CONTRACTING WITH UNKNOWN TECHNOLOGIES

NEMANJA ANTIĆ

MEDS, Kellogg School of Management, Northwestern University
nemanja.antic@kellogg.northwestern.edu

Abstract. I study contracting with moral hazard when the agent has available a known (baseline) production technology but the principal thinks that the agent may also have access to other technologies, and maximizes his worst-case expected utilities under those possible technologies. The nature of the Pareto efficient contract depends on the most unproductive distribution that the principal thinks might be available to the agent. When this lower-bound technology becomes trivial and all distributions are possible, equity is a Pareto efficient contract, generalizing existing work on robust contracting. When the lower-bound MLRP technology approaches the baseline technology, efficient contracts approach debt, providing robust foundations for debt in a classic financial contracting model. For intermediate lower-bounds, participating preferred equity contracts, mixtures of debt and equity are Pareto efficient for specific technology sets.

1. Introduction

Moral hazard, the idea that agents must be given appropriate incentives when their actions cannot be perfectly observed, occurs in a wide array of applications, e.g., insurance, franchising, employment contracts, unemployment benefits, CEO compensation, financial contracting, etc. The classic model of security design with moral hazard shows that debt contracts are efficient as long as contracts are restricted to be monotonic in cash flows (Innes, 1990)**?**. However, the efficient contract without this restriction is strikingly unrealistic[1]. The reason for this is aptly summarized by Holmström and Milgrom (1987)**?**:

Date: March 20, 2021.

Key words and phrases. Security design, Robust mechanism design, Maxmin principal.

I would like to thank Marco Battaglini, Sylvain Chassang, Olivier Darmouni, Peter Klibanoff, Stephen Morris and Andrew Newman for invaluable discussions and advice. Thank you to seminar audiences at Bocconi, Boston, Cornell, Chicago Booth, EIEF, Northwestern MEDS, Rochester, Utah and Vanderbilt for helpful comments. Thanks to Théo Durandard for outstanding research assistance.

[1]The Pareto optimal contract is a live-or-die contract, where the principal gets paid the entire profit up to some cut-off level (lives) and gets zero above that level (dies).

> "Real world incentive schemes appear to take less extreme forms than the finely tuned rules predicted by the basic theory... Agents in the real world typically face a wider range of alternatives and principals a more diffuse picture of circumstances than is assumed in the usual models."

Relative to the classical models, the present paper, building on the robust contracting literature started by Carroll (2015), relaxes the assumption that at the time of contracting the principal (or investor) knows exactly the set of technologies available to he agent (or entrepreneur) to convert effort into profits and thus faces a more diffuse picture of circumstances. A technology is a family of distributions, indexed by effort, which satisfies the monotone likelihood ratio property. I assume that the principal knows two things: (i) a specific 'baseline' technology[2] which will certainly be available to the agent, and (ii) a 'lower-bound' CDF, which is a worst-case CDF that's a starting point for all other available technologies. The principal evaluates other possible technologies with a maxmin criterion. I show that in this 'robust contracting' setting monotonic contracts emerge because the principal is concerned that the agent might have access to a technology that exploits any non-monotonicity. I show that debt is an efficient robust contract when the lower-bound technology is the same as the baseline technology. In this case, debt provides the best incentives for the agent to work hard by leaving all profits to him after a certain threshold.

However, I show that equity contracts are Pareto efficient when the principal fears that arbitrarily bad technologies could be realized. Intuitively, equity financing guarantees that the agent will not chose a technology that excessively hurts the principal, since the incentives of the two are perfectly aligned. In this case, maxmin considerations dominate the value of providing incentives, consistent with the results of the recent literature on robust contracting, e.g., Chassang (2013)**?** and Carroll (2015)**?**. In intermediate cases, a richer set of contracts, participating preferred equity (a mixture of debt and equity, which includes both as extremes) can occur. A familiar intuition from Innes is that debt contracts can be useful for incentive provision and thus pushing towards first-best effort levels. When the Knightian uncertainty of the principal is relatively small, many contracts could give the principal the required payoff and incentives for effort play a role. However, when robustness concerns are large, it is hard to look beyond the security that an equity contract provides.

The present paper makes several contributions to the literature. I first show that in a general robust contracting framework debt contracts can be efficient, in one extreme

[2]More generally, this could be a set of profit distribution and effort pairs.

case. As an intermediary step, I provide an ambiguity foundation for the monotonicity assumption commonly made in the security design literature.[3] Secondly, I show that in another extreme case equity is efficient, in line with the robust contracting literature (Chassang, 2013; Carroll, 2015). Third, I show that in intermediate environments, and for specific realizations of technology sets, efficient contracts take the form of participating preferred equity: a mixture of debt and equity, including both as special cases. The difference between the worst-case and baseline technology is the key simple parameter that determines whether the Pareto efficient contract is debt, equity or perhaps a mixture of the two. These results are due to the assumption that the principal has a non-trivial lower-bound technology and thus does not minimize over a set of distributions that is as rich as the majority of the literature focuses on.[4] Finally, I prove a technical result: in sufficiently rich maxmin contracting environments, it is without loss of generality to focus on contracts which are lower semicontinuous. There is no need for ex-ante restrictions on the set of allowable contracts;[5] and this technical result justifies the use of simple constructive techniques exhibited here.

The rest of the paper is organized as follows: Section 2 presents a simple example with an application to financial contracting; Section 3 defines the model and makes some remarks about the MLRP; Section 4 makes initial general observations which are applied throughout the analysis which follows; Section 5 considers the "smallest ambiguity" extreme case and shows the Pareto optimality of debt; section 6 considers the largest possible ambiguity case and shows the efficient contract is simple equity; Section 7 provides general results that encompass the preceding observations and shows that in general participating preferred equity is optimal; Section 8 concludes. **?**

2. Financial Contracting Example

To provide a preview of the results, we start with an example which contrasts the classic problem of contracting with moral hazard and the robust contracting model we study. The example highlights the Pareto efficiency of the live-or-die contract in the classic model

[3]This type of monotonicity assumption has been used by an array of authors, including DeMarzo & Duffie (1999)**?**, Matthews (2001)**?**, Biais & Mariotti (2005)**?**, DeMarzo (2005)**?**, DeMarzo, Kremer & Skrzypacz (2005)**?**, Inderst & Mueller (2006)**?**, Axelson (2007)**?**, Poblete & Spulber (2012)**?** and Dang, Gorton and Holmstrom (2012)**?**.

[4]Including Carroll (2015) and Barron et al. (2019)**?**, among others. This richness assumption is described in Walton and Carroll (2019) who show the optimality of linear contracts whenever it is satisfied. Of course, smaller ambiguity is very relevant and has been considered by a number of authors, albeit in very different settings; see, for example, Bergemann and Schlag (2011)**?**, Kos and Messner (2015)**?** and Madarász and Prat (2017)**?**.

[5]For example, Carroll (2015) assumes contracts are continuous.

and shows that in different robust contracting environments equity and debt contracts arise.

An entrepreneur (agent) with an idea for a project needs an investor (principal) to finance it. The investor will provide seed funding for the project, but it is up to the entrepreneur to put in effort to ensure the success of the project.

To fix ideas, suppose that the possible profit realizations are $\pi \in \{0, \underline{\pi}, \overline{\pi}\}$: either zero profit, low profit $\underline{\pi} = \frac{1}{2}$ or high profit $\overline{\pi} = 1$. Once a contract is agreed, the entrepreneur decides how much effort to put into the project, with more effort leading to better profit distributions. In particular, the entrepreneur can choose effort $e \in [0,1]$ and the technology available to him produces random profit given by the following family of probability distributions indexed by effort:

$$f_0(\pi \mid e) = \begin{cases} 1 - \frac{2}{3}e - \frac{1}{3}e^2 & \text{if } \pi = 0 \\ \frac{2}{3}e & \text{if } \pi = \underline{\pi} \\ \frac{1}{3}e^2 & \text{if } \pi = \overline{\pi} \end{cases}. \tag{2.1}$$

We shall sometimes write distributions in $\Delta(0, \underline{\pi}, \overline{\pi})$ as a triple. Observe that the above technology satisfies the monotone likelihood ratio property. Effort carries a utility cost for the entrepreneur, given by $c(e) = \frac{1}{2}e^2$.

