MORAL HAZARD, UNCERTAIN TECHNOLOGIES, AND LINEAR CONTRACTS

MARTIN DUMAV, URMEE KHAN

Abstract. We analyze a moral hazard problem where both contracting parties face non-probabilistic uncertainty (ambiguity) about how actions translate to output and seek robust performance from a contract in relation to their respective worst-case scenarios. Linear contracts that align the Principal's and the Agent's pessimistic expectations are optimal.

Keywords: Moral hazard, non-probabilistic uncertainty, ambiguity aversion, linear contracts

JEL Classification: D81, D82, D86

1. Introduction

Consider a Principal who wishes to hire an Agent, but possesses only coarse information about relevant technology. In particular, she knows exactly the set of actions Agent might take, but has less precise knowledge about possible consequences of those actions than does Agent. Out-sourcing to an 'expert' is a natural example, where presumably Principal would be less knowledgable about the technology involved, nevertheless it could be contractually possible to restrict the allowable set of actions. What kind of a contract would Principal offer?

Carroll [2015] poses a related, but different question. In his setting, Principal is unsure about the possible actions that Agent might take, although she knows, and agrees with Agent on, the consequences of the actions as quantified in terms of an unique probability distribution on outcomes. A rationalization for linear contracts is provided, as Principal seeks guarantees robust to Agent's action set.

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 the Principal's, and that is all that Principal knows.

With our formulation of non-probabilistic beliefs, we establish a key linear relation between Principal's and Agent's payoffs analogous to Carroll [2015]. Building on Carroll [2015], we show that linear contracts provide a tight 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 infer

Date: September 21, 2017. The first author gratefully acknowledges the financial support from the Ministerio Economia y Competitividad (Spain) through grants ECO2014-55953-P and MDM 2014-0431.

uniquely Agent's action choice but has to allow for multiplicity of actions consistent with her rough information.

The main contribution of the paper is therefore to establish justification for linear contracts in an environment where even though the Principal can restrict the set of allowable actions, her ambiguous perception of consequences of those actions lead her to seek guarantees. In the process of doing so, we adapt the method developed in Carroll [2015]. The key additional steps is in allowing non-probabilistic beliefs on both parties, and in particular the Agent has more precise understanding of technology than the principal. This leads in an endogenous manner to a rich set of actions the agent can choose for a given contract offer by the Principal.

Our paper fits broadly in the literature on design of contracts in ambiguous environments where a desire for robust mechanisms motivates the use of simple contracts, starting with Mukerji [1998] and more recently, for instance, Chassang [2013], Antic [2014], Garrett [2014], and Carroll [2015]. Against the backdrop of this literature we investigate whether one can rationalize linear sharing rules in a moral hazard problem with ambiguity aversion on both sides.

Holmstrom and Milgrom [1987] and Diamond [1998] offer earlier foundations for linear contracts in Bayesian models. In these papers, the optimality of linear contracts are optimal because it aligns the Principal's and the Agent's objectives, Diamond [1998] was the first to offer this explanation. The same idea has been extended to models with MaxMin criteria in various frameworks including Hurwicz and Shapiro [1978], Carroll [2015], Chassang [2013], Antic [2014]. Against this backdrop, our contribution is that the linear contracts are optimal in moral hazard relations because they align the pessimistic expectations of the contracting parties faced with ambiguity about technology. Our contribution also relates to mechanism design with ambiguity, including recent examples, Bergemann and Morris [2005], Bergemann and Schlag [2011], Garrett [2014], Frankel [2014]. Our results provide tractable approach to model richer possibilities for uncertainty to complement those contributions.

Robustness of high powered incentives have been analyzed in specific agency relations. Mukerji [1998] shows that in a vertical relationship optimal contracts feature incompleteness and low powered incentives. It remains an open question to whether Mukerji [1998]'s result is general in agency relations. Our comparative static analysis suggest that ambiguity about technology can also lead to flattening of incentive contracts. 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. The main challenge is that ambiguity averse parties can disagree about the worst-case expectations which are endogenous objects. Expanding Carroll [2015]'s formulation of moral hazard in the space of output and payoffs, we analyze the moral hazard problem with ambiguity in a tractable manner. Our formulation of ambiguity in the outcome space provides traction at no loss of generality, as Karni [2006] provide a subjective foundation for such a formulation.

