Contract Design Under Approximate Best Responses

Francesco Bacchiocchi¹ Jiarui Gan² Matteo Castiglioni¹ Alberto Marchesi¹ Nicola Gatti¹

Abstract

Principal-agent problems model scenarios where a principal incentivizes an agent to take costly, unobservable actions through the provision of payments. Such problems are ubiquitous in several real-world applications, ranging from blockchain to the delegation of machine learning tasks. In this paper, we initiate the study of hidden-action principal-agent problems under approximate best responses, in which the agent may select any action that is not too much suboptimal given the principal's payment scheme (a.k.a. *contract*). Our main result is a polynomial-time algorithm to compute an optimal contract under approximate best responses. This positive result is perhaps surprising, since, in Stackelberg games, computing an optimal commitment under approximate best responses is computationally intractable. We also investigate the learnability of contracts under approximate best responses, by providing a no-regret learning algorithm for a natural application scenario where the principal has no prior knowledge about the environment.

1. Introduction

In *hidden-action principal-agent* problems, a principal tries to steer the behavior of a self-interested agent toward favorable outcomes. The agent has to take a costly action that stochastically determines an outcome resulting in a reward for the principal. The main challenge is that the agent's action is *hidden* to the principal, who can only observe the realized outcome. Thus, the principal influences the agent's behavior by *committing to a contract*, which is an outcome dependent payment scheme whose aim is to induce the agent to take a high-cost action leading to high principal's rewards. The principal's goal is to design an *optimal* contract, namely one maximizing their utility, *i.e.*, rewards minus payments.

Nowadays, principal-agent problems find application in a terrific number of real-world settings, such as, *e.g.*, crowd-sourcing (Ho et al., 2014), online labor platforms (Kaynar & Siddiq, 2023), blockchain (Cong & He, 2019), delegation of machine learning tasks (Cai et al., 2015), and pay-forperformance healthcare (Bastani et al., 2016; 2019). Moreover, algorithmic contract design is playing a crucial role in today's world, which increasingly relies on AI agents to perform complex tasks (see, *e.g.*, (Hadfield-Menell & Hadfield, 2019; Saig et al., 2024)).

Previous works on algorithmic contract design assume that the agent always plays a *best-response action* to the principal's contract. However, if the agent actually responds (even slightly) suboptimally to the principal, then the principal's utility may substantially deteriorate. This may be the case in most of the real-world applications of interest, for several different reasons. For instance, the principal may *not* perfectly know agent's features and account for the wrong best response, the agent may *not* be powerful enough to compute an (exact) best-response action, or they may inaccurately interpret the principal's contract.

In this paper, we initiate the study of hidden-action principalagent problems *under approximate agent's best responses*. Specifically, we consider settings in which the agent may take actions that are up to $\delta \in (0, 1)$ suboptimal under the principal's contract. We do *not* make any assumption on the specific δ -best response selected by the agent, but we allow them to take *any* of such actions. Thus, we take a *worst-case approach* and consider the problem of designing contracts under the assumption that the agent selects the worst δ -best response for the principal. Contracts designed in such a way are said to be *robust*, as they guard the principal against any possible (conceivable) suboptimal behavior of the agent.

1.1. Results and Techniques

In this paper, we provide an extensive treatment of the computational and learnability aspects of the design of robust contracts. The results presented in the paper are organized into three main parts as summarized below.

The Price of Robustness In the first part of the paper, as a preliminary analysis, we provide a characterization of the maximum utility that the principal can achieve by means of

¹Politecnico di Milano ²University of Oxford. Correspondence to: Francesco Bacchiocchi <francesco.bacchiocchi@polimi.it>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

robust contracts. Specifically, we provide upper and lower bounds on this utility, as a function of a parameter $\delta \in (0, 1)$ quantifying the agent's best response suboptimality. These bounds give insights on the price (in terms of utility) that the principal incurs for being robust, as the parameter δ can be seen as a measure of the robustness level of the principal's contracts. Interestingly, our results show that, differently from what happens in general Stackelberg games (see, e.g., (Gan et al., 2023)), upper/lower bounds do not depend on the inducibility gap $\Delta > 0$ characterizing the problem instance. In order to derive the bounds, we prove that it is possible to convert any non-robust contract into a robust one by properly moving its payments in the direction of principal's rewards, and that, in an optimal robust contract, which provides the principal with a utility greater than zero, the agent's utility should be at least δ .

Computing Robust Contracts The second part of the paper addresses the problem of computing a robust contract that is optimal for the principal. Our main contribution is a polynomial-time algorithm for this problem. This is perhaps surprising, since analogous problems in Stackelberg and Bayesian persuasion settings are known to be computationally intractable (Gan et al., 2023; Yang & Zhang, 2024), despite having more amenable solution spaces (Δ_m in these problems vs. \mathbb{R}^m_+ in contract design). At a high level, our algorithm cleverly exploits a particular structure that we discover in the robustness constraints (which ensure that the agent plays the worst approximate best response for the principal). Intuitively, we show that, when the agent's best response and their (worst) approximate best response are fixed to two arbitrarily-selected actions, the problem of computing a utility-maximizing robust contract can be formulated by means of a union of n + 1 different *linear* programs (LPs). Therefore, by taking the maximum over all these LPs, one can compute a utility-maximizing robust contract once the agent's best response and their (worst) approximate best response are fixed. By further iterating this procedure for all the possible choices of the action pair, we then determine an optimal robust contract.

Learning Robust Contracts In the third and last part of the paper, we investigate the learnability of robust contracts. Specifically, we study an online learning framework similar to the one analyzed by Zhu et al. (2023), in which the features of agent's actions, *i.e.*, costs and probabilities over outcomes, depend on an agent's *type* that is sampled at each round from some (fixed) unknown probability distribution. At each round, after committing to some contract, the principal only observes the outcome (and its associated reward) realized as an effect of the approximate best response played by the agent. Our main result within this online learning framework is the design of an algorithm that achieves sublinear regret with respect to always playing an optimal robust

contract. Additionally, we show that when the parameter $\delta \in (0, 1)$ measuring the suboptimality of agent's actions is sufficiently small with respect to the time horizon, our algorithm also achieves sublinear regret with respect to always playing an optimal *non-robust* contract. Our approach presents some advantages compared to the state-of-the-art proposed by Zhu et al. (2023) for the non-robust version of the problem, while achieving similar regret guarantees. Indeed, our algorithm employs a simpler discretization of the set of contracts employed by the principal, which does *not* require that the principal knows its rewards beforehand. This is made possible by the fact that we rely on a novel "continuity" argument for the principal's expected utility (see Lemma 3), which is different from the one originally proposed by Zhu et al. (2023).

1.2. Related Works

Robustness to Approximate Best Responses Two works that are closely related to ours are (Gan et al., 2023; Yang & Zhang, 2024), which consider robustness notions analogous to ours, though in different settings. Specifically, Gan et al. (2023) initiate this research line, by studying th problem of computing robust leader's commitments in *Stackelberg* games where the follower plays an approximate best response. They show that it is NP-hard to approximate an optimal robust commitment of the leader and, in accordance to this hardness result, they provide a quasi-polynomial-time approximation scheme (QPTAS). Yang & Zhang (2024) extend the study initiated by Gan et al. (2023) to Bayesian persuasion, with the goal of computing robust signaling schemes under approximate best responses of the receiver. Similarly to Gan et al. (2023), they show that computing an approximately-optimal robust signaling scheme is NP-hard and provide a QPTAS. In sharp contrast with these works, we show that, in hidden-action principal-agent problems, an optimal robust contract can be computed efficiently.

Contract Design Contract theory has been extensively studied in economics (Holmstrom & Milgrom, 1991). However, the interest in its computational aspects is more recent. Dütting et al. (2019) analyze linear contracts, proving approximation bounds. Guruganesh et al. (2021); Castiglioni et al. (2022a;b) study the computational aspects of contract design in Bayesian settings. Alon et al. (2023) study Bayesian linear contracts, proving near-optimality under sufficient uncertainty. Other works explore combinatorial principal-agent problems (Babaioff et al., 2006; 2009) and multi-agent extensions of hidden-action problems (Castiglioni et al., 2023; Duetting et al., 2024).

Other Forms of Robustness in Contract Design Our work is also related to other research lines addressing different concepts of *robustness* in contract design. For instance,

Carroll (2015) studies settings where the principal only knows a superset of agents' actions, while Dütting et al. (2019) introduce a different notion of uncertainty in which the principal has partial knowledge of the distributions over outcomes associated with agent's actions. Both these works show that linear contracts are a sufficient class of contracts to determine the *min-max robust* optimal contract. Notice that these two frameworks differ from ours, as within our framework, when the robustness parameter $\delta \in (0, 1)$ is arbitrarily small, the problem becomes very close to the classical version of the hidden-action principal-agent problem, in which it is known that linear contracts are *not* generally optimal (see, e.g., (Dütting et al., 2019)). Recently, Bernasconi et al. (2024) study settings where uncertainty lies in the costs of agent's actions. In this framework, the principal only knows a set containing the true cost vectors, and computing an optimal min-max robust contract is APX-hard.

Learning in Principal-Agent Problems Our work is also related to online learning problems in hidden-action principal-agent problems. Zhu et al. (2023) study general hidden-action principal-agent problem instances in which the principal faces multiple agent's types. They show that it is possible to design an algorithm that achieves a regret bound of the order of $\widetilde{\mathcal{O}}(\sqrt{m} \cdot T^{1-1/(2m+1)})$ when the principal selects contracts from the hypercube $[0, 1]^m$, where m is the number of outcomes. In our work, we show that it is possible to design an algorithm achieving similar regret guarantees even when the different agent's types select approximate best responses. Our algorithm presents some advantages compared to the one proposed by Zhu et al. (2023). Specifically, our approach employs a simpler discretization of the hypercube-used during the execution of the algorithm-compared to (Zhu et al., 2023), and it does not require prior knowledge of the principal's rewards. Furthermore, Zhu et al. (2023) provide an (almost-matching) lower bound of $\Omega(T^{1-1/(m+2)})$, which holds even with a single agent's type. Some recent works have introduced additional hypothesis to overcome this negative result (see, e.g., (Bacchiocchi et al., 2024; Chen et al., 2024)).

2. Preliminaries

In this section, we first introduce the classical hidden-action principal-agent problems (Section 2.1), and then the variation studied in this paper, in which the agent plays an approximate best response (Section 2.2).

