Ambiguous Contracts[*]

Paul Dütting[†] Michal Feldman[‡] Daniel Peretz[‡]
Larry Samuelson[§]

June 20, 2024

Abstract

We explore the deliberate infusion of ambiguity into the design of contracts. We show that when the agent is ambiguity-averse and hence chooses an action that maximizes their minimum utility, the principal can strictly gain from using an ambiguous contract, and this gain can be arbitrarily high. We characterize the structure of optimal ambiguous contracts, showing that ambiguity drives optimal contracts towards simplicity. We also provide a characterization of ambiguity-proof classes of contracts, where the principal cannot gain by infusing ambiguity. Finally, we show that when the agent can engage in mixed actions, the advantages of ambiguous contracts disappear.

1 Introduction

Contracts are often ambiguous. A construction contract may require that a builder use "superior materials," a professional services contract may require that a provider

[*]Accepted for presentation at EC 2023. We thank Tomer Ezra, Uriel Feige and Yarden Rashti for insightful feedback, and the co-editor and referees for helpful comments and suggestions. The work of Feldman and Peretz has been supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 866132), the Amazon Research Award, the NSF-BSF (grant number 2020788), and by a grant from the Tel Aviv University Center for AI and Data Science (TAD).

[†]Google Research, Zurich, Switzerland. Email: `duetting@google.com`

[‡]School of Computer Science, Tel Aviv University, Israel. Email: `mfeldman@tauex.ac.il`, `danielperetz@mail.tau.ac.il`

[§]Department of Economics, Yale University, New Haven, CT, USA. Email: `larry.samuelson@yale.edu`

exert "due diligence," or "act as a fiduciary," a labor contract may require that the parties "bargain in good faith," and the promotion guidelines of a university may require that a candidate exhibit "research productivity and excellence." In each case, the meaning of these phrases may be ambiguous.

This ambiguity may reflect the impossibility of precise specification. In contrast, we explore the deliberate infusion of ambiguity as a tool that a principal may employ to increase her contracting power.

1.1 The Model

We examine a familiar, finite moral hazard problem, augmented to accommodate ambiguous contracts analogously to the treatment of mechanism design problems by Di Tillio et al. (2017). In the problem we consider, a principal (she) interacts with an agent (he). The agent can take one of n costly actions. Each action $i \in [n]$ induces a probability distribution $p_i = (p_1, \ldots, p_m)$ over m outcomes and imposes a non-negative cost c_i on the agent. Each outcome $j \in [m]$ comes with a reward r_j for the principal. The principal cannot directly observe the agent's action, and seeks to influence the agent's choice of action by paying for the stochastic outcomes of the action taken by the agent. The principal and agent are both risk neutral, and a limited liability constraint forces payments to be non-negative.

A classic contract for this setting includes a payment function $t = (t_1, \ldots, t_m)$, where t_j specifies the non-negative payment from the principal to the agent when outcome $j \in [m]$ is realized. Given a payment function t, the agent chooses an action $i \in [n]$ that maximizes his expected payment minus cost. The principal, in turn receives the expected reward of the implemented action, minus the expected payment to the agent under the chosen action.

We are interested in ambiguous contracts. The source of the ambiguity can be given various interpretations. It may be that the contract has missing provisions, vaguely worded provisions, or provisions specified in prohibitively copious fine print,[1] each case leaving the agent facing a set of possible realized contracts that he cannot refine. The resolution of the ambiguity may reflect a decision on the part of the principal, the action of a third party such as a court, or random events.

[1] Zuboff (2023) argues that pyramiding cross references can render it impossible to read all of the fine print in a typical contract.

To capture this ambiguity, we define an *ambiguous contract* to be a collection of payment functions $\tau = \{t^1, t^2, \ldots, t^k\}$. The agent evaluates each action $i \in [n]$ by the minimum utility he could receive from a payment function $t^\ell \in \tau$, and chooses an action that maximizes this minimum utility. We say that the ambiguous contract implements the selected action. We impose the consistency condition that, for the implemented action, every payment function in the support of τ gives the principal the same payoff. We motivate this requirement as ensuring that the principal's "threat" that she may choose any payment function $t \in \tau$ is credible, though we show that the requirement is without loss of generality.

1.2 Our Contribution

Section 2 sets up the model and examines the sets of implementable actions. We recall the familiar conditions (Hermalin and Katz (1991)) for an action to be implementable with a classic contract, and then characterize implementable actions under ambiguous contracts. In general, ambiguity expands the set of implementable actions.

Section 3 examines optimal ambiguous contracts. We first show that the principal can use ambiguity to her advantage, either because ambiguity allows her to implement the optimal action under classic contracts at a reduced cost, or because she exploits the ambiguity to implement a different action. Indeed, a principal wielding ambiguous contracts may optimally induce an action that is impossible to implement under classic contracts.

We then show that ambiguity drives optimal contracts towards simplicity. In general, optimal ambiguous contracts are composed of at most $\max\{n-1, m\}$ payment functions, each of which can be taken to attach a positive payment to precisely one outcome (i.e., is a single-outcome payment function). If we restrict attention to principal-agent problems satisfying the monotone likelihood ratio property (MLRP), then optimal ambiguous contracts contain at most two (single-outcome payment) contracts. If payment functions must be monotone—higher rewards engender higher payments—perhaps for reasons of fairness, regulation, or robustness, then analogous results hold, with payment functions now being step functions (with a single step) rather than paying for only a single outcome.

Section 4 uses the concept of an *ambiguity gap*—the largest possible ratio of the principal's payoff under an optimal ambiguous contract to that of an optimal classic contract—to quantify how much the principal gains by exploiting ambiguity. In general, this ratio can be arbitrarily large. When all rewards are positive, the ambiguity gap is $n - 1$, and hence grows arbitrarily large when the number of actions grows large.

Section 5 defines a class of contracts to be *ambiguity-proof* if it is impossible for the principal to implement an action at a lower expected payment with an ambiguous contract than with a classic contract. We show that a class of contracts is ambiguity-proof if and only if it is *ordered*, in the sense that for any two contracts in the class, one of them attaches a weakly higher payment to every outcome than does the other. An immediate implication of this result is that the classes of linear contracts is ambiguity-proof, among others. In contrast, many other natural classes of contracts, such as the classes of all affine, polynomial or monotone contracts, are *not* ambiguity-proof.

Section 6 shows that the advantages of ambiguity disappear if the agent can mix over actions. The ability to mix provides the agent with more alternative actions, tightening the incentive constraints enough to dissipate any advantage the principal gains from ambiguous contracts. As explained by Raiffa (1961) in his assessment of the Ellsberg (1961) paradox, mixing allows the agent to transform a situation of ambiguity into one of risk, alleviating the force of ambiguity-aversion.

1.3 Implications

We do not expect to see agents literally facing a bevy of payment functions, wondering which will actually be applied, but we do view our results as helpful in understanding three aspects of real-world contracts.

First, we believe that contracts typically are ambiguous. Indeed, it is difficult to imagine a contract that specifies beyond any doubt the implications of every outcome. The literature has focused on feasibility constraints as the source of such imprecision—it may be prohibitively expensive or impossible to anticipate, describe, or verify the various outcomes.[2] In contrast, we suggest that ambiguity may be deliberately embraced by the principal as an incentive device. As our ambiguity gap results show, the gains to the principal from doing so can be large.

Second, a theme that emerges from our results is that optimal ambiguous contracts tend to be simple. We expect circumstances will often constrain contracts to be

[2] Aghion and Holden (2011) provides a survey of a literature that has its roots in (Hart, 1988) and (Hart and Moore, 1988).

monotone. In this case, the optimal ambiguous contract features a collection of step functions, each with a single step. In practice, this would take the form of a single contract, specifying bonuses if various performance thresholds are reached, but written sufficiently imprecisely as to make the performance thresholds and bonuses ambiguous. If the technology satisfies the monotone likelihood ratio property, the number of such performance-dependent bonuses is small, namely two. We thus have a contract that requires at least adequate performance in order to elicit payment, with a bonus for superior performance, written so as to allow some ambiguity as to the precise performance thresholds and payments. We view actual contracts in many circumstances as fitting this description. For example, an assistant professor may believe that a reasonable research record engenders promotion to associate professor and a raise, with an exemplary record bringing promotion to full professor and a larger raise, with the thresholds and the amounts of the raises both ambiguous.

Third, ambiguous contracts can be a burden for agents, either because they are more difficult to evaluate and enforce or because they are a weapon for extracting surplus from agents. Circumstances may accordingly restrict attention to ambiguity-proof classes of contracts. Our results show that insisting on ambiguity proofness also drives contracts toward (a different notion of) simplicity. Ordered contracts, with linear contracts as a leading example, are straightforward to process. We note that commission contracts are common, and in their simplest form are linear.[3]

The literature has established other reasons why contracts may be imprecise or simple or linear. In each case we cannot claim to have provided "the" explanation, but our work adds another factor to the list of considerations.

1.4 Related Literature

We work with a familiar hidden-action moral hazard problem, as in Holmström (1979), Grossman and Hart (1983), and Laffont and Martimort (2009, Chapter 4), with the friction arising out of limited liability (as in Innes (1990)) rather than risk aversion. In contrast to much of the moral hazard literature, our principal offers ambiguous contracts to an ambiguity-averse agent. We implement the agent's ambiguity aversion by modeling the agent as maximizing his max-min utility (Schmeidler, 1989; Gilboa

[3]A contract may include a base payment plus a commission, making it affine rather than linear. Section 5 notes that the class of all affine contracts is not ambiguity proof, but there are ambiguity-proof classes of affine contracts.

and Schmeidler, 1993).

A flourishing literature examines design problems in the face of non-Bayesian uncertainty. One branch of this literature examines models in which the principal entertains non-Bayesian uncertainty about the agents. Bergemann and Schlag (2011) examine monopoly pricing on the part of a principal with ambiguous beliefs about buyers' valuations. Carrasco et al. (2018) examine screening problems in which the principal is only partially informed of the distribution of agent preferences. Carroll (2015), Carroll and Walton (2022) and Kambhampati (2023) examine moral hazard problems in which the principal has ambiguous beliefs about the set of actions the agent can choose from. Dai and Toikka (2022) examine a principal who writes contracts to shape the actions of a team of agents, with the principal holding ambiguous beliefs about the actions available to the agents. Dütting et al. (2019) examine moral hazard problems in which the principal has ambiguous beliefs about the distribution of outcomes induced by the agent's actions.

