

Simple versus Optimal Contracts

PAUL DÜTTING, London School of Economics, UK

TIM ROUGHGARDEN, Columbia University, USA

INBAL TALGAM-COHEN, Technion, Israel

We consider the classic principal-agent model of contract theory, in which a principal designs an outcome-dependent compensation scheme to incentivize an agent to take a costly and unobservable action. When all of the model parameters—including the full distribution over principal rewards resulting from each agent action—are known to the designer, an optimal contract can in principle be computed by linear programming. In addition to their demanding informational requirements, however, such optimal contracts are often complex and unintuitive, and do not resemble contracts used in practice.

This paper examines contract theory through the theoretical computer science lens, with the goal of developing novel theory to explain and justify the prevalence of relatively simple contracts, such as linear (pure commission) contracts. First, we consider the case where the principal knows only the first moment of each action's reward distribution, and we prove that linear contracts are guaranteed to be worst-case optimal, ranging over all reward distributions consistent with the given moments. Second, we study linear contracts from a worst-case approximation perspective, and prove several tight parameterized approximation bounds.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory and mechanism design**; • **Applied computing** → **Law; Economics**.

Additional Key Words and Phrases: principal-agent model, model uncertainty, max-min robustness

ACM Reference Format:

Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. 2019. Simple versus Optimal Contracts. In *ACM EC '19: ACM Conference on Economics and Computation (EC '19), June 24–28, 2019, Phoenix, AZ, USA*. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3328526.3329591>

1 INTRODUCTION

Classic contract theory. Many fundamental economic interactions can be phrased in terms of two parties, a *principal* and an *agent*, where the agent chooses an action and this imposes some (negative or positive) externality on the principal. Naturally, the principal will want to influence which action the agent chooses. This influence will often take the form of a *contract*, in which the principal compensates the agent contingent on either the actions or their outcomes; with the more challenging and realistic scenario being the one in which the principal cannot directly observe the agent's chosen action. Instead, the principal can only observe a stochastic outcome that results from the agent's action.

For example, consider a salesperson (agent) working for a company (principal) producing a range of products with different revenue margins. The salesperson chooses the amount of effort to spend on promoting the various products. The company may not be able to directly observe effort, but can presumably track the number of orders the salesperson generates. Assuming this number is correlated with the salesperson's actions (the harder he works, the more sales of higher margin

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '19, June 24–28, 2019, Phoenix, AZ, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6792-9/19/06...\$15.00

<https://doi.org/10.1145/3328526.3329591>

products he generates),¹ it may make sense for the company to base his pay on sales—i.e., to put him on commission—to induce him to expend the appropriate level of effort. Another example is the interaction between a car owner (agent) and an insurance company (principal). The car owner's behavior influences the risks his car is exposed to. If realized risks are borne entirely by the insurance company, the owner might not bother to maintain the asset properly (e.g., might park in a bad neighborhood). This is an example of the well-known *moral hazard* problem. Typically, such bad behavior is difficult to contract against directly. But because this behavior imposes an externality on the insurance company, companies have an interest in designing contracts that guard against it.

Both of these examples fall under the umbrella of the *hidden action principal-agent model*, arguably the central model of *contract theory*, which in turn is an important and well-developed area within microeconomics.² Perhaps surprisingly, this area has received far less attention from the theoretical computer science community than auction and mechanism design, despite its strong ties to optimization, and potential applications ranging from crowdsourcing platforms [25], through blockchain-based smart contracts [17], to incentivizing quality healthcare [8].

The model. Every instance of the model can be described by a pair (A_n, Ω_m) of n actions and m outcomes. In the salesperson example, the actions are the levels of effort and the outcomes are the revenues from ordered products. As in this example, we usually identify the (abstract) outcomes with the (numerical) rewards to the principal associated with them. The agent chooses an action $a_i \in A_n$, unobservable to the principal, which incurs a cost $c_i \geq 0$ for the agent, and results in a distribution F_i with expectation R_i over the outcomes in Ω_m . The realized outcome $x_j \geq 0$ is awarded to the principal.

The principal designs a contract that specifies a *payment* $t(x_j) \geq 0$ to the agent for every outcome x_j (since the outcomes, unlike the actions, are observable to the principal). This induces an expected payment $T_i = \mathbb{E}_{x_j \sim F_i}[t(x_j)]$ for every action a_i . The agent then chooses the action that maximizes his expected *utility* $T_i - c_i$ over all actions (“incentive compatibility”), or opts out of the contract if no action with non-negative expected utility exists (“individual rationality”).

As the design goal, the principal wishes to maximize her expected *payoff*: the expected outcome R_i minus the agent's expected payment T_i , where a_i is the action chosen by the agent. Therefore contract design immediately translates into the following *optimization problem*: given (A_n, Ω_m) , find a payment vector t that maximizes $R_i - T_i$, where a_i is incentive compatible and individually rational for the agent. We focus on the *limited liability* case, where the contract's payments t are constrained to be non-negative (i.e., money only flows from the principal to the agent).³ A detailed description of the model appears in Section 2.

Optimal contracts and their issues. With perfect knowledge of the distributions mapping actions to outcomes, it is straightforward to solve the optimization problem associated with finding the optimal contract for the principal, by solving one linear program per action. Each linear program minimizes the expected payment to the agent subject to the constraint that he prefers this action to

¹For clarity we shall address the agent as male and the principal as female.

²For example, the 2016 Nobel Prize in economics was awarded to Oliver Hart and Bengt Holmström for their contributions to contract theory. The prize announcement stresses the importance of contracts in practice, describing modern economies as “held together by innumerable contracts”, and explains how the theoretical tools created by Hart and Holmström are invaluable to their understanding.

³Without some such assumption there is a simple but unsatisfying optimal solution for the principal when the agent is risk-neutral: simply sell the project to the agent, at a price just below the maximum expected welfare $R_i - c_i$ the agent can generate by choosing an action. The agent may as well accept (and then select the welfare-maximizing action), and the principal pockets essentially the full welfare. This solution is incompatible with practical principal-agent settings, e.g., a salesperson does not typically buy the company from its owner.

any other—including opting out of the contract—and subject to the payments all being non-negative. The best of these linear programs (achieving the highest expected payoff for the principal) gives the best action to incentivize and an optimal contract incentivizing it.

In addition to their demanding informational requirements, however, such optimal contracts are often complex and unintuitive, and do not resemble contracts used in practice.⁴ Example 1.1 demonstrates this critique—it is not clear how to interpret the optimal contract nor how to justify it to a non-expert. The optimal payment scheme in this example is not even *monotone*, i.e., a better outcome for the principal can result in a lower payment for the agent! In the salesperson example, this would create an incentive for the salesperson to manipulate the outcome, for example by hiding or canceling orders.

Example 1.1. There are $m = 6$ outcomes $x = (1, 1.1, 4.9, 5, 5.1, 5.2)$, and $n = 4$ actions with the following outcome distributions and costs: $F_1 = (\frac{3}{8}, \frac{3}{8}, \frac{1}{4}, 0, 0, 0)$, $F_2 = (0, \frac{3}{8}, \frac{3}{8}, \frac{1}{4}, 0, 0)$, $F_3 = (0, 0, \frac{3}{8}, \frac{3}{8}, \frac{1}{4}, 0)$, $F_4 = (0, 0, 0, \frac{3}{8}, \frac{3}{8}, \frac{1}{4})$, and $(c_1, c_2, c_3, c_4) = (0, 1, 2, 2.2)$. The LP-based approach shows that the optimal contract in this case incentivizes action a_3 with payments

$$t \approx (0, 0, 0.15, 3.93, 2.04, 0).$$

Linear contracts as an alternative. Perhaps the simplest non-trivial contracts are *linear* contracts, where the principal commits to paying the agent an α -fraction of the realized outcome (i.e., payments are linear in the outcomes). Unlike optimal contracts, linear contracts are conceptually simple and easy to explain; their payments are automatically monotone; and they are the most ubiquitous contract form in practice.⁵ From an optimization standpoint, however, they can be suboptimal even in simple settings, as the next example demonstrates:

Example 1.2. There are $m = 2$ outcomes $x = (1, 3)$, and $n = 2$ actions a_1 and a_2 with $F_1 = (1, 0)$, $c_1 = 0$ and $F_2 = (0, 1)$, $c_2 = \frac{4}{3}$, respectively. The optimal contract incentivizes action a_2 with payments $t = (0, \frac{4}{3})$, resulting in expected payoff of $3 - \frac{4}{3} = \frac{5}{3}$ for the principal. The maximum expected payoff of any linear contract is 1 (regardless of which action is incentivized).

Simple versus optimal contracts in the economic literature. The critique that optimal contracts are unrealistically complex is well known to economists. Already in 1987 Milgrom and Holmström [26, p. 326] wrote:

“It is probably the great robustness of linear rules based on aggregates that accounts for their popularity. That point is not made as effectively as we would like by our model; we suspect that it cannot be made effectively in any traditional Bayesian model.”

