Robust Contracting under Distributional Uncertainty*

Jiangtao Li† Kexin Wang‡

November 21, 2023

Abstract

We study the design of contracts when the principal has limited statistical information about the output distributions induced by the agent's actions. In the baseline model, we consider a principal who only knows the mean of the output distribution for each action. The mean restrictions allow for a large set of profiles of output distributions, including some extreme output distributions that can be used to establish the robust optimality of monotone affine contracts. Motivated by this, we study the set of distributions that can be used to establish the robust optimality of the monotone affine contracts. This facilitates the understanding for the use of monotone affine contracts in settings with more restrictions on the output distributions. Our main result shows that the optimality of monotone affine contracts persists even if the principal has access to other information about the output distributions, such as the information that the output distribution induced by each action has full support.

Keywords: Robust mechanism design, robust contracting, distributional uncertainty, monotone affine contracts, duality approach.

*We thank Takashi Kunimoto for helpful comments.

†School of Economics, Singapore Management University, jtli@smu.edu.sg

‡School of Economics, Singapore Management University, kexinwang.2020@phdecons.smu.edu.sg

1 Introduction

In traditional models of the principal-agent problem, it is typically assumed that the principal has detailed information about the environment, such as the actions available to the agent and the exact consequence of the agent taking these actions. It has been well documented that the optimal contracts these models predict often take complicated functional forms (see, for example, Grossman and Hart (1983) and Bolton and Dewatripont (2005, Chapter 4)). Nevertheless, in many realistic settings, the principal presumably only has access to some limited statistical information about the environment, and theoretical conclusions derived from traditional models can be fragile—contracts that are optimized to perform well when the assumptions are exactly true may fail miserably in the much more frequent cases when the assumptions are untrue.

This paper studies the design of contracts when the principal has limited statistical information about the output distributions induced by the agent's actions. As in standard moral hazard models, the principal contracts with an agent, who is to take a costly action that leads to a stochastic output. The action is not observable to the principal; only the resulting output is. The principal incentives the agent using a contract that specifies the payment to the agent for each level of output, and maximizes her expected payoff—the output minus the wage paid to the agent. The set of actions, the possible levels of output, and the cost for taking each action are assumed to be common knowledge between the principal and the agent. The consequence of the agent taking these actions—the output distribution for each action—is known to the agent but not perfectly known to the principal. The principal ranks contracts according to their payoff guarantee—the worst-case expected payoff where the worst case is taken over all profiles of stochastic output that are perceived to be plausible. A contract is optimal if it generates the highest payoff guarantee.

The situation in which the principal only has limited statistical information about the output distributions is ubiquitous. For example, in the case of salesforce compensation, the amount of data that firms can access for generating the exact demand distribution is often limited. Firms face ambiguity regarding the influence of the sales agent's effort, and they often rely on accessible statistical information such as the mean rather than the exact distribution to make decisions. The recent literature on robust mechanism design also offers many motivations for limited statistical information about the prior distribution in various settings; see Carrasco, Luz, Kos, Messner, Monteiro, and Moreira (2018), Koçyiğit, Iyengar, Kuhn, and Wiesemann (2020), Brooks and Du (2021a), Suzdaltsev (2022), He and Li (2022), Zhang (2022), and Che (2022), among many others.

In the baseline model, we assume that the principal knows the mean of the output distribution for each action but does not have reliable information about other aspects of the output distributions. The principal perceives a profile of stochastic output to be plausible as long as it is consistent with the mean restrictions. In this setting, a monotone affine contract, which pays the agent some nonnegative fraction of the output plus a fixed payment, could perfectly hedge the principal's uncertainty in the sense that a monotone affine contract induces the same action for the agent and generates the same expected payoff for the principal, regardless of the actual profile of output distributions. Theorem 1 shows that it is optimal for the principal to adopt a monotone affine contract.

Since a monotone affine contract generates the same expected payoff for the principal regardless of the actual profile of stochastic output, to establish the optimality of monotone affine contracts, it suffices to identify a plausible profile of stochastic output against which a monotone affine contract is optimal in the Bayesian framework. This is because, if such a profile exists, then the payoff guarantee of any feasible contract cannot exceed its expected payoff against this profile, which in turn cannot exceed the expected payoff of the optimal monotone affine contract against this profile. It is straightforward to show that there exists a profile (where the output distribution for each action has binary support) against which an affine contract is optimal in the Bayesian framework. We proceed to show that (1) it is suboptimal to incentivize the agent to choose an action that can be implemented by an affine contract but cannot be implemented by a monotone affine contract, and (2) for any action that is implementable by monotone affine contracts, the least expected cost of implementing this action using affine contracts and monotone affine contracts are the same. Thus, against this profile, it is optimal to use a monotone affine contract in the Bayesian framework.

In the baseline model, the principal is assumed to only know the expected output for each action. This allows for a large set of profiles of stochastic output, including the binary support distributions used to prove Theorem 1. Nevertheless, there are various scenarios in which the principal may have access to additional information about the output distributions. Such information imposes more restrictions on the set of plausible profiles, which corresponds to a smaller ambiguity set that might rule out the binary support distributions as plausible.

To accommodate additional information that might be available to the principal, we conduct a more systematic analysis in Section 4. Adopting a duality approach, we show in Theorem 2 that as long as the output distribution for the action with the smallest mean satisfies a simple condition, there exists a profile consistent with the mean restriction against which a monotone affine contract is optimal in the Bayesian framework. In other words, Theorem 2 identifies a large collection of profiles that could be used to establish the optimality of monotone affine contracts when the principal only knows the mean.

Besides the simplicity of the condition, Theorem 2 serves as a toolbox to prove the optimality of monotone affine contracts when there is additional information available to the principal. In particular, we show that, very generally, even if we require the output distribution induced by each action to have full support, the optimality of monotone affine contracts remains. Clearly, if the output distribution for each action is required to have full support, then the binary support distributions used to prove Theorem 1 would no longer be perceived to be plausible. Nevertheless, applying Theorem 2, we could easily show the optimality of monotone affine contracts by identifying an output distribution for the action with the smallest mean that satisfies the condition in Theorem 2 and the full support requirement. We also present two more examples, on additional information about variance and quantile respectively, to further illustrate how Theorem 2 can be applied.

The remainder of this introduction discusses some related literature. Section 2 presents the notation, concepts, and the model. Section 3 establishes the optimality of monotone affine contracts when the principal only knows the mean. Section 4 offers a more systematic analysis that accommodates additional information that might be available to the principal. Section 5 discusses the difference of our model from Diamond (1998) and the knowledge of higher-order moments.

1.1 Related literature

This paper joins the recently growing literature exploring the design of contracts when the principal does not have detailed information about the environment, starting with

Hurwicz and Shapiro (1978), and more recently, Carroll (2015), Carroll and Meng (2016a), Carroll and Meng (2016b), Li and Kirshner (2021), Kambhampati (2023), Kambhampati, Toikka, and Vohra (2023), Rosenthal (2023), among others.

Carroll (2015) considers a moral hazard problem where the principal is uncertain as to what the agent can and cannot do: she knows some actions available to the agent, but other, unknown actions may also exist. He shows that the optimal deterministic contract is linear. Kambhampati (2023) shows that, at the same level of generality as Carroll (2015), the principal can strictly increase her payoff guarantee by randomizing over deterministic contracts. Kambhampati, Toikka, and Vohra (2023) further show that there exists an optimal contract that is simple—a uniform lottery over two linear contracts. The principal in our model faces a different kind of uncertainty—the output distributions, and we establish the optimality of monotone affine contracts. In particular, randomization will not benefit the principal in our setting.

As in this paper, Carroll and Meng (2016a), Carroll and Meng (2016b), Li and Kirshner (2021), and Rosenthal (2023) consider uncertainty about how the agent's actions translate into output. Carroll and Meng (2016b) consider a model in which the agent privately observes the realization of some shock that affects output before choosing how much effort to exert, while the principal only observes total output. The principal is uncertain about the exact distribution of the additive shock and only knows its mean. There are important differences between their model and ours. The principal in Carroll and Meng (2016b) only has uncertainty about the distribution of additive noise, and the output for all effect levels is affected by the same additive noise. In our baseline model where we consider only the mean restrictions, the output distribution under each action is only restricted to be consistent with the given mean. Furthermore, our analysis covers additional information available to the principal, such as the information that the output distribution for each action has full support. Carroll and Meng (2016a) study a moral hazard problem in which the principal has local uncertainty about how the agent's actions translate into output. Li and Kirshner (2021) consider a moral hazard problem with two-sided ambiguity, where both the principal and the agent face uncertainty about the output distributions. Rosenthal (2023) considers a principal who understands the relationship between effort and output, but does not know the agent's beliefs about the production technology. Our paper differs in (at least) two important aspects. First, while the principal in Rosenthal (2023) has her own belief about how output is distributed and uses this belief when evaluating any given contract, the principal in our model does not know how output is distributed. Second, while Rosenthal (2023) is primarily concerned with the mean restrictions, our main result, Theorem 2, can be used to establish the robust optimality of monotone affine contracts even if the principal has additional information besides the mean.

