

MORAL HAZARD, UNCERTAIN TECHNOLOGIES, AND LINEAR CONTRACTS

MARTIN DUMAV, URMEE KHAN

ABSTRACT. We analyze a moral hazard problem where both contracting parties have imprecise information (non-probabilistic uncertainty) about how actions translate to output. Agent has (weakly) more precise information than Principal, and both seek robust performance from a contract in relation to their respective worst-case scenarios. We show that linear contracts that align Principal's and Agent's pessimistic expectations are optimal. This result holds under very general conditions on the structure of information, including the case where Principal does not know exactly the extent of disagreement between her and Agent's information. Methodologically, by using only the properties of the sets of expected payoffs to derive the results, we provide a way to characterize optimal contracts without requiring such knowledge on the part of Principal. Substantively, our results provide some insights into the formal link between robustness and simplicity of contracts, in particular that non-linearity creates sub-optimal divergence between the respective worst-cases.

Keywords: Moral hazard, non-probabilistic uncertainty, robustness, linear contracts

JEL Classification: D81, D82, D86

1. INTRODUCTION

Consider a Principal who possesses only coarse information about relevant technology for a project, and wishes to hire an Agent. In particular, she knows exactly the set of actions Agent might choose from, but has less precise knowledge than does Agent about possible consequences of those actions. Out-sourcing to an 'expert' is a natural example: presumably Principal would be less knowledgeable about the technology involved, nevertheless it could be contractually possible to restrict the allowable set of actions. More generally, such issues may arise in various circumstances: dealing with brand new technology or business model usually involves a lack of past data, or information coming from conflicting sources. In such settings, it is plausible to imagine that the person to be hired, even if not himself precisely informed, might have the benefit of some prior experience to know a little bit more than his potential employer. What kind of a contract would Principal offer?

Date: January 26, 2019. The first author gratefully acknowledges the financial support from the Ministerio Economia y Competitividad (Spain) through grants ECO2017-86261-P and MDM 2014-0431.

Integral to the answer of this question is the attitude of contracting parties towards such imprecise information. We analyze the case where both parties are concerned with *robustness* of their actions to the imprecision of their own perceptions, and evaluate options according to their worst-case guaranteed value. In this setting, we provide a strikingly simple answer to the question posed above: linear contracts, i.e. simple sharing rules, do as well as any arbitrarily complicated contract.

In our model, Principal and Agent both know the common set of actions available. But for any action, the consequences are *ambiguous*. Namely, consequences are perceived as *sets* of distributions. Moreover, Principal has a broader uncertainty about these *sets*, i.e., Agent’s perceived set of distributions associated with each action is contained in Principal’s, and that is all that she knows.

With our formulation of non-probabilistic beliefs, and robustness concerns of the parties, we show that linear contracts provide a direct and monotonic relationship between the parties’ values by aligning their worst-case scenarios. The driving force behind our results is the observation that Principal faces *endogenous* uncertainty for any given contract, as she cannot uniquely infer Agent’s action choice, rather has to allow for multiplicity of actions consistent with her rough information. As a result, ex-ante alignment of incentives takes center stage, as opposed to ex-post inference.

The main contribution of the paper is therefore twofold. Formally, we establish justification for linear contracts in an environment where, even though Principal can restrict the set of allowable actions, her ambiguous perception of consequences of those actions lead her to seek guarantees. In the process, the paper also provides insights about a more general question about incentive contracting, namely, what happens to ‘high-powered incentives’ as the information available to Principal becomes more imprecise in comparison to standard Bayesian models? Our results suggest that such incentives lose their impact.

In order to contextualize our results, at this point it is useful to take a step back and recap the standard formulation of ‘moral hazard’ in a principal-agent setting. These are problems involving unobserved actions taken by Agent after the contract is agreed to, actions that affect the outcome and thus the payoff to Principal. The workhorse model in solving such contract design problems comes from Grossman and Hart (1983) (we abbreviate it as G-H from now on). The ‘riskiness’ of the environment is captured in a common prior probability, and Principal’s problem is solved in the following steps:

- calculate ex-ante expected costs of inducing Agent to undertake any available action,

- figure out which action she would like Agent to take,
- write a contract to ‘implement’ that action.

Let us take note of the key substantive elements of the G-H approach and the environment it is embedded in. First, Principal knows enough to decide, from an ex-ante perspective, what action(s) she would like to be implemented. Second, the information structure allows her to draw inference after the fact, so that implementation through ‘reward-punishment’ schemes based on some observed and verifiable signal is feasible. Third, depending on the details of a particular model, one can end up with non-linear, even non-monotone contracts (in contract theory terminology, such contracts provide ‘high-powered’ incentives). More often than not, the complicated form of these contracts is not robust to perturbation of the Bayesian information structure. To ensure even a basic property like monotonicity, assumptions that are only known to be sufficient, but not necessary, are imposed (for example, monotone likelihood ratio property for outcome distributions).

The above set of observations immediately give rise to a number of questions. What if the environment itself is not well-understood, at least not well enough for it to be plausibly modeled with a unique/precise and common prior/probability distribution? The examples we cited at the beginning are but a subset of all such possibilities. What are ‘robust’ solutions to moral hazard problems in an environment with such imprecise information? Furthermore, what happens to ‘high-powered incentives’ as the information available to Principal becomes more imprecise in comparison to standard Bayesian models? We show that simple sharing rules turn out to be optimal when we seek robustness in the face of such imprecision. As for high-powered incentives captured in non-linear contracts, they fail to improve upon simple sharing rules. The key factor behind this result is the following: for any arbitrary non-linear contract, due to ambiguity or imprecise information, Principal and Agent differ in their perceptions of worst-case scenarios, unlike in the Bayesian model. Given that the worst-case scenarios are the criteria they use to evaluate options because of robustness concerns, aligning these worst-case scenarios improves their ex-ante welfare. Linear contracts provide the best of such improvements by fully aligning the worst-case scenarios.

Both in our formal approach and our focus on robustness, Carroll (2015) is our main inspiration. Carroll poses a related, but different question. Principal is unsure about the possible actions that Agent might take. However, each action, known or unknown, has precise probabilistic consequences, as quantified in terms of an unique/precise probability distribution over outcomes. Moreover, Principal and Agent agree on the consequences of

the commonly known actions. In this sense, the model in Carroll (2015) is a hybrid one - it extends the common-prior, common-knowledge Bayesian information structure with the addition of some actions that are known to the Agent, but not to the Principal. A rationalization for linear contracts is provided, as Principal seeks guarantees robust to Agent's entire action set.

We expand on the theme introduced by Carroll (2015), by introducing a general framework to capture imprecise, non-probabilistic perceptions about consequences of actions, even if the possible set of actions are commonly known. Moreover, our formulation allows for different degrees of precision for Principal and Agent, with Agent possessing weakly more precise information. In comparison with Carroll (2015), although we adopt a similar mathematical approach based on the affine geometry of the space of outcomes and wages, the substantive differences in our primitive environment necessitate a different set of arguments leading to our results. The key differences arise from allowing non-probabilistic beliefs for both parties. This leads to, in an *endogenous* manner, a rich set of actions Agent can choose for a given contract. The intuitive justification for linearity of contracts in our setting is therefore different: it reduces the impact of this endogenous uncertainty about Agent's optimal action on Principal's ex-ante guaranteed payoff.

To see the difference in the formal aspects of our model in comparison with Carroll (2015), note the relevant sets generated in the output-wage space. For any given contract, Principal in Carroll (2015) sees the possibilities that may arise from any of the actions known to her. But given that she is aware of the existence of actions that she does not know about, the entire possibility set encompasses the whole convex hull generated by the contract in question. For the technical analysis that shows the optimality of linear contracts, it is absolutely essential that Principal envisages the entire convex hull to be possible. It is indeed the required 'richness' condition in the paper (see p. 545 in Carroll (2015)) that help identify the optimal linear contract. In other words, the lack of information on Principal's part about possible actions, combined with the modeling feature that each action generates a unique probability distribution over outcomes, means that from Principal's point of view, the overall possibility set is much larger than what she can envisage based on her information.

In the current paper, the relative imprecision of Principal's information in fact leads her to view the possibility sets in a different way. In our model, both parties may have imprecise information about consequences of actions, as opposed to actions themselves ¹. It

¹This is not merely a semantic difference. In Carroll (2015), Agent choosing an action means him choosing a distribution over outcomes. In our model, agent choosing an action merely means a set of distributions over outcomes.

is this imprecise knowledge of consequences, that generates non-trivial sets of possibilities as envisaged by the Principal. However, these sets do not need to encompass the entire convex hull and this makes our agency problem technically challenging. Our key observation for traction is that the boundaries of these sets are well-defined from the Principal’s perspective. Moreover, her possibility sets are supersets of the sets perceived by Agent, not subsets. Hence, the ‘richness’ condition required in our environment is much weaker and entirely well-defined for operational purposes. Namely, we only need that Principal attaches non-singleton, convex, compact sets of distributions to each action.

Our paper fits more broadly in the growing literature on design of contracts in ambiguous environments. A desire for robust mechanisms often motivates the use of simple contracts in this literature, starting with Hurwicz and Shapiro (1978), Mukerji (1998) and more recently, for instance, Chassang (2013), Antic (2014), Garrett (2014), in addition to Carroll (2015).