2.1. Hidden-Action Principal-Agent Problems

An instance of *hidden-action principal-agent problem* is characterized by a tuple $(\mathcal{A}, \Omega, F, r, c)$, where \mathcal{A} is a finite set of $n := |\mathcal{A}|$ actions available to the agent, Ω is a finite set of $m := |\Omega|$ possible outcomes, $F \in [0, 1]^{m \times n}$ is a matrix representing the effects of agent's action, $r \in [0, 1]^m$ is a reward vector for the principal, and $c \in [0, 1]^n$ is a vector of agent's costs. Each agent's action $a \in \mathcal{A}$ determines a probability distribution over outcomes, encoded by a column $F_a \in \Delta_{\Omega}$ of the matrix F, and it results in a cost for the agent, encoded by a component $c_a \in [0, 1]$ of vector c.¹ We denote by $F_{a,\omega} \in [0, 1]$ the probability with which action a results in outcome $\omega \in \Omega$, as prescribed by F_a . Thus, it must be the case that $\sum_{\omega \in \Omega} F_{a,\omega} = 1$ for all $a \in \mathcal{A}$. Each outcome $\omega \in \Omega$ is associated with a reward for the principal, encoded by a component $r_{\omega} \in [0, 1]$ of the vector r. Thus, whenever the agent selects an action $a \in \mathcal{A}$, the principal's expected reward can be computed as $R_a := F_a \cdot r$.²

The principal commits to an outcome-dependent payment scheme called *contract*, which is a vector $p \in \mathbb{R}^m_+$ defining a payment $p_{\omega} \ge 0$ from the principal to the agent for every outcome $\omega \in \Omega$.³ Given a contract $p \in \mathbb{R}^m_+$, the agent plays a *best-response* action that is: (i) *incentive compatible* (IC), which means that it maximizes their expected utility; and (ii) *individually rational* (IR), meaning that it has non-negative expected utility. We assume w.l.o.g. that there always exists at least one *opt-out* action $a \in \mathcal{A}$ with null cost, *i.e.*, $c_a = 0$, and for which the principal's expected reward is equal to zero, *i.e.*, $F_a \cdot r = 0$. Since the agent's utility is at least zero when playing the *opt-out* action, any IC action is also IR, allowing us to focus on incentive compatibility only.

Whenever the principal commits to a contract $p \in \mathbb{R}^m_+$ and the agent responds by playing an action $a \in \mathcal{A}$, the agent's and the principal's expected utilities are, respectively,

$$u^{\mathsf{A}}(p,a) \coloneqq F_a \cdot p - c_a, \text{ and } u^{\mathsf{P}}(p,a) \coloneqq F_a \cdot (r-p).$$

The set $A(p) \subseteq A$ of agent's best responses in a contract $p \in \mathbb{R}^m_+$ is defined as follows:

$$A(p) \coloneqq \arg \max_{a \in \mathcal{A}} \left\{ F_a \cdot p - c_a \right\}.$$

In classical (non-robust) hidden-action principal-agent problems, the agent breaks ties in favor of the principal when having multiple best responses available (see, *e.g.*, (Dütting et al., 2019)). We denote by $a(p) \in A(p)$ the action played by the agent in a given contract $p \in \mathbb{R}^m_+$. This is an action $a \in A(p)$ that maximizes the principal's utility $F_a \cdot (r - p)$. Formally, $a(p) \in \arg \max_{a \in A(p)} F_a \cdot (r - p)$. Then, the goal of the principal is to design a contract $p \in \mathbb{R}^m_+$ that maximizes their expected utility $u^P(p, a(p))$. We say that a contract $p^* \in \mathbb{R}^m_+$ is a (non-robust) *optimal* contract if it holds $p^* \in \arg \max_{p \in \mathbb{R}^m_+} u^P(p, a(p))$. In the following, we define

¹We denote by Δ_X the set of all probability distributions over the set X. Given $n \in \mathbb{N}_{>0}$, we write $[n] := \{1, \ldots, n\}$.

²We denote by $x \cdot y$ the dot product of two vectors $x, y \in \mathbb{R}^d$.

³As customary in contract theory (Carroll, 2015), we assume that the agent has *limited liability*, meaning that the payments can only be from the principal to the agent, and *not viceversa*.

the principal's utility in an optimal (non-robust) contract as OPT := $\max_{p \in \mathbb{R}^m_+} u^{\mathbb{P}}(p, a(p))$, while we let the value of the social welfare be SW := $\max_{a \in \mathcal{A}} \{F_a \cdot r - c_a\}$.

2.2. Robust Contracts and Approximate Best Responses

In this paper, we study a variation of the classical hiddenaction principal-agent problem, where the agent plays an action that is an *approximate* best response. Given $\delta \in (0, 1)$, we define the set $A^{\delta}(p) \subseteq A$ of agent's δ -best responses in a given contract $p \in \mathbb{R}^m_+$ as follows:

$$A^{\delta}(p) \coloneqq \left\{ a \in \mathcal{A} \mid F_a \cdot p - c_a > \max_{a' \in \mathcal{A}} \left\{ F_{a'} \cdot p - c_{a'} \right\} - \delta \right\}.$$

We adopt an adversarial robust approach, in the sense that, whenever the principal commits to a contract $p \in \mathbb{R}^m_+$, the agent selects a δ -best response that minimizes principal's expected utility, namely an action $a^{\delta}(p) \in A^{\delta}(p)$ such that $a^{\delta}(p) \in \arg \min_{a \in A^{\delta}(p)} F_a \cdot (r-p)$. We refer to the utility of a δ -robust contract p as $u^{\mathsf{P}}(p, a^{\delta}(p))$.

Given $\delta \in (0, 1)$, the principal's goal is to design an *optimal* δ -robust contract $p^* \in \mathbb{R}^m_+$, which is formally defined as:

$$p^{\star} \in \underset{p \in \mathbb{R}^m_+}{\arg \max} \min_{a \in A^{\delta}(p)} u^{\mathsf{P}}(p, a) = \underset{p \in \mathbb{R}^m_+}{\arg \max} \Psi(p), \quad (1)$$

where $\Psi(p) := \min_{a \in A^{\delta}(p)} u^{\mathsf{P}}(p, a)$ denotes the principal's expected utility in a δ -robust contract $p \in \mathbb{R}^{m}_{+}$. Notice that analogous " δ -robust solution concepts" have been already introduced in similar settings, namely Stackelberg games (Gan et al., 2023) and Bayesian persuasion (Yang & Zhang, 2024). Our δ -robust contracts are their analogous in the context of hidden-action principal-agent problems.

In the following, given $\delta \in (0, 1)$, we denote the expected utility of the principal in an optimal δ -robust contract p^* as $OPT(\delta) := u^P(p^*, a^{\delta}(p^*))$. We remark that $OPT(\delta)$ is always well defined, as an optimal δ -robust contract, according to the definition in Eq. (1), always exists. Intuitively, this is due to the strict inequality in the definition of the set $A^{\delta}(p)$, as observed by Gan et al. (2023) for Stackelberg games. Thus, in order to prove the existence of an optimal δ -robust contract, it is possible to employ the same argument used to prove Proposition 1 in (Gan et al., 2023).

3. The Price of Robustness

We start by providing a characterization of how the value OPT(δ) of an optimal δ -robust contract varies as a function of the parameter $\delta \in (0, 1)$, which controls the suboptimality of the agent's best response. The goal of our analysis is to quantify how much is the *price (in terms of utility) that the principal incurs for being robust* to agent's approximate best responses. Indeed, the parameter δ can be interpreted as

a measure of the robustness level of the principal's contract, with higher δ values indicating higher levels of robustness.

We first establish upper and lower bounds that identify a suitable region in which the values of $OPT(\delta)$ are contained. Such bounds only depend on the parameter $\delta \in (0, 1)$, the value of an optimal non-robust contract OPT, and the social welfare SW achievable with non-robust contracts.

Proposition 1 (Upper and lower bounds). *Given an instance of hidden-action principal-agent problem:*

1. For every $\delta \in (0, 1)$, it holds:

$$OPT(\delta) \ge OPT_{LB}(\delta) \coloneqq OPT - 2\sqrt{\delta} + \delta.$$

2. For every $\delta \in (0, 1)$, it holds:

$$OPT(\delta) \le OPT_{UB}(\delta) \coloneqq \max\{0, SW - \delta\}.$$

In order to prove the first point in Proposition 1, we show that it is always possible to convert a non-robust contract into a robust one by suitably moving it towards the direction of the principal's reward vector. In order to prove the second point of Proposition 1, we observe that, if an optimal δ robust contract p^* provides the principal with an expected utility strictly larger than zero, then the opt-out action does *not* belong to the set of δ -best responses for such a contract. Therefore, in an optimal δ -robust contract that provides the principal with an expected utility strictly larger than zero, the agent's utility is at least $\delta > 0$, and thus the principal's expected utility is at most SW $-\delta$.

The following proposition shows that the upper and lower bounds in Proposition 1 are tight. Formally:

Proposition 2. Given $\delta \in (0,1)$, for every integer $n > n(\delta)$, there is an instance of hidden-action principal-agent problem (parametrized by δ) with 2n + 1 actions, where:

$$OPT(\delta) - OPT_{LB}(\delta) \le \mathcal{O}\left(\frac{1}{n}\right).$$

Furthermore, for any $\delta \in (0, 1)$, there exists an instance of hidden-action principal-agent problem in which it holds:

$$OPT(\delta) = OPT_{UB}(\delta).$$

Proposition 2 shows that our upper bound $OPT_{UB}(\delta)$ is (strictly) tight, while our lower bound $OPT_{LB}(\delta)$ is tight up to an additive term that linearly goes to zero as the number of agent's actions *n* increases.

Finally, thanks to the previous propositions, we can show the following property about the value of an optimal δ -robust contracts as a function of $\delta \in (0, 1)$.

Proposition 3. The function $(0,1) \ni \delta \mapsto \text{OPT}(\delta)$ is continuous and non-increasing in δ . Moreover, $\lim_{\delta \to 0^+} \text{OPT}(\delta) = \text{OPT}$ and $\lim_{\delta \to 1^-} \text{OPT}(\delta) = 0$.



Figure 1. The blue area corresponds to the region in which $OPT(\delta)$ is bounded as a function of $\delta \in (0, 1)$ in an instance of hiddenaction principal-agent problem with SW = 0.9 and OPT = 0.7.

Figure 1 shows an example of the region in which the value of $OPT(\delta)$ is bounded, as defined by the lower and upper bounds in Propositions 1 and 3 as functions of δ .

