)}$ with $g(\theta) \in \Theta_2$ such that $\bar{p}^\theta = \bar{p}^{g(\theta)}$, the condition in Equation (12) must be checked only for types $\hat{\theta} \in \Theta$. In particular, for every pair of agent's types $\theta, \hat{\theta} \in \Theta_2$, the following holds:

$$\begin{aligned} \sum_{\omega \in \Omega} F_{\theta, \bar{a}^\theta, \omega} \bar{p}_\omega^\theta - c_{\theta, \bar{a}^\theta} &\geq \sum_{\omega \in \Omega} \left(F_{\theta, \bar{a}^\theta, \omega} p_\omega^\theta - \frac{\delta^2}{16m} \right) - c_{\theta, \bar{a}^\theta} \\ &= \sum_{\omega \in \Omega} F_{\theta, b^\theta(p^\theta), \omega} p_\omega^\theta - c_{\theta, b^\theta(p^\theta)} - \frac{\delta^2}{16} \\ &\geq \sum_{\omega \in \Omega} F_{\theta, b^\theta(p^{\hat{\theta}}), \omega} p_\omega^{\hat{\theta}} - c_{\theta, b^\theta(p^{\hat{\theta}})} - \frac{\delta^2}{16} \\ &\geq \sum_{\omega \in \Omega} F_{\theta, b^\theta(\bar{p}^{\hat{\theta}}), \omega} p_\omega^{\hat{\theta}} - c_{\theta, b^\theta(\bar{p}^{\hat{\theta}})} - \frac{\delta^2}{16} \\ &\geq \sum_{\omega \in \Omega} F_{\theta, b^\theta(\bar{p}^{\hat{\theta}}), \omega} \bar{p}_\omega^{\hat{\theta}} - c_{\theta, b^\theta(\bar{p}^{\hat{\theta}})} - \frac{\delta^2}{16}, \end{aligned}$$

where the first inequality follows from Equation (B.1), the second one holds by the DSIC property of $P = (p^\theta)_{\theta \in \Theta}$, the third one from the definition of best response, while the last inequality follows from the fact that $\bar{p}_\omega^{\theta'} \leq p_\omega^{\theta'}$ for every $\theta \in \Theta_2$ and $\omega \in \Omega$.

Now, we prove that the principal's expected utility in the $\frac{\delta^2}{16}$ -approximate menu of deterministic contracts $(\bar{p}^\theta, \bar{a}^\theta)_{\theta \in \Theta}$ is at least $OPT - \frac{\delta}{2}$. For every agent's type $\theta \in \Theta_1$, we have that

$$\begin{aligned} \sum_{\omega \in \Omega} F_{\theta, \bar{a}^\theta, \omega} \bar{p}_\omega^\theta - c_{\theta, \bar{a}^\theta} &= \sum_{\omega \in \Omega} F_{\theta, b^\theta(\bar{p}^{g(\theta)}), \omega} \bar{p}_\omega^{g(\theta)} - c_{\theta, b^\theta(\bar{p}^{g(\theta)})} \\ &\leq \sum_{\omega \in \Omega} F_{\theta, b^\theta(\bar{p}^{g(\theta)}), \omega} p_\omega^{g(\theta)} - c_{\theta, b^\theta(\bar{p}^{g(\theta)})} \\ &\leq \sum_{\omega \in \Omega} F_{\theta, b^\theta(p^\theta), \omega} p_\omega^\theta - c_{\theta, b^\theta(p^\theta)}, \end{aligned}$$

where the first inequality follows from the fact that $\bar{p}_\omega^\theta \leq p_\omega^\theta$ for every $\theta \in \Theta_2$ and $\omega \in \Omega$, while the second one from the DSIC properties of $P = (p^\theta)_{\theta \in \Theta}$. Thus, we can conclude that $\sum_{\theta \in \Theta} F_{\theta, \bar{a}^\theta, \omega} \bar{p}_\omega^\theta \leq \sum_{\theta \in \Theta} F_{\theta, b^\theta(p^\theta), \omega} p_\omega^\theta + 1$. Moreover, the principal's expected reward of contract p^θ is $\sum_{\theta \in \Theta} F_{\theta, b^\theta(p^\theta), \omega} r_\omega \leq 1$. Thus, the revenue that the principal extracts from an agent of type $\theta \in \Theta_1$ decreases by at most 2, since, by changing the contract from p^θ to \bar{p}^θ , the reward decreases by at most 1 and the payment increases by at most 1. On the other hand, for every agent's type $\theta \in \Theta_2$, the principal's expected utility does not decrease, since the agent plays the same action and the payments decrease. Hence, the principal's expected utility decreases by at most $2 \sum_{\theta \in \Theta_1} \mu_\theta \leq 2 \frac{\delta}{4} = \frac{\delta}{2}$, where the inequality follows from Lemma 6.