Diamond (1998, Section 5) establishes the optimality of linear contracts in the environment in which the agent can either choose one action that produces a low expected output or the other action that produces a higher expected output, and the agent can freely distribute the probabilities in any way that preserves the expected output of each action. The key difference of our paper from Diamond (1998) is that, while Diamond (1998) allows the agent to choose probability distributions to maximize his own expected payoff, Nature chooses probability distributions to minimize the principal's expected payoff in our model. Section 5.1 presents a concrete example to illustrate the difference between the agent-maximizing problem in Diamond (1998) and the nature-minimizing problem in our model.

More broadly, our paper is related to the literature on robust mechanism design. Among others, Chung and Ely (2007), Du (2018), Chen and Li (2018), Brooks and Du (2021b), Yamashita and Zhu (2022), and Brooks and Du (2023) study a revenue maximizing designer who knows the joint distribution of the agents' payoff-relevant information but has non-Bayesian uncertainty about their beliefs. To the best of our knowledge, Carrasco, Luz, Kos, Messner, Monteiro, and Moreira (2018) offer the first systematic analysis on the implications of moment conditions—they study the revenue maximization problem of a seller who is partially informed about the distribution of the buyer's valuation, only knowing a finite number of moments. Neeman (2003), Koçyiğit, Iyengar, Kuhn, and Wiesemann (2020), Brooks and Du (2021a), Suzdaltsev (2022), He and Li (2022), Zhang (2022), and Che (2022) consider the robust auction design problem where the auctioneer has limited statistical information about the join distribution of the bidders' valuations.

2 Preliminaries

A principal (she) contracts with an agent (he), who is to take a costly action that leads to a stochastic output. Only the resulting output is observable to the principal and can be contracted upon. Any contract must specify how much the agent is paid for each level of output. We assume one-sided limited liability: the agent can never be paid less than zero. Both parties are financially risk-neutral.

There are $K \geq 2$ actions available to the agent, each corresponding to a different level of effort the agent can exert. Let $A = \{a_1, a_2, \ldots, a_K\}$ denote the set of actions. There are $N \geq 3$ levels of output that might be realized, denoted $q_1, q_2, \ldots, q_N$, where $0 \leq q_1 < q_2 < \ldots < q_N$. If the agent takes the action a_k, then the agent incurs a cost $c(a_k) \geq 0$, and the vector $p(a_k) = (p_1(a_k), p_2(a_k) \ldots, p_N(a_k))$ specifies the resulting probability distribution over output, where $p_i(a_k) \geq 0$ for all $i = 1, 2, \ldots, N$ and $\sum_{i=1}^{N} p_i(a_k) = 1$. The agent has an outside option a_0, which is costless for the agent and generates zero output.

The set of actions A, the possible levels of output $\{q_1, q_2, \ldots, q_N\}$, and the cost function $c : A \to \mathbb{R}_+$ are common knowledge between the two parties. The profile of stochastic output $P = \{p(a_1), p(a_2), \ldots, p(a_K)\}$ is known to the agent but not perfectly known to the principal. In the baseline model, we consider a principal who knows the mean $m(a_k)$ of the output distribution for each action a_k and has non-Bayesian uncertainty about other aspects of the profile of stochastic output. Let $m = (m(a_1), m(a_2), \ldots, m(a_K))$. For ease of exposition, we assume that $q_1 < m(a_1) < m(a_2) < \ldots < m(a_K) < q_N$. From the principal's perspective, a profile of stochastic output is perceived to be plausible as long as it is consistent with m. We write

$$\mathcal{P}(m) = \Big\{ P = \{p(a_1), p(a_2), \ldots, p(a_K)\} \Big| \; p(a_k) \in \mathbb{R}_+^N, \sum_{i=1}^{N} p_i(a_k) = 1, \text{ and } \sum_{i=1}^{N} p_i(a_k) q_i = m(a_k), \forall k = 1, 2, \ldots, K \Big\}.$$

to denote the collection of profiles of stochastic output that are consistent with m.

A contract w is a vector of N nonnegative numbers, specifying how much the agent is paid for each level of output. The nonnegativity requirement captures the one-sided limited liability constraint. For any given contract w, the agent selects an action that maximizes his expected payoff under this contract. Formally, if the profile of stochastic output is P, then the set of actions that the agent is willing to choose is

$$A^*(w|P) = \arg\max_{a \in A} \Big(\sum_{i=1}^{N} w_i p_i(a) - c(a) \Big) \quad \text{subject to } \sum_{i=1}^{N} w_i p_i(a) - c(a) \geq 0.$$

If the agent is indifferent among several actions, as is standard in the literature, we assume he maximizes the principal's expected payoff. Thus, if the profile of stochastic output is P, the principal's expected payoff under a given contract w is

$$V(w|P) = \max_{a \in A^*(w|P)} \sum_{i=1}^{N} (q_i - w_i) p_i(a).$$

Finally, since the principal only knows the mean of the output distribution for each action, she ranks contracts according to their payoff guarantee—the worst-case expected payoff where the worst case is taken over all profiles of stochastic output that are consistent with m:

$$V(w) = \inf_{P \in \mathcal{P}(m)} V(w|P).$$

We focus on the principal's problem, namely to maximize $V(w)$. Denote by V^* the principal's highest payoff guarantee.

3 The baseline model

In the baseline model, we consider a principal who has no additional information about the actual profile of stochastic output beyond the mean of the output distribution for each action.

3.1 The Bayesian framework

As a preparation, we consider below the Bayesian framework in which the profile of stochastic output P is also known to the principal. In this case, the principal solves the following maximization problem:

$$\begin{aligned} & \max_{a \in A, w} \sum_{i=1}^{N} (q_i - w_i) p_i(a) && \text{(grand max)} \\ \text{subject to} \quad & \sum_{i=1}^{N} w_i \, p_i(a) - c(a) \geq \sum_{i=1}^{N} w_i p_i(a') - c(a'), \ \forall a' \in A \backslash \{a\}, \\ & \sum_{i=1}^{N} w_i p_i(a) - c(a) \geq 0, \\ & w_i \geq 0, \ \forall i = 1, 2, \ldots, N. \end{aligned}$$

Let $V^*(P)$ denote the principal's maximized payoff solved from (grand max).

Following Grossman and Hart (1983), we could solve the principal's problem (grand max) in two steps. We say that an action a is implementable against the profile P if there exists some contract w such that

$$\sum_{i=1}^{N} w_i p_i(a) - c(a) \geq \sum_{i=1}^{N} w_i p_i(a') - c(a'), \forall a' \in A \backslash \{a\},$$
$$\sum_{i=1}^{N} w_i p_i(a) - c(a) \geq 0.$$

Let A_P denote the set of actions implementable against the profile P. We consider first, for each implementable action, the least expected cost of implementing this action. For each action $a \in A_P$, let $\Psi(a|P)$ be the least expected cost of implementing a against the profile P. Formally, $\Psi(a|P)$ is solved from the following minimization problem:

$$\min_{w} \sum_{i=1}^{N} w_i p_i(a) \tag{min}$$
$$\text{subject to} \quad \sum_{i=1}^{N} w_i p_i(a) - c(a) \geq \sum_{i=1}^{N} w_i p_i(a') - c(a'), \forall a' \in A \backslash \{a\},$$
$$\sum_{i=1}^{N} w_i p_i(a) - c(a) \geq 0,$$
$$w_i \geq 0, \forall i = 1, 2, \ldots, N.$$

We then consider which action should be implemented, that is, we solve for the action that generates the highest expected payoff for the principal. The principal's maximized payoff can be derived as follows:

$$V^*(P) = \max_{a \in A_P} \{m(a) - \Psi(a|P)\}$$

3.2 Payoff guarantee of monotone affine contracts

A special class of contracts is the class of affine contracts. An affine contract takes the form $w_i = \alpha q_i - \beta$ for constants $\alpha, \beta \in \mathbb{R}$ for all i. We are particularly interested in affine contracts that have nonnegative slopes ($\alpha \geq 0$). We call such contracts monotone affine contracts.

Notably, any affine contract generates the same expected payoff to the principal, regardless of the actual profile of stochastic output, as long as it is consistent with m. To wit, fix an arbitrary affine contract $w_i = \alpha q_i - \beta$. For any $P \in \mathcal{P}(m)$, the agent's expected payoff from an action a is

$$\sum_{i=1}^{N} w_i p_i(a) - c(a) = \sum_{i=1}^{N} (\alpha q_i - \beta)\, p_i(a) - c(a) = \alpha m(a) - \beta - c(a),$$

which depends on the profile of stochastic output only through m. Thus, the set of actions that the agent is willing to choose is the same for any $P \in \mathcal{P}(m)$ and given by[1]

$$\begin{aligned} A^*(w) = &\arg\max_{a\in A} \Big(\alpha m(a) - \beta - c(a)\Big) \\ &\text{subject to } \alpha m(a) - \beta - c(a) \geq 0. \end{aligned}$$

Consequently, regardless of the actual profile of stochastic output, the principal would obtain the same expected payoff:

$$V(w) = V(w|P) = \max_{a\in A^*(w)} \Big((1-\alpha)m(a) + \beta\Big),\ \forall P \in \mathcal{P}(m).$$

If the principal can only choose from the class of affine contracts, then her payoff guarantee maximization problem reduces to:

$$\begin{aligned} &\max_{a\in A,\alpha,\beta} (1-\alpha)m(a) + \beta && \text{(grand max-Affine)} \\ \text{subject to}\quad &\alpha m(a) - c(a) \geq \alpha m(a') - c(a'),\ \forall a' \in A\backslash\{a\}, \\ &\alpha m(a) - \beta - c(a) \geq 0, \\ &\alpha q_1 - \beta \geq 0, \\ &\alpha q_N - \beta \geq 0. \end{aligned}$$

If the principal can only choose from the class of monotone affine contracts, then her payoff guarantee maximization problem reduces to:

$$\begin{aligned} &\max_{a\in A,\alpha,\beta} (1-\alpha)m(a) + \beta && \text{(grand max-Monotone)} \\ \text{subject to}\quad &\alpha m(a) - c(a) \geq \alpha m(a') - c(a'),\ \forall a' \in A\backslash\{a\}, \\ &\alpha m(a) - \beta - c(a) \geq 0, \\ &\alpha q_1 - \beta \geq 0, \end{aligned}$$

[1] Here, we write $A^*(w)$ rather than $A^*(w|P)$ to highlight that this set of actions does not depend on the actual profile of stochastic output.

$$\alpha \geq 0.$$

Denote by V_A^* (resp. V_M^*) the principal's highest payoff guarantee from choosing an affine contract (resp. a monotone affine contract). Clearly, $V^* \geq V_A^* \geq V_M^*$.