A contract specifies how the profit is split between the investor and entrepreneur; it is a function, B, which maps profit realizations $\pi \mapsto B(\pi) \in [0, \pi]$ to the dollars of profit given to the investor.[6] The entrepreneur gets the remaining $\pi - B(\pi)$ dollars. In the example, a contract is just two numbers $\underline{b} = B(\underline{\pi}) \in \left[0, \frac{1}{2}\right]$ and $\overline{b} = B(\overline{\pi}) \in [0,1]$; if zero profit is realized neither party can get paid by the limited liability assumption. The goal is to identify the best contract for the entrepreneur, subject to the investor getting a required return of $\frac{1}{48}$.

Classic Model

In the textbook financial contracting model (Innes, 1990), the investor perfectly knows the agent's technology at the ex-ante contracting stage. In this case, the problem we want to solve is:

$$\begin{aligned} &\max_{\underline{b}, \overline{b}} \max_{e} \left\{ f_0(\underline{\pi} \mid e)\left(\frac{1}{2} - \underline{b}\right) + f_0(\overline{\pi} \mid e)\left(1 - \overline{b}\right) - \frac{1}{2}e^2 \right\} \\ &\text{s.t.} \quad f_0(\underline{\pi} \mid e^*)\,\underline{b} + f_0(\overline{\pi} \mid e^*)\,\overline{b} \geq \frac{1}{48}, \end{aligned}$$

[6]Observe that $B(\pi) \in [0, \pi]$, so that the entrepreneur cannot commit to paying the investor more than the entire profit and the investor is not liable for more than the initial investment.

where e^* is the entrepreneurs choice of effort for a given contract. In this example, for an arbitrary contract B, this is:

$$e(B) = \frac{1-2\underline{b}}{1+2\overline{b}}.$$

The above problem then becomes (after simplifying):

$$\max_{\underline{b},\overline{b}} \frac{(1-2\underline{b})^2}{6+12\overline{b}}$$
$$\text{s.t.} \quad \frac{\overline{b}+2\underline{b}-4\left(1+\overline{b}\right)\underline{b}^2}{3\left(1+2\overline{b}\right)^2} \geq \frac{1}{48}.$$

The solution to this program is $\underline{b} = \frac{1}{4} - \frac{\sqrt{3}}{8}$ and $\overline{b} = 0$. Under this contract, the agent's effort choice is $e(1/8, 0) = \frac{1}{2} + \frac{\sqrt{3}}{4} \simeq 0.933 < 1$, which is a little lower than the first-best level of effort $e^* = 1$. This contract induces a payoff of $\frac{1}{48}$ for the principal and approximately 0.145 for the agent.

This is a discrete version of what Innes (1990) terms a live-or-die contract: the investor gets all of the profit up to a point (lives) and gets paid nothing if profit is high (dies).

However, a live-or-die contract is sensitive to the investor knowing exactly the technology available to the entrepreneur: if the entrepreneur had a more productive technology than f_0, there would be a higher chance of $\overline{\pi}$ being realized, in which case the investor gets $\overline{b} = 0$ with a larger probability. These contracts are therefore not robust to assumption that the investor perfectly knows the production technology ex-ante.

Robust Financial Contracting

I study a model, based on Carroll (2015), in which the investor does not know the production technology available at the time of contracting. Instead assume the investor knows two things: (i) a 'baseline' technology, denoted f_0, and (ii) a 'lower-bound' CDF, denoted G. The baseline technology will be available to the agent for sure, but there could be other, unknown technologies also available. The lower-bound CDF can be thought of as the worst (in terms of first-order stochastic dominance) possible profit distribution that the agent can choose. The investor has maxmin preferences: she wants to guarantee herself a payoff of $\frac{1}{48}$ under every possible unknown technology that the agent might choose.

To keep the examples tractable, let the baseline technology be f_0, as defined in equation 2.1. We will consider what happens for different choices of the lower-bound CDF G. Intuitively, worse lower-bounds imply more ambiguity, i.e., larger sets of possible distributions for the principal to minimize over.

Largest Ambiguity. Suppose initially that the lower-bound CDF, denoted G, is arbitrarily bad, e.g., a Dirac mass on 0, denoted $\delta_0 = (1,0,0)$. Thus the investor thinks that any technology could be available in addition to f_0.

The investor can guarantee herself a positive profit in this setting, since she knows that whatever technology choice the entrepreneur makes must give the entrepreneur at least as much utility as using the optimal effort level under the baseline technology f_0. This level of utility, denoted $\upsilon(B)$, is:

$$\upsilon(B) = \max_e f_0(\underline{\pi} \mid e)\left(\frac{1}{2} - \underline{b}\right) + f_0(\overline{\pi} \mid e)\left(1 - \overline{b}\right) - \frac{1}{2}e^2.$$

The investor still requires a return of $\frac{1}{48}$ and evaluates a contract B as follows:

$$\begin{aligned}&\min_{f_m} f_m(\underline{\pi} \mid e)\,\underline{b} + f_m(\overline{\pi} \mid e)\,\overline{b}\\ &\text{s.t. } f_m(\underline{\pi} \mid e)\left(\frac{1}{2} - \underline{b}\right) + f_m(\overline{\pi} \mid e)\left(1 - \overline{b}\right) - \frac{1}{2}e^2 \geq \upsilon(B),\end{aligned}$$

where e is the effort the agent would choose if given technology f_m. Clearly the minimizing technology f_m depends on the contract B. Let $V_P(B \mid f_0)$ be the value function of the above program. Note that the constraint in the minimization problem is relaxed if e is smaller[7] and thus the worst-case occurs when $e = 0$. We can therefore re-write the principal's problem as follows:

$$\begin{aligned}V_P(B \mid f_0) &= \min_{f_m} f_m(\underline{\pi})\,\underline{b} + f_m(\overline{\pi})\,\overline{b} \\ &\quad \text{s.t. } f_m(\underline{\pi})\left(\frac{1}{2} - \underline{b}\right) + f_m(\overline{\pi})\left(1 - \overline{b}\right) \geq \upsilon(B).\end{aligned} \tag{2.2}$$

Fixing $B = \left(\underline{b}, \overline{b}\right)$, this program has a linear objective and constraint and therefore has corner solutions unless $\underline{b} = \frac{1}{2}\overline{b}$, in which case any f that satisfies the constraint is a minimizer and $V_P(B \mid f_0) = \frac{\upsilon(B)\overline{b}}{1-\overline{b}}$.

To allow for a ready comparison with the preceding example, assume the agent only has access to technology f_0, so that we want to solve the same Pareto problem as in the textbook model, except with the robust objective function for the principal:

$$\begin{aligned}&\max_B \max_e \left\{ f_0(\underline{\pi} \mid e)\left(\frac{1}{2} - \underline{b}\right) + f_0(\overline{\pi} \mid e)\left(1 - \overline{b}\right) - \frac{1}{2}e^2 \right\}\\ &\text{s.t. } V_P(B \mid f_0) \geq \frac{1}{48},\end{aligned}$$

[7] This is a familiar result from Carroll (2015): one could simply lower the effort levels associated with any distribution and this "new" technology would be better for the agent.

Assume by way of contradiction that in the optimal contract for the entrepreneur $\underline{b} > \frac{1}{2}\bar{b}$. Then $f_m(\underline{\pi}) = 0$ and $f_m(\bar{\pi}) = \frac{\upsilon(B)}{1-\bar{b}}$, and hence $V_P(B \mid f_0) = \frac{\upsilon(B)\bar{b}}{1-\bar{b}} \geq \frac{1}{16}$. However, since $\underline{b} > \frac{1}{2}\bar{b} \geq 0$ this contract cannot be best for the entrepreneur as $\underline{b}$ could be decreased at no loss to the investor (and for all $e > 0$, $f(\underline{\pi}) > 0$, so that the entrepreneur values this). A similar argument applies if $\underline{b} < \frac{1}{2}\bar{b}$, but a more formal treatment is in main text.

