Simple versus Optimal Contracts

Paul Dütting* Tim Roughgarden† Inbal Talgam-Cohen‡

Abstract

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, 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.

*Department of Mathematics, London School of Economics, Houghton Street, London WC2A 2AE, UK. Email: p.d.duetting@lse.ac.uk.

†Computer Science Department, Stanford University, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

‡Computer Science Department, Technion – Israel Institute of Technology, Haifa, Israel 3200003. Email: italgam@cs.technion.ac.il.

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 working for a company producing a range of products with different revenue margins. The salesperson chooses the amount of effort spent 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 products he generates), 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.[1] Another example is the interaction between a car owner and an insurance company. The car owner's behavior influences the risks his car is exposed to. If costs are borne only 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 [23], through blockchain-based smart contracts [14], 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 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 nonnegative 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

[1] We shall address the agent as male and the principal as female.

[2] For example, the 2016 Nobel Prize in economics was awarded to Oliver Hart and Bengt Holmström for their contributions to contract theory. The prize announcement stated: "Modern economies are held together by innumerable contracts... [T]ools created by Hart and Holmström are valuable to the understanding of real-life contracts".

flows from the principal to the agent).[3] A detailed description of the model appears in Section 2.

Optimal contracts and their issues. It is straightforward to solve the optimization problem associated with finding the optimal contract maximizing the principal's expected payoff, 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 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. However, this is not the end of the story: this approach can result in contracts that are *nothing like the contracts used in practice*.[4] Example 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. *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: $F_1 = (3/8, 3/8, 1/4, 0, 0, 0)$, $F_2 = (0, 3/8, 3/8, 1/4, 0, 0)$, $F_3 = (0, 0, 3/8, 3/8, 1/4, 0)$, $F_4 = (0, 0, 0, 3/8, 3/8, 1/4)$, 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 $t \approx (0, 0, 0.15, 3.93, 2.04, 0)$. The analysis appears for completeness in Appendix A.1.*

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 the most ubiquitous contract form in practice;[5] they are conceptually simple and easy to explain; payments are automatically monotone; and the agent is guaranteed non-negative utility with probability one. From an optimization standpoint, however, they can be suboptimal even in simple settings, as the next example demonstrates:

Example 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 $F_2 = (0, 1), c_2 = 4/3$, respectively. The optimal contract incentivizes action a_2 with payments $t = (0, 4/3)$, resulting in expected payoff of $3 - 4/3 = 5/3$ 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 complexity critique of optimal contracts is well known to economists. The dominant paradigm in the economics literature for addressing this is to justify simple contracts such as linear ones on robustness grounds.[6] Several works have attempted to characterize linear contracts as optimally robust in various max-min senses—see, e.g., [15, 13]; for a survey see [12]; for concurrent work see [18, 41] and Appendix C.2.

Perhaps most notable among these is the recent work of Carroll [11]. 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 A is fully known to the principal (that is, she is fully aware

[3] Without 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 $R_i - c_i$ that 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.

[4] A 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., [21]).

[5] To 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 [22]. Our results do not qualitatively change for such contracts—see Appendix F. We thank S. Matthew Weinberg for bringing this contract form to our attention.

[6] For example, in their classic paper on linear contracts in dynamic environments, Holmström and Milgrom [24] write: "It is probably the great robustness of linear rules based on aggregates that accounts for their popularity."

of the distributions the actions induce over outcomes as well as of their costs). See [11, p. 546] for "[o]ne way to make sense of this combination of non-quantifiable and quantifiable uncertainty", i.e., completely known and completely unknown actions. The main result is that a linear contract is max-min optimal in the worst case over all possible sets of actions B that contain the known set A (where B can be anything). Carroll sees the main contribution not 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" [12, 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. Our robustness result fits within the family of max-min characterizations championed by economic theory, but our setup and assumptions are more natural from an optimization perspective. 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. This is similar in spirit to previous work in algorithmic game theory on prior-independent parametric auctions, which sought max-min robust mechanisms in the worst case over all distributions with known first and second moments (see [2], and also the robust optimization approach of [7]). 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. For each of the four main parameters of the principal-agent model—number of actions n, number of outcomes m, range of expected rewards R_i, and range of costs c_i—we give tight approximation guarantees in that parameter, which apply uniformly across the other three parameters:

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*

(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 standard regularity condition to the principal-agent settings, called the monotone likelihood ratio property (MLRP) (see Appendix E.1). In Appendix F we show an extension of our lower bounds to all *monotone* (not necessarily linear) contracts: we show that among principal-agent settings with n actions, the worst-case ratio between the expected principal payoff under an optimal contract and under the best *monotone* contract can be $n-1$.

Discussion. We view parts (a)–(c) of the above theorem 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.

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 [38, 20, 35, 25], the excellent introductions of [10, 40], the comprehensive textbooks [29, 9] (see also [30, Chapters 13-14]), and the scientific background to the 2016 Nobel Prize in Economics [36].

Computational approaches to contract design. To our knowledge, decidedly computational approaches to contract design have appeared so far only in the work of [6] (see also follow-ups [4, 5]), the work of [3], and the work of [23] (see also follow-up [28]). The first paper [6] 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 [3] 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 [23] 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 [39, 31, 32]. This is not the case, and the relationship is no closer than that between auction theory (screening) and signaling. As Dughmi [16] 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, 34, 17].

Concurrent work on algorithmic delegation. Two recent papers [26, 27] 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 very different from contract design.

[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" [16, p. 1].

2 The Hidden Action Principal-Agent Model

The principal-agent model. An instance is described by a pair (A_n, Ω_m) of *n actions* and *m outcomes*.

- Outcomes: We identify the jth 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., $0 \leq x_1 \leq x_2 \leq \cdots \leq x_m$.
- Actions: 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. 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 action is hidden*: the principal observes the outcome x_j but not the action a.

Notation and terminology 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. 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 $c_a = 0$.

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 nonnegative 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 nonnegative for every j is referred to in the literature as *limited liability* [11], and it plays the same role as the standard *risk averseness* assumption in ruling out trivial solutions where a contract is not actually required [20]. 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 nonnegative 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

[8] The idea is that one could perturb the payment schedule slightly to make the desired action uniquely optimal for the agent. For further discussion see [10, p. 8].

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 agent.

Given an action a, the linear program (LP) appearing in Appendix A either finds a payment scheme that implements a at minimum expected payment to the agent, or establishes that a is not implementable. The optimal contract can be found by solving n such LPs—one for each action—and comparing the principal's expected payoff in each case after paying the agent. The LP and its dual can also be used to characterize implementable actions and to show that the optimal contract will have at most $n-1$ nonnegative payments—see Appendices A.2-A.3 for details.

Linear/monotone contracts. In addition to optimal contracts we consider the following simple classes.

Definition 1. *A contract is* linear *if 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 nonnegative 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:

Definition 2. *A contract is* monotone *if its payments are nondecreasing in the outcomes, i.e., $t_j \leq t_{j'}$ for $j < j'$.*

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.

Definition 3. *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 $\mathbb{E}_{x \sim F_a}[x] = R_a$.*

Definition 4. *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 $F_1, \ldots, F_n$ over the outcomes compatible with the actions.*

(A 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 Appendix G.)

In ambiguous settings, it is appropriate to apply a worst-case performance measure to evaluate contracts:

Definition 5. *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 [11] in making the following assumption, which simplifies but does not qualitatively affect the results in Section 4.[9]

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 $\rho(\mathcal{A}) := \max_{(A_n, \Omega_m) \in \mathcal{A}} \frac{OPT(A_n, \Omega_m)}{ALG(A_n, \Omega_m)}$.

[9] As explained in [11, 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 $x_1 = 0$ 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.

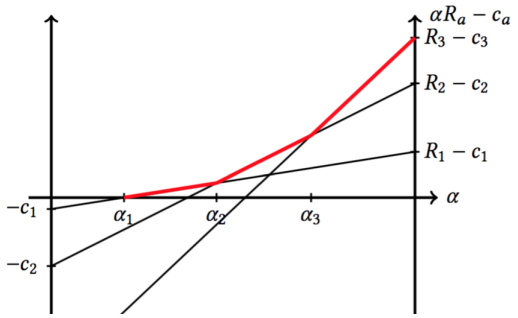

Figure 1: Linearly-implementable actions via upper envelope.

3 Properties and Geometry of Linearly-Implementable Actions

Our goal in this section is to establish a geometric characterization of *linearly-implementable* actions and to derive from it several useful consequences. See Appendix B for details and missing proofs.

Definition 6. *In a principal-agent setting (A_n, Ω_m), an action $a \in A_n$ is* linearly-implementable *if there exists a linear contract with parameter $\alpha \leq 1$ that implements a.*[10]

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 1 establish their equivalence.

- *Linear-implementability mapping $a(\cdot)$.* Denote by $a(\cdot) : [0,1] \to I_N \cup \{\varnothing\}$ 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—for completeness we state and prove this in Appendix B), or to $\varnothing$ 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 $a(\alpha) = a_i$.
- *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 nonnegative portion. Let $u(\cdot) : [0,1] \to A_n \cup \{\varnothing\}$ be the mapping from α to either $\varnothing$ 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.

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

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

Lemma 1 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., $R_1 - c_1 \leq R_2 - c_2 \leq \cdots \leq R_N - c_N$; (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 $\alpha_{i-1,i} = \frac{c_i - c_{i-1}}{R_i - R_{i-1}}$ (Observation 6 in Appendix B). Using this notation, we can rewrite the third implication as: $\alpha_i = \alpha_{i-1,i}$ for every $i \in [N]$.