Finally, by Lemma 5, starting from the $\frac{\delta^2}{16}$ -approximate menu of deterministic contracts $(\bar{p}^\theta, \bar{a}^\theta)_{\theta \in \Theta}$, we can build a DSIC menu of deterministic contracts $\hat{P} = (\hat{p}^\theta)_{\theta \in \Theta}$ providing the principal with a utility at least of $APX - 2\sqrt{\frac{\delta^2}{16}} = APX - \frac{\delta}{2}$, where APX is the expected utility of $(\bar{p}^\theta, \bar{a}^\theta)_{\theta \in \Theta}$. Hence, the principal's expected utility is at least $APX - \frac{\delta}{2} = OPT - \delta$. Moreover, we also have that $\hat{p}_\omega^\theta \in \{(1 - \eta)^i M(\omega) (1 - \frac{\delta}{4}) + \frac{\delta}{4} r_\omega\}_{i \in I} \cup \{\frac{\delta}{4} r_\omega\}$ for every $\theta \in \Theta$ and $\omega \in \Omega$. This holds since the menu of contracts built in Lemma 5 uses contracts defined as $(1 - \sqrt{\epsilon})\bar{p}^\theta + \sqrt{\epsilon}r$ for $\theta \in \Theta$. Thus, there are at most

$$|I + 1|^m = \left(\frac{\log \frac{\delta^3}{64m}}{\log \left(1 - \frac{\delta^3}{64m} \right)} + 3 \right)^m \leq \left(\frac{64m}{\delta^3} \log \frac{64m}{\delta^3} \right)^m = O \left(\left(\frac{m}{\delta^3} \log \frac{m}{\delta} \right)^m \right)$$

different contracts, where in the first inequality we use $-\log(1 - x) \geq x$ for $x \leq 1$. \square

Theorem 4. *There exists an additive PTAS for the problem of computing an optimal DSIC menu of deterministic contracts in Bayesian principal-agent instances with a constant number of outcomes.*

Proof. Let $\delta > 0$ be a desired additive approximation factor for the PTAS. By Lemma 7, in order to find an approximate DSIC menu of deterministic contracts, it is sufficient to optimize over all the menus using at most $k = k(\delta, m)$ different contracts. Recall that a menu that uses at most k different contracts can be represented by a matrix $T \in \mathcal{T} = \mathbb{R}_+^{k \times m}$ and a function $f : \Theta \rightarrow \{1, \dots, k\}$.

The PTAS works by organizing the space of possible matrices representing menus of deterministic contracts, namely $\mathbb{R}_+^{k \times m}$, into polyhedra, depending on which function f can be paired with the matrix to define a DSIC menu and which best responses are induced by the contract. In particular, given a function $f : \Theta \rightarrow \{1, \dots, k\}$ and a set of best responses $(a^\theta)_{\theta \in \Theta}$, we define $\mathcal{Q}^{f, (a^\theta)_{\theta \in \Theta}} \subseteq \mathbb{R}_+^{k \times m}$ as the polyhedra defined by the following set of linear inequalities enforcing DSIC:

$$\sum_{\omega \in \Omega} F_{\theta, a^\theta, \omega} T_{f(\theta), \omega} - c_{\theta, a^\theta} \geq \sum_{\omega \in \Omega} F_{\theta, a, \omega} T_{i, \omega} - c_{\theta, a} \quad \forall \theta \in \Theta, \forall a \in A, \forall i \in \{1, \dots, k\}.$$