Given that $\underline{b} = \frac{1}{2}\bar{b}$, we can interpret this as an equity contract that gives the investor a $\bar{b}$ percentage stake in the project. In this example we have $\bar{b} = \frac{1}{4}$, $\underline{b} = \frac{1}{8}$, so that the investor gets 25% equity. This implies an effort choice under the technology f_0 of $e(1/8, 1/4) = 0.5 < 0.933$; not surprisingly less efficient (further below the efficient effort of $e^* = 1$) than under the live-or-die contract described in the previous subsection. This equity contract, gives the agent a payoff of $\upsilon(B) = \frac{1}{16} < 0.145$.

To check that the principal's robust objective attains the required level of utility given this equity contract, we note that one worst-case distribution is (like in the case $\underline{b} > \frac{1}{2}\bar{b}$) to set $f_m(\underline{\pi}) = 0$ and $f_m(\bar{\pi}) = \frac{\upsilon(B)}{1-\bar{b}} = \frac{1}{16}\frac{4}{3} = \frac{1}{12}$ (implying that $f_m(0) = \frac{11}{12}$). This indeed generates the required return of $V_P(B \mid f_0) = \frac{\upsilon(B)\bar{b}}{1-\bar{b}} = \frac{1}{48}$.

Contracts which are good for motivating effort tend to backload rewards for agents under the MLRP assumption, i.e., they have $\underline{b} > \frac{1}{2}\bar{b}$. With this amount of ambiguity however, these contracts are ruled out by possible actions for the agent which put close to probability 1 on zero profit and just enough probability on the highest possible profit outcome to make them attractive. In the example above that meant a worst-case distribution $f_m = \left(\frac{11}{12}, 0, \frac{1}{12}\right)$. It raises the question about what we can say in environments where the principal is not facing such large ambiguity, i.e., where the lower-bound distribution does not permit such pessimism.

Small Ambiguity. Suppose now that the lower-bound CDF does not permit any measure, e.g., let $G = \left(\frac{5}{6}, 1, 1\right)$. This lower-bound implies that the principal puts at most probability $\frac{5}{6}$ on the 0 profit outcome. Indeed, this is typical in the worst-case for the principal and so the question is how to split the remaining $\frac{1}{6}$ probability between the low and high profit outcomes. Given this, the principal's objective from 2.2 can be rewritten as:

$$V_P(B \mid f_0) = \min_{p \in \left[0, \frac{1}{6}\right]} p\underline{b} + \left(\frac{1}{6} - p\right)\bar{b} \qquad (2.3)$$
$$\text{s.t. } p\left(\frac{1}{2} - \underline{b}\right) + \left(\frac{1}{6} - p\right)\left(1 - \bar{b}\right) \geq \upsilon(B).$$

If $\overline{b} < \underline{b}$, the solution to the minimization problem has $f(\underline{\pi}) = 0$ and $f(\overline{\pi}) = \frac{\upsilon(B)}{1-\overline{b}}$. This can never be optimal since we could decrease $\underline{b}$ (to the level of $\overline{b}$) and the investor's payoff would be unaffected. Such a change in the contract is obviously good for the entrepreneur, strictly so if there is a positive probability of profit $\underline{\pi}$ being realized. Thus, it is without loss of generality to consider monotonic contracts, i.e., contracts for which $\overline{b} \geq \underline{b}$.

Given that $\overline{b} \geq \underline{b}$, the solution to the unconstrained version of the principal's minimization problem is $V_P(B \mid f_0) = \underline{b}/6$. We will later check that the agent's incentive constraint is satisfied for the relevant choice of $\underline{b}$ and $\overline{b}$.

Now, to allow comparison to the previous examples, suppose that the only available technology to the agent is f_0.[8] We want to solve the Pareto problem given by:

$$\max_{B,} \max_{e} \left\{ f_0(\underline{\pi} \mid e)\left(\frac{1}{2} - \underline{b}\right) + f_0(\overline{\pi} \mid e)\left(1 - \overline{b}\right) - \frac{1}{2}e^2 \right\}$$
$$\text{s.t.} \quad \frac{1}{6}\underline{b} \geq \frac{1}{48}, \ \underline{b} \leq \overline{b}.$$

Unsurprisingly, the solution to this is $\underline{b} = 1/8$ and $\overline{b} = 1/8$; a "debt" contract. The effort choice under this contract is

$$e^* = \frac{1 - 2\underline{b}}{1 + 2\overline{b}} = 0.6,$$

which is higher than 0.5, the effort choice under the equity contract in the largest ambiguity case.[9] This is still below the first-best efficient effort level $e^* = 1$, but moves in the direction of efficiency. Unsurprisingly, the agent's overall utility is higher as well. For this contract the agent's utility from the f_0 technology is $\upsilon(B) = 3/40 > 1/16$; decreasing the principal's ambiguity makes the agent better off.

Lastly, one of the minimizing probability distributions for the principal in expression 2.3 is $\left(\frac{5}{6}, 0, \frac{1}{6}\right)$; this action is feasible since it gives the agent utility $\frac{1}{6}\frac{7}{8} = \frac{7}{48} > \frac{3}{40}$ and thus satisfies his incentive constraint.[10] If the lower-bound G was further improved debt contracts remain optimal in this example, but the repayment level the investor needs to meet her required return decreases.

[8] Our choice of G technically rules out some low effort levels for the agent from technology f_0. We disregard this and later verify that these effort levels would not have been chosen in a subsequent footnote.

[9] Given that the effort chosen for the debt contract is $e = 0.6$, this induces the probability distribution over profits of $\left(\frac{12}{25}, \frac{2}{5}, \frac{3}{25}\right)$, which indeed first-order stochastically dominates G.

[10] There are other minimizing distributions for the principal which make the agent's incentive constraint tight.

3. Model

I study a moral hazard model where the agent may have technologies which are unknown to the principal at the ex-ante contracting stage. A principal (she) contracts with an agent (he), who is to take a costly, private action which will randomly produce a publicly observable profit outcome $\pi \in [0, \overline{\pi}] =: \Pi$.

More formally, an action is a pair $(e, F) \in [0, \overline{e}] \times \Delta(\Pi)$, where $e \in [0, \overline{e}]$ is interpreted as a level of effort, F is a cumulative distribution function (CDF) over profit outcomes and $\Delta(\Pi)$ is the set of Borel measures over Π, which we endow with the topology of weak convergence. The function mapping effort levels to utility cost for the agent, $c\colon [0, \overline{e}] \to \mathbb{R}_+$, is common knowledge and increasing. We normalize c so that $c(0) = 0$. A technology for the agent is a method for converting effort into random profit outcomes, i.e., a technology is a function $F : [0, \overline{e}] \to \Delta(\Pi)$. Instead of writing $(F(e))(\pi)$ we write $F(\pi \mid e)$. Since functions can be represented by their graphs, we can think of technology F as the graph of F:

$$\Gamma(F) = \{(e, F(\cdot \mid e)) \in [0, \overline{e}] \times \Delta(\Pi) : e \in [0, \overline{e}]\},$$

that is, technology F is simply a set of actions (where effort levels are not repeated). Where it causes little confusion we will abuse notation and denote $\Gamma(F)$ by F. If we also assume that F is continuous in e and satisfies a stochastic concavity property[11] and that c is strictly increasing and convex, this would guarantee that an optimal effort choice for the agent exists and that first-order conditions are sufficient for describing it (Jewitt, 1988; Athey, 2000). These technical assumptions are common in the classic moral hazard literature, however are not generally necessary for our purposes.