The three implications above are formulated as Corollaries 4, 5 and 6, respectively, and appear with their proofs in Appendix B. We shall use Lemma 1 in Sections 4-5 and its corollaries in Section 5.

[10] The requirement $\alpha \leq 1$ is w.l.o.g.

4 Robust Optimality of Linear Contracts

In this section we establish a robust optimality result for linear contracts. All deferred proofs appear in Appendix C. Our main result in this section is that a linear contract maximizes the principal's expected payoff in ambiguous settings, *in the worst case over the unknown distributions*.

Theorem 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 1, "an optimal linear contract" is well-defined: For a (non-ambiguous) principal-agent setting, a linear contract is optimal if 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, thus an optimal one is still defined as maximizing the principal's expected payoff (see also Corollary 7 in Appendix B). In the remainder of the section we prove Theorem 1.

4.1 Main Lemma for Robust Optimality

The key step in the proof of Theorem 1 is to show that we may restrict the search for optimally robust contracts to affine contracts.

Lemma 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 $0 = x_1 < \cdots < x_m$ to payments $t_1, \ldots, t_m \geq 0$. 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 $F'_{i,m} := {}^{R_i}/_{x_m}$ (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 2 holds for the case of $t_1 > t_m$.*

Proof of Claim 1. A proof of this claim appears in Appendix C. □

Assume from now on that $t_1 \leq t_m$. If t is affine, this means that its slope parameter α_1 must be nonnegative. Similarly, we can write $t_1 = \alpha_0 + \alpha_1 x_1$ and plug in our assumption that $x_1 = 0$ to get $t_1 = \alpha_0$, and so by limited liability ($t \geq 0$), α_0 must also be nonnegative. Thus if t is affine, Lemma 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 2a and 2b. 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 Appendix C. □

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 2a and 2b). 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 nonnegative 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 $F''_{i,j} := {}^{R_i}/_{x_j}$ (by assumption this is a valid probability), and

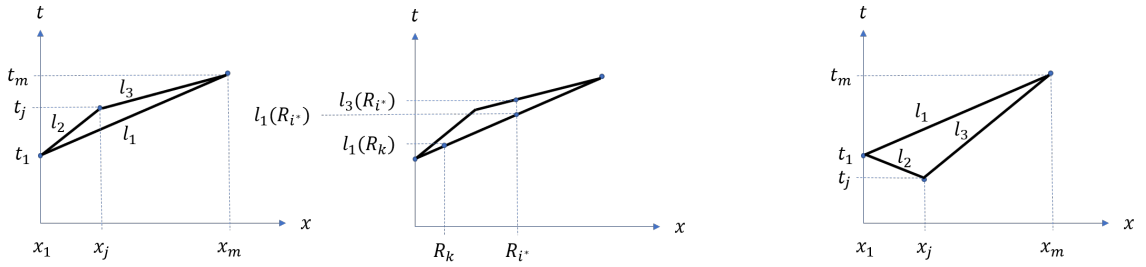

(a) Point (x_j,t_j) is strictly above line l_1.

(b) Point (x_j,t_j) is strictly below line l_1.

Figure 2: Case distinction

$F''_{i,1} := 1 - F''_{i,j}$. 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 2a and 2b.

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

Observation 1. *Consider a contract t and two outcomes x,x'; let $l : \mathbb{R} \to \mathbb{R}$ 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 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 1 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 1 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. □

Claim 4. *Lemma 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 Appendix C. □

This completes the analysis of the cases depicted in Figures 2a and 2b, thus proving Lemma 2. □

4.2 Proof of Theorem 1

Having established Lemma 2, Theorem 1 is now easy to prove using the following observation:

Observation 2. *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 1. *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 1. Consider an ambiguous principal-agent setting. For every limited liability contract t, by Lemma 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 1, the optimal linear contract has even higher worst-case expected payoff, completing the proof. □

5 Approximation Guarantees of Linear Contracts

In this section we study linear contracts and their approximation guarantees. We show tight bounds on the approximation guarantee of linear contracts in all relevant parameters of the model. Our upper bounds in fact apply to the stronger benchmark of optimal welfare. Our lower bounds continue to hold under standard regularity assumptions (see Appendix E.1).

Tight approximation guarantee in number of actions. Our first pair of results provides tight bounds on the approximation guarantee of linear contracts, as parametrized by the number of actions n.

Theorem 2. *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 3. *For every n and $\varepsilon > 0$, there is a principal-agent setting (A_n, Ω_m) with n actions and $m = n$ outcomes, such that the multiplicative loss in the principal's expected payoff from using a linear contract rather than an arbitrary one is at least $n - \varepsilon$.*

Tight approximation guarantee in 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 4. *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 2. *For every range $[1, H)$, there is a principal-agent setting (A_n, Ω_m) with $n = \log H$ for which $\forall a \in A_n : R_a \in [1, H)$, such that the multiplicative loss in the principal's expected payoff from using a linear contract rather than an arbitrary one is at least $\Omega(\log H)$*

Tight approximation guarantee in 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 $[1,2), [2,4), \ldots$ etc., and let L be the number of non-empty buckets.

Theorem 5. *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 3. *For every range $[1, C)$ such that $\log(2C) \geq 3$, there is a principal-agent setting (A_n, Ω_m) with $n = \log(2C)$ for which $\forall a \in A_n : c_a \in [1, C)$, such that the multiplicative loss in the principal's expected payoff from using a linear contract rather than an arbitrary one is at least $\frac{1}{2}n = \Omega(\log C)$.*

Additional tight bounds and optimality among all monotone contracts. In Appendix E we strengthen Theorem 3 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. In Appendix F we show that the approximation guarantee of n provided by linear contracts is asymptotically optimal among all monotone contracts.

5.1 Proofs for Selected Approximation Guarantees

To give a gist of our techniques, we now give the proofs of the upper and lower bounds in the number of actions (Theorems 2 and 3), and in the range of rewards (Theorem 4 and Corollary 2).

The upper bound in Theorem 5, the stronger version of Theorem 3 that holds for $m = 3$, and the strengthening of Theorem 3 to monotone contracts pose further technical challenges, and the proofs are deferred to Appendices D, E and F, respectively.

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

5.2 Proofs of Upper Bounds in Theorem 2 and Theorem 4

The key tools in proving the upper bounds is the following observation and the two lemmas below, which rely on the geometric insights from Section 3.

Observation 3. *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*

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

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. □

Below we shall apply Observation 3 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 6, in this contract the agent is indifferent among actions a_i and a_{i-1}. The left-hand side of (1) 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 (1) is the principal's expected payoff, and so the inequality holds.

The next lemma uses Observation 3 to upper-bound the expected welfare of the k-th linearly-implementable action.

Lemma 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., $R_{k-1} - c_{k-1} \leq \sum_{i=1}^{k-1} (1 - \alpha_{i-1,i}) R_i$ (*). By Corollary 5, 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 3 to actions $a = a_k, a' = a_{k-1}$ and get $(R_k - c_k) - (R_{k-1} - c_{k-1}) \leq (1 - \alpha_{k-1,k}) R_k$ (**). Adding inequality (*) to (**) completes the proof for k. □

The next lemma shows an upper bound on the payoff that the principal can achieve with an optimal (unconstrained) contract.

Lemma 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 4, $a(1) = a_N$ 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. □

With Lemma 3 and Lemma 4 at hand it is now relatively straightforward to show the upper bounds.

Proof of Theorem 2. To prove the approximation guarantee of $N \leq n$ claimed in the theorem, observe that

$$OPT \leq R_N - c_N \leq \sum_{i \leq N} (1-\alpha_{i-1,i}) R_i = \sum_{i \leq N} (1-\alpha_i) R_i \leq N \cdot \max_{i \leq N} \{(1-\alpha_i) R_i\} = N \cdot ALG, \tag{2}$$

where the first inequality holds by Lemma 4, the second inequality holds by Lemma 3, and the equality holds by Corollary 6. □

Proof of Theorem 4. 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 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 2. Let $\alpha_{h(k-1),h(k)} = (c_{h(k-1)} - c_{h(k)})/(R_{h(k-1)} - R_{h(k)})$. Observation 3 applies to actions in I_K, and so we can apply a version of Lemma 3 to get

$$OPT \leq R_{h(K)} - c_{h(K)} \leq \sum_{k \leq K} (1-\alpha_{h(k-1),h(k)}) R_{h(k)}. \tag{3}$$

Our goal is now to upper-bound the right-hand side of (3).

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

$$\alpha_{h(k-1),h(k)} R_{h(k)} - c_{h(k)} = \alpha_{h(k-1),h(k)} R_{l(k)-1} - c_{l(k)-1}. \tag{4}$$

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 (4) gives $\alpha_{l(k)} R_{h(k)} - c_{h(k)} > \alpha_{l(k)} R_{l(k)-1} - c_{l(k)-1}$. By Corollary 6, $\alpha_{l(k)} = \alpha_{l(k)-1,l(k)}$, 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. □

Applying Claim 5 to (3) we get the following chain of inequalities:

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

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

□

5.3 Proofs of Lower Bounds in Theorem 3 and Corollary 2

We conclude this section by showing matching lower bounds for the upper bounds established above.

