
Incentivizing the Dynamic Workforce: Learning Contracts in the Gig-Economy

Alon Cohen
Technion—Israel Institute of Technology
{alon.cohen, ko}@technion.ac.il

Moran Koren

Argyrios Deligkas
University of Liverpool
argyrios.deligkas@liverpool.ac.uk

Abstract

In principal-agent models, a principal offers a contract to an agent to preform a certain task. The agent exerts a level of effort that maximizes her utility. The principal is oblivious to the agent’s chosen level of effort, and conditions her wage only on possible outcomes. In this work, we consider a model in which the principal is unaware of the agent’s utility and action space. She sequentially offers contracts to identical agents, and observes the resulting outcomes. We present an algorithm for learning the optimal contract under mild assumptions. We bound the number of samples needed for the principal obtain a contract that is within ϵ of her optimal net profit for every $\epsilon > 0$.

1 Introduction

Recent technological advances have had a profound impact over the relationship between firms and employees. The rapid change in employees’ required skill set combined with the availability of off-shore, qualified, cheap hiring alternatives drove employers to adopt employees on a short-term, task specific bases. This hiring model, dubbed *dynamic workforce* or *gig-economy*, has been growing exponentially over the past couple of years. Bughin et al. [2016] found that 20-30% of US workers are employed independently, at least partly, and predicts this trend will continue with popularity of the “Lean Startup” business model, and the appearance of platforms such as “Amazon Mechanical Turk”.

The aforementioned changes in the labor market have a profound effect on the information available to firms when it making hiring decision. In the traditional workspace, hiring is often a long-term, expensive procedure. Nowadays hiring is cheap and short-termed. As a result, an employer has less information on the character and quality of her employees, thus lacking the knowledge on how to best align their incentives with those of the firm. In this work we study a model more suitable to the new “dynamic workforce”, and, using this model, we try to learn the best wages an employer should offer to her employees.

Economists use the term *agency problems* to describe models like the above. In such models one party, the *principal*, offers a *contract* to another party, the *agent*, to perform a task [Laffont and Martimort, 2009]. In his seminal paper, Ross [1973] introduce the term *principal-agent* to capture these models. In the basic model, the outcome of the task is chosen randomly from a distribution determined by the *level of effort* invested by the agent. A higher effort level induces a distribution in which the probability of a better outcome is higher than in lower effort levels. On the other hand, a higher level causes the agent greater disutility. Ross [1973] assume that the principal is *risk-neutral* while the agent is *risk-averse*, therefore the incentives of both parties are *not* fully aligned.* Hence, the effort level that is optimal from the perspective of the principal,

*When given the choice between participating in some lottery X or receiving $E(X)$ with probability one, a risk-averse agent will strictly prefer the latter while a risk-neutral agent is indifferent. In the Von Neumann

is not necessarily the optimal one for the agent. To bridge this gap, the principal offers the agent a contract in which the agent is rewarded for any additional effort. The lion’s share of previous work assume that the principal, while oblivious to the agent’s choice, has full information about the agent’s utility structure *and* the set of levels of effort she can choose from *and* the probability distribution associated with every level of effort [Ross, 1973, Grossman and Hart, 1983, Holmstrom, 1979, Gershkov and Perry, 2012, Holmström, 2017]. These assumptions seem confining, especially when considering the motivating scenario of a dynamic workplace.

Holmström [2017] identified this gap in theory as one of the main challenges in current research. Ho et al. [2016] were among the first who theoretically tried to bridge it. They consider a setting where there are several types of agents, each one with her own utility function, set of effort levels and effort costs, and probability distributions over the outcomes; all of which unknown to the principal. They considered a repeated setting with T rounds, where in every rounds the type of the agent is chosen i.i.d. from an unknown probability distribution. Their goal was to find, before round T , the optimal contract from a predefined finite set S . They followed a multi-armed bandit [Robbins, 1985] approach and derive an algorithm that finds an approximately-optimal contract. However, the contract they find is approximately-optimal *only* with respect to the best contract in S . They provide no theoretical guarantees on approximating the optimal contract overall, but only for the case of one effort level for the agent, i.e. reject the contract or accept it, known as *posted-price auctions*.*

In this paper we follow a similar route to Ho et al. [2016]. We assume the principal has *zero* information about the agent; we assume unknown utility for the agent, unknown set of effort levels and their associated costs, and unknown probability distribution for any effort level. The only knowledge the principal has is the set of outcomes and their corresponding profits.