We say that an action a is implementable by affine contracts (resp. monotone affine contracts) if there exists some affine contract (resp. monotone affine contract) such that

$$\alpha m(a) - c(a) \geq \alpha m(a') - c(a'), \forall a' \in A \backslash \{a\},$$
$$\alpha m(a) - \beta - c(a) \geq 0.$$

Let A_A (resp. A_M) denote the set of actions implementable by affine contracts (resp. monotone affine contracts). The above arguments also imply that the least expected cost of implementing an action using affine contracts or monotone affine contracts depends on the profile of stochastic output only through m. Thus, for notational simplicity, we drop the dependence on P. For each action $a \in A_A$, let $\Psi_A(a)$ denote the least expected cost of implementing a using affine contracts, solved from the following cost minimization problem:

$$\begin{aligned} & \min_{\alpha,\beta} \alpha m(a) - \beta && \text{(min-Affine)} \\ \text{subject to} \quad & \alpha m(a) - c(a) \geq \alpha m(a') - c(a'), \forall a' \in A \backslash \{a\}, \\ & \alpha m(a) - c(a) - \beta \geq 0, \\ & \alpha q_1 - \beta \geq 0, \\ & \alpha q_N - \beta \geq 0. \end{aligned}$$

For each action $a \in A_M$, let $\Psi_M(a)$ be the least expected cost of implementing a using monotone affine contracts, solved from the following cost minimization problem:

$$\begin{aligned} & \min_{\alpha,\beta} \alpha m(a) - \beta && \text{(min-Monotone)} \\ \text{subject to} \quad & \alpha m(a) - c(a) \geq \alpha m(a') - c(a'), \forall a' \in A \backslash \{a\}, \\ & \alpha m(a) - c(a) - \beta \geq 0, \\ & \alpha q_1 - \beta \geq 0, \\ & \alpha \geq 0. \end{aligned}$$

Following Grossman and Hart (1983), the highest payoff guarantee from choosing an affine contract and a monotone affine contract can be derived as follows:

$$V_A^* = \max_{a \in A_A} \{m(a) - \Psi_A(a)\}, \text{ and } V_M^* = \max_{a \in A_M} \{m(a) - \Psi_M(a)\}.$$

3.3 Optimality of monotone affine contracts

Since for any monotone affine contract, the principal would obtain the same expected payoff, regardless of the actual profile of stochastic output, one might expect that there exists a monotone affine contract that maximizes V. Theorem 1 below shows exactly this.

Theorem 1. *There exists a monotone affine contract that maximizes V. Formally,*

$$V^* = V_M^*.$$

To prove Theorem 1, it suffices to identify a particular profile of stochastic output $P \in \mathcal{P}(m)$ against which a monotone affine contract is optimal. This is because, if there exists such a profile of stochastic output, then the payoff guarantee of any contract obviously cannot exceed its expected payoff against this particular profile of stochastic output, which in turn cannot exceed the expected payoff of the optimal monotone affine contract against this particular profile of stochastic output. But the expected payoff of the optimal monotone affine contract against this particular profile of stochastic output is simply its payoff guarantee, since any monotone affine contract generates the same expected payoff for the principal against every $P \in \mathcal{P}(m)$.

The structure of the proof is as follows. Step (1) identifies a profile of stochastic output against which an affine contract is optimal. This step is rather straightforward and follows from the following observation. Since the principal perceives any profile $P \in \mathcal{P}(m)$ to be plausible, from the principal's perspective, it could well be the case that for any action a, only two levels of output q_1 and q_N can be realized. Against this profile of stochastic output, only the wages following q_1 and q_N matter for the agent's incentive. Therefore, the principal can simply choose an affine contract. Formally, Step (1) shows $V^* = V_A^*$. We are left to show that $V_A^* = V_M^*$. Recall that

$$V_A^* = \max_{a \in A_A} \{m(a) - \Psi_A(a)\}, \text{ and } V_M^* = \max_{a \in A_M} \{m(a) - \Psi_M(a)\}.$$

Step (2) shows that it is suboptimal to incentivize the agent to choose an action that can be implemented by an affine contract but cannot be implemented by a monotone affine contract. Step (3) shows that for any action that is implementable by monotone affine contracts, the least expected costs of implementing this action using affine contracts and monotone affine contracts are the same. Thus, $V_A^* = V_M^*$. The detailed proof can be found in Appendix A.

4 Main results

In the baseline model, the principal is assumed to only know the expected output for each action. This allows for a large set of profiles of stochastic output—any $P \in \mathcal{P}(m)$ is perceived to be plausible, including the binary support distributions used in the proof of Theorem 1. Nevertheless, there are various scenarios in which the principal may have access to additional information about the output distributions. Such information imposes more restrictions on the set of plausible profiles, which corresponds to a smaller ambiguity set that might rule out the binary support distributions as plausible.

In this section, we extend our analysis in the baseline model and provide more systematic analysis to show the optimality of monotone affine contracts for a worst-case minded principal when she is uncertain about the actual profile of stochastic output. This analysis provides a foundation for the use of monotone affine contracts even when there is additional information available to the principal.

In Section 4.1, we adopt a duality approach and establish conditions under which there exists a monotone affine contract that maximizes V. In Section 4.2 and Section 4.3, we illustrate how this condition can be applied under different kinds of additional information. Section 4.2 shows that, very generally, even if we require the output distribution induced by each action to have full support, the optimality of monotone affine contracts remains. In Section 4.3, we present two more examples, on additional information about variance and quantile respectively, to further illustrate how the condition can be applied.

Throughout this section, we make the following assumption about the cost function $c(\cdot)$ and the mean restriction $m(\cdot)$.

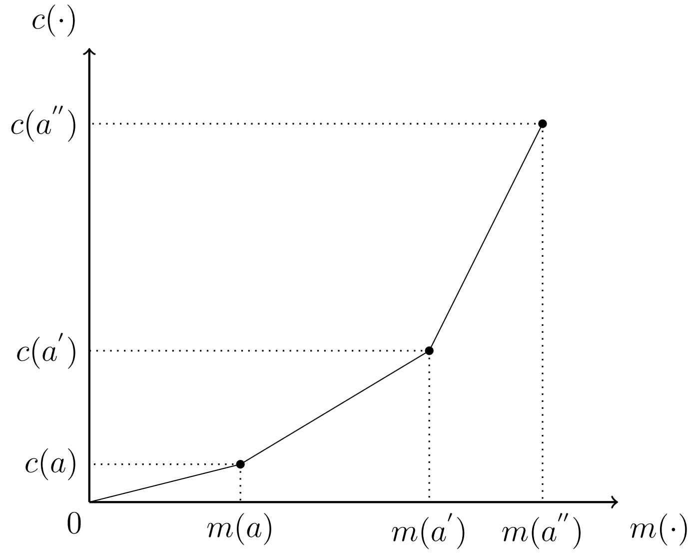

Figure 1: Assumption 1

Assumption 1. *(1) For any $a, a' \in A$, we have*

$$\frac{c(a') - c(a)}{m(a') - m(a)} > 0, \tag{1}$$

(2) For any $a, a', a'' \in A \cup \{a_0\}$ such that $m(a'') > m(a') > m(a)$, we have

$$\frac{c(a'') - c(a')}{m(a'') - m(a')} > \frac{c(a') - c(a)}{m(a') - m(a)}. \tag{2}$$

Loosely speaking, Assumption 1 says that the cost function $c(a)$ can be viewed as a strictly increasing and strictly convex function of $m(a)$; see Figure 1 for an illustration.