A second branch of the literature examines settings in which the agent has ambiguous beliefs that the principal can potentially exploit. Beauchêne et al. (2019), Cheng (2020), and Cheng et al. (2024) examine Bayesian persuasion problems in which the sender exploits the ambiguity aversion of the receiver. Bodoh-Creed (2012) and Di Tillio et al. (2017) examine screening problems with agents who have max-min preferences. Bose and Renou (2014) examine mechanism design problems in which agents have max-min preferences. Lopomo et al. (2011) examine moral hazard problems with agents who have Bewley preferences. Bose et al. (2006) consider auctions in which the seller and bidders may both be ambiguity averse. Our paper is distinguished from the two branches above by examining moral hazard problems in which the agent faces ambiguity concerning the payments attached to outcomes.

The paper most closely related to our work is Di Tillio et al. (2017), who conduct a parallel exercise in the context of a screening model, with a seller allocating an object to an ambiguity-averse buyer. They find that the seller can gain from offering an ambiguous mechanism, consisting of a set of simple mechanisms, just as our principal can benefit from offering an ambiguous contract. In each case, the gain comes from using the agent's ambiguity aversion to relax incentive constraints. Di Tillio et al. (2017) show that an optimal ambiguous mechanisms contains at most $N-1$ mechanisms, where N is the number of agent types (as opposed to the number of actions, in our case), each of which takes a simple form reminiscent of our single-outcome payment functions. Finally, they show that as the number of agent types grows arbitrarily large, the principal comes arbitrarily close to extracting all of the surplus.

Di Tillio et al. (2017) impose a consistency condition on ambiguous mechanisms, analogous to the consistency condition we impose on ambiguous contracts. In their case this restriction is substantive rather than sacrificing no generality, reflecting the differing structure of the incentive constraints that arise in screening and moral hazard problems. Di Tillio et al. (2017) do not have counterparts of our findings that the number of payment functions in an optimal ambiguous contract is precisely two in the presence of the MLRP condition, though they maintain throughout the screening counterpart of this assumption, in the form of a single-crossing condition.

An implication of our results is that in the context of moral hazard problems, ambiguity and max-min utility drive optimal designs towards simplicity. We thus join a literature, with Holmstrom and Milgrom (1987) as a key early entry, endeavoring to explain why actual contracts in moral hazard settings tend to be simple, in contrast to their theoretical counterparts. Carroll (2015), Carroll and Walton (2022), and Dai and Toikka (2022) show that a principal who is uncertain of the actions available to an agent and who has max-min preferences will optimally choose a linear contract. Dütting et al. (2019) show that the same holds for a principal who is uncertain about the technology by which actions turn into outcomes and who has max-min preferences. These papers thus show that ambiguity aversion on the part of the principal can lead to linear contracts, whereas we find linear contracts may be attractive as a device for preventing the principal from exploiting the agent's ambiguity aversion. In our setting, exploiting the agent's ambiguity aversion leads the principal to an alternative class of simple contracts, consisting of single-outcome payment functions or step functions (when payment functions must be monotone), and including only two such functions when the MLRP condition holds.

2 The Model

Our starting point is the classic hidden-action principal-agent model of contract theory (e.g., Holmström (1979), Grossman and Hart (1983), and Laffont and Martimort (2009, Chapter 4)). A principal (she) induces an agent (he) to take a costly, unobservable action by writing a contract that attaches payments to the observable, stochastic outcomes of the chosen action. We augment the model by introducing the notion of an *ambiguous contract* and examining optimal ambiguous contracts.

2.1 The Principal-Agent Model

The basic ingredients of the principal-agent model apply to both classic and ambiguous contracts.

Definition 1 (Instance)**.** *An instance (c, r, p) of the principal-agent problem with n actions and m outcomes is specified by:*

- *For each action $i \in [n]$ a non-negative cost $c_i \in \mathbb{R}_+$. We write $c = (c_1, \ldots, c_n)$ for the vector of costs, and sort the actions so that $c_1 \leq c_2 \leq \ldots \leq c_n$.*
- *For each outcome $j \in [m]$, a* reward *$r_j \in \mathbb{R}$. We write $r = (r_1, \ldots, r_m)$ for the vector of rewards, and sort outcomes so that $r_1 \leq r_2 \leq \ldots \leq r_m$.*
- *For each action $i \in [n]$, a probability distribution $p_i \in \Delta^m$. We use p_{ij} to denote the probability of outcome j under action i.*

We use $R_i = \sum_{j=1}^{m} p_{ij} r_j$ to denote the expected reward of action i, and write $W_i = R_i - c_i$ for action i's *expected welfare*.

The agent retains the option of not participating. To capture this, we assume throughout that action 1 is a zero-cost action, that leads with probability 1 to an outcome that we interpret as the status-quo outcome. As we explain in Sections 2.2–2.3, this does not limit the generality of the model, but leads to a unified treatment of the incentive compatibility and individual rationality constraints.

The literature often focuses on instances that satisfy the monotone likelihood ratio property:

Definition 2 (MLRP)**.** *An instance (c, r, p) satisfies the* MLRP *(monotone likelihood ratio property) if for any two actions i, i' such that $c_i > c_{i'}$, it holds for all $j' > j$ that:*

$$p_{ij'} p_{i'j} \geq p_{ij} p_{i'j'}.$$

Intuitively, the MLRP condition requires that more costly actions are more likely to yield high outcomes.

The specification (c, r, p) is known to both the principal and the agent. The agent's action is known only by the agent, while realized outcomes are observed by both the principal and agent.

A payment function $t : [m] \to \mathbb{R}_+$ identifies a payment made by the principal to the agent upon the realization of each outcome, with the payment $t(j)$ made in response to outcome j typically denoted by t_j. Payments are attached to outcomes rather than actions because the principal can observe only the former. We assume that payments are non-negative. This is a standard limited liability assumption.

In many cases, considerations of fairness, regulation, or robustness may restrict attention to monotone payment functions.

Definition 3. *The payment function t is* monotone *if outcomes generating larger rewards engender (at least weakly) larger payments: $r_j \geq r_{j'} \implies t_j \geq t_{j'}$ for all $j, j' \in [m]$.*

2.2 Classic Contracts

We first describe the classic setting, in which a contract, denoted by $\langle t, i \rangle$, is a payment function t and a recommended action $i \in [n]$. The interpretation is that the principal posts a contract, the agent observes the contract and chooses an action and bears the attendant cost, an outcome is drawn from the distribution over outcomes induced by that action, and the principal makes the payment to the agent specified by the contract. The inclusion of a recommended action in the contract allows us to capture the common presumption that the agent "breaks ties in favor of the principal".

More precisely, an agent who chooses action i' when facing a contract $\langle t, i \rangle$ garners expected utility

$$U_A(i' \mid t) \quad = \quad \sum_{j=1}^{m} p_{i'j} t_j - c_{i'} \quad = \quad T_{i'}(t) - c_{i'},$$

given by the difference between the expected payment $T_{i'}(t)$ and the cost $c_{i'}$. The resulting principal utility is $U_P(i' \mid t) = R_{i'} - T_{i'}(t)$.

Definition 4 (IC contract)**.** *A contract $\langle t, i \rangle$ is incentive compatible (IC) if*

$$i \in \operatorname*{argmax}_{i' \in [n]} U_A(i' \mid t),$$

in which case we say that contract $\langle t, i \rangle$ implements action i.

Because payments are non-negative and action 1 has zero cost, incentive compatibility ensures that the agent secures an expected utility of at least zero, and hence implies individual rationality.

We assume the agent follows the recommendation of an incentive-compatible contract. If the principal posts the incentive compatible contract $\langle t, i \rangle$, the payoffs to the principal and agent are then

$$\begin{aligned} U_P(\langle t, i \rangle) &= U_P(i \mid t) &= R_i - T_i(t) \\ U_A(\langle t, i \rangle) &= U_A(i \mid t) &= T_i(t) - c_i. \end{aligned}$$

It is without loss of generality to restrict the principal to incentive compatible contracts. The idea is that an agent facing contract $\langle t, i \rangle$ will choose an action that maximizes her expected utility given t, and hence the principal might as well name such an action in the contract. The optimal classic contract implementing each action i can be identified by solving a linear programming problem, presented in Figure 8 in Appendix A. The optimal (incentive compatible) classic contract will inevitably induce indifference on the part of the agent, as some incentive constraint will bind.

2.3 Ambiguous Contracts

An *ambiguous* contract $\langle \tau, i \rangle = \langle \{t^1, \ldots, t^k\}, i \rangle$ is a set of payment functions and a recommended action. If $t \in \{t^1, \ldots, t^k\}$, then we say that t is in the support of τ.

The principal now posts an ambiguous contract, the agent observes the ambiguous contract and chooses an action and bears the attendant cost, a payment function is selected from the support of the ambiguous contract, an outcome is drawn from the distribution over outcomes induced by that action, and the principal makes the payment to the agent specified by the selected contract.

The agent is a max-min expected utility maximizer (Schmeidler, 1989; Gilboa and Schmeidler, 1993), and so evaluates each action i according to the payment function that minimizes the expected payment of the action.

Formally, given an ambiguous contract $\langle \tau, i \rangle$, the agent's utility for action $i' \in [n]$ is

$$U_A(i' \mid \tau) = \min_{t \in \tau} U_A(i \mid t).$$

Definition 5 (IC ambiguous contract)**.** *An ambiguous contract* $\langle\tau, i\rangle$ *is incentive compatible (IC) if*

$$i \in \operatorname*{argmax}_{i' \in [n]} U_A(i' \mid \tau),$$

in which case we say that ambiguous contract $\langle\tau, i\rangle$ *implements action* i*.*

As in the case of a classic contract, the incentive compatibility constraint implies individual rationality. It is again without loss to restrict the principal to incentive compatible ambiguous contracts.

If the principal's expected utility $U_P(i \mid t)$ under payment scheme $t \in \tau$ is strictly higher than $U_P(i \mid t')$ for some $t' \in \tau$, then the principal's "threat" that she may use any contract in τ may not be credible—an agent facing ambiguous contract $\langle\tau, i\rangle$ may fear the principal will contrive to invariably select the payment function t rather than t'. As in Di Tillio et al. (2017), we accordingly restrict the principal to consistent ambiguous contracts.