Given the ubiquity of linear sharing rules, they have received considerable attention from contract theorists from the very early days. Holmstrom and Milgrom (1987) and Diamond (1998) offer foundations for linear contracts in the standard Bayesian setting. The intuition provided in Diamond (1998) is close to the spirit of our paper: it aligns Principal’s and Agent’s objectives. In our setting with imprecise non-probabilistic beliefs, linear contracts are optimal in moral hazard relations because they align the pessimistic expectations of the contracting parties.

Our paper also relates to the literature on robust mechanism design with ambiguity, including Bergemann and Morris (2005), Bergemann and Schlag (2011), Garrett (2014), Frankel (2014). Robustness of high powered incentives have also been analyzed in specific applications. Mukerji (1998) shows that in a vertical relationship optimal contracts feature incompleteness and low powered incentives. Ghirardato (1994) shared a similar motivation. His model of moral hazard problem with ambiguity modeled with Choquet capacities illustrates the difficulties involved in analyzing the moral hazard problem with ambiguity by standard method of considering expected benefit of an action and the agency cost of implementing it, as in G-H framework. In our setting, the main challenge is that contracting parties can disagree about the worst-case scenarios, which themselves are endogenous objects. In confronting this challenge we directly cast the contracting problem in output and payoff space and analyze the complications with ambiguity in a tractable manner.

Lastly, we want to make a remark about our use of the MaxMin criteria as a formalization of attitude towards imprecise information or non-probabilistic beliefs. One class of axiomatizations of MaxMin preferences come from an approach that takes the state space as a

primitive. We do our analysis on the space of outcomes and payments. In general, the state-space approach towards ambiguity is not conceptually amenable to agency problems, as the probabilities of outcomes in the latter are conditional on actions (see Karni (2006) and Karni (2009) for discussions on this issue). We therefore do not invoke such axiomatizations as a foundation for our formulation, instead we see this simply as a plausible way to capture the idea of robustness concerns when designing contracts. Our take is that such concerns are more appropriately modeled as evaluations of outcome prospects.

2. MODEL

2.1. Outcomes and Actions. Let $\mathcal{Y} \subseteq \mathbb{R}_+$ denote a compact convex set of outcomes, $\Delta(\mathcal{Y})$ the set of Borel distributions on \mathcal{Y} with the weak* topology and $\mathbb{K}_{\Delta(\mathcal{Y})}$ denote the class of non-empty, compact, convex subsets of $\Delta(\mathcal{Y})$.

Agent chooses an unobservable action a from a compact set \mathcal{A} ²

Let $g : \mathcal{A} \mapsto \mathbb{R}_+$ be a continuous, bounded function that describes the cost of effort to Agent.

Assumption 1. *There exists $a \in \mathcal{A}$ such that $g(a) = 0$.*³

2.2. Information. Principal's information about technology is characterized as an upper-hemicontinuous set-valued mapping from set of actions \mathcal{A} to $\mathbb{K}_{\Delta(\mathcal{Y})}$ given by $Q^P(\cdot) : \mathcal{A} \mapsto \mathbb{K}_{\Delta(\mathcal{Y})}$. We interpret this set-valued mapping as follows: for each action $a \in \mathcal{A}$, Principal believes that this action induces a convex and compact set of distributions over outcomes denoted by $Q^P(a) \in \mathbb{K}_{\Delta(\mathcal{Y})}$. Similarly, Agent's information is characterized by another upper-hemicontinuous mapping, $Q^A(\cdot) : \mathcal{A} \mapsto \mathbb{K}_{\Delta(\mathcal{Y})}$. The key assumption we make is that Agent has (weakly) more precise knowledge about technology than Principal.

Assumption 2. $Q^A(a) \subseteq Q^P(a)$ for each $a \in \mathcal{A}$.

Let $\mathcal{Q}^A(a) = \{M \in \mathbb{K}_{\Delta(\mathcal{Y})} : M \subset Q^P(a)\}$ denote the collection of all possible sets of distributions induced by action a that is consistent with Principal's information. The family of all such collections of sets, for all possible actions, denoted by $\mathcal{Q}^A = \{\mathcal{Q}^A(a) : a \in \mathcal{A}\}$

²This plays no more of a role than that of an index set. A closed, bounded set in \mathbb{R} , including a finite set, works.

³This zero cost action does not play any major role in our analysis, but it does help us formalize the case where it is optimal for Principal to offer a zero contract. We interpret it as the option of doing nothing.

describes Principal's perception of Agent's information about technological possibilities. Q^A thus contains the range of any possible mapping $Q^A(\cdot)$ that Principal might imagine for Agent to possess.

2.3. Payoffs and Timing. We assume both Principal and Agent evaluate contracts by their worst-case expected payoff. In our primary analysis, we maintain the assumption that both are risk-neutral over monetary payoffs. In an extension, we show that introducing risk aversion for Agent does not qualitatively change our results.

A contract is a continuous, bounded function $w : \mathcal{Y} \rightarrow \mathbb{R}_+$ that specifies output contingent payments and protects Agent with limited liability (i.e. $w(0) \geq 0$). In particular, we allow the zero contract: $w(y) = 0, \forall y \in \mathcal{Y}$.

The timing of the contracting game is as follows:

- (1) Principal offers a contract w ;
- (2) Agent, knowing Q^A , chooses action $a \in \mathcal{A}$;
- (3) output y is realized;
- (4) payoffs are received: $y - w(y)$ to Principal and $w(y) - g(a)$ to Agent.

3. SYMMETRIC AMBIGUITY

We begin with the benchmark case in which Principal and Agent both have the same, ambiguous information. Formally, $Q^A(a) = Q^P(a), \forall a \in \mathcal{A}$. In this symmetric case, we assume that none of the sets $Q^i(\cdot)$ are singletons, so as not to let the problem collapse into the unique and common prior case. This is an important assumption for our technical approach to work, so we note it here.

Assumption 3. $Q^A(a) = Q^P(a) = Q(a), \forall a \in \mathcal{A}$, with each $Q(a)$ non-singleton.

Consider $\mathcal{C}(w) = \text{co} \{(y, w(y)) : y \in \mathcal{Y}\}$, where co denotes convex hull. Given that $w(y)$ is a continuous, bounded function on convex, compact \mathcal{Y} , the graph of $w(y)$, denoted by $gr(w)$, is closed and compact. Hence $\mathcal{C}(w)$ is closed and compact as the convex hull of $gr(w)$.

Consider also any action $a \in \mathcal{A}$ that Agent might choose while working under $w(y)$. Let

$$\mathcal{S}_a(w|Q) = \{(\mathbb{E}_q y, \mathbb{E}_q w(y)) | q \in Q(a), y \in \mathcal{Y}\}$$

Clearly, $\mathcal{S}_a(w|Q)$ is fully contained in $\mathcal{C}(w)$. We note also that $\mathcal{S}_a(w|Q)$ is a convex, compact set in \mathbb{R}^2 .

We are going to use the two sets defined above extensively. Our first lemma is about the dimensionality of $\mathcal{S}_a(w|Q)$.

Lemma 1. *For an arbitrary non-linear contract $w(y)$, $\mathcal{S}_a(w|Q)$ can be characterized in one of the following two ways:*

- (i) *it is a convex compact subset of \mathbb{R}^2 with non-empty relative interior in \mathbb{R}^2 ; hence \mathbb{R}^2 is its affine hull and it has dimension 2.*
- (ii) *it is a line segment, with dimension 1.*

Proof. (i) follows from the fact that $\mathcal{S}_a(w|Q)$ is a convex compact set, hence by Minkowski's theorem is the convex hull of its extreme points.⁴ For any $Q(a)$ that has more than two points in its support, a 2-dimensional $\mathcal{S}_a(w|Q)$ would arise.

As for (ii), relative interior in \mathbb{R}^2 of $\mathcal{S}_a(w|Q) = \emptyset$ implies that $\mathcal{S}_a(w|Q)$ is a linear segment, which is only possible for a non-linear $w(y)$, and convex, compact $Q(a)$ only if $Q(a)$ contains distributions with only a two-point support.⁵ \square

We have an immediate corollary of case (i) in Lemma 1. Define Π as the projection from \mathbb{R}_2 to \mathbb{R}_1 , with $\Pi(u, v) = u$. Define the *slice* of $\mathcal{S}_a(w|Q)$ along u as

$$\sigma(u) = \{v \in \mathbb{R}_+ | (u, v) \in \mathcal{S}_a(w|Q)\}$$

Lemma 2. *For any non-linear $w(y)$, and any $Q(a)$ that has more than two points in the support, the slices of $\mathcal{S}_a(w|Q)$ along y are not singletons, unless at an extreme point of $\mathcal{S}_a(w|Q)$.*

Proof. The implication follows directly from the observation that slices along y being singletons except at extreme points would imply $\mathcal{S}_a(w|Q)$ is a line segment. \square

3.1. Agent's Optimization Problem. Agent solves the following problem:

$$\max_{a \in \mathcal{A}} \left(\min_{q \in Q(a)} \mathbb{E}_q[w(y)] - g(a) \right)$$

⁴For most of the definitions and results on convex analysis and affine geometry used in this paper, we refer to Hiriart-Urruty and Lemaréchal (2012). See Ch 4 for this theorem and discussion.