Assumption 1 ensures that every action can be implemented using a monotone affine contract. Recall that an action a_k is implementable by monotone affine contracts if there exists some constants $\alpha \geq 0, \beta$ such that

$$\begin{aligned}
&\alpha m(a) - c(a) \geq \alpha m(a') - c(a'), \forall a' \in A \backslash \{a\}, \\
&\alpha m(a) - c(a) - \beta \geq 0.
\end{aligned}$$

Assumption 1 implies that for any $a, a', a'' \in A \cup \{a_0\}$ such that $m(a'') > m(a') >$

$m(a)$, we have[2]

$$\frac{c(a'') - c(a')}{m(a'') - m(a')} > \frac{c(a'') - c(a)}{m(a'') - m(a)} > \frac{c(a') - c(a)}{m(a') - m(a)}. \tag{3}$$

It is straightforward to verify that any action a_k with $k < K$ can be implemented by the monotone affine contract

$$\alpha = \frac{c(a_{k+1}) - c(a_k)}{m(a_{k+1}) - m(a_k)} > 0, \ \beta = 0,$$

and the action a_K can be implemented by the monotone affine contract

$$\alpha = \frac{c(a_K) - c(a_{K-1})}{m(a_K) - m(a_{K-1})} > 0, \ \beta = 0.$$

4.1 Optimality of monotone affine contracts

The implementability of any action $a \in A$ using monotone affine contracts greatly simplifies the comparison of the principal's optimization problems between using any feasible contract and using only monotone affine contracts. Against a given profile of stochastic output P, the principal's maximized payoff $V^*(P)$ solved from (grand max) can be derived as follows:

$$V^*(P) = \max_{a \in A} \{m(a) - \Psi(a|P)\}.$$

The optimal monotone affine contract generates the following expected payoff:

$$V_M^* = \max_{a \in A} \{m(a) - \Psi_M(a)\}.$$

Theorem 2 below presents a condition on the profile of stochastic output $\tilde{P}$ that guarantees the equivalence of the least expected costs $\Psi(a|\tilde{P})$ and $\Psi_M(a)$ for all $a \in A$. Subsequently, we can establish the optimality of monotone affine contract against $\tilde{P}$ in the Bayesian framework.

Theorem 2. *Let $\tilde{\mathcal{P}}$ denote the collection of profiles constructed as follows: (1) $\tilde{p}(a_1)$*

---

[2]This is purely algebraic. For any $x, y, z, w > 0$,

$$\frac{x}{y} > \frac{z}{w} \implies \frac{x}{y} > \frac{x+z}{y+w} > \frac{z}{w}.$$

To get (3), set $x = c(a'') - c(a')$, $y = m(a'') - m(a')$, $z = c(a') - c(a)$, and $w = m(a') - m(a)$.

is a probability distribution that satisfies

$$\sum_{i=1}^{N} q_i \tilde{p}_i(a_1) = m(a_1),$$
$$\tilde{p}_1(a_1) \geq \frac{m(a_K) - m(a_1)}{m(a_K) - q_1}, \text{ and}$$

(2) $\tilde{p}(a_k)$ *for any* $k > 1$ *is derived from* $\tilde{p}(a_1)$ *as follows:*

$$\tilde{p}_1(a_k) = \frac{m(a_k) - q_1}{m(a_1) - q_1} \tilde{p}_1(a_1) - \frac{m(a_k) - m(a_1)}{m(a_1) - q_1},$$
$$\tilde{p}_i(a_k) = \frac{m(a_k) - q_1}{m(a_1) - q_1} \tilde{p}_i(a_1), \ \forall i = 2, 3, \ldots, N.$$

We have $\tilde{\mathcal{P}} \subseteq \mathcal{P}(m)$*. Under Assumption 1,* $V^*(\tilde{P}) = V_M^*$ *for any* $\tilde{P} \in \tilde{\mathcal{P}}$*.*[3]

Theorem 2 has several appealing features. First, the condition is extremely easy to check because it merely imposes certain requirements on the output distribution induced by a_1 (the output distributions for the other actions are completely pinned down by the output distribution induced by a_1). Second, we can use Theorem 2 as a toolbox to handle various kinds of additional information beyond the mean restriction. To wit, suppose that in addition to the mean restriction, the principal has additional information that the actual profile of stochastic output must also lie in some set $\bar{\mathcal{P}}$. We can readily establish the optimality of monotone affine contracts if $\tilde{\mathcal{P}} \cap \bar{\mathcal{P}} \neq \emptyset$. Indeed, as we show in the next section, if the principal has additional information that the output distribution induced by each action has full support, then a monotone affine contract maximizes the principal's payoff guarantee. We further illustrate how to utilize this theorem when the principal has other kinds of additional information such as variance and quantile information in Section 4.3.

4.2 Full support output distributions

Suppose that, in addition to the mean restriction, the principal knows that the output distribution $p(a)$ for each action $a \in A$ has full support. That is, the principal perceives any $P \in \mathcal{P}(m) \cap \mathcal{P}(F)$ to be plausible, where

$$\mathcal{P}(F) = \Big\{ P = \{p(a_1), p(a_2), \ldots, p(a_K)\} \Big|$$

[3]Clearly, the binary support distributions used to prove Theorem 1 is contained in $\tilde{\mathcal{P}}$.

$$p(a_k) \in \mathbb{R}_+^N,\ \sum_{i=1}^{N} p_i(a_k) = 1,\ \forall k = 1, 2, \ldots, K,$$
$$p_i(a_k) > 0,\ \forall k = 1, 2, \ldots, K,\ \forall i = 1, 2, \ldots, N \Big\}.$$

Obviously, the binary support distributions used to prove Theorem 1 is no longer plausible. Nevertheless, applying Theorem 2, we could establish the optimality of monotone affine contracts by showing that $\tilde{\mathcal{P}} \cap \mathcal{P}(F) \neq \emptyset$.

Theorem 3. *Under Assumption 1, there exists a profile of stochastic output* $\tilde{P} \in \mathcal{P}(m) \cap \mathcal{P}(F)$ *such that* $V^*(\tilde{P}) = V_M^*$.

By Theorem 2, it suffices to show that $\tilde{\mathcal{P}} \cap \mathcal{P}(F) \neq \emptyset$. We explicitly construct a probability distribution $\tilde{p}(a_1)$ (contained in Appendix C) such that

$$\begin{aligned}
&\sum_{i=1}^{N} q_i \tilde{p}_i(a_1) = m(a_1), \\
&\tilde{p}_1(a_1) > \frac{m(a_K) - m(a_1)}{m(a_K) - q_1}, \\
&\tilde{p}_i(a_1) > 0,\ \forall i = 2, 3, \ldots, N.
\end{aligned}$$

Since $\tilde{p}(a_1)$ satisfies the condition in Theorem 2, the profile $\tilde{P}$ constructed as follows is contained in $\tilde{\mathcal{P}}$:

$$\begin{aligned}
\tilde{p}_1(a_k) &= \frac{m(a_k) - q_1}{m(a_1) - q_1} \tilde{p}_1(a_1) - \frac{m(a_k) - m(a_1)}{m(a_1) - q_1},\ \forall k > 1 \\
\tilde{p}_i(a_k) &= \frac{m(a_k) - q_1}{m(a_1) - q_1} \tilde{p}_i(a_1),\ \forall k > 1,\ \forall i = 2, 3, \ldots, N.
\end{aligned}$$

Since $\tilde{p}_1(a_1) > \frac{m(a_K) - m(a_1)}{m(a_K) - q_1}$,

$$\tilde{p}_1(a_k) = \frac{m(a_k) - q_1}{m(a_1) - q_1} \tilde{p}_1(a_1) - \frac{m(a_k) - m(a_1)}{m(a_1) - q_1} > \frac{m(a_K) - m(a_k)}{m(a_K) - q_1} \geq 0, \forall k > 1.$$

Since $\tilde{p}_i(a_1) > 0,\ \forall i = 2, 3, \ldots, N$, we have

$$\tilde{p}_i(a_k) = \frac{m(a_k) - q_1}{m(a_1) - q_1} \tilde{p}_i(a_1) > 0,\ \forall k > 1,\ \forall i = 2, 3, \ldots, N.$$

Thus, $\tilde{P} \in \tilde{\mathcal{P}} \cap \mathcal{P}(F)$.

4.3 Additional information on variance/ quantile

In many settings, the principal could have access to some statistical features of the output distributions. In this section, we present two examples to further illustrate how to apply Theorem 2 to establish the optimality of monotone affine contracts with different kinds of additional information. Example 1 considers the scenario in which the principal has additional knowledge on the upper bound of the variance of the output distribution for each action, and Example 2 considers the scenario in which there are some restrictions on the quantile of each output distribution.

The two examples share the same basic environment. There are three output levels: $q_1 = 0.5, q_2 = 1$, and $q_3 = 2$. There are two actions available to the agent: $c(a_1) = 0.25, m(a_1) = 0.7$ and $c(a_2) = 1.65, m(a_2) = 1.8$.