Our contribution. We introduce a novel set of contracts we call *monotone-smooth* contracts; a large subset of monotone contracts set studied in [Ho et al., 2016]. We complement this definition with a suitable discretization of the contract space. Unlike Ho et al. [2016], we show that for any monotone-smooth contract there exists a contract in the discretized space for which the principal’s expected net profit is ϵ -approximated. As far as we are aware, this is the first work to do so! Moreover, our result does not assume any specific agent utility function, but rather only mild assumptions. This allows to apply machineries from multi-armed bandit theory to find a contract ϵ -optimal against any monotone-smooth contract. Finally, we present two important cases in which economic theory suggests the learned contract is ϵ -optimal against *any* contract the principal may offer.

Further related work. Our work lies in the intersection of principal-agent models and multi-armed bandit theory; Ho et al. [2016] provide an excellent overview of literature in the field. Sannikov [2008, 2011], Williams [2004] study a repeated setting where the principal interacts with the agent for multiple periods. Conitzer and Garera [2006] empirically compare several learning algorithms in a setting similar to [Ho et al., 2016].

The Lipschitz Bandit problem [Agrawal, 1995] is a generalization of the multi-armed bandit problem in which the set of arms comes from some compact space, and the expected reward of each arm is a Lipschitz function of the arm. It has received much attention from the bandit theory community over the years [Kleinberg, 2005, Auer et al., 2007, Bubeck et al., 2011, Kleinberg et al., 2013, Magureanu et al., 2014].

The technique was used in the seminal paper of Kleinberg and Leighton [2003] for learning posted-price auctions, which can be seen as a principal-agent problem with only one outcome. There is a flourish of studies on learning auctions [Blum et al., 2004, Elkind, 2007, Cole and Roughgarden, 2014, Balcan et al., 2016, Gonczarowski and Nisan, 2016, Morgenstern and Roughgarden, 2016, Roughgarden and Schrijvers, 2016, Syrgkanis, 2017, Bubeck et al., 2017].

2 Preliminaries

In what follows, $[m^*] := \{0, 1, \dots, m\}$ and $[m] := \{1, \dots, m\}$ for every natural m .

Morgenstern utility theory, if an agent is risk-averse, then she has a concave utility function; if she is risk-neutral, then her utility is linear.

* As they highlight, it is not generally unclear whether the best contract from S can provide a good theoretical guarantees for the general problem of dynamic contract design.

We study principal-agent problems with k outcomes and n effort levels. Let $0 < \pi(1) < \pi(2) < \dots < \pi(k)$ denote the value the principal gets under outcome $i \in [k]$. We assume that $\pi(k) = H$, hence the values are bounded.

A contract $w = (w(1), \dots, w(k))$ specifies a positive payment to the agent for every outcome; namely, $w(i)$ is the wage the principal pays the agent for outcome i .

Upon receiving a contract, the agent chooses an effort level $e \in [n^*]$. Every effort level is associated with a probability distribution f_e over the set of outcomes and a cost $c(e)$; $f_e(j)$ is the probability of realizing outcome j when the agent chooses the effort level e . In effort level 0, the agent rejects the contract, her utility is zero under any contract, and by convention the value for the principal is zero. We assume that the effort levels are ordered, i.e., $1 \prec \dots \prec n$. We follow the literature and assume that the agent has a von Neumann-Morgenstern utility. For a contract $w = (w(1), \dots, w(k))$ she chooses effort level of $\hat{e}(w)$ as to maximize her utility, defined as

$$U(w, e) = \sum_{j=1}^k f_e(j) \cdot u(w(j)) - c(e),$$

where u is a monotonically-increasing concave function. Hence, $\hat{e}(w) = \arg \max_{e \in [n^*]} U(w, e)$.

The principal is risk-neutral; when she offers contract w to the agent, her expected net profit from the contract is,

$$V(w) = \sum_{j=1}^k f_{\hat{e}(w)}(j) \cdot (\pi(j) - w(j)).$$

To ensure that higher effort levels yield higher expected profit for the principal, the literature commonly lays down some assumptions about the outcome distributions.