Comparison With Stackelberg Games The characterization results derived in this section exhibit some crucial differences compared to analogous results derived for Stackelberg games by Gan et al. (2023). Indeed, in such settings, Gan et al. (2023) show that the value of an optimal δ -robust commitment crucially depends on a parameter $\Delta > 0$ that represents the *inducibility gap* of the problem instance. Intuitively, the inducibility gap encodes how easy it is for the leader to induce the follower to play any action; see (Gan et al., 2023) for a formal definition. Specifically, in Stackelberg games, the leader's expected utility $OPT(\delta)$ in an optimal δ -robust commitment is a Lipschitz function lower bounded by OPT $-\delta/\Delta$ if $\delta < \Delta$, whereas OPT(δ) may *not* be even a continuous function if $\delta > \Delta$. In contrast, in hidden-action principal-agent problems, the value of an optimal δ -robust contract is a continuous function with respect to $\delta \in (0, 1)$, regardless of the inducibility gap of the instance. Furthermore, the value of $OPT(\delta)$ is either zero or it is upper bounded by SW $-\delta$, showing that for large values of δ , the maximum utility that the principal can achieve may be particularly small. Intuitively, this is because the principal must provide the agent with a large expected payment to induce them to take desirable actions rather than the opt-out one. This upper bound does not generally hold in Stackelberg games in which the principal may achieve a large utility even for large values of δ .

4. Computing Optimal Robust Contracts

Now, we present our main result: a polynomial-time algorithm to compute an optimal δ -robust contract.

4.1. Characterizing an Optimal Contract

We begin by presenting an optimization problem that characterizes an optimal δ -robust contract. Consider an arbitrary optimal δ -robust contract p^* , as well as two arbitrary actions $a^* \in A(p^*)$ and $a^{\delta} \in \arg \min_{a \in A^{\delta}(p^*)} u^{\mathsf{P}}(p^*, a)$. By fixing a^* and a^{δ} , we show that the following optimization problem, over the variable $p \in \mathbb{R}^m_+$, characterizes an optimal δ -robust contract p^* .

$$\max_{p \in \mathbb{R}^m_+} \quad u^{\mathsf{P}}(p, a^{\delta}) \tag{2}$$

subject to the following *disjunctive* constraints, which must hold for every agent's action $a \in A$:

$$\left(\!u^{\mathsf{A}}(p,a) \le u^{\mathsf{A}}(p,a^{\star}) - \delta\right) \lor \left(\!u^{\mathsf{P}}(p,a) \ge u^{\mathsf{P}}(p,a^{\delta})\right)\!.$$
(2a)

Intuitively, Eq. (2a) requires that each action a is either *not* a δ -best response (first inequality) or no worse than a^{δ} for the principal (second inequality). This ensures that the objective function $u^{\mathsf{P}}(p, a^{\delta})$ captures the principal's utility in a δ -robust contract p. More formally, we establish the following lemma (recall that $\Psi(p) = \min_{a \in A^{\delta}(p)} u^{\mathsf{P}}(p, a)$ denotes the principal's utility in a δ -robust contract p).

Lemma 1. Every optimal solution $p \in \mathbb{R}^m_+$ to Problem (2) is an optimal δ -robust contract, i.e., $\Psi(p) = \Psi(p^*)$.

Proof. First, observe that p^* is a feasible solution to Problem (2). Indeed, if an action a is *not* a δ -best response to p^* , then by definition this means $u^A(p^*, a) \leq u^A(p^*, a^*) - \delta$; otherwise, it must be that $u^P(p^*, a) \geq u^P(p^*, a^{\delta})$ as a^{δ} is by definition the worst δ -best response in p^* . As a result, Eq. (2a) holds for p^* for every $a \in A$.

Furthermore, notice that, according to the definition of a^{δ} :

$$\Psi(p^{\star}) = \min_{a \in A_{\delta}(p^{\star})} u^{\mathsf{P}}(p^{\star}, a) = u^{\mathsf{P}}(p^{\star}, a^{\delta}).$$

In order to complete the proof, it is sufficient to show that $u^{\mathsf{P}}(p, a^{\delta}) \leq \Psi(p)$ for every feasible solution $p \in \mathbb{R}^m_+$ to Problem (2). Indeed, if the condition above holds, for any arbitrary optimal solution $p' \in \mathbb{R}^m_+$ to Problem (2), we have:

$$\Psi(p^{\star}) = u^{\mathsf{P}}(p^{\star}, a^{\delta}) \le u^{\mathsf{P}}(p', a^{\delta}) \le \Psi(p') \le \Psi(p^{\star}),$$

where $u^{\mathsf{P}}(p^{\star}, a^{\delta}) \leq u^{\mathsf{P}}(p', a^{\delta})$ since p^{\star} is a feasible solution to Problem (2) and p' is optimal, while $\Psi(p') \leq \Psi(p^{\star})$ since p^{\star} is an optimal δ -robust contract by definition. Then, it must be the case that $\Psi(p') = \Psi(p^{\star})$.

Now, to complete the proof, consider any feasible solution $p \in \mathbb{R}^m_+$ to Problem (2). We show that $u^{\mathsf{P}}(p, a^{\delta}) \leq \Psi(p)$. Pick an arbitrary best-response action $a' \in A(p)$ and consider any $a \in A^{\delta}(p)$. By definition,

$$u^{\mathsf{A}}(p,a) > u^{\mathsf{A}}(p,a') - \delta$$

= $\max_{a \in \mathcal{A}} u^{\mathsf{A}}(p,a) - \delta \ge u^{\mathsf{A}}(p,a^{\star}) - \delta.$

Hence, for Eq. (2a) to hold for action a, it must be that $u^{\mathsf{P}}(p, a) \ge u^{\mathsf{P}}(p, a^{\delta})$. Since the choice of a is arbitrary, this holds for every $a \in A^{\delta}(p)$. Consequently,

$$\Psi(p) = \min_{a \in A^{\delta}(p)} u^{\mathsf{P}}(p, a) \ge u^{\mathsf{P}}(p, a^{\delta}). \qquad \Box$$

The above lemma implies that we can effectively "guess" a^* and a^{δ} , fixing these actions in Problem (2) and solving the optimization problem to obtain p^* (or possibly a different, but still optimal, δ -robust contract). Since there are only $O(n^2)$ possible combinations of the values of a^* and a^{δ} , the approach is efficient as long as Problem (2) can be solved efficiently. A correct guess yields a contract p such that $\Psi(p) = \Psi(p^*)$, whereas $\Psi(p) \leq \Psi(p^*)$ for incorrect guesses. Thus, by comparing the Ψ values, we can identify a correct guess and a corresponding optimal δ -robust contract.

It remains to show how to efficiently solve Problem (2).

4.2. Solving Problem (2)

Problem (2) does *not* directly admit any efficient solution algorithm due to the non-convex constraint in Eq. (2a). To deal with this issue, we rewrite Eq. (2a) as follows:

$$\left(F_a \cdot p \le c_a + u^{\mathsf{A}}(p, a^{\star}) - \delta \right) \lor$$

$$\left(F_a \cdot p \le F_a \cdot r - u^{\mathsf{P}}(p, a^{\delta}) \right), \quad (3)$$

by expanding the utilities as $u^{A}(p, a) = F_{a} \cdot p - c_{a}$ and $u^{P}(p, a) = F_{a} \cdot (r - p)$ and rearranging the terms.

Hence, Eq. (2a) is satisfied for action $a \in A$ if and only if $F_a \cdot p$ is smaller than the maximum of the right-hand sides of the two inequalities in Eq. (3). The constraint effectively reduces to the second inequality for all $p \in \mathbb{R}^m_+$ such that

$$c_a + u^{\mathsf{A}}(p, a^{\star}) - \delta \le F_a \cdot r - u^{\mathsf{P}}(p, a^{\delta}), \qquad (4)$$

while for the other $p \in \mathbb{R}^m_+$, it reduces to the first inequality.

Consequently, we can partition the contract space based on the satisfiability of Eq. (4), considering all $a \in A$. Within each subspace in the partition, only one inequality in Eq. (3) is active for every $a \in A$. So, effectively, Eq. (3) reduces to a linear constraint for each action a, and the optimization problem to an LP. It then suffices to solve an LP for every subspace, each generating an optimal contract within its corresponding subspace. Among these contracts, the one providing the highest utility is an optimal solution to Problem (2).

Now, if the linear inequalities in Eq. (4) (one for each action $a \in A$) were n arbitrary inequalities, the above partition may consist of exponentially many subspaces, making the approach inefficient. Fortunately, the partition is much more

well-structured, since the hyperplanes corresponding to the inequalities are *parallel*, as the coefficients of p in Eq. (4) are invariant with respect to a. As a result, they partition the space into only O(n) subspaces.

Next, we show how to exploit the above observation to solve Problem (2), by solving O(n) suitable subproblems instead.

4.3. Formulating the Subproblems

Let us first rearrange Eq. (4) as follows:

$$u^{\mathsf{A}}(p, a^{\star}) + u^{\mathsf{P}}(p, a^{\delta}) - \delta \le \nu_a \coloneqq F_a \cdot r - c_a, \quad (5)$$

where ν_a is exactly the social welfare generated by action *a* (which is independent of the specific contract adopted).

Then, we can re-order agent's actions a_1, \ldots, a_n in such a way that $\nu_{a_1} \leq \nu_{a_2} \leq \cdots \leq \nu_{a_n}$. For simplicity, we write $\nu_j = \nu_{a_j}$, and we let $\nu_0 = -\infty$ and $\nu_{n+1} = +\infty$. Then, the following lemma is straightforward.

Lemma 2. For every contract $p \in \mathbb{R}^m_+$, if it holds that $\nu_{j-1} \leq u^{\mathsf{A}}(p, a^*) + u^{\mathsf{P}}(p, a^{\delta}) - \delta \leq \nu_j$, then:

- $F_a \cdot p \leq F_a \cdot r u^{\mathsf{P}}(p, a^{\delta}) \iff Eq.$ (3) holds for all actions $a \in \{a_{\ell} \mid \ell \leq j-1\}$; and
- $F_a \cdot p \leq c_a + u^{\mathsf{A}}(p, a^*) \delta \iff Eq. (3)$ holds for all actions $a \in \{a_\ell \mid j \leq \ell\}$.

In other words, the condition in the lemma defines a suitable subspace of contracts \mathcal{P}_j for each $j \in \{1, \ldots, n+1\}$. The following LP solves for a $p \in \mathbb{R}^m_+$ that is optimal within \mathcal{P}_j .

$$\max_{p \in \mathbb{R}^m_+} \quad u^{\mathsf{P}}(p, a^{\delta}) \tag{6}$$

subject to the following constraints:

$$\nu_{j-1} \le u^{\mathsf{A}}(p, a^{\star}) + u^{\mathsf{P}}(p, a^{\delta}) - \delta \le \nu_j \tag{6a}$$

$$F_a \cdot p \le F_a \cdot r - u^{\mathsf{P}}(p, a^{\delta}) \quad \forall a \in \{a_\ell \mid \ell \le j - 1\}$$
 (6b)

$$F_a \cdot p \le c_a + u^{\mathsf{A}}(p, a^*) - \delta \qquad \forall a \in \{a_\ell \mid j \le \ell\}.$$
 (6c)

Since $\bigcup_{j=1}^{n+1} \mathcal{P}_j = \mathbb{R}^m_+$, solving the LP in Problem (6) for all $j \in \{1, \dots, n+1\}$ and picking the best among the obtained solutions gives an optimal solution to Problem (2).