Let $\tilde{p}(a_1)$ be a probability distribution consistent with $m(a_1)$. Since $\sum_{i=1}^{3} \tilde{p}_i(a_1) = 1$ and $\sum_{i=1}^{3} \tilde{p}_i(a_1) q_i = m(a_1)$, we have

$$\begin{aligned}
\tilde{p}_1(a_1) &= \frac{1}{q_2 - q_1}\Big((q_3 - q_2)\tilde{p}_3(a_1) + q_2 - m(a_1)\Big), \\
\tilde{p}_2(a_1) &= \frac{1}{q_2 - q_1}\Big(m(a_1) - q_1 - (q_3 - q_1)\tilde{p}_3(a_1)\Big).
\end{aligned}$$

For $\tilde{p}(a_1)$ to satisfy the condition in Theorem 2, we must have

$$\max\left\{0, \frac{(m(a_1) - q_1)(m(a_2) - q_2)}{(m(a_2) - q_1)(q_3 - q_2)}\right\} \leq \tilde{p}_3(a_1) \leq \frac{m(a_1) - q_1}{q_3 - q_1}.$$

For any such $\tilde{p}_3(a_1)$, the profile $\tilde{P}$ constructed as follows is contained in $\tilde{\mathcal{P}}$:

$$\begin{aligned}
\tilde{p}_1(a_1) &= \frac{1}{q_2 - q_1}\Big((q_3 - q_2)\tilde{p}_3(a_1) + q_2 - m(a_1)\Big), \\
\tilde{p}_2(a_1) &= \frac{1}{q_2 - q_1}\Big(m(a_1) - q_1 - (q_3 - q_1)\tilde{p}_3(a_1)\Big), \\
\tilde{p}_1(a_2) &= \frac{m(a_2) - q_1}{m(a_1) - q_1}\,\tilde{p}_1(a_1) - \frac{m(a_2) - m(a_1)}{m(a_1) - q_1}, \\
\tilde{p}_i(a_2) &= \frac{m(a_2) - q_1}{m(a_1) - q_1}\,\tilde{p}_i(a_1),\ \forall i = 2, 3.
\end{aligned}$$

Example 1 (Upper bound on variance)**.** Besides the mean restriction, the principal has additional information that for each action, the variance of the output distribution cannot exceed an upper bound: $var(a_1) \leq 0.25, var(a_2) \leq 0.2$. Let $\mathcal{P}_{Var}$ denote the collection of profiles that are consistent with these variance requirements.

It is easy to check that the binary support distributions used in the proof of Theorem 1 would no longer be plausible. Nevertheless, by Theorem 2, we can establish the optimality of monotone affine contracts by showing that $\tilde{\mathcal{P}} \cap \mathcal{P}_{Var} \neq \emptyset$. For the profile $\tilde{P}$ constructed above to be contained in $\tilde{\mathcal{P}}$, we need $\frac{8}{65} \leq \tilde{p}_3(a_1) \leq \frac{2}{15}$. For the profile $\tilde{P}$ to be contained in $\mathcal{P}_{Var}$, we need $\tilde{p}_3(a_1) \leq \frac{19}{150}$. Therefore, $\tilde{\mathcal{P}} \cap \mathcal{P}_{Var} \neq \emptyset$. For instance, one such profile is as follows:

$$
\begin{aligned}
(\tilde{p}_1(a_1), \tilde{p}_2(a_1), \tilde{p}_3(a_1)) &= \left(\frac{17}{20}, \frac{1}{40}, \frac{1}{8}\right), \\
(\tilde{p}_1(a_2), \tilde{p}_2(a_2), \tilde{p}_3(a_2)) &= \left(\frac{1}{40}, \frac{13}{80}, \frac{13}{16}\right).
\end{aligned}
$$

Example 2 (Quantile)**.** Besides the mean restriction, the principal also has access to some quantile information. In particular, the probability of the highest output level p_3 satisfies that $p_3(a_1) \geq 0.13$ and $p_3(a_2) \leq 0.91$. Let $\mathcal{P}_Q$ denote the collection of profiles that are consistent with these quantile requirements.

It is easy to check that the binary support distributions used in the proof of Theorem 1 would no longer be plausible. Nevertheless, we show that $\tilde{\mathcal{P}} \cap \mathcal{P}_Q \neq \emptyset$. By Theorem 2, this further implies that a monotone affine contract maximizes the principal's payoff guarantee. For the profile $\tilde{P}$ to be contained in $\tilde{\mathcal{P}}$, we need $\frac{8}{65} \leq \tilde{p}_3(a_1) \leq \frac{2}{15}$. For the profile $\tilde{P}$ to be contained in $\mathcal{P}_Q$, we need $0.13 \leq \tilde{p}_3(a_1) \leq 0.14$. Therefore, $\tilde{\mathcal{P}} \cap \mathcal{P}_Q \neq \emptyset$. For instance, one such profile is as follows:

$$
\begin{aligned}
(\tilde{p}_1(a_1), \tilde{p}_2(a_1), \tilde{p}_3(a_1)) &= (0.86, 0.01, 0.13), \\
(\tilde{p}_1(a_2), \tilde{p}_2(a_2), \tilde{p}_3(a_2)) &= (0.09, 0.065, 0.845).
\end{aligned}
$$

5 Discussion

5.1 Comparison with Diamond (1998)

Diamond (1998, Section 5) considers a model in which the agent can either take one action that produces a low expected output m_l without cost or the other action that produces a higher expected output m_h with a cost $c > 0$. It is assumed that only inducing m_h is worthwhile for the principal. Given a contract w, for each action, the agent can freely distribute the probabilities in any way that preserves the expected output of this action. Therefore, the agent chooses among all possible probability distributions that are consistent with the given mean to maximize his own expected payoff. That is, the agent solves the following maximizing problem:

$$
\begin{aligned}
& \max_{p} \sum_{i=1}^{N} w_i p_i - c \\
\text{subject to} \quad & \sum_{i=1}^{N} p_i = 1, \\
& \sum_{i=1}^{N} q_i p_i = m_h, \\
& p_i \geq 0, \forall i = 1, 2, \ldots, N.
\end{aligned}
$$

In both Diamond (1998, Section 5) and our baseline model, the principal is assumed to know only the mean of the output distribution induced by each action. While Diamond (1998, Section 5) allows the agent to choose among all plausible probability distributions to maximize his own expected payoff, in our model, Nature chooses among all plausible probability distributions to minimize the principal's expected payoff. Example 3 below demonstrates the difference between the agent-maximizing problem in Diamond (1998, Section 5) and the nature-minimizing problem in our model.

Example 3. There are three output levels: $q_1 = 0, q_2 = 1, q_3 = 2$, and there are two actions available to the agent: $c(a_1) = 0, m(a_1) = \frac{1}{2}$ and $c(a_2) = 0.1, m(a_2) = 1$. Consider the following contract: $w_1 = w_2 = 0, w_3 = 1$.

We first consider the agent-maximizing problem. Since the payment to the agent is positive at a single output level q_3, it is easy to see that the agent will choose the action a_2 and the distribution that puts probability $\frac{1}{2}$ on q_1 and probability $\frac{1}{2}$ on q_3. The agent obtains an expected payoff of $\frac{2}{5}$. The principal obtains an expected payoff of $\frac{1}{2}$.

Next, we consider the nature-minimizing problem. We show that the principal's payoff cannot exceed $\frac{1}{4}$. To see this, consider the strategy of nature such that (1) under action a_1, the distribution puts probability $\frac{3}{4}$ on q_1 and probability $\frac{1}{4}$ on q_3, and (2) under action a_2, the distribution puts probability 1 on q_2. It is straightforward to check that in this case, the agent will choose action a_1. The agent's expected payoff is $\frac{1}{4}$, and the principal's expected payoff is $\frac{1}{4}$.

Thus, the principal's expected payoff is $\frac{1}{2}$ in the agent-maximizing problem, and the principal's payoff guarantee cannot exceed $\frac{1}{4}$ in the nature-minimizing problem. This demonstrates that the two problems are different.

This should not come as a surprise. In the nature-minimizing problem, the principal and Nature have extremely opposing interests, whereas in the agent-maximizing problem, their interests could be partially aligned.

5.2 Higher-order moments

The optimal contracts that traditional models of the principal-agent problem predict often take complicated functional forms (see, for example, Grossman and Hart (1983) and Bolton and Dewatripont (2005, Chapter 4)). Nevertheless, our baseline model establishes the optimality of monotone affine contracts when the principal only knows the mean of the output distribution for each action. We could interpret these two settings as two extreme cases of information available to the principal. One may wonder whether monotone affine contracts still generate the highest payoff guarantee when the principal has access to information about some higher-order moments of the output distributions. The following example illustrates that even when the principal has information about only the first two moments, monotone affine contracts might not be optimal.

Example 4. Suppose that there are four levels of output: $q_1 = 0.5, q_2 = 1, q_3 = 2$, and $q_4 = 2.5$. There are two actions $\{a_1, a_2\}$ available to the agent with $c(a_1) = 0.2$ and $c(a_2) = 1.5$. The principal knows the mean and variance of the output distribution for each action. Let $m(a_1) = 0.8, m(a_2) = 2.2$ and $\sigma^2(a_1) = 0.85, \sigma^2(a_2) = 5.05$, where $\sigma^2(a) = \sum_{i=1}^{N} q_i^2 p_i(a)$ is the second-order moment of the output distribution $p(a)$ for the action a. The principal perceives any profile of stochastic output to be plausible as long as it is consistent with these mean and second-order moment restrictions. The principal chooses a contract to maximize her payoff guarantee.

As before, any monotone affine contract generates the same expected payoff for the principal, regardless of the actual profile of stochastic output, as long as it is consistent with m. It is straightforward to calculate that the least expected costs of implementing a_1 and a_2 using monotone affine contracts are $\Psi_M(a_1) = 0.2$ and $\Psi_M(a_2) = 1.58$ respectively. Thus, the highest payoff guarantee from using a monotone affine contract is $m(a_2) - \Psi_M(a_2) = 0.62$.