Definition 6 (Consistency)**.** *An ambiguous contract* $\langle\tau, i\rangle = \langle\{t^1, \ldots, t^k\}, i\rangle$ *is* consistent *if for any* $\ell, \ell' \in [k]$,

$$U_P(i \mid t^{\ell}) = U_P(i \mid t^{\ell'}). \tag{1}$$

If the principal posts the consistent, incentive compatible ambiguous contract $\langle\tau, i\rangle = \langle\{t^1, \ldots, t^k\}, i\rangle$, then the induced payment $T_i(\tau)$ can be defined as

$$\begin{aligned} T_i(\tau) = T_i(t^{\ell}) = T_i(t^{\ell'}) \qquad & \forall \ell, \ell' \in [k], \quad \text{and hence} \\ U_A(i \mid t^{\ell}) = U_A(i \mid t^{\ell'}) \qquad & \forall \ell, \ell' \in [k], \end{aligned}$$

with the first directly implied by consistency and the second following from $U_A(i \mid t) = T_i(t) - c_i$. The expected utilities $U_P(\langle\tau, i\rangle)$ and $U_A(\langle\tau, i\rangle)$ of the principal and agent are then given by, for any $t \in \tau$,

$$\begin{aligned} U_P(\langle\tau, i\rangle) &= U_P(i \mid t) = R_i - T_i(\tau), \quad \text{and} \\ U_A(\langle\tau, i\rangle) &= U_A(i \mid t) = T_i(\tau) - c_i. \end{aligned}$$

It is without loss of generality to restrict the principal to consistent ambiguous contracts:

Lemma 1. *Suppose $\langle \tau, i \rangle$ is incentive compatible. Then there exists a consistent, incentive compatible ambiguous contract $\langle \tau', i \rangle$ from which the principal obtains expected payoff at least $\max_{t \in \tau} U_P(i \mid t)$.*

Proof. Consider an incentive compatible ambiguous contract $\langle \tau, i \rangle$ that is not consistent. Let the payment functions in τ be numbered so that

$$U_P(i \mid t^1) = \max_{t \in \tau} U_P(i \mid t).$$

Suppose

$$U_P(i \mid t^1) > U_P(i \mid t^2).$$

Then it must be that

$$\sum_{j=1}^{m} p_{ij} t_j^1 < \sum_{j=1}^{m} p_{ij} t_j^2.$$

Let $\theta \in [0, 1)$ satisfy

$$\sum_{j=1}^{m} p_{ij} \theta t_j^2 = \sum_{j=1}^{m} p_{ij} t_j^1,$$

and consider the ambiguous contract $\langle \tau', i \rangle = \langle \{t^1, \theta t^2, \ldots, t^k\}, i \rangle$, constructed from $\langle \tau, i \rangle$ by replacing payment function t^2 with θt^2. We have $U_P(i \mid t^1) = U_P(i \mid \theta t^2) > U_P(i \mid t^2)$, and so the principal's payoff is at least as high under $\langle \tau', i \rangle$ as under $\langle \tau, i \rangle$. In addition, we have

$$\begin{array}{rclcll} U_A(i \mid \theta t^2) & = & U_A(i \mid t^1) & < & U_A(i \mid t^2) & \text{and} \\ U_A(i' \mid \theta t^2) & = & \theta T_{i'}(t^2) - c_{i'} & \leq & U_A(i' \mid t^2), & \end{array}$$

which imply

$$\begin{array}{rcll} \min_{t \in \tau'} U_A(i \mid t) & = & \min_{t \in \tau} U_A(i \mid t) & \text{and} \\ \min_{t \in \tau'} U_A(i' \mid t) & \leq & \min_{t \in \tau} U_A(i' \mid t) & \forall i' \in [n], \end{array}$$

which establishes that $\langle \tau', i \rangle$ is incentive compatible. Applying a similar argument to payment functions $t^3, \ldots, t^k$ yields the result. □

An implication of this result is that we can equivalently view the principal as selecting the contract to be implemented, or as max-min preferences over the outcome of a selection by a third party.

2.4 Implementability

This section characterizes the actions that are implementable with classic and ambiguous contracts. The following result for classic contracts is standard (e.g., Hermalin and Katz (1991, Proposition 2)).

Proposition 1. *Action $i \in [n]$ is implementable with a classic contract if and only if there does not exist a convex combination $\lambda_{i'} \in [0,1]$ of the actions $i' \neq i$ that yields the same distribution over rewards $\sum_{i' \neq i} \lambda_{i'} p_{i'j} = p_{ij}$ for all j but at a strictly lower cost $\sum_{i'} \lambda_{i'} c_{i'} < c_i$.*

For completeness, we provide a proof in Appendix A. In contrast, the conditions for implementing an action with ambiguous contracts are more permissive:

Proposition 2. *Action $i \in [n]$ is implementable with an ambiguous contract if and only if there is no other action $i' \neq i$ such that $p_{i'} = p_i$ and $c_{i'} < c_i$.*

Proof. We first show that if there exists an action $i' \neq i$ such that $p_{i'} = p_i$ and $c_{i'} < c_i$, then it is impossible to implement action i with an ambiguous contract. For the sake of contradiction, suppose that ambiguous contract $\langle \tau, i \rangle = \langle \{t^1, \ldots, t^k\}, i \rangle$ implements action i. In this case, since $p_i = p_{i'}$, we have $T_i(t^\ell) = T_{i'}(t^{\ell'})$ for all $\ell, \ell' \in [k]$. But then $U_A(i' \mid \tau) = \min_{\ell \in [k]} T_{i'}(t^\ell) - c_{i'} > \min_{\ell \in [k]} T_i(t^\ell) - c_i = U_A(i \mid \tau)$, contradicting the fact that $\langle \tau, i \rangle$ is incentive compatible.

Next we show that if there is no action $i' \neq i$ such that $p_{i'} = p_i$ and $c_{i'} < c_i$, then action i can be implemented with an ambiguous contract. In this case, for each action $i' \neq i$, either (i) $p_{i'} \neq p_i$ or (ii) $p_{i'} = p_i$ and $c_{i'} \geq c_i$. Let A be the actions of type (i). If A is empty, then i must be a zero-cost action. A (consistent) ambiguous contract for implementing that action is $\langle \tau, i \rangle$ with $\tau = \{(0, \ldots, 0)\}$.

Assume A is nonempty. We construct an ambiguous contract $\langle \tau, i \rangle$ for implementing action i that has one contract $t^{i'}$ for each action $i' \neq i$ of type (i). For each action $i' \in A$, let $j(i')$ be an outcome j such that $p_{ij}/p_{i'j}$ is maximal. Note that $p_{ij(i')}/p_{i'j(i')} > 1$. Let

$$T = \max_{i' \in A} \left\{ \min \left\{ x \geq 0 \;\middle|\; p_{ij(i')} \cdot \frac{x}{p_{ij(i')}} - c_i \geq p_{i'j(i')} \cdot \frac{x}{p_{ij(i')}} - c_{i'} \right\} \right\}.$$

rewards:	$r_1 = 0$	$r_2 = 2$	$r_3 = 2$	costs
action 1:	1	0	0	$c_1 = 0$
action 2:	1/2	1/2	0	$c_2 = 1$
action 3:	1/2	0	1/2	$c_3 = 1$
action 4:	0	1/2	1/2	$c_4 = 3$

Figure 1: Instance (c, r, p) for Example 1.

For each $i' \in A$, let $t^{i'}_{j(i')} = T/p_{ij(i')}$ and $t^{i'}_{j'} = 0$ for $j' \neq j(i')$.

We conclude by verifying that $\langle \tau, i \rangle = \langle \{t^{i'} \mid i' \in A\}, i \rangle$ is a (consistent) ambiguous contract that implements action i. It is easy to check consistency.

To see that $\langle \tau, i \rangle$ is incentive compatible, first consider actions $i' \neq i$ of type (ii). For these actions we have

$$U_A(i' \mid \tau) = U_A(i \mid \tau) + c_i - c_{i'} \leq U_A(i \mid \tau),$$

where we used that $p_{i'} = p_i$ and $c_{i'} \geq c_i$.

Next consider actions $i' \neq i$ of type (i). For these actions, there must be a $T_{i'} \geq 0$ with $T_{i'} \leq T$ such that

$$p_{ij(i')} \cdot \frac{T_{i'}}{p_{ij(i')}} - c_i \geq p_{i'j(i')} \cdot \frac{T_{i'}}{p_{ij(i')}} - c_{i'}.$$

Since $T_{i'} \leq T$ and $p_{ij(i')} > p_{i'j(i')}$ this implies

$$\begin{aligned} U_A(i \mid \tau) = p_{ij(i')} \cdot \frac{T}{p_{ij(i')}} - c_i &\geq p_{i'j(i')} \cdot \frac{T}{p_{ij(i')}} - c_{i'} \\ &= \min_{i'' \in A} \left(p_{i'j(i'')} \cdot \frac{T}{p_{ij(i'')}} - c_{i'} \right) = U_A(i' \mid \tau), \end{aligned}$$

where the first equality holds by consistency, the second equality holds by definition of $j(i')$, and the final equality holds by definition. □

The following example presents an action that is implementable with an ambiguous contract but not a classic contract.

Example 1 (Action implementable with ambiguous but not classic contract)**.** Consider the instance in Figure 1. A half/half combination of actions 2 and 3 gives

rewards:	$r_1 = 0$	$r_2 = 2$	$r_3 = 2$	costs
action 1:	1	0	0	$c_1 = 0$
action 2:	1/2	1/2	0	$c_2 = 1/4$
action 3:	1/2	0	1/2	$c_3 = 1/4$
action 4:	0	1/2	1/2	$c_4 = 1$

Figure 2: Instance (c, r, p) for Example 2.

the same distribution over outcomes as action 4, but at a lower cost, and hence no classic contract can implement action 4. However, the ambiguous contract $\langle \tau, 4 \rangle = \langle \{t^1, t^2\}, 4 \rangle$ with $t^1 = (0, 6, 0)$ and $t^2 = (0, 0, 6)$ implements action 4 with the minimum possible expected payment of $T_4(\tau) = c_4 = 3$). ■

3 Optimal Ambiguous Contracts

We first show, in Section 3.1, that the principal can benefit by employing ambiguous rather than classic contracts. This observation motivates us to study the structure of optimal ambiguous contracts. Section 3.2 establishes that optimal ambiguous contracts taking a particularly simple form always exist. Section 3.3 establishes the counterparts of these results for monotone contracts.

3.1 The Advantage of Optimal Ambiguous Contracts

In the following simple variation of Example 1, the principal gains from ambiguity by implementing an action different from the one that she would implement with a classic contract.

Example 2 (Strict improvement)**.** Consider the instance shown in Figure 2. The best classic contract implements action 2 (or, by symmetry, action 3), and yields the principal an expected utility of 3/4. One way to achieve this maximum utility is via contract $\langle (0, 1/2, 0), 2 \rangle$. In contrast, the cheapest way to implement action 4 with a classic contract is via contract $\langle (0, 3/2, 3/2), 4 \rangle$. This gives the principal an expected utility of $1/2 < 3/4$.