Once a function f and a tuple of actions $(a^\theta)_{\theta \in \Theta}$ are fixed, by restricting the attention to menus of deterministic contracts in $\mathcal{Q}^{f, (a^\theta)_{\theta \in \Theta}}$, an optimal menu can be obtained as an optimal solution to the following LP:

$$\max_{T \in \mathbb{R}_+^{k \times m}} \sum_{\theta \in \Theta} \mu_\theta \sum_{\omega \in \Omega} F_{\theta, a^\theta, \omega} (r_\omega - T_{f(\theta), \omega}) \quad \text{s.t.} \tag{B.2a}$$

$$\sum_{\omega \in \Omega} F_{\theta, a^\theta} T_{f(\theta), \omega} - c_{\theta, a^\theta} \geq \sum_{\omega \in \Omega} F_{\theta, a, \omega} T_{i, \omega} - c_{\theta, a} \quad \forall \theta \in \Theta, \forall a \in A, \forall i \in \{1, \dots, k\} \tag{B.2b}$$

$$T_{i, \omega} \geq 0 \quad \forall i \in \{1, \dots, k\}, \forall \omega \in \Omega. \tag{B.2c}$$

Since the objective of the LP is a linear function with non-positive coefficients and the feasible region is $\mathbb{R}_+^{k \times m}$, the LP admits an optimal solution which is one of the vertexes of the polyhedron defined by Constraints (B.2b) and (B.2c). In order to conclude the proof, notice that, for every possible function f and tuple of actions $(a^\theta)_{\theta \in \Theta}$, the feasible region of the LP is defined by a subset of the following $\ell n^2 k^2 + km$ constraints:

$$\sum_{\omega \in \Omega} F_{\theta, a'} T_{j, \omega} - c_{\theta, a'} \geq \sum_{\omega \in \Omega} F_{\theta, a, \omega} T_{i, \omega} - c_{\theta, a} \quad \forall \theta \in \Theta, \forall a \in A, \forall a' \in A, \forall i, j \in \{1, \dots, k\}$$

$$T_{i, \omega} \geq 0 \quad \forall i \in \{1, \dots, k\}, \forall \omega \in \Omega.$$

Since each vertex is at the intersection of km linearly independent hyperplanes representing the inequalities and there are at most $\ell n^2 k^2 + km$ different hyperplanes, there are at most $\binom{\ell n^2 k^2 + km}{km}$ possible vertexes. Finally, the vertexes can be enumerated in polynomial time. \square

Appendix C. Proofs omitted from Section 6

Lemma 8. *In Bayesian principal-agent instances (Θ, A, Ω) with $|\Omega| = 2$, there exists a contract having the same principal's expected utility as an optimal DSIC menu of deterministic contracts.*

Proof. Let $\Omega = \{\omega_0, \omega_1\}$ be the set made by the two outcomes. Moreover, w.l.o.g., let ω_0 be the outcome with smaller reward, i.e., $r_{\omega_0} \leq r_{\omega_1}$. We show that, given a menu of deterministic contracts $P = (p^\theta)_{\theta \in \Theta}$, we can build a contract $\hat{p} \in \mathbb{R}_+^m$ providing the principal with at least the same principal's expected utility. Let $\hat{\theta} \in \arg \max_{\theta \in \Theta} p_{\omega_1}^\theta$. Then, we define the contract \hat{p} such that $\hat{p}_{\omega_1} = \min \{p_{\omega_1}^{\hat{\theta}}, r_{\omega_1} - r_{\omega_0}\}$ and $\hat{p}_{\omega_0} = 0$. We show that, by replacing the menu $P = (p^\theta)_{\theta \in \Theta}$ with the contract \hat{p} , for each agent's type $\theta \in \Theta$, the principal's expected utility does *not* decrease.

Take any type $\theta \in \Theta$. As a preliminarily step, notice that

$$F_{\theta, b^\theta(p^\theta), \omega_1} \hat{p}_{\omega_1} \leq F_{\theta, b^\theta(p^\theta), \omega_1} p_{\omega_1}^{\hat{\theta}} + (1 - F_{\theta, b^\theta(p^\theta), \omega_1}) p_{\omega_0}^{\hat{\theta}}$$

$$\leq F_{\theta, b^\theta(p^\theta), \omega_1} p_{\omega_1}^\theta + (1 - F_{\theta, b^\theta(p^\theta), \omega_1}) p_{\omega_0}^\theta, \tag{C.1}$$