Remark 1. The left-hand side of Eq. (5) is roughly (up to a δ difference) the social welfare of an optimal contract. Thus, Lemma 2 can be interpreted as follows. For low-social-welfare actions a_1, \ldots, a_{j-1} , yielding a sufficiently high utility for the principal automatically provides a low utility for the agent. This fulfills the first inequality in Eq. (3) (and Eq. (2a)). Conversely, for high-social-welfare actions a_j, \ldots, a_n , a sufficiently low utility for the agent automatically gives a high utility for the principal. This fulfills the second inequality in Eq. (3) (and Eq. (2a)).

Algorithm 1 Compute an optimal δ -robust contract1: $p^* \leftarrow \text{null}, \ \psi^* \leftarrow -\infty$ 2: for all $(a^*, a^{\delta}) \in \mathcal{A} \times \mathcal{A}$ do3: for all $j = 1, \dots, n+1$ do4: Solve Problem (6) instantiated with (a^*, a^{δ}) s.t.
Eqs. (6a) to (6c), and let an optimal solution be p'5: if $\Psi(p') > \psi^*$ then6: $p^* \leftarrow p'$ 7: $\psi^* \leftarrow \Psi(p')$ 8: return p^*

4.4. Putting All Together

We summarize our results in this section into Algorithm 1 and the following main theorem.

Theorem 1. Algorithm 1 computes an optimal δ -robust contract in polynomial time.

Proof Sketch. The polynomial runtime of Algorithm 1 is obvious as it enumerates $O(n^3)$ value combinations and solves an LP for each of them. To see the correctness of the algorithm, note that the inner for-loop of Algorithm 1 effectively solves Problem (2), for the pair (a^*, a^{δ}) enumerated in the outer for-loop. Now, if a^* and a^{δ} happen to be the agent's exact and δ -best responses, respectively, under some optimal contract, then according to Lemma 1, the outer loop produces an optimal contract p', with $\Psi(p') \geq \Psi(p)$ for all $p \in \mathbb{R}^m_+$. By comparing the Ψ values, the algorithm identifies and outputs such an optimal contract. \Box

5. Learning Robust Contracts

The algorithm in Section 4 works under the assumption that the principal has full knowledge of all the payoff-relevant information about the agent and, thus, they can compute an optimal δ -robust contract. In this section, we address the case in which the principal has no such knowledge, and, thus, they have to learn an optimal δ -robust contract in an online fashion, by repeatedly interacting with the agent.

5.1. Learning Interaction

We consider an online learning framework similar to the one studied by Zhu et al. (2023), in which the features of agent's actions, *i.e.*, costs and probabilities over outcomes, depend on an agent's *type* that is sampled at each round from some (fixed) unknown probability distribution. Before formally defining the online learning framework, we introduce some notation that is needed to deal with hidden-action principalagent problems in which the agent can be of different types.

Hidden-Action Principal-Agent Problems With Types We let Θ be the finite set of possible agent's types, while $F_{\theta,a}$ and $c_{\theta,a}$ denote the probability distribution over outcomes and the cost, respectively, of action $a \in A$, when the agent has type $\theta \in \Theta$. The agent's type is drawn from an unknown probability distribution $\lambda \in \Delta_{\Theta}$, with $\lambda_{\theta} \in [0, 1]$ being the probability of type $\theta \in \Theta$. We extend all the notation introduced in Section 2 to account for agent's types. Specifically, given $\delta \in (0,1)$, we let $A^{\theta,\delta}(p) \subseteq \mathcal{A}$ be the set of δ -best responses to a contract $p \in \mathbb{R}^m_+$ for type $\theta \in \Theta$. Moreover, we denote by $a^{\theta,\delta}(p)$ the δ -best response that is played by the agent under a δ -robust contract, namely the worst one for the principal. Formally, $a^{\theta,\delta}(p) \in \arg\min_{a \in A^{\delta}(p)} F_{\theta,a} \cdot (r-p)$. Similarly, we use $A^{\theta}(p)$ and $a^{\theta}(p)$ for (exact) best responses. Finally, the principal's expected utility when committing to a contract $p \in \mathbb{R}^m_+$ against an agent of type $\theta \in \Theta$ playing an action $a \in \mathcal{A}$ is $u^{\mathsf{P}}(p, a, \theta) \coloneqq F_{\theta, a} \cdot (r - p).$

We study an online learning framework in which the principal and the agent repeatedly interact over T > 0 rounds. Each round involves a repetition of the same instance of hidden-action principal-agent problem, with only the agent's type changing from round to round. Specifically, at each round $t \in [T]$, the principal-agent interaction is as follows:

- 1. The agent's type $\theta^t \in \Theta$ is sampled from the distribution λ . Notice that θ^t is *not* observed by the principal.
- 2. The principal commits to a contract $p^t \in \mathcal{C} := [0, 1]^m$.
- 3. After observing the contract p^t , the agent plays a δ -best response $a^t \coloneqq a^{\theta^t,\delta}(p^t)$. Notice that the action a^t is *not* observed by the principal.
- 4. The principal observes the outcome $\omega^t \sim F_{a^t}$ that is realized as an effect of the agent's action a^t .

We remark that, in the interaction described above, the principal's contract design space is assumed to be limited to the hypercube $C := [0, 1]^m$. Restricting the contract space to a bounded set is standard in the literature on online learning in contract design (see, *e.g.*, (Ho et al., 2014; Zhu et al., 2023; Chen et al., 2024)) and it is motivated by the negative result proved in Theorem 1 in (Bacchiocchi et al., 2024). Therefore, we need to introduce a formal definition of the principal's expected utility when committing to an optimal δ -robust contract restricted to such a space. Formally:

$$\operatorname{OPT}(\mathcal{C}, \delta) := \max_{p \in \mathcal{C}} \sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p, a^{\theta, \delta}(p), \theta).$$

The goal of the principal is to minimize the (*cumulative*) *pseudo-regret*, or simply *regret*, which can be defined as:

$$R_T(\mathcal{C},\delta) := T \operatorname{OPT}(\mathcal{C},\delta) - \mathbb{E}\left[\sum_{t \in [T], \theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p^t, a^{\theta,\delta}(p^t), \theta)\right],$$

where the expectation is over the randomness of the algorithm and the environment that generates the feedback received by the learner at each round. Our goal is to design a *no-regret* algorithm for the principal, namely one achieving $R_T(\mathcal{C}, \delta) = o(T)$.

5.2. A No-Regret Algorithm

Before introducing our no-regret learning algorithm (Algorithm 2), we need to prove a key lemma about the "continuity" of the principal's expected utility over the hypercube. In particular, given a δ -robust contract $p \in \mathbb{R}^m_+$, we show that it is possible to build another contract $p' \in \mathbb{R}^m_+$ that is $(\delta + \epsilon)$ -robust and provides the principal with utility at most $2\sqrt{\epsilon}$ worse than the one of contract p. Formally:

Lemma 3. Given any $\delta, \epsilon \in (0, 1)$ and a contract $p \in \mathbb{R}^m_+$, the following holds for every $\theta \in \Theta$:

$$u^{\mathsf{P}}(p', a^{\theta, \delta+\epsilon}(p'), \theta) \ge u^{\mathsf{P}}(p, a^{\theta, \delta}(p), \theta) - 2\sqrt{\epsilon},$$

where $p' \coloneqq (1 - \sqrt{\epsilon})p + \sqrt{\epsilon}r$.

We observe that the idea of shifting payments towards the principal's reward vector, as in Lemma 3, was first adopted by Dutting et al. (2021) in non-robust settings, in order to deal with approximately incentive-compatible contracts.

Thanks to Lemma 3, it is possible to show that there always exists a contract $p \in \mathcal{B}_{\epsilon}$ that satisfies the following condition, where $\mathcal{B}_{\epsilon} \subseteq C$ is a uniform grid of the hypercube C, built with step size $\epsilon > 0$ (see also Algorithm 2).

$$\sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p, a^{\theta, \delta}(p), \theta) \ge \mathsf{OPT}(\mathcal{C}, \delta) - \mathcal{O}(\sqrt{\epsilon}).$$

Consequently, by suitably choosing $\epsilon > 0$ as a function of the time horizon T, and by instantiating a no-regret algorithm with set of arms \mathcal{B}_{ϵ} , we can upper bound the regret suffered by Algorithm 2 as follows.

Theorem 2. The regret suffered by Algorithm 2 can be upper bounded as follows:

$$R_T(\mathcal{C},\delta) \le \widetilde{\mathcal{O}}\left(T^{1-\frac{1}{2(m+1)}}\right)$$

We remark that the regret lower bound introduced by Zhu et al. (2023) also holds in our setting. This is because, when the parameter δ is arbitrarily small, if we consider the same instances used by Zhu et al. (2023) in their lower bound, the principal is still required to enumerate an exponential number of regions, thus suffering $\Omega(T^{1-1/(m+2)})$ regret. This confirms that an exponential dependence on the number of outcomes is unavoidable in the regret suffered by any algorithm. Furthermore, Theorem 2 shows regret guarantees

Algorithm 2 Regret minimizer for δ -robust contracts **Require:** T > 0

1: Set $\epsilon \leftarrow T^{-\frac{1}{m+1}}$

2: $\mathcal{B}_{\epsilon} \leftarrow \{ p \in [0,1]^m \mid p_{\omega} \in \{0,\epsilon,2\epsilon,...,1\} \forall \omega \in \Omega \}$

3: Run UCB1 with \mathcal{B}_{ϵ} as set of arms

similar to those obtained by Zhu et al. (2023) in the non-robust version of the problem. Indeed, they achieve an upper bound of the order of $\widetilde{\mathcal{O}}(\sqrt{m} \cdot T^{1-1/(2m+1)})$ on the regret.