Now consider the class of quadratic contracts, which take the form $w_i = \alpha q_i^2 +$

$\beta q_i - \gamma$ for constants $\alpha, \beta, \gamma \in \mathbb{R}$ for all i. Since the principal knows $m(\cdot)$ and $\sigma^2(\cdot)$, we can apply the same arguments as in Section 3.2 and show that, for any quadratic contract, the agent always exerts the same action and the principal obtains the same expected payoff for all plausible profiles of stochastic output. Consider a particular quadratic contract

$$w_i = 0.262q_i^2 + 0.143q_i - 0.137, \forall i.$$

Given this contract, the agent chooses the action a_2 and the principal obtained the expected payoff

$$-0.262\sigma^2(a_2) + (1 - 0.143)m(a_2) + 0.137 = 0.6993,$$

which is strictly higher than the payoff guarantee from the optimal monotone affine contract.

This example shows that when there is more information on higher-order moments available to the principal, monotone affine contracts may not attain the highest payoff guarantee. This is not surprising since the knowledge on higher moments greatly reduce the nature's freedom to choose the profile of stochastic output. Besides affine contracts, there are other kinds of contracts (such as quadratic contracts in this example) that can perfectly hedge the principal's uncertainty.

A Proof of Theorem 1

Step (1) *There exists a profile $P \in \mathcal{P}(m)$ against which $V^*(P) = V_A^*$. Subsequently, $V^* = V_A^*$.*

When comparing the two maximizations problems (grand max) and (grand max-Affine), we observe that if the output distribution for any action only puts positive probability on q_1 and q_N, then these two maximization problems are essentially the same. Indeed, consider the following profile of stochastic output $\tilde{P} \in \mathcal{P}(m)$: for each $a \in A$,

$$\tilde{p}_1(a) = \frac{q_N - m(a)}{q_N - q_1}, \quad \tilde{p}_N(a) = \frac{m(a) - q_1}{q_N - q_1}, \quad \tilde{p}_i(a) = 0, \forall i \neq 1, N.$$

Against $\tilde{P}$, the principal's maximization problem (grand max) reduces to:

$$\begin{aligned} &\max_{a\in A,\tilde{\alpha},\tilde{\beta}} (1-\tilde{\alpha})m(a)+\tilde{\beta} \\ \text{subject to} \quad &\tilde{\alpha}m(a)-c(a)\geq \tilde{\alpha}m(a')-c(a') \;\; \forall a'\in A\backslash\{a\}, \\ &\tilde{\alpha}m(a)-\tilde{\beta}-c(a)\geq 0, \\ &\tilde{\alpha}q_1-\tilde{\beta}\geq 0, \\ &\tilde{\alpha}q_N-\tilde{\beta}\geq 0, \end{aligned}$$

where

$$\tilde{\alpha}=\frac{w_N-w_1}{q_N-q_1} \text{ and } \tilde{\beta}=\frac{q_1w_N-q_Nw_1}{q_N-q_1}.$$

Thus, without loss of optimality, the principal chooses from the class of affine contracts. We have $V^*(\tilde{P})=V_A^*$. It follows that $V^*\leq V^*(\tilde{P})=V_A^*$. Since clearly $V_A^*\leq V^*$, we have $V^*=V_A^*$.

We still have to show that $V_A^*=V_M^*$. Recall that A_A (resp. A_M) is the set of actions implementable by affine contracts (resp. monotone affine contracts), and

$$V_A^*=\max_{a\in A_A}\{m(a)-\Psi_A(a)\},\; V_M^*=\max_{a\in A_M}\{m(a)-\Psi_M(a)\}.$$

Step (2). *If $a\in A_A\setminus A_M$, then there exists some $\hat{a}\in A_M$ such that $m(\hat{a})-\Psi_A(\hat{a})>m(a)-\Psi_A(a)$.*

Intuitively, if an action is implementable by affine contracts but not implementable by monotone affine contracts, then implementing this action cannot be optimal for the principal. Since the agent can only be incentivized to take this action by an affine contract with a negative slope, under which he prefers to generate low levels of output, the principal's expected payoff presumably will improve if she implements some other action. In what follows, we prove this intuition rigorously.

Let a_e be the action with the lowest cost among all the actions.[4] In Step (2.1), we show that a_e is implementable by monotone affine contracts, and the principal's highest expected payoff from implementing a_e using affine contracts or monotone affine contracts is $m(a_e)-c(a_e)$. Fix an action a that is implementable by affine contracts but not implementable by monotone affine contracts. Step (2.2) shows that $m(a)<m(a_e)$,

[4]If there are multiple actions with the lowest cost, let a_e be the action with the highest mean among these actions.

and Step (2.3) shows that the highest expected payoff from implementing any such a using an affine contract must be strictly less than the highest expected payoff from implementing a_e using an affine contract.

Step (2.1) *$a_e \in A_M$ and $\Psi_A(a_e) = \Psi_M(a_e) = c(a_e)$.*

Since $c(a_e) \leq c(a)$ for any $a \in A \setminus \{a_e\}$, the monotone affine contract $\alpha = 0$, $\beta = -c(a_e)$, which pays the agent a constant wage $c(a_e)$ regardless of his action, can be used to implement a_e. Clearly, the least expected cost of implementing the action a_e using any feasible contract is at least $c(a_e)$. Thus, $\Psi_A(a_e) = \Psi_M(a_e) = c(a_e)$, and the associated expected payoff for the principal is $m(a_e) - c(a_e)$.

Step (2.2) *If $a \in A_A \setminus A_M$, then $m(a) < m(a_e)$.*

Fix some action $a \in A_A$. We show that if $m(a) > m(a_e)$, then $a \in A_M$. This is because any affine contract (α, β) that implements a must satisfy the following constraint in the implementation problem

$$\alpha m(a) - c(a) \geq \alpha m(a_e) - c(a_e) \implies \alpha \geq \frac{c(a) - c(a_e)}{m(a) - m(a_e)} \geq 0.$$

Thus, a is also implementable by monotone affine contracts.

Step (2.3) $m(a_e) - \Psi_A(a_e) > m(a) - \Psi_A(a)$ for any $a \in A_A \setminus A_M$.

Fix an action $a \in A_A \setminus A_M$. By construction, $c(a_e) \leq c(a)$. From Step (2.2), we know that $m(a_e) > m(a)$. It follows that $m(a_e) - c(a_e) > m(a) - c(a)$. Since the least expected cost of implementing a is at least $c(a)$, the principal's highest expected payoff from implementing a using any feasible contract cannot exceed $m(a) - c(a)$. From Step (2.1), we know that the principal's highest expected payoff from implementing a using affine contracts is $m(a_e) - c(a_e)$.

This completes the proof of Step (2). We have

$$V_A^* = \max_{a \in A_A} \{m(a) - \Psi_A(a)\} = \max_{a \in A_M} \{m(a) - \Psi_A(a)\}$$

Step (3). *If $a \in A_M$, then $\Psi_A(a) = \Psi_M(a)$.*

Fix an action a that is implementable by monotone affine contracts. We show that the least expected cost of implementing such an action does not increase if the principal is restricted to choosing from the class of monotone affine contracts rather than from the class of affine contracts.

First, if there exists some action a' such that $m(a') < m(a)$ and $c(a') \leq c(a)$, then any affine contract (α, β) that implements a must be monotone, since (α, β) necessarily satisfies the following constraint in the implementation problem

$$\alpha m(a) - c(a) \geq \alpha m(a') - c(a') \implies \alpha \geq \frac{c(a) - c(a')}{m(a) - m(a')} \geq 0.$$

Thus, $\Psi_A(a) = \Psi_M(a)$.

Now consider the case in which any action with a lower mean than $m(a)$ has a higher cost than $c(a)$. Since a is implementable by monotone affine contracts, it must also be that $c(a) \leq c(a')$ for any a' with a higher mean than $m(a)$. Otherwise, there exists some action a' with $m(a') > m(a)$ and $c(a') < c(a)$, and any affine contract (α, β) that implements a must satisfy the following constraint in the implementation problem

$$\alpha m(a) - c(a) \geq \alpha m(a') - c(a') \implies \alpha \leq \frac{c(a') - c(a)}{m(a') - m(a)} < 0,$$

which contracts that a is implementable by monotone affine contracts. Since the action a has the smallest cost among all the actions, we know from Step (2.1) that $\Psi_A(a) = \Psi_M(a)$.

Thus,

$$V^* = V_A^* = \max_{a \in A_M} \{m(a) - \Psi_A(a)\} = \max_{a \in A_M} \{m(a) - \Psi_M(a)\} = V_M^*,$$

where the first equality follows from Step (1), the second equality follows from Step (2), and the third equality follows from Step (3). This completes the proof of Theorem 1.

B Proof of Theorem 2

It is straightforward to verify that $\tilde{\mathcal{P}} \subseteq \mathcal{P}(m)$. This step is purely algebraic and we omit the details.