where the first inequality follows from the definition of \hat{p} and the second one from DSIC. Then, we consider two cases. First, suppose that $b^\theta(\hat{p}) = b^\theta(p^\theta)$. Then, the expected payment under the contract p for an agent of type θ is

$$F_{\theta, b^\theta(\hat{p}), \omega_1} \hat{p}_{\omega_1} \leq F_{\theta, b^\theta(p^\theta), \omega_1} p_{\omega_1}^\theta + (1 - F_{\theta, b^\theta(p^\theta), \omega_1}) p_{\omega_0}^\theta$$

where the inequality comes from Equation (C.1) and $b^\theta(\hat{p}) = b^\theta(p^\theta)$. Hence, the payment decreases while the revenue does not, implying that the principal's expected utility does *not* decrease.

Second, suppose that $b^\theta(\hat{p}) \neq b^\theta(p^\theta)$. By IC conditions, we have

$$F_{\theta, b^\theta(\hat{p}), \omega_1} \hat{p}_{\omega_1} - c_{\theta, b^\theta(\hat{p})} \geq F_{\theta, b^\theta(p^\theta), \omega_1} \hat{p}_{\omega_1} - c_{\theta, b^\theta(p^\theta)},$$

and

$$F_{\theta, b^\theta(\hat{p}), \omega_1} p_{\omega_1}^\theta + (1 - F_{\theta, b^\theta(\hat{p}), \omega_1}) p_{\omega_0}^\theta - c_{\theta, b^\theta(\hat{p})} \leq F_{\theta, b^\theta(p^\theta), \omega_1} p_{\omega_1}^\theta + (1 - F_{\theta, b^\theta(p^\theta), \omega_1}) p_{\omega_0}^\theta - c_{\theta, b^\theta(p^\theta)}.$$

By summing the two inequalities, we get

$$F_{\theta, b^\theta(\hat{p}), \omega_1} (\hat{p}_{\omega_1} - p_{\omega_1}^\theta + p_{\omega_0}^\theta) \geq F_{\theta, b^\theta(p^\theta), \omega_1} (\hat{p}_{\omega_1} - p_{\omega_1}^\theta + p_{\omega_0}^\theta). \tag{C.2}$$

We consider two cases. If $p_{\omega_1}^\theta \leq r_{\omega_1} - r_{\omega_0}$, then $\hat{p}_{\omega_1} \geq p_{\omega_1}^\theta$, which implies $F_{\theta, b^\theta(\hat{p}), \omega_1} \geq F_{\theta, b^\theta(p^\theta), \omega_1}$ by Equation (C.2). Then, the principal's expected utility when the agent plays action $b^\theta(\hat{p})$ is

$$\begin{aligned} F_{\theta, b^\theta(\hat{p}), \omega_1} (r_{\omega_1} - \hat{p}_{\omega_1}) + (1 - F_{\theta, b^\theta(\hat{p}), \omega_1}) r_{\omega_0} &\geq F_{\theta, b^\theta(p^\theta), \omega_1} (r_{\omega_1} - \hat{p}_{\omega_1}) + (1 - F_{\theta, b^\theta(p^\theta), \omega_1}) r_{\omega_0} \\ &\geq F_{\theta, b^\theta(p^\theta), \omega_1} (r_{\omega_1} - p_{\omega_1}^\theta) + (1 - F_{\theta, b^\theta(p^\theta), \omega_1}) (r_{\omega_0} - p_{\omega_0}^\theta), \end{aligned}$$

where the first inequality comes from $r_{\omega_1} - \hat{p}_{\omega_1} \geq r_{\omega_0}$ (by construction) and $F_{\theta, b^\theta(\hat{p}), \omega_1} \geq F_{\theta, b^\theta(p^\theta), \omega_1}$, while the second one comes from Equation (C.1). Hence, the principal's utility does not decrease.