Finally, we also show that, when the parameter $\delta \in (0, 1)$ is sufficiently small (with respect to the time horizon T > 0), Algorithm 2 achieves sublinear regret with respect to an optimal non-robust contract within C. Formally, we define the value of such a contract as follows:

$$OPT(\mathcal{C}) := \max_{p \in \mathcal{C}} \sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p, a^{\theta}(p), \theta).$$

Furthermore, when OPT(C) is chosen as baseline, the regret definition becomes the following:

$$R_T(\mathcal{C}) \coloneqq T \operatorname{OPT}(\mathcal{C}) - \mathbb{E}\left[\sum_{t \in [T], \theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p^t, a^{\theta, \delta}(p^t), \theta)\right],$$

where the expectation is over the randomness of the algorithm and the environment that generates the feedback received by the learner at each round. We observe that the baseline in the regret definition above coincides with the one introduced by Zhu et al. (2023). Then, we can prove that the following corollary of Theorem 2 holds.

Corollary 1. *The regret suffered by Algorithm 2 can be upper bounded as follows:*

$$R_T(\mathcal{C}) \leq \widetilde{\mathcal{O}}\left(T^{1-\frac{1}{2(m+1)}}\right) + 2\sqrt{\delta}T.$$

We remark that, if we set $\delta = 1/T^{\alpha}$ with $\alpha > 0$ in Corollary 1, then the regret with respect to an optimal non-robust contract within the hypercube C is sublinear.

Remark 2. Our algorithm can be extended to deal with settings in which the agent can play any δ -best response within the set $A^{\theta^t,\delta}(p^t)$. In such a setting, the principal's utility is *not* fully stochastic. However, our Algorithm 2 can be easily extended by instantiating an adversarial no-regret algorithm instead of UCB1 (Auer et al., 2002).

Comparison With (Zhu et al., 2023) We observe that our algorithm provides some advantages compared to the state-of-the-art algorithm proposed by Zhu et al. (2023). First, our approach employs as set of contracts a simple discretization of the hypercube, while the approach by Zhu et al. (2023) requires determining the minimum set of contracts $\mathcal{V}_{\epsilon}(\mathcal{C})$ —for some suitable $\epsilon > 0$ —such that, for every $p \in \mathcal{C}$, there

exists a $p' \in \mathcal{V}_{\epsilon}(\mathcal{C})$ that satisfies $(r-p) \cdot p' \ge \cos(\epsilon)$. However, the need of designing such a set of contracts makes the approach by Zhu et al. (2023) challenging to be employed in practice compared to ours. Second, our algorithm does *not* require *apriori* knowledge of the principal's reward, which is instead required by the algorithm of Zhu et al. (2023).

Acknowledgments

This work was supported by the Italian MIUR PRIN 2022 Project "Targeted Learning Dynamics: Computing Efficient and Fair Equilibria through No-Regret Algorithms", by the 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), and by the EU Horizon project ELIAS (European Lighthouse of AI for Sustainability, No. 101120237).

Impact Statement

This paper presents theoretical results and has the goal of advancing the field of Machine Learning. There are no potential societal consequences of our work which we feel must be highlighted here.

References

- Alon, T., Duetting, P., Li, Y., and Talgam-Cohen, I. Bayesian analysis of linear contracts. In *Proceedings of the 24th* ACM Conference on Economics and Computation, EC '23, pp. 66, 2023.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47:235–256, 2002.
- Babaioff, M., Feldman, M., and Nisan, N. Combinatorial agency. In Proceedings of the 7th ACM Conference on Electronic Commerce, pp. 18–28, 2006.
- Babaioff, M., Feldman, M., and Nisan, N. Free-riding and free-labor in combinatorial agency. In *International Symposium on Algorithmic Game Theory*, pp. 109–121. Springer, 2009.
- Bacchiocchi, F., Castiglioni, M., Marchesi, A., and Gatti, N. Learning optimal contracts: How to exploit small action spaces. In *The Twelfth International Conference* on Learning Representations, 2024.
- Bastani, H., Bayati, M., Braverman, M., Gummadi, R., and Johari, R. Analysis of medicare pay-for-performance contracts. *Available at SSRN 2839143*, 2016.
- Bastani, H., Goh, J., and Bayati, M. Evidence of upcoding

in pay-for-performance programs. *Management Science*, 65(3):1042–1060, 2019.

- Bernasconi, M., Castiglioni, M., and Marchesi, A. Regretminimizing contracts: Agency under uncertainty, 2024. URL https://arxiv.org/abs/2402.13156.
- Cai, Y., Daskalakis, C., and Papadimitriou, C. Optimum statistical estimation with strategic data sources. In *Conference on Learning Theory*, pp. 280–296. PMLR, 2015.
- Carroll, G. Robustness and linear contracts. *American Economic Review*, 105(2):536–563, 2015.
- Castiglioni, M., Marchesi, A., and Gatti, N. Bayesian agency: Linear versus tractable contracts. *Artificial Intelligence*, 307:103684, 2022a.
- Castiglioni, M., Marchesi, A., and Gatti, N. Designing menus of contracts efficiently: The power of randomization. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp. 705–735, 2022b.
- Castiglioni, M., Marchesi, A., and Gatti, N. Multi-agent contract design: How to commission multiple agents with individual outcomes. In *Proceedings of the 24th* ACM Conference on Economics and Computation, pp. 412–448, 2023.
- Chen, Y., Chen, Z., Deng, X., and Huang, Z. Are bounded contracts learnable and approximately optimal? *arXiv preprint arXiv:2402.14486*, 2024.
- Cong, L. W. and He, Z. Blockchain disruption and smart contracts. *The Review of Financial Studies*, 32(5):1754– 1797, 2019.
- Duetting, P., Ezra, T., Feldman, M., and Kesselheim, T. Multi-agent combinatorial contracts, 2024.
- Dütting, P., Roughgarden, T., and Talgam-Cohen, I. Simple versus optimal contracts. In *Proceedings of the 2019* ACM Conference on Economics and Computation, pp. 369–387, 2019.
- Dutting, P., Roughgarden, T., and Talgam-Cohen, I. The complexity of contracts. *SIAM Journal on Computing*, 50(1):211–254, 2021.
- Dütting, P., Feldman, M., Talgam-Cohen, I., et al. Algorithmic contract theory: A survey. *Foundations and Trends*® *in Theoretical Computer Science*, 16(3-4):211–412, 2024.
- Gan, J., Han, M., Wu, J., and Xu, H. Robust stackelberg equilibria. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, pp. 735, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701047. doi: 10. 1145/3580507.3597680. URL https://doi.org/ 10.1145/3580507.3597680.

- Guruganesh, G., Schneider, J., and Wang, J. R. Contracts under moral hazard and adverse selection. In *Proceed*ings of the 22nd ACM Conference on Economics and Computation, pp. 563–582, 2021.
- Hadfield-Menell, D. and Hadfield, G. K. Incomplete contracting and ai alignment. In *Proceedings of the 2019* AAAI/ACM Conference on AI, Ethics, and Society, pp. 417–422, 2019.
- Ho, C.-J., Slivkins, A., and Vaughan, J. W. Adaptive contract design for crowdsourcing markets: Bandit algorithms for repeated principal-agent problems. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pp. 359–376, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450325653. doi: 10.1145/2600057.2602880.
- Holmstrom, B. and Milgrom, P. Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design. *The Journal of Law, Economics, and Organization*, 7(special_issue):24–52, 1991.
- Kaynar, N. and Siddiq, A. Estimating effects of incentive contracts in online labor platforms. *Management Science*, 69(4):2106–2126, 2023.
- Lattimore, T. and Szepesvári, C. Bandit algorithms. Cambridge University Press, 2020.
- Saig, E., Einav, O., and Talgam-Cohen, I. Incentivizing quality text generation via statistical contracts. *arXiv* preprint arXiv:2406.11118, 2024.
- Yang, K. and Zhang, H. Computational aspects of bayesian persuasion under approximate best response. *arXiv* preprint arXiv:2402.07426, 2024.
- Zhu, B., Bates, S., Yang, Z., Wang, Y., Jiao, J., and Jordan, M. I. The sample complexity of online contract design. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, pp. 1188, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701047. doi: 10. 1145/3580507.3597673. URL https://doi.org/ 10.1145/3580507.3597673.

Appendix

Proposition 1 (Upper and lower bounds). Given an instance of hidden-action principal-agent problem:

1. For every $\delta \in (0, 1)$, it holds:

$$OPT(\delta) \ge OPT_{LB}(\delta) \coloneqq OPT - 2\sqrt{\delta} + \delta.$$

2. For every $\delta \in (0, 1)$, it holds:

$$OPT(\delta) \le OPT_{UB}(\delta) \coloneqq \max\{0, SW - \delta\}.$$

Proof. We prove the two points separately.

1. Let $p' = (1 - \sqrt{\delta})p + \sqrt{\delta}r$ for any contract $p \in \mathbb{R}^m_+$. We start by observing that the following inequality holds:

$$\sum_{\omega \in \Omega} F_{a(p),\omega} p_{\omega} - c_{a(p)} \ge \sum_{\omega \in \Omega} F_{a^{\delta}(p'),\omega} p_{\omega} - c_{a^{\delta}(p')},$$

because $a(p) \in \mathcal{A}(p)$. Furthermore, we have:

$$\sum_{\omega \in \Omega} F_{a^{\delta}(p'),\omega} p'_{\omega} - c_{a^{\delta}(p')} > \max_{a' \in \mathcal{A}} \sum_{\omega \in \Omega} F_{a',\omega} p_{\omega} - c_{a'} - \delta \ge \sum_{\omega \in \Omega} F_{a(p),\omega} p'_{\omega} - c_{a(p)} - \delta,$$

thanks to the definition of $\mathcal{A}^{\delta}(p')$. Then, by employing the two above inequalities and the definition of p', we have:

$$\begin{split} \delta &\geq \sum_{\omega \in \Omega} F_{a(p),\omega} p'_{\omega} - c_{a(p)} - \left(\sum_{\omega \in \Omega} F_{a^{\delta}(p'),\omega} p'_{\omega} - c_{a^{\delta}(p')} \right) \\ &= \sum_{\omega \in \Omega} F_{a(p),\omega} p_{\omega} - c_{a(p)} - \left(\sum_{\omega \in \Omega} F_{a^{\delta}(p'),\omega} p_{\omega} - c_{a^{\delta}(p')} \right) + \sqrt{\delta} \left(\sum_{\omega \in \Omega} \left(F_{a(p),\omega} - F_{a^{\delta}(p),\omega} \right) (r_{\omega} - p_{\omega}) \right) \\ &\geq \sqrt{\delta} \left(\sum_{\omega \in \Omega} \left(F_{a(p),\omega} - F_{a^{\delta}(p),\omega} \right) (r_{\omega} - p_{\omega}) \right). \end{split}$$

Thus, by rearranging the latter inequality we get:

$$\sqrt{\delta} \ge \left(\sum_{\omega \in \Omega} \left(F_{a(p),\omega} - F_{a^{\delta}(p),\omega} \right) \left(r_{\omega} - p_{\omega} \right) \right).$$
(7)

Finally, we can show that:

$$u^{\mathsf{P}}(p, a(p)) - u^{\mathsf{P}}(p', a^{\delta}(p')) = \sum_{\omega \in \Omega} F_{a(p),\omega}(r_{\omega} - p_{\omega}) - \left(\sum_{\omega \in \Omega} F_{a^{\delta}(p'),\omega}(r_{\omega} - p'_{\omega})\right)$$
$$= \sum_{\omega \in \Omega} F_{a(p),\omega}(r_{\omega} - p_{\omega}) - \left(1 - \sqrt{\delta}\right) \left(\sum_{\omega \in \Omega} F_{a^{\delta}(p'),\omega}(r_{\omega} - p_{\omega})\right)$$
$$\leq \left(1 - \sqrt{\delta}\right) \left(\sum_{\omega \in \Omega} (F_{a(p),\omega} - F_{a^{\delta}(p'),\omega})(r_{\omega} - p_{\omega})\right) + \sqrt{\delta}$$
$$\leq 2\sqrt{\delta} - \delta,$$

where the second equality holds thanks to the definition of p', the first inequality because $u^{\mathsf{P}}(p, a(p)) \leq 1$ for each $p \in \mathbb{R}^m_+$ and the second inequality because of Eq. (7). Finally, let $p \in \mathbb{R}$ be an optimal (non-robust) contract, then we have:

$$OPT - 2\sqrt{\delta} + \delta \le u^{\mathsf{P}}(p', a^{\delta}(p')) \le OPT(\delta),$$

concluding the first part of the proof.