Assumption 1 (First-order Stochastic Dominance (FOSD)). *A probability distribution associated with higher effort first order stochastically dominates a probability distribution associated with lower effort. Formally, if $e \succ e'$, then for every $j \in [k]$ it holds that $\sum_{i=j}^k f_e(i) \geq \sum_{i=j}^k f_{e'}(i)$.*

Note that the assumption is equivalent to the following. For every pair of effort levels $e \succ e'$ and for every sequence of real numbers $a(1) \leq \dots \leq a(k)$,

$$\sum_{i=1}^k f_e(i) \cdot a(i) \geq \sum_{i=1}^k f_{e'}(i) \cdot a(i). \quad (1)$$

Additionally, to break ties between effort levels we assume the following.

Assumption 2. *When indifferent between two or more levels of effort, the agent will choose the higher effort.*

In this work, the principal is faced with a stream of agents. The agents are all different but identical—they share a common utility function, effort levels, costs from effort, and outcome distributions associated with each effort level. The principal proceeds in rounds $t = 1, 2, \dots$. On round t , the principal offers a contract w_t to the agent associated with this round. The agent privately chooses effort level $\hat{e}(w_t)$ unbeknown to the principal. The principal observes only the outcome i_t independently drawn from $f_{\hat{e}(w_t)}$, and consequently gets a net profit of $\pi(i_t) - w_t(i_t)$.

In what follows, for $\epsilon > 0$, the goal of the principal is to find an ϵ -optimal contract in the minimum number of rounds. A contract w is ϵ -optimal if $V(w) \geq V(w') - \epsilon$ for every $w' \in W$, for a set of contracts W to be defined in the sequel.

2.1 Multi-armed Bandit

In the multi-armed bandit problem [Robbins, 1985], a decision maker sequentially collects rewards from a given set of arms. In each round, the decision maker chooses a single arm, and observes an independent sample from a reward distribution associated with that arm. In our case, the goal of the decision maker is, after a predetermined number of rounds, to select an ϵ -optimal arm; that is, an arm whose expected reward is at most ϵ less than the expected reward of any arm.

When the set of arms is finite, of size N , and the rewards are bounded in $[0, B]$, the seminal work of Even-Dar et al. [2006] presents an algorithm called MEDIANELIMINATION with the following guarantee.

Theorem 1 (Even-Dar et al. [2006]). *The MEDIANELIMINATION(ϵ, δ) returns an ϵ -optimal arm with probability at least $1 - \delta$ after $O((NB^2/\epsilon^2) \cdot \log(1/\delta))$ rounds.*

In our problem, each contract can be seen as an arm. The expected reward of each arm is exactly the principal's utility associated with this contract. It is then expected that the principal would simply execute MEDIANELIMINATION on the space of contracts to obtain an ϵ -optimal one. However, the space of contracts is not finite which is crucial for MEDIANELIMINATION to run. In the sequel we show how to overcome this difficulty by discretizing the space of contracts, and running MEDIANELIMINATION over the discretization.

3 Main Result

In this section we present our algorithm and analyze its sample complexity, but before doing so let us first define the space of contracts W that we can learn. The algorithm is presented in Section 3.2.

3.1 Learnable Contracts

Let $w_0 > 0$ be a minimum wage for any outcome.

Definition 1 (B -bounded contract). *A contract w is B -bounded if $w_0 \leq w(i) \leq B$ for every $i \in [k]$.*

For a bounded contract, together with the assumption that the principal's profits are bounded, ensures that the principal's expected net profit $V(w)$ can be estimated statistically.

Definition 2 (Monotone-smooth contract). *A contract w is monotone-smooth if for every $i \in [k-1]$ it holds that $0 \leq w(i+1) - w(i) \leq \pi(i+1) - \pi(i)$.*

For a monotone-smooth contract, Eq. (1) ensures that, keeping the contract fixed, the principal's utility cannot decrease if the agent increases her effort level. In the sequel, this property allows us to bound the difference in the principal's utility between two similar contracts.

We define W as follows.

$$W = \{w : w \text{ is monotone-smooth and } H\text{-bounded}\} .$$

We are aware that this set seems restrictive at first glance, yet we argue that in some important special cases, the principal's optimal net profit is achieved by a contract from this set. For example, when there are only two outcomes or the utility of the agent is linear (see Section 4).

3.2 Algorithm

Let $w^* \in W$ be an optimal contract in W , that is $V(w^*) \geq V(w)$ for all $w \in W$. The goal of our algorithm is to find an ϵ -optimal contract w , namely a contract for which $V(w^*) \leq V(w) + \epsilon$ within a predetermined number of rounds. However, we conjecture that it cannot be done in general, and to alleviate this issue we make the following assumption.