An optimal ambiguous contract is $\langle \tau, 4 \rangle = \langle \{t^1, t^2\}, 4 \rangle$, with $t^1 = (0, 2, 0)$ and $t^2 = (0, 0, 2)$. The worst payment function in τ for action 2 is t^2, giving the agent an expected payment of 0. Similarly, the worst payment function for action 3 is t^1, for an

rewards:	$r_1 = 0$	$r_2 = 32$	$r_3 = 33$	$r_4 = 34$	costs
action 1:	1	0	0	0	$c_1 = 0$
action 2:	0.4	0.6	0	0	$c_2 = 1$
action 3:	0	0.5	0.5	0	$c_3 = 11$
action 4:	0	0	0.6	0.4	$c_4 = 12$

Figure 3: Instance (c, r, p) for Example 3.

expected payment of 0. Thus, both actions 2 and 3 give the agent negative utilities. In contrast, the expected payment for action 4 is 1 under both t^1 and t^2, giving the agent an expected utility of 0. The ambiguous contract $\langle \tau, 4 \rangle$ thus implements action 4, with an expected payment of 1, and an expected utility for the principal of 1 strictly higher than her optimal utility under a classic contract. ■

We next show that the same phenomenon can occur in instances satisfying the MLRP condition. In this example, the optimal classic contract and the optimal ambiguous contract implement the same action.

Example 3 (Strict improvement under MLRP)**.** Consider the instance shown in Figure 3. In this instance, the principal can implement both action 2 and action 4 with a classic contract, with an expected payment equal to the agent's respective cost. Possible contracts that achieve this include $\langle (0, 5/3, 0, 0), 2 \rangle$ and $\langle (0, 0, 0, 30), 4 \rangle$. The resulting principal utility is $R_2 - c_2 = 18.2$ for action 2 and $R_4 - c_4 = 21.4$ for action 4. The cheapest way to implement action 3 (e.g., by solving the LP in Figure 8 in Appendix A) is via contract $\langle t = (0, 25/12, 245/12, 0), 3 \rangle$, yielding a utility of $R_3 - T_3(t) = 21.25$. The maximal utility the principal can achieve with a classic contract is thus 21.4. Contrast this with the optimal ambiguous contract $\langle \{t^1, t^2\}, 3 \rangle$, consisting of the two payment functions $t^1 = (0, 22, 0, 0)$ and $t^2 = (0, 0, 22, 0)$. This ambiguous contract implements action 3 with an expected payment equal to the agent's cost, for a principal utility of $R_3 - c_3 = 21.5$. ■

In our next example, the principal gains from using an ambiguous contract to implement an action that *cannot* be implemented with a classic contract.

Example 4 (Action optimal with ambiguous contract but not implementable with classic contract)**.** Consider the instance shown in Figure 4. Action 6 cannot be implemented by a classic contract, with the half/half combination of actions 4 and 5 giving

rewards:	$r_1 = -200$	$r_2 = 0$	$r_3 = 21$	$r_4 = 21$	costs
action 1:	0	1	0	0	$c_1 = 0$
action 2:	0.1	0	0.9	0	$c_2 = 8$
action 3:	0.1	0	0	0.9	$c_3 = 8$
action 4:	0	0	1	0	$c_4 = 10$
action 5:	0	0	0	1	$c_5 = 10$
action 6:	0	0	0.5	0.5	$c_6 = 11$

Figure 4: Instance (c, r, p) for Example 4.

rewards:	$r_1 = 0$	$r_2 = 9$	$r_3 = 9$	costs
action 1:	1	0	0	$c_1 = 0$
action 2:	0.6	0.3	0.1	$c_2 = 0.6$
action 3:	0.6	0.1	0.3	$c_3 = 0.6$
action 4:	0.2	0.4	0.4	$c_4 = 3$

Figure 5: Instance (c, r, p) for Example 5.

the same distribution over outcomes at a lower cost. Actions 2 and 3 have a negative expected welfare, and so will never be optimal for the principal. Actions 4 and 5 can both be implemented with a classic contract, and yield the same maximal utility for the principal. Optimal classic contracts for these actions include $\langle(0, 0, 20, 0), 4\rangle$ and $\langle(0, 0, 0, 20), 5\rangle$, each giving the principal an expected utility of 1. In contrast, the optimal ambiguous contract $\langle\{(0, 0, 22, 0), (0, 0, 0, 22)\}, 6\rangle$ implements action 6, for an expected utility of 10. ■

We conclude with an example showing that an ambiguous contract may benefit *both* the principal and the agent. Clearly, this can only happen when the optimal action under an ambiguous contract differs from the optimal action under classic contracts.

Example 5 (Ambiguous contracts may benefit both principal and agent)**.** Consider the instance shown in Figure 5. An optimal classic contract is $\langle(0, 2, 0), 2\rangle$, implementing action 2 with utilities 0 and 3 to the agent and principal. The ambiguous contract $\langle\{(0, 8, 0), (0, 0, 8)\}, 4\rangle$ implements action 4 with utilities 0.2 and 4 to the agent and principal. ■

3.2 The Structure of Optimal Ambiguous Contracts

We now investigate the structure of optimal ambiguous contracts. We first introduce the simplicity notion of a *single-outcome payment* contract.

Definition 7 (SOP payment function). *A payment function $t = (t_1, t_2, \ldots, t_m)$ is a* single-outcome payment function *if there exists an outcome $j \in [m]$ such that $t_j > 0$, and for any outcome $j' \neq j$, $t_{j'} = 0$.*

An ambiguous contract is a single-outcome payment *contract* if all of its payment functions have this property. Proposition 2 used SOP contracts to establish the sufficiency of conditions for implementation under ambiguous contracts.

The following theorem shows that it is without loss of generality to consider ambiguous SOP contracts, and that at most $\min\{m, n-1\}$ payment functions are needed. In Proposition 5 in Appendix B we show that this bound is tight.

Theorem 1 (Optimal ambiguous contracts). *For every IC ambiguous contract $\langle \tau, i \rangle$, there exists an IC ambiguous contract $\langle \tau', i \rangle$, containing at most $\min\{m, n-1\}$ payment functions, such that:*

1. *$T_i(\tau') = T_i(\tau)$ (i.e., both contracts have the same expected payment and hence same expected payoff to the principal).*

2. *For every $t' \in \tau'$, t' is an SOP payment function.*

Proof. Let the ambiguous contact τ implement action i. Let $J = \{j \in [m] \mid p_{ij} > 0\}$. For every $j \in J$, consider the SOP payment function with payment $\frac{T_i(\tau)}{p_{ij}}$ for outcome j. Let τ' be the ambiguous contract consisting of these SOP payment functions. By construction, τ' satisfies properties (1)–(2). We show that τ' implements i. Consider an action $i' \neq i$. Because τ implements i, there exists $t \in \tau$ with

$$c_i - c_{i'} \;\leq\; T_i(\tau) - \sum_j t_j p_{i'j} \;=\; \sum_j t_j p_{ij} - \sum_j t_j p_{i'j}.$$

To show that τ' implements i, it suffices to show

$$c_i - c_{i'} \;\leq\; T_i(\tau) - \min_{j \in J} p_{i'j} \frac{T_i(\tau)}{p_{ij}} \;=\; \sum_j t_j p_{ij} - \min_{j \in J} \frac{p_{i'j}}{p_{ij}} \sum_j t_j p_{ij}.$$

Combining these, it suffices to show

$$\min_{j \in J} \frac{p_{i'j}}{p_{ij}} \sum_j t_j p_{ij} \;\le\; \sum_j t_j p_{i'j},$$

which is equivalent to the obvious statement that

$$\min_{j \in J} \frac{p_{i'j}}{p_{ij}} \;\le\; \frac{\sum_j t_j p_{i'j}}{\sum_j t_j p_{ij}}.$$

Notice that τ' consists of at most m SOP payment functions (in fact, at most $|J|$ payment functions). If $m > n-1$, one can eliminate from τ' every payment function that does not minimize the expected payoff to one of the alternatives $i' \neq i$, leaving at most $n-1$ SOP payment functions. □

An implication of this result is that if the agent has only two feasible actions, then there exists an optimal ambiguous contract with at most $n-1 = 1$ payment functions, which is a classic contract. Hence, with only two actions, ambiguous contracts cannot improve on classic contracts for the principal.

With the help of Theorem 1, we can show that the contract $\langle \tau, i \rangle$ that we constructed to establish the sufficiency of the conditions for implementation (in the proof of Proposition 2) is optimal.

Proposition 3. *Suppose action $i \in [n]$ is implementable by an ambiguous contract. Let $A = \{i' \neq i \mid p_{i'} \neq p_i\}$. If $A = \emptyset$, then $c_i = 0$, and the IC contract $\langle \{(0, \ldots, 0)\}, i \rangle$ is optimal for action i. Otherwise, for each $i' \in A$ let $j(i')$ be an outcome such that $p_{ij(i')}/p_{i'j(i')}$ is maximal. Let*

$$T = \max_{i' \in A} \left\{ \min \left\{ x \geq 0 \;\middle|\; p_{ij(i')} \cdot \frac{x}{p_{ij(i')}} - c_i \geq p_{i'j(i')} \cdot \frac{x}{p_{ij(i')}} - c_{i'} \right\} \right\}.$$

For each $i' \in A$, let $t^{i'}_{j(i')} = T/p_{ij(i')}$ and $t^{i'}_{j'} = 0$ for $j' \neq j(i')$. Then the IC contract $\langle \tau, i \rangle = \langle \{t^{i'} \mid i' \in A\}, i \rangle$ is optimal for action i.

The proof of Proposition 3 in Appendix C relies on arguments similar to those used to establish that in classic contracts, with two actions, it is optimal to pay only for the maximum likelihood-ratio outcome (e.g., Laffont and Martimort (2009, Chapter 4.5.1), Dütting et al. (2019, Full version, Proposition 5)).

We next show that optimal ambiguous contracts for instances satisfying the MLRP condition admit an even simpler structure, namely an ambiguous contract composed of only two SOP payment functions.

Theorem 2 (Optimal ambiguous contracts under MLRP)**.** *Let (c, r, p) be an instance that satisfies the MLRP condition. For every IC ambiguous contract $\langle \tau, i \rangle$, there exists an IC ambiguous contract $\langle \tau', i \rangle = \langle \{t^1, t^k\}, i \rangle$, such that:*

1. *$T_i(\tau') = T_i(\tau)$ (i.e., both contracts have the same expected payment and hence the same expected payoff to the principal).*
2. *t^1 and t^k are SOP payment functions, where $t^1_\ell = \frac{T_i(\tau)}{p_{i\ell}}$ for $\ell = \min\{j \in [m] \mid p_{ij} > 0\}$, and $t^k_h = \frac{T_i(\tau)}{p_{ih}}$ for $h = \max\{j \in [m] \mid p_{ij} > 0\}$.*

The proof of Theorem 2, which is deferred to Appendix D, combines the structural properties established in Theorem 1 with the MLRP condition to argue that two SOP payment functions, introduced in the proof of Theorem 1, suffice. Payment function t^k "defeats" all actions with smaller costs and payment function t^1 "defeats" all actions with higher costs.