Indeed, for several decades the economic literature struggled to find the right tools to explain the prevalence of simple, in particular, linear contracts in practice.

A breakthrough was the recent paper by Carroll [13], in which he proposes “a forthrightly non-Bayesian model of robustness” (p. 537). In this work, a key change to the standard principal-agent model is introduced: the set of actions available to the agent is *completely unknown* to the principal. Because in this new setting no guarantee is possible, Carroll relaxes the new model by adding the assumption that *some* set of actions \mathcal{A}_0 is fully known to the principal (that is, she is fully aware of the distributions the actions in \mathcal{A}_0 induce over outcomes as well as of their costs). Carroll discusses two possible ways in which to “make sense of this combination of non-quantifiable and quantifiable

⁴A similar issue arises in auction theory: linear programs can be used to characterize optimal auctions, which often turn out to be impractically complicated and unintuitive (see, e.g., [23]).

⁵To our knowledge, the only other common contract form according to the economics literature is “debt contracts,” which are similar to linear contracts except with zero payments for a set of the lowest outcomes [24]. Our results do not qualitatively change for such contracts—see Section 6.

uncertainty” (i.e., completely unknown and known actions—see p. 546). The main result is that a linear contract is max-min optimal in the worst case over all possible sets of actions \mathcal{A} that contain the known set \mathcal{A}_0 (where $\mathcal{A} \setminus \mathcal{A}_0$ can be anything). Carroll sees the main contribution not necessarily in a literal interpretation of the model, but rather in finding “a formalization of a robustness property of linear contracts—a way in which one can make guarantees about the principal’s payoff with very little information” [15, Sec. 2.1].

1.1 Our Contributions

Our main goal is *to initiate the study of simple contracts and their guarantees through the lens of theoretical computer science*. We utilize tools and ideas from worst-case algorithmic analysis and prior-independent auction design to make contributions in two main directions, both justifying and informing the use of linear contracts.

Max-min robustness of linear contracts. Our first contribution is a new notion of robustness for linear contracts. Instead of assuming that there are completely unknown actions available to the agent alongside a fully known action set, we assume that the principal has partial knowledge of all actions—she knows their costs and has *moment information* about their reward distributions. Assuming moment information is natural from an optimization perspective (the idea can at least be traced back to Scarf’s seminal 1958 paper on distributionally-robust stochastic programming [37]), and it’s a standard assumption in prior-independent auction design (see, e.g., [2, 7, 12]). The assumption is particularly natural in the context of contracts and the principal-agent model, where it is the minimal assumption that allows for optimization over linear contracts. With this assumption, the optimal linear contract is a natural contract to focus on, and our analysis shows that it is in fact *max-min optimal* across all possible contracts (Theorem 4.1 in Section 4).⁶ Our result thus offers an alternative formulation of the robustness property of linear contracts, in a natural model of moment information that is easy to interpret.

THEOREM (SEE SECTION 4). *For every outcome set, action set, action costs, and expected action rewards, a linear contract maximizes the principal’s worst-case expected payoff, where the worst case is over all reward distributions with the given expectations.*

Approximation guarantees. Our second contribution is to conduct *the first study of simple contracts from an approximation perspective*. Studying the worst-case approximation guarantees of classic microeconomic mechanisms—linear contracts in this case—has been a fruitful approach in other areas of algorithmic game theory. Applying this approach, we achieve a *complete and tight* analysis of the approximation landscape for linear contracts. Our analysis shows that *linear contracts are approximately optimal except in pathological cases, which have many actions, a big spread of expected rewards among actions, and a big spread of costs among actions, simultaneously*. Concretely, in the pathological cases we construct for our lower bounds (see e.g. Theorem 5.2 in Section 5), the expected rewards of actions a_1, a_2, a_3, \dots are $1, 1/\epsilon, 1/\epsilon^2, \dots$ for vanishing ϵ , and the costs are likewise “exploding” and tailored such that the difference between the exponentially large expected reward and exponentially large cost of the i th action is roughly i . The following summarizes our findings more formally:

THEOREM (SEE SECTION 5). *Let ρ denote the worst-case ratio between the expected principal payoff under an optimal contract and under the best linear contract. Then*

⁶Our max-min optimality result is conceptually related to work by Carroll [14], Gravin and Lu [21], and Bei et al. [9]. These papers show the max-min optimality of selling items separately when the marginals of all items are held fixed but the joint distribution is unknown (i.e., adversarial). Knowing the marginals is precisely what is needed for optimizing over mechanisms that sell each item separately.

- (a) Among principal-agent settings with n actions, $\rho = n$.
- (b) Among settings where the ratio of the highest and lowest expected rewards is H , $\rho = \Theta(\log H)$.
- (c) Among settings where the ratio of the highest and lowest action costs is C , $\rho = \Theta(\log C)$.
- (d) Among settings with $m \geq 3$ outcomes, ρ can be arbitrarily large in the worst case.

The upper bounds summarized in the above theorem hold even with respect to the strongest-possible benchmark of the optimal expected welfare (rather than merely the optimal principal expected payoff); they thus answer the natural question of how much of the “first-best” welfare a linear contract can extract. The matching lower bounds in the above theorem all apply even when we add a strong regularity condition to the principal-agent settings, known in the literature as the *monotone likelihood ratio property* (MLRP). In Section 6, we show an extension of our lower bounds to all *monotone* (not necessarily linear) contracts: we prove that among principal-agent settings with n actions, the ratio between the expected principal payoff under an optimal contract and under the best *monotone* contract can be $n - 1$ in the (pathological) worst case.

Discussion. We view parts (a)–(c) of the above theorem on approximation guarantees as surprisingly positive results. A priori, it is completely unclear which of the many parameters of principal-agent settings, if any, governs the performance of simple contracts. Our results show that there is no finite approximation bound that holds uniformly over *all* of the model parameters, but that we can obtain the next best thing: by fixing *just one* of the model’s ingredients (either the number of actions, the range of the outcomes, or the range of the costs, as preferred), it is possible to obtain an approximation guarantee that holds uniformly over all other parameters. Our theorem shows that linear contracts are far from optimal only when the number of actions is large, *and* there is a huge spread in expected rewards, *and* there is a huge spread of action costs. Few if any of the most popular instantiations of the principal-agent model have all three of these properties.

Organization. Section 2 introduces the model. Section 3 establishes several basic but fundamental properties of linear contracts. Sections 4 and 5 contain our main results on their performance, and Section 6 goes beyond linear to monotone contracts. Section 7 concludes.

1.2 Further Related Work

Contract theory is one of the pillars of microeconomic theory. We refer the interested reader to the classic papers of [22, 27, 36, 38], the coverage in [32, Chapters 13–14], the excellent introductions of [11, 40], and the comprehensive textbooks [10, 31].

Computational approaches to contract design. To our knowledge, decidedly computational approaches to contract design have appeared so far only in the work of [5] (and follow-ups [3, 4]), the work of [6], and the work of [25] (and follow-up [16]). The first paper [5] initiates the study of a related but different model known as *combinatorial agency*, in which combinations of agents replace the single agent in the classic principal-agent model. The challenge in the new model stems from the need to incentivize multiple agents, while the action structure of each agent is kept simple (effort/no effort). The focus of this line of work is on complex combinations of agents’ efforts influencing the outcomes, and how these determine the subsets of agents to contract with. The second paper [6] introduces a notion of contract complexity based on the number of different payments specified in the contract, and studies this complexity measure in an n -player normal-form game framework. In their framework there are no hidden actions, making our model very different from theirs. The third paper [25] develops a model of *dynamic* contract design: in each sequential round, the principal determines a contract, an agent arrives and chooses an action (effort level), and the principal receives a reward. Agents are drawn from an unknown prior distribution that dictates their available actions. The problem thus reduces to a multi-armed bandit variant with

each arm representing a potential contract. The main focus of this line of work is on implicitly learning the underlying agent distribution to minimize the principal's regret over time.

(Non)relation to signaling. Since one of the main features of the principal-agent model is the information asymmetry regarding the chosen action (the agent knows while the principal is oblivious), and due to the “principal” and “agent” terminology, on a superficial level contract theory may seem closely related to signaling [33, 34, 39]. This is not the case, and the relationship is no closer than that between auction theory and signaling. As Dughmi [18] explains, the heart of signaling is in creating the right information structure, whereas the heart of contract design is in setting the right payment scheme.⁷ Put differently, in signaling, it is the more-informed party that faces an economic design problem; in hidden-action contract theory, it is the less-informed party (i.e., the principal). For more on signaling from a computational perspective the reader is referred to [19, 20, 35].