The textbook models of moral hazard, starting with the classic paper by Holmström (1979)**?**, assume that there is a single profit technology, F_0, which is common knowledge. This literature requires further assumptions on the technology to deliver general results; in particular, these papers assume that F_0 satisfies the monotone likelihood ratio property (MLRP). MLRP is a natural regularity condition on the profit technology which formalizes the idea that more effort should lead to better profit distributions. Consistent with this literature, we will often be interested in cases where the agent is choosing from technologies F that satisfy the MLRP. In many places this assumption can be relaxed, as will become obvious when the problem is simplified.

We will need a more general version of the MLRP than is commonly used, as allow a rich set of probability measures and in particular measures which do not have densities.

[11] F satisfies stochastic concavity, i.e., for all π, $-\int_0^{\pi} F(\pi' \mid e)\ d\pi'$ is concave in e.

The generalization of MLRP, due to Athey (2002)**?**, which extends the usual definition using Radon-Nikodym derivatives is discussed in the online appendix.

I consider a robust moral hazard problem in which the assumption that there is a single common knowledge profit technology, F_0, is relaxed. In particular, the principal knows that some baseline technology F_0 is available to the agent, but there could be other, unknown, profit technologies also available. This builds on the model of Carroll (2015)**?** and the assumption of a known technology plays the same role here—it guarantees a minimum utility level for the agent can attain under any contract. This utility allows the principal to constrain the set of technologies that the agent would be willing to choose and thus arbitrarily bad technologies which produce no profit can be ruled out. We simplify the analysis and notation without losing this qualitative property by assuming that F_0 is a constant technology (i.e., $F_0(\cdot|e) = F_0(\cdot|e')$ for all e, e'); we shall therefore write F_0 for the CDF $F_0(\cdot|0)$.[12]

On top of the baseline technology F_0, we assume that the principal knows a lower-bound CDF,[13] G, such that any realized technology (first-order) stochastically dominates G. This is all the set of all possible technologies as far as the principal is concerned be:

$$\mathcal{D}_G := \left\{ F \in \Delta(\Pi)^{[0,\bar{e}]} : \Gamma(F) \text{ compact, } F(\cdot \mid e) \leq G \text{ for all } e \right\}.$$

If $G = \delta_0$, then the constraint holds trivially for any CDF. As G approaches F_0, the amount of ambiguity the principal faces is diminishing. We start with these two extremes and will later consider a generic lower-bound CDF, G.

A contract, $B\colon [0,\overline{\pi}] \to \mathbb{R}_+$, specifies the payment made to the principal as a function of the realized profit. We assume B is Borel measurable with respect to the usual topology on $[0,\overline{\pi}] \subset \mathbb{R}$ and $B(\pi) \in [0,\pi]$ for all π (i.e., the investor's liability is limited to the initial investment and the entrepreneur's liability is limited to his entire profit).[14]

[12]Note that in the financial contracting example in section 2 we had a non-constant F_0. More generally, the F_0 CDF can be thought of as the distribution that would have been chosen by the agent from a non-constant technology. This would however complicate the analysis somewhat as changing a contract may also change the chosen CDF, without affecting the substantive role of F_0, which is to provide a bound on the agent's utility.

[13]We could assume that the principal knows a lower-bound technology. As we will see, the relevant bound for the principal's worst-case analysis is a profit distribution the agent can costlessly induce. As such, we can replace this assumption by a lower-bound technology. If the technology is sufficiently unproductive (a lower-bound on how effort gets converted into marginal benefit in terms of profit distributions), the analysis is unchanged.

[14]The literature often makes stronger assumptions, e.g., Carroll (2015) and Walton and Carroll (2019) assume contracts are continuous. Given that live-or-die contracts are discontinuous and optimal in the classic setting, we prefer to allow for a more general class of contracts.

The agent is a risk-neutral expected utility maximizer: given the set of technologies available to him, $\mathcal{A} = \{F_0, F_1, ..., F_N\} \subset \mathcal{D}_G$, and a contract, B, he solves:

$$\sup_{(e,F)\in\Gamma(\mathcal{A})} \int_0^{\overline{\pi}} (\pi - B(\pi)) \, \mathrm{d}F(\pi \mid e) - c(e), \tag{3.1}$$

where $\Gamma(\mathcal{A}) = \Gamma(F_0) \cup \Gamma(F_1) \cup ...\Gamma(F_N)$ is a set of actions representing the union of the possible actions under (or graphs of) the various available technologies. We sometimes abuse notation and write $(e, F) \in \mathcal{A}$, where we use $\mathcal{A}$ to denote $\Gamma(\mathcal{A})$. Let $V_A(B \mid \mathcal{A})$ denote the value function of the above. The assumption that $\mathcal{A}$ is a finite union of technologies ensures that $\Gamma(\mathcal{A})$ is compact, given our previous assumptions.[15] We say that the agent's technology set $\mathcal{A}$ satisfy MLRP if F_i respects the MLRP order with respect to effort for $i = 1, ...N$.

Even after the regularity assumptions we have made, note that the supremum in the above problem may not be attained unless we further restrict the set of permissible contracts B. Although restrictions are common in the literature, e.g., Carroll (2015) considers continuous contracts, one of the first technical results in this paper is that it is without loss of generality to assume B is lower semicontinuous, which implies that the supremum in equation 3.1 is attained. Thus, it is sensible to describe the arguments which maximize the agent's utility, $A^*(B \mid \mathcal{A}) \subset \Gamma(\mathcal{A})$.

Principals are extremely ambiguity averse about the potential technologies available to the agent, but are risk-neutral with respect to risks they understand. In particular, the principal's utility is

$$V_P(B \mid F_0) = \inf_{\substack{\mathcal{A} \text{ s.t.} \\ F_0 \in \mathcal{A} \subset \mathcal{D}_G}} \inf_{(e,F)\in A^*(B|\mathcal{A})} \int_0^{\overline{\pi}} B(\pi) \, \mathrm{d}F(\pi \mid e).$$

The assumption that the principal is getting the worst possible outcome when the agent is indifferent is largely inconsequential, since the worst-case $\mathcal{A}$ will usually have a single minimizing action.[16] Furthermore, when we show that restricting to lower semicontinuous contracts is without loss of generality, we will have that the infimum above is attained and therefore we may think of it as a minimum.

[15]Weaker assumptions on $\mathcal{A}$ also suffice; in particular, one could show that if $\mathcal{A}$ is compact in the right topology then $\Gamma(\mathcal{A})$ is also compact.

[16]Brooks (2014)**?** makes the same assumption as above, while Carroll (2015) assumes the agent maximizes the principal's utility when indifferent. The only instance in which the above is consequential is when we have a contract B and a baseline technology F_0, such that at the lowest effort level under F_0 the agent is obtaining the maximum possible profit he can get given B. Carroll (2015) rules these out by requiring contracts to be "eligibile". I make the assumption above predominantly because it avoids special cases and streamlines proofs.

We want to characterize Pareto efficient contracts in this environment. It is not immediate how a Pareto problem should be posed in this case, since the agent perfectly knows the technology set $\mathcal{A}$, while the principal faces Knightian uncertainty. The idea is to give the agent any extra utility that results from the realized $\mathcal{A}$, while satisfying a robust utility constraint for principal; this is natural since the principal does not express a preference over a specific $\mathcal{A}$, but rather takes a minimum over possible $\mathcal{A}$. As such, for each technology set, $\mathcal{A}$, we want to solve for the Pareto frontier,[17] given by the following problem:

$$\begin{aligned} &\max_{B} V_A\left(B \mid \mathcal{A}\right) \\ &\text{s.t. } V_P\left(B \mid F_0\right) \geq R, \end{aligned} \tag{3.2}$$

where $R \in [0, R_{\max}]$ denotes the location on the frontier and $R_{\max}$ is the maximum payment the principal can be guaranteed (the point at which the agent's participation constraint binds).

Since our $\mathcal{A}$ is very general, and in particular does not inherit the MLRP from individual technologies,[18] we will typically need to assume some additional structure to be able to solve the above Pareto problem. In problems of this type in the classical literature, starting with Holmström (1979)**?**, without the MLRP assumption we cannot hope to provide general results. The same thing is true in the robust contracting problem, unless the robustness of the principal's preferences simplifies the problem significantly. While this indeed happens in the largest ambiguity case, it is not generally true when the amount of ambiguity is limited.