The ability to restrict attention to SOP contracts allows us to show that optimal ambiguous contracts are relatively easy to identify. Appendix H shows that there exists an algorithm capable of computing the optimal ambiguous contract in time $O(n^2m)$ (and time $O(n^2 + m)$ under the MLRP condition). The key implications of these results is that computation time increases polynomially rather than exponentially in the size of the instance.

3.3 Optimal Monotone Ambiguous Contracts

In some scenarios it is desired or even required, for reasons of fairness, robustness, or regulation, to restrict attention to monotone payment functions. A contract whose payment functions are monotone is a monotone contract. The monotonicity requirement rules out SOP contracts. The following is a natural alternative simplicity notion for monotone contracts.

Definition 8 (Step payment function)**.** *A payment function $t = (t_1, t_2, \ldots, t_m)$ is a* step *payment function if there exists an outcome $k \in [m]$ and some $x \geq 0$, such that $t_j = 0$ for every outcome $j < k$, and $t_j = x$ for every outcome $j \geq k$.*

A contract composed of step payment functions is a step contract.

The following theorem shows that it is without loss of generality to consider ambiguous monotone contracts that are composed of step payment functions. In that sense, step contracts are the analogue of SOP contracts for monotone contracts.

Theorem 3 (Optimal monotone ambiguous contracts)**.** *For every IC monotone ambiguous contract $\langle \tau, i \rangle$, there exists an IC monotone ambiguous contract $\langle \tau', i \rangle$ consisting of at most* $\min\{m, n-1\}$ *contracts, such that:*

1. *$T_i(\tau') = T_i(\tau)$ (i.e., both contracts have the same expected payment and hence the same payoff to the principal).*

2. *Every payment function in τ' is a step payment function.*

Clearly, the theorem holds if we replace "monotone ambiguous contract" with "step ambiguous contract."

Proof of Theorem 3. We construct τ' as follows. For every action $i' \neq i$ there must be a monotone payment function $t^{i'} \in \tau$, such that $U_A(i' \mid t^{i'}) \leq U_A(i \mid t^{i'}) = U_A(i \mid \tau)$. By the same arguments as in the proof of Theorem 1 it is now sufficient to show that there is a step payment function $\hat{t}^{i'}$ to put into τ' such that (i) $T_i(\hat{t}^{i'}) = T_i(t^{i'})$ and (ii) $T_{i'}(\hat{t}^{i'}) \leq T_{i'}(t^{i'})$.

For actions $i' \neq i$ such that there exists an outcome $\hat{j} \in [m]$ for which $\sum_{\ell=\hat{j}}^{m} p_{i\ell} > 0$ and $\sum_{\ell=\hat{j}}^{m} p_{i'\ell} = 0$, set $\hat{t}^{i'}_j = \frac{T_i(t^{i'})}{\sum_{\ell=\hat{j}}^{m} p_{i\ell}}$ for all $j \geq \hat{j}$ and $\hat{t}^{i'}_j = 0$ for all $j < \hat{j}$. To see that Condition (i) is satisfied, observe that

$$T_i(\hat{t}^{i'}) = \sum_{j=1}^{m} p_{ij} \cdot \hat{t}^{i'}_j = \sum_{j=\hat{j}}^{m} p_{ij} \cdot \frac{T_i(t^{i'})}{\sum_{\ell=\hat{j}}^{m} p_{i\ell}} = T_i(t^{i'}) \cdot \frac{\sum_{j=\hat{j}}^{m} p_{ij}}{\sum_{\ell=\hat{j}}^{m} p_{i\ell}} = T_i(t^{i'}).$$

Condition (ii) is satisfied, because $\sum_{\ell=\hat{j}}^{m} p_{i'\ell} = 0$ and $\hat{t}_j = 0$ for $j < \hat{j}$ imply that $T_{i'}(\hat{t}^{i'}) = 0$, while $T_{i'}(t^{i'}) \geq 0$ by limited liability.

Consider next actions $i' \neq i$ such that for all $\hat{j} \in [m]$ where $\sum_{\ell=\hat{j}}^{m} p_{i\ell} > 0$ it holds that $\sum_{\ell=\hat{j}}^{m} p_{i'\ell} > 0$. Let

$$\hat{j} \in \operatorname*{argmax}_{j' \in [m] : \sum_{\ell=j'}^{m} p_{i\ell} > 0} \frac{\sum_{\ell=j'}^{m} p_{i\ell}}{\sum_{\ell=j'}^{m} p_{i'\ell}}.$$

Define $\hat{t}^i$ as follows: Let $\hat{t}^{i'}_j = \frac{T_i(t^{i'})}{\sum_{\ell=\hat{j}}^m p_{i\ell}}$ for all $j \geq \hat{j}$, and $\hat{t}^{i'}_j = 0$ for all $j < \hat{j}$. It is again easy to verify that Condition (i) holds. Condition (ii) follows from noting that

$$\begin{aligned} T_{i'}(\hat{t}^{i'}) &= \sum_{j=\hat{j}}^{m} p_{i'j} \frac{T_i(t^{i'})}{\sum_{\ell=\hat{j}}^m p_{i\ell}} = \frac{\sum_{j=\hat{j}}^m p_{i'j}}{\sum_{\ell=\hat{j}}^m p_{i\ell}} \sum_{j=1}^{m} p_{ij} t'{i}_j \\ &= \min_{j' \in [m]: \sum_{\ell=j'}^m p_{i\ell} > 0} \frac{\sum_{j=j'}^m p_{'j}}{\sum_{j=j'}^m p_{i\ell}} \sum_{j=1}^{m} p_{i,j} t^{i'}_j \leq \sum_{j=1}^{m} p_{ij} t^{i'}_j = T_{i'}(t^{i'}), \end{aligned}$$

where the inequality follows by observing that one possible choice for j' is the smallest index such that $p_{ij'} > 0$.

We still have to show that $|\tau'| \leq \min\{m, n-1\}$. Clearly, $|\tau'| \leq n-1$ by the construction of τ' (we add one payment function for every $i' \neq i$). The fact that $|\tau'| \leq m$ follows by the consistency of τ', combined with the fact that any payment function $t \in \tau'$ is a step payment function, i.e., $T_i(\tau)$ combined with the outcome in which the step of a payment function occurs, uniquely defines the payment function. $\square$

We next show that optimal monotone ambiguous contracts for instances satisfying the MLRP condition admit an even simpler structure. Namely, similar to the unrestricted case where we did not impose monotonicity, the MLRP condition implies that an optimal ambiguous contract consists of only two payment functions.

Theorem 4 (Optimal monotone ambiguous contracts under MLRP)**.** *Let (c, r, p) be an instance that satisfies the MLRP condition. For every IC monotone ambiguous contract $\langle \tau, i \rangle$, there exists an IC monotone ambiguous contract $\langle \tau', i \rangle = \langle \{t^1, t^k\}, i \rangle$ such that:*

1. *$T_i(\tau') = T_i(\tau)$ (i.e., both contracts have the same expected payment and hence the same payoff to the principal).*

2. *t^1 and t^k are step payment functions: $t^1_j = T_i(\tau)$ for all $j \geq \ell$, where $\ell = \min\{j \in [m] \mid p_{ij} > 0\}$, and $t^k_j = \frac{T_i(\tau)}{p_{ih}}$ for all $j \geq h$, where $h = \max\{j \in [m] \mid p_{ij} > 0\}$.*

The proof of Theorem 4 appears in Appendix E, and proceeds by verifying that the contract stated in the second bullet satisfies the condition in the first bullet and implements action i. As in the case of Theorem 2, t^k protects against all actions with lower cost and t^1 protects against all actions with higher cost.

Example 6 (Strict improvement under MLRP with monotone contracts)**.** We return to the instance shown in Figure 3, but now require contracts to be monotone. The payment functions for the optimal classic monotone contracts implementing each of the various actions, and the attendant payoffs for the principal, are:

$$\begin{array}{lrr} \text{action } 1: & (0,0,0,0) & 0 \\ \text{action } 2: & (0,5/3,5/3,5/3) & 18.2 \\ \text{action } 3: & (0,25,25,25) & 7.5 \\ \text{action } 4: & (0,0,0,30) & 21.4. \end{array}$$

An ambiguous contract allows us to improve on the cost of implementing (only) action 3. An optimal ambiguous contract implements action 3 with the two step payment functions

$$\begin{aligned} t^1 &= (0,11,11,11), \qquad \text{and} \\ t^k &= (0,0,22,22). \end{aligned}$$

for a payoff to the principal of 21.5. ■

4 The Ambiguity Gap

Section 3.1 confirmed that it can be advantageous for the principal to offer ambiguous contracts. To quantify the extent of the potential gains, we introduce the notion of the *ambiguity gap*, defined as the worst-case ratio between the principal's utility with and without ambiguity.

We restrict attention throughout this section to instances for which the optimal classic contract induces a non-negative utility for the principal. Let $\mathcal{C}(c,r,p)$ and $\mathcal{A}(c,r,p)$ be the sets of incentive compatible classic contracts and incentive compatible ambiguous contracts, for an instance (c,r,p).

Definition 9 (Ambiguity gap)**.** *The ambiguity gap $\rho(c,r,p)$ of a given instance (c,r,p) and the ambiguity gap $\rho(\mathcal{I})$ of a class of instances $\mathcal{I}$, are*

$$\rho(c,r,p) = \frac{\max_{\langle \tau,i\rangle \in \mathcal{A}(c,r,p)} U_P(\langle \tau,i\rangle)}{\max_{\langle t,i\rangle \in \mathcal{C}(c,r,p)} U_P(\langle t,i\rangle)} \qquad \textit{and} \qquad \rho(\mathcal{I}) = \sup_{(c,r,p)\in\mathcal{I}} \rho(c,r,p).$$

rewards:	$r_1 = -r$	$r_2 = -r$	$r_3 = 0$	$r_4 = r$	costs
action 1:	0	0	1	0	$c_1 = 0$
action 2:	0.5	0	0	0.5	$c_2 = 10$
action 3:	0	0.5	0	0.5	$c_3 = 10$
action 4:	0.2	0.2	0	0.6	$c_4 = 20$

Figure 6: Instance (c, r, p) for Example 7

4.1 Unbounded Ambiguity Gap in General

The following example shows that the ambiguity gap can be arbitrarily large.