Assumption 3 (Bounded Risk-Aversion). *The agent's utility from wage u is twice continuously-differentiable and satisfies $-u''(w)/u'(w) \leq 1/w$ for all $w > 0$. This is equivalent to $w \mapsto w \cdot u'(w)$ being monotone-nondecreasing in w .*

Intuitively, the assumption ensures that making small changes to a contract does not produce drastically different behavior by the agent.*

Our algorithm works as follows. The principal initially constructs a cover W_η of W , and then run MEDIANELIMINATION on W_η . Indeed, the main technical difficulty in this paper is in defining W_η properly so that the following result holds.

*An equivalent formulation of the assumption is that the agent's Arrow-Pratt relative risk aversion measure is smaller than unity, i.e., $\frac{-x \cdot u''(x)}{u'(x)} < 1$. Arrow [1965], Pratt [1964]. When considering finite outcomes, this assumption is not very restrictive (see discussion in Arrow [1965], Choi and Menezes [1992]).

Theorem 2. *Suppose that Assumptions 1 to 3 hold. Let $\eta < 1/4k$. There exists a contract space W_η such that for every contract $w \in W$, there is a contract $w' \in W_\eta$ for which $V(w) \leq V(w') + 2kH\eta$. Moreover, the size of W_η is at most $M = (\log(2H/w_0)/\eta)^k$, and W_η can be constructed in time $O(kM)$.*

The proof of the theorem is found in Section 3.3. Finally, we have our main result.

Theorem 3. *Suppose that Assumptions 1 to 3 hold. Let $\eta = \epsilon/4kH$. Executing $\text{MEDIANELIMINATION}(\epsilon/2, \delta)$ on the set W_η produces the following guarantee. With probability at least $1 - \delta$, the algorithm outputs an ϵ -optimal contract after*

$$O\left(\left(\frac{4kH \log(H/w_0)}{\epsilon}\right)^{k+2} \log(1/\delta)\right)$$

rounds.

Proof. By Theorem 2 and by the choice of η , there is a $w' \in W_\eta$ for which $V(w') \leq V(w^*) + \epsilon/2$. By Theorem 1, with probability $1 - \delta$, MEDIANELIMINATION returns a contract $\hat{w} \in W_\eta$ such that $V(\hat{w}) \leq V(w') + \epsilon/2$. Combining both results we get

$$V(\hat{w}) \leq V(w') + \epsilon/2 \leq V(w^*) + \epsilon/2 + \epsilon/2 = V(w^*) + \epsilon,$$

as required. Moreover, MEDIANELIMINATION is done in the following number of rounds:

$$O\left(\frac{|W_\eta|H^2}{(\epsilon/2)^2} \log(1/\delta)\right) = O\left(\left(\frac{4kH \log(H/w_0)}{\epsilon}\right)^{k+2} \log(1/\delta)\right). \quad \square$$

3.3 Discretization of the Contract Space

In this section we prove Theorem 2. We start by defining the notion of a coarse contract.

Definition 3 (η -coarse contract). *A contract W is η -coarse if there exists natural numbers l_0, l_1, \dots, l_{k-1} such that $w(1) = w_0 \exp(\eta \cdot l_0)$, and for $i \in [k-1]$, $w(i+1) = w(i) \exp(\eta \cdot l_i)$.*

That is, a coarse is a contract in which the ratios between wages of consecutive outcomes come from a discrete set of options. We define:

$$W_\eta = \{w : w \text{ is } \eta\text{-coarse and } 2H\text{-bounded}\}.$$

We prove the following.

Lemma 4. *The size of W_η is at most $M = (\log(2H/w_0)/\eta)^k$. Moreover, W_η can be constructed in time $O(kM)$.*

Proof. The wage of outcome i has the form $w_0 \exp(\eta \cdot l)$ for a natural number l , and satisfies $w(i) \leq 2H$. Therefore, the number of choices for $w(i)$ is at most $\log(2H/w_0)/\eta$. Since there are k outcomes, there must be at most M such contracts. To construct W_η we can go over all of its elements one-by-one, which takes $O(M)$ time. \square

Finally let $w \in W$, we need to show that there is $w' \in W_\eta$ such that $V(w) \leq V(w') + 2kH\eta$. We construct $w' \in W_\eta$ as follows. We let $l_0 = \lceil \log(w(1)/w_0)/\eta \rceil$, and for $i \in [k-1]$, $l_i = \lceil \log(w(i+1)/w(i))/\eta \rceil$. We can immediately observe that, by construction, $w'(1) \geq w(1)$ and for $i \in [k-1]$, we have $w'(i+1)/w'(i) \geq w(i+1)/w(i)$. Moreover, it is clear that w' is η -coarse and that $w'(i) \geq w_0$, yet it remains to show that $w'(i) \leq 2H$. For that, we have the following lemma.