Proof of Theorem 3. For every n, consider a family of principal-agent instances $\{(A_n^\varepsilon, \Omega_n) \mid \varepsilon > 0\}$, each with n actions and $m = n$ outcomes. For every $i \in [n]$, the ith action $a_i = (F_i, c_i) \in A_n^\varepsilon$ has $F_{i,i} = 1$, i.e., deterministically leads to the ith 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 (Definition 8) holds. We define action a_i's expectation R_i (equal to outcome x_i) and its cost c_i recursively:

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

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

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

$$R_i = \tfrac{1}{\varepsilon^{i-1}}, \quad c_i = \tfrac{1}{\varepsilon^{i-1}} - i + \varepsilon(i-1), \tag{5}$$

$$R_i - c_i = i - \varepsilon(i-1). \tag{6}$$

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 (5) into $\alpha_{i-1,i} = \frac{c_i - c_{i-1}}{R_i - R_{i-1}}$ we get

$$\alpha_{i-1,i} = 1 - \varepsilon^{i-1},$$

$$(1 - \alpha_{i-1,i}) R_i = 1. \tag{7}$$

$\alpha_{i-1,i}$ 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 $(A_n^\varepsilon, \Omega_n)$. 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 (6) we thus get

$$OPT^\varepsilon = \max_i \{i - \varepsilon(i-1)\} = n - \varepsilon(n-1) \underset{\varepsilon \to 0}{\longrightarrow} n.$$

In Lemma 7 in Appendix D, we analyze linear-implementability in the setting $(A_n^\varepsilon, \Omega_n)$, showing that $\alpha_i = \alpha_{i-1,i}$ for every $i \in [n]$. Thus from (7) it follows that

$$ALG^\varepsilon \leq 1,$$

completing the proof. □

Proof of Corollary 2. We use the same construction as in the proof of Theorem 3, but set $\varepsilon = 1/2$. Since $n = \log H$, it indeed holds that $R_n = 2^{n-1} < H$. We know OPT can achieve at least $n - \varepsilon(n-1) > \frac{1}{2}n$ while ALG can't do better than 1, completing the proof. □

6 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 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

The authors wish to thank Michal Feldman and S. Matthew Weinberg for enlightening conversations.

This research was supported by the Israel Science Foundation (ISF grant No. 11/633), and has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 740282).

A Linear Programming Formulation and Implications

The LP for incentivizing action a at minimum expected payment has m payment variables $\{t_j\}$, which by limited liability must be nonnegative, 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 nonnegative. The LP is:

$$\begin{array}{lll} \min & \sum_{j\in[m]} F_{a,j}t_j & \\ \text{s.t.} & \sum_{j\in[m]} F_{a,j}t_j - c_a \geq \sum_{j\in[m]} F_{a',j}t_j - c_{a'} & \forall a' \neq a, a' \in A_n, \\ & t_j \geq 0 & \forall j \in [m]. \end{array} \tag{8}$$

The dual of LP (8) has $n-1$ nonnegative variables, one for every action other than a:

$$\begin{array}{lll} \max & \sum_{a'\neq a} \lambda_{a'}(c_a - c_{a'}) & \\ \text{s.t.} & \sum_{a'\neq a} \lambda_{a'}(F_{a,j} - F_{a',j}) \leq F_{a,j} & \forall j \in [m], \\ & \lambda_{a'} \geq 0 & \forall a' \neq a, a' \in A_n. \end{array} \tag{9}$$

In the following subsections, we first use the LP-based approach to provide additional details for Example 1. We then show two implications of the LP-based approach. First, the LP and its dual can be used to characterize if an action is implementable or not, whether by an arbitrary contract or by a monotone one. Second, there always exists an optimal contract with at most $n-1$ positive payments.

A.1 Nonmonotonicity of the Optimal Contract

We analyze Example 1 in Section 1, which demonstrates nonmonotonicity of the optimal contract.

For payment profile $t \approx (0,0,0.15,3.93,2.04,0)$ all IC constraints in the LP are tight, i.e., the agent's utility is the same for all actions. The agent tie-breaks in favor of action a_3, which has the highest expected payoff of $2.95 = 4.99 - 2.04$ for the principal (where 4.99 is the expected outcome and 2.04 the expected payment).

No other action can achieve expected payoff for the principal as high as a_3. Moreover, if the payments are constrained to be monotone, or if the number of positive payments is constrained to be < 3, then a_3 can no longer be implemented for expected payment of merely 2.04.

A.2 Implementable and Monotonically-Implementable Actions

The following propositions characterize when actions are implementable, both by an arbitrary contract and by a monotone one.

Proposition 1 (Implementability). *An action a is implementable (up to tie-breaking) if and only if there is no convex combination of the other actions that results in the same distribution $\sum_{a'\neq a} \lambda_{a'} F_{a'} = F_a$ but lower cost $\sum_{a'\neq a} \lambda_{a'} c_{a'} < c_a$.*

Proposition 1 is immediate from linear programming duality of LP (8) with the objective replaced by "min 0".

Proposition 2 (Monotonic implementability). *An action a is implementable by a monotone contract (up to tie-breaking) if and only if there is no convex combination of the other actions that results in a first-order stochastically dominating distribution $\sum_{a'\neq a} \lambda_{a'} F_{a'}$ with lower cost $\sum_{a'\neq a} \lambda_{a'} c_{a'} < c_a$.*

Proposition 2 is immediate from linear programming duality of the following LP:

$$\begin{array}{lll} \min & 0 & \\ \text{s.t.} & \sum_{j\in[m]} F_{a,j}t_j - c_a \geq \sum_{j\in[m]} F_{a',j}t_j - c_{a'} & \forall a' \in A_n \setminus \{a\}, \\ & t_j \geq t_{j-1} & 2 \leq j \leq m, \\ & t_j \geq 0 & \forall j \in [m]. \end{array}$$

The dual has $n-1$ nonnegative λ-variables, one for every action other than a, and $m-1$ nonnegative μ-variables $\mu_2, \ldots, \mu_m$. It can be written w.l.o.g. as:

$$\begin{array}{lll} \max & c_a - \sum_{a'\neq a} \lambda_{a'} c_{a'} & \\ \text{s.t.} & F_{a,1} \leq \sum_{a'\neq a} \lambda_{a'} F_{a',1} + \mu_2 & , \\ & F_{a,j} \leq \sum_{a'\neq a} \lambda_{a'} F_{a',j} + \mu_{j+1} - \mu_j & 2 \leq j \leq m-1, \\ & F_{a,m} \leq \sum_{a'\neq a} \lambda_{a'} F_{a',j} - \mu_m & , \\ & \sum_{a'\neq a} \lambda_{a'} = 1 & , \\ & \lambda_{a'} \geq 0 & \forall a' \neq a, a' \in A_n. \\ & \mu_j \geq 0 & 2 \leq j \leq m. \end{array}$$

A.3 Number of Nonzero Payments in Optimal Contract

The next lemma implies the existence of an optimal contract with at most $n-1$ nonzero payments.

Lemma 5. *Consider a principal-agent setting (A_n, Ω_m) with n actions and m outcomes. For every implementable action a, there is an implementing contract with minimum expected payment, such that its payment scheme is positive for $\leq n-1$ outcomes.*

Proof. Because we assume action a is implementable, LP (8) is feasible and bounded, and so has an optimal basic feasible solution for which m constraints are tight [33]. There are only $n-1$ constraints other than non-negativity constraints, so at least $m-n+1$ of the non-negativity constraints are tight, meaning that the corresponding payments equal zero. Thus at most $n-1$ payments can be positive. □

B Appendix for Section 3

B.1 Basic Properties of Linearly-Implementable Actions

This section contains three observations on linear contracts and the incentives they create for the agent.

Observation 4. *If a linear contract with parameter α implements action a^*, then the agent's expected utility and the principal's expected payoff are, respectively,*

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

Observation 5. *Under Assumptions A1-A3, a linear contract implements at most one action.*

Proof. A contract implements more than one action only if two actions have the same maximum expected utility for the agent, and the same expected payoff for the principal (since the agent tie-breaks in favor of the principal). In a linear contract with parameter $\alpha < 1$, the principal's expected payoff is $(1-\alpha)R_a$, and as $R_a \neq R_{a'}$ for any two actions by Assumption A1, the necessary condition cannot hold. In a linear contract with parameter $\alpha = 1$, the agent's expected utility is $R_a - c_a$, and so by Assumption A2 there is a unique action that maximizes this, completing the proof. □

Observation 6. *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).*

For brevity we use $\alpha_{i-1,i}$ to denote the parameter α_{a_{i-1},a_i} that makes the agent indifferent among actions a_{i-1} and a_i.

Proof of Observation 6. The agent's expected utility from action a is $\alpha_{a,a'}R_a - c_a$, which is equal by definition of $\alpha_{a,a'}$ to action a''s utility $\alpha_{a,a'}R_{a'} - c_{a'}$. For an example in which α_{a',a_2} incentivizes neither a' nor a_2, consider the following. If the actions are $a_1 = (F_1, c_1) = ((1,0,0),\ 0)$, $a' = (F', c') = ((0,1,0),\ 1)$, and $a_2 = (F_2, c_2) = ((0,0,1),\ 2)$ and the outcomes are $(x_1, x_2, x_3) = (1,3,6)$, then only a_1 and a_2 can be implemented. □

B.2 Proof of Key Structural Lemma and its Implications

Our goal in this section is to prove Lemma 1 and derive implications.

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 $\{R_a - c_a\}$.