Example 7 (Unbounded gap with negative rewards)**.** Consider the instance shown in Figure 6, where $r \geq 0$ is a parameter we will allow to vary. Actions 2 and 3 generate negative welfare, and hence only action 4 is capable (depending on r) of producing positive welfare. Welfare is given by

$$\max\{0, 0.2r - 20\},$$

and is positive if and only if $r > 100$. An optimal classic contract implementing action 4 is $\langle t, 4\rangle = \langle(0, 0, 0, 100), 4\rangle$, which gives

$$U_P(\langle t, 4\rangle) = 0.2r - 60,$$

which is positive if and only if $r > 300$. An optimal ambiguous contract implementing action 4 is $\langle \tau, 4\rangle = \langle\{t^1, t^2\}, 4\rangle = \langle\{(100, 0, 0, 0), (0, 100, 0, 0)\}, 4\rangle$, giving

$$U_P(\langle \tau, 4\rangle) = 0.2r - 20,$$

which is positive if and only if $r > 100$. Hence for $r \in (100, 300]$, the best classic contract generates a payoff of 0, while the best ambiguous contract generates a positive payoff, yielding an infinite ambiguity gap. ■

4.2 Tight Ambiguity Gap under Non-Negative Rewards

In contrast to the unbounded ambiguity gap in general instances, we next show that for instances in which all rewards are non-negative the ambiguity gap is at most $n-1$

rewards:	$r_1 = 0$	$r_2 = 0$	$r_3 = \frac{1}{\gamma^{n-2}}$	costs
action 1:	1	0	0	$c_1 = 0$
action $2 \leq i \leq n-1$:	0	$1 - \gamma^{n-i}$	γ^{n-i}	$c_i = \frac{1}{\gamma^{i-2}} - (i-1) + (i-2)\gamma$
action n:	δ	0	$1 - \delta$	$c_n = \frac{1}{\gamma^{n-2}} - (n-1) + (n-2)\gamma$

Figure 7: Instance (c, r, p) used in the proof of Proposition 4.

and this is tight.

Proposition 4. *Fix $n \geq 2$. Let $\mathcal{I}_n^+$ denote the class of all instances with n actions and non-negative rewards. The ambiguity gap of $\mathcal{I}_n^+$ is*

$$\rho(\mathcal{I}_n^+) = n - 1.$$

The upper-bound direction of the argument makes use of the notion of a linear contract, which is used to state the following lemma.

Definition 10 (Linear contract)**.** *Consider an instance $(c, r, p) \in \mathcal{I}_n^+$. A (classic) contract $\langle t, i \rangle$ is* linear *if $t = (\alpha r_1, \ldots, \alpha r_m)$ for some $\alpha \geq 0$.*

For an instance $(c, r, p) \in \mathcal{I}_n^+$, denote by $\mathcal{L}(c, r, p)$ the set of all incentive compatible linear contracts. Let $W = \max_{i \in [n]} W_i$ denote the maximum welfare.

The following lemma shows that in instances in which the status-quo outcome has a reward of zero, the principal can achieve a $1/(n-1)$ fraction of the optimal welfare as utility with a linear contract.

Lemma 2. *Consider instance $(c, r, p) \in \mathcal{I}_n^+$ in which action 1 has a cost of $c_1 = 0$ and invariably leads to reward $r_1 = 0$. Then there exists a subset of actions $A \subseteq [n]$ with $|A| \leq n - 1$ and a scalar $\alpha_i \geq 0$ for each $i \in A$ such that each linear contract $\langle t, i \rangle$ with $t = (\alpha_i r_1, \ldots, \alpha_i r_m)$ is IC and*

$$\max_{\langle t,i \rangle \in \mathcal{L}(c,r,p)} U_P(\langle t, i \rangle) = \max_{i \in A} \, (1 - \alpha_i) R_i \geq \frac{W}{n-1}. \tag{2}$$

The proof of Lemma 2 is similar to arguments in (Dütting et al., 2019). For completeness, we provide a proof of this lemma in Appendix F. We are now ready to prove Proposition 4.

Proof of Proposition 4. We first show the upper bound on the ambiguity gap. To this end, fix any $n \geq 2$ and any instance $(c, r, p) \in \mathcal{I}_n^+$. Clearly, the maximum utility the principal can achieve with an ambiguous contract satisfies $\max_{\langle \tau, i \rangle \in \mathcal{A}(c,r,p)} U_P(\langle \tau, i \rangle) \leq W$. Thus, in order to prove the upper bound on the ambiguity gap, it suffices to show that the maximum utility the principal can achieve with a classic contract satisfies $\max_{\langle t,i \rangle \in \mathcal{C}(c,r,p)} U_P(\langle t, i \rangle) \geq W/(n-1)$. Consider using a contract of the form $\langle t, i \rangle$ with $t = (\alpha(r_1 - r_1), \ldots, \alpha(r_m - r_1))$, and let (c, r', p) be a modified instance in which $r'_j = r_j - r_1$ for all $j \in [m]$. Note that from the agent's perspective applying contract $\langle t, i \rangle$ in the original instance, is equivalent to applying contract $\langle t', i \rangle$ with $t' = (\alpha r'_1, \ldots, r'_m)$ in the modified instance.

Let R'_i and W'_i for $i \in [n]$ denote the expected reward and welfare of action i in the modified instance, and let $W' = \max_{i \in [n]} W'_i$. Note that $R_i = r_1 + R'_i$ and $W_i = r_1 + W'_i$ for all $i \in [n]$, and hence also $W = r_1 + W'$. Applying Lemma 2 to the modified instance, we know that there exists a set $A \subseteq [n]$ and scalars α_i for $i \in A$ such that

$$
\begin{aligned}
\max_{\langle t,i \rangle \in \mathcal{C}(c,r,p)} U_P(\langle t, i \rangle) &\geq \max_{i \in A} \big(r_1 + (1 - \alpha_i) R'_i \big) \\
&= r_1 + \max_{i \in A} \big((1 - \alpha_i) R'_i \big) \\
&\geq r_1 + \frac{1}{n-1} W' \\
&\geq \frac{1}{n-1} \big(r_1 + W' \big) = \frac{1}{n-1} W,
\end{aligned}
$$

where in the last step we used that $r_1 \geq 0$. This completes the proof of the upper bound on the ambiguity gap.

We next show the lower bound on the ambiguity gap. To this end, we vary a lower bound construction due to Dütting et al. (2021). For $n = 2$ there is nothing to show, so fix any $n \geq 3$. Let $\gamma, \epsilon \in (0, 1)$ and let $\delta = \epsilon \cdot \gamma^{n-2}$. Consider the parameterized instance (c, r, p) with n actions depicted in Figure 7. Lemma 3 in Appendix G shows that the maximal utility the principal can achieve in this instance with a classic contract is at most 1. The argument proceeds by showing an upper bound of 1 for each action $i \in [n]$. For actions $i = 1, 2$ the upper bound is immediate, as the welfare of these actions is $W_1 \leq W_2 \leq 1$. For actions $i \in \{3, \ldots, n\}$ the upper bound can be shown by considering only a subset of the IC constraints.

The proof is completed by observing that with an ambiguous contract the principal can implement action n, with an expected payment equal to c_n. This is enough to show the claim, as the welfare from that action is $W_n = (n-1) - (n-2)\gamma - \epsilon$, and $W_n \to n-1$ as $\gamma, \epsilon \to 0$. The ambiguous contract that achieves this is $\langle \{t^1, t^2\}, n \rangle$ with $t^1 = (\frac{c_n}{\delta}, 0, 0)$ and $t^2 = (0, 0, \frac{c_n}{1-\delta})$. It is easy to verify that this contract is consistent, and entails an expected payment of c_n for action n. It is IC, because for all actions $i \neq n$ it gives a minimum payment of zero. □

5 Ambiguity Proofness

In this section, we explore which classes of contracts are amenable to improvements via ambiguous contracts. We phrase our results in terms of properties of payment functions.

It simplifies the exposition to restrict attention to payment functions with the property that two outcomes that induce the same reward also induce the same payment. We can thus think of a payment function as mapping from $\mathbb{R}$ into $\mathbb{R}_+$. Within this setting, a *class of payment functions* $\mathcal{T}$ is a set (possibly of infinite size) of payment functions $t : \mathbb{R} \to \mathbb{R}_+$.

We first give the definition of an ambiguity-proof class of payment functions. For a given instance (c, r, p), let $\mathcal{C}_\mathcal{T}(c, r, p)$ denote the set of all incentive compatible classic contracts with payment functions from $\mathcal{T}$. Let $\mathcal{A}_\mathcal{T}(c, r, p)$ be the analogous definition for ambiguous contracts.

Definition 11 (Ambiguity-proof)**.** *A class of payment functions $\mathcal{T}$ is* ambiguity-proof *if for any instance (c, r, p) and any action $i \in [n]$ it holds that*

$$\max_{\langle \tau, i \rangle \in \mathcal{A}_\mathcal{T}(c,r,p)} U_P(\langle \tau, i \rangle) \leq \max_{\langle t, i \rangle \in \mathcal{C}_\mathcal{T}(c,r,p)} U_P(\langle t, i \rangle),$$

i.e., the principal cannot gain from implementing any action i with an ambiguous rather than a classic contract.

For example, the principal-agent setting in Example 2 shows that the contract class of all contracts is not ambiguity-proof.

Our condition for ambiguity-proofness will be the following:

Definition 12 (Ordered class of payment functions). *A class of payment functions* $\mathcal{T}$ *is* ordered *if for any two payment functions* $t, t' \in \mathcal{T}$ *it holds that:*

$$t(x) \geq t'(x) \quad \textit{for all } x \in \mathbb{R} \quad \textit{or} \quad t(x) \leq t'(x) \quad \textit{for all } x \in \mathbb{R}.$$

Theorem 5 (Ambiguity-proofness characterization). *A class of payment functions* $\mathcal{T}$ *is ambiguity-proof if and only if it is ordered.*

Proof. We first show that an ordered class of payment functions is ambiguity-proof. Suppose $\mathcal{T}$ is ordered. Consider a (consistent) incentive compatible ambiguous contract $\langle \tau, i \rangle = \langle \{t^1, \ldots, t^k\}, i \rangle \in \mathcal{A}_{\mathcal{T}}(c, r, p)$ with $k \geq 2$. Since $\mathcal{T}$ is ordered, there must exist a payment function in τ, say t^1, with the property that $t^1_j \leq t_j$ for all $j \in [m]$ and all $t \in \tau$. Hence, for every action $i' \in [n]$, we have

$$U_A(i' \mid t^1) = \sum_{j=1}^{m} p_{i'j} t^1_j - c_{i'} = \min_{t \in \tau} \sum_{j=1}^{m} p_{i'j} t_j - c_{i'} = \min_{t \in \tau} U_A(i' \mid t) = U_A(i' \mid \tau),$$

which implies that the classic contract $\langle t^1, i \rangle \in \mathcal{C}_{\mathcal{T}}(c, r, p)$ is incentive compatible. Moreover, by consistency of $\langle \tau, i \rangle$, it also holds that

$$U_P(\langle t^1, i \rangle) = U_P(i \mid t^1) = U_P(i \mid \tau) = U_P(\langle \tau, i \rangle).$$

Hence, the classic contract $\langle t^1, i \rangle$ implements action i at the same cost as does $\langle \tau, i \rangle$, and so $\mathcal{T}$ is ambiguity-proof.