⁵or if $w(y)$ has a linear segment, which is an uninteresting case in our context. In general, when we refer to a contract as being *non-linear*, we mean it to be strictly so.

Compactness of $Q(a)$ in weak* topology ensures the \min is well defined, and our upper-hemicontinuity assumption on the correspondence Q ensures the same for the \max .⁶

Given symmetric ambiguity Principal can fully infer Agent's decision rule⁷:

$$a^*(w|Q) = \operatorname{argmax}_{a \in \mathcal{A}} \left(\min_{q \in Q(a)} \mathbb{E}_q[w(y)] - g(a) \right) \quad (1)$$

The best of the worst-case scenarios perceived by Agent is thereby $q_A^*(w|Q) \in Q(a^*(w|Q))$.⁸

The next lemma characterizes the max-min value for Agent in terms of a supporting hyperplane of $\mathcal{S}_{a^*}(w|Q)$.

Lemma 3. *Let $V_A(w|Q)$ be the value of the problem defined in (3.1). This value defines a supporting hyperplane for $\mathcal{S}_{a^*}(w|Q)$*

Proof. Let $H_{A,w}$ be the hyperplane defined as

$$H_{A,w} = \{(u, v) \in \mathbb{R}^2 \mid v = V_A(w|Q)\} \quad (2)$$

It exposes the face of $\mathcal{S}_{a^*}(w|Q)$ with $\mathbb{E}_{q^*}(w(y)) = V_A(w|Q) + g(a^*)$ and is thereby a supporting hyperplane for $\mathcal{S}_{a^*}(w|Q)$. \square

3.2. Principal's Optimization Problem. For an arbitrary contract w , Principal can infer agent's decision rule given the symmetry of ambiguous perceptions. Hence Principal's ex-ante guaranteed value

$$V_P(w|Q) = \min_{q \in Q(a^*(w|Q))} \mathbb{E}_q[y - w(y)] \quad (3)$$

His worst-case scenario corresponds to those distributions that attain the minimum, denoted by $q_P^*(w|Q) \in Q(a^*(w|Q))$.⁹

Observation: Principal and Agent do not have to agree on the worst-case scenario, i.e., $q_P^*(w|Q)$ and $q_A^*(w|Q)$ can, and typically will, differ. This point turns out to be central to much of our subsequent analysis.

⁶See Appendix for a detailed argument.

⁷For simplicity, in the main body of the paper we take this argmax to be a singleton. There are several ways of dealing with case of multiplicity, as we discuss in the Appendix, but our qualitative results do not change.

⁸If there are multiple such distributions, Agent would be payoff-indifferent among them.

⁹Once again, we can ignore the issue of multiplicity without loss of generality, because of payoff-indifference. In any case, it would be a compact set of distributions.

Principal offers a contract w that solves

$$\sup_w V_P(w|Q) \quad \text{subject to} \quad V_A(w|Q) \geq 0$$

For the subsequent analysis we assume that contracting relation is viable: for some contract w , we have that $V_P(w|Q) \geq 0$ and $V_A(w|Q) \geq 0$. Since we allow zero contracts, and zero cost actions, the set of contracts satisfying the two constraints is non-empty.

The next set of lemmas characterize another supporting hyperplane for $\mathcal{S}_{a^*}(w|Q)$ representing Principal's guaranteed value associated with $Q(a)$.

Define $H_{P,w}(\lambda)$ as the hyperplane

$$H_{P,w}(\lambda) = \{(u, v) \in \mathbb{R}^2 | u - v = \lambda\} \quad (4)$$

A family of these hyperplanes can be used to represent Principal's iso-profit curves corresponding to expected profit level λ , i.e.

$$H_{P,w}(\lambda) = \{(\mathbb{E}y, \mathbb{E}w(y)) \in \mathbb{R}^2 | \mathbb{E}y - \mathbb{E}w(y) = \lambda\} \quad (5)$$

Geometrically in \mathbb{R}^2 , these are rays through the origin defined by the equation $y - w(y) = \lambda$. In particular, Principal's guaranteed expected payoff corresponding to any a^* chosen by Agent under a contract $w(y)$ is represented by $H_{P,w}(V_P(w|Q))$. Recall that Principal's guaranteed value corresponds to the worst-case scenario related to a^* , hence it is represented by an extreme point or face of $\mathcal{S}_{a^*}(w|Q)$. In other words, we have the following.

$$H_{P,w}(\pi) = \left\{ (u, v) \in \mathbb{R}^2 | v - u = \pi \right\} \quad (6)$$

Lemma 4. $H_{P,w}(V_P(w|Q))$ is a supporting hyperplane of $\mathcal{S}_{a^*}(w|Q)$.

Proof. Consider the family of all half-spaces $\{(u, v) \in \mathbb{R}^2 | u - v \geq \lambda\}$ that contains $\mathcal{S}_{a^*}(w|Q)$. The half-space generated by the intersection of this family of half-spaces corresponds to the smallest λ for which $\mathcal{S}_{a^*}(w|Q)$ is fully contained in it. The defining hyperplane has $\lambda = V_P(w|Q)$, and is by construction a supporting hyperplane for $\mathcal{S}_{a^*}(w|Q)$ \square

We now have an affine set containing all $(\mathbb{E}y, \mathbb{E}w(y))$ pairs that would be incentive compatible for implementing a^* , while guaranteeing Principal at least $V_P(w|Q)$, and also a relationship between $V_A(w|Q)$ and $V_P(w|Q)$ under any contract $w(y)$ that induces a^* as the optimal action.

Lemma 5. $\mathcal{S}_{a^*}(w|Q)$ is contained in the 2-dimensional affine convex cone generated by $H_{A,w}$ and $H_{P,w}$ hyperplanes.

Proof. $H_{A,w}$ and $H_{P,w}$ are supporting hyperplanes of $\mathcal{S}_a(w|Q)$; the intersection of the two associated half-spaces that both contain $\mathcal{S}_a(w|Q)$ is an affine convex cone. \square

As an immediate implication we have the following:

Lemma 6. For any contract $w(y)$ that induces a^* as the optimal action,

$$V_P(w|Q) = \min_{q \in Q(a^*)} \mathbb{E}_q[y - w(y)] \quad \text{subject to} \quad \min_{q \in Q(a^*)} [\mathbb{E}_q w(y) - g(a^*)] \geq V_A(w|Q)$$

Proof. By construction, for any contract that implements a^* as the optimal action, $\mathcal{S}_{a^*}(w|Q)$ contains all feasible $\{\mathbb{E}y, \mathbb{E}w(y)\}$ pairs. By definition, $V_P(w|Q)$ and $V_A(w|Q)$ are the worst-cases respectively for Principal and Agent over this feasible set with expectations taken over their respective worst-case distributions. \square

Note that the solution to the problem specified in Lemma 6 exists, given a compact \mathcal{Y} and bounded $w(y)$. It is essentially a linear programming problem in the $(y, w(y))$ space. The program feasible and bounded by construction.

Next we establish that a linear contract can be found that will implement a^* , while making Principal and Agent both as at least as well-off as under $w(y)$. In fact, within the family of linear contracts, there is an optimal one.

3.3. Optimality of Linear Contracts. Note that whatever optimal action a^* Agent chooses, with the associated ‘best worst-case’ for Agent $q_A^* \in Q(a^*)$, Agent’s expected payments satisfies

$$\mathbb{E}_{q_A^*} w(y) \geq \mathbb{E}_{q_A^*} w(y) - g(a^*) \equiv V_A(w|Q) \tag{7}$$

A linear contract ℓ is of the form $\ell(y) = \alpha y$ with $\alpha \in (0, 1]$. For this contract Principal’s ex-post payoffs are

$$y - \ell(y) = \frac{1 - \alpha}{\alpha} \ell(y). \quad (8)$$

The next Proposition lays out a key inequality characterizing the relationship between the guaranteed values of Principal and Agent for a linear contract.

Proposition 1. *For a linear contract $\ell(y) = \alpha y$, under symmetric ambiguity, the maxmin guaranteed values for Principal and Agent are related as follows:*

$$\mathbb{E}_{q_P^*}[y - \ell(y)] = \frac{1 - \alpha}{\alpha} \mathbb{E}_{q_P^*} \ell(y) \geq \frac{1 - \alpha}{\alpha} \mathbb{E}_{q_A^*} \ell(y) \geq \frac{1 - \alpha}{\alpha} V_A(\ell|Q) \quad (9)$$

Proof. The first equality follows by taking expectation of both sides in equation (8).¹⁰

The second inequality relation follows because the worst-case of expected payments *made* according to Principal's pessimism is at least as severe than Agent's conservative perception of worst-case expected payments *received*.

Formally, for an arbitrary contract $w(y)$, any a and $Q(a)$, Agent's worst-case scenario minimizes expected payments received: $q_A = \arg \min_{q \in Q(a)} \mathbb{E}_q[w(y)]$. Equivalently, $\mathbb{E}_{q_A}[w(y)] \leq \mathbb{E}_q[w(y)]$ for all $q \in Q(a)$.