Lemma 5. *We have for all $i \in [k-1]$, $w'(i) \leq e^{\eta i} w(i)$.*

Proof. By construction, for each $i \in [k-1]$ we have $w'(i+1)/w'(i) \leq e^\eta \cdot w(i+1)/w(i)$. From this we entail that $w'(i)/w'(1) \leq \exp((i-1) \cdot \eta) \cdot w(i)/w(1)$. Since also by construction $w'(1) \leq \exp(\eta) \cdot w(1)$, we get that $w'(i) \leq \exp(i \cdot \eta) \cdot w(i)$ as required. \square

With the lemma at hand, and by assumption that $\eta < 1/2k$ we obtain

$$w'(k) \leq \exp(\eta \cdot k) \cdot w(k) \leq e^{1/2} H \leq 2H.$$

Therefore we have $w' \in W_\eta$.

We now show that, compared to w , under w' the agent's effort cannot decrease. Then, we use this fact to bound the difference in the principal's utility between w and w' .

In order to prove our first claim, we will make use of the following lemma.

Lemma 6 (Grossman and Hart [1983]). *Let w^1 and w^2 be contracts. Then,*

$$\sum_{i=1}^k (f_{\hat{e}(w^1)}(i) - f_{\hat{e}(w^2)}(i)) \cdot (u(w^1(i)) - u(w^2(i))) \geq 0.$$

Lemma 7. *The effort level the agent chooses can only increase from w to w' .*

Proof. First, notice that since the wages only increase, had the agent accepted the contract w , i.e., chose an effort level different than 0, she would also accept contract w' . So, let $e' = \hat{e}(w')$ and $e = \hat{e}(w)$ be the effort levels the agent chooses under contracts w' and w respectively.

If we apply Lemma 6 with w and w' , we obtain

$$\sum_{i=1}^k (f_{e'}(i) - f_e(i)) \cdot (u(w'(i)) - u(w(i))) \geq 0.$$

Assume for now that $u(w'(i)) - u(w(i))$ is monotone nondecreasing in i . Given this, we will show by contradiction that $e' \succ e$. So, for the sake of contradiction assume that $e' \prec e$. From the fact that f_e dominates $f_{e'}$, Eq. (1) implies that

$$\sum_{i=1}^k (f_{e'}(i) - f_e(i)) \cdot (u(w'(i)) - u(w(i))) \leq 0.$$

Thus,

$$\sum_{i=1}^k (f_{e'}(i) - f_e(i)) \cdot u(w(i)) = \sum_{i=1}^k (f_{e'}(i) - f_e(i)) \cdot u(w'(i)). \quad (2)$$

Therefore, by optimality of e and e' under contracts w and w' respectively, we obtain

$$\sum_{i=1}^k (f_{e'}(i) - f_e(i)) \cdot u(w(i)) \leq c(e') - c(e), \text{ and } \sum_{i=1}^k (f_{e'}(i) - f_e(i)) \cdot u(w'(i)) \geq c(e') - c(e).$$

Combining Eq. (2) with the two inequalities above, we obtain that

$$U(w', e') = \sum_{i=1}^k f_{e'}(i) \cdot u(w'(i)) - c(e') = \sum_{i=1}^k f_e(i) \cdot u(w'(i)) - c(e) = U(w', e).$$

This means that the agent is indifferent between effort levels e and e' under contract w' . Since, by Assumption 2, the agent chooses the highest effort in this case, we must have $e \prec e'$ — a contradiction.