We next show that ambiguity-proofness implies ordering, by proving the contrapositive. Suppose $\mathcal{T}$ violates ordering. Then there exist $t, t' \in \mathcal{T}$ and $x_1, x_2 \in \mathbb{R}$ such that $t(x_1) > t'(x_1)$ and $t(x_2) < t'(x_2)$. Letting $\delta_1 = t(x_1) - t'(x_1) > 0$, $\delta_2 = t(x_2) - t'(x_2) < 0$, $q_1 = \frac{-\delta_2}{\delta_1 - \delta_2}$ and $q_2 = \frac{\delta_1}{\delta_1 - \delta_2}$, we obtain values $q_1, q_2 > 0$ with $q_1 + q_2 = 1$ satisfying

$$q_1 t(x_1) + q_2 t(x_2) = q_1 t'(x_1) + q_2 t'(x_2). \tag{3}$$

Let $\kappa = \min_{j=1,2} \min\{t(x_j), t'(x_j)\}$. Note that $\kappa \geq 0$ by limited liability. Now consider the following instance with 2 outcomes $r_1 = x_1$ and $r_2 = x_2$ and 3 actions as follows:

- Action $i \in \{1, 2\}$: $p_{ii} = 1$ and $c_i = \min\{t(r_i), t'(r_i)\} - \kappa$.
- Action 3: for $j \in \{1, 2\}$, $p_{3j} = q_j$, and $c_3 = \sum_{j=1}^{2} q_j t(r_j) - \kappa$.

Note that this construction ensures that $c_i \geq 0$ for all $i \in \{1, 2, 3\}$, and that $c_i = 0$ for some $i \in \{1, 2\}$, which we can take to be the default action.

We argue that action 3 can be implemented by an ambiguous contract and cannot be implemented by any classic contract. We first show that action 3 can be implemented with the ambiguous contract $\langle \tau, 3 \rangle$ where $\tau = \{t, t'\}$, with an expected payment equal to $c_3 + \kappa$. To see that $\langle \tau, 3 \rangle$ is consistent, note that

$$T_3(t) = \sum_{j=1}^{2} q_j \cdot t(r_j) \overset{(3)}{=} \sum_{j=1}^{2} q_j \cdot t'(r_j) = T_3(t').$$

Since t and t' have the same expected payment for action 3, the agent's expected utility for taking action 3 under t and t' is the same, and equals

$$U_A(3 \mid t) = U_A(3 \mid t') = \sum_{j=1}^{2} q_j \cdot t(r_j) - c_3 = \kappa.$$

It remains to show that for any action $i \in \{1, 2\}$, the agent's utility under $\langle \tau, 3 \rangle$ is at most κ. If $t(r_i) < t'(r_i)$, then $c_i = t(r_i) - \kappa$, and the agent's expected utility for action i is $U_A(i \mid \tau) = U_A(i \mid t) = t(r_i) - c_i = \kappa$. Similarly, if $t'(r_i) \leq t(r_i)$, then $c_i = t'(r_i) - \kappa$, and the agent's utility for action i is $U_A(i \mid \tau) = U_A(i \mid t') = t'(r_i) - c_i = \kappa$. So in either case, the agent's utility from the ambiguous contract is at most κ.

We complete the argument by showing that action 3 cannot be implemented by any classic contract (even if we don't restrict the classic contract to come from class $\mathcal{T}$). We have

$$\sum_{j=1}^{2} q_j c_j = \left(\sum_{j=1}^{2} q_j \cdot \min(t(r_j), t'(r_j)) \right) - \kappa < \left(\sum_{j=1}^{2} q_j \cdot t(r_j) \right) - \kappa = c_3.$$

The convex combination of actions $1, 2$ via vector (q_1, q_2) thus yields the same distribution over rewards as action 3, but at a strictly lower cost. By Proposition 1, this means that action 3 is not implementable by any classic contract. □

As an immediate corollary of our characterization we obtain that for any fixed $d \in \mathbb{R}_+$ and $\beta \leq r_1$, the class of payment functions $\mathcal{T}_d(\beta) = \{t(x) = \alpha \cdot (x - \beta)^d \mid \alpha \geq 0\}$ is ambiguity-proof. If all rewards are non-negative, then we can set $\beta = 0$ to see that the class of linear contracts $\mathcal{T}_1(0)$ is ambiguity-proof. For general rewards, if $\beta < 0$, then $\mathcal{T}_1(\beta)$ describes an ambiguity-proof class of affine contracts. This is cast in the following corollary.

Corollary 1. *For any fixed $d \in \mathbb{R}_+$, and $\beta \leq r_1$, the class of payment functions $\mathcal{T}_d(\beta) = \{t(x) = \alpha \cdot (x - \beta)^d \mid \alpha \geq 0\}$ is ambiguity-proof. In particular, when all rewards are positive, the class of linear payment functions (corresponding to $\mathcal{T}_1(0)$) is ambiguity-proof.*

On the other hand, our characterization implies that many natural classes of contracts, such as the class of all affine contracts, all polynomial contracts, or the class of all monotone contracts, fail to be ambiguity-proof.

6 Mixing Hedges Against Ambiguity

In this section we explore the power of ambiguity when the agent is allowed to select a mixed action and the principal is allowed to implementing mixed actions. Our main result (Theorem 6) is that in this case, the principal cannot gain from using an ambiguous contract. Bade (2023) obtains a similarly-spirited result in a mechanism design context, showing that if agents are dynamically consistent, meaning that they update their beliefs in response to information so as to make it optimal to continue with their ex-ante optimal plan of action, then ambiguity does not expand the set of implementable social choice functions. In contrast, Kambhampati (2023) shows that a principal who entertains ambiguous beliefs about the actions available to an agent can typically improve her payoff by offering a random contract to the agent.

6.1 Extension to Mixed Actions

The definition of a payment function remains the same, $t = (t_1, \ldots, t_m) \in \mathbb{R}^m_+$ defines a non-negative transfer t_j for each outcome $j \in [m]$. A mixed action of the agent, denoted by $\psi \in \Delta^n$, is a convex combination over actions $i \in [n]$, so that ψ_i denotes the probability with which the agent chooses action i. A pure action is the special case of a mixed action in which $\psi_i = 1$ for some i and $\psi_{i'} = 0$ for all other $i' \neq i$. The expected reward of mixed action ψ, is $R_\psi = \sum_{i=1}^n \psi_i R_i$. We write $T_\psi(t) = \sum_{i=1}^n \psi_i T_i(t)$ for the expected payment for mixed action ψ under payment function t.

The agent's expected utility for mixed action ψ' under classic contract $\langle t, \psi \rangle$ is $U_A(\psi' \mid t) = \sum_{i=1}^n \psi'_i U_A(i \mid t)$. Classic contract $\langle t, \psi \rangle$ is incentive compatible if for any mixed action ψ', it holds that $U_A(\psi' \mid t) \leq U_A(\psi \mid t)$. We use $U_A(\langle t, \psi \rangle)$ for the agent's utility under incentive compatible contract $\langle t, \psi \rangle$, and $U_P(\langle t, \psi \rangle) = \sum_i \psi_i U_P(i \mid t)$ for the principal's utility.

We define the agent's expected utility for mixed action ψ' under ambiguous contract $\langle \tau, \psi \rangle$ to be the minimum utility under any payment function $t \in \tau$. That is, the agent's expected utility is $U_A(\psi' \mid \tau) = \min_{t \in \tau} U_A(\psi' \mid t)$.

We say that ambiguous contract $\langle \tau, \psi \rangle$ implements mixed action ψ if for every mixed action ψ' it holds that $U_A(\psi' \mid \tau) \leq U_A(\psi \mid \tau)$. An incentive compatible contract $\langle \tau, \psi \rangle$ is consistent if for any two contracts $t, t' \in \tau$ it holds that $U_P(\psi \mid t) = U_P(\psi \mid t')$ and, thus, $T_\psi(t) = T_\psi(t')$. For a (consistent) incentive compatible contract $\langle \tau, \psi \rangle$, we write $U_A(\langle \tau, \psi \rangle)$ for the agent's utility under ψ, and $T_\psi(\tau)$ for the resulting expected payment. The principal's expected utility under (consistent) incentive compatible contract $\langle \tau, \psi \rangle$ is $U_P(\langle \tau, \psi \rangle) = U_P(\psi \mid \tau) = R_\psi - T_\psi(\tau)$.

In what follows, without loss, we restrict attention to consistent incentive compatible contracts.

6.2 Mixing Hedges Against Ambiguity

When the agent can choose mixed actions, the principal cannot gain by employing an ambiguous contract. To prove this result we make use of the min-max theorem applied to a suitably defined zero-sum game.

Theorem 6. *Consider an incentive compatible ambiguous contract $\langle \tau, \psi \rangle$, with payoffs $U_A(\psi \mid \tau)$ and $U_P(\psi \mid \tau)$. Then there exists an incentive compatible classic contract $\langle \hat{t}, \psi \rangle$ with the same agent and principal payoffs.*

Proof. Let $\langle \tau, \psi \rangle$ be an incentive compatible ambiguous contract with payoffs $U_A(\psi \mid \tau)$ and $U_P(\psi \mid \tau)$.

Consider a zero-sum game played by the agent and the principal. The agent's (convex and compact) strategy set is the set of mixed actions $\psi' \in \Delta^n$, while the principal's (also convex and compact) strategy set is the set of payment functions $\hat{\mathcal{T}} = \{t \in \mathbb{R}^m_+ : \sum_{i=1}^n \sum_{j=1}^m \psi_i p_{ij} t_j = T_\psi(\tau)\}$. This is the set of payment functions that preserve the principal's payoff $U_P(\psi \mid \tau)$ when the agent plays action ψ. Note that $\hat{\mathcal{T}}$ is non-empty since every payment function $t \in \tau$ has the same expected payment under ψ (namely $T_\psi(\tau)$). The agent's payoff in this game is $\sum_{i=1}^n \sum_{j=1}^m \psi_i (p_{ij} t_j - c_i)$, while the principal's payoff is the negative of this quantity. Note that this way the agent's payoff in this game is precisely the agent's utility in the principal-agent setting, while the principal's payoff differs from that in the principal-agent setting.