In particular, this holds for $q_P \in Q(a)$ – Principal's worst-case distribution minimizing expected profits $\mathbb{E}_q[y - w(y)]$ over $Q(a)$. Setting $\ell(y) = w(y)$, we get the inequality.

The last inequality is derived by applying equation (7) to the linear case:

$$\frac{1 - \alpha}{\alpha} \mathbb{E}_{q_A^*} \ell(y) \geq \mathbb{E}_{q_A^*} \ell(y) - g(a^*) \equiv V_A(w|Q) \quad (10)$$

□

The second inequality in (9) plays an important role in our analysis.

Let V_p be the value of Principal's problem as defined in (3.2). We then have the following theorem establishing optimality of a linear contract.

Theorem 1. *Suppose $Q^A(a) = Q^P(a)$ for all $a \in \mathcal{A}$. There exists a linear contract that maximizes V_p .*

¹⁰For a linear contract $\ell(y)$, for any action $a \in \mathcal{A}$, $\mathcal{S}_a(\ell|Q)$ is a linear segment [case (ii) in Lemma 1]

Proof. We make the argument in several steps.

- (a) Let a^* be the optimal action chosen by Agent under $w(y)$, that provides a guarantee $V_P(w|Q) = \lambda$ to Principal. Let $\gamma(a^*)$ be the associated $\mathbb{E}_{q_A^*} w(y)$. First we establish that there is an affine contract that implements a^* and guarantees Principal at least λ .

Construct the proposed affine contract as follows:

- i. If $V_P(w|Q) < 0$, choose $w(y) = 0$
- ii. If $V_P(w|Q) \geq 0$, let $w'(y) = \beta + \alpha y$, with $\alpha \in (0, 1)$, and $\beta = (1 - \alpha)\gamma - \alpha\lambda \geq 0$

It is straightforward to check that the line $\beta + \alpha y$ passes through the vertex of the affine convex cone containing $\mathcal{S}_{a^*}(w|Q)$, and thereby has non-empty intersection with $\mathcal{S}_{a^*}(w|Q)$. In fact, the line segment defined by $gr(w'(y)) \cap \mathcal{S}_{a^*}(w|Q)$ represents the set $\{(\mathbb{E}_q y, \mathbb{E}_q w'(y)) | q \in Q(a^*)\}$. Recall also that any of these pairs $(\mathbb{E}_q y, \mathbb{E}_q w'(y))$ fall within the region representing both $y - w(y) \geq \lambda$ and $w(y) \geq \gamma$ [see Lemma 6]. In other words, a^* is still a possible optimal choice for Agent, while Principal is at least as well-off as before.

To finish this step of the argument, we need to check for any possible deviation from a^* by Agent under the new contract $w'(y)$. We first note that Agent would not deviate to any action a' with $\min_{q \in Q(a')} \mathbb{E}_q w'(y) < \min_{q \in Q(a^*)} \mathbb{E}_q w(y)$. To see this, recall that Agent's indifference sets in the $(y, w(y))$ space are of the form $\{(u, v) | v \text{ constant}\}$ [geometrically, these are horizontal lines, representing a particular value of $\mathbb{E}w(y)$]. By revealed preference, a^* had $\min_{q \in Q(a^*)} \mathbb{E}_q w(y) > \min_{q \in Q(a')} \mathbb{E}_q w(y)$ for all a' with $g(a') < g(a^*)$, and possibly even some actions with $g(a') > g(a^*)$. With $\min_{q \in Q(a^*)} \mathbb{E}_q w'(y) \geq \min_{q \in Q(a^*)} \mathbb{E}_q w(y)$, we must have, if Agent were to deviate to some a' under w' ,

$$\min_{q \in Q(a')} \mathbb{E}_q w'(y) \geq \min_{q \in Q(a^*)} \mathbb{E}_q w'(y) \geq \min_{q \in Q(a^*)} \mathbb{E}_q w(y) \quad (11)$$

Hence we conclude that with this new affine contract w' , Agent either still chooses a^* , or deviates to an a' with $\min_{q \in Q(a')} \mathbb{E}_q w'(y) \geq \min_{q \in Q(a^*)} \mathbb{E}_q w'(y)$. Given the affine nature of w' , any such deviation would also improve Principal's worst-case. We have,

$$y - w'(y) = \frac{1 - \alpha}{\alpha} w'(y) - \frac{1 + \alpha}{\alpha} \beta \quad (12)$$

Taking expectations on both sides, and then taking minima, establishes the last claim. Since this is true for any a^* , we have found an affine contract that dominates $w(y)$.

- (b) We next show that, for any affine contract $w'(y)$, there is a linear contract $\ell(y)$ that does at least as well, or strictly better, for Principal, and implements the same optimal action as $w'(y)$.

If $\beta = 0$, then $w'(y)$ is already linear and we are done. If $\beta > 0$, then note that $w'(0) = \beta$. Define $\ell(y) = \alpha y = w'(y) - \beta$. This improves Principal's payoff, with Agent's incentives unchanged.

- (c) The previous steps, taken together, imply that any action that is implemented by an arbitrary contract, $w(y)$, can be implemented by a linear contract, that makes Principal and Agent both at least as well off as under $w(y)$. In other words, it is enough to only optimize within the class of linear contracts.

From (a), we can strengthen (9) in Proposition 1 to equality:

$$\mathbb{E}_{q_P^*}[y - \ell(y)] = \frac{1 - \alpha}{\alpha} \mathbb{E}_{q_P^*} \ell(y) = \frac{1 - \alpha}{\alpha} \mathbb{E}_{q_A^*} \ell(y) \quad (13)$$

This follows from noting that with a linear contract, for any action a , Principal's and Agent's worst-cases are perfectly aligned, i.e. $q_P^* = q_A^*$.

In other words, with a linear contract, the best worst-case guarantees for Principal and Agent are related via a linear relationship. Principal's value V_p is continuous in $\alpha \in [0, 1]$, and an $\alpha^*(a^*)$ exists that maximizes it.¹¹

□

4. ASYMMETRIC AMBIGUITY

In this section we will take up the case where Principal has strictly less precise information than Agent, i.e. $Q^A(a) \subset Q_P(a)$ for all $a \in \mathcal{A}$, and that is all that Principal knows. In particular, Principal does not know exactly what Agent knows, i.e. the exact subsets $Q^A(a) \subset Q_P(a)$ are not known to Principal. With a slight abuse of notation, we are going to use Q^i to denote the range of the mapping $Q^i(\cdot)$ in the following analysis. As we shall see, with a few tweaks, the results derived for the symmetric case go through for the most part.

Let us recap here the essential definitions and notations related to Principal's and Agent's optimization problems and suitably modify them for the asymmetric case. Given an output contingent contract $w(y)$, a risk neutral, ambiguity averse Agent with technology $Q^A \subset \mathcal{Q}^A$ chooses an action that solves his optimization problem:

$$a^*(w|Q^A) = \operatorname{argmax}_a \left[\min_{q \in Q^A(a)} \mathbb{E}_q w(y) - g(a) \right] \quad (14)$$

¹¹There are some subtleties regarding $\alpha = 0$, that we take up in Section 4.2

Let $q_A^*(w|Q^A) \in Q^A(a^*(w|Q^A))$ be Agent's 'best worst-case' distribution facing w and with information Q^A . The associated guaranteed value to Agent:

$$V_A(w|Q^A) = \mathbb{E}_{q_A^*(w|Q^A)} w(y) - g(a^*(w|Q^A)) \quad (15)$$

Relative to the information Q^P , Principal's value :

$$V_P(w|Q^P) = \min_{p \in Q^P(a^*(w|Q^P))} \mathbb{E}_p[y - w(y)] \quad (16)$$

Notice here that we now have a possible discrepancy between $a^*(w|Q^A)$ and $a^*(w|Q^P)$, as Principal can only use information available to him. The key point here is that any action chosen by Agent would guarantee at least as much expected payoff to him as $a^*(w|Q^P)$ does.

Now, given that $Q^A(a) \subset Q^P(a)$ for all $a \in \mathcal{A}$, Principal's worst-case expected payoff over all information mappings Agent might possess,

$$V_P(w) := \inf_{Q \subset Q^P} V_P(w|Q) \leq V_P(w|Q^P) \leq V_P(w|Q^A)$$

The last two terms refer to the cases when Principal has information $Q = Q^P$ and $Q = Q^A$ respectively, and the second inequality follows as $Q^A(a) \subset Q^P(a)$ for all a , since more precise information would improve the guaranteed value.¹²

The first inequality follows from the fact that for any given contract Principal cannot infer precisely Agent's decision rule and has *endogenous* ambiguity about the induced set of probabilities over outputs. This observation plays a key role in the determination of Principal's guarantee below.