Lemma 6. *For every two parameters $0 \le \alpha < \alpha' \le 1$, either $u(\alpha) = \varnothing$, 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) $R_a - c_a \le R_{a'} - c_{a'}$.*

Proof. Notice that for every action a and corresponding line $\alpha R_a - c_a$, the slope R_a is non-negative. A key fact is that an upper envelope of affine functions with non-negative slopes is convex [37]. From convexity it follows that the upper envelope crosses the x-axis at most once, so $u(\alpha) \in A_n \implies u(\alpha') \in A_n$. Also from convexity, the slopes of the line segments forming the upper envelope are increasing in α, so $R_a < R_{a'}$. By Assumption A2 it follows that $c_a < c_{a'}$. Assume for contradiction that $R_a - c_a > R_{a'} - c_{a'}$. But this means that segment ℓ_a intersects $\alpha = 0$ at a higher point than segment $\ell_{a'}$ (as $-c_a > -c_{a'}$), and also intersects $\alpha = 1$ at a higher point (as $R_a - c_a > R_{a'} - c_{a'}$). Segment ℓ_a thus completely overshadows $\ell_{a'}$, in contradiction to the fact that $u(\alpha') = a'$ and so $\ell_{a'}$ is part of the upper envelope. □

Lemma 1 states that the linear-implementability mapping and the upper envelope mapping defined in Section 3 are equivalent. We now prove this lemma.

Proof of Lemma 1. Fix $\alpha \in [0,1]$ and consider the linear contract with parameter α. Action a is IR when $\alpha R_a - c_a \ge 0$, i.e., if and only if its corresponding segment ℓ_a is at or above the x-axis at α. Action a is IC when $\alpha R_a - c_a \ge \alpha R_{a'} - c_{a'}$ for every a', i.e., if and only if its segment ℓ_a participates in the upper envelope at α. Thus both $a(\alpha)$ and $u(\alpha)$ return $\varnothing$ when all segments at α are below the x-axis, equiv., no action is IR given the linear contract with parameter α. Otherwise, both return the action a whose segment ℓ_a forms the upper envelope at α above the x-axis, equiv., the IC and IR action given the linear contract with parameter α. In case of a tie, both break the tie in favor of the action with the highest expected outcome R_a—mapping $u(\cdot)$ does so by definition and mapping $a(\cdot)$ since this is the action that maximizes the principal's expected payoff $(1-\alpha)R_a$. This completes the proof. □

We state and prove three corollaries of Lemma 1.

Corollary 4. *For every $\alpha \in [0,1]$,*

$$\begin{aligned} \forall i \in [N-1] \ : \ a(\alpha) = a_i \quad &\iff \quad \alpha_i \le \alpha < \alpha_{i+1}, \\ a(\alpha) = a_N \quad &\iff \quad \alpha_N \le \alpha \le 1. \end{aligned} \tag{10}$$

Proof. Recall that mapping $a(\cdot)$ is onto I_N. Lemma 1 shows that $a(\cdot)$ is equivalent to the upper envelope mapping $u(\cdot)$. In the upper envelope, every segment appears once. This means that $a(\cdot)$ maps to action $a_i \in I_N$ for a consecutive range of αs, starting at α_i (by its definition as the smallest α such that $a(\alpha) = a_i$). In the upper envelope, the segments $\{\ell_a\}$ are ordered by their expected outcomes $\{R_a\}$ (Lemma 6). Since the actions in I_N are also ordered by their expected outcomes (i.e., $R_1 < \cdots < R_N$), the range of αs mapping to a_i is immediately followed by the range mapping to a_{i+1}, establishing (10). The final range of αs ending at 1 maps to action a_N, completing the proof. □

Corollary 5. *The welfare $R_i - c_i$ of linearly-implementable action $a_i \in I_N$ is increasing in i.*

Proof. Follows directly from Corollary 4 and Lemma 6. □

Corollary 6. *For every $i \in [N]$, $\alpha_i = \alpha_{i-1,i}$.*

Proof. For every $i \in [N]$, denote by ℓ_i the segment corresponding to action a_i. By Lemma 1 and Corollary 4, parameter α_i is precisely the intersection point between ℓ_{i-1} and ℓ_i for every $i \geq 2$. Observe that the intersection between $\alpha R_i - c_i$ and $\alpha R_{i-1} - c_{i-1}$ is at point $\alpha_{i-1,i}$. It remains to consider the case of $i = 1$, and in this case α_1 is the intersection point between ℓ_1 and the x-axis, which occurs at $\alpha = c_i/R_i$. □

The final corollary in this section highlights the fact that in terms of linear-implementability, two principal-agent settings whose actions have the same expected outcomes and costs are equivalent. The distributions, outcome values and even number of outcomes matter for linear-implementability only as far as determining the expected outcome of each action. This is special to linear contracts—optimal contracts can depend on the details of the distributions beyond just expected outcomes, adding to their complexity.

Corollary 7. *Consider two principal-agent settings (A_n, Ω_m), $(A'_n, \Omega'_{m'})$ for which there exists a bijection $b : A_n \to A'_n$ 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.*

Proof. The first part of the corollary follows from Lemma 1, which establishes equivalence between the linearly-implementability mapping and the upper envelope mapping, and from the fact that the upper envelope depends only on lines $\alpha R_a - c_a$ parameterized by R_a, c_a. The second part of the corollary follows from the fact that the principal's expected payoffs are $(1-\alpha)R_{a(\alpha)}$ and $(1-\alpha)R_{a'(\alpha)}$, respectively, and we have that $R_{a'(\alpha)} = R_{b(a(\alpha))} = R_{a(\alpha)}$. □

C Appendix for Section 4

C.1 Deferred Proofs

Proof of Claim 1. In this case, let $F_i = F'_i$ for every action a_i. We argue that the linear contract with parameter $\alpha = 0$ has expected payoff at least as high as that of contract t. Observe that since $t_1 > t_m$, the expected payments for the actions are decreasing in the actions' expected outcomes: if $R_i < R_k$ then $F_{i,m} = F'_{i,m} < F'_{k,m} = F_{k,m}$ and so $F_{i,1} > F_{k,1}$; thus $F_{i,1}t_1 + F_{i,m}t_m > F_{k,1}t_1 + F_{k,m}t_m$. Consider the zero-cost action a_1 (which exists by Assumption A3), and let a_{i^*} be the action incentivized by contract t. The expected outcome of action a_{i^*} must be (weakly) lower than that of action a_1—its cost is (weakly) higher so its expected payment must be (weakly) higher. Since the agent's choice of action a_{i^*} is IR, its expected outcome is an upper bound on contract t's expected payoff to the principal. But the linear contract with parameter $\alpha = 0$ incentivizes an action with (weakly) higher expected outcome at no cost to the principal,

thus guaranteeing (weakly) higher expected payoff to the principal. This completes the proof for the case of $t_1 > t_m$. □

Proof of Claim 2. A characterization of affine mappings is that they map every 3 collinear points to points that are themselves collinear. Thus there must exist 3 points $(x,t(x)),(x',t(x')),(x'',t(x''))$ where (w.l.o.g.) $x < x' < x''$ such that these points are not collinear. Now consider the line between the 2 points (x_1,t_1) and (x_m,t_m). It cannot be the case that the 3 points $(x,t(x)),(x',t(x')),(x'',t(x''))$ are all on this line. Thus we have shown the existence of an index j as required. □

Proof of Claim 4. We prove the claim by showing that the affine contract with parameters α_0,α_1 (corresponding to l_1) has 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 $k \neq i^*$. A similar argument as in the proof of Claim 3 establishes that for distributions $\{F_i\}_{i=1}^n$ as defined, contract t will also incentivize action a_{i^*}, and at the same expected payment to the agent as the affine contract. This is because the expected payment for a_{i^*} is the same in contract t as in the affine contract, while the expected payments for all other actions (weakly) decrease compared to the affine contract. □

Proof of Observation 2. Fix $\alpha_0 \geq 0$, and consider the mapping from α_1 to the action implemented by the affine contract with parameters α_0,α_1. This mapping is identical to the mapping from α to the action implemented by the linear contract with parameter $\alpha = \alpha_1$. This follows from the analysis in Section 3 and in particular Lemma 1,[11] and by observing that the line segments forming the upper envelope for the affine contract are the same as those of the linear contract, with an additional vertical shift of magnitude α_0. So for every α_1, the linear contract with $\alpha = \alpha_1$ implements the same action as the affine contract; but its expected payment to the agent is lower by $\alpha_0 \geq 0$ than that of the affine contract. Thus its expected payoff to the principal is higher by $\alpha_0 \geq 0$, completing the proof. □

C.2 Further Robustness Models in the Literature

In this section we briefly explain how our model in Section 4 differs from two other robustness models for linear contracts considered in previous and concurrent economic literature. While in our model distributions are adversarial, in Diamond [15] one can predict what the distributions will be given the contract (and the focus is on only $n = 2$ actions). Two concurrent working papers [18, 41] crucially assume uncertainty about the distributions coupled with max-min behavior *on the agent side too*. In [18], the principal does not know what the agent knows; this makes their model very different from ours, in which the principal knows that the agent has knowledge of the distributions.

D Appendix for Section 5

D.1 Auxiliary Lemma used in the Proof of Theorem 3

Lemma 7. *Consider the principal-agent settings $(A_n^\varepsilon,\Omega_n)$ defined in the proof of Theorem 3. Then, $\alpha_i = \alpha_{i-1,i}$.*

Proof. Recall from Lemma 1 that a linear contract with parameter α implements the action whose segment forms the upper envelope at α. The next claim shows that for every $\alpha > \alpha_{i-1,i}$, the segment of every action $a_{i'}$ such that $i' < i$ is (weakly) below that of action a_i. It can similarly be shown that for every $\alpha < \alpha_{i,i+1}$ (or $\alpha \leq 1$ for $i = n$), the segment of every action $a_{i'}$ such that $i' > i$ is below that of action a_i.

[11] Note that we are not limiting α_1 to be ≤ 1; this is not an issue since everything in Section 3 technically holds for $\alpha_1 > 1$ (of course, $\alpha_1 > 1$ does not make sense for the principal since it leaves her with negative expected payoff).