The min-max theorem implies:

$$\max_{\psi' \in \Delta^n} \min_{t \in \hat{\mathcal{T}}} \sum_{i=1}^{n} \sum_{j=1}^{m} \psi'_i (p_{ij} t_j - c_i) = \min_{t \in \hat{\mathcal{T}}} \max_{\psi' \in \Delta^n} \sum_{i=1}^{n} \sum_{j=1}^{m} \psi'_i (p_{ij} t_j - c_i). \tag{4}$$

Note that $\hat{\mathcal{T}}$ is just a set of payment functions, so we can interpret it as a (potentially infinite-size) set that constitutes the payment functions within an ambiguous contract. Moreover, by the construction of $\hat{\mathcal{T}}$, we have $U_A(\psi \mid \hat{\mathcal{T}}) = U_A(\psi \mid \tau)$, since ψ ensures utility $U_P(\psi \mid \tau)$ against every element of $\hat{\mathcal{T}}$ to the principal, thus utility $U_A(\psi \mid \tau)$ against every element of $\hat{\mathcal{T}}$ to the agent.

We next show that the value of the zero-sum game is also equal to this quantity. To this end, we show that

$$U_A(\psi \mid \tau) \geq \max_{\psi' \in \Delta^n} \min_{t \in \hat{\mathcal{T}}} \sum_{i=1}^{n} \sum_{j=1}^{m} \psi'_i (p_{ij} t_j - c_i) \geq U_A(\psi \mid \hat{\mathcal{T}}).$$

To see this, observe that

$$\begin{aligned} U_A(\psi \mid \tau) &= \min_{t \in \tau} \sum_{i=1}^{n} \sum_{j=1}^{m} \psi_i (p_{ij} t_j - c_i) = \max_{\psi' \in \Delta^n} \min_{t \in \tau} \sum_{i=1}^{n} \sum_{j=1}^{m} \psi'_i (p_{ij} t_j - c_i) && (5) \\ &\geq \max_{\psi' \in \Delta^n} \min_{t \in \hat{\mathcal{T}}} \sum_{i=1}^{n} \sum_{j=1}^{m} \psi'_i (p_{ij} t_j - c_i) \geq \min_{t \in \hat{\mathcal{T}}} \sum_{i=1}^{n} \sum_{j=1}^{m} \psi_i (p_{ij} t_j - c_i) && (6) \\ &= U_A(\psi \mid \hat{\mathcal{T}}). && (7) \end{aligned}$$

The second equality in (5) holds by the fact that τ implements ψ, the first inequality in (6) follows from comparing feasible sets for the corresponding minimizations, and the following inequality holds because ψ is feasible in the maximization.

Consider the payment function $\hat{t} \in \mathrm{argmin}_{t \in \hat{\mathcal{T}}} \max_{\psi' \in \Delta^n} \sum_{i=1}^{n} \sum_{j=1}^{m} \psi'_i (p_{ij} t_j - c_i)$. From (4), there is no action in Δ^n giving the agent a payoff against $\hat{t}$ higher than $\max_{\psi' \in \Delta^n} \min_{t \in \hat{\mathcal{T}}} \sum_{i=1}^{n} \sum_{j=1}^{m} \psi'_i (p_{ij} t_j - c_i) = U_A(\psi \mid \tau)$. By construction, mixed action ψ gives this payoff when facing the payment function $\hat{t}$. Hence, $\langle \hat{t}, \psi \rangle$ is an incentive compatible classic contract with payoffs $U_A(\psi \mid \tau)$ to the agent and $U_P(\psi \mid \tau)$ to the principal. □

The ability to mix provides the agent with more alternative actions, tightening the incentive constraints enough to dissipate any advantage the principal gains from ambiguous contracts. The following example illustrates this.

Example 8. *Return to Example 2.* Suppose the agent is restricted to pure actions. As we have seen, the uniquely optimal ambiguous contract is $\langle \tau, 4 \rangle$, with $\tau = \{(0, 2, 0), (0, 0, 2)\}$. The payoffs to actions 1, 2, 3 and 4 under τ are 0, $-1/4$, $-1/4$ and 0. Pure actions 1, 2, and 3 are thus strictly inferior to action 4 for the agent.

Now suppose the agent can choose a mixed action. If the payoffs to the pure actions had been generated by classical contracts, then no mixture over actions 2 and 3 could give a payoff higher than $-1/4$, and hence no such mixture could be superior to action 4. Under ambiguity, this familiar property of mixed actions breaks down. The mixture that places probability 1/2 on each of actions 2 and 3 gives an expected payoff of 1/4, strictly larger than the payoffs to any pure strategy in the support of the mixture, ensuring that the ambiguous contract $\langle \tau, 4 \rangle$ no longer implements action 4. Indeed, given the agent's ability to choose this mixture, action 4 cannot be implemented at any payment less than 3/2, matching the payment under classic contracts. ■

6.3 Relation to Ellsberg Paradox and Raiffa's Critique

Ellsberg (1961) pioneered the argument that humans tend to prefer choices with quantifiable risks over those with unquantifiable, incalculable risks, giving rise to the ambiguity-aversion literature. We have shown that the principal can take advantage of the agent's ambiguity-aversion if, but only if, the agent is restricted to pure actions.

A critique of Ellsberg's experiments, raised by Raiffa (1961), is that when faced with the Ellsberg urns, a player could mentally flip a coin and implement a mixed action that induces an objective probability distribution over outcomes. Doing so removes all of the ambiguity from the decision, and with it any need for ambiguity aversion. Raiffa's argument highlights the potential power an ambiguity-averse agent can derive from engaging in mixed actions. Indeed, Theorem 6 shows that engaging in mixed strategies completely eliminates the principal's power stemming from ambiguous contracts.

Raiffa's argument has given rise to a discussion, centered around the question of whether such mental coin flips are indeed effective in banishing uncertainty (e.g., Ke and Zhang (2020); Saito (2015)). One way of formalizing this "critique of Raiffa's critique," is Bade (2023)'s notion of dynamic semi-consistency. According to this behavioral assumption, agents should not update their beliefs in response to signals that are independent of the environment (such as an independent coin flips). This notion rules out the type of mixing required for Theorem 6, but aligns with our behavioral assumptions underlying the results in Sections 2–5.

References

Aghion, P. and Holden, R. (2011). Incomplete contracts and the theory of the firm: What have we learned over the past 25 years? *Journal of Economic Perspectives*.

Bade, S. (2023). Ambiguity by design. Working paper.

Beauchêne, D., Li, J., and Li, M. (2019). Ambiguous persuasion. *Journal of Economic Theory*, 179:312–365.

Bergemann, D. and Schlag, K. (2011). Robust monopoly pricing. *Journal of Economic Theory*, 146:2527–2543.

Bodoh-Creed, A. L. (2012). Ambiguous beliefs and mechanism design. *Games and Economic Behavior*, 75:518–537.

Bose, S., Ozdenoren, E., and Pape, A. D. (2006). Optimal auctions with ambiguity. *Theoretical Economics*, 1:411–438.

Bose, S. and Renou, L. (2014). Mechanism design with ambiguous communication devices. *Econometrica*, 82(5):1853–1872.

Carrasco, V., Farinha Luz, V., Kos, N., Messner, M., Monteiro, P., and Moreira, H. (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):536–63.

Carroll, G. and Walton, D. (2022). A general framework for robust contracting models. *Econometrica*, 90:2129–2159.

Cheng, X. (2020). Ambiguous persuasion: An ex-ante formulation. *arXiv preprint arXiv:2010.05376*.

Cheng, X., Klibanoff, P., Mukerji, S., and Renou, L. (2024). Persuasion with ambiguous communication. Technical report, Florida State University, Northwestern University, and Queen Mary University.

Dai, T. and Toikka, J. (2022). Robust incentives for teams. *Econometrica*, 90(4):1583–1613.

Di Tillio, A., Kos, N., and Messner, M. (2017). The design of ambiguous mechanisms. *The Review of Economic Studies*, 84:237–276.

Dütting, P., Roughgarden, T., and Talgam-Cohen, I. (2019). Simple versus optimal contracts. In *EC'19: The 20th ACM Conference on Economics and Computation, Phoenix, AZ, USA, June 24-28, 2019*, pages 369–387. Full version available from: `https://arxiv.org/abs/1808.03713`.

Dütting, P., Roughgarden, T., and Talgam-Cohen, I. (2021). The complexity of contracts. *SIAM Journal on Computing*, 50:211–254.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics*, 75:643–669.

Gilboa, I. and Schmeidler, D. (1993). Updating ambiguous beliefs. *Journal of Economic Theory*, 59:33–49.

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

Hart, O. (1988). Incomplete contracts and the theory of the firm. *Journal of Law, Economics, and Organization*, 4:119–139.

Hart, O. and Moore, J. (1988). Incomplete contracts and renegotiation. *Econometrica*, 56:755–785.

Hermalin, B. E. and Katz, M. L. (1991). Moral hazard and verifiability: The effects of renegotiation in agency. *Econometrica*, 59:1735–1753.

Holmström, B. (1979). Moral hazard and observability. *The Bell Journal of Economics*, 10(1):74–91.

Holmstrom, B. and Milgrom, P. (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica: Journal of the Econometric Society*, pages 303–328.

Innes, R. D. (1990). Limited liability and incentive contracting with ex-ante action choices. *Journal of economic theory*, 52(1):45–67.

Kambhampati, A. (2023). Randomization is optimal in the robust principal-agent problem. *Journal of Economic Theory*, 207:105585.

Ke, S. and Zhang, Q. (2020). Randomization and ambiguity aversion. *Econometrica*, 88(3):1159–1195.

Laffont, J.-J. and Martimort, D. (2009). *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press.

Lopomo, G., Rigotti, L., and Shannon, C. (2011). Knightian uncertainty and moral hazard. *Journal of Economic Theory*, 146(3):1148–1172.

Raiffa, H. (1961). Risk, ambiguity and the savage axioms: A comment. *Quarterly Journal of Economics*, 75:690—694.

Saito, K. (2015). Preferences for flexibility and randomization under uncertainty. *American Economic Review*, 105(3):1246–1271.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57:571–587.

Zuboff, S. (2023). The age of surveillance capitalism. In *Social Theory Re-Wired*, pages 203–213. Routledge.