Clause for Concern: Contract Design with a Second Opinion

TAMAR GARBUZ, Tel Aviv University
ARIEL D. PROCACCIA, Harvard University
EDEN SAIG, The Technion
INBAL TALGAM-COHEN, Tel Aviv University
JAMIE TUCKER-FOLTZ, Harvard University

In many economic settings, an initial evaluation may be uncertain or imperfect, leading decision-makers to pay for a secondary evaluation that refines judgments and mitigates risk. Classic contract design theory does not capture this possibility, creating an opportunity for economic mechanisms that leverage this structure. In this work, we develop a principal-agent contract model in which the principal can pay to acquire additional information about the agent's hidden action by inspecting its outcome. The principal's objective is to incentivize a rational agent to choose a reward-maximizing action, yielding a joint optimization problem over both the inspection policy (i.e., when to request additional information), and the corresponding monetary transfers. We show that this problem is NP-hard in general, but becomes tractable under natural structural assumptions, or when core problem dimensions are small. We further extend our model to incorporate probabilistic inspection and show that contract design problems with probabilistic outcome inspection do not admit a Stackelberg equilibrium, in contrast to classical contract design problems. Finally, we examine the economic implications of our results, and propose methods to restore equilibrium guarantees in probabilistic inspection settings.

Contents

1 Introduction

Contract theory examines the strategic interaction between a principal and an agent, focusing on how to incentivize individuals to exert costly effort. A principal aims to motivate an agent to take a hidden action that incurs a cost, and leads to a probability distribution over possible outcomes. Since the principal derives utility from these outcomes, they offer a payment scheme that links realized outcomes to financial compensation. The goal is to design a contract that maximizes the principal's expected utility while ensuring that the agent is incentivized to take the desired action; the payments can depend only on observed outcomes rather than on the hidden action.

The classical contract theory model can be seen as assuming that the principal receives a single signal regarding the action taken by the agent, in the form of the observed stochastic outcome. However, in many real-world situations it is possible to acquire secondary evaluations to refine initial assessments and decrease the information gap. Everyday scenarios—from medical diagnoses to home improvement estimates—routinely incorporate second opinions to ensure more accurate and reliable assessments.

We are especially interested in situations where obtaining a second opinion is costly. As an example, consider a setting where the principal is paying a generative AI service for code generation, as discussed by Saig et al. [25]. Different levels of effort by the agent correspond to using increasingly sophisticated (and computationally intensive) AI models, say GPT-3.5, GPT-4, and GPT-4o. The principal cannot directly observe which model is being used. [1] A coarse (and essentially free) signal of the quality of the code generation can be obtained by compiling it. A second opinion can then be obtained by hiring a proficient human to inspect the code; this provides another, more informative signal but incurs a cost.

1.1 Our Contributions

To capture such settings, we extend the standard contract design framework by introducing the option to pay for a second opinion, adding a new layer of complexity to the design of optimal contracts. In our model, each action initially generates a stochastic signal, which can then be further inspected by paying a third party for access to additional stochastic outcome. To maximize utility, the principal commits to both an inspection policy (i.e., specifying which signals to inspect) and a payment scheme (i.e., specifying how much to pay), balancing the benefits of reduced information gaps against the expense of acquiring extra information. Notably, our model exhibits unique economic properties absent in prior contract theory; for instance, there exist scenarios where the first-best is strictly positive, yet the principal's optimal utility is exactly zero. This contrasts with previous work, where even linear contracts are known to approximate the first-best.

Our first set of results characterize the computational complexity of finding optimal contracts with second opinion. In particular, we demonstrate that the contract optimization problem admits a polynomial-time solution in two key practical settings: (i) when the number of actions available to the agent is constant (see Theorem 3.5), and (ii) when both the Independent Second Opinion Property (ISOP) and the Inspection Monotone Likelihood Ratio Property (Inspection-MLRP) regularity assumptions hold (see Theorem 3.7). For example, in the context of AI-generated code, the number of models (or actions) available to the agent is typically small. This scenario aligns with the first setting, allowing for the design of an optimal contract that engages a human expert to review the code in a tractable manner. The second setting captures cases where the second opinion is obtained independently of the initial signal and where higher inspection costs lead to better outcomes. Interestingly, we also show that this combination of assumptions is *tight*: the contract design

[1]While one can trust OpenAI to use a specific model, this may not be the case in the future, as generative AI services proliferate. For example, consider the skepticism with which the release of DeepSeek-R1 was met [7].

problem becomes NP-hard when each assumption holds individually (see Theorems 3.8 and 3.9). Our positive results are achieved by combining sparsity arguments with a nested reduction to the classical contract design problem, leveraging existing results on the structure of optimal contracts as a critical sub-component of our analysis (see Section 3.1).

We extend our model to include second-opinion contracts with random inspection. In the code generation example, this corresponds to randomly selecting a subset of code modules for in-depth review, balancing inspection costs against the risk of undetected vulnerabilities. Intriguingly, we prove that this natural probabilistic extension induces a game that does not admit a Stackelberg equilibrium, in contrast to classical contract design settings (Theorem 4.1). The key insight is that under probabilistic inspection, the principal can continually reduce expected inspection costs by promising increasingly high payments with vanishing probabilities, thereby undermining the existence of an equilibrium. Motivated by this negative property, we propose methods to restore equilibrium guarantees by imposing additional constraints on the principal. In particular, we consider two practically-motivated variants: (1) disallowing commitment to inspection probabilities, and (2) requiring that the payment for inspected outcomes never exceeds the payment for corresponding uninspected signals. We prove that each of these restrictions—individually and in combination—restores equilibrium guarantees (Theorem 4.4).

We conclude by discussing open questions and directions for further inquiry.

1.2 Related Work

Our work contributes to the growing frontier in algorithmic game theory focused on optimizing the efforts of others [1–5, 8, 10, 12–14, 16–18, 23]. See [15] for a recent survey.

Contracts and inspections. Our results complement a recent line of work on *action* (rather than *outcome*) inspection. Fallah and Jordan [20] study the problem of a single principal interacting with multiple agents, aiming to maximize utility by promoting safety-compliant actions. The principal uses a combination of payments and a limited budget for action inspections to incentivize compliance. Ezra et al. [19] relax the hidden-action assumption and introduce a combinatorial model in which the principal is allowed to inspect sets of actions at a cost. The principal proposes an inspection scheme described by a suggested action, a payment, and an inspection distribution over subsets of actions, and permitted to withhold payment if the agent is caught not performing the agreed-upon action. The principal's goal is to find the IC inspection scheme (where the suggested action is the best response for the agent), such that the principal's utility is maximized. Our model is inspired by theirs, but is largely complementary. One special case common to both models occurs when we consider an instance with binary signals (representing their binary outcomes), and costly inspections that fully disclose the agent's action. But in general, we introduce a complementary setting to action inspection by allowing the principal to seek a second opinion regarding the *outcome* of the agent's action. This perspective opens new avenues for optimizing contract structures, balancing the cost of additional information with the benefits of enhanced decision-making accuracy. Additional works that consider the effect of information on contract design include papers by Babichenko et al. [6], Castiglioni and Chen [9] and Garrett et al. [21].

Contracts, generative AI agents, and trust. Immorlica et al. [24] discuss generative AI as economic agents. Saig et al. [26] study contracts for incentivizing AI agents to "exert effort" (carry out their task using a strong model; see also [27]). While strategic misreporting of the model in LLMs has not yet been officially acknowledged, it has been shown that LLMs are volatile, with fluctuating performance on the same tasks at different points of time [11]. In related industries, there are certainly examples of trust violations from companies who provide black box services — an extreme

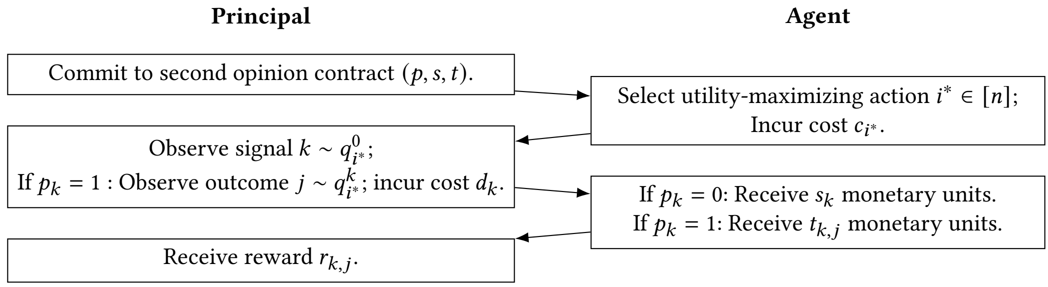

Fig. 1. Schematic diagram of the interaction model described in Section 2.

example is Theranos, and another example is throttling by internet service providers. In contracts, allowing outcome inspections could significantly improve this trust.

2 Model

Setting. The model involves a principal interacting with an agent. The agent has a set of actions $[n]$, and the action *costs* for the agent are $0 \le c_1 \le \cdots \le c_n$. The agent's action stochastically generates a *signal* $k \in [\ell]$ according to a probability matrix $\mathbf{q}^0 = \{q^0_{i,k}\}_{i\in[n],k\in[\ell]}$, where $q^0_{i,k}$ denotes the probability of observing signal k given action i. Upon receiving a signal $k \in [\ell]$, the principal decides whether to *inspect* it. Inspection incurs a cost $d_k \ge 0$ for the principal, and reveals an *outcome* $j \in [m_k]$ (possible outcomes associated with signal k), which is drawn according to a probability matrix $\mathbf{q}^k = \{q^k_{i,j}\}_{i\in[n],j\in[m_k]}$. Here, $q^k_{i,j}$ represents the probability of obtaining outcome $j \in [m_k]$ given action i and signal k. We note that an outcome is always drawn regardless of whether the principal inspects the signal. The inspection decision solely determines whether the principal observes the outcome. Hence, payments to the agent may depend on the observed signal and, if inspected, the observed outcome. Thus, an instance of this model can be described by the tuple $(\mathbf{q}, c, d)$, where $\mathbf{q} = (\mathbf{q}^0, \{\mathbf{q}^k\}_{k\in[\ell]})$ represents the signal and outcome probability matrices, $c = (c_1, \ldots, c_n)$ denotes the agent's action costs, and $d = (d_1, \ldots, d_\ell)$ specifies the inspection costs for each signal.

Contracts with a Second Opinion. A contract with a second opinion specifies the principal's commitment to an inspection policy and a payment scheme. The inspection policy is defined by a vector p, where p_k represents the probability that the principal inspects signal k upon observing it. We begin by focusing on deterministic inspections, where $p_k \in \{0, 1\}$. Later, we expand the model to incorporate probabilistic inspections, where $p_k \in [0, 1]$. Payment commitments are described by monetary transfer vectors (s, t). The principal pays $s_k \ge 0$ when signal $k \in [\ell]$ is observed but not inspected, and $t_{k,j} \ge 0$ when signal k is inspected and outcome $j \in [m_k]$ is observed. Thus, the contract is fully characterized by the tuple (p, s, t).

Agent's Best Response. For a contract (p, s, t), we denote the expected monetary transfer from the principal to the agent when taking action i by $T_i = \mathbb{E}_{k\sim q^0_i}\left[(1 - p_k)s_k + p_k\, \mathbb{E}_{j\sim q^k_i}[t_{k,j}]\right]$. The agent's expected utility for taking action i is given by $U_A(i \mid p, s, t) = T_i - c_i$. We assume the agent is rational and selects an action that maximizes their expected utility: $i^*(p, s, t) \in \arg\max_{i\in[n]} U_A(i \mid p, s, t)$. When it is clear from context, we omit the contract and write i^* instead of $i^*(p, s, t)$. Following standard conventions in the literature, we assume that ties are broken in favor of the principal.