In summary: the key features of the above assumptions is that (1) there is common knowledge of a lower-bound CDF and a constant baseline technology that the agent can choose and (2) we will characterize solutions to the Pareto problem, as stated in program 3.2, and mostly focus on the case where the agent is choosing from an MLRP set of technologies $\mathcal{A}$.

4. Preliminary Analysis

This section makes several preliminary observations, which greatly simplify the proofs of the major results. Several results in this section may be of independent interest. Some of these results are familiar from the literature studying robustness in mechanism design,

[17]Note that the notion of a Pareto frontier in the textbook setting also makes reference to a specific technology; in that case there is a single technology which is common knowledge.

[18]We will describe this in detail in section 6.

but since the literature is focused on the largest ambiguity case, i.e., $G = \delta_0$, proofs are included in the online appendix. The two most novel results of this section are that robust contracts are (i) lower semicontinuous and (ii) monotonic in profit (see next subsection).

A first observation, familiar from the literature, is that in finding the principal's worst-case scenario we can, without loss of generality, assume this occurs with zero effort from the agent. This is simply because the principal's worst-case analysis only considers CDFs that the agent would choose, i.e., by revealed preference, distributions $F(\pi \mid e)$ which satisfy:

$$\int_0^{\overline{\pi}} (\pi - B(\pi))dF(\pi \mid e) - c(e) \geq V_A(B \mid F_0),$$

but since $c(\cdot)$ is increasing this constraint is most permissive when $e = 0$. As such, a principal who faces unknown technologies can still bound her payoff. Aside from this revealed-preference constraint, she also has the knowledge that the chosen distribution must first-order stochastically dominate the lower-bound G.

Lemma 4.1. *For any Borel contract B, we have*

$$V_P(B \mid F_0) = \inf_{F \leq G} \left\{ \int_0^{\overline{\pi}} B(\pi)\, dF(\pi) \,:\, \int_0^{\overline{\pi}} (\pi - B(\pi))dF(\pi) \geq V_A(B \mid F_0) \right\}. \tag{4.1}$$

If B is lower semicontinuous,

$$V_P(B \mid F_0) = \min_{F \leq G} \left\{ \int_0^{\overline{\pi}} B(\pi)\, dF(\pi) \,:\, \int_0^{\overline{\pi}} (\pi - B(\pi))dF(\pi) \geq V_A(B \mid F_0) \right\}. \tag{4.2}$$

Proof. See appendix A.1. □

The first part of the lemma is a generalization of similar observations made in theorem 1 in Chassang (2013) and lemma 2.2 in Carroll (2015). In a moral hazard setting both of these papers find that the principal can essentially only bound her utility by the knowledge that the agent will not choose a worse outcome than what he is guaranteed under the known technology. Madarász and Prat (2017)**?** exploit a similar argument in a screening setting. The main difference between the above proof and earlier literature arises from the fact that the principal has a second way of bounding the worst-case outcome, the lower-bound technology G. The main difference in the proof above to the previous results is in dealing with the case when $G \neq \delta_0$ and the fact that, because B is only assumed to be Borel, we have to deal with the infimum problem for the principal (and supremum problem for the agent). The second part of the lemma is a technical result which follows from a careful application of the theorem of the maximum.

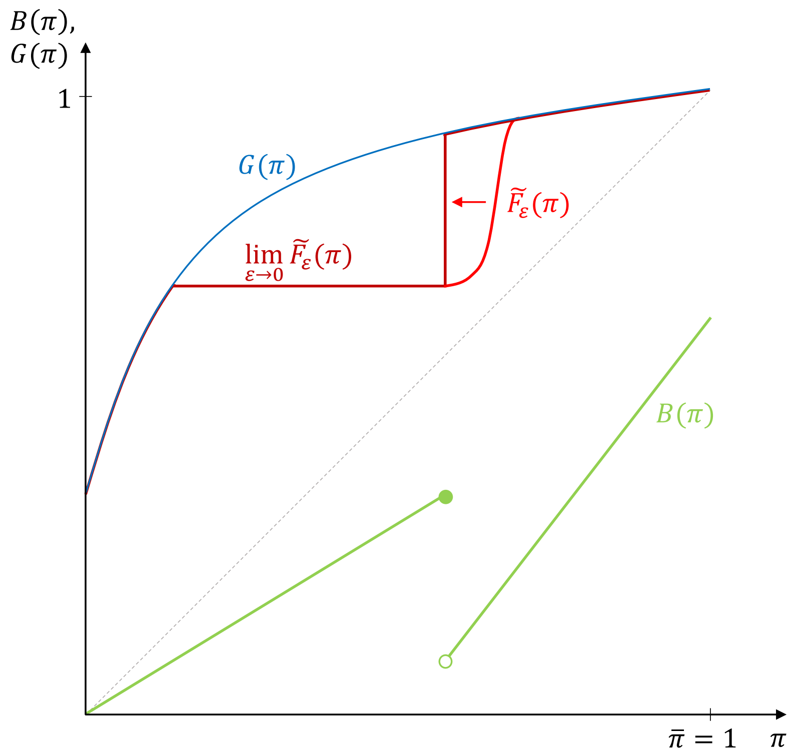

FIGURE 4.1. Proof idea for lemma 4.2.

Secondly, without loss of generality, we can restrict attention from all Borel contracts B to lower semicontinuous contracts. Since Innes (1990) finds that optimal contracts are not continuous (live-or-die contracts are not continuous, given our definitions), we do not wish to ex ante restrict our analysis to continuous contracts in the moral hazard problem presented above.[19] Let $\widehat{B}$ denote the lower semicontinuous hull of B, i.e., $\widehat{B}$ is the greatest lower semicontinuous function majorized by B.

Lemma 4.2. *For any Borel contract B and any $G \in \Delta(\Pi)$, we have that $V_P(B \mid F_0) \leq V_P(\hat{B} \mid F_0)$, i.e., the principal prefers the lower semicontinuous contract. Furthermore, for any $\mathcal{A} \subseteq \mathcal{D}_G$, $V_A(\hat{B} \mid \mathcal{A}) \geq V_A(B \mid \mathcal{A})$, i.e., the agent also prefers the lower semicontinuous contract.*

Proof. See online appendix B.2. □

The proof of this lemma is involved, since it uses several approximation arguments. The above, combined with lemma 4.1 implies that the principal's preferences are given by equation 4.2, a fact we will continually use.

[19]For example, Carroll (2015) and Walton and Carroll (2019) assume continuous contracts.

The intuition behind the assertion that $V_P(B \mid F_0) \leq V_P\left(\widehat{B} \mid F_0\right)$ is represented in figure 4.1. The figure plots both CDFs and contracts on the same axis, which is aided by assuming $\overline{\pi} = 1$. The curve in blue is the lower-bound CDF G, and the 45° line is in dashed grey. A proposed contract, B, is in green; B is not lower semicontinuous. The infimum sequence of CDFs, represented in red, denoted $\widetilde{F}_\varepsilon$, puts mass on π ever closer to the jump point of B, as figure 4.1 shows. However, we cannot shift the mass all the way, since this limiting CDF would result in a higher payoff to the principal (since the B under consideration is upper semicontinuous). Clearly, when the limiting CDF is considered with the lower semicontinuous hull of B (which in this case just involves lowering the function value at the jump point) we obtain the same payoff as the infimum of CDFs. The figure also illustrates the significance of this lemma—we are able to look at a single minimizing CDF (the limiting one) instead of having to use sequences.