Claim 6. *For every $\alpha \geq \alpha_{i-1,i}$, the agent's expected utility from action a_i is at least his expected utility from any "previous" action $a_{i'}$ where $i' < i$.*

Proof of Claim 6. The claim holds trivially for the base case $i = 1$. Assuming the claim holds for $i-1$, to establish it for i it is sufficient to show that for every $\alpha \geq \alpha_{i-1,i}$, the agent's expected utility from action a_i is at least his expected utility from action a_{i-1}. For every $\alpha = 1 - \varepsilon^{i-1} + \delta$ where $\delta \geq 0$, action a_i has expected utility $\alpha R_i - c_i = (i-1)(1-\varepsilon) + \frac{\delta}{\varepsilon^{i-1}}$ for the agent, whereas action a_{i-1} has expected utility $\alpha R_{i-1} - c_{i-1} = (i-1)(1-\varepsilon) + \frac{\delta}{\varepsilon^{i-2}}$, which is lower as required to establish Claim 6. □

We conclude that for every $\alpha \in [\alpha_{i-1,i}, \alpha_{i,i+1})$ (or $\alpha \in [\alpha_{n-1,n}, 1)$ for $i = n$), the linear contract with parameter α implements action a_i, and so $\alpha_i = \alpha_{i-1,i}$. This completes the proof of Lemma 7. □

D.2 Proof of Theorem 5 and Corollary 3

Proof of Theorem 5. Consider the set I_N of linearly implementable actions. Denote by $C = \max_{a \in I_N} c_a$ the highest cost of any of the linearly implementable actions. We will bucket the set of linearly implementable actions into $L = \lceil \log_2(C) \rceil$ buckets $B_1, \ldots, B_L$ such that

$$B_i = \{a \mid 2^{i-1} \leq c_a < 2^i\}.$$

Note that this bucketing ensures that every implementable actions is in some bucket. By Lemma 6 within each bucket actions are sorted simultaneously by expected outcome, cost, and welfare.

As in the proof of Theorem 4 let $h(k)$ denote the action $a \in B_k$ with the highest R_a (and hence highest c_a), and let $l(k)$ denote the action $a \in B_k$ with the lowest R_a (and hence lowest c_a). Now by the same argument as in Theorem 4

$$OPT \leq R_{h(k)} - c_{h(k)} \leq \sum_{k \leq L} (1 - \alpha_{h(k),h(k)-1}) R_{h(k)}. \tag{11}$$

Claim 7. *For every bucket B_k either (C1) $R_{l(k)} \geq R_{h(k)}/4$ or (C2) there exist an action $a_i \in B_k$ such that $R_i \geq R_{h(k)}/2$ and $\alpha_i \leq 1/2$.*

Before we prove this claim let's see how it implies a logarithmic approximation guarantee. The high level idea is that in each bucket B_k we will identify a linearly implementable action $\tau(k)$ whose expected payoff to the principal is at least one quarter of that bucket's contribution to the sum in on the RHS of inequality (11).

Consider a fixed bucket B_k. We say that bucket B_k is of Type 1 if it meets condition (C1) and of Type 2 if it meets condition (C2). For Type 1 buckets we choose $\tau(k) = l(k)$, and for Type 2 buckets we choose $\tau(k) = i$.

Then for Type 1 buckets:

$$(1 - \alpha_{h(k),h(k)-1}) R_{h(k)} \leq (1 - \alpha_{l(k)}) R_{h(k)} \leq 4 \cdot (1 - \alpha_{l(k)}) R_{l(k)} = 4(1 - \alpha_{\tau(k)}) R_{\tau(k)}$$

And for Type 2 buckets:

$$(1 - \alpha_{h(k),h(k)-1}) R_{h(k)} \leq R_{h(k)} \leq 4(1 - \alpha_i) R_i = 4(1 - \alpha_{\tau(k)}) R_{\tau(k)}$$

Where for the derivation of the inequalities for Type 1 buckets we used Claim 1 in the proof of Theorem 4, which shows that $\alpha_{h(k),h(k)-1} \geq \alpha_{l(k)}$, and for Type 2 buckets we used that $\alpha_{h(k),h(k)-1} \geq 0$.

We thus get,

$$\begin{aligned} OPT &\leq \sum_{k\in[L]} (1-\alpha_{h(k),h(k)-1})R_{h(k)} \\ &\leq L\cdot \max_{k\in[L]}\{(1-\alpha_{h(k),h(k)-1})R_{h(k)}\} \\ &\leq 4L\cdot \max_{k\in[L]}\{(1-\alpha_{\tau(k)})R_{\tau(k)}\} \\ &\leq 4L\cdot ALG \end{aligned}$$

To complete the proof of Theorem 5 it remains to show Claim 7.

Claim 7. Fix a bucket B_k. Consider the actions $l(k),\ldots,h(k)$ in B_k. Note that $\alpha_{l(k)} \leq \alpha_{l(k)+1} \leq \cdots \leq \alpha_{h(k)}$. If Condition (C1) is met, we are done. So assume Condition (C1) is *not* met. That is, assume that $R_{l(k)} < R_{h(k)}/4$. Note that this is possible only if $l(k) \neq h(k)$ and, thus, $l(k) < h(k)$. Further note that if $\alpha_{h(k)} \leq 1/2$ then Condition (C2) would be met by $i = h(k)$. So the only cases left are those where for a non-empty suffix of the indices $l(k),\ldots,h(k)$ it holds that $\alpha_j > 1/2$.

We claim that it can't be that $\alpha_j > 1/2$ for all $j \in \{l(k),\ldots,h(k)\}$. Indeed, if this was the case, we could use that for all $j' = l(k)+1,\ldots,h(k)$ by the definition of $\alpha_{j'}$

$$R_{j'} - R_{j'-1} = \frac{1}{\alpha_{j'}}(c_{j'} - c_{j'-1}) \leq 2\cdot(c_{j'} - c_{j'-1}).$$

Summing this inequality over all $j' = l(k)+1,\ldots,h(k)$ would give us

$$R_{h(k)} - R_{l(k)} = \sum_{j'=l(k)+1}^{h(k)} (R_{j'} - R_{j'-1}) \leq \sum_{j'=l(k)+1}^{h(k)} 2\cdot(c_{j'} - c_{j'-1}) = 2\cdot(c_{h(k)} - c_{l(k)}).$$

Since $c_{h(k)} < 2c_{l(k)}$ and $c_{l(k)} \leq R_{l(k)}$ this would show

$$R_{l(k)} \geq \frac{1}{3}\cdot R_{h(k)},$$

but this would contradict our assumption that Condition (C1) is not met.

So there must be a largest index i with $l(k) \leq i < h(k)$ for which it holds that $a_i \leq 1/2$ and $\alpha_{i'} > 1/2$ for all $i' > i$. We claim that this i satisfies Condition (C2). It certainly has $a_i \leq 1/2$. For the expected outcome we can use the same argument that we used when we assumed that all the α_j's are strictly positive to conclude that

$$R_{h(k)} - R_i = R_{h(k)} - R_{i'-1} \leq 2\cdot(c_{h(k)} - c_{i'-1}) = 2\cdot(c_{h(k)} - c_i).$$

Because the actions in each bucket are sorted by costs, $c_i \geq c_{l(k)}$. Also, as we have argued before, $c_{h(k)} < 2c_{l(k)}$ and $c_{l(k)} \leq R_{l(k)}$. So,

$$R_{h(k)} - R_i \leq 2\cdot(c_{h(k)} - c_i) \leq 2c_{l(k)} \leq 2R_{l(k)}.$$

But now because Condition (C1) is not met

$$R_{h(k)} - R_i \leq 2R_{l(k)} \leq \frac{1}{2}R_{h(k)},$$

which shows that $R_i \geq R_{h(k)}/2$ as claimed. This established Claim 7. □

□

Proof of Corollary 3. Consider the construction of the lower bound in Theorem 3.3 (Workshop version) with $\varepsilon = 1/2$. Then $c_i = \frac{1}{2}(2^i - i - 1)$ for all $1 \le i \le n$. In particular, $c_n \in (\frac{1}{2} \cdot 2^{n-1}, \frac{1}{2} \cdot 2^n) = (\frac{C}{2}, C)$ where we used $\log(2C) \ge 3$. With an arbitrary contract the principal can guarantee himself an expected payoff of $OPT = n - \varepsilon(n-1) > \frac{n}{2}$, while with a linear contract he can achieve at most $ALG \le 1$. So $OPT/ALG \ge \frac{n}{2} = \Omega(\log C)$. □

E Strengthened Lower Bound for Linear Contracts

In this appendix we present a lower bound construction for linear contracts that requires only three outcomes. The proof of this lower bound also demonstrates that many actions can be linearly-implementable even if there are just a few outcomes.

E.1 The Regularity Assumption of MLRP

The economic literature (see, e.g., [20]) 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 7 (MLR). *Let F, G be two distributions over m values $v_1, \ldots, v_m$. The likelihood ratio F_j/G_j is monotonically increasing in j if*

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

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 8 (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 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.*

We demonstrate MLRP and non-MLRP through the following examples:

Example 3 (Two outcomes). *Assume there are $m = 2$ outcomes $\ell < h$, and that a higher action cost means higher expected outcome. Then the probability $F_{i,h}$ of action i to achieve the high outcome is strictly increasing in i. Thus MLRP holds.*

Example 4 (Spanning condition [20]). *Assume that the agent has two basic actions, such as "effort" (costly) and "no effort" (cost zero), for which MLR holds, and the agent can interpolate among these (this is known as a setting satisfying the "spanning condition"). Then the resulting action set satisfies MLRP.*

Example 5 (Binomial distributions). *Assume that higher cost means more effort on behalf of the agent, that the level of effort determines the probability of the agent's success in a Bernoulli trial, and that the outcome is the Binomially distributed number of successful trials out of a total of $m-1$ trials. Then the action set satisfies MLRP: one can verify that*

$$\frac{\binom{m-1}{j-1} p^{j-1}(1-p)^{m-j}}{\binom{m-1}{j-1} q^{j-1}(1-q)^{m-j}}$$

is increasing in j when $p > q$.