Principal's goal. Each signal-outcome pair (k, j) is associated with a reward $r_{k,j} \geq 0$ for the principal. Accordingly, the expected reward for action i is given by $R_i = \mathbb{E}_{k \sim q_i^0} \mathbb{E}_{j \sim q_i^k} [r_{k,j}]$. We denote the expected inspection cost given inspection policy p and action i by $D_i = \mathbb{E}_{k \sim q_i^0} [p_k d_k]$. Consequently, for a given contract (p, s, t), the principal's utility is $U_P(p, s, t) = R_{i^*(p,s,t)} - T_{i^*(p,s,t)} - D_{i^*(p,s,t)}$. The principal aims to design a contract that maximizes their utility: $(p^*, s^*, t^*) \in \arg\max_{p,s,t} U_P(p, s, t)$.

The interaction model described in this section is illustrated in Figure 1.

2.1 Regularity Assumptions

While our general model does not impose specific assumptions on the costs or on the signal and outcome distributions, real-world contract design problems often exhibit inherent *structure*. In this section, we introduce structural assumptions motivated by practical applications. Our first assumption is based on the notion of *independence*. A contract design satisfies the Independent Second Opinion Property (ISOP) if inspecting any signal yields an independent random draw from the same distribution:

Definition 2.1 (Independent Second Opinion Property (ISOP)). A contract design setting (q, c, d) has the *Independent Second Opinion Property (ISOP)* if $m_k = \ell$ for all $k \in [\ell]$, and it holds that $q_{i,k}^{k'} = q_{i,k}^0$ for all $i \in [n], k, k' \in [\ell]$.

Although the second opinion is drawn independently from the same distribution, ISOP still permits this second opinion to carry varying weight, as the reward for each signal-outcome pair can be arbitrary. However, in certain settings, the sequence in which these opinions are received should not affect the principal's utility. We formalize this restriction as follows:

Definition 2.2 (Symmetric-ISOP). A contract design setting (q, c, d) satisfies *Symmetric-ISOP* if it satisfies ISOP and, additionally, the rewards are symmetric, i.e., $r_{k,j} = r_{j,k}$ for all $j, k \in [\ell]$.

Whether this symmetry condition is a reasonable assumption is a matter of taste and application domain. It turns out not to affect complexity as far as this paper is concerned: Our positive result about ISOP is for the general case, while our hardness result holds even for Symmetric-ISOP.

Our final assumption is that the *ordering* of signals and outcomes carries information about the hidden action. Intuitively, a classical contract design setting satisfies the Monotone Likelihood Ratio Property (MLRP) when higher outcomes are more likely to be produced by a high-cost action than a low-cost one. We extend the classical definition to settings with a second opinion by asserting that MLRP holds for the signal distribution $q_{i,k}^0$ and the outcome distribution $q_{i,j}^k$ for each signal:

Definition 2.3 (Inspection-MLRP). A contract design setting (q, c, d) satisfies *Inspection-MLRP* if for every pair of actions $i, i' \in [n]$ such that $c_i < c_{i'}$, the likelihood ratio $q_{i',k}^0 / q_{i,k}^0$ is increasing in k, and the likelihood ratio $q_{i',j}^k / q_{i,j}^k$ is increasing in j for all $k \in [\ell]$.

As we will show, these assumptions affect tractability in a *tight* sense: In Section 3.3, we demonstrate that the problem is tractable when both assumptions hold; In Section 3.4, we show that the problem becomes computationally hard if even one of these assumptions is not satisfied.

2.2 First Best

Finally, we note that the second opinion model exhibits distinct economic properties compared to classical contract design. In a delegation setting, the *first-best* represents the maximum utility achievable in the non-strategic version of the problem where incentive constraints are ignored and players cooperate to maximize the common good (i.e., the total "size of the pie" that can be divided among them). This is given by the highest expected reward minus the corresponding cost of any action, formally expressed as $\max_{i \in [n]} \{R_i - c_i\}$.

In the classical contract design literature, Dütting et al. [16] have shown that even linear contracts can guarantee a fraction of the first-best (specifically at least the first best divided by the number of actions). In contrast, under the second opinion model we identify instances where the first-best is strictly positive, yet the principal's optimal utility is exactly zero.

For example, consider the following instance, with two actions, one signal, and two outcomes, defined as follows:

$$q^1_{1,j} = \begin{pmatrix} 1 & 0 \end{pmatrix}; \quad q^1_{2,j} = \begin{pmatrix} 0 & 1 \end{pmatrix}; \quad c_i = \begin{pmatrix} 0 & 1 \end{pmatrix}; \quad d_1 = 1 \quad r_{1,j} = \begin{pmatrix} 0 & 2 \end{pmatrix}.$$

This implies that the first-best is $\max_i\{R_i - c_i\} = 1$ (attained for action $i = 2$). To incentivize the first action, the optimal contract is not to inspect the signal and pay 0. Under this contract, the agent's best response is to choose the first action, and the principal's utility is 0. To incentivize the second action, the optimal contract is to inspect the signal and pay $t_{1,2} = 1$ and 0 otherwise. Under this contract, the agent's best response, assuming tie-breaking in favor of the principal, is to choose the second action, and the principal's utility is also 0.

The key observation is that the signal is uninformative. As a result, an action with a positive cost cannot be incentivized without inspection. Furthermore, the cost of inspection exactly offsets the reward minus the payment. Thus, the principal's optimal utility is 0, even though cooperation can lead to a positive value.

3 Computing Optimal Contracts

We first consider two settings in which finding the optimal principal utility turns computationally tractable, namely, when the number of actions is constant, or when ISOP and Inspection-MLRP both hold. We then show that the problem becomes hard with a non-constant number of actions when either ISOP or Inspection-MLRP are relaxed, completing the picture of the complexity landscape.

We seek the action i that maximizes R_i minus the expected inspection cost and payments required to incentivize it. This amounts to solving the following Quadratically-Constrained Quadratic Program (QCQP) for each action $i \in [n]$:

$$\begin{array}{lll}
\textbf{minimize} & \sum_{k\in[\ell]} q^0_{i,k} \left((1-p_k)s_k + p_k\left(d_k + \sum_{j\in[m_k]} q^k_{i,j} t_{k,j}\right)\right) & \\
& & \\
& s_k \geq 0 & \forall k \in [\ell] \\
& t_{k,j} \geq 0 & \forall k \in [\ell],\, j \in [m_k] \\
\textbf{subject to} & 0 \leq p_k \leq 1 & \forall k \in [\ell] \\
& \sum_{k\in[\ell]} q^0_{i,k} \left((1-p_k)s_k + p_k \sum_{j\in[m_k]} q^k_{i,j} t_{k,j}\right) - c_i \geq & \\
& \sum_{k\in[\ell]} q^0_{i',k} \left((1-p_k)s_k + p_k \sum_{j\in[m_k]} q^k_{i',j} t_{k,j}\right) - c_{i'} & \forall i' \in [n]
\end{array} \qquad (1)$$

We denote the last constraint above as the *IC constraint.*

Note that in the deterministic inspection setting, which is the primary focus of our paper, we additionally require each p_k to take values in $\{0, 1\}$. However, it will be useful to consider the more general problem without this restriction. Even when allowing p_k to take any value in the interval $[0, 1]$, the program remains generally non-convex.

If we fix the values of p and then optimize the remaining s and t variables, the problem reduces to a linear program, as elaborated in the next subsection. This will allow us to treat Equation (1) as a nested optimization problem, and provide tractability guarantees.

3.1 Combined Outcome Space and Distribution

A key ingredient in our tractability proofs is a correspondence between second opinion contracts and classical contracts for any fixed inspection policy p. Given a second opinion contract setting,

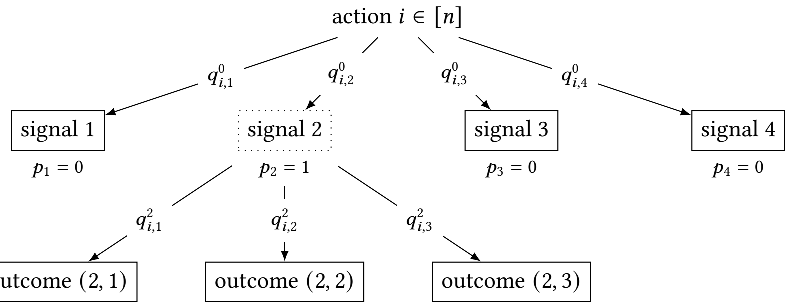

Fig. 2. Schematic diagram of the combined outcome space Ω, as defined in Section 3.1. When the inspection policy p is fixed, sampling from the combined outcome distribution $f_{i,\omega}$ (Equation (2)) is equivalent to a random walk on the inspection tree according to the signal and outcome probabilities ($q^0_{i,j}$ and $q^k_{i,j}$, respectively), starting from the root and proceeding until a leaf is reached.

we define the combined outcome space and corresponding distributions, which provide a more unified perspective on the signals and outcomes. For $k \in \{0, \dots, \ell\}$, we denote the set of signals by Ω_0 (such that $|\Omega_0| = \ell$), and the sets of outcomes by Ω_k for $k \geq 1$ (such that $|\Omega_k| = m_k$). We assume that elements in all sets are uniquely labeled, and therefore the sets do not intersect. The *combined outcome space* Ω is the union:

$$\Omega = \Omega_0 \cup \cdots \cup \Omega_\ell.$$

Given inspection probabilities $p \in [0, 1]^\ell$, for any action-outcome pair (i, ω) the combined outcome distribution is:

$$f_{i,\omega}(p) = \begin{cases} q^0_{i,\omega}(1 - p_\omega) & \omega \in \Omega_0 \\ q^0_{i,k} p_k q^k_{i,\omega} & \omega \in \Omega_k \text{ for } k > 0 \end{cases} \tag{2}$$

For brevity, we denote $f_{i,\omega} = f_{i,\omega}(p)$ when p is clear from context. We note that for any action i, the vector $f_{i,\omega}$ is a probability distribution since plugging in the definition and rearranging we get:

$$\sum_{\omega\in\Omega} f_{i,\omega} = \sum_{\omega\in\Omega_0} f_{i,\omega} + \sum_{k\in[\ell]} \sum_{\omega\in\Omega_k} f_{i,\omega} = \sum_{k\in[\ell]} (1 - p_k) q^0_{i,k} + \sum_{k\in[\ell]} p_k q^0_{i,k} \underbrace{\sum_{\omega\in\Omega_k} q^k_{i,\omega}}_{=1} = \sum_{k\in[\ell]} q^0_{i,k} = 1$$

Similarly, the combined payments vector is:

$$v_\omega = \begin{cases} s_\omega & \omega \in \Omega_0 \\ t_{k,\omega} & \omega \in \Omega_k \text{ for } k > 0 \end{cases} \tag{3}$$

Figure 2 provides a schematic representation of the combined outcome space. Using these notations, we reformulate the contract design problem:

LEMMA 3.1. *Equation* (1) *is equivalent to the following optimization problem:*

$$\begin{aligned} \textbf{\textit{minimize}} \quad & \textstyle\sum_{\omega\in\Omega} f_{i,\omega} v_\omega + D_i \\ \textbf{\textit{subject to}} \quad & v \geq 0 \\ & p \in [0, 1]^\ell \\ & \textstyle\sum_{\omega\in\Omega} f_{i,\omega} v_\omega - c_i \geq \sum_{\omega\in\Omega} f_{i',\omega} v_\omega - c_{i'} \quad \forall i' \neq i \end{aligned} \tag{4}$$

Proof. By definition of $f_{i,\omega}$ (Equation (2)), D_i (Section 2), and v_ω (Equation (3)).
The objective of Equation (1) is:

$$\sum_{k\in[\ell]} q^0_{i,k}\left((1-p_k)s_k + p_k\left(d_k + \sum_{j\in[m_k]} q^k_{i,j}t_{k,j}\right)\right)$$