If $p_{\omega_1}^\theta \geq r_{\omega_1} - r_{\omega_0}$, then the expected reward is

$$\begin{aligned} F_{\theta, b^\theta(\hat{p}), \omega_1} (r_{\omega_1} - \hat{p}_{\omega_1}) + F_{\theta, b^\theta(\hat{p}), \omega_0} r_{\omega_0} &\geq F_{\theta, b^\theta(\hat{p}), \omega_1} r_{\omega_0} + F_{\theta, b^\theta(\hat{p}), \omega_0} r_{\omega_0} \\ &= r_{\omega_0} \geq F_{\theta, b^\theta(p^\theta), \omega_1} (r_{\omega_1} - p_{\omega_1}^\theta) + F_{\theta, b^\theta(p^\theta), \omega_0} (r_{\omega_0} - p_{\omega_0}^\theta) \end{aligned}$$

where the first inequality holds since $\hat{p}_{\omega_1} \leq r_{\omega_1} - r_{\omega_0}$ by construction and the last inequality holds since both $r_{\omega_0} - p_{\omega_0}^\theta$ and $r_{\omega_1} - p_{\omega_1}^\theta$ are smaller than r_{ω_0} . Thus, the principal's utility does not decrease. This concludes the proof. \square

Theorem 6. *In Bayesian principal-agent instances with a constant number of agent's types, there exists a polynomial-time algorithm that computes an optimal DSIC menu of deterministic contracts.*

Proof. Consider an optimal menu of contracts $(p^\theta)_{\theta \in \Theta}$. This optimal menu of contracts induces an action $a^\theta = b^\theta(p^\theta)$ for each possible receiver's type $\theta \in \Theta$. Once the action profile $(a^\theta)_{\theta \in \Theta}$ is determined, the optimal menu of contracts can be found by solving the following LP of polynomial size.

$$\min_{p \in \mathbb{R}_+^{\ell \times m}} \sum_{\theta} \mu_{\theta} \sum_{\omega \in \Omega} p_{\omega} F_{\theta, a^\theta, \omega} \quad \text{s.t.} \tag{C.4a}$$

$$\sum_{\omega \in \Omega} p_{\omega}^\theta F_{\theta, a^\theta, \omega} - c_{\theta, a^\theta} \geq \sum_{\omega \in \Omega} p_{\omega}^{\theta'} F_{\theta, a, \omega} - c_{\theta, a} \quad \forall \theta, \theta' \in \Theta, a \in A \tag{C.4b}$$

We can enumerate all the actions' profiles and, for each action profile, find an optimal menu of contracts that incentivize this action profile. Since there are $O(n^\ell)$ action profiles, the algorithm works in polynomial time if the number of types is fixed. \square

References

[1] M. Castiglioni, A. Marchesi, N. Gatti, Designing menus of contracts efficiently: the power of randomization, in: Proceedings of the 23rd ACM Conference on Economics and Computation, EC '22, Association for Computing Machinery, New York, NY, USA, ISBN 9781450391504, 2022, pp. 705–735.
 [2] C.J. Ho, A. Slivkins, J.W. Vaughan, Adaptive contract design for crowdsourcing markets: bandit algorithms for repeated principal-agent problems, *J. Artif. Intell. Res.* 55 (2016) 317–359.
 [3] L.W. Cong, Z. He, Blockchain disruption and smart contracts, *Rev. Financ. Stud.* 32 (2019) 1754–1797.
 [4] H. Bastani, M. Bayati, M. Braverman, R. Gummadi, R. Johari, Analysis of medicare pay-for-performance contracts, Available at SSRN 2839143, 2016.
 [5] G. Guruganesh, J. Schneider, J.R. Wang, Contracts under moral hazard and adverse selection, in: Proceedings of the 22nd ACM Conference on Economics and Computation, EC '21, Association for Computing Machinery, New York, NY, USA, ISBN 9781450385541, 2021, pp. 563–582.
 [6] M. Castiglioni, A. Marchesi, N. Gatti, Bayesian agency: linear versus tractable contracts, *Artif. Intell.* 307 (2022) 103684.
 [7] T. Alon, P. Dütting, I. Talgam-Cohen, Contracts with private cost per unit-of-effort, in: Proceedings of the 22nd ACM Conference on Economics and Computation, 2021, pp. 52–69.
 [8] P. Dütting, T. Roughgarden, I. Talgam-Cohen, Simple versus optimal contracts, in: Proceedings of the 2019 ACM Conference on Economics and Computation, 2019, pp. 369–387.
 [9] S. Shavell, Risk sharing and incentives in the principal and agent relationship, *Bell J. Econ.* (1979) 55–73.
 [10] S.J. Grossman, O.D. Hart, An analysis of the principal-agent problem, *Econometrica* 51 (1) (1983) 7–46.
 [11] W.P. Rogerson, Repeated moral hazard, *Econometrica* (1985) 69–76.
 [12] B. Holmstrom, P. Milgrom, Multitask principal-agent analyses: incentive contracts, asset ownership, and job design, *J. Law Econ. Organ.* 7 (1991) 24.
 [13] A. Mas-Colell, M.D. Whinston, J.R. Green, et al., *Microeconomic Theory*, vol. 1, Oxford University Press, New York, 1995.
 [14] P. Bolton, M. Dewatripont, et al., *Contract Theory*, MIT Press, 2005.