Hence, in order to prove the lemma it suffices to prove that $u(w'(i)) - u(w(i))$ is monotone nondecreasing in i . This is equivalent to showing that $u(w'(i+1)) - u(w'(i)) \geq u(w(i+1)) - u(w(i))$. Denote $c = w'(i)/w(i)$, for $c \geq 1$. By construction, $w'(i+1)/w(i+1) \geq c$. Consequently, since u is monotone nondecreasing, it suffices to show $u(c \cdot w(i+1)) - u(c \cdot w(i)) \geq u(w(i+1)) - u(w(i))$.

Thus, the proof boils down to showing $c \mapsto u(c \cdot w(i+1)) - u(c \cdot w(i))$ is monotone nondecreasing in c . Taking the derivative with respect to c , we need to show

$$u'(c \cdot w(i+1)) \cdot w(i+1) - u'(c \cdot w(i)) \cdot w(i) \geq 0.$$

However, since the agent is BRA (Assumption 3), we have $u'(c \cdot w(i+1)) \cdot c \cdot w(i+1) \geq u'(c \cdot w(i)) \cdot c \cdot w(i)$. The proof is complete by recalling that $c \geq 1$. \square

We can now bound the loss the principal suffers when she offers w' instead of w .

Lemma 8. *It holds that $V(w) \leq V(w') + 2kH\eta$.*

Proof. Since we focus on scenarios where the optimal contract is monotone-smooth, we get that the net profit of the principal at the optimal contract, $\pi(i) - w(i)$, is nondecreasing in i . Furthermore, from Eq. (1), keeping w fixed, the principal only benefits from an increase of the agent's effort level. Denote $e = \hat{e}(w)$ and $e' = \hat{e}(w')$. From Lemma 7 we know that $e' \succ e$. We have,

$$\begin{aligned} V(w) &= \sum_{i=1}^k f_e(i) \cdot (\pi(i) - w(i)) \\ &\leq \sum_{i=1}^k f_{e'}(i) \cdot (\pi(i) - w(i)) \\ &= V(w') + \sum_{i=1}^k f_{e'}(i) \cdot (w'(i) - w(i)). \end{aligned} \quad (3)$$

Now by Lemma 5, $w'(i) \leq e^{i\eta} \cdot w(i)$ for all $i \in [k]$. Since $e^x - 1 \leq 2x$ for all $x \in [0, 1]$ and since $\eta < 1/k$, we obtain

$$w'(i) - w(i) \leq (e^{i\eta} - 1) \cdot w(i) \leq 2 \cdot \eta \cdot i \cdot w(i) \leq 2 \cdot \eta \cdot k \cdot H.$$

Combining the latter with Eq. (3), we get $V(w) \leq V(w') + 2 \cdot \eta \cdot k \cdot H$. \square

4 Applications

In this section we highlight two cases that received attention in the past. For each of them, when Assumptions 1 to 3 hold, the optimal contract will be in the set W , and thus by learning an ϵ -optimal contract in W , we approximate the best contract the principle could have offered the agent had she known her utility function, effort levels and costs, and the distributions they induce over outcomes.

Two outcomes. Firstly, we focus on the case where there are only two outcomes and show that the optimal contract is in W .

Lemma 9. *When there are only two outcomes, the optimal contract is monotone-smooth and $2H$ -bounded.*

Proof. By Grossman and Hart [1983], the optimal contract in the two outcome case is of the shape: $w(2) = w(1) + a(\pi(2) - \pi(1))$ for some $a \in [0, 1]$. By plugging this expression into the inequality in Definition 2, we get that the optimal contract is monotone-smooth. To see that the optimal contract is $2H$ -bounded, let e denote her chosen effort level. Since the principal's utility at the optimal contract is nonnegative, $f_e(1)w(1) + f_e(2)w(2) \leq f_e(1)\pi(1) + f_e(2)\pi(2) \leq H$, and in particular $w(1) \leq H$. Now, $w(2) = w(1) + a(\pi(2) - \pi(1)) \leq H + 1 \cdot H = 2H$. \square

Thus, by Theorem 3, applied to our discretized contract space, MEDIANELIMINATION finds an ϵ -optimal contract under Assumptions 1 to 3. Note, the FOSD assumption (Assumption 1) is standard in the literature, hence our result essentially requires only the bounded risk-averse assumption (Assumption 3).

Risk neutral agent. Carroll [2015] studies a setting in the agent is risk-neutral, and the principal has only partial knowledge of the agent's action space. Had the principal known the complete action space of the agent, the optimal contract would have been linear. In this setting, the principal can derive the optimal *linear* contract with respect to only the known actions of the agent. Carroll [2015] show that her profit from the actual action taken by the agent (which can be one that the principal is unaware of) can only be higher than the principal's expectations.