Rearranging the sums, we obtain:

$$\underbrace{\sum_{k\in[\ell]} q^0_{i,k}(1-p_k)s_k}_{=\sum_{\omega\in\Omega_0} f_{i,\omega}v_\omega} + \underbrace{\sum_{k\in[\ell]}\sum_{j\in[m_k]} q^0_{i,k}p_k q^k_{i,j}t_{k,j}}_{=\sum_{k\in[\ell]}\sum_{\omega\in\Omega_k} f_{i,\omega}v_\omega} + \underbrace{\sum_{k\in[\ell]} q^0_{i,k}p_k d_k}_{=D_i}$$

Which is equal to the objective $\sum_{\omega\in\Omega} f_{i,\omega}v_\omega + D_i$.
Similarly, the (IC) constraint in Equation (1) is:

$$\sum_{k\in[\ell]} q^0_{i,k}\left((1-p_k)s_k + p_k\sum_{j\in[m_k]} q^k_{i,j}t_{k,j}\right) - c_i \geq \sum_{k\in[\ell]} q^0_{i',k}\left((1-p_k)s_k + p_k\sum_{j\in[m_k]} q^k_{i',j}t_{k,j}\right) - c_{i'}$$

Using the same arguments, we rearrange and obtain equivalence to the inequality as required:

$$\sum_{\omega\in\Omega} f_{i,\omega}v_\omega - c_i \geq \sum_{\omega\in\Omega} f_{i',\omega}v_\omega - c_{i'}.$$

□

We make the following observation about Equation (4):

Lemma 3.2. *For any fixed p, the optimal v in Equation (4) is a min-pay optimal contract.*

Proof. In the objective of Equation (4), $\sum_\omega f_{i,\omega}v_\omega$ is the expected pay of the contract v, and D_i is an additive term which does not depend on v. When p in fixed, f_i and D_i are also fixed. In that case, the constant term D_i does not affect the optimization, and the optimization problem is therefore equivalent to the classic Stackelberg min-pay problem. □

3.2 Tractability With a Constant Number of Actions

We show that the contract problem is tractable when the number of actions is constant. Intuitively, our proof relies on a sparsity argument: When the number of actions is bounded by a constant, there exists an optimal contract in which the number of inspected signals is also bounded by that constant, significantly reducing the search space. Furthermore, we demonstrate that if a contract inspects a signal but assigns zero payment for all its outcomes, eliminating this inspection while appropriately adjusting the remaining payments preserves incentive compatibility and weakly improves the principal's utility. Combining these insights leads to an efficient enumeration-based algorithm that computes an optimal contract in polynomial time.

Proposition 3.3. *Let (p, s, t) be a contract. If there exists an inspected signal $k_0 \in [\ell]$ such that $t_{k_0,j} = 0$ for all $j \in [m_{k_0}]$, then there exists a contract (p', s', t') with $p'_{k_0} = 0$ and $p'_k = p_k$ for all $k \neq k_0$, which yields weakly higher utility for the principal.*

Proof. Define a modified contract (p', s', t') such that:

$$p'_k = \begin{cases} 0 & k = k_0 \\ p_k & \text{otherwise} \end{cases}, \quad s'_k = \begin{cases} (1-p_{k_0})s_{k_0} & k = k_0 \\ s_k & \text{otherwise} \end{cases}, \quad t'_{k,j} = t_{k,j}$$

For any action $i \in [n]$, the expected monetary transfer from the principal to the agent satisfies:

$$\begin{aligned}
T_i(p,s,t) &= \sum_{k\in[\ell]} q^0_{i,k}\left((1-p_k)s_k + p_k \sum_{j\in[m_k]} q^k_{i,j} t_{k,j}\right) \\
&= \sum_{k\in[\ell]\setminus\{k_0\}} q^0_{i,k}\left((1-p_k)s_k + p_k \sum_{j\in[m_k]} q^k_{i,j} t_{k,j}\right) + q^0_{i,k_0}\left((1-p_{k_0})s_{k_0} + p_{k_0} \sum_{j\in[m_{k_0}]} q^{k_0}_{i,j} \underbrace{t_{k_0,j}}_{=0}\right) \\
&= \sum_{k\in[\ell]\setminus\{k_0\}} q^0_{i,k}\left((1-p_k)s_k + p_k \sum_{j\in[m_k]} q^k_{i,j} t_{k,j}\right) + q^0_{i,k_0} \underbrace{(1-p_{k_0})s_{k_0}}_{=s'_{k_0}} \\
&= \sum_{k\in[\ell]} q^0_{i,k}\left((1-p'_k)s'_k + p'_k \sum_{j\in[m_k]} q^k_{i,j} t'_{k,j}\right) = T_i(p',s',t')
\end{aligned}$$

The two contracts transfer identical amounts in expectation for any action, and therefore the modified contract (p', s', t') implements the same action. The modified contract (p', s', t') also yields weakly higher utility for the principal, as it does not inspect signal k_0. □

Lemma 3.4. *Consider a contract design setting (q, c, d) with ℓ signals and n actions. There is an optimal contract which implements action i^* while inspecting at most $n-1$ signals (with at most $n-1$ nonzero payments).*

Proof. By contradiction. Denote by p^* the inspection policy of an optimal contract, and denote the corresponding inspection set by $S = \{k \mid p^*_k > 0\}$. Assume that p^* has a minimal inspection set size $|S|$ among all optimal contracts, and assume by contradiction that $|S| \geq n$.

By Lemma 3.2, for any fixed inspection policy $p \in [0,1]^\ell$, and in particular for the given p^*, the optimal signal and outcome payments s^*, t^* are equivalent to a min-pay contract over the combined outcome space $\Omega = \Omega_0 \cup \cdots \cup \Omega_\ell$. Denote by v^*_ω the representation of s^*, t^* in the combined outcome space, as given by Equation (3). By [16], a min-pay contract design problem over n actions has an optimal solution with at most $n-1$ nonzero payments, and therefore we assume that v^*_ω has at most $n-1$ nonzero entries. By definition of the combined space Ω, there exist at most $n-1$ signals $k \in [\ell]$ for which there exists $\omega \in \Omega_k$ such that $v^*_\omega > 0$. Equivalently, in the s^*, t^* representation, there exist at most $n-1$ signals $k \in [\ell]$ for which $t^*_{k,j} > 0$ for some $j \in [m_k]$.

Since $|S| \geq n$, there exists at least one inspected signal k' for which $t^*_{k',j} = 0$ for all $j \in [m_{k'}]$. By Claim 3.3, there exists a weakly-better contract (p', s', t') with a smaller inspection set, contradicting the minimality of $|S|$. Therefore, there exist an optimal contract inspecting at most $n-1$ signals. □

Theorem 3.5. *Consider the family of contract design problems where the number of actions is constant. There exists an algorithm which computes the optimal deterministic-inspection contract in $O(\mathrm{poly}(\ell, m))$ time for any contract design problem (q, c, d) with ℓ signals and $m_k \leq m$ outcomes for each signal.*

Proof. Denote the set of actions by $[n]$, and consider the following algorithm: For each subset of signals $S \subseteq [\ell]$ such that $|S| \leq n-1$, compute the optimal payments s, t to implement action i^* under the constraint $p_k = \mathbb{1}_{k\in S}$. Output the contract (p^*, s^*, t^*) yielding the minimal expected pay.

By Lemma 3.2, for any fixed inspection policy $p \in [0,1]^\ell$, the optimization problem for the signal and outcome payments s, t is equivalent to a min-pay contract design problem, and therefore the optimal contract in each iteration can be computed in polynomial time by solving a linear program. There are $O(\ell^n)$ subsets of signals of size smaller than n, and since $n = O(1)$, it holds that $O(\ell^n) = O(\text{poly}(\ell))$. Multiplying the time complexity of each iteration by the total number of iterations, we obtain that the algorithm described above runs in $O(\text{poly}(\ell, m))$ time in total.

By Claim 3.4, there exists an optimal contract inspecting at most $n-1$ signals. Since the algorithm enumerates all subsets of signals up to this size, it is guaranteed to encounter the optimal subset of signals, and thus return an optimal result. □

3.3 Tractability When Both ISOP and Inspection-MLRP Hold

We demonstrate that tractability also emerges when the ISOP and Inspection-MLRP regularity assumptions hold jointly. Intuitively, we prove that these assumptions, when holding simultaneously, limit the number of binding constraints in the optimization problem, reducing computational complexity. In Section 3.4, we prove that the problem becomes NP-Hard when any of these two assumptions are relaxed.

Proposition 3.6. *Let (q, c, d) be a contract design setting satisfying ISOP and Inspection-MLRP, and targeting the highest-cost action n. If the action is implementable, then any optimal contract pays for at most one signal, and for at most one outcome.*

Proof. Denote an optimal contract by (p^*, s^*, t^*). Since the design setting satisfies ISOP, we denote for brevity $q_{i,k} = q^0_{i,k} = q^{k'}_{i,k}$. We consider two cases:

If the contract doesn't inspect any signal (i.e., $p^*_k = 0$ for all $k \in [\ell]$), then by Lemma 3.2, the optimal pay for signals s^* is a min-pay contract over the signal distribution $q_{i,k}$ and costs c_i. The contract design problem $(q_{i,k}, c_i)$ satisfies MLRP, and therefore by [16] there exists an optimal contract which only pays for the highest signal $k = \ell$.

Otherwise, the contract inspects at least one signal. By Lemma 3.2, the transfers s^*, t^* are an optimal solution to a min-pay contract design problem over the combined outcome space:

$$\begin{aligned}
\textbf{minimize} \quad & \textstyle\sum_{\omega\in\Omega} f_{i,\omega} v_\omega \\
\textbf{subject to} \quad & v \geq 0 \\
& \textstyle\sum_{\omega\in\Omega} f_{n,\omega} v_\omega - c_n \geq \sum_{\omega\in\Omega} f_{i,\omega} v_\omega - c_i \quad \forall i < n
\end{aligned}$$

Where $f_{i,\omega}$ is given by Equation (2), and v_ω is given by Equation (3). By [16], the dual of the min-pay LP is:

$$\begin{aligned}
\textbf{maximize} \quad & \textstyle\sum_{i<n} \lambda_i (c_n - c_i) \\
\textbf{subject to} \quad & \lambda \geq 0 \\
& \textstyle\sum_{i<n} \lambda_i \left(f_{n,\omega} - f_{i,\omega}\right) \leq f_{n,\omega} \quad \forall \omega \in \Omega
\end{aligned}$$

Plugging in the definitions of $f_{i,\omega}$ and Ω we get:

$$\begin{aligned}
\textbf{maximize} \quad & \textstyle\sum_{i<n} \lambda_i (c_n - c_i) \\
\textbf{subject to} \quad & \lambda \geq 0 \\
& \textstyle\sum_{i<n} \lambda_i \left(q_{n,k} - q_{i,k}\right) \leq q_{n,k} && \forall k \text{ s.t. } p^*_k = 0 \\
& \textstyle\sum_{i<n} \lambda_i \left(q_{n,k} q_{n,j} - q_{i,k} q_{i,j}\right) \leq q_{n,k} q_{n,j} && \forall k \text{ s.t. } p^*_k = 1,\ j \in [\ell]
\end{aligned}$$

Constraints for which $q_{n,k} = 0$ or $q_{n,k}q_{n,j} = 0$ represents signals and outcomes that cannot be reached. They are satisfied for any λ, and are therefore redundant.