Example 6 (No MLRP [10]). *Let $n = 2$ and $m = 3$. We define a principal-agent setting (A_n, Ω_m) where $A_n = (a_1, a_2)$ and $\Omega_m = (x_1, x_2, x_3) = (0, 1, 2)$. The actions are defined as follows:*

$$\begin{aligned} a_1 = (F_1, c_1) &= ((1/3, 1/3, 1/3),\ 0), \\ a_2 = (F_2, c_2) &= ((1/3, 1/6, 1/2),\ 1). \end{aligned}$$

These distributions can be viewed as convex combinations of distributions with MLR (analogously to combinations of regular distributions leading to irregularity in auction theory). Namely: $F_i = \frac{1}{3}F_i^1 + \frac{2}{3}F_i^2$*, where* $F_1^1 = (1,0,0)$*,* $F_1^2 = (0,1/2,1/2)$*,* $F_2^1 = (1,0,0)$*,* $F_2^2 = (0,1/4,3/4)$*, and the likelihood ratios* $F_{2,j}^1/F_{1,j}^1$ *and* $F_{2,j}^2/F_{1,j}^2$ *are both monotonically increasing.*[12]

In addition to MLRP, there are other less natural regularity assumptions in the literature, such as CDFP (see Definition 9).

E.2 Lower Bound Statement and Proof

Theorem 6. *For every number of actions n, there is a principal-agent setting* (A_n'', Ω_3) *with* $m = 3$ *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.*

Proof. Let $(A_n^\varepsilon, \Omega_m)$ be the principal-agent setting defined in the proof of Theorem 3, where $R_i = 1/\varepsilon^{i-1}$ and $c_i = R_i - i + \varepsilon(i-1)$. Let *ALG* be the best achievable expected payoff to the principal using a linear contract, and let *OPT* be the same using an arbitrary contract. From the proof of Theorem 3 we know that $ALG = 1$ and $OPT \to n$ as $\varepsilon \to 0$. To prove Proposition 6, we define two additional principal-agent settings, (A_n', Ω_2) and (A_n'', Ω_3), with ALG', OPT' and ALG'', OPT'' denoting their optimal linear and optimal expected payoffs, respectively. Our settings will be such that MLRP holds, and $ALG'' \to 1$ while $OPT'' \to n$, thus establishing the proposition.

Auxiliary setting (A_n', Ω_2)**.** Let Ω_2 be an outcome space with 2 outcomes $x_1 = 0, x_2 = R_n$. For every action $a_i \in A_n^\varepsilon$, we define a corresponding action $b(a_i)$ with the same cost c_i, which leads to outcome x_2 with probability R_i/R_n and to outcome x_1 otherwise. The expected outcome of action $b(a_i)$ is thus R_i. Let $A_n' = \{b(a_i)\}_{i\in[n]}$ be the collection of all actions corresponding to those in A_n^ε. Observe that MLRP holds for A_n', since for every $i' > i$, the likelihood ratio of outcome x_2 is $R_{i'}/R_i > 1$, and the likelihood ratio of outcome x_1 is $(R_n - R_{i'})/(R_n - R_i) < 1$.

Invoking Corollary 7 for principal-agent settings $(A_n^\varepsilon, \Omega_m)$ and (A_n', Ω_2), we get that the principal's expected payoff from any linear contract is the same in both settings, so $ALG' = ALG = 1$. While action $b(a_n)$ has welfare approaching n as $\varepsilon \to 0$, we are not in a full information setting and thus OPT' may be much lower.

Setting (A_n'', Ω_3)**.** Our goal now is to define a principal-agent setting (A_n'', Ω_3) for which $ALG'' \approx ALG'$, $OPT'' \approx n$, and MLRP still holds. We start from (A_n', Ω_2) and add an outcome $x_3 = x_2 + 1 = R_n + 1$ to get the new outcome set Ω_3. We change action $b(a_n)$ such that it leads to outcome x_3 with some small probability δ (to be determined below); the probabilities over the other outcomes are renormalized by factor $(1-\delta)$. We denote the resulting action by a_n'', and its expected outcome by $R_n'' = (1-\delta)R_n + \delta x_3 = R_n + \delta$. We change every other action $b(a_i)$ only by adding zero probability that it leads to outcome x_3, and denote the new action by a_i''. The new action set A'' is $\{a_i''\}_{i\in[n]}$. Observe that MLRP still holds, since the likelihood ratio of action a_n'' and any other action a_i'' for outcome x_3 is ∞.

In the new setting, $OPT'' = R_n'' - c_n = R_n + \delta - c_n = OPT + \delta$, by paying c_n/δ for outcome x_3 and zero for any other outcome, thus incentivizing the agent to choose action a_n'' while paying c_n in expectation. As for ALG'', the only change relative to the original setting $(A_n^\varepsilon, \Omega_m)$ and the auxiliary setting (A_n', Ω_2) is that action a_n'' becomes linearly-implementable by a contract with a smaller parameter α than the original action a_n. Denoting this parameter by α_n'', we have that $ALG'' = (1-\alpha_n'')R_n''$, since linearly-implementing any other

[12] A possible economic story behind this example could be that the agent chooses between "no effort" (action a_1) and "effort" (action a_2). Without effort, the distribution is $(1/2,1/2)$ over outcomes (x_2,x_3), and with effort the distribution is $(1/4,3/4)$. However, regardless of the agent's effort level, with probability $1/3$ some exogenous bad event occurs (e.g., the market adopts a different technology as the industry standard, causing sales to drop), resulting in an outcome of $x_1 = 0$.

action has inferior expected payoff of 1. By Corollary 6, $\alpha''_n = (c_n - c_{n-1})/(R''_n - R_{n-1})$. So

$$\begin{aligned} ALG'' &= (1-\alpha''_n)R''_n \\ &= (1-\frac{c_n - c_{n-1}}{R_n+\delta-R_{n-1}})(R_n+\delta) \\ &= \frac{\varepsilon^{n-1}-\varepsilon^n+\delta\varepsilon^{n-1}}{1-\varepsilon+\delta\varepsilon^{n-1}}(\varepsilon^{-(n-1)}+\delta) \\ &\leq (\varepsilon^{n-1}+\delta\varepsilon^{n-1})(\varepsilon^{-(n-1)}+\delta) \\ &= 1+\delta+\varepsilon^{n-1}\delta(1+\delta). \end{aligned}$$

By letting $\delta \to 0$ and $\varepsilon \to 0$ we get $OPT''/ALG'' \to n$, completing the proof. □

F Lower Bound for Monotone Contracts

In this appendix we show that a similar lower bound to our bound for linear contracts (Theorem 3) applies to all *monotone* contracts (Definition 2); the only difference is that our lower bound of n (the number of actions) is slightly relaxed to $n-1$. That is, we construct an (MLRP – see Appendix E.1) 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. Since the class of monotone contracts captures in particular debt contracts (see Footnote 5), this shows our results do not qualitatively change for this alternative family of simple contracts.

Theorem 7. *For every number of actions n, there is a principal-agent setting $(A_n^{\varepsilon,\delta}, \Omega_m)$ parameterized by $\varepsilon \gg \delta > 0$ 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 $\varepsilon, \delta \to 0$.*

The instance we construct to prove Theorem 7 is based upon the construction of Proposition 6 for $n-1$ actions, with an additional nth high-cost action. Intuitively, the IC constraint with this extra action together with monotonicity enforce relatively homogeneous payments over the highest outcomes, whereas the optimal contract requires a single high payment for the second-highest outcome.

Proof of Theorem 7. Let $\varepsilon, \delta, \gamma > 0$ be vanishingly small, and define an m-outcome vector and $n \times m$ distribution matrix as follows, where $m = 4$:

$$\begin{aligned} x &= \left(0 \quad \tfrac{1}{\varepsilon^{n-2}} \quad \tfrac{1}{\varepsilon^{n-2}}+\gamma \quad \tfrac{1}{\varepsilon^{n-2}}+2\gamma\right), \\ F &= \begin{pmatrix} 1-\varepsilon^{n-2} & \varepsilon^{n-2} & 0 & 0 \\ 1-\varepsilon^{n-1} & \varepsilon^{n-1} & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 1-\varepsilon^{n-i-1} & \varepsilon^{n-i-1} & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 1-\varepsilon & \varepsilon & 0 & 0 \\ 0 & 1-\delta & \delta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \end{aligned}$$

The costs are $c_i = \frac{1}{\varepsilon^{i-1}} - i + \varepsilon(i-1)$ for $i \leq n-1$, and $c_n = \frac{1}{\varepsilon^{n-2}}$.

Observe that in the above setting, MLRP holds; the expected outcomes are $R_i = \frac{1}{\varepsilon^{i-1}}$ for $i \leq n-2$, $R_{n-1} = \frac{1}{\varepsilon^{n-2}} + \delta\gamma$, and $R_n = \frac{1}{\varepsilon^{n-2}} + 2\gamma$; the expected welfares are $R_i - c_i = i - \varepsilon(i-1)$ for $i \leq n-2$, $R_{n-1} - c_{n-1} = n-1-\varepsilon(n-2)+\delta\gamma$, and $R_n - c_n = 2\gamma$. We now analyze several possible contracts to establish the theorem.