We show that for any $\tilde{P} \in \tilde{\mathcal{P}}$, $\Psi(a|\tilde{P}) = \Psi_M(a)$ for any action a. Clearly, we have $\Psi_M(a_1) = \Psi(a_1|P) = c(a_1)$ for any $P \in \mathcal{P}(m)$. This is because (1) since a_1 has the smallest cost among all the actions, the monotone affine contract with $\alpha = 0$ and $\beta = -c(a_1)$, which pays the agent a constant wage $c(a_1)$ regardless of his action, can be used to implement a_1, and (2) the least expected cost of implementing a_1 using any feasible contract is at least $c(a_1)$ regardless of the profile of stochastic output. Without loss, in what follows, we assume that $a \neq a_1$.

Given the cost minimization problems (min) and (min-Monotone) for implementing an action $a_k \in A, k > 1$, we shall work with their corresponding dual problems. The strong duality holds, since any action $a \in A$ is implementable by monotone affine contracts and the least expected costs from (min) and (min-Monotone) are both finite. Thus, each dual problem induces the same optimal value as the corresponding primal problem.

We first consider the least expected cost of implementing $a_k \in A, k > 1$ using monotone affine contracts. We derive the dual maximization problem (dual-Monotone) as follows:

$$\max_{\lambda_M, \eta_M, \nu_M} \sum_{k' \neq k} \lambda_M(a_k, a_{k'}) \Big(c(a_k) - c(a_{k'}) \Big) + \eta_M c(a_k) \qquad \text{(dual-Monotone)}$$

$$\text{subject to} \quad \sum_{k' \neq k} \lambda_M(a_k, a_{k'}) \Big(m(a_k) - m(a_{k'}) \Big) + \eta_M m(a_k) + \nu_M q_1 \leq m(a_k), \qquad (4)$$

$$\eta_M + \nu_M = 1, \qquad (5)$$

$$\lambda_M(a_k, a_{k'}) \geq 0, \ \forall k' \neq k,$$

$$\eta_M \geq 0,$$

$$\nu_M \geq 0,$$

where $\lambda_M(a_k, a_{k'})$ is the multiplier on the constraint of preferring a_k to $a_{k'}$, η_M is the multiple on the participating constraint, and ν_M is the multiplier on the limited liability constraint in the primal problem (min-Monotone).

We simplify this maximization problem using a series of observations. We claim that (4) must bind at an optimum. Suppose to the contrary, (4) holds with strict inequality at an optimum $(\lambda_M^*, \eta_M^*, \nu_M^*)$. Then we can increase the value of $\lambda_M^*(a_k, a_1)$ by a sufficiently small $\epsilon > 0$, while keeping the other variables unchanged. The new set of variables would still satisfy all the constraints, and the value of the objective function increases. This contradicts the optimality of $(\lambda_M^*, \eta_M^*, \nu_M^*)$.

Furthermore, all the multipliers on the constraints of preferring a_k to $a_{k'}$ for $k' > k$ must be zero at an optimum $(\lambda_M^*, \eta_M^*, \nu_M^*)$. Suppose to the contrary, there exists some $k_1 > k$ such that $\lambda_M^*(a_k, a_{k_1}) > 0$. Note that (4) (which must be binding at an optimum) and (5) imply that

$$\sum_{k'\neq k} \lambda_M^*(a_k, a_{k'})\Big(m(a_k) - m(a_{k'})\Big) = \nu_M^*\Big(m(a_k) - q_1\Big) \geq 0. \tag{6}$$

It follows from (6) that there exists some $k_2 < k$ such that $\lambda_M^*(a_k, a_{k_2}) > 0$. Let

$$\bar{\lambda}_M^*(a_k, a_{k_1}) = \lambda_M^*(a_k, a_{k_1}) - \epsilon \text{ and } \bar{\lambda}_M^*(a_k, a_{k_2}) = \lambda_M^*(a_k, a_{k_2}) + \frac{m(a_k) - m(a_{k_1})}{m(a_k) - m(a_{k_2})}\epsilon$$

for some sufficiently small $\epsilon > 0$, while keeping all other variables unchanged. The new set of variables would still satisfy all the constraints, and the change in the value of the objective function is

$$-\epsilon\left(c(a_k) - c(a_{k_1})\right) + \epsilon\frac{m(a_k) - m(a_{k_1})}{m(a_k) - m(a_{k_2})}\left(c(a_k) - c(a_{k_2})\right).$$

It follows from Assumption 1 that this change is positive, contradicting the optimality of $(\lambda_M^*, \eta_M^*, \nu_M^*)$.

With the above observations, the maximization problem (dual-Monotone) reduces to the following:

$$\begin{aligned} &\max_{\lambda_M, \eta_M, \nu_M} \sum_{k'<k} \lambda_M(a_k, a_{k'})\Big(c(a_k) - c(a_{k'})\Big) + \eta_M c(a_k) \\ \text{subject to} \quad &\sum_{k'<k} \lambda_M(a_k, a_{k'})\Big(m(a_k) - m(a_{k'})\Big) + \eta_M\Big(m(a_k) - q_1\Big) = m(a_k) - q_1, \\ &\lambda_M(a_k, a_{k'}) \geq 0, \ \forall k' < k, \\ &0 \leq \eta_M \leq 1. \end{aligned}$$

This is isomorphic to a utility maximization problem of a consumer who gains marginal utility $c(a_k) - c(a_{k'})$ for consuming good $k' < k$ and marginal utility $c(a_k)$ for consuming good k, faces the price $m(a_k) - m(a_{k'})$ for good $k' < k$ and the price $m(a_k) - q_1$ for good k, and has a budget constraint of $m(a_k) - q_1$. It follows from Assumption 1 that at an optimum, it must be that $\lambda_M^*(a_k, a_{k'}) = 0$ for any $k' \neq k - 1$. We still have two cases to consider.

Case 1. If $\frac{c(a_k) - c(a_{k-1})}{m(a_k) - m(a_{k-1})} > \frac{c(a_k)}{m(a_k) - q_1}$, then

$$\lambda_M^*(a_k, a_{k-1}) = \frac{m(a_k) - q_1}{m(a_k) - m(a_{k-1})}, \ \lambda_M^*(a_k, a_{k'}) = 0, \ \forall k' \neq k - 1, \text{ and } \eta_M^* = 0.$$

Case 2. If $\frac{c(a_k)-c(a_{k-1})}{m(a_k)-m(a_{k-1})} \le \frac{c(a_k)}{m(a_k)-q_1}$, then one optimum is

$$\eta_M^* = 1, \text{ and } \lambda_M^*(a_k, a_{k'}) = 0, \forall k'.$$

We then consider the least expected cost of implementing $a_k \in A, k > 1$ using any feasible contract. For a given profile of stochastic output P, the dual of the cost minimization problem (min) is derived as follows:

$$\begin{aligned} &\max_{\lambda,\eta} \sum_{k' \neq k} \lambda(a_k, a_{k'})\Big(c(a_k) - c(a_{k'})\Big) + \eta c(a_k) && \text{(dual)} \\ \text{subject to } &\sum_{k' \neq k} \lambda(a_k, a_{k'})\Big(p_i(a_k) - p_i(a_{k'})\Big) + \eta p_i(a_k) \le p_i(a_k), \forall i = 1, 2, \dots, N, \\ &\lambda(a_k, a_{k'}) \ge 0, \forall k' \neq k, \\ &\eta \ge 0, \end{aligned}$$

where $\lambda(a_k, a_{k'})$ is the multiplier on the constraint of preferring a_k to $a_{k'}$, η is the multiplier on the participating constraint in the primal problem (min-Monotone).

We now identify conditions under which, for any action $a_k \in A$, $k > 1$, the solution for (dual-Monotone) is feasible for (dual). This guarantees the equivalence of $\Psi(a|\tilde{P})$ and $\Psi_M(a)$ for all $a \in A$.

Fix an action a_k, $k > 1$. If $\frac{c(a_k)-c(a_{k-1})}{m(a_k)-m(a_{k-1})} \le \frac{c(a_k)}{m(a_k)-q_1}$, then the optimal solution

$$\eta_M^* = 1, \text{ and } \lambda_M^*(a_k, a_{k'}) = 0, \forall k'$$

is feasible for the maximization problem (dual) for all P. If $\frac{c(a_k)-c(a_{k-1})}{m(a_k)-m(a_{k-1})} > \frac{c(a_k)}{m(a_k)-q_1}$, then for the optimal solution

$$\lambda_M^*(a_k, a_{k-1}) = \frac{m(a_k) - q_1}{m(a_k) - m(a_{k-1})}, \ \lambda_M^*(a_k, a_{k'}) = 0, \forall k' \neq k-1, \text{ and } \eta_M^* = 0$$

to be feasible for the maximization problem, $\tilde{P}$ must be such that

$$\frac{m(a_k) - q_1}{m(a_k) - m(a_{k-1})}\Big(\tilde{p}_i(a_k) - \tilde{p}_i(a_{k-1})\Big) \le \tilde{p}_i(a_k), \forall i = 1, 2, \dots, N.$$

Suppose that $\tilde{P}$ satisfies

$$\frac{m(a_k) - q_1}{m(a_k) - m(a_{k-1})}\Big(\tilde{p}_i(a_k) - \tilde{p}_i(a_{k-1})\Big) \leq \tilde{p}_i(a_k),\ \forall k > 1,\ \forall i = 1, 2, \ldots, N.$$