- 2. We split the proof into two parts.
 - (a) if $OPT(\delta) = 0$. Then, we trivially have :

$$0 = OPT(\delta) \le \max\left(0, \max_{a \in \mathcal{A}} \sum_{\omega \in \Omega} F_{a,\omega} r_{\omega} - c_a - \delta\right) = \max\left(0, SW - \delta\right)$$

(b) if OPT(δ) > 0. We define p^{*} ∈ ℝ^m₊ as an optimal δ-robust contract. Then, we notice that a₁ ∉ A^δ(p^{*}), where a₁ is the opt-out action. Indeed, if a₁ ∈ A^δ(p^{*}), then the agent may select the action a₁ as a δ-robust best-response, which provides zero or negative utility to the principal and contradicts the fact that OPT(δ) > 0. Therefore, in an optimal robust contract p^{*}, we must have:

$$\sum_{\omega \in \Omega} F_{a',\omega} p_{\omega}^{\star} - c_{a'} \ge \sum_{\omega \in \Omega} F_{a_1,\omega} p_{\omega}^{\star} - c_{a_1} + \delta \ge \sum_{\omega \in \Omega} F_{a_1,\omega} p_{\omega}^{\star} + \delta \ge \delta,$$

for some $a' \in \mathcal{A}^{\delta}(p^{\star})$. Thus, we have:

$$\begin{aligned} \operatorname{OPT}(\delta) &= \sum_{\omega \in \Omega} F_{a^{\delta}(p),\omega}(r_{\omega} - p_{\omega}^{\star}) \\ &\leq \sum_{\omega \in \Omega} F_{a',\omega}(r_{\omega} - p_{\omega}^{\star}) \\ &\leq \sum_{\omega \in \Omega} F_{a',\omega}r_{\omega} - c_{a'} - \delta \\ &\leq \max_{a \in \mathcal{A}} \sum_{\omega \in \Omega} F_{a,\omega}r_{\omega} - c_{a} - \delta \\ &\leq \max\left(0, \max_{a \in \mathcal{A}} \sum_{\omega \in \Omega} F_{a,\omega}r_{\omega} - c_{a} - \delta\right) = \max\left(0, \operatorname{SW} - \delta\right). \end{aligned}$$

The first inequality above follows from the fact that $a' \in \mathcal{A}^{\delta}(p^{\star})$, while the second inequality holds due to the previous observation.

The two above points conclude the proof.

Proposition 2. Given $\delta \in (0, 1)$, for every integer $n > n(\delta)$, there is an instance of hidden-action principal-agent problem (parametrized by δ) with 2n + 1 actions, where:

$$OPT(\delta) - OPT_{LB}(\delta) \le \mathcal{O}\left(\frac{1}{n}\right)$$

Furthermore, for any $\delta \in (0, 1)$ *, there exists an instance of hidden-action principal-agent problem in which it holds:*

$$OPT(\delta) = OPT_{UB}(\delta).$$

Proof. We prove the two points separately.

1. We consider an instance parametrized by $\delta > 0$, with $|\Omega| = 2$ and $|\mathcal{A}| = 2n + 1$ for some $n \in \mathbb{N}$, to be defined below. We let:

$$\kappa = \min\left\{i \in \mathbb{N}_{>0} \,|\, \sqrt{\delta} < \frac{i-1}{i}\right\}.$$

Furthermore, we introduce the following decreasing sequence:

$$\gamma_i = \begin{cases} \frac{i}{i-1} & i = \kappa, \dots, n\\ \frac{2n+1-i}{2n+2-i} & i = n+2, \dots, 2n, \end{cases}$$

with $\gamma_{n+1} = 1$ and $\gamma_{2n+1} = 0$, where $n \in \mathbb{N}$ is such that $n > \kappa$. The distributions over the set of outcomes for the different actions and their corresponding costs are given by:

$$\begin{cases} F_{a_i,\omega_1} = 0 & c_{a_i} = 0, \ i = 1, \dots, \kappa - 1, \\ F_{a_i,\omega_1} = 1 - \gamma_i \sqrt{\delta} & c_{a_i} = 0, \ i = \kappa, \dots, n, \\ F_{a_{2n+1,\omega_1}} = 1 & c_{a_{2n+1}} = 0. \end{cases}$$

and the principal's reward is r = (1, 0). Notice that the distribution over outcomes of the different actions are always well defined because of the definition of $\kappa > 0$. Furthermore, since all the agent's actions a_i with i < k are coincident, we assume for the sake of presentation that the agent always selects a_1 . It is easy to verify that the value of an optimal (non-robust) contract is OPT = 1 since the agent breaks ties optimistically and $c_{2n+1} = 0$.

With a similar argument to the one proposed by (Dütting et al., 2024) in Proposition 3.9, an optimal δ -robust contract is such that $p^* = (0, \alpha)$ for some $\alpha \in \mathbb{R}_+$. Consequently, in the rest of the proof, we focus on determining the maximum utility achievable in a linear, δ -robust contract.

For every $a_i \in A$, with $\kappa < i \le 2n + 1$, if $\alpha \in \mathbb{R}_+$ is such that $a^{\delta}(\alpha r) = a_i$, then the two following conditions hold.

All the actions a_j ∉ A^δ(αr) for each j < i. This is because u^P(α, a_j) ≤ u^P(α, a_i) for each α ∈ ℝ₊ since {γ_i}_{i∈[2n+1]} is a decreasing sequence. Thus, the latter observation implies that:

$$u^{\mathsf{A}}(\alpha, a_{j}) \leq u^{\mathsf{A}}(\alpha, a_{i-1}) = \alpha R_{a_{i-1}} - c_{a_{i-1}}$$
$$\leq \max_{i \in [2n+1]} u^{\mathsf{A}}(\alpha, a_{i}) - \delta$$
$$= u^{\mathsf{A}}(\alpha, a_{2n+1}) - \delta$$
$$= \alpha R_{a_{2n+1}} - c_{a_{2n+1}} - \delta$$
$$= \alpha - \delta.$$

Therefore, we have that:

$$\alpha R_{i-1} \le \alpha - \delta_i$$

and, thus, $\alpha \geq \sqrt{\delta}/\gamma_{i-1}$.

• The action $a_i \in \mathcal{A}^{\delta}(\alpha r)$. Thus,

$$\alpha R_{a_i} - c_{a_i} > \max_{i \in [2n+1]} u^{\mathsf{A}}(\alpha, a_i) - \delta = \alpha - \delta,$$

which implies that $\alpha < \sqrt{\delta}/\gamma_i$.

The two observation above shows that a_i , with $i > \kappa$, is a δ -best response for all the values of α such that:

$$\sqrt{\delta}/\gamma_{i-1} \le \alpha < \sqrt{\delta}/\gamma_i.$$

Thus, the smallest value of $\alpha \in \mathbb{R}_+$ such that $a^{\delta}(\alpha r) = a_i$ is $\alpha_i \coloneqq \sqrt{\delta}/\gamma_{i-1}$, when $i > \kappa$. With the same argument above, it is possible to show that $a_1 = a^{\delta}(\alpha r)$ for all $\alpha \in [\alpha_1, \alpha_{\kappa})$ and $a_{\kappa} = a^{\delta}(\alpha r)$ for all $\alpha \in [\alpha_{\kappa}, \alpha_{\kappa+1})$, with $\alpha_1 = 0$ and $\alpha_{\kappa} = \delta$.

We also observe that the principal's utility in each α_i , with $i \in \{\kappa + 1, \ldots, 2n + 1\}$, is such that:

$$u^{\mathsf{P}}(\alpha_{i}, a_{i}) = (1 - \alpha_{i})R_{a_{i}} = \left(1 - \frac{\sqrt{\delta}}{\gamma_{i-1}}\right)\left(1 - \gamma_{i}\sqrt{\delta}\right)$$
$$= \left(1 - \frac{\sqrt{\delta}}{\gamma_{i-1}}\right)\left(1 - \gamma_{i}\sqrt{\delta}\right)$$
$$= \left(1 - \left(\gamma_{i} + \frac{1}{\gamma_{i-1}}\right)\right)\sqrt{\delta} + \frac{\gamma_{i}}{\gamma_{i-1}}\delta.$$

Therefore, using the latter formula, we have that the following holds.

• If $i = \kappa + 1, \ldots, n$, we have:

$$u(\alpha_i) \le 1 - 2\sqrt{\delta} + \frac{\gamma_i}{\gamma_{i-1}}\delta \le 1 - 2\sqrt{\delta} + \delta_i$$

since $\{\gamma_i\}_{i>\kappa}$ is a decreasing sequence.

• If i = n + 1, n + 2, we have:

$$u(\alpha_i) \le 1 - \left(1 + \frac{n-1}{n}\right)\sqrt{\delta} + \frac{n-1}{n}\delta \le 1 - 2\sqrt{\delta} + \delta + \frac{\sqrt{\delta}}{n}.$$

• If i = n + 3, ..., 2n + 1, we have:

$$u(\alpha_i) \le 1 - 2\sqrt{\delta} + \frac{\gamma_i}{\gamma_{i-1}}\delta \le 1 - 2\sqrt{\delta} + \delta,$$

since $\{\gamma_i\}_{i>\kappa}$ is a decreasing sequence.