Optimal contract. The optimal contract incentivizes action $n-1$ by setting $t_3 = \frac{c_{n-1}}{\delta}$ (such that actions a_1, a_{n-1} both have expected utility 0 for the agent, all other actions have negative expected utilities, and tie-breaking is in favor of the principal). The expected payoff to the principal is $R_{n-1} - \delta t_3 \approx n-1$ (action a_{n-1}'s welfare).

Optimal monotone contract incentivizing $a_i \neq a_{n-1}$. We claim that the payoff that the principal can achieve by incentivizing any action a_i with $i \neq n-1$ is at most 1, so that it suffices to consider action a_{n-1}. For action a_1 the payoff to the principal is upper bounded by the welfare this action obtains, which is 1. For every action a_i where $2 \leq i \leq n-2$, dis-incentivizing deviation from a_i to a_{i-1} necessitates

$$\begin{aligned} & (1-\varepsilon^{n-i-1})t_1 + \varepsilon^{n-i-1}t_2 - c_i \geq (1-\varepsilon^{n-i})t_1 + \varepsilon^{n-i}t_2 - c_{i-1} \\ \Leftrightarrow \quad & (1-\varepsilon)\varepsilon^{n-i-1}t_2 \geq \big(c_i - c_{i-1}\big) + (1-\varepsilon)\varepsilon^{n-i-1}t_1 \\ \Rightarrow \quad & t_2 \geq \frac{1}{(1-\varepsilon)\varepsilon^{n-i-1}}\big(c_i - c_{i-1}\big), \end{aligned}$$

where

$$c_i - c_{i-1} = \left(\frac{1}{\varepsilon^{i-1}} - i + \varepsilon(i-1)\right) - \left(\frac{1}{\varepsilon^{i-2} - (i-1) + \varepsilon(i-2)}\right) = (1-\varepsilon)\left(\frac{1}{\varepsilon^{i-1}} - 1\right).$$

We get that

$$t_2 \geq \frac{1}{\varepsilon^{n-i-1}}\left(\frac{1}{\varepsilon^{i-1}} - 1\right),$$

and the payoff to the principal is at most

$$\frac{1}{\varepsilon^{i-1}} - \varepsilon^{n-i-1}\frac{1}{\varepsilon^{n-i-1}}\left(\frac{1}{\varepsilon^{i-1}-1}\right) = 1.$$

The expected welfare of a_n is almost zero.

Optimal monotone contract incentivizing a_{n-1}. It is w.l.o.g. to set $t_1 = 0$. To dis-incentivize deviations from a_{n-1} to a_{n-2} and to a_n:

$$(1-\delta)t_2 + \delta t_3 - c_{n-1} \geq \varepsilon t_2 - c_{n-2}, \tag{12}$$

$$(1-\delta)t_2 + \delta t_3 - c_{n-1} \geq t_4 - c_n. \tag{13}$$

From (13) and monotonicity we have that $(1-\delta)t_2 + \delta t_3 \geq t_4 - (c_n - c_{n-1}) \geq t_3 - (c_n - c_{n-1})$. So:

$$t_2 \geq t_3 - \frac{c_n - c_{n-1}}{1-\delta}.$$

Combining this with (12) we get $(1-\delta-\varepsilon)t_2 \geq c_{n-1} - c_{n-2} - \delta t_3 \geq c_{n-1} - c_{n-2} - \delta t_2 - \frac{\delta(c_n - c_{n-1})}{1-\delta}$, and after rearranging,

$$(1-\varepsilon)t_2 \geq c_{n-1} - c_{n-2} - \frac{\delta(c_n - c_{n-1})}{1-\delta}.$$

Since $c_{n-1} - c_{n-2} = \frac{1-\varepsilon}{\varepsilon^{n-2}} - (1-\varepsilon)$,

$$t_2 \geq \frac{1}{\varepsilon^{n-2}} - 1 - \frac{\delta(c_n - c_{n-1})}{(1-\delta)(1-\varepsilon)}.$$

So the expected payment in the optimal monotone contract incentivizing action a_{n-1} is at least $\frac{1}{\varepsilon^{n-2}} - 1$ minus a term that is vanishing with δ, leaving an expected payoff of ≈ 1 for the principal. This completes the proof of Theorem 7. □

G Special Cases in Which Simple Contracts are Optimal

In several special cases of interest, a simple contract with a single nonzero payment is optimal. An obvious such case (given Lemma 5, which bounds the number of non-negative payments by $n-1$) is that of $n=2$ actions. We state here for completeness the optimal contract in this case:

Proposition 4. *For $n=2$ actions, if the optimal contract incentivizes the nonzero-cost action a_2 rather than a_1 with cost 0, then there is an optimal contract that pays only for the outcome that maximizes the likelihood ratio $F_{2,j}/F_{1,j}$, and the payment is $\frac{c_2}{F_{2,j}-F_{1,j}}$.*

Proof. Consider an optimal solution with a single positive payment (Lemma 5). The constraint $\sum_j F_{1,j}p_j - c = \sum_j F_{0,j}p_j$ must be tight. For every j, if $p_j > 0$ and the rest of the payments are zero, then $p_j = \frac{c}{F_{1,j}-F_{0,j}}$ and the expected payment is $F_{1,j}p_j = \frac{c}{1-F_{0,j}/F_{1,j}}$. To minimize this expected payment, $F_{0,j}/F_{1,j}$ must be minimized, completing the proof. □

The next proposition shows additional cases of interest in which such a simple contract is optimal. We use the following standard definition from the contract theory literature:

Definition 9. *An action a satisfies the* concavity of distribution function property (CDFP) *if for every two actions such that a's cost c_a is a convex combination of their costs, it holds that a's distribution over outcomes first-order stochastically dominates the convex combination of their distributions. A principal-agent setting satisfies CDFP if it holds for every action.*

To establish Proposition 5 we also use a lemma (Lemma 8) whose statement and proof is deferred to Section G.1.

Proposition 5. *Consider a principal-agent setting (A_n, Ω_m) with n actions and m outcomes. Then there exists an optimal contract with a single nonzero payment for the highest outcome x_m if either of the following holds:*

- *The setting satisfies MLRP and there exists an optimally-implementable action satisfying CDFP.*
- *There are $m=2$ outcomes.*
- *The setting satisfies MLRP and the actions have strictly increasing welfare.*

We divide the proof into three claims.

Claim 8. *Proposition 5 holds for CDFP.*

Proof of Claim 8. The proof follows that of Proposition 12 in [10]. Assume there exists an optimal contract with payment profile t implementing a CDFP action a_i at expected cost T. Denote the principal's expected payoff by OPT. We show there is simple contract of the required format implementing a_i at the same expected cost (thus achieving OPT).

Consider a new setting with actions $\{a_k \mid k \le i\}$. Action a_i is implementable in this setting (e.g., by transfer profile t) and MLRP holds. Thus by Lemma 8, there is a contract implementing a_i at minimum expected cost with nonzero payment only for x_m. Denote this payment by t'_m, and let T' be its expected cost. Since removing actions could not have increased the cost of implementing a_i, $T' \le T$. Denote the principal's expected payoff by OPT' and observe $OPT' \ge OPT$. Let $i' < i$ be such that the IC constraint is binding for action $a_{i'}$ (such an action must exist or T' could have been lowered).

Now add back the actions $\{a_k \mid k > i\}$. We argue that the same simple contract still implements a_i. Indeed, assume for contradiction that the simple contract implements $a_{i''}$ rather than a_i, where $i'' > i$. We will use CDFP to show that in this case, the principal's payoff is $> OPT'$, in contradiction to the optimality of the contract achieving $OPT \le OPT'$.

Since $c_{i'} < c_i < c_{i''}$, we can write $c_i = \lambda c_{i'} + (1-\lambda)c_{i''}$ for $\lambda \in [0,1]$. By CDFP, the distribution of a_i first-order stochastically dominates that of the mixed action $\lambda a_{i'} + (1-\lambda)a_{i''}$. Since the simple contract

only pays for the highest outcome, this means that the agent's utility from a_i is at least his utility from $\lambda a_{i'} + (1-\lambda)a_{i''}$. By tightness of the IC constraint for $a_{i'}$ and by the agent's preference for $a_{i''}$, the IC constraint must be tight for $a_{i''}$ as well. We know that the agent breaks ties in favor of the principal, and so the principal's payoff from $a_{i''}$ must exceed her payoff of OPT' from a_i. □

Claim 9. *Proposition 5 holds for $m = 2$.*

Proof of Claim 9. Assume there are 2 outcomes, low (ℓ) and high (h). There are n actions numbered in (strictly) increasing order of their expected outcome R_i. Observe that the probability $F_{i,h}$ of action i to achieve the high outcome is also strictly increasing. This implies that MLRP holds. Let a_i be an optimally-implementable action. If $i = 1$ (so a_i is the zero-cost action), the trivial contract with no payments is optimal; if $i = n$, the claim follows from Lemma 8. We now show that if $1 < i < n$, action a_i must satisfy CDFP, and so the proof is complete by Claim 8.