Ignoring redundant constraints, we divide the first set of constraints by $q_{n,k}$ and the second set of constraints by $q_{n,k}q_{n,j}$ to obtain:

$$\begin{array}{ll} \textbf{maximize} & \sum_{i<n} \lambda_i (c_n - c_i) \\ & \\ & \lambda \geq 0 \\ \textbf{subject to} & \sum_{i<n} \lambda_i \left(1 - \frac{q_{i,k}}{q_{n,k}}\right) \leq 1 \qquad \forall k \text{ s.t. } p_k^* = 0 \\ & \sum_{i<n} \lambda_i \left(1 - \frac{q_{i,k}}{q_{n,k}} \cdot \frac{q_{i,j}}{q_{n,j}}\right) \leq 1 \quad \forall k \text{ s.t. } p_k^* = 1,\ j \in [\ell] \end{array}$$

From the Inspection-MLRP assumption, the ratio $\frac{q_{i,k}}{q_{n,k}}$ is decreasing in k. We denote $k_0^* = \max\{k \mid p_k^* = 0\}$. For any $i < n$, $k < k_0^*$, and $\lambda_i \geq 0$, it holds that:

$$\lambda_i \left(1 - \frac{q_{i,k}}{q_{n,k}}\right) \leq \lambda_i \left(1 - \frac{q_{i,k_0^*}}{q_{n,k_0^*}}\right)$$

and therefore the first set of constraints satisfies:

$$\sum_{i<n} \lambda_i \left(1 - \frac{q_{i,k_0^*}}{q_{n,k_0^*}}\right) \leq 1 \quad \Rightarrow \quad \sum_{i<n} \lambda_i \left(1 - \frac{q_{i,k}}{q_{n,k}}\right) \leq 1 \tag{5}$$

Similarly, we denote $k_1^* = \max\{k \mid p_k^* = 1\}$. For any $i < n$, $k < k_1^*$,$j \in [\ell]$, and $\lambda_i \geq 0$, it holds that:

$$\lambda_i \left(1 - \frac{q_{i,k}}{q_{n,k}} \cdot \frac{q_{i,j}}{q_{n,j}}\right) \leq \lambda_i \left(1 - \frac{q_{i,k_1^*}}{q_{n,k_1^*}} \cdot \frac{q_{i,\ell}}{q_{n,\ell}}\right)$$

and therefore the second set of constraints satisifes:

$$\sum_{i<n} \lambda_i \left(1 - \frac{q_{i,k_1^*}}{q_{n,k_1^*}} \cdot \frac{q_{i,\ell}}{q_{n,\ell}}\right) \leq 1 \quad \Rightarrow \quad \sum_{i<n} \lambda_i \left(1 - \frac{q_{i,k}}{q_{n,k}} \cdot \frac{q_{i,j}}{q_{n,j}}\right) \leq 1 \tag{6}$$

From equations (5) and (6), the dual LP has at most two binding constraints, and therefore from complementary slackness, the optimal solution for the primal LP pays for at most one signal, and at most one outcome. □

Remark 3.1. *Unlike the classical result of [16] where the optimal contract pays only for the highest outcome, in our case, the optimal contract does not necessarily follow this structure (restrict payment to the highest outcome of the highest signal). Instead, it pays for at most the highest not inspected signal and at most the highest outcome of the highest inspected signal. This distinction separates our model from the classical setting.*

Theorem 3.7. *An optimal deterministic inspection policy for a contract design setting satisfying ISOP and Inspection-MLRP can be found in polynomial time.*

Proof. By Claim 3.6, any optimal contract (p^*, s^*, t^*) pays for one signal at most, and one outcome at most. Therefore, by Claim 3.3 it can be assumed without loss of generality that in the optimal contract there exists at most one $k \in [\ell]$ such that $p_k = 1$. All single-signal inspection policies can be enumerated in linear time, and an optimal contract for each policy can be computed in polynomial time, yielding a polynomial-time algorithm in total. □

3.4 Hardness in the General Case

We next establish that the contract design problem is computationally hard in the general case. Specifically, we show that computing the optimal principal utility is NP-hard when the number of actions is not constant and at most one of the structural properties Inspection-MLRP and ISOP holds. (In fact, we even show hardness for Symmetric-ISOP under the further restriction that all signals cost the same amount to inspect.) This contrasts with the tractable cases where either the number of actions is constant or both structural properties hold simultaneously, completely characterizing the precise boundary between tractable and intractable instances within our framework.

Theorem 3.8. *It is NP-hard to compute the optimal principal utility under a deterministic contract, even in settings satisfying Inspection-MLRP.*

Proof. We reduce from Vertex Cover. Given an input graph $G = (V, E)$, we define an instance with $n := |E| + 1$ actions, $\ell := |V|$ signals, and two outcomes for each signal, which we will refer to as the *good outcome* and *bad outcome* for that signal. Each good outcome will have a reward of ℓ, and each bad outcome will have a reward of zero. The cost to the principal to inspect each signal k is $d_k = 1$. We arbitrarily number the vertices $v_1, v_2, \ldots, v_\ell$ and edges $e_1, e_2, \ldots, e_{n-1}$ so that they naturally correspond to the signals indexed from 1 to ℓ and actions indexed from 1 to $n-1$, which we call *bad actions*. Each of these bad actions costs the agent nothing. Action n is the *good action*, and costs the agent $c_n = \frac{1}{3\ell^2}$.

If the agent takes the good action, a uniformly random signal is drawn, and the outcome is the good outcome for that signal. If the agent takes bad action i, again a uniformly random signal k is drawn, and the bad outcome is realized if edge e_i is incident to vertex v_k; otherwise the good outcome is realized. Accordingly, the expected reward from the good action is $R_n = \ell$ and from any bad action i is $R_i = \ell \cdot \left(1 - \frac{2}{\ell}\right) = \ell - 2$.

We note that this instance satisfies the definition of Inspection-MLRP. The cost of the good action is the only one that is positive. When the agent takes the good action, the outcome is the good outcome for a uniformly random signal. It is impossible to get a bad outcome under the good action and therefore $q^k_{n,good} = 1$ for all $k \in [\ell]$. Consequently, for all $i < n$, it holds that $\frac{q^k_{n,bad}}{q^k_{i,bad}} \leq \frac{q^k_{n,good}}{q^k_{i,good}}$ for all $k \in [\ell]$ as required. In addition regardless of the action, a uniformly random signal is drawn, which means that $\frac{q^0_{n,k}}{q^0_{i,k}} = 1$ for all $i < n$, and thus, the required monotonicity in k is also satisfied.

We claim that, for any positive integer y, it is possible for the principal to achieve an expected utility of at least $\ell - \frac{y+1/2}{\ell}$ if and only if G has a vertex cover of size at most y. For the backward direction, given a vertex cover C of size y, we consider the contract that inspects each signal k such that $v_k \in C$ and pays $\frac{1}{2\ell}$ unless a signal is inspected and the bad outcome is realized, in which case the payment is zero. If the agent takes a bad action i, they will get an expected payoff of at most $\frac{1}{2\ell} \cdot (1 - \frac{1}{\ell})$, since there is at least a $\frac{1}{\ell}$ chance that the signal corresponds to some vertex in C covering e_i, in which case the agent will get no payment. On the other hand, if they take the good action, they will get a deterministic payoff of $\frac{1}{2\ell}$. Thus, the marginal payoff for taking the good action rather than any bad action is at least $\frac{1}{2\ell^2}$, which is more than the cost of taking the good action, $c_n = \frac{1}{3\ell^2}$. So the agent will always take the good action. The principal's reward is $R_n = \ell$, and their expected cost is $\frac{y}{\ell}$ for inspection plus at most $\frac{1}{2\ell}$ for paying the agent. The principal's expected utility is therefore at least $\ell - \frac{y+1/2}{\ell}$.

For the forward direction, suppose the principal attains this level of utility with some deterministic contract. We will show that

$$C := \{v_k \mid \text{the principal inspects signal } k\}$$

is a vertex cover of G of size at most y.

To see that $|C| \leq y$, observe that inspection costs the principal an expected $\frac{y}{\ell}$ for each vertex in C. Hence, supposing toward a contradiction that the principal inspects more than y (i.e., at least $y+1$) of the k signals, they incur an expected cost of $\frac{y+1}{\ell}$, for a maximum possible payoff of $\ell - \frac{y+1}{\ell}$. This contradicts our assumption that they receive a payoff of at least $\ell - \frac{y+1/2}{\ell}$, so $|C| \leq y$.

Finally, we show that C is a vertex cover. Suppose toward a contradiction that some edge e_i is uncovered. Since the signals for both of its incident vertices are never inspected, the principal will never be able to tell whether the agent took action i or action n; these are the only two signals that can distinguish actions i and n. Hence, the agent will never take action n, as it costs more than action i and leads to the same expected payment. Thus, the principal's reward is at most $\ell - 2 < \ell - \frac{y+1/2}{\ell}$, which again is a contradiction. □

The key observation is that unless the principal inspects a sufficiently large subset of signals, corresponding to a vertex cover, the agent can always choose a bad action that goes undetected. Consequently, minimizing the number of inspected signals aligns directly with solving the minimum vertex cover problem, proving that computing the optimal contract is NP-hard.

Theorem 3.9. *It is NP-hard to compute the optimal principal utility under a deterministic contract, even in settings satisfying Symmetric-ISOP, even when all inspection costs are the same.*

Proof sketch. This proof is considerably more involved, so is deferred to Appendix A. At a high-level, we reduce from a variant of Set Cover through an intermediate problem that we call *Row-Sparsest Matrix (RSM)*. The objective of RSM is to fill in an $n \times n$ matrix with nonnegative entries in a way that satisfies a series of linear constraints while having as few rows with nonzero entries as possible. These matrix entries ultimately become the payments in the optimal contract. The types of linear constraints in RSM take a very restricted form. Most notably, they are symmetric with respect to transposing the matrix, which is why we are able to relate this problem to Symmetric-ISOP. □

4 Random Inspection

The optimality of nondeterministic mechanisms is a ubiquitous phenomenon in economic theory. From selling multiple goods, to fairly dividing resources, to aggregating complex preferences, randomization is frequently able to overcome impossibilities and improve utility.

In terms of our contract design model, a natural place to use randomness is in deciding whether to obtain a second opinion. If we allow the principal to commit to a given inspection probability for each signal k, we derive a Quadratically-Constrained Quadratic Program (QCQP) from Equation (1), with the constraints $p_k \in [0, 1]$ for all $k \in [\ell]$. However, we show that permitting probabilistic inspection may preclude the existence of a Stackelberg equilibrium. Specifically, for any design setting with positive inspection costs and a contract with a nonzero inspection probability, we show that the principal can scale down the probability and accordingly scale up certain payments in a way that preserves the IC constraint. This adjustment results is a new contract incentivizing the same action as before, but at a strictly lower expected inspection cost. Thus, it may be the case that no equilibrium exists under probabilistic inspection! Moreover, the contracts which achieve close-to-optimal principal utilities are impractical, as they involve payments with vanishing probabilities.

Obviously, such contracts are not encountered in practice. In response to this challenge, we study two practical assumptions that can be imposed to eliminate this potential deviation by the principal (see Figure 3). We begin by introducing these assumptions formally and discussing the general computational approach for solving them. In Theorems 4.1 and 4.3, we show that having at least one of the two assumptions is both necessary and sufficient for equilibria to exist outside of

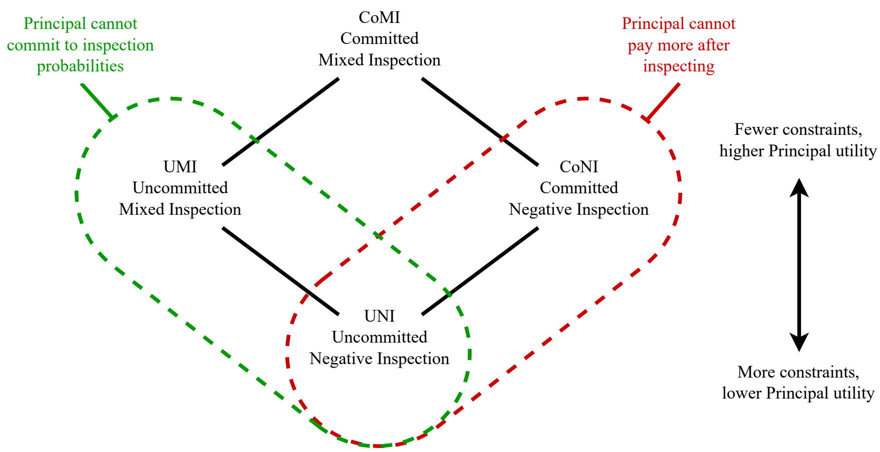

Fig. 3. The four nondeterministic contract design problem variants we consider, partially ordered by optimal principal utility.

trivial cases. We then give some initial results on the complexity of the contract design problem and give an upper bound on the cost savings afforded by randomized inspection. Surprisingly, we show in Theorem 4.6 that there are settings in which randomized inspection can yield a higher principal utility even when the principal is unable to commit to inspection probabilities, and thus must be indifferent about whether to pay for a second opinion.