Concurrent work on principal-agent settings with no money. Two recent papers [28, 30] study algorithmic aspects of another loosely related but distinct problem called *optimal delegation* [1]. In this problem, a principal has to search for and decide upon a solution, and wishes to delegate the search to an agent with misaligned incentives regarding which solution to choose. Crucially, there are *no monetary transfers*, making the problem quite different from ours. A third recent paper [29] studies a principal and agent setting where the agent has a “budget” of effort to spread across different actions (rather than a cost per level of effort), and any effort spent within the budget does not lower the agent's utility. There are again no monetary transfers from the principal to the agent, and the principal incentivizes the agent by mapping his outcomes to academic grades or other classifications.

2 THE HIDDEN ACTION PRINCIPAL-AGENT MODEL

The model. An instance of the principal-agent model is described by a pair (A_n, Ω_m) of n actions and m outcomes. We identify the j th outcome for every $j \in [m]$ with its reward x_j to the principal, and assume w.l.o.g. that the outcomes are increasing, i.e., $0 \leq x_1 \leq x_2 \leq \dots \leq x_m$.

Each action is a pair $a = (F_a, c_a)$, in which F_a is a distribution over the m outcomes and $c_a \geq 0$ is a cost. Denote by $F_{a,j}$ the probability of action a to lead to outcome x_j ; we assume w.l.o.g. that every outcome has some action leading to it with positive probability. The agent chooses an action $a \in A_n$ and bears the cost c_a , whereas the principal receives a random reward x_j drawn from F_a . Crucially, *the chosen action is hidden*: the principal observes the outcome x_j but not the action a .

Denote the *expected outcome* (i.e., reward to the principal) from action a by $R_a = \mathbb{E}_{j \sim F_a}[x_j] = \sum_{j \in [m]} F_{a,j} x_j$. The difference $R_a - c_a$ is the expected *welfare* from choosing action a . When an action a_i is indexed by i , we write for brevity $R_i, F_{i,j}, c_i$ (rather than $R_{a_i}, F_{a_i,j}, c_{a_i}$).

Standard assumptions. Unless stated otherwise we assume:

- A1** There are no “dominated” actions, i.e., every two actions a, a' have distinct expected outcomes $R_a \neq R_{a'}$, and the action with the higher expected outcome $R_a > R_{a'}$ also has higher cost $c_a > c_{a'}$.
- A2** There is a unique action a with maximum welfare $R_a - c_a$.
- A3** There is a zero-cost action a with $c_a = 0$.

⁷“There are two primary ways of influencing the behavior of self-interested agents: by providing incentives, or by influencing beliefs. The former is the domain of traditional mechanism design, and involves the promise of tangible rewards such as [...] money. The latter [...] involves the selective provision of payoff-relevant information to agents through strategic communication” [18, p. 1].

Assumption A1 means there is no action with lower expected outcome and higher cost than some other action, although we emphasize that there *can* be an action with lower *welfare* and higher cost (in fact, incentivizing the agent to avoid such actions is a source of contract complexity). Our main results in Sections 4-5 do not require this assumption (see Section 5 for details). Assumptions A2 and A3 are for the sake of expositional simplicity. In particular, Assumption A3 means we can assume the agent does not reject a contract with non-negative payments, since there is always an individually rational choice of action; alternatively, individual rationality could have been imposed directly.

Contracts. A contract defines a *payment scheme* t with a payment (transfer) $t_j \geq 0$ from the principal to the agent for every outcome x_j . We denote by T_a the expected payment $\mathbb{E}_{j \sim F_a}[t_j] = \sum_j F_{a,j} t_j$ for action a , and by T_i the expected payment for a_i . Note that the payments are contingent only on the outcomes as the actions are not observable to the principal. The requirement that t_j is non-negative for every j is referred to in the literature as *limited liability* [13], and it plays the same role as the standard *risk averseness* assumption in ruling out trivial solutions where a contract is not actually required [22]. Limited liability (or its parallel agent risk averseness) is the second crucial feature of the classic principal-agent model, in addition to the actions being hidden from the principal.

Implementable actions. The agent's expected *utility* from action a given payment scheme t is $T_a - c_a$. The agent chooses an action that is: (i) *incentive compatible (IC)*, i.e., maximizes his expected utility among all actions in A_n ; and (ii) *individually rational (IR)*, i.e., has non-negative expected utility (if there is no IR action the agent refuses the contract). We adopt the standard assumption that the agent tie-breaks among IC, IR actions in favor of the principal.⁸ We say a contract *implements* action a^* if given its payment scheme t , the agent chooses a^* ; if there exists such a contract we say a^* is *implementable*.

Optimal contracts and LPs. The principal seeks an *optimal* contract: a payment scheme t that maximizes her expected *payoff* $R_a - T_a$, where a is the action implemented by the contract (i.e., a is both IC and IR for the agent, with ties broken to maximize the expected payoff of the principal). Notice that summing up the agent's expected utility $T_a - c_a$ with the principal's expected payoff $R_a - T_a$ results in the contract's expected *welfare* $R_a - c_a$. A contract's payment scheme thus determines both the size of the pie (expected welfare), and how it is split between the principal and the agent.

An optimal contract can be found through linear programming: For each action determine the minimum expected payment at which this action can be implemented. The best of these linear programs (achieving the highest expected payoff for the principal) gives the best action to incentivize and an optimal contract incentivizing it.

The LP for incentivizing action a at minimum expected payment has m payment variables $\{t_j\}$, which by limited liability must be non-negative, and $n - 1$ IC constraints ensuring that the agent's expected utility from action a is at least his expected utility from any other action. Note that by Assumption A3, there is no need for an IR constraint to ensure that the expected utility is non-negative. The LP is:

$$\begin{aligned} \min \quad & \sum_{j \in [m]} F_{a,j} t_j \\ \text{s.t.} \quad & \sum_{j \in [m]} F_{a,j} t_j - c_a \geq \sum_{j \in [m]} F_{a',j} t_j - c_{a'} \quad \forall a' \neq a, a' \in A_n, \\ & t_j \geq 0 \quad \forall j \in [m]. \end{aligned}$$

⁸The idea is that one could perturb the payment schedule slightly to make the desired action uniquely optimal for the agent. For further discussion see [11, p. 8].

Fairly little is known about the *structure* of the optimal contracts that come out of this approach. There are a few notable exemptions from this general rule—e.g., settings with just $n = 2$ actions—which we discuss in more detail in the full version of this paper.

This LP and close variants together with their duals can also be used to characterize implementable actions, where the implementing contract is either arbitrary or required to be *monotone* as defined below. The characterizations hinge on whether the action's distribution can be emulated or dominated by a combination of other actions at lower cost. See full version for details.

Linear/monotone contracts. In a *linear* contract the payment scheme is a linear function of the outcomes, i.e., $t_j = \alpha x_j \geq 0$ for every $j \in [m]$. We refer to α as the linear contract's *parameter*, which is ≥ 0 due to limited liability. A natural generalization is a *degree- d polynomial* contract, in which the payment scheme is a non-negative degree- d polynomial function of the outcomes: $t_j = \sum_{k=0}^d \alpha_k x_j^k \geq 0$ for every $j \in [m]$. If $d = 1$ we get an *affine* contract; such contracts play a role in Section 4. Linear and affine contracts are *monotone*, where a contract is monotone if its payments are nondecreasing in the outcomes, i.e., $t_j \leq t_{j'}$ for $j < j'$.

Regularity assumptions. The economic literature (see, e.g., [22]) introduces a regularity assumption called the *monotone likelihood ratio property* (MLRP) for principal-agent settings. Intuitively, the assumption asserts that the higher the outcome, the more likely it is to be produced by a high-cost action than a low-cost one (in a relative sense). We adapt the standard definition to accommodate for zero probabilities, as follows:

Definition 2.1 (MLR). Let F, G be two distributions over m values v_1, \dots, v_m . The likelihood ratio F_j/G_j is monotonically increasing in j if

$$F_j/G_j \leq F_{j'}/G_{j'}$$

for every $j < j'$ such that at least one of F_j, G_j is positive, and at least one of $F_{j'}, G_{j'}$ is positive.

Definition 2.2 (MLRP). A principal-agent problem satisfies *MLRP* if for every pair of actions a, a' such that $c_a < c_{a'}$, the likelihood ratio $F_{a',j}/F_{a,j}$ is monotonically increasing in j .

PROPOSITION 2.3 (MLR \implies FOSD [40]). *If the likelihood ratio F_j/G_j is monotonically increasing in j , then F first-order stochastically dominates G . The converse does not hold.*

Max-min evaluation and approximation. We apply two approaches to evaluate the performance of simple contracts: max-min in Section 4, and approximation in Section 5. We now present the necessary definitions, starting with the max-min approach.