In the following lemma we show that when the agent is risk-neutral, the optimal theoretical contract is in W , and thus our algorithm learns a contract that approximates it.

Lemma 10. *If the agent is risk-neutral then the optimal contract is H -bounded and monotone-smooth.*

Proof. When the agent is risk neutral, the optimal contract is of the shape $w(i) = \pi(i) - \alpha$ for some constant $\alpha \in \mathbb{R}_+$ (Proposition 14.B.2 on page 482 in Mas-Colell et al. [1995]). As $\pi(i+1) > \pi(i)$, this contract is monotone-smooth. And as $\pi(k) \leq H$ it is also H -bounded. \square

5 Conclusions

In this paper we studied the principal-agent problem when the principal has zero information about the the agent. We focus on the class of monotone-smooth contracts and show that when the optimal contract is monotone-smooth and the agent is bounded risk-averse, then we can learn an approximately optimal contract. We complemented this result with a multi-armed bandit algorithm that finds an approximately optimal contract and we provided bounds on the number of samples it needs.

When the output space of the task is binary, or when the agent is risk-neutral, economic theory suggests that the optimal contract *is* monotone-smooth. Thus the net profit of the principal generated by the resulting contract approximates the optimal net profit she can achieve in general.

Several intriguing questions remain. It is interesting to understand whether the assumption of bounded risk-aversion is needed to guarantee learning of monotone-smooth contracts. The answer to this question is not obvious even when there are only two outcomes. Furthermore, we wish to find other conditions and assumptions that allow learning. On the other hand, we conjecture that there exist cases in which learning is not possible at all. Lower bounds, or even partial characterizations of such cases would be of great interest.

References

- Rajeev Agrawal. The continuum-armed bandit problem. *SIAM journal on control and optimization*, 33(6):1926–1951, 1995.
- Kenneth J Arrow. *Aspects of the Theory of Risk-Bearing*. Yrjö Jahnsson Foundation, Helsinki, 1965.
- Peter Auer, Ronald Ortner, and Csaba Szepesvári. Improved rates for the stochastic continuum-armed bandit problem. In *International Conference on Computational Learning Theory*, pages 454–468. Springer, 2007.
- Maria-Florina F Balcan, Tuomas Sandholm, and Ellen Vitercik. Sample complexity of automated mechanism design. In *Advances in Neural Information Processing Systems*, pages 2083–2091, 2016.
- Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. *Soda*, 324 (2-3 SPEC. ISS.):137–146, 2004. ISSN 03043975. doi: 10.1016/j.tcs.2004.05.012.
- Sébastien Bubeck, Gilles Stoltz, and Jia Yuan Yu. Lipschitz bandits without the lipschitz constant. In *International Conference on Algorithmic Learning Theory*, pages 144–158. Springer, 2011.
- Sebastien Bubeck, Nikhil R Devanur, Zhiyi Huang, and Rad Niazadeh. Online auctions and multi-scale online learning. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 497–514. ACM, 2017.
- Jacques Bughin, James Manyika, Jonathan Woetzel, Eric Labaye, chairman Michael Chui, Susan Lund, Andres Cadena, Richard Dobbs, Katy George, Rajat Gupta, Eric Hazan, Acha Leke, Scott Nyquist, Gary Pinkus, and Shirish Sankhe. Independent work: Choice, necessity, and the gig economy. Technical report, McKinsey & Company 2016, 2016. URL <https://www.mckinsey.com/{~}/media/McKinsey/GlobalThemes/EmploymentandGrowth/IndependentworkCh>
- Gabriel Carroll. Robustness and linear contracts. *American Economic Review*, 105(2):536–563, 2015. ISSN 00028282. doi: 10.1257/aer.20131159.