Lemma 4.2 therefore shows that replacing a contract by it's lower semicontinuous hull is preferred by both the principal and agent and thus we focus on lower semicontinuous contracts without loss of generality. Recall that $\widehat{B}$ denotes the lower semicontinuous hull of B. For a lower semicontinuous contract B, and a fixed technology set $\mathcal{A}$, we write the agent's preferences as follows

$$\max_{\substack{F\in\mathcal{A}\\ e\in[0,\overline{e}]}} \int_0^{\overline{\pi}} (\pi - B(\pi))\ \mathrm{d}F(\pi \mid e) - c(e), \tag{4.3}$$

and denote by $V_A(B \mid \mathcal{A})$ and $A^*(B \mid \mathcal{A})$ the value function and argmax of the above, respectively. Note that these are well defined since $\pi - B(\pi)$ is upper semicontinuous.

4.1. Robustness of Monotone Contracts

Given the above preliminaries, the key assertion of this subsection is that robustness considerations lead to monotonic contracts. The intuition for this is that a principal facing a non-monotonic contract will assume that a productive technology which exploits the non-monotonicity will be available to the agent and therefore disregard any non-monotonic aspects of the contract.

Theorem 4.3. *For any $G \in \Delta(\Pi)$ and any non-monotonic contract $B(\pi)$ there exists a monotonic contract $B_m(\pi)$ such that*

$$V_P(B \mid F_0) \leq V_P(B_m \mid F_0)$$

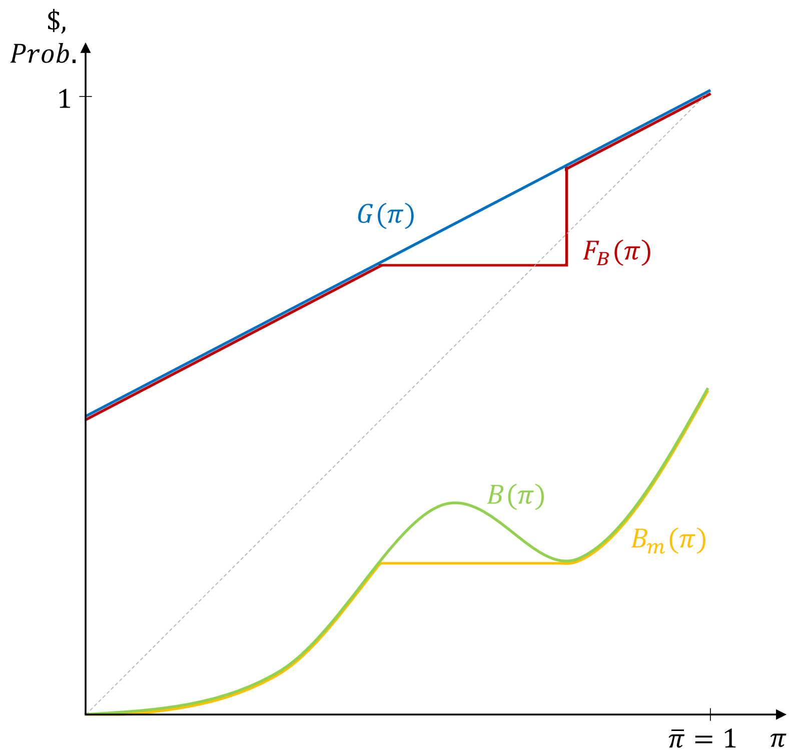

FIGURE 4.1. Proof idea for theorem 4.3.

i.e., the principal (weakly) prefers the monotone contract, and since $B(\pi) \geq B_m(\pi)$*, we have that for all* $\mathcal{A}$*,* $V_A(B \mid \mathcal{A}) = V_A(B_m \mid \mathcal{A})$*, i.e., the agent prefers the monotonic contract.*

Proof. See appendix A.2. □

The intuition for the above theorem is given in figure 4.1. The idea is that if a principal is offered a non-monotonic contract B, the green line in figure 4.1, she would discount the non-monotonic part, since in the worst-case analysis she thinks that nature will endow the agent with a technology which puts no mass on the non-monotonic part. From the perspective of the principal, we can thus replace this contract by its lower monotone hull, B_m, the yellow line in the figure, if we assume that the feasible set of distributions for the agent does not change. However, when B is replaced by the B_m, the agent's utility from the baseline technology, F_0, increases (which decreases the set of possible distributions the principal is minimizing over). While the change from B to B_m, also affects the way the agent compares contracts this turns out not to introduce complications for the principal, as the same sorts of distributions end up being feared.

5. Smallest Ambiguity, $F_0(\cdot \mid e) = G$

This section and the next consider the two extreme cases of the model and build intuition for the results. This section considers the smallest ambiguity case, where the lower-bound (G) and reference (F_0) technologies are the same. The main result is the Pareto optimality of debt contracts when the agent's realized technology MLRP dominates $G = F_0$.

Overall, the proof is very much related to the argument in Innes (1990), but there are two major differences. First, the restriction to monotonic contracts is a result of theorem 4.3, not an assumption. Second, we need to generalize the definitions and key lemmas to allow for non-differentiability of CDFs.

5.1. Result

The main result of this section is that Pareto optimal contracts take the form of debt.

Theorem 5.1. *For any $\mathcal{A} \subset \mathcal{D}_G$, which MLRP dominates G a solution to:*

$$\max_B V_A(B \mid \mathcal{A}),$$
$$\textit{subject to} \quad V_P(B \mid F_0) \geq R,$$

is $B_z^D(\pi) := \min(\pi, z)$ for some $z \in [0, \overline{\pi}]$.

The proof relies on an important property of the monotone likelihood ration order is summarized in the following lemma (from Innes, 1990).

Lemma 5.2. *Let $\phi(\pi)$ be a function such that $\phi(\pi) \geq 0$ for $\pi \leq \pi_B$, $\phi(\pi) \leq 0$ for $\pi \geq \pi_B$ and either:*

(1) $\int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) = 0$, *or*
(2) $\int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) \leq 0$ *and $\phi(\pi)$ decreasing for $\pi \geq \pi_B$.*

Then, for any $e_H > e_L$ and any MLRP family F, we have that $\int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) \geq \int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_H)$.

This is a generalization of lemma 1 from Innes (1990). The line of argument in the proof is similar, but we need to take care of technical difficulties arising from the non-existence of densities. The lemma is key in the proof of the main theorem, since it says that replacing generic monotone contracts by debt contracts implies higher marginal returns to effort.

The proof of the theorem proceeds by showing that when a monotonic non-debt contract, B, is replaced by an appropriately chosen debt contract the agent and principal are both (weakly) better off. Working out the debt contract to use is a two-step process:

there is the standard step, familiar from Innes (1990), where a debt contract is chosen to replace a monotone contract B to make the agent indifferent given the effort he was choosing from a MLRP family F. However, we also find a debt contract that makes the principal indifferent under the worst-case distribution (which in this case is $G = F_0$). The latter debt contract involves a (weakly) lower repayment level, since whatever the distribution the agent was using first-order stochastically dominates G, and thus replacing B with it makes the agent better off, while the principal is still guaranteed return R.

Corollary 5.3. *The repayment level, z, in the optimal contract, $B_z^D(\pi) = \min(\pi, z)$, is increasing in R and decreasing in G.*[20]

The above corollary follows since the level of debt is chosen so as to guarantee the principal the required utility R under the worst-case scenario where G and only G is available. This implies that the level of repayment z is increasing in R. Furthermore if $G \leq G'$, the level of repayment required under G' would be greater than under G.

5.2. Numerical Example

To demonstrate the efficiency of debt contracts, let us consider a simple example. Let $\Pi = [0,1]$, $e = [0,1]$, $c(e) = \frac{1}{10}e^2$ and $G = F_0 = U[0,1]$. Fix a level of principal utility R. As discussed, the worst-case scenario for the principal is that only the (constant) technology F_0 is available to the agent[21]. The principal is thus indifferent between many contracts. In particular, the principal is indifferent between an equity and debt contract defined as follows:

$$\begin{array}{ll} B_\alpha^E(\pi) = \alpha\pi, & \text{with } \alpha = 2R \\ B_z^D(\pi) = \min(\pi, z), & \text{with } z = 1 - \sqrt{1-2R}, \text{ and} \end{array},$$

since:

$$\int_0^1 B_z^D(\pi)\ \mathrm{d}G(\pi) = 2R\int_0^1 \pi\ \mathrm{d}\pi = R,$$

and:

$$\begin{aligned} \int_0^1 B_z^D(\pi)\ \mathrm{d}G(\pi) &= \int_0^{1-\sqrt{1-2R}} \pi\ \mathrm{d}\pi + \left(1-\sqrt{1-2R}\right)\left(1 - G\left(1-\sqrt{1-2R}\right)\right) \\ &= \frac{\left(1-\sqrt{1-2R}\right)^2}{2} + \left(1-\sqrt{1-2R}\right)\sqrt{1-2R} \\ &= 1-\sqrt{1-2R} - R \quad + \quad 2R + \sqrt{1-2R} - 1 = R. \end{aligned}$$

[20]If we think of potential G CDFs as being ordered by first-order stochastic dominance.