Note also that whatever optimal action a^* Agent chooses, with the associated 'best worst-case' for Agent $q_A^* \in Q^A(a^*)$, Agent's expected payments recieved satisfies

$$E_{q_A^*}[w(y)] \geq E_{q_A^*}[w(y)] - g(a^*) = V_A(w|Q^A) \geq V_A(w|Q^P). \quad (17)$$

The last term denotes Agent's guaranteed value if he were to possess the coarser information that Principal possesses; Principal in fact would compute Agent's guaranteed value this

¹²Formally, $V_P(w|Q^A)$ is the minimal value over smaller sets of distributions than $V_P(w|Q^P)$.

way. The second inequality holds because having more precise information can only make a pessimistic Agent (weakly) better off.¹³

The above analysis shows that in case of asymmetric ambiguity, Principal faces a substantial inference problem in anticipating ex-ante agent's optimal action choice, and also has a more pessimistic estimation of his maxmin guarantee than he would have if had the same information that Agent does. Reduction of this disadvantage is where linear contracts pay their role.

4.1. Linear Contracts. Suppose Principal offers a linear contract: $\ell(y) = \alpha y$ with $\alpha \in (0, 1]$.

For the linear contract Principal's ex-post payoffs are:

$$y - w(y) = \frac{1 - \alpha}{\alpha} w(y). \quad (18)$$

Combining this with (17) gives a lower bound on Principal's expected payoff

$$\mathbb{E}_{q_P^*}[y - \ell(y)] = \frac{1 - \alpha}{\alpha} \mathbb{E}_{q_P^*}[\ell(y)] \geq \frac{1 - \alpha}{\alpha} E_{q_A^*}[\ell(y)] \geq \frac{1 - \alpha}{\alpha} V_A(w|Q^P), \quad (19)$$

where q_P^* is the worst-case scenario for Principal that minimizes expected profit $\mathbb{E}_q[y - \ell(y)]$ over the set of probability distributions in $Q^A(a^*)$.

Using this observation for a linear contract, we then derive a relationship between the guarantees to Principal and to Agent. Since (19) holds regardless of the technology Q^A available to Agent, taking infimum over all such technologies, Principal's worst-case expected value satisfies:

$$V_p(\ell) \geq \frac{1 - \alpha}{\alpha} V_A(w|Q^P).$$

This shows how to obtain a payoff guarantee from a linear contract. Turns out, the optimal guarantee to Principal comes from a linear contract:

Theorem 2. *Suppose $Q^A(a) \subset Q^P(a)$ for all $a \in \mathcal{A}$. There exists a linear contract that maximizes V_p .*

¹³Again, the comparison is between minima over smaller as opposed to bigger sets.

Proof. The analysis of the contract design is driven by Principal's coarse information. For a given contract, unlike in the symmetric case Principal is not able to uniquely infer Agent's decision rule. Instead, Principal can infer the set of "rationalizable" actions that the Agent might choose consistent with Principal's coarse information. The proof of this Theorem is analogous to that of Theorem 1, with some subtle adjustments. We begin with a series of Lemmas.

Let $q_A^*(w|Q^P) \in Q^P(a^*(w|Q^P))$ be Agent's 'best worst-case' distribution facing w from Principal's perspective, that is consistent with her coarse understanding according to Q^P . The associated guaranteed value to Agent:

$$V_A(w|Q^P) = \mathbb{E}_{q_A^*(w|Q^P)} w(y) - g(a^*(w|Q^P)) \quad (20)$$

Step 1. Construction of Agent's Supporting Hyperplane

Using Lemma 3 with Q^P in the role of Q the guarantee value of $V_A(w|Q^P)$ Agent's problem defined in (20) yields a supporting hyperplane for $\mathcal{S}_{a^*}(w|Q^P)$. Analogous to (2) in Lemma 3 let $\widehat{H}_{A,w}$ be the hyperplane defined as

$$\widehat{H}_{A,w} = \{(u, v) \in \mathbb{R}^2 | v = V_A(w|Q^P)\} \quad (21)$$

It exposes the face of $\mathcal{S}_{a^*}(w|Q^P)$ with $\mathbb{E}_{q_A^*(w|Q^P)} w(y) = V_A(w|Q^P) + g(a^*)$ and is thereby a supporting hyperplane for $\mathcal{S}_{a^*}(w|Q^P)$.

Step 2. Construction of Principal's Supporting Hyperplane

From Principal's perspective the set of rationalizable actions $\hat{\mathcal{A}}$ are those that she can rationalize as optimal choice by Agent on the basis of his relatively more precise information, which can take arbitrary forms consistent with Principal's information. Formally

$$\hat{\mathcal{A}} = \{a \in \mathcal{A} : \mathcal{S}_a(w|Q^P) \cap V_A(w|Q^P)^+ \neq \emptyset\} \quad (22)$$

where $V_A(w|Q^P)^+ = \{(y, v) \in \mathcal{C}(w) : v \geq V_A(w|Q^P)\}$.

Relative to the information Q^P , Principal's guarantee value is the worst-case value from the set of rationalizable actions $\hat{\mathcal{A}}$:

$$V_p(w|Q^P) = \min_{a \in \hat{\mathcal{A}}} \min_{p \in Q^P(a)} \mathbb{E}_p[y - w(y)] \quad (23)$$

Let \hat{a} be the action in $\hat{\mathcal{A}}$ that attains the minimum. Recalling the definition of the hyperplanes for Principal's isoprofit curves in (6), Principal's guarantee in the case of asymmetric perceptions $V_p(w|Q^P)$ is alternatively characterized as:

Lemma 7. $\widehat{H}_{P,w}(V_p(w|Q^P))$ is a supporting hyperplane of $\mathcal{S}_{\hat{a}}(w|Q^P)$.

Proof follows from Lemma 4 taking Q^P and \hat{a} in the roles of Q and a^* , respectively.

Proof. Consider the family of all half-spaces $\{(u, v) \in \mathbb{R}^2 | u - v \geq \lambda\}$ that contains $\mathcal{S}_{\hat{a}}(w|Q^P)$. The half-space generated by the intersection of this family of half-spaces corresponds to the smallest λ for which $\mathcal{S}_{\hat{a}}(w|Q^P)$ is fully contained in it. The defining hyperplane has $\lambda = V_p(w|Q^P)$, and is by construction a supporting hyperplane for $\mathcal{S}_{\hat{a}}(w|Q^P)$. \square

Step 3. 2-dimensional affine cone generated by $\widehat{H}_{A,w}$ and $\widehat{H}_{P,w}$ hyperplanes for $Q = Q^P$ and the relationship between the values.

Observe that from the definition of Principal's value in (23) as the minimum value over the actions in $\hat{\mathcal{A}}$ the half-space defined by the hyperplane $\widehat{H}_{P,w}(V_p(w|Q^P))$ supports $\mathcal{S}_{a^*}(w|Q^P)$ and contains all of the $\mathcal{S}_a(w|Q^P)$ for all $a \in \hat{\mathcal{A}}$.

Let \hat{a}^* be the optimal action that Agent with more precise information chooses. Since his optimal choice gives Agent at least $V_A(w|Q^P)$ consistent with coarse information, the optimal choice is amongst the rationalizable actions, that is $\hat{a}^* \in \hat{\mathcal{A}}$.

We now have an affine set containing all $(\mathbb{E}y, \mathbb{E}w(y))$ pairs that would be give Agent the value at least $V_A(w|Q^P)$ from Step 1 while guaranteeing Principal at least $V_P(w|Q)$ from Step 2, and also a relationship between $V_A(w|Q)$ and $V_P(w|Q)$ for any $Q \subset Q^P$ under any contract $w(y)$ that induces a rationalizable action $\hat{a}^* \in \hat{\mathcal{A}}$ as the optimal action.

Lemma 8. $\mathcal{S}_{\hat{a}^*}(w|Q)$ is contained in the 2-dimensional affine convex cone generated by $\widehat{H}_{A,w}$ and $\widehat{H}_{P,w}$ hyperplanes.

Proof. $\widehat{H}_{A,w}$ and $\widehat{H}_{P,w}$ are supporting hyperplanes of $\mathcal{S}_{a^*}(w|Q^P)$ and $\mathcal{S}_{\hat{a}}(w|Q^P)$, respectively; the intersection of the two associated half-spaces that both contain $\mathcal{S}_{\hat{a}^*}(w|Q^A)$ is an affine convex cone. \square

As an immediate implication of Step 3 we have the following relationship between the values from any given contract:

Lemma 9. For any contract $w(y)$ that induces \hat{a}^* as the optimal action,

$$V_P(w|Q) = \min_{q \in Q(\hat{a}^*)} \mathbb{E}_q[y - w(y)] \quad \text{subject to} \quad \min_{q \in Q(\hat{a}^*)} [\mathbb{E}_q w(y) - g(a^*)] \geq V_A(w|Q)$$

Proof. By construction, for any contract that induces Agent to choose \hat{a}^* as the optimal action amongst the rationalizable actions $\hat{\mathcal{A}}$, $\mathcal{S}_{\hat{a}^*}(w|Q)$ contains all feasible $\{\mathbb{E}y, \mathbb{E}w(y)\}$ pairs. By definition, $V_P(w|Q)$ and $V_A(w|Q)$ are the worst-cases

respectively for Principal and Agent over this feasible set with expectations taken over their respective worst-case distributions. \square

We note that the existence of the solution to the problem specified in Lemma 9 follows from the analogous arguments given for the existence of the solution to the problem in Lemma 6.

Next we establish that a linear contract can be found that will induce a^* as a rationalizable action, while making Principal and Agent both as at least as well-off as under $w(y)$. In fact, within the family of linear contracts, there is an optimal one. We develop the analogous steps as in Theorem 1 to conclude.