A *distribution-ambiguous* action is a pair $a = (R_a, c_a)$, in which $R_a \geq 0$ is the action's expected outcome and $c_a \geq 0$ is its cost. Distribution F_a over outcomes $\{x\}$ is *compatible* with distribution-ambiguous action a if $\mathbb{E}_{x \sim F_a}[x] = R_a$. A principal-agent setting (A_n, Ω_m) is *ambiguous* if it has $m \geq 3$ outcomes and n distribution-ambiguous actions, and there exist distributions F_1, \dots, F_n over the outcomes compatible with the actions.⁹ In ambiguous settings, it is appropriate to apply a worst-case performance measure to evaluate contracts: Given an ambiguous principal-agent setting, a contract's *worst-case expected payoff* is its infimum expected payoff to the principal over all distributions $\{F_i\}_{i=1}^n$ compatible with the known expected outcomes $\{R_i\}_{i=1}^n$. We follow [13] in making the following assumption, which simplifies the results in Section 4 but does not qualitatively affect them:¹⁰

⁹A setting with $m = 2$ outcomes cannot be ambiguous since the expectation determines the distribution; moreover the conundrum of “optimal but complex” vs. “suboptimal but ubiquitous” never arises as the optimal contract has a simple form—see the full version of this paper.

¹⁰As explained in [13, Footnote 2], Assumption A4 is simply an additive normalization of the principal's payoffs. Without this assumption, a robustly optimal contract would take the form $t_j = \alpha(x_j - x_1)$. Further justification for assuming $x_1 = 0$

A4 In ambiguous principal-agent settings, the outcome 0 belongs to Ω_m , i.e., $x_1 = 0$.

In Section 5, we are interested in bounding the potential loss in the principal's expected payoff if she is restricted to use a linear contract. Formally, let \mathcal{A} be the family of principal-agent settings. For $(A_n, \Omega_m) \in \mathcal{A}$, denote by $OPT(A_n, \Omega_m)$ the optimal expected payoff to the principal with an arbitrary contract, and by $ALG(A_n, \Omega_m)$ the best possible expected payoff with a contract of the restricted form (we omit (A_n, Ω_m) from the notation where clear from context). We seek to bound $\rho(\mathcal{A}) := \max_{(A_n, \Omega_m) \in \mathcal{A}} \frac{OPT(A_n, \Omega_m)}{ALG(A_n, \Omega_m)}$.

3 PROPERTIES AND GEOMETRY OF LINEARLY-IMPLEMENTABLE ACTIONS

We begin by developing a geometric characterization of *linearly-implementable* actions, and deriving several useful consequences, which we apply in Section 4 and especially in Section 5. See the full version of this paper for details and missing proofs.

Consider a principal-agent setting (A_n, Ω_m) . We say that an action $a \in A_n$ is *linearly-implementable* if there exists a linear contract with parameter $\alpha \leq 1$ that implements a . Note that $\alpha \leq 1$ is w.l.o.g. because the expected payoff to the principal would be negative otherwise.

OBSERVATION 1. *If a linear contract with parameter α implements action a^* , then the agent's expected utility and the principal's expected payoff are, respectively,*

$$\alpha R_{a^*} - c_{a^*}; (1 - \alpha)R_{a^*}.$$

OBSERVATION 2. *Under Assumptions A1-A3, a linear contract implements at most one action.*

OBSERVATION 3. *Let a, a' be a pair of actions such that $R_{a'} > R_a$ and $c_{a'} > c_a$. Then a linear contract with parameter $\alpha_{a,a'} = \frac{c_{a'} - c_a}{R_{a'} - R_a}$ makes the agent indifferent among actions a' and a (but does not necessarily incentivize either of these actions).*

Let N denote the number of linearly-implementable actions, and let $I_N \subseteq A_n$ denote the set of such actions. Index the actions in I_N in order of their expected outcomes, i.e., for $a_i, a_{i'} \in I_N$, $i < i' \iff R_i < R_{i'}$. We now define two different mappings, and in Lemma 3.2 establish their equivalence.

(1) *Linear-implementability mapping $a(\cdot)$.* Denote by $a(\cdot) : [0, 1] \rightarrow I_N \cup \{\emptyset\}$ the mapping of α to either the action implemented by the linear contract with parameter α (observe there is at most one such action under our assumptions), or to \emptyset if there is no such action. So mapping $a(\cdot)$ is onto I_N by definition. Denote by α_i the smallest $\alpha \in [0, 1]$ such that action $a_i \in I_N$ is implemented by a linear contract with parameter α , then α_i is the smallest α such that $a(\alpha) = a_i$.

(2) *Upper envelope mapping $u(\cdot)$.* For every action $a \in A_n$, consider the line $\alpha R_a - c_a$ and let ℓ_a denote the segment between $\alpha = 0$ and $\alpha = 1$; these segments appear in Figure 1, where the x -axis represents the possible values of α from 0 to 1. Take the upper envelope of the n segments $\{\ell_a\}_{a \in A_n}$ and consider its non-negative portion. Let $u(\cdot) : [0, 1] \rightarrow A_n \cup \{\emptyset\}$ be the mapping from α to either \emptyset if the upper envelope is negative at α , or to the action whose segment forms the upper envelope at α otherwise. If there is more than one such action, let $u(\alpha)$ map to the one with the highest expected outcome R_a .

We begin by showing the following monotonicity property of the upper envelope mapping u . Let $\underline{\alpha}$ denote the smallest α at which the upper envelope intersects the x -axis (or $\underline{\alpha} = 1$ if no such α exists). As α goes from $\underline{\alpha}$ to 1, $u(\alpha)$ maps to actions with increasingly higher expected outcomes $\{R_a\}$, costs $\{c_a\}$ and expected welfares $\{R_a - c_a\}$.

is that the principal may have ambiguity not just with respect to the action distributions but also as to her possible rewards, and she prefers a contract robust to the possibility (however slim) of receiving a zero reward.




Fig. 1. Linearly-implementable actions via upper envelope.

LEMMA 3.1. For every two parameters $0 \leq \alpha < \alpha' \leq 1$, either $u(\alpha) = \emptyset$, or it holds that $u(\alpha) = a \in A_n$ and $u(\alpha') = a' \in A_n$ such that (i) $R_a < R_{a'}$; (ii) $c_a < c_{a'}$; and (iii) $R_a - c_a \leq R_{a'} - c_{a'}$.

Our main structural insight in this section is that the upper envelope mapping precisely captures linear implementability:

LEMMA 3.2. For every $\alpha \in [0, 1]$, $a(\alpha) = u(\alpha)$.

Lemma 3.2 has three useful implications: (1) The actions $\{a_i \in I_N\}$ appear on the upper envelope in the order in which they are indexed (i.e., sorted by increasing expected outcome). (2) These actions are also sorted by increasing welfare, i.e., $R_1 - c_1 \leq R_2 - c_2 \leq \dots \leq R_N - c_N$. (3) The smallest α that incentivizes action a_i (which we refer to as α_i) is the same α that makes the agent indifferent between action a_i and action a_{i-1} . We denote this “indifference α ” by $\alpha_{i-1,i}$ and observe that $\alpha_{i-1,i} = \frac{c_i - c_{i-1}}{R_i - R_{i-1}}$ from Observation 3. Using this notation, we can rewrite the third implication as $\alpha_i = \alpha_{i-1,i}$. Formally:

COROLLARY 3.3. For every $\alpha \in [0, 1]$,

$$\begin{aligned} \forall i \in [N-1] : a(\alpha) = a_i &\iff \alpha_i \leq \alpha < \alpha_{i+1}, \\ a(\alpha) = a_N &\iff \alpha_N \leq \alpha \leq 1. \end{aligned} \quad (1)$$

COROLLARY 3.4. For every $i \in [N-1]$, $R_i - c_i \leq R_{i+1} - c_{i+1}$.

COROLLARY 3.5. For every $i \in [N]$, $\alpha_i = \alpha_{i-1,i}$.

We apply Corollaries 3.3-3.5 to derive our approximation bounds in Section 5.

The final corollary in this section, which we apply to establish robustness in Section 4, highlights the fact that in terms of linear-implementability, two principal-agent settings whose actions have the same expected outcomes and costs are in effect equivalent. The distributions, outcome values, and even number of outcomes matter for linear-implementability only to the extent of determining the expected outcome of each action. This property is special to linear contracts; in contrast, optimal contracts depend on the details of the distributions beyond their expected outcomes, and this adds to their complexity.

COROLLARY 3.6. Consider two principal-agent settings (A_n, Ω_m) , (A'_n, Ω'_m) for which there exists a bijection $b : A_n \rightarrow A'_n$ between the action sets, such that actions a and $b(a)$ have the same expected outcome $R_a = R_{b(a)}$ and cost $c_a = c_{b(a)}$ for every $a \in A_n$. Let a, a' be the linearly-implementability mappings of the two settings, respectively. Then for every parameter $\alpha \in [0, 1]$, $b(a(\alpha)) = a'(\alpha)$, and the principal's expected payoff from a linear contract with parameter α is the same in both settings.