2. Model

Let $\mathcal{Y} \subseteq \mathbb{R}^k$ denote a compact convex set of outcomess, $\Delta(\mathcal{Y})$ 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 $\Delta(\mathcal{Y})$. Principal's information about technology is characterized as a single-valued mapping from a compact set of actions $\mathcal{A}$ to $\mathbb{K}_{\Delta(\mathcal{Y})}$ given by $Q^P(\cdot) : \mathcal{A} \mapsto \mathbb{K}_{\Delta(\mathcal{Y})}$. For each action $a \in \mathcal{A}$, Principal believes that this action induces a convex and compact set of distributions over outcomes denoted by $Q^P(a) \in \mathbb{K}_{\Delta(\mathcal{Y})}$. Agent's information is characterized by $Q^A(\cdot) : \mathcal{A} \mapsto \mathbb{K}_{\Delta(\mathcal{Y})}$. Agent has more precise knowledge about technology than Principal: $Q^A(a) \subset Q^P(a)$ for each $a \in \mathcal{A}$. Moreover, Principal does not know the exact $Q^A(a)$.

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 $\mathcal{Q}^A = \{\mathcal{Q}^A(a) : a \in \mathcal{A}\}$ describes Principal's perception about Agent's information about technological possibilities. $\mathcal{Q}^A$ thus contains the range of any possible mapping $Q^A(\cdot)$ that Principal might imagine for Agent to possess. Let $g : \mathcal{A} \mapsto \mathbb{R}^+$ be any continuous function that describes the cost of effort.

A contract is a continuous function $w : \mathcal{Y} \to \mathbb{R}^+$ that specifies output contingent payments and protects Agent with limited liability. We also assume that Principal and Agent can benefit from a contractual relationship, i.e. there exists eligible contracts: $V_P(w) > V_P(0)$.

We assume both Principal and Agent are ambiguity averse, they evaluate contracts by their worst-case expected payoff. Since expected utility is linear in probabilities, continuity of w and compactness of each set in $\mathbb{K}_{\Delta(\mathcal{Y})}$ ensure that these values and the worst-cases are well-defined for each $M \in \mathcal{Q}^A$.

3. Analysis

With a slight abuse of notation, we are going to use Q^i to denote the range of $Q^i(\cdot)$ in the following analysis.

Given an output contingent contract $w(y)$ a risk neutral, ambiguity averse Agent with technology $Q \subset \mathcal{Q}^A$ chooses an action that solves his optimization problem:

$$a^*(w|Q) = \operatorname*{argmax}_a \left[\min_{q \in Q(a)} E_q(w(y)) - g(a) \right]$$

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

$$V_A(w|Q) = E_{q^*_A(w|Q)} w(y) - g(a^*(w|Q))$$

Relative to the same technology Q, Principal's value :

$$V_p(w|Q) = \min_{p \in Q(a^*(w|Q))} E_p[y - w(y)]$$

Now, given that $Q^P \subset \mathcal{Q}^A$, Principal's worst-case expected payoff over all technologies possibly available to Agent,

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

The last two terms refer to the cases when both have common 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. 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.

Suppose Principal offers a linear contract: $w(y) = \alpha y$ with $\alpha \in (0, 1]$. Note 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 satisfies

$$E_{q_A^*}[w(y)] \geq E_{q_A^*}[w(y)] - g(a^*) = V_A(w|Q^A) \geq V_A(w|Q^P). \tag{1}$$

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

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

$$y - w(y) = \frac{1-\alpha}{\alpha} w(y). \tag{2}$$

Combining this with (1) gives a lower bound on Principal's expected payoff

$$E_{q_P^*}[y - w(y)] = \frac{1-\alpha}{\alpha} E_{q_P^*}[w(y)] \geq \frac{1-\alpha}{\alpha} E_{q_A^*}[w(y)] \geq \frac{1-\alpha}{\alpha} V_A(w|Q^P), \tag{3}$$

where q_P^* is the worst-case scenario for Principal that minimizes expected profit $E_q[y - w(y)]$ over the set of probability distributions in $Q^A(a^*)$.