We also observe that the principal's utility in α_{κ} is such that:

$$u(\alpha_{\kappa}) = (1 - \delta)(1 - \gamma_{\kappa}\sqrt{\delta})$$

since $\gamma_k > 1$ and $\delta > 0$. Thanks to the definition of κ , with a simple calculation, it possible to show that:

$$\kappa = \begin{cases} \left\lceil \frac{1}{1-\sqrt{\delta}} \right\rceil & \text{ if } \frac{1}{1-\sqrt{\delta}} \not \in \mathbb{N} \\ \frac{1}{1-\sqrt{\delta}} + 1 & \text{ if } \frac{1}{1-\sqrt{\delta}} \in \mathbb{N}. \end{cases}$$

Thus, when $1/(1-\sqrt{\delta}) \notin \mathbb{N}$, we have:

$$u(\alpha_{\kappa}) = (1-\delta)(1-\gamma_{\kappa}\sqrt{\delta})$$

= $(1-\delta)\left(1-\frac{\left\lceil\frac{1}{1-\sqrt{\delta}}\right\rceil}{\left\lceil\frac{1}{1-\sqrt{\delta}}\right\rceil-1}\sqrt{\delta}\right)$
 $\leq (1-\delta)\left(1-2\sqrt{\delta}+\delta\right).$

Similarly, it is possible to show that the same relation holds even when $1/(1 - \sqrt{\delta}) \in \mathbb{N}$. Finally, by putting all together, we have that:

$$\begin{aligned} \mathsf{OPT}(\delta) - \mathsf{OPT}_{\mathsf{LB}}(\delta) &\leq \max_{i \in [2n+1]} u(\alpha_i) - (1 - 2\sqrt{\delta} + \delta) \\ &= u(\alpha_{n+1}) - (1 - 2\sqrt{\delta} + \delta) \leq \frac{\sqrt{\delta}}{n}, \end{aligned}$$

concluding the first part of the proof.

2. We consider an instance with $|\Omega| = |\mathcal{A}| = 2$ and r = (0, 1). The distributions over the set of outcomes for the different actions and their corresponding costs are given by:

$$\begin{cases} F_{a_1} = (1,0) & c_{a_1} = 0 \\ F_{a_2} = (0,1) & c_{a_2} = 0 \end{cases}$$

It is easy to verify that $SW = R_{a_2} - c_{a_2} = 1$. Furthermore, with a similar argument to the one proposed by (Dütting et al., 2024) in Proposition 3.9, an optimal δ -robust contract is such that $p^* = (0, \alpha)$ for some $\alpha \in \mathbb{R}_+$.

We show that $OPT(\delta) = 1 - \delta$. Let $a_1 \in \mathcal{A}$ be the opt-out action. We observe that $a_1 \notin A^{\delta}(\alpha r)$ for all the $\alpha \in \mathbb{R}_+$ that satisfy:

$$\alpha R_{a_2} - c_{a_2} = \alpha \ge \alpha R_{a_1} - c_{a_1} + \delta = \delta.$$

Thus, $a_1 \notin A^{\delta}(\alpha r)$ for all $\alpha \geq \delta$. Therefore, the smallest $\alpha \in \mathbb{R}_+$ such that the agent selects the action a_2 coincides with $\delta > 0$. Thus, the largest utility the principal can achieve satisfies $OPT(\delta) = 1 - \delta = SW - \delta$.

The two points above conclude the proof.

Proposition 3. The function $(0,1) \ni \delta \mapsto OPT(\delta)$ is continuous and non-increasing in δ . Moreover, $\lim_{\delta \to 0^+} OPT(\delta) = OPT$ and $\lim_{\delta \to 1^-} OPT(\delta) = 0$.

Proof. We prove the claims separately.

We first show that OPT(δ) is a continuous function. Let p^{*} ∈ ℝ^m₊ be a δ-robust and p' = (1 − √ε)p^{*} + √εr. Then, by Lemma 3, we have:

$$OPT(\delta + \epsilon) \ge u^{\mathsf{P}}(p', a^{\delta + \epsilon}(p')) \ge u^{\mathsf{P}}(p^{\star}, a^{\delta}(p^{\star})) - 2\sqrt{\epsilon} + \epsilon = OPT(\delta) - 2\sqrt{\epsilon} + \epsilon.$$

Thus, rewriting the latter inequality and taking the limit, we have that:

$$0 = \lim_{\epsilon \to 0} (2\sqrt{\epsilon} - \epsilon) \ge \lim_{\epsilon \to 0} \operatorname{OPT}(\delta) - \operatorname{OPT}(\delta + \epsilon).$$

Let $p^* \in \mathbb{R}^m_+$ be a δ' -robust and $p' = (1 - \sqrt{\epsilon})p^* + \sqrt{\epsilon}r$ with $\delta' = \delta - \epsilon$. Then, by Lemma 3, we have:

$$OPT(\delta) = OPT(\delta' + \epsilon) \ge u^{P}(p', a^{\delta' + \epsilon}(p')) \ge u^{P}(p^{\star}, a^{\delta'}(p^{\star})) - 2\sqrt{\epsilon} + \epsilon = OPT(\delta - \epsilon) - 2\sqrt{\epsilon} + \epsilon.$$

Thus, by taking the limit, we have that:

$$0 = \lim_{\epsilon \to 0} (2\sqrt{\epsilon} - \epsilon) \ge \lim_{\epsilon \to 0} \operatorname{OPT}(\delta - \epsilon) - \operatorname{OPT}(\delta).$$

Since the latter considerations hold for every $\delta \in (0, 1)$, we can easily show that $OPT(\delta)$ is continuous in $\delta \in (0, 1)$.

• We now prove that $OPT(\delta)$ is non-increasing. Let $\delta, \delta' \in (0, 1)$ such that $\delta' < \delta$. Furthermore, let $p^* \in \mathbb{R}^m_+$ be an optimal δ -robust contract. Therefore, we have:

$$OPT(\delta) = u(a^{\delta}(p^{\star}), p^{\star}) \le u(a^{\delta'}(p^{\star}), p^{\star}) \le OPT(\delta'),$$

since, for each $p \in \mathbb{R}^m_+$, we have $A^{\delta'}(p) \subseteq A^{\delta}(p)$.

• By Proposition 1, we have:

$$\lim_{\delta \to 0^+} \operatorname{OPT}(\delta) \ge \lim_{\delta \to 0^+} \operatorname{OPT} - 2\sqrt{\delta} + \delta = \operatorname{OPT},$$

and,

$$\lim_{\delta \to 1^{-}} \operatorname{OPT}(\delta) \le \lim_{\delta \to 1^{-}} 1 - \delta = 0.$$

The two points above conclude the proof.

Lemma 3. Given any $\delta, \epsilon \in (0, 1)$ and a contract $p \in \mathbb{R}^m_+$, the following holds for every $\theta \in \Theta$:

$$u^{\mathsf{P}}(p', a^{\theta, \delta+\epsilon}(p'), \theta) \ge u^{\mathsf{P}}(p, a^{\theta, \delta}(p), \theta) - 2\sqrt{\epsilon},$$

where $p' := (1 - \sqrt{\epsilon})p + \sqrt{\epsilon}r$.

Proof. For the sake of the presentation, we avoid the dependence on the type $\theta \in \Theta$. We split the proof in two cases:

1. If $a^{\delta+\epsilon}(p') \notin A^{\delta}(p)$, then $a^{\delta+\epsilon}(p')$ is not a δ -best-response in p. Therefore, we have that:

$$\sum_{\omega \in \Omega} F_{a^{\delta}(p),\omega} p_{\omega} - c_{a^{\delta}(p)} \ge \sum_{\omega \in \Omega} F_{a^{\delta+\epsilon}(p'),\omega} p_{\omega} - c_{a^{\delta+\epsilon}(p')} + \delta,$$

and,

$$\sum_{\omega \in \Omega} F_{a^{\delta + \epsilon}(p'),\omega} p'_{\omega} - c_{a^{\delta + \epsilon}(p')} \ge \sum_{\omega \in \Omega} F_{a^{\delta}(p),\omega} p'_{\omega} - c_{a^{\delta}(p)} - \delta - \epsilon$$

Then, by taking the summation of the above quantities and employing the definition of p' we get:

$$\sqrt{\epsilon} \ge \sum_{\omega \in \Omega} (F_{a^{\delta}(p),\omega} - F_{a^{\delta + \epsilon}(p'),\omega})(r_{\omega} - p_{\omega})$$

Therefore, we can prove that the following holds:

$$u^{\mathsf{P}}(p, a^{\delta}(p)) - u^{\mathsf{P}}(p', a^{\delta+\epsilon}(p')) = \sum_{\omega \in \Omega} F_{a^{\delta}(p),\omega}(r_{\omega} - p_{\omega}) - \left(\sum_{\omega \in \Omega} F_{a^{\delta+\epsilon}(p'),\omega}(r_{\omega} - p'_{\omega})\right)$$
$$= \sum_{\omega \in \Omega} F_{a^{\delta}(p),\omega}(r_{\omega} - p_{\omega}) - \left(1 - \sqrt{\epsilon}\right) \left(\sum_{\omega \in \Omega} F_{a^{\delta+\epsilon}(p'),\omega}(r_{\omega} - p_{\omega})\right)$$
$$\leq \left(1 - \sqrt{\epsilon}\right) \left(\sum_{\omega \in \Omega} (F_{a^{\delta}(p),\omega} - F_{a^{\delta+\epsilon}(p'),\omega})(r_{\omega} - p_{\omega})\right) + \sqrt{\epsilon}$$
$$\leq 2\sqrt{\epsilon} - \epsilon.$$

2. if $a^{\delta+\epsilon}(p') \in A^{\delta}(p)$, then one of the following hold.