Fig. 2. **Visualization of the proof of Lemma 4.2—Case I.** In the case depicted here, point (x_j, t_j) is strictly above line l_1 . This case is addressed in Claim 3.

4 ROBUST OPTIMALITY OF LINEAR CONTRACTS

In this section we establish our robust optimality result for linear contracts. We show that a linear contract maximizes the principal's expected payoff in ambiguous settings, *in the worst case over the unknown distributions*. All deferred proofs appear in the full version of this paper.

THEOREM 4.1 (ROBUST OPTIMALITY). *For every ambiguous principal-agent setting, an optimal linear contract has maximum worst-case expected payoff among all limited liability contracts.*

In Theorem 4.1, “an optimal linear contract” is well-defined: For a (non-ambiguous) principal-agent setting, a linear contract is optimal when it maximizes the principal's expected payoff over all linear contracts. For an *ambiguous* principal-agent setting, a linear contract has the same expected payoff to the principal over all compatible distributions—see Corollary 3.6. Thus an optimal linear contract can still be defined as maximizing the principal's expected payoff. In the remainder of the section we prove Theorem 4.1.

4.1 Main Lemma for Robust Optimality

The key step in the proof of Theorem 4.1 is to show that we may restrict the search for optimally robust contracts to affine contracts.

LEMMA 4.2. *Consider an ambiguous principal-agent setting (A_n, Ω_m) . For every limited liability contract with payment scheme t , there exist compatible distributions $\{F_i\}_{i=1}^n$ and an affine contract with parameters $\alpha_0, \alpha_1 \geq 0$, such that the affine contract's expected payoff is at least that of contract t for distributions $\{F_i\}_{i=1}^n$.*

PROOF. The payment scheme t maps the outcomes $0 = x_1 < \dots < x_m$ to payments $t_1, \dots, t_m \geq 0$. Consider the two extreme outcomes x_1, x_m and their corresponding payments t_1, t_m . We begin by defining simple compatible distributions $\{F'_i\}_{i=1}^n$ whose support is the extreme outcomes, as follows. For every distribution-ambiguous action a_i , set $F'_{i,m} := \frac{R_i}{x_m}$ (this is a valid probability since $R_i \leq x_m$; otherwise there could not have been compatible distributions). Set $F'_{i,1} := 1 - F'_{i,m}$ and let the other probabilities be zero. The expected outcome of distribution F'_i is $F'_{i,1}x_1 + F'_{i,m}x_m = R_i$. The defined distributions already enable us to prove the lemma for the case of $t_1 > t_m$:

CLAIM 1. *Lemma 4.2 holds for the case of $t_1 > t_m$.*

PROOF OF CLAIM 1. A proof of this claim appears in the full version. △

Assume from now on that $t_1 \leq t_m$. If t is affine, this means that its slope parameter α_1 must be non-negative. Similarly, we can write $t_1 = \alpha_0 + \alpha_1 x_1$ and plug in our assumption that $x_1 = 0$ to get




Fig. 3. **Visualization of the proof of Lemma 4.2—Case II.** In the case depicted here, point (x_j, t_j) is strictly below line l_1 . This case is addressed in Claim 4.

$t_1 = \alpha_0$, and so by limited liability ($t \geq 0$), α_0 must also be non-negative. Thus if t is affine, Lemma 4.2 holds. We focus from now on on the case that t is non-affine; this guarantees the existence of a point (x_j, t_j) as appears in Figures 2 and 3. We argue this formally and then proceed by case analysis.

CLAIM 2. *If t is non-affine, there exists an index $1 < j < m$ such that the 3 points (x_1, t_1) , (x_j, t_j) and (x_m, t_m) on the Euclidean plane are non-collinear.*

PROOF OF CLAIM 2. A proof of this claim appears in the full version. \triangle

We introduce the following notation – denote the line between (x_1, t_1) and (x_m, t_m) by l_1 , the line between (x_1, t_1) and (x_j, t_j) by l_2 , and the line between (x_j, t_j) and (x_m, t_m) by l_3 (see Figures 2 and 3). We denote the parameters of line l_1 by α_0 and α_1 (i.e., $t_1 = \alpha_0 + \alpha_1 x_1$ and $t_m = \alpha_0 + \alpha_1 x_m$). These naturally give rise to a corresponding affine contract. As argued above, since $t_1 \leq t_m$ we have $\alpha_1 \geq 0$, and since $x_1 = 0$ and $t_1 \geq 0$ we have $\alpha_0 \geq 0$, so the affine contract has non-negative parameters.

Recall that the support of compatible distributions $\{F'_i\}_{i=1}^n$ is the endpoints of l_1 . We define alternative compatible distributions $\{F''_i\}_{i=1}^n$, whose support is either the endpoints of l_2 or of l_3 , as follows: For every distribution-ambiguous action a_i , if $R_i \leq x_j$ set $F''_{i,j} := \frac{R_i}{x_j}$ (by assumption this is a valid probability), and $F''_{i,1} := 1 - F''_{i,j}$. If $R_i > x_j$ set $F''_{i,m} := \frac{R_i - x_j}{x_m - x_j}$ and $F''_{i,j} := 1 - F''_{i,m}$. All other probabilities are set to zero. Observe that in either case, the expected outcome of distribution F''_i is R_i . With the two sets of compatible distributions $\{F'_i\}_{i=1}^n$ and $\{F''_i\}_{i=1}^n$ at hand, the analysis proceeds by addressing separately the two cases depicted in Figures 2 and 3.

The following simple observation will be useful in the case analysis:

OBSERVATION 4. *Consider a contract t and two outcomes x, x' ; let $l : \mathbb{R} \rightarrow \mathbb{R}$ be the line determined by points $(x, t(x))$ and $(x', t(x'))$. If action a_i induces an outcome distribution F_i over the support $\{x, x'\}$ with expectation R_i , then its expected payment to the agent is $l(R_i)$.*

CLAIM 3. *Lemma 4.2 holds for the case that (x_j, t_j) is strictly above l_1 .*

PROOF OF CLAIM 3. We prove the claim by showing that the affine contract with parameters α_0, α_1 (corresponding to l_1) has at least as high expected payoff to the principal as that of contract t for distributions $\{F_i\}_{i=1}^n$, which are defined as follows: Let a_{i^*} be the action incentivized by the affine contract with parameters α_0, α_1 . Set $F_{i^*} = F''_{i^*}$ and $F_k = F'_k$ for every $k \neq i^*$. We argue that for distributions $\{F_i\}_{i=1}^n$ as defined, contract t will also incentivize action a_{i^*} , but at a (weakly) higher expected payment to the agent compared to the affine contract.

Consider first the affine contract. By definition, for every action a_i with expected outcome R_i , the expected payment to the agent for choosing a_i is $\alpha_0 + \alpha_1 R_i$. Now consider contract t . For every action a_k where $k \neq i^*$, by Observation 4 the expected payment to the agent taken over F'_k is $l_1(R_k) = \alpha_0 + \alpha_1 R_k$, i.e., the same as in the affine contract. For action a_{i^*} , by Observation 4 the expected payment to the agent taken over F''_k is either $l_2(R_{i^*})$ or $l_3(R_{i^*})$. In either case, since l_2 and l_3 are above l_1 , the expected payment (weakly) exceeds $l_1(R_{i^*})$. Since a_{i^*} maximizes the agent's expected utility in the affine contract, and the expected payment for a_{i^*} only increases in contract t while staying the same for other actions, this completes the proof of Claim 3. \triangle

CLAIM 4. *Lemma 4.2 holds for the case that (x_j, t_j) is strictly below l_1 .*

PROOF OF CLAIM 4. A proof of this claim appears in the full version. \triangle

This completes the analysis of the cases depicted in Figures 2 and 3, thus proving Lemma 4.2. \square

4.2 Proof of Theorem 4.1

Having established Lemma 4.2, Theorem 4.1 is now easy to prove using the following observation:

OBSERVATION 5. *Consider an affine contract with parameters $\alpha_0, \alpha_1 \geq 0$. For any distributions $\{F_i\}_{i=1}^n$, the expected payoff to the principal from the affine contract is at most the expected payoff from the linear contract with parameter $\alpha = \alpha_1$.*

COROLLARY 4.3. *For every affine contract with parameters $\alpha_0, \alpha_1 \geq 0$, there is a linear contract with (weakly) higher worst-case expected payoff.*

PROOF OF THEOREM 4.1. Consider an ambiguous principal-agent setting. For every limited liability contract t , by Lemma 4.2 there exist compatible distributions $\{F_i\}_{i=1}^n$ and an affine contract with parameters $\alpha_0, \alpha_1 \geq 0$ such that the worst-case expected payoff of contract t is at most the expected payoff of the affine contract for distributions $\{F_i\}_{i=1}^n$. But the expected payoff of an affine contract is identical for all compatible distributions, and so for every limited liability contract t there exists an affine contract with parameters $\alpha_0, \alpha_1 \geq 0$ and higher worst-case expected payoff. By Corollary 4.3, the optimal linear contract has even higher worst-case expected payoff, completing the proof. \square

5 APPROXIMATION GUARANTEES OF LINEAR CONTRACTS

In this section we study linear contracts and their approximation guarantees. We show tight bounds on the approximation guarantees of linear contracts in all relevant parameters of the model—number of actions and outcomes (Section 5.1), spread of rewards (Section 5.2) and costs (Section 5.3). Our upper bounds in fact apply to the stronger benchmark of optimal welfare. Our lower bounds continue to hold if the instances are required to satisfy MLRP. See the full version of this paper for all proofs omitted from this section.