The inequality relation (3) is a key step in our analysis. Here the first inequality follows because the worst case of expected payments *made* according to Principal's pessimism is more severe than Agent's conservative perception of worst-case expected payments *received*.[1] For a linear contract, this inequality holds as equality. To see this, let us consider the graph of any arbitrary contract, $w(y)$. Given any contract and for each action a, Agent's worst-case expected payoff lies on the horizontal hyperplane that supports the convex set $\mathcal{H}(a) = \{(E_q[y], E_q[w(y)] : q \in Q(a)\}$ from below (corresponding to Agent's lowest indifference curve). The Principal's worst-case distribution lies on the highest 45^0-degree hyperplane supporting $\mathcal{H}(a)$, representing Principal's iso-profit curve that minimizes expected profit within $\mathcal{H}(a)$. For an arbitrary contract, $\mathcal{H}(a)$ has possibly non-empty interior and hence

[1] For any a and Q, the Agent's worst case scenario minimizes expected payments received $q_A = \arg\min_{q \in Q(a)} E_q[w(y)]$ and equivalently

$$E_{q_A}[w(y)] \leq E_q[w(y)] \text{ for all } q \in Q(a).$$

In particular, this holds for $q_P \in Q(a)$ – the Principal's worst case minimizing expected profits $E_q[y - w(y)]$ over $Q(a)$.

Principal and Agent do not agree on the worst-case distribution. This divergence of worst-case scenarios is not beneficial for Principal. With a linear contract the common worst-case lies in the intersection of Agent's and Principal's worst-case hyperplanes and hence Principal and Agent agree on the worst-case distribution; for each a^*, $q_P^* = q_A^*$ in $Q^A(a^*)$. Using this observation for a linear contract, we then derive a tight relationship between the guarantees to Principal and to Agent. Since (3) holds regardless of the technology Q^A available to Agent, taking infimum over all such technologies, Principal's worst-case expected value satisfies:

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

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 1. *There exists a linear contract that maximizes V_p.*

To prove this Theorem, we begin with a series of Lemmas. Given any contract $w(y)$ Agent's guarantee from Principal's perspective:

$$V_A(w|Q^P) = \max_{a \in A} \left(\min_{q \in Q^P(a)} E_q[w(y)] - g(a) \right) \quad \text{(G)}$$

Principal's guarantee taking Agent's guarantee in (G) into account:

$$V_P(w) = \min E_q\big[y - w(y)\big] \quad \text{over } \big\{q \in \Delta(Y) : E_q[w(y)] \geq V_A(w|Q^P)\big\}$$

In this program, Principal considers a rich set of distributions that Agent's incentive compatible action choice can induce. Since Agent's information takes arbitrary forms consistent with $Q^A \subset Q^P$ for all a, any distribution induced by agent's optimal action would at least provide Principal the guarantee according to her coarse information. For any contract w, Principal therefore considers possible the subset of expected payoff pairs that is consistent with Agent's worst case value $V_A(w|Q^P)$. Principal's guarantee V_P is then the minimum of expected profits over that set. More formally:

Lemma 1. *Let w be any eligible contract, different from the zero contract. Then,*

$$V_P(w|Q^P) = \min E_q[y - w(y)] \text{ over } q \in \Delta(Y) \text{ s.t. } E_q[w(y)] \geq V_A(w|Q^P) \quad (*)$$

Moreover, for any q attaining the minimum, the constraint holds with equality.