Assume for contradiction that a_i is not CDFP. So there exist two actions $a_{i'}, a_{i''}$ where $i' < i < i''$, and $\lambda \in [0,1]$, such that

$$\begin{aligned} c_i &= \lambda c_{i'} + (1-\lambda)c_{i''}, \\ F_{i,h} &< \lambda F_{i',h} + (1-\lambda)F_{i'',h}. \end{aligned}$$

This implies existence of $\lambda < \lambda' < 1$ such that

$$\begin{aligned} c_i &> \lambda' c_{i'} + (1-\lambda')c_{i''}, \\ F_{i,h} &= \lambda' F_{i',h} + (1-\lambda')F_{i'',h}. \end{aligned}$$

But by Proposition 1 this means that action a_i is not implementable, contradiction. □

Claim 10. *Proposition 5 holds for MLRP and increasing welfare.*

Proof of Claim 10. Assume there is an optimal contract with payment profile t incentivizing a_i. Let $j < m$ be the lowest outcome such that $t_j > 0$. Due to MLRP we can move weight from t_j to t_m at such a ratio that it will weakly decrease (resp., increase) the utility of all actions below (resp., above) a_i, and not change the utility of a_i. If at any point the utilities of a_i and a_k where $k > i$ become the same, then we have found a payment profile that incentivizes a_k with the same utility to the agent as in the optimal contract that incentivizes a_i. But since the welfare of a_k is larger, the principal's payoff must be larger in contradiction to the optimality of the contract. Thus we can move weight until t_j becomes 0. We conclude that with MLRP and increasing welfare, there is always an optimal contract that pays only for the highest outcome. □

G.1 Statement and Proof of Lemma 8

Lemma 8. *Consider a principal-agent setting (A_n, Ω_m) with n actions and m outcomes, for which MLRP holds. If the highest-cost action a_n is implementable, then there is an implementing contract with minimum expected payment that is a 2-partition contract. Moreover, this contract has a single nonzero payment, which is rewarded for the highest outcome x_m.*

Proof. Recall the implementability primal LP from Appendix A. Its variables are the payments $t_1, \ldots, t_m$. We need to show that there is an optimal solution to this LP for action a_n that is a simple 2-partition contract, with a single nonzero variable t_m. We achieve this by creating a *reduced* version of the primal LP with only one variable t_m, and showing that its optimal objective value is no worse (no larger) than that of the original LP. Our argument uses the dual LP, in which there is a constraint for every one of the m outcomes. Our proof proceeds as follows:

Step 1. Create a reduced dual by dropping all the constraints except for the one corresponding to the maximum outcome m.

Step 2. Solve the relaxed dual LP to optimality.

Step 3. Verify that the resulting solution is feasible (and hence optimal) for the original dual LP.

These steps are sufficient to complete the proof: The reduced dual has the same optimal objective value as the original one. Dualizing back, the optimal value of the reduced primal LP is the same as that of the original primal LP, as required. Hence there is an optimal primal solution that only uses t_m.

Step 1 The reduced dual is:

$$\begin{array}{ll} \max & \sum_{i<n} \lambda_i (c_n - c_i) \\ \text{s.t.} & \sum_{i<n} \lambda_i (F_{n,m} - F_{i,m}) \leq F_{n,m} \quad , \\ & \lambda_i \geq 0 \qquad\qquad \forall i < n. \end{array}$$

Step 2 To solve the reduced dual, we note that $(c_n - c_i) \geq 0$ for every $i < n$ (using that a_n is the highest-cost action). Thus, all coefficients in our objective are nonnegative. Also, since MLRP implies stochastic dominance, $(F_{n,m} - F_{i,m}) \geq 0$ for every $i < n$. Thus, all coefficients in our (sole) dual constraint are nonnegative. The optimal solution is then to "max out" on the action with the maximum "bang-per-buck," meaning an action in $\arg\max_i (c_n - c_i)/(F_{n,m} - F_{i,m})$. To make the dual constraint tight, we set $\lambda_i = F_{n,m}/(F_{n,m} - F_{i,m})$ (and other variables to 0).

Step 3 In this step we need to verify feasibility of the original dual constraints. Pick an arbitrary outcome $j < m$. With our choice of λ_i, feasibility becomes $F_{n,m}(F_{n,j} - F_{i,j})/(F_{n,m} - F_{i,m}) \leq F_{n,j}$. After clearing denominators and canceling terms, this reduces to MLRP (i.e., $F_{n,m}/F_{i,m} \geq F_{n,j}/F_{i,j}$), which holds by assumption. □

References

[1] M. Armstrong and J. Vickers. A model of delegated project choice. *Econometrica*, 78(1):213–244, 2010.

[2] P. D. Azar, C. Daskalakis, S. Micali, and S. M. Weinberg. Optimal and efficient parametric auctions. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 596–604, 2013.

[3] M. Babaioff and E. Winter. Contract complexity. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, page 911, 2014.

[4] M. Babaioff, M. Feldman, and N. Nisan. Free-riding and free-labor in combinatorial agency. In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory (SAGT)*, pages 109–121, 2009.

[5] M. Babaioff, M. Feldman, and N. Nisan. Mixed strategies in combinatorial agency. *Journal of Artificial Intelligence Research*, 38:339–369, 2010.

[6] M. Babaioff, M. Feldman, N. Nisan, and E. Winter. Combinatorial agency. *Journal of Economic Theory*, 147(3):999–1034, 2012.

[7] C. Bandi and D. Bertsimas. Optimal design for multi-item auctions: A robust optimization approach. *Mathematics of Operations Research*, 39(4):1012–1038, 2014.

[8] H. Bastani, M. Bayati, M. Braverman, R. Gummadi, and R. Johari. Analysis of medicare pay-for-performance contracts. Working paper, 2018.

[9] P. Bolton and M. Dewatripont. *Contract Theory*. The MIT Press, Cambridge, MA, 2004.

[10] B. Caillaud and B. E. Hermalin. Hidden-information agency. Lecture notes, 2000.

[11] G. Carroll. Robustness and linear contracts. *American Economic Review*, 105(2):536–563, 2015.

[12] G. Carroll. Robustness in mechanism design and contracting. Survey forthcoming in Annual Reviews, 2018.

[13] S. Chassang. Calibrated incentive contract. *Econometrica*, 81(5):1935–1971, 2013.

[14] L. W. Cong and Z. He. Blockchain disruption and smart contracts. Forthcoming in Review of Financial Studies, 2018.

[15] P. Diamond. Managerial incentives: On the near linearity of optimal compensation. *Journal of Political Economy*, 106(5):931–957, 1998.

[16] S. Dughmi. Algorithmic information structure design: A survey. *SIGecom Exchanges*, 15(2):2–24, 2017.

[17] S. Dughmi, N. Immorlica, and A. Roth. Constrained signaling in auction design. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1341–1357, 2014.

[18] M. Dumav and U. Khan. Moral hazard, uncertain technologies, and linear contracts. Working paper, 2017.

[19] Y. Emek, M. Feldman, I. Gamzu, R. P. Leme, and M. Tennenholtz. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computation*, 2(2):5:1–5:19, 2014.

[20] S. J. Grossman and O. D. Hart. An analysis of the principal-agent problem. *Econometrica*, 51(1): 7–45, 1983.

[21] J. D. Hartline. Mechanism design and approximation. Working book, 2017.

[22] B. Hébert. Moral hazard and the optimality of debt. *The Review of Economic Studies*, pages 55–73, 2017.

[23] C. Ho, A. Slivkins, and J. W. Vaughan. Adaptive contract design for crowdsourcing markets: Bandit algorithms for repeated principal-agent problems. *J. Artif. Intell. Res.*, 55:317–359, 2016.

[24] B. Holmström and P. Milgrom. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, 55(2):303–328, 1987.

[25] B. Holmström and P. Milgrom. Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, & Organization*, 7:24–52, 1991.

[26] A. Khodabakhsh, E. Pountourakis, and S. Taggart. Algorithmic delegation. Working paper, 2018.

[27] J. M. Kleinberg and R. Kleinberg. Delegated search approximates efficient search. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 287–302, 2018. doi: 10.1145/3219166.3219205. URL `http://doi.acm.org/10.1145/3219166.3219205`.

[28] M. Koren and A. Cohen. Incentivizing the dynamic workforce: Learning contracts in the gig-economy. Working paper, 2018.

[29] J.-J. Laffont and D. Martimort. *The Theory of Incentives*. Princeton University Press, Princeton, NJ, 2002.

[30] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, Oxford, UK, 1995.

[31] E. Maskin and J. Tirole. The principal-agent relationship with an informed principal: The case of private values. *Econometrica*, 58(2):379–409, 1990.

[32] E. Maskin and J. Tirole. The principal-agent relationship with an informed principal, II: Common values. *Econometrica*, 60(1):1–42, 1992.

[33] J. Matousek and B. Gärtner. *Understanding and Using Linear Programming*. Springer-Verlag, Berlin Heidelberg, 2007.

[34] P. B. Miltersen and O. Sheffet. Send mixed signals: Earn more, work less. In *Proceedings of the 13th ACM Conference on Economics and Computation (EC)*, pages 234–247, 2012.

[35] W. P. Rogerson. Repeated moral hazard. *Econometrica*, 53:69–76, 1985.

[36] Royal Swedish Academy of Sciences. Oliver Hart and Bengt Holmström: Contract theory. Scientific background on the Nobel prize in economic sciences, 2016.

[37] A. M. Rubinov. Global optimization: Envelope representation. In C. A. Floudas and P. M. Pardalos, editors, *Encyclopedia of Optimization*. Springer, Boston, MA, 2001.

[38] S. Shavell. Risk sharing and incentives in the principal and agent relationship. *Bell Journal of Economics*, 10:55–73, 1979.

[39] A. M. Spence. Job market signaling. *Quarterly Journal of Economics*, 87:355–374, 1973.

[40] S. Tadelis and I. Segal. Lectures in contract theory. Lecture notes, 2005.

[41] Y. Yu and X. Kong. Robust contract designs: Linear contracts for sales force compensation. Working paper, 2018.