4.1 Committed Mixed Inspection

In the *Committed Mixed Inspection* (CoMI) setting, the principal commits in advance to a random inspection strategy, by assigning each signal a fixed probability of inspection, and then deciding whether to inspect based on a biased coin toss. Formally, optimal CoMI contracts are optimal solutions to Equation (1), under the constraint $p_k \in [0, 1]$ for all $k \in [\ell]$, with p_k interpreted as the probability of inspecting signal k upon observing it. For instance, consider an organization that employs AI to generate code but remains concerned about potential bugs. To mitigate risk without incurring prohibitive inspection costs, the organization might commit to further inspect at random a fixed percentage (e.g., 10%) of the code modules that pass initial evaluation. Despite the intuitive appeal of this randomized, pre-committed approach as a natural extension of the deterministic second-opinion model, we show that the CoMI setting does not admit a Stackelberg equilibrium:

Theorem 4.1. *Given any CoMI instance X, let $\widehat{X}$ be the same as X, except with zero inspection costs. Then $\widehat{X}$ always has an optimal solution with some value (expected principal utility) y, and X has an optimal solution if and only if one of the optimal solutions to $\widehat{X}$ incentivizes an action i while never inspecting any signal k such that $q^0_{i,k} d_k > 0$; otherwise, there is a sequence of solutions to X with values converging to y.*

Proof in Appendix B. As an immediate corollary, we can resolve the complexity of CoMI.

Corollary 4.2. *There is a polynomial-time algorithm to solve CoMI (in the sense of computing the supremum of possible expected principal utilities).*

Proof. Given any CoMI instance X, consider the instance $\widehat{X}$ from Theorem 4.1. Since inspection costs zero, it is without loss of generality to set each $p_k = 1$. Then Equation (1) becomes a linear program, so we can solve $\widehat{X}$ in polynomial time to obtain the supremum utility. □

Intuitively, the core argument in the proof of Theorem 4.1 is that for any contract which inspects signals with some probability, the principal can lower their expected payment by decreasing inspection probabilities and proportionally increasing payment for inspected outcomes – Implying that no optimal contract exists under these conditions, and providing a series of contracts that converge towards the optimal value y. Moreover, we note that the resulting contracts approaching the optimal value y become impractical: Their worst-case payment tends to infinity while the probability of receiving any payment tends to zero, exceeding the budget limits of any practical principal, and deterring agents with the slightest degree of risk-aversion. In light of these issues, the next section demonstrates that equilibrium guarantees can be recovered by imposing additional restrictions on the principal.

4.2 Restoring Equilibrium Guarantees

In this section, we show that equilibrium guarantees can be restored by imposing additional restrictions on the principal. Specifically, we consider two variants: (1) disallowing commitment to inspection probabilities, and (2) requiring that the payment for inspected outcomes always remains at most that of corresponding uninspected signals. In Theorem 4.3, we prove that each of these restrictions—individually or in combination—restores equilibrium guarantees.

4.2.1 Uncommitted Mixed Inspection (UMI). One possible way to restore equilibrium guarantees is to consider settings which generalize beyond the traditional Stackelberg framework. Specifically, we introduce the Uncommitted Mixed Inspection (UMI) variant, in which the principal only commits to payments (s, t), but does not commit in advance to an inspection policy p. Unlike the Stackelberg approach taken by classical contract design theory and by the settings we discussed thus far, the objective in UMI is to design a contract (p, s, t) which induces a *subgame-perfect Bayes-Nash equilibrium* of the principal-agent interaction that is optimal for the principal. By inducing such an equilibirium, the principal can declare an intention to inspect a particular signal k with an intermediate probability $0 < p_k < 1$, which the agent will accept as believable. Subgame perfection necessitates that this is credible, meaning the principal must be genuinely indifferent between inspecting and not inspecting the signal.

In view of the QCQP in Equation (1) which can be used to solve CoMI, we may handle UMI by imposing the following additional subgame-perfection constraints:

- For each $k \in [\ell]$ such that $p_k < 1$,

$$s_k \leq d_k + \sum_{j \in [m_k]} q_{i,j}^k t_{k,j}. \tag{7}$$

- For each $k \in [\ell]$ such that $p_k > 0$,

$$s_k \geq d_k + \sum_{j \in [m_k]} q_{i,j}^k t_{k,j}. \tag{8}$$

Observe that it is without loss of generality to enforce the constraint given by Equation (7), even when $p_k = 1$. This is because the variable s_k is irrelevant to both the objective function and the other constraints, so we may satisfy this constraint by setting $s_k := d_k + \sum_{j \in [m_k]} q_{i,j}^k t_{k,j}$. We would like to similarly expand the constraint given by Equation (8). However, when $p_k = 0$ it is *not* without loss of generality to assume constraint (8) holds, because if we were to apply the same trick it could require setting some $t_{k,j}$ to be negative. To circumvent this issue, suppose we fix a set of signals

S_0 with the additional rule that $p_k = 0$ for all $k \in S_0$, compatible with a computational approach of enumerating all subsets of signals. Then we may enforce constraint (8) for all k, as long as we remove the stipulation that $t_{k,j} \geq 0$ for $k \in S_0$. This lets us combine constraints (7) and (8) into a single equation:

$$s_k = d_k + \sum_{j \in [m_k]} q^k_{i,j} t_{k,j}.$$

This equality must hold for all $k \in [\ell]$, which allows us to simplify both the objective function and the constraints, and remove all occurrences of the s_k variables. Unfortunately, it does not let us remove all of the quadratic terms in the IC constraint. The result is thus a Quadratically-Constrained Linear Program (QCLP), parameterized by an action i and a set of signals S_0:

$$\begin{array}{lll}
\textbf{minimize} & \sum_{k \in [\ell]} q^0_{i,k} d_k + \sum_{k \in [\ell]} q^0_{i,k} \sum_{j \in [m_k]} q^k_{i,j} t_{k,j} & \\
& & \\
& \sum_{j \in [m_k]} q^k_{i,j} t_{k,j} \geq -d_k & \text{for all } k \in S_0 \\
& t_{k,j} \geq 0 & \text{for all } k \in \overline{S_0},\ j \in [m_k] \\
& p_k = 0 & \text{for all } k \in S_0 \\
\textbf{subject to} & 0 \leq p_k \leq 1 & \text{for all } k \in \overline{S_0} \\
& \sum_{k \in [\ell]} q^0_{i,k} \left((1 - p_k) d_k + \sum_{j \in [m_k]} q^k_{i,j} t_{k,j} \right) - c_i \geq & \\
& \sum_{k \in [\ell]} q^0_{i',k} \Big((1 - p_k) \left(d_k + \sum_{j \in [m_k]} q^k_{i,j} t_{k,j} \right) & \text{for all } i' \in [n] \\
& \qquad + p_k \sum_{j \in [m_k]} q^k_{i',j} t_{k,j} \Big) - c_{i'} &
\end{array}$$

By enumerating all sets of signals S_0, we may effectively optimize small UMI instances (as we do for the proof of Theorem 4.6).

4.2.2 Committed Negative Inspection (CoNI). Another possible way to guarantee the existence of an equilibrium is to impose that the payment for an inspected outcome never exceeds the payment for the corresponding uninspected signal. This condition aligns well with scenarios where inspections may uncover negative information, justifying reduced pay. For example, in the code generation use-case discussed in the introduction, further inspection may uncover security vulnerabilities that reduce the code's value. Formally, this is represented by constraints requiring that $t_{k,j} \leq s_k$ for all $k \in [\ell]$ and $j \in [m_k]$, ensuring that inspection can only reduce or maintain the monetary transfer.

4.2.3 Uncommitted Negative Inspection (UNI). Finally, we consider a problem variant which integrates both the no-commitment and negative inspection assumptions. Here, the principal does not commit to an inspection policy, and is also subject to the constraint $t_{k,j} \leq s_k$ for all k and j.

4.2.4 Equilibrium Guarantees. Not every QCQP or QCLP has an optimal value. Hence, it is not clear that any of the three variant games actually have an equilibrium, as there may not be one optimal contract for the principal to choose. The following result, proved in Appendix B, shows that optimal contracts exist in all variants:

THEOREM 4.3. *Any UMI, CoNI, or UNI instance has an optimal solution.*

4.3 Deterministic vs. Optimal Inspection

We now turn to the question of how powerful randomized inspection policies can be, in terms of what actions they can incentivize and at what cost. The following theorem characterizes implementability and bounds the cost savings of nondeterminism for CoMI, CoNI, and UMI. The relationships between the principal's minimal payment in each variant are depicted in Figure 4.

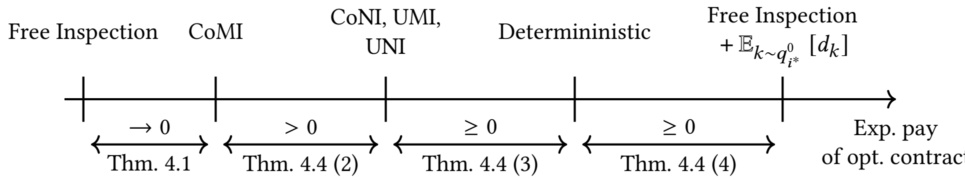

Fig. 4. Relationships between expected payments in optimal deterministic versus nondeterministic contracts. Note that the strict inequality holds only when second opinions are ever used in the optimal contract.

Theorem 4.4. *For any contract design setting and any action i:*

1. *In each of CoMI, CoNI, UMI, and UNI, it is possible to incentivize an action i with a nondeterministic contract if and only if it is possible to incentivize i with a deterministic contract.*
2. *In CoNI, UMI, or UNI, if an optimal contract to incentivize i ever pays for inspection, the principal can incentivize i with a strictly smaller expected payment in CoMI; otherwise, the minimum expected payments are the same.*
3. *In CoNI, UMI, or UNI, the principal's minimal payment to incentivize i is weakly smaller than their minimum expected payment required under a deterministic contract.*
4. *The infimum expected payment required to incentivize i in CoMI is within $\mathbb{E}_{k\sim q^0_{i^*}}[d_k]$ of the minimum expected payment required to incentivize i in a deterministic contract.*

Proof. We will show that it is possible to transform each of a deterministic, CoNI, UMI, and CoMI contract into one another while incentivizing the same action i, which will prove (1). We will also bound the gaps between the minimum expected payments of each transformation, proving the other three statements.

First suppose there is some deterministic contract incentivizing action i. We will show that there is an equivalent contract in UNI (and thus in CoNI and UMI as well). Specifically, for each signal k, we will transform the s_k and $t_{k,j}$ variables in a way that satisfies the commitment/negativity constraints and does not change the principal's expected payment. For signals k such that $p_k = 0$, the $t_{k,j}$ variables are payoff-irrelevant. Thus, we may set $t_{k,j} := s_k$ to satisfy both the negativity constraint and the relevant commitment constraint, namely Equation (7). Likewise, for all signals k such that $p_k = 1$, the s_k variables are payoff-irrelevant, so we may set $s_k = d_k + \max_j t_{k,j}$ to both the negativity constraint and the relevant commitment constraint, namely Equation (8). Thus, we can transform any deterministic contract into an UNI contract with the same expected payment, so the optimal CoNI/UMI/UNI contracts have weakly lower expected payments. This proves (3).