- (a') Let a^* be the optimal action chosen by Agent under $w(y)$ according to Principal's coarse information Q^P , that provides a guarantee $V_P(w|Q^P) = \lambda$ to Principal. Let $\gamma(a^*)$ be the associated $\mathbb{E}_{q^*} w(y)$. First we establish that there is an affine contract that still rationalize a^* and guarantees Principal at least λ .

Construct the proposed affine contract as follows:

- i.' If $V_P(w|Q^P) < 0$, choose $w(y) = 0$
- ii.' If $V_P(w|Q^P) \geq 0$, let $w'(y) = \beta + \alpha y$, with $\alpha \in (0, 1)$, and $\beta = (1 - \alpha)\gamma - \alpha\lambda \geq 0$

It is straightforward to check that the line $\beta + \alpha y$ passes through the vertex of the affine convex cone generated by $\widehat{H_{A,w}}$ and $\widehat{H_{P,w}}$ hyperplanes. In light of the analysis in Step 3 the cone contains $\mathcal{S}_{a^*}(w|Q^P)$, and thereby has non-empty intersection with $\mathcal{S}_{a^*}(w|Q^P)$. In fact, the line segment defined by $gr(w'(y)) \cap \mathcal{S}_{a^*}(w|Q^P)$ represents the set $\{(\mathbb{E}_q y, \mathbb{E}_q w'(y)) | q \in Q^P(a^*)\}$. Recall also that any of these pairs $(\mathbb{E}_q y, \mathbb{E}_q w'(y))$ fall within the region representing both $y - w(y) \geq \lambda$ and $w(y) \geq \gamma$ [see Lemma 8]. In other words, a^* is still a rationalizable choice for Agent, while Principal is at least as well-off as before.

To finish this step of the argument, as in the symmetric case we consider Agent optimal choice \widehat{a}^* from his rationalizable set of actions in $\widehat{\mathcal{A}}$, which might be different from a^* since Agent's precise information consistent with $Q^A \subset Q^P$ can take arbitrary forms.

To check for any possibly different optimal choice from a^* by Agent under the new contract $w'(y)$. By analogous 'revealed preference' argument given in the symmetric case, we first note that Agent would only choose an action in the set of rationalizable actions, $\widehat{\mathcal{A}}$: Agent would not choose any action a' with $\min_{q \in Q(a')} \mathbb{E}_q w'(y) < \min_{q \in Q(a^*)} \mathbb{E}_q w(y)$. We must have, if Agent were to choose \widehat{a}^* in $\widehat{\mathcal{A}}$ under w' , for some $Q \subset Q^P$

$$\min_{q \in Q(\widehat{a}^*)} \mathbb{E}_q w'(y) \geq \min_{q \in Q(a^*)} \mathbb{E}_q w'(y) \geq \min_{q \in Q(a^*)} \mathbb{E}_q w(y) \quad (24)$$

Hence we conclude that with this new affine contract w' , Agent either still chooses a^* , or instead optimally chooses a different action \hat{a}^* with $\min_{q \in Q(\hat{a}^*)} \mathbb{E}_q w'(y) \geq \min_{q \in Q(a^*)} \mathbb{E}_q w'(y)$. Given the affine nature of w' , any such deviation would also improve Principal's worst-case. We have,

$$y - w'(y) = \frac{1 - \alpha}{\alpha} w'(y) - \frac{1 + \alpha}{\alpha} \beta \quad (25)$$

Taking expectations on both sides, and then taking minima, establishes the last claim. Since this is true for any a^* , we have found an affine contract that dominates $w(y)$.

- (b') We next show that, for any affine contract $w'(y)$, there is a linear contract $\ell(y)$ that does at least as well, or strictly better, for Principal, and still rationalize the same action as $w'(y)$.

If $\beta = 0$, then $w'(y)$ is already linear and we are done. If $\beta > 0$, then note that $w'(0) = \beta$. Define $\ell(y) = \alpha y = w'(y) - \beta$. This improves Principal's payoff and leaves Agent's incentives unchanged.

- (c') The previous steps, taken together, imply that any action that is rationalized by an arbitrary contract, $w(y)$, can be also rationalized by an appropriately specified linear contract, that makes Principal and Agent both at least as well off as under $w(y)$. In other words, it is enough for Principal to only optimize within the class of linear contracts.

From (a'), we can strengthen (19) to equality:

$$\mathbb{E}_{q_P^*}[y - \ell(y)] = \frac{1 - \alpha}{\alpha} \mathbb{E}_{q_P^*} \ell(y) = \frac{1 - \alpha}{\alpha} \mathbb{E}_{q_A^*} \ell(y) \quad (26)$$

This follows from noting that with a linear contract, for any action a , Principal and Agent agree on the worst-case $q_P^* = q_A^*$, even if Principal is not able to perfectly infer \hat{a}^* the linear contract ensures that Agent's optimal action choice maximizes Principal's guarantee.

In other words, with a linear contract, the best worst-case guarantees for Principal and Agent are related via a linear relationship described in (26). Now using this linear relationship we establish the existence of optimal contract in linear form $\ell(y) = \alpha y$. From the linear contract αy Principal's guarantee value is $V_P(\alpha) = \max_{a \in \mathcal{A}} [(1 - \alpha) E_{q^*(a)}[y] - \frac{1 - \alpha}{\alpha} g(a)]$. Continuity of V_P in α follows from the Theorem of Maximum by taking α in the role of the parameter, and an $\alpha^*(a^*)$ exists in $[0, 1]$ that maximizes it.¹⁴

□

¹⁴There are some subtleties regarding $\alpha = 0$, for which we provide more details below in Section 4.2.

4.2. Optimal Linear Contract. We have established the optimality of linear contracts. Next we identify the optimal linear contract. Using Theorem 2 and especially the linear relationship between the values in (26) the optimal share maximizes

$$\max_{a \in \mathcal{A}} \left[(1 - \alpha) E_{q^*(a)}[y] - \frac{1 - \alpha}{\alpha} g(a) \right] \quad (27)$$

jointly over $a \in \mathcal{A}$ and $\alpha \in [0, 1]$. Here we use the observation that $q_p^*(a)$, the lower envelope of $Q^P(a)$ is the common worst-case that minimizes the expected output in that set. Our approach is to find the optimal α for any given action a , then find the maximum over \mathcal{A} . If the worst-case expected output from an action a is not high enough to cover its effort cost, that is if $E_{q_p^*(a)}[y] < g(a)$, then the optimum α is zero, corresponding to zero contract. Otherwise, the optimal share α for action a that solves (27) is equal to $\alpha(a) = \sqrt{g(a)/E_{q_p^*(a)}[y]}$.

Using $\alpha(a)$ in (27) and by continuity of the objective function in a over a compact \mathcal{A} we solve for the optimal action a^* and pick as the share $\alpha^* = \alpha(a^*) = \sqrt{g(a^*)/E_{q^*(a^*)}[y]}$. Notice that as Principal's perception of technology becomes more ambiguous $Q^P(a) \subset \hat{Q}^P(a)$ the expected worst-case becomes worse $E_{\hat{q}^*(a^*)}[y] \leq E_{q^*(a^*)}[y]$ and the share of output paid to Agent increases (this observation is true for *uniform* increase in ambiguity, i.e. when $Q^P(a)$ becomes larger for each a).

5. AMBIGUITY AND RISK AVERSE AGENT

We extend our analysis of the contracting problem with two-sided symmetric ambiguity to allow risk-averse attitudes on the part of the the agent. Let the agent's risk attitudes be represented with an increasing and concave vNM utility function u .

From an arbitrary contract w , the agent's guaranteed expected utility, or his MaxMin expected utility, is given by

$$v_A^u := V_A^u(w|Q) = E_{q_A(w|Q)}[u(w(y))] - g(a^*(w|Q))$$

where $a^*(w|Q)$ is the agent's MaxMin optimal action:

$$a^*(w|Q) = \operatorname{argmax}_{a \in \mathcal{A}} \min_{q \in Q(a)} \mathbb{E}_q[u(w(y)) - g(a)],$$

and $q_A(w|Q) \in Q(a^*(w|Q))$ is the associated worst-case distribution perceived by the agent. Since the principal and the agent have symmetric perception of ambiguity for each action, the principal infers uniquely the agent's decision rule for his MaxMin optimal action, and

her principal's guaranteed value from the contract w is given by

$$v_P^u := V_P^u(w|Q) = \min_{q \in Q(a^*(w|Q))} E_q[y - w(y)]$$

His associated worst-case scenario is $q_P(w|Q) \in Q(a^*(w|Q))$.

The principal offers a contract w that maximizes

$$V_P^u(w|Q) \text{ subject to } V_A^u(w|Q) \geq 0$$

We assume that contracting relation is viable: for some contract w , we have that $V_P^u(w|Q) \geq 0$ and $V_A^u(w|Q) \geq 0$

Theorem 3. Suppose $Q^A(a) = Q^P(a)$ for all $a \in \mathcal{A}$. There exists an affine contract that maximizes V_P .

This result is similar to that obtained in Theorem 1 except that with risk-averse agent, affine contracts do the job.