5.1 Bounds in the Number of Actions

Our first pair of results provides tight bounds on the approximation guarantees of linear contracts, as parameterized by the number of actions n .

THEOREM 5.1. *Consider a principal-agent setting (A_n, Ω_m) with n actions and $N \leq n$ linearly implementable actions. Then the multiplicative loss in the principal's expected payoff from using a linear contract rather than an arbitrary one is at most N .*

THEOREM 5.2. *For every n and $\varepsilon > 0$, there is a principal-agent setting (A_n, Ω_m) with n actions and $m = n$ outcomes for which MLRP holds, such that the multiplicative loss in the principal's expected payoff from using a linear contract rather than an arbitrary one is at least $n - \varepsilon$.*

The optimal contract in the construction used to prove Theorem 5.2 is monotone. It follows that the gap between the best linear contract and the optimal monotone contract can be n . In the full version we strengthen Theorem 5.2 by showing that it applies even if $m = 3$, thus implying that the approximation ratio as parametrized by the number of outcomes m can be unbounded.

Proof of upper bound. Our proof of Theorem 5.1 exploits the geometric insights developed in Section 3. The key tools are the following observation (Observation 6) and the two lemmas (Lemma 5.3 and Lemma 5.4) below.

Recall that I_N denotes the set of $N \leq n$ linearly-implementable actions, indexed such that their expected outcomes are increasing, i.e., $I_N = \{a_1, \dots, a_N\}$ and $R_1 < \dots < R_N$. Note that by Assumption A1, $c_1 < \dots < c_N$, and recall that this does *not* imply that $R_1 - c_1 \leq R_2 - c_2 \leq \dots \leq R_N - c_N$.

OBSERVATION 6. Consider two actions a, a' such that a has higher expected outcome and (weakly) higher welfare, i.e., $R_a > R_{a'}$ and $R_a - c_a \geq R_{a'} - c_{a'}$. Let $\alpha_{a',a} = \frac{c_a - c_{a'}}{R_a - R_{a'}}$. Then

$$(R_a - c_a) - (R_{a'} - c_{a'}) \leq (1 - \alpha_{a',a})R_a. \quad (2)$$

PROOF. Since $R_a - c_a \geq R_{a'} - c_{a'}$ we have $R_a - R_{a'} \geq c_a - c_{a'}$. Using that $R_a - R_{a'} > 0$ we get $\alpha_{a',a} = \frac{c_a - c_{a'}}{R_a - R_{a'}} \leq 1$. So we can write $R_{a'} - c_{a'} \geq \alpha_{a',a}R_{a'} - c_{a'} = \alpha_{a',a}R_a - c_a$, where the equality follows from our definition of $\alpha_{a',a}$. Hence, $(R_a - c_a) - (R_{a'} - c_{a'}) \leq (R_a - c_a) - (\alpha_{a',a}R_a - c_a) = (1 - \alpha_{a',a})R_a$, as required. \square

Below we shall apply Observation 6 to actions $a_i, a_{i-1} \in I_N$. In this context, the intuition behind the observation is as follows: Consider the linear contract with parameter $\alpha_{i-1,i} \in [0, 1]$. By Observation 3, in this contract the agent is indifferent among actions a_i and a_{i-1} . The left-hand side of (2) is the increase in expected welfare by switching to action a_i from a_{i-1} . For the agent to get the same expected utility from a_i and a_{i-1} , the principal must get this entire welfare increase as part of her expected payoff. The right-hand side of (2) is the principal's expected payoff, and so the inequality holds.

LEMMA 5.3. For every $k \in [N]$ and linearly-implementable action $a_k \in I_N$, $R_k - c_k \leq \sum_{i=1}^k (1 - \alpha_{i-1,i})R_i$.

PROOF. The proof is by induction on k . For $k = 1$, recall that $\alpha_{0,1} = 0$ by definition, and it trivially holds that $R_1 - c_1 \leq R_1$. Now assume the inequality holds for $k - 1$, i.e.,

$$R_{k-1} - c_{k-1} \leq \sum_{i=1}^{k-1} (1 - \alpha_{i-1,i})R_i. \quad (3)$$

By Corollary 3.4, the welfare of a_k is at least that of a_{k-1} , and we know that $R_k > R_{k-1}$. We can thus apply Observation 6 to actions $a = a_k, a' = a_{k-1}$ and get

$$(R_k - c_k) - (R_{k-1} - c_{k-1}) \leq (1 - \alpha_{k-1,k})R_k. \quad (4)$$

Adding inequality (4) to (3) completes the proof for k . \square

LEMMA 5.4. Consider a principal-agent setting (A_n, Ω_m) with linearly-implementable action set $I_N \subseteq A_n$. The expected payoff of an optimal (not necessarily linear) contract is at most $R_N - c_N$.

PROOF. In a linear contract with parameter $\alpha = 1$, the agent's expected utility for any action a is its welfare $R_a - c_a$. Thus an action is implemented by such a contract if and only if it maximizes welfare among all actions A_n . By Corollary 3.3, $a(1) = a_N$ and so a_N must be the welfare-maximizing action. In every contract, the IR property ensures that the agent's expected payment covers the cost

c_a of the implemented action a , and so the principal's expected payoff is always upper-bounded by $R_a - c_a$. We conclude that $OPT \leq \max_{a \in A_n} \{R_a - c_a\} = R_N - c_N$, as required. \square

With Lemma 5.3 and Lemma 5.4 at hand we can now prove Theorem 5.1.

PROOF OF THEOREM 5.1. To prove the approximation guarantee of $N \leq n$ claimed in the theorem, observe that

$$OPT \leq R_N - c_N \leq \sum_{i \leq N} (1 - \alpha_{i-1,i}) R_i = \sum_{i \leq N} (1 - \alpha_i) R_i \leq N \cdot \max_{i \leq N} \{(1 - \alpha_i) R_i\} = N \cdot ALG, \quad (5)$$

where the first inequality holds by Lemma 5.4, the second inequality holds by Lemma 5.3, and the equality holds by Corollary 3.5. \square

Proof of lower bound. Our proof of Theorem 5.2 relies on a recursively defined principal-agent setting with full information, in which the optimal contract can extract full welfare of n while linear contracts can extract at most 1. The construction requires the expected outcome and the costs to tend to infinity at different speeds.

PROOF OF THEOREM 5.2. For every n , consider a family of principal-agent instances $\{(A_n^\epsilon, \Omega_n) \mid \epsilon > 0\}$, each with n actions and $m = n$ outcomes. For every $i \in [n]$, the i th action $a_i = (F_i, c_i) \in A_n^\epsilon$ has $F_{i,i} = 1$, i.e., deterministically leads to the i th outcome $x_i \in \Omega_n$. Every principal-agent instance is thus a full information setting in which the outcome indicates the action, and for which MLRP holds. We define action a_i 's expectation R_i (equal to outcome x_i) and its cost c_i recursively:

$$R_{i+1} = \frac{R_i}{\epsilon}, \quad c_{i+1} = c_i + (R_{i+1} - R_i) \left(1 - \frac{1}{R_{i+1}}\right),$$

where $R_1 = 1$ and $c_1 = 0$.

We establish several useful facts about instance (A_n^ϵ, Ω_n) : For every $i \in [n]$, it is not hard to verify by induction that

$$R_i = \frac{1}{\epsilon^{i-1}}, \quad c_i = \frac{1}{\epsilon^{i-1}} - i + \epsilon(i-1), \quad (6)$$

$$R_i - c_i = i - \epsilon(i-1). \quad (7)$$

Observe that the actions are ordered such that R_i , c_i , and $R_i - c_i$ are strictly increasing in i . Plugging the values in (6) into $\alpha_{i-1,i} = \frac{c_i - c_{i-1}}{R_i - R_{i-1}}$ we get

$$\begin{aligned} \alpha_{i-1,i} &= 1 - \epsilon^{i-1}, \\ (1 - \alpha_{i-1,i}) R_i &= 1. \end{aligned} \quad (8)$$

$\alpha_{i-1,i}$ is also strictly increasing in i .