Proof of Lemma 1. Consider any technology Q consistent with Principal's knowledge Q^P. Agent's payoff is at least $V_A(w|Q) \geq V_A(w|Q^P)$. Moreover, for his optimal action a^* Principal's and Agent's worst-case payments made and received, respectively, are related by the linear relation:

$$E_{q_P^*}[w(y)] \geq E_{q_A^*}[w(y)] \geq E_{q_A^*}[w(y)] - g(a^*) \geq V_A(w|Q^P)$$

The first inequality follows because the worst case for Principal's expected payments made is not better than Agent's conservative perception of the expected payments received, as we observed above. Here Principal and Agent potentially disagree on the worst case and this inequality can be strict. Principal's payoff $V_P(w|Q) = E_{q_P}[y - w(y)]$ is then at least the minimum given in $(*)$ since q_P^* is in the constrained set. Thus, Principal's worst-case payoff $V_P(w)$ is no lower than that given by $(*)$. Tightness of the inequality follows from the fact that according to Principal's perspective, multiplicity of Agent's optimal action covers the convex hull of the graph of $w(y)$, (i.e. pairs of expected output and expected wage payments from the contract w). □

Notice that for a linear contract the parties agree on the worst case $q_A^* = q_P^*$ and Lemma 1 gives a tight linear relation between the values:

Lemma 2. *For any $\alpha \in (0,1]$, if the linear contract $w(y) = \alpha y$ is eligible, then*

$$V_P(w) = \frac{1-\alpha}{\alpha} V_A(w|Q^P) = \max_{a \in A} \left[(1-\alpha) \min_{q \in Q^P(a)} E_q[y] - \frac{1-\alpha}{\alpha} g(a) \right]$$

With a tractable characterization of Principal's guarantee at hand one can proceed to a separation argument and the linearity of the optimal contract.

Proof of Theorem 1. Consider the following two sets for a separation argument: $S = co(\{(w(y), y - w(y)) : y \in Y\})$ and let T be the set of all pairs (u, v) such that $u > V_A(w|Q^P)$ and $v < V_P(w)$. Lemma 1 implies that S and T are disjoint. Applying Lemma 3 through Lemma 6 in Carroll [2015] establishes existence and optimality of the linear contracts. □

We next identify the optimal share in the optimal contract. Using Lemma 2, the optimal share maximizes

$$\max_a [(1-\alpha) E_{q_p^*(a)}[y] - \frac{1-\alpha}{\alpha} g(a)] \tag{4}$$

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. For any given interior action, the optimal α that maximizes $(1-\alpha)E_{q_p^*(a)}[y] - \frac{1-\alpha}{\alpha} g(a)$ is equal to $\alpha = \sqrt{g(a)/E_{q_p^*(a)}[y]}$. Using this in (4) we can solve for the optimal action a^* and pick as the share $\alpha^* = \sqrt{g(a^*)/E_{q^*(a^*)}[y]}$. Notice that as the Principal's perception of technology becomes more ambiguous $Q^P(a) \subset \hat{Q}^P(a)$ the expected wort-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.

4. Discussion and Work in Progress

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. In the process, we retain similar conclusions, while moving away from the assumption of common and unique probabilistic beliefs about the consequences of actions, as in [Carroll, 2015, §II]. We view our results to be complementary to those in Carroll [2015], as neither model nests the other.

We conclude with a few remarks about our assumptions. It is important that Principal has sufficiently coarser understanding of technology relative to Agent's. This assumption ensures in Lemma 1 that there is rich enough uncertainty in the set of distributions that Agent induces and hence Principal's worst-case value lies on the boundary of the convex hull of the graph of $w(\cdot)$. If instead we assume that Principal and Agent have completely symmetric understanding of the technology, then the worst-case value would not necessarily be on the boundary and the improving linear contract may not be well-defined. Our assumption on the coarseness of the Principal's understanding allows for richness in the Agent's action choice. 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). In that case, Principal is not able to uniquely identify a lower bound on the Agent's value for an arbitrary contract. Characterizing the optimal contract in such a case is our next step in this research project.

References

Nemanja Antic. Contracting with unknown technologies. *unpublished paper, Princeton University*, 2014.

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.

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. Agency theory with maxmin expected utility players. Technical report, mimeo, Johns Hopkins University, available at http://www. econ. jhu. edu/People/Karni, 2006.

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

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