Next take an arbitrary contract in CoNI or UMI incentivizing action i. Since such a contract is also a valid solution to CoMI, the principal's minimum expected payment in CoMI is weakly smaller. If, additionally, the contract sometimes pays for a second opinion, there must be some signal k such that $q^0_{i,k} > 0$, $p_k > 0$, and $d_k > 0$. By Lemma B.1 (1) (stated and proved in Appendix B), we can decrease the inspection probability of signal k and adjust payments accordingly to get a strictly smaller payment in CoMI. This proves (2).

Finally, to complete the cycle, take an arbitrary CoMI contract (p, s, t) incentivizing action i. Consider the alternative contract (p', s', t') where each $p'_k = 1$, s'_k is defined arbitrarily, and

$$t'_{k,j} = (1 - p_k)s_k + p_k t_{k,j}.$$

In other words, (p', s', t') simulates (p, s, t) by always inspecting every signal and sometimes simply ignoring the outcome and using the old s_k payments. The agent's expected utilities for each action are clearly the same in (p', s', t') as in (p, s, t), so (p', s', t') still incentivizes action i. The additional expected payment is

$$\sum_{k\in[\ell]} q^0_{i,k}(1-p_k)d_k \le \sum_{k\in[\ell]} q^0_{i,k}d_k = \mathbb{E}_{k\sim q^0_i}[d_k],$$

which proves (4). □

This result has complexity implications as well. One might hope that non-deterministic contracts are easier to find since the domain of each p_k variable is the convex set $[0,1]$ rather than the discrete set $\{0,1\}$. However, combining Theorem 4.4 (3) and an observation about the proof of Theorem 3.8, we can conclude that at least two of our nondeterministic variants are still NP-hard:

Theorem 4.5. *It is NP-hard to find the optimal principal utility in either UMI or UNI.*

Proof. We observe that the reduction from Vertex Cover in Theorem 3.8 also works in both UMI and UNI:

- For the forward direction, Theorem 4.4 (3) implies that the principal's utility can only be greater in UMI or UNI.
- For the backward direction, we define C to be the set of vertices v_k such that the principal inspects signal k with nonzero probability. Then the same argument as before shows that C is a vertex cover. To prove that $|C| < y$, we observe that, since the principal is unable to commit to probabilities, the principal must weakly prefer paying the inspection cost for the signal of each vertex in C. Thus, we still have that it costs an expected $\frac{y}{\ell}$ for each vertex in C, and the proof follows. □

Finally, we show that the middle inequality from Figure 4 is *sometimes* strict. In other words, randomized inspection can improve principal's optimal utility in some cases, even with either the commitment or negativity constraints imposed.

Theorem 4.6. *There exist instances of UMI and CoNI where nondeterministic inspection is strictly optimal.*

Proof. All numerical claims in this proof have been verified computationally, using Gurobi [22] to solve the various non-convex programs discussed previously. Consider an instance with three actions, two signals, and two outcomes per signal, with

$$q^0_{i,k} = \begin{pmatrix} 0.5 & 0.5 \\ 0.6 & 0.4 \\ 0.6 & 0.4 \end{pmatrix}; \quad q^1_{i,j} = q^2_{i,j} = \begin{pmatrix} 0.6 & 0.4 \\ 0.4 & 0.6 \\ 0.6 & 0.4 \end{pmatrix}; \quad c_i = \begin{pmatrix} 0 & 0 & 1 \end{pmatrix}; \quad d_k = \begin{pmatrix} 1 & 1 \end{pmatrix}.$$

Suppose the rewards are such that the principal wishes to incentivize the costly action 3. Since action 3 is positively correlated with signal 1 and outcome 1, the optimal deterministic contract inspects signal 1 and only pays for outcome 1. The necessary payment is $16 + \frac{2}{3}$, and the total expected cost (including inspection) is 6.6. However, with a nondeterministic contract, it is preferable to save on inspection cost by inspecting randomly. In CoNI, the optimal inspection probability is 0.625, for an expected cost of 6.375; in UMI, the optimal probability is 0.525, for an expected cost of 6.315. □

5 Discussion

By extending the classical principal-agent framework to include outcome inspections, our model reveals new layers of complexity and rich economic behavior absent from traditional models. Under specific assumptions such as a constant number of agent actions or when ISOP and Inspection-MLRP hold simultaneously we have shown that the contract design problem becomes tractable, and that computational hardness emerges when these assumptions are violated. Additionally, we have shown that classical equilibrium guarantees do not hold under probabilistic inspection, motivating the introduction of new methods to restore equilibrium in these settings. These insights open avenues for future work on designing contracts that effectively balance the cost of acquiring additional information with the benefits of improved decision-making.

In the realm of nondeterministic inspection, our work has opened several intriguing questions about the structure of optimal contracts. We have shown that randomized inspection policies can incentivize desirable actions at strictly lower cost, even under either of the two practical constraints we consider. However, we do not have a tight upper-bound on just how large the cost savings can be. Additionally, computational experiments suggest that randomization is not helpful at all when there are only two actions. We are not able to offer any explanation for this phenomenon.

Complexity questions regarding nondeterministic inspection remain as well. We are able to resolve the complexity (in the general case) for three of our four problem variants, but it is unknown whether the final CoNI variant can be solved in polynomial time. It is also open whether any of the regularity assumptions we consider in this paper for deterministic contracts can likewise make these problems tractable.

Another open question concerns the complexity of contracts involving more than two opinions. While our model extends classical contract design to accommodate a second opinion, the computational landscape for settings with three or more opinions remains unexplored. It is unclear whether specific structural properties could render such problems tractable. Addressing this could shed new light on the interaction between information acquisition and contract design.

References

[1] Tal Alon, Paul Dütting, and Inbal Talgam-Cohen. 2021. Contracts with private cost per unit-of-effort. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 52–69.

[2] Tal Alon, Inbal Talgam-Cohen, Ron Lavi, and Elisheva Shamash. 2024. Incomplete information VCG contracts for common agency. *Operations Research* 72, 1 (2024), 288–299.

[3] Moshe Babaioff, Michal Feldman, and Noam Nisan. 2006. Combinatorial agency. In *Proceedings of the 7th ACM Conference on Electronic Commerce*. 18–28.

[4] Moshe Babaioff, Michal Feldman, and Noam Nisan. 2009. Free-riding and free-labor in combinatorial agency. In *International Symposium on Algorithmic Game Theory*. Springer, 109–121.

[5] Moshe Babaioff, Michal Feldman, and Noam Nisan. 2010. Mixed strategies in combinatorial agency. *Journal of Artificial Intelligence Research* 38 (2010), 339–369.

[6] Yakov Babichenko, Inbal Talgam-Cohen, Haifeng Xu, and Konstantin Zabarnyi. 2024. Information Design in the Principal-Agent Problem. In *Proc. of EC 2024*.

[7] Diane Brady. 2025. Is DeepSeek a DeepFake or real? Doubts and questions emerge on Wall Street and in Silicon Valley. *Fortune* (2025). https://fortune.com/2025/01/28/deepseek-deepfake-or-real/.

[8] Federico Cacciamani, Martino Bernasconi, Matteo Castiglioni, and Nicola Gatti. 2024. Multi-agent contract design beyond binary actions. *arXiv preprint arXiv:2402.13824* (2024).

[9] Matteo Castiglioni and Junjie Chen. 2025. Hiring for An Uncertain Task: Joint Design of Information and Contracts. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1758–1794.

[10] Matteo Castiglioni, Alberto Marchesi, and Nicola Gatti. 2023. Multi-agent contract design: How to commission multiple agents with individual outcomes. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 412–448.

[11] Lingjiao Chen, Matei Zaharia, and James Zou. 2024. How Is ChatGPT's Behavior Changing Over Time? *Harvard Data Science Review* 6, 2 (2024).

[12] Paul Düetting, Tomer Ezra, Michal Feldman, and Thomas Kesselheim. 2025. Multi-agent combinatorial contracts. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1857–1891.

[13] Paul Dütting, Tomer Ezra, Michal Feldman, and Thomas Kesselheim. 2022. Combinatorial contracts. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 815–826.

[14] Paul Dütting, Tomer Ezra, Michal Feldman, and Thomas Kesselheim. 2023. Multi-agent contracts. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 1311–1324.

[15] Paul Dütting, Michal Feldman, and Inbal Talgam-Cohen. 2024. Algorithmic contract theory: A survey. *Foundations and Trends in Theoretical Computer Science* 16, 3-4 (2024), 211–412.

[16] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. 2019. Simple versus optimal contracts. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. 369–387.

[17] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. 2021. The complexity of contracts. *SIAM J. Comput.* 50, 1 (2021), 211–254.
[18] Yuval Emek and Michal Feldman. 2012. Computing optimal contracts in combinatorial agencies. *Theoretical Computer Science* 452 (2012), 56–74.
[19] Tomer Ezra, Stefano Leonardi, and Matteo Russo. 2024. Contracts with Inspections. *arXiv preprint arXiv:2402.16553* (2024).
[20] Alireza Fallah and Michael Jordan. 2024. Contract design with safety inspections. In *Proceedings of the 25th ACM Conference on Economics and Computation*. 616–638.
[21] Daniel F Garrett, George Georgiadis, Alex Smolin, and Balázs Szentes. 2023. Optimal technology design. *Journal of Economic Theory* 209 (2023), 105621.
[22] Gurobi Optimization, LLC. 2024. Gurobi Optimizer Reference Manual. https://www.gurobi.com
[23] Guru Guruganesh, Yoav Kolumbus, Jon Schneider, Inbal Talgam-Cohen, Emmanouil-Vasileios Vlatakis-Gkaragkounis, Joshua R Wang, and S Matthew Weinberg. 2024. Contracting with a learning agent. *arXiv preprint arXiv:2401.16198* (2024).
[24] Nicole Immorlica, Brendan Lucier, and Aleksandrs Slivkins. 2024. Generative AI as Economic Agents. *SIGecom Exch.* 22, 1 (2024), 93–109.
[25] Eden Saig, Ohad Einav, and Inbal Talgam-Cohen. 2024. Incentivizing Quality Text Generation via Statistical Contracts. In *Advances in Neural Information Processing Systems 38*.
[26] Eden Saig, Ohad Einav, and Inbal Talgam-Cohen. 2024. Incentivizing Quality Text Generation via Statistical Contracts. *arXiv preprint arXiv:2406.11118* (2024).
[27] Eden Saig, Inbal Talgam-Cohen, and Nir Rosenfeld. 2023. Delegated Classification. In *NeurIPS 2023*.

A Proof of Theorem 3.9 (Hardness of Symmetric-ISOP)

We show hardness via a series of two reductions from a variant of Set Cover. The intermediate problem is a promise problem which we call *Row-Sparsest Matrix*, or *RSM* for short. An instance of RSM is a tuple (G, A, y), where $G = (V, E)$ is a simple, undirected graph on $|V| = n$ vertices numbered $1, 2, \ldots, n$ with no isolated vertices, with $A \subseteq V$ and $y \in [n]$. The instance is promised to fall into one of the following two cases, with the objective being to determine which one holds.