Proof. Fix a viable contract w , and let a^* be the associated optimal action choice, its dependence on the contract w and the nature of ambiguity Q is suppressed to ease notation. We make the argument in several steps.

- (M-1) To represent the set of possible values from an arbitrary contract, we use the space of output and utility level pairs (instead of output-payment space used in the analysis with the risk-neutral agent) in the analysis of the optimal contract. On this space we start with observing that the agent's MaxMin optimal choice $a^*(w|Q)$ induces a compact convex set denoted by $\mathcal{S}_{a^*}^u(w|Q)$ analogous to that in Section 3.
- (M-2) In the output-utility space, the agent's indifference curves are represented by horizontal lines, as in the analysis with risk-neutral agent. Therefore, Lemma 1 through Lemma 3 applies analogously in this space and from the contract w the agent's guaranteed value $V_A^u(w|Q)$ defines a supporting hyperplane for the closed convex set $\mathcal{S}_{a^*}^u(w|Q)$.
- (M-3) In the output-utility space the principal's iso-profit curves are represented by concave curves that are not linear, unlike in the analysis of the contracting problem with risk-neutral parties. The latter implies that Lemma 4 does not directly apply in the current model. However, we first observe below that the principal's better than sets are closed and convex due to continuity and concavity of the agent's utility function. It in turn allows us to extend Lemma 4 and characterize with Lemma 10 'supporting'

epigraph for $\mathcal{S}_{a^*}^u(w|Q)$ representing the principal's guaranteed value associated with $Q(a^*)$. The supporting epigraph is well-defined as the smallest epigraph that contains the set $\mathcal{S}_{a^*}^u(w|Q)$.

More formally, for any given profits levels $k \geq 0$ let the set $H_P^u(k)$ denote the locus of output-utility pairs that gives the principal profits k :

$$H_P^u(k) = \{(y, v) : v = u(y - k)\}.$$

Notice that $H_P^u(k)$ in the output-utility space graphs utility values that yields profit level k for the principal and it is obtained via a right-shift in the agent's utility function u by amount k .

Let $\mathcal{B}(H_P^u(k))$ denote the hypograph which is the better-than set of $H_P^u(k)$ for the principal and it is defined by:

$$\mathcal{B}(H_P^u(k)) = \{(y, v) : v \leq u(y - k)\}. \quad (28)$$

For any viable contract this set is non-empty as it contains zero, and it is closed and convex by continuity and concavity of u , respectively. Moreover, since utility function is monotone, the better-than-sets shrink as the principal's guaranteed value increases: for $\hat{k} < k$, $\mathcal{B}(H_P^u(k)) \subset \mathcal{B}(H_P^u(\hat{k}))$, and this implies the following property

Lemma 10. Hypograph $\mathcal{B}(H_P^u(k))$ at $k = V_P^u(w|Q)$ supports $\mathcal{S}_{a^*}^u(w|Q)$.

Proof. Viability of the contract w implies that $\mathcal{B}(H_P^u(0))$ fully contains all of $\mathcal{S}_{a^*}^u(w|Q)$. For otherwise, the principal would expect to make a loss from the contract and instead she would rather offer a 'zero-contract' that ensures a non-negative guarantee.

Now notice that the hypograph in the statement of the lemma is the intersection of all $\mathcal{B}(H_P^u(\hat{k}))$ such that $\hat{k} < k$ and that all fully contain $\mathcal{S}_{a^*}^u(w|Q)$. Since hypograph $\mathcal{B}(H_P^u(k))$ is a closed convex set and $\mathcal{B}(H_P^u(k)) \subset \mathcal{B}(H_P^u(\hat{k}))$ for $0 \leq \hat{k} < k$, the defining hypograph $\mathcal{B}(H_P^u(k))$ at $k = V_P^u(w|Q)$ supports $\mathcal{S}_{a^*}^u(w|Q)$. \square

We next characterize a closed and convex set \mathcal{E} containing all $(\mathbb{E}[y], \mathbb{E}[u(w(y))])$ pairs that yields for the principal at least her original guaranteed value $V_P(w|Q)$, and for the agent at least $V_A(w|Q)$. This set is analogous to the convex cone constructed in Lemma 5 and it contains all of $\mathcal{S}_{a^*}^u(w|Q)$, and more specifically

Lemma 11. The closed convex set $\mathcal{S}_{a^*}^u(w|Q)$ is contained in the 2-dimensional closed convex set \mathcal{E} with its boundaries generated by hyperplane $H_A^u(\lambda)$ at $\lambda = V_A^u(w|Q)$ and curve $H_P^u(k)$ at $k = V_P^u(w|Q)$.

Proof. The horizontal line $H_A^u(\lambda)$ at $\lambda = V_A^u(w|Q)$ is a supporting hyperplane of $\mathcal{S}_{a^*}^u(w|Q)$ as we have shown by analogous arguments as in Lemma 3; and the curve

$H_P^u(k)$ at $k = V_P^u(w|Q)$ supports $\mathcal{S}_{a^*}^u(w|Q)$ by Lemma 10. The defining set \mathcal{E} is determined by taking the intersection of the two better-than-sets that are both closed convex sets and both contain $\mathcal{S}_{a^*}^u(w|Q)$: one for the agent that is epigraph of the horizontal line $H_A^u(\lambda)$ at $\lambda = V_A^u(w|Q)$ and the other for the principal that is hypograph $\mathcal{B}(H_P^u(k))$ at $k = V_P^u(w|Q)$. By construction, the set \mathcal{E} is bounded below by the horizontal line and above by the curve. \square

As an immediate implication we have the following:

Lemma 12. For any contract $w(y)$ that induces a^* as the optimal action,

$$V_P^u(w|Q) = \min_{q \in Q(a^*)} \mathbb{E}_q[y - w(y)] \quad \text{subject to} \quad \min_{q \in Q(a^*)} \mathbb{E}_q[u(w(y)) - g(a^*)] \geq V_A^u(w|Q)$$

Proof. By construction, for any contract that induces a^* as the optimal action, the closed convex set $\mathcal{S}_{a^*}^u(w|Q)$ contains all feasible expected output and expected utility pairs $\{\mathbb{E}_q y, \mathbb{E}_q u(w(y))\}$ induced by this action a^* . Over this feasible set, by definition, $V_P^u(w|Q)$ and $V_A^u(w|Q)$ are the worst-case expected values for the principal and the agent, respectively, with expectations taken over with respect to their respective worst-case distributions. \square

Note that the solution to the problem specified in Lemma 12 exists, given a compact \mathcal{Y} and bounded $w(y)$. This is because, it is a linear programming problem in the $(y, u(w(y)))$ space that is compact convex by construction.

(M-4) We next show that for an arbitrary contract w we can construct an affine contract that ensures the parties higher guaranteed values than the contract w does.

In what follows, we construct an affine contract $\ell \in \mathcal{L}$ such that from the action a^* the principal and the agent both obtain guaranteed values at least their respective original guarantees $V_P^u(w|Q)$ and $V_A^u(w|Q)$. Furthermore, if the agent chooses a different action, that improves the guaranteed values for both of them.

For this, we start with the affine contract $\ell_0(y) = \beta_0$ that makes a fixed payment $\beta_0 := u^{-1}(V_A^u(w|Q))$ to the agent and nothing else. This contract has the property that its set of expected output and utility pairs $\mathcal{S}_{a^*}^u(\ell_0|Q) := \{(\mathbb{E}_q[y], V_A^u(w|Q)) : q \in Q(a^*)\}$ is fully contained in the set \mathcal{E} given in Lemma 11. The latter implies that the contract ℓ_0 yields the agent his original guarantee since $u(\beta_0) = V_A^u(w|Q)$, and yields the principal better guarantee $V_P^u(\ell_0|Q)$ than his original $V_P^u(w|Q)$ if the agent still chooses the action a^* given the contract ℓ_0 . Since the constant wage is independent of the output, and hence it does not necessarily incentivize the agent to choose that action, for instance, the agent can deviate and choose a less costly action than a^* that could reduce the principal's guarantee. However, by offering an affine contract

with $\alpha > 0$ and the same fixed payment β_0 , the action a^* guarantees strictly better value for the agent than $V_A^u(w|Q)$. Furthermore, the principal can do better using an affine contract.

From any affine contract $\ell = \alpha y + \beta$ the guaranteed values for the principal and the agent $V_P^u(\ell|Q)$ and $V_A^u(\ell|Q)$, respectively, are both continuous in α and β ; and the former is monotone decreasing in the parameters of affine contract while the latter is monotone increasing. Now for an affine contract $\hat{\ell}$ with its $\hat{\beta}$ is sufficiently close to β_0 from below, and its $\hat{\alpha}$ is close enough to zero from above, the principal and the agent both obtains better guarantees if the agent chooses action a^* : $V_P^u(\hat{\ell}|Q) > V_P^u(w|Q)$ and $V_A^u(\hat{\ell}|Q) > V_A^u(w|Q)$.