- E Kwan Choi and Carmen F Menezes. Is relative risk aversion greater than one? *International Review of Economics & Finance*, 1(1):43–54, 1992. ISSN 1059-0560. doi: 10.1016/1059-0560(92)90005-W. URL <http://www.sciencedirect.com/science/article/B6W4V-45DN2KK-5/2/4b861bdff2a7a2ffa0c4a39a9564f8c>
- Richard Cole and Tim Roughgarden. The Sample Complexity of Revenue Maximization. *Stoc*, pages 0–25, 2014. ISSN 07378017. doi: 10.1145/2591796.2591867.
- Vincent Conitzer and Nikesh Garera. Learning algorithms for online principal-agent problems (and selling goods online). In *Proceedings of the 23rd international conference on Machine learning*, pages 209–216. ACM, 2006.
- Edith Elkind. Designing and learning optimal finite support auctions. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 736–745. Society for Industrial and Applied Mathematics, 2007.
- Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(Jun):1079–1105, 2006.
- A. Gershkov and M. Perry. Dynamic Contracts with Moral Hazard and Adverse Selection. *The Review of Economic Studies*, 79(1):268–306, jan 2012. ISSN 0034-6527. doi: 10.1093/restud/rdr026. URL <https://academic.oup.com/restud/article-lookup/doi/10.1093/restud/rdr026>.
- Yannai A. Gonczarowski and Noam Nisan. Efficient Empirical Revenue Maximization in Single-Parameter Auction Environments. pages 1–39, 2016. ISSN 07378017. doi: 10.1145/3055399.3055427. URL <http://arxiv.org/abs/1610.09976>.
- Sanford J Grossman and Oliver D Hart. An Analysis of the Principal-Agent Problem. *Econometrica*, 51(1):7–45, 1983. ISSN 00129682. doi: 10.2307/1912246.
- Chien Ju Ho, Aleksandrs Slivkins, and Jennifer Wortman Vaughan. Adaptive contract design for crowdsourcing markets: Bandit algorithms for repeated principal-agent problems. *Journal of Artificial Intelligence Research*, 55:317–359, 2016. ISSN 10769757. doi: 10.1145/2600057.2602880.
- Bengt Holmstrom. Moral Hazard and Observability. *The Bell Journal of Economics*, 10(1):74, 1979. ISSN 0361915X. doi: 10.2307/3003320. URL <http://www.jstor.org/stable/3003320?origin=crossref>.
- Bengt Holmström. Pay for performance and beyond. *American Economic Review*, 107(7):1753–1777, 2017. ISSN 1473-7167. doi: 10.1586/14737167.7.4.351. URL <http://www.tandfonline.com/doi/full/10.1586/14737167.7.4.351>.
- Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 594–605. IEEE, 2003.
- Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. Bandits and experts in metric spaces. *arXiv preprint arXiv:1312.1277*, 2013.
- Robert D Kleinberg. Nearly tight bounds for the continuum-armed bandit problem. In *Advances in Neural Information Processing Systems*, pages 697–704, 2005.
- Jean-Jacques Laffont and David Martimort. *The theory of incentives: the principal-agent model*. Princeton university press, 2009.
- Stefan Magureanu, Richard Combes, and Alexandre Proutiere. Lipschitz bandits: Regret lower bounds and optimal algorithms. *arXiv preprint arXiv:1405.4758*, 2014.
- Andreu Mas-Colell, Michael Dennis. Whinston, and Jerry R. Green. *Microeconomic theory*. Oxford University Press, 1995. ISBN 0195073401.

- Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *Conference on Learning Theory*, pages 1298–1318, 2016.
- John W. Pratt. Risk Aversion in the Small and in the Large. *Econometrica*, 32(1/2):122, jan 1964. ISSN 00129682. doi: 10.2307/1913738. URL <https://www.jstor.org/stable/1913738?origin=crossref>.
- Herbert Robbins. Some aspects of the sequential design of experiments. In *Herbert Robbins Selected Papers*, pages 169–177. Springer, 1985.
- Stephen A Ross. The Economic Theory of Agency : The Principal’s Problem. *American Economic Review*, 63(2):134–139, 1973.
- Tim Roughgarden and Okke Schrijvers. Ironing in the dark. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 1–18. ACM, 2016.
- Y Sannikov. Contracts: The theory of dynamic principal–agent relationships and the continuous-time approach. 1:89–124, 01 2011.
- Yuliy Sannikov. A continuous- time version of the principal: Agent problem. *The Review of Economic Studies*, 75(3):957–984, 2008. ISSN 00346527, 1467937X. URL <http://www.jstor.org/stable/20185061>.
- Vasilis Syrgkanis. A sample complexity measure with applications to learning optimal auctions. In *Advances in Neural Information Processing Systems*, pages 5358–5365, 2017.
- Noah Williams. On dynamic principal-agent problems in continuous time. Technical report, UCLA Department of Economics, 2004.