YES instance. There exists an $n \times n$ matrix M of nonnegative real numbers such that:

(1) The average value in M is 1.
(2) At most y rows of M contain nonzero values.
(3) For all $\{i, j\} \in E$, we have $M_{i,i} = M_{i,j} = M_{j,i} = M_{j,j} = 0$.
(4) For all $k \in A$, $\sum_{i \in [n]} M_{i,k} + \sum_{i \in [n]} M_{k,i} \geq 10n$

NO instance. There does *not* exist an $n \times n$ matrix M of nonnegative real numbers such that:

(1) The average value in M is between 1 and $\frac{12}{11}$.
(2) At most y rows of M contain values greater than or equal to $\frac{5}{11}$.
(3) For all $\{i, j\} \in E$, we have $M_{i,i}, M_{i,j}, M_{j,i}, M_{j,j} \leq \frac{5}{11}$.
(4) For all $k \in A$, $\sum_{i \in [n]} M_{i,k} + \sum_{i \in [n]} M_{k,i} \geq n$.

Lemma A.1. *The problem RSM is NP-hard.*

Proof. We reduce from the variant of Set Cover where the number of elements is restricted to be exactly $\frac{1}{9}$ the number of sets. (Set Cover remains NP-hard with this restriction since we can duplicate elements and/or sets until the equality is satisfied.) The input to Set Cover consists of a collection of elements $x_1, x_2, \ldots, x_m$, a collection of sets of elements which we number $S_{m+1}, S_{m+2}, \ldots, S_n$ and a target value $y \in [n - m]$. The objective is to determine whether a collection of at most y of the sets covers all m elements. By our assumption on the number of elements versus the number of sets, we have $n = 10m$.

Given such an instance, let H be the bipartite graph where there is an edge between $i \in [n]$ and $j \in [n]$ if $i \leq m$, $j > m$, and element x_i is contained in set S_j. We then output the RSM instance (G, A, y) where G is the complement of H and $A = [m]$. We will show that, if a set cover of size y exists, then (G, A, y) is YES instance; otherwise, it is a NO instance.

First suppose C is a set cover of size y. For each element x_i, let $M_{j,i} = 10n$ for one arbitrary j such that $x_i \in S_j \in C$ (which must exist since C is a set cover). Let all other entries of M be zero. Observe that M satisfies all four YES instance properties:

(1) The sum of all nonzero elements is $10n \cdot m = n^2$, so the average value is 1.
(2) Only the y rows corresponding to the sets in the cover have nonzero entries.
(3) We have a nonzero value at (i, j) only when $\{i, j\}$ is an edge in H. This does not include any edges of G (since G is the complement of H) or diagonal entries.
(4) For each $i \in [m]$, we know that, for some j, $M_{j,i} = 10n$.

Conversely, suppose (G, A, y) is *not* a NO instance. This means there *does* exists a matrix M satisfying all of the NO instance properties. Let I be the set of indices $i \in [m]$ such that the i^{th} row of M contains values greater than or equal to $\frac{1}{2}$; Let J be the set of such indices from $m + 1$ to n. For each $i \in I$, choose an arbitrary column j such that $M_{i,j}$ is at least $\frac{1}{2}$. Let $f : I \longrightarrow [n]$ be the mapping of arbitrary choices for each such index i. Consider the collection of sets

$$C := \{S_{f(i)} \mid i \in I\} \cup \{S_j \mid j \in J\}.$$

Note that $|C| \le |I| + |J| \le y$ from property (2). We claim that C is a set cover. Consider an arbitrary element x_i. We know that there must exist some $j \in [n]$ such that $M_{j,i} \ge \frac{1}{2}$ or $M_{i,j} \ge \frac{1}{2}$, for otherwise we would have

$$\sum_{j \in [n]} M_{j,i} + \sum_{j \in [n]} M_{i,j} < 2n \cdot \frac{1}{2} = n,$$

violating property (4). Furthermore, from the definition of G, we know that j must be an index greater than m, and S_j contains x_i, for otherwise both cases $M_{j,i} \ge \frac{1}{2}$ and $M_{i,j} \ge \frac{1}{2}$ would violate property (3). In the former case ($M_{j,i} \ge \frac{1}{2}$), we have $j \in J$, so $x_i \in S_j \in C$. In the latter case ($M_{i,j} \ge \frac{1}{2}$), we have $i \in I$, so $x_i \in S_{f(i)} \in C$. □

Proof of Theorem 3.9. We reduce from RSM, which is NP-hard by Lemma A.1. Given an instance (G, A, y) where $G = (V, E)$ has n vertices, we define an instance with $n+1$ signals/outcomes, numbered $0, 1, 2, \dots, n$. Let

$$\varepsilon := 1 - \sqrt{\frac{(n-1)^2(n+2)}{n^3}} > 0. \tag{9}$$

The reward for any realization (k, j), where $k, j \in \{0, 1, 2, \dots, n\}$, is $\frac{1}{\varepsilon}$ times the number of vertices realized (i.e. nonzero values among k and j). All signals cost $\frac{1}{11}$ to get a second opinion, which will be second independent sample from the same probability distribution.

There are three kinds of actions the agent can take, enumerated as follows.[2]

- Action g (the *good action*), which costs 1 and leads to a uniformly random vertex.
- For each vertex $k \in A$, an action a_k which costs 1. With probability ε, outcome 0 is realized; and with probability $1 - \varepsilon$, a uniformly random vertex *other than* k is realized.
- For each edge of G, between vertex i and vertex j, an action $e_{i,j}$, which costs 0. With probability ε, outcome 0 is realized; with probability $\frac{1-\varepsilon}{2}$, i is realized; and with probability $\frac{1-\varepsilon}{2}$, vertex j is realized.

We will show that, if (G, A, y) is a YES instance, there exists a contract where the principal can attain expected utility at least

$$U := \frac{2}{\varepsilon} - 1 - \frac{y}{11n};$$

Whereas if it is a NO instance, there does not exist such a contract. We derive an equivalent interpretation of this utility target as follows. If the agent takes any action other than the good action g, a vertex is realized each draw with probability at most $1 - \varepsilon$, so the expected reward of the principal is at most

$$2 \cdot (1 - \varepsilon) \cdot \frac{1}{\varepsilon} = \frac{2}{\varepsilon} - 2 < U.$$

Hence, it is not possible for the principal to attain expected utility at least U if the agent is taking any action besides g. On the other hand, when the agent takes action g, the probability of realizing a vertex is 1 each draw, so the principal's expected reward is

$$2 \cdot \frac{1}{\varepsilon} = U + 1 + \frac{y}{11n}.$$

Thus, the principal can obtain expected utility at least U if and only if it is possible to incentivize the agent to take action g by paying at most $1 + \frac{y}{11n}$ in expectation.

For the forward direction, suppose (G, A, y) is a YES instance. Consider the contract t that pays $M_{i,j}$ when vertex i is realized on the first draw and vertex j is realized on the second draw. Any time outcome 0 is observed, the payment is zero. By property (2), all but y rows of M are all-zero,

[2]To maintain consistency in indices used for vertices throughout the proof, we break with the convention that i is for actions, j is for outcomes, and k is for signals.

meaning they do not require inspection: Upon observing a vertex i for which the i^{th} row of M is all-zero, the principal will simply pay zero and not inspect. Thus the total inspection cost is

$$(\text{num inspected outcomes}) \cdot (\text{probability of outcome}) \cdot (\text{inspection cost}) = y \cdot \frac{1}{n} \cdot \frac{1}{11} = \frac{y}{11n}.$$

Additionally, the total expected transfer from the principal to the agent under the good action g is

$$T_g = \sum_{i \in [n]} \sum_{j \in [n]} \frac{1}{n^2} M_{i,j} = 1.$$

Hence, when the agent takes the good action g, the principal's total expected payment is $1 + \frac{y}{11n}$ as desired. It remains to check the incentive constraints, that g is optimal for the agent.

Each of the $e_{i,j}$ actions yields a transfer of 0, as the only possible outcomes are combinations of i, j, and 0, which all result in zero payment by property (3). Thus, $e_{i,j}$ costs one less than g but also earns one less, so the agent weakly prefers g.

Next consider an action a_k. We may lower-bound the difference between the expected transfer T_g if the agent chooses g and T_{a_k} if the agent chooses a_k as follows:

$$\begin{aligned}
T_g - T_{a_k} &= \frac{1}{n^2} M_{k,k} + \sum_{i \in [n] \setminus \{k\}} \left(\frac{1}{n^2}\right) \left(M_{i,k} + M_{k,i}\right) + \sum_{i,j \in [n] \setminus \{k\}} \left(\frac{1}{n^2} - \frac{(1-\varepsilon)^2}{(n-1)^2}\right) M_{i,j} \\
&= \frac{1}{n^2} M_{k,k} + \left(\frac{1}{n^2}\right) \sum_{i \in [n] \setminus \{k\}} \left(M_{i,k} + M_{k,i}\right) + \left(-\frac{2}{n^3}\right) \sum_{i,j \in [n] \setminus \{k\}} M_{i,j} \\
&\qquad \text{(rearranging Equation (9))} \\
&\geq \frac{1}{n^2} M_{k,k} + \frac{1}{5n^2} \sum_{i \in [n] \setminus \{k\}} \left(M_{i,k} + M_{k,i}\right) - \frac{2}{n^3} \sum_{i,j \in [n]} M_{i,j} \\
&\qquad \left(\text{since } \frac{4}{5n^2} \geq \frac{2}{n^3} \text{ for } n \geq 3\right) \\
&\geq \frac{2}{5n^2} M_{k,k} + \frac{1}{5n^2} \sum_{i \in [n] \setminus \{k\}} \left(M_{i,k} + M_{k,i}\right) - \frac{2}{n^3} \sum_{i,j \in [n]} M_{i,j} \\
&= \frac{1}{5n^2} \left(\sum_{i \in [n]} M_{i,k} + \sum_{i \in [n]} M_{k,i} \right) - \frac{2}{n^3} \sum_{i,j \in [n]} M_{i,j} \\
&\geq \frac{1}{5n^2} (10n) - \frac{2}{n^3} \sum_{i,j \in [n]} M_{i,j} \qquad \text{(from property (4))} \\
&= \frac{2}{n} \left(1 - \frac{1}{n^2} \sum_{i,j \in [n]} M_{i,j} \right) = 0,
\end{aligned}$$

where in the final equality we have used property (1). Thus, the agent earns weakly more from taking action g than action a_k. Since the two actions cost the same, the agent weakly prefers g.

For the backward direction, suppose there is a contract in which the principal incentivizes action g by paying at most $1 + \frac{y}{11n}$. Our objective is to show that (G, A, y) is *not* a NO instance. Let M be the matrix where $M_{i,j}$ is the payment upon realizing vertex i from the first draw and j from the second draw. Note that M is constant on rows corresponding to vertex indices that are not inspected. We will show that M satisfies all four NO instance properties.

First, observe that, since the contract incentivizes the costly action g, it must transfer at least $T_g \geq 1$ in expectation from the principal to the agent, otherwise the agent is better off taking one of the actions that costs 0. It follows that the principal can only inspect y actions, for otherwise the total cost exceeds $1 + \frac{y}{11n}$. On the other hand,

$$T_g \leq 1 + \frac{y}{11n} \leq 1 + \frac{1}{11} = \frac{12}{11}.$$

This proves property (1), since T_g is precisely the average value of the matrix M.

For any edge $\{i, j\}$ and any entry $z \in \{M_{i,i}, M_{i,j}, M_{j,i}, M_{j,j}\}$, we have

$$\begin{aligned}
z &< \frac{5(1-\varepsilon)^2}{4} z \qquad \text{(for large enough } n\text{)} \\
&\leq \frac{5(1-\varepsilon)^2}{4} (M_{i,i} + M_{i,j} + M_{j,i} + M_{j,j}) \\
&= 5T_{e_{i,j}} \\
&\leq 5(T_g - 1) \qquad \text{(since the contract incentivizes } g\text{)} \\
&\leq 5\left(\frac{12}{11} - 1\right) \qquad \text{(from the inequality above)} \\
&= \frac{5}{11}.
\end{aligned}$$

This establishes property (3). Furthermore, since each row i is constant if i is not inspected, and $M_{i,i} < \frac{5}{11}$ (each diagonal entry appears as some z in the argument above because we assume G has no isolated vertices), we have that the entire row i must be less than $\frac{5}{11}$ if i is not inspected. As only y rows are inspected, property (2) follows.