If, on the other hand, the agent deviates from a^* and chooses a different action under the affine contract ℓ , this improves the guarantees for both of them. To see this, Lemma 11 applied to the affine contract $\hat{\ell}$ implies that the set $S_{a^*}^u(\hat{\ell}|Q) := \{(\mathbb{E}_q[y], \mathbb{E}_q[u((1 - \hat{\alpha})y + \hat{\beta})]) : q \in Q(a^*)\}$ is fully contained in the closed convex set $\hat{\mathcal{E}}$ which is determined by the intersection of the hypograph of $u(\hat{\ell})$ and the epigraph of the hyperplane $H_A^u(\lambda)$ at $\lambda = V_A^u(\hat{\ell}|Q)$. Now, under the affine contract $\hat{\ell}$ if the agent deviates from a^* and chooses a different action, say \hat{a} , the corresponding set $S_{\hat{a}}^u(\hat{\ell}|Q)$ is fully contained in the closed convex set $\hat{\mathcal{E}}$ and increases the guarantees of both parties.

(M-5) The arguments thus far have shown that whenever a contract w generates a two-dimensional set expected output and utility pairs $S_{a^*}^u(w|Q)$ then there is an affine contract that does better. In the output-utility space, one-dimensional set of expected output and utility pairs arises only from the contracts that are linear in utility: that is, there are the contracts w of the form that $u(w(y)) = ay$, or equivalently, $w(y) = u^{-1}(ay)$. We next show that affine contracts ℓ does better than utility-linear contracts, mainly because the former provides with the principal a flexible contracting set to smooth the consumption profile of the agent and hence minimize on the agency costs.

Consider an arbitrary utility-linear contract w that includes a^* as the optimal action. For such a contract the agent and the principal agree on the worst-case scenario: the distribution q_* in $Q(a^*)$ that minimizes the expected output. The fact that distribution minimizes the agent's expected value follows from utility-linearity of the contract. On the other hand, the principal's worst-case perception is determined as follows

$$\min_{q \in Q(a^*)} \mathbb{E}_q[y - w(y)] = \mathbb{E}_{q^*}[y - w(y)].$$

Here the equality follows from the fact that $y - w(y)$ is strictly concave function as $w(y) = u^{-1}(y)$, and for this concave function between the two extreme points $q_* = \operatorname{argmin}_{q \in Q(a^*)} \mathbb{E}[y]$ and $q^* = \operatorname{argmax}_{q \in Q(a^*)} \mathbb{E}[y]$, the former attains the minimum value.

Now consider an affine contract $\hat{\ell}$ such that $\hat{\beta}$ arbitrarily close to and less than the agent's guarantee from w and $\hat{\alpha}$ close to zero, so that by continuity the affine contract $\hat{\ell}$ provides the same guarantee to the agent. Since the affine contract $\hat{\ell}$ provides the same guarantee to the risk-averse agent for smaller variability in the agent's utility levels, it saves the principal's cost of inducing the action a^* .

□

6. DISCUSSION

We illustrate optimality of linear contracts in a moral hazard problem when both parties have ambiguous perceptions of technology, and Principal faces more ambiguity than Agent. The main result derives from the ability of linear contracts to align the worst-case scenarios for both parties, thereby reducing the impact of the endogenous uncertainty of Agent's choice of action on Principal's payoff.

We conclude with a few remarks about our assumptions. It is important that Principal has coarser understanding of technology relative to Agent's. This assumption ensures that Principal's worst-case value, relative to her own coarser information is well-defined (it lies on the boundary of a certain convex subset of the convex hull of the graph of $w(\cdot)$). Our assumption on the coarseness of Principal's understanding allows for richness in Agent's action choice as perceived by Principal. In this sense, it generates in an endogenous manner Carroll (2015) richness condition on the set of actions available to Agent.

Thus, our analysis provides intuition as to why a similar argument for linearity may not work if the asymmetry of uncertainty is reversed, i.e., if Principal is better informed about technology than Agent (for example, when an experienced Principal hires a rookie Agent).¹⁵ In that case, Principal is not able to uniquely identify a lower bound on Agent's value for an arbitrary contract. Characterizing the optimal contract in such a case is our next step in this research project.

¹⁵See Lopomo et al. (2011) for a treatment of such a problem in a different formal setting; a two-part contract with a flat payment and a bonus is proposed.

REFERENCES

- Charalambos Aliprantis and Kim Border. *Infinite dimensional analysis*. Springer, 2006.
- Nemanja Antic. Contracting with unknown technologies. *unpublished paper, Princeton University*, 2014.
- Lawrence M Ausubel and Raymond J Deneckere. A generalized theorem of the maximum. *Economic Theory*, 3(1):99–107, 1993.
- Dirk Bergemann and Stephen Morris. Robust mechanism design. *Econometrica*, 73(6):1771–1813, 2005.
- Dirk Bergemann and Karl Schlag. Robust monopoly pricing. *Journal of Economic Theory*, 146(6):2527–2543, 2011.
- Gabriel Carroll. Robustness and linear contracts. *The American Economic Review*, 105(2):536–563, 2015.
- Sylvain Chassang. Calibrated incentive contracts. *Econometrica*, 81(5):1935–1971, 2013.
- Peter Diamond. Managerial incentives: on the near linearity of optimal compensation. *Journal of Political Economy*, 106(5):931–957, 1998.
- Alexander Frankel. Aligned delegation. *The American Economic Review*, 104(1):66–83, 2014.
- Daniel F Garrett. Robustness of simple menus of contracts in cost-based procurement. *Games and Economic Behavior*, 87:631–641, 2014.
- Paolo Ghirardato. Agency theory with non-additive uncertainty. *Unpublished working paper, Cal. Tech*, 1994.
- Sanford J Grossman and Oliver D Hart. An analysis of the principal-agent problem. *Econometrica: Journal of the Econometric Society*, pages 7–45, 1983.
- Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Fundamentals of convex analysis*. Springer Science & Business Media, 2012.
- Bengt Holmstrom and Paul Milgrom. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica: Journal of the Econometric Society*, pages 303–328, 1987.
- Leonid Hurwicz and Leonard Shapiro. Incentive structures maximizing residual gain under incomplete information. *The Bell Journal of Economics*, pages 180–191, 1978.
- Edi Karni. Subjective expected utility theory without states of the world. *Journal of Mathematical Economics*, 42(3):325–342, 2006.
- Edi Karni. A reformulation of the maxmin expected utility model with application to agency theory. *Journal of Mathematical Economics*, 45(1):97–112, 2009.
- Giuseppe Lopomo, Luca Rigotti, and Chris Shannon. Knightian uncertainty and moral hazard. *Journal of Economic Theory*, 146(3):1148–1172, 2011.

Sujoy Mukerji. Ambiguity aversion and incompleteness of contractual form. *American Economic Review*, pages 1207–1231, 1998.

APPENDIX A. EXISTENCE OF SOLUTIONS TO AGENT’S OPTIMIZATION PROBLEM

Assume that the correspondence that maps a in compact set \mathcal{A} to $Q(a)$ in compact convex set of probabilities over output levels $\mathbb{K}_{\Delta(\mathcal{Y})}$ is upper hemi-continuous.

We want to show that for a given continuous contract w , the maxmin action $a^* \in \mathcal{A}$ is well-defined.

We start with characterizing Agent’s decision rule for the effort choice for a given contract. Notice that for each action a Agent effort choice solves:

$$V_A(a; w) = \min_{q \in Q(a)} E_q[w(y)] - g(a)$$

and the value function is continuous in a . To see this, consider effort level a in the role of a parameter in the constrained optimization problem. In particular, the objective function is continuous in effort a and the feasible set correspondence $Q(a)$ is uhc. Therefore, by the (generalized) Maximum Theorem from Ausubel and Deneckere (1993), that uses upper-hemicontinuity of the feasible set of alternatives, the worst-case value $V_A(a; w)$ is continuous in the parameter, or in this case the action a .

For the contract w Agent’s optimal effort $a^*(w)$ which maximizes his continuous value function $V_A(\cdot; w)$ over a compact set \mathcal{A} is well-defined by the Extreme Value Theorem (Aliprantis and Border (2006)).

By symmetry, Principal infers Agent’s decision rule a^* and if the optimal action is unique Principal uses it to compute his worst-value from the contract:

$$V_p(w) = \min_{q \in Q(a^*(w))} E_q[y - w(y)]$$

If Agent’s decision rule assigns multiple actions as optimal to a contract w , from such a contract Principal evaluates his guarantee according to the worst case. That is, if $A^*(w)$ is Agent’s set of optimizers, Principal’s guarantee considers the worst of these actions from her perspective:

$$V_p(w) = \min_{a \in A^*(w)} \min_{q \in Q(a)} E_q[y - w(y)]$$

Since $A^*(w)$ compact, being a closed subset of compact \mathcal{A} , Principal's guarantee is again well-defined. An alternative approach to the decision rules with multiple actions is that Agent when indifferent picks the best action for Principal. This alternative approach which is typically used in standard moral hazard problems is easy to adopt in our guarantee specification above by taking max over $A^*(w)$ instead of min, and without affecting the analysis and the results qualitatively.

DEPARTMENT OF ECONOMICS, UNIVERSIDAD CARLOS III DE MADRID, e-mail: mdumav@gmail.com, AND
DEPARTMENT OF ECONOMICS, UNIVERSITY OF CALIFORNIA AT RIVERSIDE, e-mail: urmeek@ucr.edu