- [15] J.J. Laffont, D. Martimort, *The Theory of Incentives: the Principal-Agent Model*, Princeton University Press, 2009.
- [16] M. Babaioff, M. Feldman, N. Nisan, Combinatorial agency, in: *Proceedings of the 7th ACM Conference on Electronic Commerce*, 2006, pp. 18–28.
- [17] M. Babaioff, M. Feldman, N. Nisan, E. Winter, Combinatorial agency, *J. Econ. Theory* 147 (3) (2012) 999–1034.
- [18] M. Babaioff, M. Feldman, N. Nisan, Free-riding and free-labor in combinatorial agency, in: *International Symposium on Algorithmic Game Theory*, Springer, 2009, pp. 109–121.
- [19] M. Babaioff, M. Feldman, N. Nisan, Mixed strategies in combinatorial agency, *J. Artif. Intell. Res.* 38 (2010) 339–369.
- [20] P. Dütting, T. Roughgarden, I.T. Cohen, The complexity of contracts, in: *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, 2020, pp. 2688–2707.
- [21] P. Dütting, T. Ezra, M. Feldman, T. Kesselheim, Combinatorial contracts, in: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, 2022, pp. 815–826.
- [22] M. Babaioff, E. Winter, Contract complexity, in: *EC '14*, 2014, p. 911.
- [23] T. Alon, P. Dütting, I. Talgam-Cohen, Contracts with private cost per unit-of-effort, arXiv preprint, arXiv:2111.09179, 2021.
- [24] O. Kadan, P.J. Reny, J.M. Swinkels, Existence of optimal mechanisms in principal-agent problems, *Econometrica* 85 (3) (2017) 769–823, <http://www.jstor.org/stable/44955140>.
- [25] R.B. Myerson, Optimal coordination mechanisms in generalized principal-agent problems, *J. Math. Econ.* 10 (1) (1982) 67–81.
- [26] J. Gan, M. Han, J. Wu, H. Xu, Optimal coordination in generalized principal-agent problems: a revisit and extensions, <https://arxiv.org/abs/2209.01146>, 2022, <https://doi.org/10.48550/ARXIV.2209.01146>.
- [27] M. Castiglioni, A. Marchesi, N. Gatti, Bayesian persuasion meets mechanism design: going beyond intractability with type reporting, in: *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 2022, pp. 226–234.
- [28] G. Carroll, Robustness and linear contracts, *Am. Econ. Rev.* 105 (2) (2015) 536–563.
- [29] Y. Shoham, K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge University Press, 2008.
- [30] Y. Ye, Recovering optimal basic variables in Karmarkar's polynomial algorithm for linear programming, *Math. Oper. Res.* 15 (3) (1990) 564–572.
- [31] N. Alon, U. Feige, A. Wigderson, D. Zuckerman, Derandomized graph products, *Comput. Complex.* 5 (1) (1995) 60–75, <https://doi.org/10.1007/BF01277956>.
- [32] L. Trevisan, Non-approximability results for optimization problems on bounded degree instances, in: *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, STOC '01, Association for Computing Machinery, New York, NY, USA, ISBN 1581133499, 2001, pp. 453–461.
- [33] T. Alon, P. Dütting, Y. Li, I. Talgam-Cohen, Bayesian analysis of linear contracts, <https://arxiv.org/abs/2211.06850>, 2022, <https://doi.org/10.48550/ARXIV.2211.06850>.
- [34] D. Bertsimas, J.N. Tsitsiklis, *Introduction to Linear Optimization*, vol. 6, Athena Scientific Belmont, MA, 1997.