[21]This is because contracts have to be monotonic and the agent gets his "promised" utility under G.

Consider now an agent with the following technology set $\mathcal{A}$:

$$F(\pi \mid e) = \pi^{e+1}, \text{ for } e \in [0,1].$$

Note that this is an MLRP technology set and that $F(\pi \mid e) \leq G(\pi)$ for all e. Figure 5.1 plots the utilities of the agent under the two contracts above, given different possible reservation utilities of the principal R.

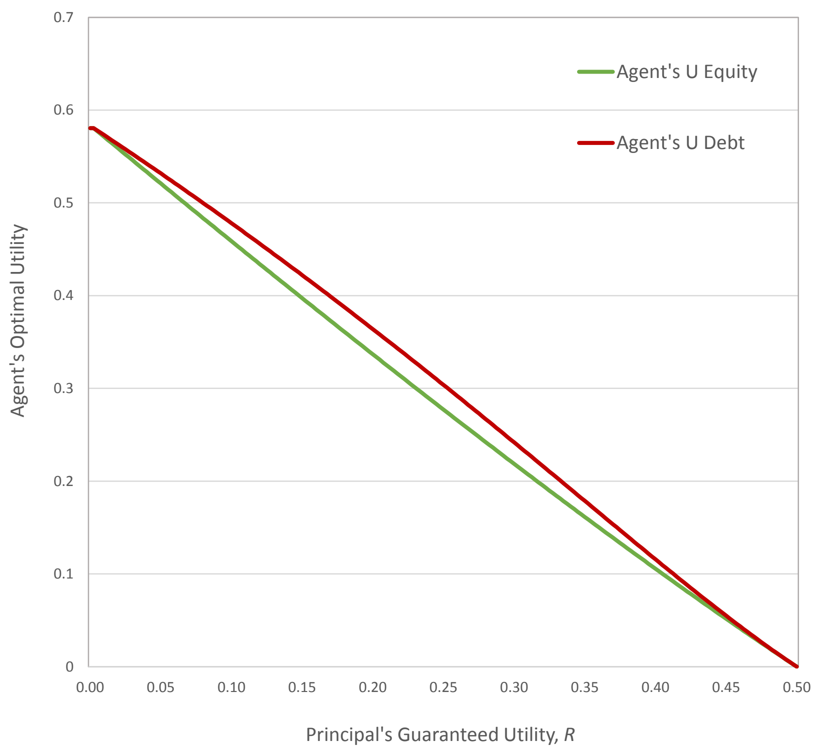

FIGURE 5.1. Numerical example illustrating Pareto Efficiency

We see in figure 5.1 that although the principal is indifferent between the contracts, the agent clearly prefers the debt contract for all $R \in (0, R_{\max})$. Note that in this case $R_{\max} = 1/2$. When $R = 0$ or $R = R_{\max}$ the debt and equity contracts are the same—they either award all profit to the agent or principal.

6. Largest Ambiguity, $G = \delta_0$

We now consider the case where the lower-bound CDF is trivial, so that the principal is minimizing over the largest set of (possibly) available technologies. We also do not

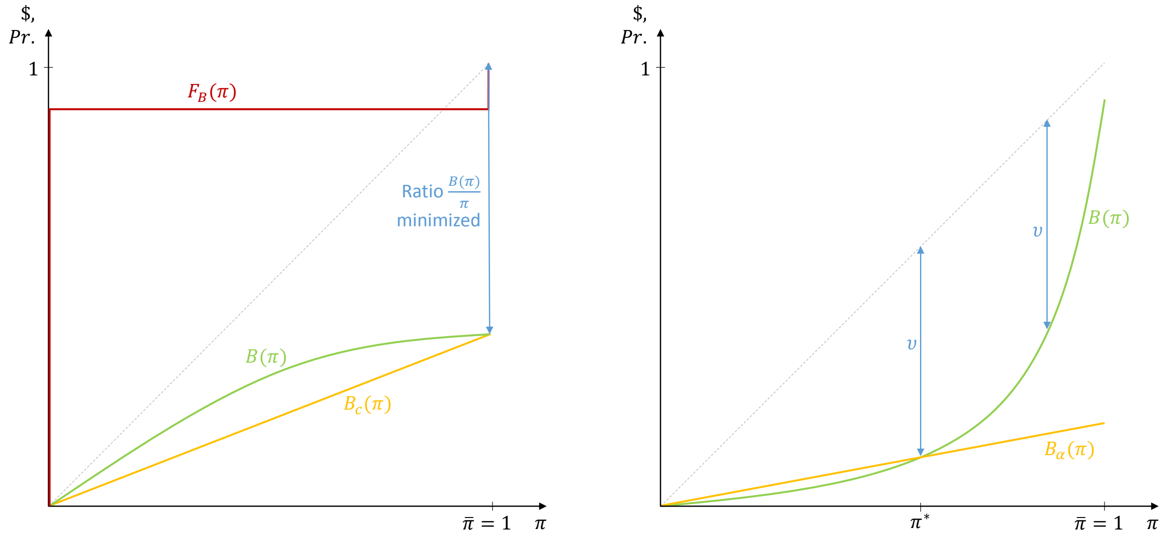

Optimal contract is (weakly) convex. Optimal contract is linear.

FIGURE 6.1. Proof idea for theorem 6.1.

need to assume that the agent is choosing from an MLRP family of technologies for this results.

Theorem 6.1. *For any $\mathcal{A} \subset \mathcal{D}_{\delta_0}$, a solution to:*

$$\max_B V_A (B \mid \mathcal{A}),$$
$$\textit{subject to} \quad V_P (B \mid F_0) \geq R,$$

is $B_\alpha(\pi) = \alpha\pi$ for some $\alpha \in [0, 1]$, i.e., a linear/equity contract.

The intuition for this proof is that an extremely uncertain principal places a huge premium on having preferences perfectly aligned with the agent, which is what happens when the contract is linear. Even if there are efficiency gains from providing stronger incentives for the agent at the upper end of profit outcomes, as is the case when $\mathcal{A}$ is an MLRP set, this benefit is over-ridden by the principal's pessimism.

The proof of theorem 6.1 is illustrated in figure 6.1. The left-hand panel gives the intuition for why contracts have to be (weakly) convex. In particular, consider a concave contract B (in green). In performing her worst-case analysis, the principal is wants to find the worst way (for her) that the agent can gain exactly the utility guaranteed by F_0, $v := V_A (B \mid F_0)$. Given that the set of CDFs she can minimize over is unrestricted, she will put mass on just two points: there will be a lot of mass on 0, since this gives her no payoff, and just enough mass on the point which minimizes $\frac{B(\pi)}{\pi}$, i.e., the point which minimizes what the principal gets relative to what the agent gets. In this case "just enough" means to make the agent choose this constructed CDF (at zero effort cost) over whatever was optimal in F_0. This worst-case CDF is illustrated by F_B in the figure. Now, consider replacing B by the lower convex hull, B_c. Note that at the worst-case the principal is indifferent between B and B_c. Furthermore, since B_c is linear, it satisfies a "no-weak-point" constraint, so that the minimizing CDF for the principal is any CDF which delivers the required utility to the agent—including F_B. This replacement therefore makes the principal no worse off, but makes the agent weakly (and generally strictly) better off.