Finally, for each $k \in A$, from the incentive constraint that the agent prefers action g to a_k, we have

$$\begin{aligned}
0 &\leq T_g - T_{a_k} \\
&\leq \frac{1}{n^2} M_{k,k} + \left(\frac{1}{n^2}\right) \sum_{i \in [n] \setminus \{k\}} \left(M_{i,k} + M_{k,i}\right) + \left(-\frac{2}{n^3}\right) \sum_{i,j \in [n] \setminus \{k\}} M_{i,j} \\
&\qquad \text{(the inequality is only due to the fact that the contract may pay for outcome 0)} \\
&\leq \frac{4}{n^2} M_{k,k} + \frac{2}{n^2} \sum_{i \in [n] \setminus \{k\}} \left(M_{i,k} + M_{k,i}\right) - \frac{2}{n^3} \sum_{i,j \in [n]} M_{i,j} \qquad \left(\text{since } \frac{2}{n^2} \geq \frac{2}{n^3}\right) \\
&= \frac{2}{n^2} \left(\sum_{i \in [n]} M_{i,k} + \sum_{i \in [n]} M_{k,i} - nT_g \right) \\
&\leq \frac{2}{n^2} \left(\sum_{i \in [n]} M_{i,k} + \sum_{i \in [n]} M_{k,i} - n \right).
\end{aligned}$$

This implies property (4):

$$\sum_{i \in [n]} \left(M_{i,k} + M_{k,i}\right) \geq n.$$

□

B Proofs of Theorems 4.1 and 4.3 (Characterization of Equilibrium Existence)

We will require the following lemma for both proofs.

Lemma B.1. *Fix any feasible solution* (p, s, t) *to a QCQP from CoMI and a signal* $k \in [\ell]$ *such that* $p_k \in (0, 1]$. *For any* $p'_k \in (0, p_k)$, *there is some value of* s'_k *and vector* $t'_k = (t'_{k,1}, t'_{k,2}, \ldots, t'_{k,m_k})$ *giving a feasible instance* (p', s', t'), *where* p', s', *and* t' *are the same as* p, s, *and* t *except on indexes involving signal* k. *Furthermore:*

(1) If $q^0_{i,k} d_k > 0$, *then* (p', s', t') *has a strictly better objective value than* (p, s, t)*; specifically, the principal's expected payment changes by* $q^0_{i,k}(p'_k - p_k)d_k < 0$.

(2) The map $p'_k \mapsto (s'_k, t'_k)$ *is continuous over the open interval* $(0, p_k)$.

Proof of Lemma B.1. For each $j \in [m_k]$, we define

$$s'_k := \frac{1 - p_k}{1 - p'_k} \cdot s_k, \qquad\qquad t'_{k,j} := \frac{p_k}{p'_k} \cdot t_{k,j}.$$

This is clearly continuous over the open interval $(0, p_k)$. Observe that, for all $i' \in [n]$ (including $i' = i$, where i is the specified action the principal is trying to incentivize), we have

$$\begin{aligned}(1 - p'_k)s'_k + p'_k \sum_{j\in[m_k]} q^k_{i',j} t'_{k,j} &= (1 - p'_k)\frac{1 - p_k}{1 - p'_k} s_k + p'_k \sum_{j\in[m_k]} q^k_{i',j} \frac{p_k}{p'_k} t_{k,j} \\ &= (1 - p_k)s_k + p_k \sum_{j\in[m_k]} q^k_{i',j} t_{k,j}\end{aligned} \tag{10}$$

Note that this equality holds for all signals (by definition), not just the specific k for which the variables changed. It follows that the IC constraint still holds. All of the other constraints obviously continue to hold as well.

In what follows, we use a signal index variable k' to avoid clashing with the specific signal k from the theorem statement. The objective value of (p', s', t') is

$$\begin{aligned}&\sum_{k'\in[\ell]} q^0_{i,k'} \left((1 - p'_{k'})s'_{k'} + p'_{k'} \left(d_{k'} + \sum_{j\in[m_{k'}]} q^{k'}_{i,j} t'_{k',j} \right)\right) \\ &= \sum_{k'\in[\ell]} q^0_{i,k'} \left((1 - p'_{k'})s'_{k'} + p'_{k'} \sum_{j\in[m_{k'}]} q^{k'}_{i,j} t'_{k',j} \right) + \sum_{k'\in[\ell]} q^0_{i,k'} p'_{k'} d_{k'} \\ &= \sum_{k'\in[\ell]} q^0_{i,k'} \left((1 - p_{k'})s_{k'} + p_{k'} \sum_{j\in[m_{k'}]} q^{k'}_{i,j} t_{k',j} \right) + \sum_{k'\in[\ell]} q^0_{i,k'} p'_{k'} d_{k'} \qquad \text{(by Equation (10) above)} \\ &= \sum_{k'\in[\ell]} q^0_{i,k'} \left((1 - p_{k'})s_{k'} + p_{k'} \sum_{j\in[m_{k'}]} q^{k'}_{i,j} t_{k',j} \right) + \sum_{k'\in[\ell]} q^0_{i,k'} p_{k'} d_{k'} + \sum_{k'\in[\ell]} q^0_{i,k'} (p'_{k'} - p_{k'}) d_{k'} \\ &= \sum_{k'\in[\ell]} q^0_{i,k'} \left((1 - p_{k'})s_{k'} + p_{k'} \left(d_{k'} + \sum_{j\in[m_{k'}]} q^{k'}_{i,j} t_{k',j} \right)\right) + \sum_{k'\in[\ell]} q^0_{i,k'} (p'_{k'} - p_{k'}) d_{k'}.\end{aligned}$$

Since the first sum in the final line above is the objective value of (p, s, t), we see that the difference is

$$\sum_{k'\in[\ell]} q^0_{i,k'} (p'_{k'} - p_{k'}) d_{k'}.$$

Each of the terms in this sum is zero by definition, except for $k' = k$. Thus, the difference in objective values is precisely $q^0_{i,k}(p'_k - p_k)d_k$ as claimed. □

Proof of Theorem 4.1. Fix a CoMI instance X. First note that the instance $\widehat{X}$ with zero inspection costs has an optimal solution because it is without harm to inspect all signals, so we may set $p_k = 1$ for all signals, obtaining a linear program.

We first prove the second part of the claim, that there is a sequence of solutions to X whose value converges to the optimal value of $\widehat{X}$. Let (p, s, t) be an optimal solution with value y. Let S be the set of signals $k \in [\ell]$ such that $p_k > 0$, and suppose that the minimum nonzero p_k value is δ. In X, which has a different objective function but the same feasible set, the value of (p, s, t) is $y + \sum_{k \in S} q^0_{i,k} p_k d_k$. For any $0 < \varepsilon < \delta$, we repeatedly apply Lemma B.1 to each signal $k \in S$ such that, obtaining a new solution which improves the objective value by an additive $\sum_{k \in S} q^0_{i,k}(\varepsilon - p_k) d_k$. Thus, the new solution, call it $(p^\varepsilon, s, t^\varepsilon)$, has objective value

$$y + \sum_{k \in S} q^0_{i,k} p_k d_k + \sum_{k \in S} q^0_{i,k}(\varepsilon - p_k) d_k = y + \varepsilon \sum_{k \in S} q^0_{i,k} d_k$$

Sending $\varepsilon \longrightarrow 0$, we see that the value of $(p^\varepsilon, s, t^\varepsilon)$ converges to y, as desired.

We now prove the characterization of when X has an optimal solution. Clearly, if $\widehat{X}$ has an optimal solution with $p_k = 0$ for all k such that $q^0_{i,k} d_k > 0$, then that same solution yields the same objective value in $\widehat{X}$. This is optimal because the minimum objective value in X is necessarily at least the minimum objective value of $\widehat{X}$. Conversely, suppose that X has an optimal solution (p, s, t) of some value y'. Then notice that $y' = y$ by the claim proved in the previous paragraph, since there are certainly solutions to X of value arbitrarily close to y. This means that (p, s, t) is an optimal solution to $\widehat{X}$, and it cannot possibly inspect any signal k for which $q^0_{i,k} d_k > 0$ with nonzero probability, for otherwise its objective value would be different in X and $\widehat{X}$. □

Proof of Theorem 4.3. In all three problem variants, every QCQP/LCQP has a continuous objective function. Furthermore, the feasible set is closed, as all constraints are specified by weak inequalities. Hence, as long as the feasible set is also bounded, an optimal solution exists, as it is the result of minimizing a continuous function over a compact set. Note that all relevant variables are bounded below, and the p_k variables are bounded above (by one). Thus, an optimal solution is guaranteed to exist as long as each s_k and $t_{k,j}$ is bounded above. We will show that, in each of UMI, CoNI, and UNI, we may impose additional constraints upper-bounding these variables without harming the value of the optimal solution.

We begin with UMI and UNI, where we saw that it is possible to rewrite each QCLP so that it only involves variables $t_{k,j}$ and not any s_k. Fix an action $i \in [n]$ and a set of non-inspected signals $S_0 \subseteq [\ell]$. For each $k \in [\ell]$ and $j \in [m_k]$, there are two cases to consider. If $q^0_{i,k} \cdot q^k_{i,j} = 0$, then variable $t_{k,j}$ irrelevant to the game, so we may set it to zero without loss of generality. Otherwise, if $q^0_{i,k} \cdot q^k_{i,j} > 0$, then we may upper-bound $t_{k,j}$ by $(\max_{i' \in [n]} R_{i'})/(q^0_{i,k} q^k_{i,j})$, for if $t_{k,j}$ is greater than this quantity, then the expected payment from the principal to the agent is more than $(\max_{i' \in [n]} R_{i'})$. This means the principal's utility is negative, so the principal would have been better off with all-zero payments. Thus, in either case, we may upper bound all variables without harming the optimal objective value.

We next consider CoNI. By similar reasoning, we first claim that we may upper bound each $t_{k,j}$ by 0 (in the case where $q^0_{i,k} \cdot q^k_{i,j} = 0$ for the given action i) or $(\max_{i' \in [n]} R_{i'})/(q^0_{i,k} q^k_{i,j})$ (otherwise). To see why the latter bound still holds, observe that, if it is violated, then the objective function contains the term

$$q^0_{i,k}\left((1 - p_k)s_k + p_k\left(d_k + q^k_{i,j} t_{k,j}\right)\right) \geq q^0_{i,k}\left((1 - p_k)t_{k,j} + p_k q^k_{i,j} t_{k,j}\right) \geq q^0_{i,k} q^k_{i,j} t_{k,j} > \max_{i' \in [n]} R_{i'},$$

where in the first inequality we have used the negative inspection constraint and dropped the non-negative d_k term, and in the second inequality we have used the fact that $q^k_{i,j} \leq 1$. So as before, the principal would have been better off with all-zero payments.

Having established that the $t_{k,j}$ variables are bounded, we next proceed to bound the s_k variables. Specifically, we claim that it is without harm to the objective value to bound

$$s_k \leq \max_{j \in [m_k]} U_{k,j},$$

where $U_{k,j}$ is our previous upper bound on $t_{k,j}$. Suppose this inequality is violated. If $p_k = 0$, then we may easily again derive that the principal's payment is more than $\max_{i' \in [n]} R_{i'}$. Otherwise, Lemma B.1 implies we can improve our objective value by locally decreasing p_k and increasing several $t_{k,j}$. Since s_k is strictly larger than the largest $t_{k,j}$, a small enough perturbation will still respect the negative inspection constraint. □