Let OPT^ϵ (resp., ALG^ϵ) denote the optimal expected payoff from an arbitrary (resp., linear) contract in the principal-agent setting (A_n^ϵ, Ω_n) . In a full information setting, the principal can extract the maximum expected welfare. This can be achieved by paying only for the outcome that indicates the welfare-maximizing action, and only enough to cover its cost. From (7) we thus get

$$OPT^\epsilon = \max_i \{i - \epsilon(i-1)\} = n - \epsilon(n-1) \xrightarrow{\epsilon \rightarrow 0} n.$$

By analyzing linear-implementability in the setting (A_n^ϵ, Ω_n) one can show that $\alpha_i = \alpha_{i-1,i}$ for every $i \in [n]$ (see full version for details). Thus from (8) it follows that

$$ALG^\epsilon \leq 1,$$

completing the proof. \square

5.2 Bounds in the Range of Expected Rewards

Our second pair of results is parametrized by the range of the expected outcomes $\{R_i\}$ normalized such that $R_i \in [1, H]$ for all $a_i \in A_n$. Consider bucketing these actions by their expected outcomes into $\lceil \log H \rceil$ buckets with ranges $[1, 2), [2, 4), [4, 8)$, etc. Let K be the number of non-empty buckets.

THEOREM 5.5. *Consider a principal-agent setting (A_n, Ω_m) such that for every action $a \in A_n$, its expected outcome R_a is in $[1, H)$. The multiplicative loss in the principal's expected payoff from using a linear contract rather than an arbitrary one is at most $2K = O(\log H)$.*

COROLLARY 5.6. *For every range $[1, H)$, there is a principal-agent setting (A_n, Ω_m) with $n = \log H$ actions and $m = n$ outcomes for which MLRP holds and $\forall a \in A_n : R_a \in [1, H)$, such that the multiplicative loss in the principal's expected payoff from using a linear contract rather than an arbitrary one is at least $\Omega(\log H)$.*

Our proof of Theorem 5.5 applies an argument similar to the one used in Theorem 5.1 to the set of actions with the highest expected outcome in each bucket.

PROOF OF THEOREM 5.5. Recall the bucketing of actions in I_N by their expected outcome. For every bucket $k \leq K$, let $h(k)$ be the action with the highest R_i in the bucket, and let $l(k)$ be the action with the lowest R_i . The bucketing is such that $R_{h(k)}/2 < R_{l(k)} \leq R_{h(k)}$. Since the actions in I_N are ordered by their expected outcome, then $h(k)$ and $l(k)$ are increasing in k , and $h(k-1) + 1 = l(k) \leq h(k)$. For the last bucket K we have that $h(K) = a_N$, and by Lemma 5.4 this implies $OPT \leq R_{h(K)} - c_{h(K)}$.

Consider the subset of linearly-implementable actions $I_K = \{a_{h(k)} \mid k \in [K]\} \subseteq I_N$. These will play a similar role in our proof as actions I_N in the proof of Theorem 5.1. Let $\alpha_{h(k-1), h(k)} = (c_{h(k-1)} - c_{h(k)}) / (R_{h(k-1)} - R_{h(k)})$. Observation 6 applies to actions in I_K , and so we can apply a version of Lemma 5.3 to get

$$OPT \leq R_{h(K)} - c_{h(K)} \leq \sum_{k \leq K} (1 - \alpha_{h(k-1), h(k)}) R_{h(k)}. \quad (9)$$

Our goal is now to upper-bound the right-hand side of (9).

CLAIM 5. $\alpha_{h(k-1), h(k)} \geq \alpha_{l(k)}$.

PROOF OF CLAIM 5. Assume for contradiction that $\alpha_{h(k-1), h(k)} < \alpha_{l(k)}$. By definition of $\alpha_{h(k-1), h(k)}$ we have that $\alpha_{h(k-1), h(k)} R_{h(k)} - c_{h(k)} = \alpha_{h(k-1), h(k)} R_{h(k-1)} - c_{h(k-1)}$. Substituting $h(k-1) = l(k) - 1$ we get

$$\alpha_{h(k-1), h(k)} R_{h(k)} - c_{h(k)} = \alpha_{h(k-1), h(k)} R_{l(k)-1} - c_{l(k)-1}. \quad (10)$$

Since the expected outcomes of actions in I_K are strictly increasing, it holds that $R_{h(k)} > R_{l(k)-1}$, and so replacing $\alpha_{h(k-1), h(k)}$ with the larger $\alpha_{l(k)}$ in (10) gives $\alpha_{l(k)} R_{h(k)} - c_{h(k)} > \alpha_{l(k)} R_{l(k)-1} - c_{l(k)-1}$. By Corollary 3.5, $\alpha_{l(k)} = \alpha_{l(k)-1, l(k)}$, and so the right-hand side $\alpha_{l(k)} R_{l(k)-1} - c_{l(k)-1}$ equals $\alpha_{l(k)} R_{l(k)} - c_{l(k)}$. We conclude that $\alpha_{l(k)} R_{h(k)} - c_{h(k)} > \alpha_{l(k)} R_{l(k)} - c_{l(k)}$, i.e., in a linear contract with parameter $\alpha_{l(k)}$, action $h(k)$ has higher expected utility for the agent than action $l(k)$. But by definition, parameter $\alpha_{l(k)}$ implements action $l(k)$, and so we have reached a contradiction. \triangle

Applying Claim 5 to (9) we get the following chain of inequalities:

$$OPT \leq \sum_{k \leq K} (1 - \alpha_{l(k)}) R_{h(k)} < 2 \sum_{k \leq K} (1 - \alpha_{l(k)}) R_{l(k)} \leq 2K \cdot \max_{k \leq K} \{(1 - \alpha_{l(k)}) R_{l(k)}\} \leq 2K \cdot ALG,$$

where the strict inequality follows from $R_{h(k)}/2 < R_{l(k)}$. \square

Corollary 5.6 can be proven using a similar recursively defined instance as used in the proof of Theorem 5.2, except that instead of letting the expected outcomes and costs tend to infinity we double them.

PROOF OF COROLLARY 5.6. We use the same construction as in the proof of Theorem 5.2, but set $\epsilon = 1/2$. Since $n = \log H$, it indeed holds that $R_n = 2^{n-1} < H$. We know *OPT* can achieve at least $n - \epsilon(n - 1) > \frac{1}{2}n$ while *ALG* can't do better than 1, completing the proof. \square

5.3 Bounds in the Range of Costs

Our final pair of results concerns the costs. As in the case of expected outcomes, suppose costs are normalized such that $c_i \in [1, C]$ for all $a_i \in A_n$, consider bucketing these into $\lceil C \rceil$ buckets $[1, 2), [2, 4), \dots$ etc., and let L be the number of non-empty buckets.

THEOREM 5.7. *Consider a principal-agent setting (A_n, Ω_m) such that for every action $a \in A_n$, its cost c_a is in $[1, C]$. The multiplicative loss in the principal's expected payoff from using a linear contract rather than an arbitrary one is at most $4L = O(\log C)$.*

COROLLARY 5.8. *For every range $[1, C]$ such that $\log(2C) \geq 3$, there is a principal-agent setting (A_n, Ω_m) with $n = \log(2C)$ actions and $m = n$ outcomes for which MLRP holds and $\forall a \in A_n : c_a \in [1, C]$, such that the multiplicative loss in the principal's expected payoff from using a linear contract rather than an arbitrary one is at least $\frac{1}{2}n = \Omega(\log C)$.*

6 BEYOND LINEAR CONTRACTS

We conclude by showing that a similar lower bound to our bound for linear contracts (Theorem 5.2) applies to all *monotone* contracts; the only difference is that our lower bound of n (the number of actions) is slightly relaxed to $n - 1$. That is, we construct a MLRP instance with n actions in which the best monotone contract cannot guarantee better than a $\frac{1}{n-1}$ -approximation to the optimal contract's expected payoff.

THEOREM 6.1. *For every number of actions n , there is a principal-agent setting $(A_n^{\epsilon, \delta}, \Omega_m)$ parameterized by $\epsilon \gg \delta > 0$ for which MLRP holds, such that the multiplicative loss in the principal's expected payoff from using a monotone contract rather than an arbitrary one approaches $n - 1$ as $\epsilon, \delta \rightarrow 0$.*

The class of monotone contracts captures in particular *debt* contracts (mentioned already in Footnote 5). Theorem 6.1 thus implies that our results do not qualitatively change for this alternative family of simple contracts, which are common, e.g., in startup venture funding.¹¹ On the upper bound side, as we show in the full version, there always exists a debt contract which achieves a $\frac{1}{n-1}$ -approximation to the optimal contract's expected payoff.