Thus, for any $k > 1$,

$$\epsilon_i(a_k) := \frac{m(a_k) - q_1}{m(a_{k-1}) - q_1}\tilde{p}_i(a_{k-1}) - \tilde{p}_i(a_k) \geq 0,\ \forall i = 1, 2, \ldots, N.$$

The requirement that $\tilde{P} \in \mathcal{P}(m)$ imposes the following restrictions on $\epsilon_i(a_k)$: $\forall k > 1$,

$$\sum_{i=1}^{N} \epsilon_i(a_k) = \frac{m(a_k) - m(a_{k-1})}{m(a_{k-1}) - q_1},\ \text{and}$$

$$\sum_{i=1}^{N} q_i\epsilon_i(a_k) = q_1\frac{m(a_k) - m(a_{k-1})}{m(a_{k-1}) - q_1}.$$

Thus, $\sum_{i=1}^{N} q_i\epsilon_i(a_k) = \sum_{i=1}^{N} q_1\epsilon_i(a_k)$, $\forall k > 1$. This can only be true if

$$\epsilon_1(a_k) = \frac{m(a_k) - m(a_{k-1})}{m(a_{k-1}) - q_1},\ \text{and}\ \epsilon_i(a_k) = 0,\ \forall k > 1,\ \forall i = 2, 3, \ldots, N.$$

Through some algebraic manipulation, we can represent $\tilde{p}(a)$ for each action $a \in A$ in terms of only $\tilde{p}(a_1)$: for any $k > 1$,

$$\tilde{p}_1(a_k) = \frac{m(a_k) - q_1}{m(a_1) - q_1}\,\tilde{p}_1(a_1) - \frac{m(a_k) - m(a_1)}{m(a_1) - q_1}, \tag{7}$$

$$\tilde{p}_i(a_k) = \frac{m(a_k) - q_1}{m(a_1) - q_1}\,\tilde{p}_i(a_1),\ \forall i = 2, 3, \ldots, N. \tag{8}$$

Therefore, $\tilde{p}(a_1)$ must satisfy the following conditions (besides being a probability distribution):

$$\sum_{i=1}^{N} q_i\tilde{p}_i(a_1) = m(a_1), \tag{9}$$

$$\tilde{p}_1(a_1) \geq \frac{m(a_K) - m(a_1)}{m(a_K) - q_1}, \tag{10}$$

where (10) is derived from the requirement that $\tilde{p}(a_k)$ is a probability distribution for any $k > 1$.

For any probability distribution $\tilde{p}(a_1)$ such that (9) and (10) are satisfied, using

(7) and (8), we could construct a profile of stochastic output $\tilde{P} \in \mathcal{P}(m)$ such that $\Psi(a|\tilde{P}) = \Psi_M(a)$ for all $a \in A$. Subsequently, $V^*(\tilde{P}) = V_M^*$. This completes the proof of Theorem 2.

C Proof of Theorem 3

As argued in the main text, to prove Theorem 3, it suffices to show that there exists a probability distribution $\tilde{p}_i(a_1)$ such that

$$\begin{aligned}
&\sum_{i=1}^{N} q_i \tilde{p}_i(a_1) = m(a_1), \\
&\tilde{p}_1(a_1) > \frac{m(a_K) - m(a_1)}{m(a_K) - q_1}, \\
&\tilde{p}_i(a_1) > 0, \ \forall i = 2, 3, \ldots, N.
\end{aligned}$$

Since $\sum_{i=1}^{N} \tilde{p}_i(a_1) = 1$ and $\sum_{i=1}^{N} q_i \tilde{p}_i(a_1) = m(a_1)$, we have

$$\begin{aligned}
\tilde{p}_1(a_1) &= \frac{1}{q_2 - q_1}\Big(\sum_{i=3}^{N}(q_i - q_2)\tilde{p}_i(a_1) + q_2 - m(a_1)\Big), \\
\tilde{p}_2(a_1) &= \frac{1}{q_2 - q_1}\Big(m(a_1) - q_1 - \sum_{i=3}^{N}(q_i - q_1)\tilde{p}_i(a_1)\Big).
\end{aligned}$$

For $\tilde{p}_1(a_1)$ to be larger than $\frac{m(a_K)-m(a_1)}{m(a_K)-q_1}$, we need

$$\sum_{i=3}^{N}(q_i - q_2)\frac{\tilde{p}_i(a_1)}{m(a_1) - q_1} > \frac{m(a_K) - q_2}{m(a_K) - q_1}.$$

For $\tilde{p}_2(a_1)$ to be larger than 0, we need

$$\sum_{i=3}^{N}(q_i - q_1)\frac{\tilde{p}_i(a_1)}{m(a_1) - q_1} < 1.$$

Lastly, we need

$$\tilde{p}_i(a_1) > 0, \ \forall i = 3, \ldots, N.$$

These three conditions can be satisfied by the following output distribution $p(a_1)$:

$$
\begin{aligned}
p_1(a_1) &= \frac{1}{q_2 - q_1}\left(\sum_{i=3}^{N}(q_i - q_2)p_i(a_1) + q_2 - m(a_1)\right),\\
p_2(a_1) &= \frac{1}{q_2 - q_1}\left(m(a_1) - q_1 - \sum_{i=3}^{N}(q_i - q_1)p_i(a_1)\right),\\
p_i(a_1) &= \frac{(m(a_1) - q_1)(q_2 - q_1)(q_N - m(a_K))}{4(m(a_K) - q_1)(q_N - q_2)\sum_{i=3}^{N-1}(q_i - q_1)}, \ \forall i = 3, 4, \ldots, N-1,\\
p_N(a_1) &= \frac{(m(a_1) - q_1)(m(a_K) - q_2)}{2(m(a_K) - q_1)(q_N - q_2)} + \frac{m(a_1) - q_1}{2(q_N - q_1)}.
\end{aligned}
$$

This completes the proof of Theorem 3.

References

Bolton, P., and M. Dewatripont (2005): *Contract Theory.* The MIT Press.

Brooks, B., and S. Du (2021a): "Maxmin Auction Design with Known Expected Values," working paper.

——— (2021b): "Optimal Auction Design with Common Values: An Informationally-Robust Approach," *Econometrica*, 89(3), 1313–1360.

——— (2023): "On the Structure of Informationally Robust Optimal Mechanisms," working paper.

Carrasco, V., V. F. Luz, N. Kos, M. Messner, P. Monteiro, and H. Moreira (2018): "Optimal Selling Mechanisms under Moment Conditions," *Journal of Economic Theory*, 177, 245–279.

Carroll, G. (2015): "Robustness and Linear Contracts," *American Economic Review*, 105(2), 536–563.

Carroll, G., and D. Meng (2016a): "Locally Robust Contracts for Moral Hazard," *Journal of Mathematical Economics*, 62, 36–51.

——— (2016b): "Robust Contracting with Additive Noise," *Journal of Economic Theory*, 166, 586–604.

Che, E. W. (2022): "Distributionally Robust Optimal Auction Design under Mean Constraints," working paper.

Chen, Y.-C., and J. Li (2018): "Revisiting the Foundations of Dominant-Strategy Mechanisms," *Journal of Economic Theory*, 178, 294–317.

Chung, K.-S., and J. C. Ely (2007): "Foundations of Dominant-Strategy Mechanisms," *Review of Economic Studies*, 74(2), 447–476.

Diamond, P. (1998): "Managerial Incentives: On the Near Linearity of Optimal Compensation," *Journal of Political Economy*, 106(5), 931–957.

Du, S. (2018): "Robust Mechanisms under Common Valuation," *Econometrica*, 86(5), 1569–1588.

Grossman, S. J., and O. D. Hart (1983): "An Analysis of the Principal-Agent Problem," *Econometrica*, 51(1), 7–46.

He, W., and J. Li (2022): "Correlation-Robust Auction Design," *Journal of Economic Theory*, 200, 105403.

Hurwicz, L., and L. Shapiro (1978): "Incentive Structures Maximizing Residual Gain under Incomplete Information," *Bell Journal of Economics*, 9(1), 108–191.

Kambhampati, A. (2023): "Randomization is Optimal in the Robust Principal-Agent Problem," *Journal of Economic Theory*, 207, 105585.

Kambhampati, A., J. Toikka, and R. Vohra (2023): "Randomization and the Robustness of Linear Contracts," working paper.

Koçyiğit, Ç., G. Iyengar, D. Kuhn, and W. Wiesemann (2020): "Distributionally Robust Mechanism Design," *Management Science*, 66(1), 159–189.

Li, Z., and S. N. Kirshner (2021): "Salesforce Compensation and Two-sided Ambiguity: Robust Moral Hazard with Moment Information," *Production and Operations Management*, 30(9), 2944–2961.

Neeman, Z. (2003): "The Effectiveness of English Auctions," *Games and Economic Behavior*, 43(2), 214–238.

Rosenthal, M. (2023): "Simple Incentives and Diverse Beliefs," working paper.

Suzdaltsev, A. (2022): "Distributionally Robust Pricing in Independent Private Value Auctions," *Journal of Economic Theory*, 206, 105555.

Yamashita, T., and S. Zhu (2022): "On the Foundations of Ex Post Incentive-Compatible Mechanisms," *American Economic Journal: Microeconomics*, 14(4), 494–514.

Zhang, W. (2022): "Correlation-Robust Optimal Auctions," working paper.