(a) if $a^{\delta}(p) \in A^{\delta+\epsilon}(p')$, then either $a^{\delta}(p) \equiv a^{\delta+\epsilon}(p')$ or $a^{\delta}(p)$ provides the same principal's utility as $a^{\delta+\epsilon}(p')$ in p'. Indeed, suppose by contradiction that the following holds:

$$\sum_{\omega \in \Omega} F_{a^{\delta + \epsilon}(p'), \omega}(r_{\omega} - p'_{\omega}) < \sum_{\omega \in \Omega} F_{a^{\delta}(p), \omega}(r_{\omega} - p'_{\omega}).$$

Then, we have:

$$\sum_{\omega \in \Omega} F_{a^{\delta+\epsilon}(p'),\omega}(r_{\omega} - p'_{\omega}) = (1 - \sqrt{\epsilon}) \sum_{\omega \in \Omega} F_{a^{\delta+\epsilon}(p'),\omega}(r_{\omega} - p_{\omega})$$
$$\geq (1 - \sqrt{\epsilon}) \sum_{\omega \in \Omega} F_{a^{\delta}(p),\omega}(r_{\omega} - p_{\omega})$$
$$= \sum_{\omega \in \Omega} F_{a^{\delta}(p),\omega}(r_{\omega} - p'_{\omega}).$$

where the inequality is a consequence of the fact that both $a^{\delta}(p)$ and $a^{\delta+\epsilon}(p')$ belongs to $A^{\delta}(p)$. This, reaches a contradiction with the initial assumption. Thus, we have that:

$$u^{\mathsf{P}}(p', a^{\delta + \epsilon}(p')) = \sum_{\omega \in \Omega} F_{a^{\delta + \epsilon}(p'), \omega}(r_{\omega} - p'_{\omega})$$

$$= \sum_{\omega \in \Omega} F_{a^{\delta}(p), \omega}(r_{\omega} - p'_{\omega})$$

$$= (1 - \sqrt{\epsilon}) \sum_{\omega \in \Omega} F_{a^{\delta}(p), \omega}(r_{\omega} - p_{\omega})$$

$$\geq \sum_{\omega \in \Omega} F_{a^{\delta}(p), \omega}(r_{\omega} - p_{\omega}) - \sqrt{\epsilon}$$

$$\geq u^{\mathsf{P}}(p, a^{\delta}(p)) - \sqrt{\epsilon}.$$

(b) if $a^{\delta}(p) \not\in A^{\delta+\epsilon}(p')$, then the following holds:

$$\sum_{\omega \in \Omega} F_{a^{\delta}(p'),\omega} p_{\omega} - c_{a^{\delta}(p)} \ge \sum_{\omega \in \Omega} F_{a^{\delta} + \epsilon(p'),\omega} p_{\omega} - c_{a^{\delta} + \epsilon(p')} - \delta,$$

and,

$$\sum_{\omega \in \Omega} F_{a^{\delta + \epsilon}(p'), \omega} p'_{\omega} - c_{a^{\delta + \epsilon}(p')} \ge \sum_{\omega \in \Omega} F_{a^{\delta}(p), \omega} p'_{\omega} - c_{a^{\delta}(p)} + \delta + \epsilon.$$

Then, by taking the summation of the above quantities and employing the definition of p' we get:

$$0 \ge \sum_{\omega \in \Omega} (F_{a^{\delta}(p),\omega} - F_{a^{\delta+\epsilon}(p'),\omega})(r_{\omega} - p_{\omega}).$$

Then, using the same argument employed at point (a) we can show that:

$$u^{\mathsf{P}}(p, a^{\delta}(p)) - u^{\mathsf{P}}(p', a^{\delta + \epsilon}(p')) \le \sqrt{\epsilon},$$

concluding the proof.

Theorem 2. The regret suffered by Algorithm 2 can be upper bounded as follows:

$$R_T(\mathcal{C},\delta) \leq \widetilde{\mathcal{O}}\left(T^{1-\frac{1}{2(m+1)}}\right).$$

Proof. In the following, we let:

- 1. p^* be such that $\sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p^*, a^{\theta, \delta}(p^*), \theta) = \mathsf{OPT}(\mathcal{C}, \delta).$
- 2. $p' \coloneqq (1 \sqrt{2\epsilon})p^{\star} + \sqrt{2\epsilon}r$. Notice that $p' \in [0, 1]^m$ since both $p, r \in \mathcal{C}$ and \mathcal{C} is convex.
- 3. $\overline{p} \in \mathcal{B}_{\epsilon}$ be such that $\|\overline{p} p'\|_{\infty} \leq \epsilon$. Notice that there always exists at least once contract \overline{p} satisfying the latter condition because of the definition of \mathcal{B}_{ϵ} and the fact that $p' \in [0, 1]^m$.

We start by observing that, for each $\theta \in \Theta$, it holds $a^{\theta,\delta}(\overline{p}) \in A^{\theta,\delta+2\epsilon}(p')$. Indeed, we have:

$$\sum_{\omega \in \Omega} F_{\theta, a^{\theta, \delta}(\overline{p}), \omega} p'_{\omega} - c_{\theta, a^{\theta, \delta}(\overline{p})} \ge \sum_{\omega \in \Omega} F_{\theta, a^{\theta, \delta}(\overline{p}), \omega} \overline{p}_{\omega} - c_{\theta, a^{\theta, \delta}(\overline{p})} - \epsilon$$
$$> \sum_{\omega \in \Omega} F_{\theta, a, \omega} \overline{p}_{\omega} - c_{\theta, a} - \epsilon - \delta$$
$$\ge \sum_{\omega \in \Omega} F_{\theta, a, \omega} p'_{\omega} - c_{\theta, a} - 2\epsilon - \delta,$$

for every actions $a \in A$. The above inequalities follows since:

$$\sum_{\omega \in \Omega} F_{\theta, a, \omega}(\overline{p}_{\omega} - p'_{\omega}) \le \|F_{\theta, a}\|_1 \|\overline{p} - p'\|_{\infty} \le \epsilon,$$

for every actions $a \in \mathcal{A}$ and $\theta \in \Theta$.

Therefore, we can prove that the following holds:

$$\sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(\overline{p}, a^{\theta, \delta}(\overline{p}), \theta) \geq \sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p', a^{\theta, \delta}(\overline{p}), \theta) - \epsilon$$
$$\geq \sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p', a^{\theta, \delta+2\epsilon}(p'), \theta) - \epsilon$$
$$\geq \sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p^{\star}, a^{\theta, \delta}(p^{\star}), \theta) - 2\sqrt{2\epsilon}$$
$$= \mathsf{OPT}(\mathcal{C}, \delta) - 2\sqrt{2\epsilon}, \tag{8}$$

where the first inequality above holds because $\sum_{\omega \in \Omega} F_{\theta,a,\omega}(\overline{p}_{\omega} - p'_{\omega}) \leq ||F_{\theta,a}||_1 ||\overline{p} - p'||_{\infty} \leq \epsilon$ for every actions $a \in \mathcal{A}$ and type $\theta \in \Theta$. The second inequality holds because of the definition of δ -best response, while the third inequality follows from Lemma 3.

At this point, we can decompose the cumulative regret as follows:

$$R_T(\mathcal{C},\delta) = T \cdot \operatorname{OPT}(\mathcal{C},\delta) - \mathbb{E}\Big[\sum_{t=1}^T \sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p^t, a^{\theta,\delta}(p^t), \theta)\Big]$$

$$\leq \underbrace{T \cdot \operatorname{OPT}(\mathcal{C}, \delta) - T \cdot \max_{p \in \mathcal{B}_{\epsilon}} \sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p, a^{\theta, \delta}(p), \theta)}_{(a)}}_{(a)} + \underbrace{T \cdot \max_{p \in \mathcal{B}_{\epsilon}} \sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p, a^{\theta, \delta}(p), \theta) - \mathbb{E}\Big[\sum_{t \in [T], \theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p^{t}, a^{\theta, \delta}(p^{t}), \theta)\Big]}_{(b)}.$$
(9)

We focus on bounding term (a) in Eq. (9). Let $\bar{p} \in [0,1]^m$ be the contract defined at point 3. Then, we have:

$$T \cdot \operatorname{OPT}(\mathcal{C}, \delta) - T \max_{p \in \mathcal{B}_{\epsilon}} \sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p^{t}, a^{\theta, \delta}(p^{t}), \theta) \leq T \cdot \operatorname{OPT}(\mathcal{C}, \delta) - T \sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(\bar{p}, a^{\theta, \delta}(\bar{p}), \theta) \\ \leq 2\sqrt{2\epsilon} T,$$

where the last inequality holds because of Eq. (8). We focus on bounding term (b) in Eq. (9). By employing the same analysis to prove the regret bound in UCB1 (see, *e.g.*, (Lattimore & Szepesvári, 2020)), we obtain:

$$T \cdot \max_{p \in \mathcal{B}_{\epsilon}} \sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p, a^{\theta, \delta}(p), \theta) - \mathbb{E}\Big[\sum_{t \in [T]} \sum_{\theta \in \Theta} \lambda_{\theta} u^{\mathsf{P}}(p^{t}, a^{\theta, \delta}(p^{t}), \theta)\Big] \le \mathcal{O}\left(\sqrt{T|\mathcal{B}_{\epsilon}|\log(T)}\right).$$
(10)

Finally, observing that $|\mathcal{B}_{\epsilon}| \leq \mathcal{O}((1/\epsilon)^m)$, by putting (a) and (b) in Eq. (9) together and setting $\epsilon = T^{-\frac{1}{m+1}}$, we have:

$$R_T(\mathcal{C},\delta) \leq \widetilde{\mathcal{O}}\left(T^{1-\frac{1}{2(m+1)}}\right),$$

concluding the proof.

Corollary 1. The regret suffered by Algorithm 2 can be upper bounded as follows:

$$R_T(\mathcal{C}) \leq \widetilde{\mathcal{O}}\left(T^{1-\frac{1}{2(m+1)}}\right) + 2\sqrt{\delta}T.$$

Proof. Let p^* be an optimal non-robust contract inside C. Then we have:

$$\begin{aligned} R_T(\mathcal{C}) &= T \cdot \operatorname{OPT}(\mathcal{C}) - \mathbb{E} \Big[\sum_{t=1}^T \sum_{\theta \in \Theta} \lambda_\theta u^{\mathsf{P}}(p^t, a^{\theta, \delta}(p^t), \theta) \Big] \\ &= T \cdot (\operatorname{OPT}(\mathcal{C}) - \operatorname{OPT}(\mathcal{C}, \delta)) + T \cdot \operatorname{OPT}(\mathcal{C}, \delta) - \mathbb{E} \Big[\sum_{t=1}^T \sum_{\theta \in \Theta} \lambda_\theta u^{\mathsf{P}}(p^t, a^{\theta, \delta}(p^t), \theta) \Big] \\ &\leq 2\sqrt{\delta}T + T \cdot \operatorname{OPT}(\mathcal{C}, \delta) - \mathbb{E} \Big[\sum_{t=1}^T \sum_{\theta \in \Theta} \lambda_\theta u^{\mathsf{P}}(p^t, a^{\theta, \delta}(p^t), \theta) \Big] \\ &\leq 2\sqrt{\delta}T + \widetilde{\mathcal{O}} \left(T^{1 - \frac{1}{2(m+1)}} \right), \end{aligned}$$

where the first inequality holds by employing the same argument needed to prove Proposition 1, and the second inequality holds thanks to Theorem 2, concluding the proof. \Box