7 CONCLUSION

One of the major contributions of theoretical computer science to economics has been the use of approximation guarantees to systematically explore complex economic design spaces, and to identify “sweet spots” of the design space where there are plausibly realistic solutions that simultaneously enjoy rigorous performance guarantees. For example, in auction and mechanism design, years of fruitful work by dozens of researchers has clarified the power and limitations of ever-more-complex mechanisms in a wide range of settings. Contract theory presents another huge

¹¹The first rewards generated by the startup CEO (the agent) go entirely to the principal to pay back for the initial loan that funded the company, so the contractual payments for the lower outcomes are zero; any further rewards are split linearly among the principal and agent.

opportunity for expanding the reach of the theoretical computer science toolbox, and we believe that this paper takes a promising first step in that direction.

ACKNOWLEDGMENTS

We wish to thank Michal Feldman and S. Matthew Weinberg for enlightening conversations. This research was supported by the National Science Foundation (grants CCF-1524062 and CCF-1813188), the European Research Council (grant 740282), the Israel Science Foundation (grant 11/633), and the British Academy (grant SRG1819/191601).

REFERENCES

- [1] Mark Armstrong and John Vickers. 2010. A Model of Delegated Project Choice. *Econometrica* 78, 1 (2010), 213–244.
- [2] Pablo Daniel Azar, Constantinos Daskalakis, Silvio Micali, and S. Matthew Weinberg. 2013. Optimal and Efficient Parametric Auctions. In *Proceedings of the 24th ACM-SIAM Symposium on Discrete Algorithms (SODA'13)*. 596–604.
- [3] Moshe Babaioff, Michal Feldman, and Noam Nisan. 2009. Free-Riding and Free-Labor in Combinatorial Agency. In *Proceedings of the 2nd Symposium on Algorithmic Game Theory (SAGT'09)*. 109–121.
- [4] Moshe Babaioff, Michal Feldman, and Noam Nisan. 2010. Mixed Strategies in Combinatorial Agency. *Journal of Artificial Intelligence Research* 38 (2010), 339–369.
- [5] Moshe Babaioff, Michal Feldman, Noam Nisan, and Eyal Winter. 2012. Combinatorial Agency. *Journal of Economic Theory* 147, 3 (2012), 999–1034.
- [6] Moshe Babaioff and Eyal Winter. 2014. Contract Complexity. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC'14)*. 911.
- [7] Chaithanya Bandi and Dimitris Bertsimas. 2014. Optimal Design for Multi-Item Auctions: A Robust Optimization Approach. *Mathematics of Operations Research* 39, 4 (2014), 1012–1038.
- [8] Hamsa Bastani, Mohsen Bayati, Mark Braverman, Ramki Gummadi, and Ramesh Johari. 2016. Analysis of Medicare Pay-for-Performance Contracts. SSRN:2839143. Retrieved from: <https://ssrn.com/abstract=2839143>.
- [9] Xiaohui Bei, Nick Gravin, Pinyan Lu, and Zhihao Gavin Tang. 2019. Correlation-Robust Analysis of Single Item Auction. In *Proceedings of the 30th ACM-SIAM Symposium on Discrete Algorithms (SODA'19)*. 193–208.
- [10] Patrick Bolton and Mathias Dewatripont. 2004. *Contract Theory*. The MIT Press, Cambridge, MA.
- [11] Bernard Caillaud and Benjamin E. Hermalin. 2000. Hidden-Information Agency. Retrieved from: <http://faculty.haas.berkeley.edu/hermalin/mechread.pdf>.
- [12] Vinicius Carrasco, Vitor Farinha Luz, Nenad Kos, Matthias Messner, Paulo Monteiro, and Humberto Moreira. 2017. Optimal Selling Mechanisms under Moment Conditions. *Journal of Economic Theory* 177 (2017), 245–279.
- [13] Gabriel Carroll. 2015. Robustness and Linear Contracts. *American Economic Review* 105, 2 (2015), 536–563.
- [14] Gabriel Carroll. 2017. Robustness and Separation in Multidimensional Screening. *Econometrica* 85, 2 (2017), 453–488.
- [15] Gabriel Carroll. 2019. Robustness in Mechanism Design and Contracting. *Annual Review of Economics* (2019). Forthcoming.
- [16] Alon Cohen, Moran Koren, and Argyrios Deligkas. 2018. Incentivizing the Dynamic Workforce: Learning Contracts in the Gig-Economy. arXiv:1811.06736. Retrieved from: <https://arxiv.org/abs/1811.06736>.
- [17] Lin William Cong and Zhiguo He. 2019. Blockchain Disruption and Smart Contracts. *Review of Financial Studies* 32, 5 (2019), 1754–1797.
- [18] Shaddin Dughmi. 2017. Algorithmic Information Structure Design: A Survey. *SIGecom Exchanges* 15, 2 (2017), 2–24.
- [19] Shaddin Dughmi, Nicole Immorlica, and Aaron Roth. 2014. Constrained Signaling in Auction Design. In *Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA'14)*. 1341–1357.
- [20] Yuval Emek, Michal Feldman, Iftah Gamzu, Renato Paes Leme, and Moshe Tennenholtz. 2014. Signaling Schemes for Revenue Maximization. *ACM Transactions on Economics and Computation* 2, 2 (2014), 5:1–5:19.
- [21] Nick Gravin and Pinyan Lu. 2018. Separation in Correlation-Robust Monopolist Problem with Budget. In *Proceedings of the 29th ACM-SIAM Symposium on Discrete Algorithms (SODA'18)*. 2069–2080.
- [22] Sanford J. Grossman and Oliver D. Hart. 1983. An Analysis of the Principal-Agent Problem. *Econometrica* 51, 1 (1983), 7–45.
- [23] Jason D. Hartline. 2017. Mechanism Design and Approximation. Retrieved from: <http://jasonhartline.com/MDnA/>.
- [24] Benjamin Hébert. 2017. Moral Hazard and the Optimality of Debt. *The Review of Economic Studies* 85, 4 (2017), 55–73.
- [25] Chien-Ju Ho, Aleksandrs Slivkins, and Jennifer Wortman Vaughan. 2016. Adaptive Contract Design for Crowdsourcing Markets: Bandit Algorithms for Repeated Principal-Agent Problems. *Journal of Artificial Intelligence Research* 55 (2016), 317–359.

- [26] Bengt Holmström and Paul Milgrom. 1987. Aggregation and Linearity in the Provision of Intertemporal Incentives. *Econometrica* 55, 2 (1987), 303–328.
- [27] Bengt Holmström and Paul Milgrom. 1991. Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design. *Journal of Law, Economics, & Organization* 7 (1991), 24–52.
- [28] Ali Khodabakhsh, Emmanouil Pountourakis, and Samuel Taggart. 2018. Algorithmic Delegation. Working paper.
- [29] Jon Kleinberg and Manish Raghavan. 2019. How Do Classifiers Induce Agents To Invest Effort Strategically. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC'19)*. Forthcoming.
- [30] Jon M. Kleinberg and Robert Kleinberg. 2018. Delegated Search Approximates Efficient Search. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC'18)*. 287–302.
- [31] Jean-Jacques Laffont and David Martimort. 2002. *The Theory of Incentives*. Princeton University Press, Princeton, NJ, USA.
- [32] Andreu Mas-Colell, Michael Dennis Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. Oxford University Press, Oxford, UK.
- [33] Eric Maskin and Jean Tirole. 1990. The Principal-Agent Relationship with an Informed Principal: The Case of Private Values. *Econometrica* 58, 2 (1990), 379–409.
- [34] Eric Maskin and Jean Tirole. 1992. The Principal-Agent Relationship with an Informed Principal, II: Common Values. *Econometrica* 60, 1 (1992), 1–42.
- [35] Peter Bro Miltersen and Or Sheffet. 2012. Send Mixed Signals: Earn More, Work Less. In *Proceedings of the 13th ACM Conference on Economics and Computation (EC'13)*. 234–247.
- [36] William P. Rogerson. 1985. Repeated Moral Hazard. *Econometrica* 53 (1985), 69–76.
- [37] Herbert E. Scarf. 1958. A Min-Max Solution of an Inventory Problem. In *Studies in the Mathematical Theory of Inventory and Production*, Kenneth J. Arrow, Samuel Karlin, and Herbert E. Scarf (Eds.). Stanford University Press, Stanford, CA, USA.
- [38] Steven Shavell. 1979. Risk Sharing and Incentives in the Principal and Agent Relationship. *Bell Journal of Economics* 10 (1979), 55–73.
- [39] A. Michael Spence. 1973. Job Market Signaling. *Quarterly Journal of Economics* 87 (1973), 355–374. Issue 3.
- [40] Steve Tadelis and Ilya Segal. 2005. Lectures in Contract Theory. (2005). Retrieved from: http://faculty.haas.berkeley.edu/stadelis/Econ_206_notes_2006.pdf.