The right-hand panel in figure 6.1 provides intuition for why contracts have to be linear. In particular, consider the principal's worst-case analysis when faced with a convex contract B, where the agent is guaranteed some level of utility υ. Jensen's inequality implies that the worst-case scenario is a Dirac distribution δ_{π^*} at the lowest level of profit which gives the agent exactly utility υ. One can replace B by a linear contract B_α that goes through $(\pi^*, B(\pi^*))$ and we again note that the principal is no worse off. It is not immediate that the agent likes this replacement however, since there is an interval, $[0, \pi^*]$, on which $B_\alpha > B$. The agent does like this replacement however—since the agent's average payoff under whatever technology he was choosing from $\mathcal{A}$ is at least υ, it cannot be the case that the agent is putting much mass on $[0, \pi^*]$ relative to the mass this CDF puts on $[\pi^*, \overline{\pi}]$. Another application of Jensen's inequality ensures that this replacement indeed gives the agent higher utility (and strictly higher if the agent's chosen distribution is not δ_{π^*}).

We say that $\mathcal{A}$ has full support, if for all $F_i \in \mathcal{A}$ and $e \in [0, \overline{e}]$, $supp(F_i(\cdot \mid e)) = [0, \overline{\pi}]$.

Corollary 6.2. *Equity is the unique solution to the above problem if $R \in (0, R_{\max})$ and $\mathcal{A}$ has full support.*

The equity contract is the unique efficient contract if the agent's technologies have full support and if the principal is not requiring a 0 return or would only accept 100% equity.

The robust contracting framework of Carroll (2015) maps closely to the largest ambiguity case analyzed above. One difference is that Carroll (2015) focuses on unknown actions, as opposed to technologies, and does not require MLRP (and indeed our proof above did not require it either); but this is in fact minor. The key difference is that Carroll (2015) focuses on the principal-optimal problem and the main result of the paper is stated below.

Theorem 6.3 (Carroll, 2015)**.** *A solution to:*

$$\max_{B} V_P(B \mid F_0),$$
$$\text{subject to} \qquad V_A(B \mid \mathcal{A}) \geq 0,$$

is $B_\alpha(\pi) = \alpha\pi$ *for some* $\alpha \in (0,1)$*.*

The above is a linear contract or, in our security-design-inspired language, the solution to the principal problem is an equity contract. As we illustrated in the discussion in section 5, this does not necessarily imply that the equity contract is efficient. However, Carroll (2015) also shows a uniqueness result: under the same conditions as in corollary 6.2, equity is the unique principal-optimal contract. This implies that equity must also be the efficient contract under these assumptions. Relative to Carroll (2015), the novelty of theorem 6.1 is therefore to show that equity is an efficient contract even when the uniqueness result fails. A contribution of the above is also the constructive nature of the proof of the result, which highlights the sorts of distributions feared by the principal (Carroll; 2015 used an elegant separating hyperplanes argument).

7. Intermediate Ambiguity

We now combine the insights from the study of the extreme cases in the two preceding sections to say something about intermediate levels of ambiguity. Let $B^P_{\alpha,z}(\pi) = \min(\pi, z + \alpha\pi)$ for some $z \in [0, \overline{\pi}]$, $\alpha \in [0,1]$ with $z + \alpha\overline{\pi} < \overline{\pi}$; this is also known as participating preferred equity.

Theorem 7.1. *For any contract* B*, there exists an MLRP technology set* $\mathcal{A}$*, such that the agent prefers* $B^P_{\alpha,z}(\pi)$ *to* B *and* $V_P\left(B^P_{\alpha,z} \mid F_0\right) \geq R$*, if* R *is sufficiently low.*

The above theorem is more restrictive than previous results, because it only looks at specific realized technology sets $\mathcal{A}$ that the agent may have access to, as opposed to general sets which the previous two results allowed. This is to be expected, as in general many contracts will satisfy the principal's robust constraint and the actual technologies available to the agent will impact the sort of contract that is Pareto efficient. However, the above suggests that participating preferred equity $B^P_{\alpha,z}$ has features that may make it optimal in a class of technology sets.

Participating preferred equity contracts can be thought of as a mixture of debt and equity. An investor issues a debt component and an equity component—the investor is entitled to all profit up to the repayment level of $\frac{z}{1-\alpha}$ and is then entitled to an additional

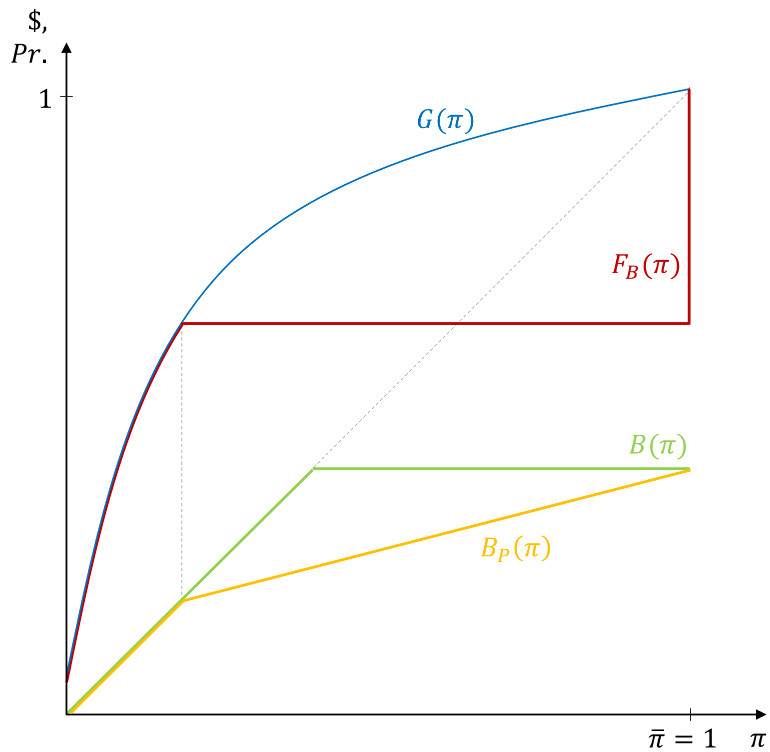

FIGURE 7.1. Proof idea for theorem 7.1.

α share of any profit above this level. The class of preferred equity contracts includes the simple debt and equity contracts we proved were efficient in previous sections. In particular, $\alpha = 0$ implies that $B^P_{\alpha,z} = B^D_z$ or simple debt, while $z = 0$ implies $B^P_{\alpha,z} = B^E_\alpha$ or simple equity.

We have already seen that debt contracts are good for incentive provision when the agent is choosing from an MLRP set of technologies. However, we now need to be careful that the replacement of a generic contract by a debt one does not pose a problem for the principal's utility and is still preferred by the agent, given the agent's realized technology set. This was straightforward for arbitrary technology sets in the case where ambiguity was small, since the principal's worst case distribution was $G = F_0$.

Now, when ambiguity is intermediate and thus $V_A(B \mid F_0)$ is sufficiently bigger than $V_A(B \mid G)$ we can make an unambiguous improvement if we are starting from the debt contract, as shown in figure 7.1.

Take any debt contract B, shown in green in figure 7.1, and consider the minimizing CDF subject to some arbitrary G. The worst-case CDF, F_B, for contract B is shown in red in the figure. This CDF has the feature that it puts mass on an interval of small profit realizations and on $\overline{\pi}$, as $\overline{\pi}$ minimizes the ratio of what the principal gets relative to the agent and thus this is the most costly way for the agent to get at least utility $V_A(B \mid F_0)$. The same logic was used in section 6 when we deduced the Pareto optimality of equity contracts.

Given the minimizing CDF, F_B, we see that B can be replaced by B_P, a preferred equity contract that is the lower convex hull of contract B on the region where F_B had no support. Note that a minimizing CDF for contract B_P is still F_B, thus the principal is indifferent to this change. The agent clearly prefers contract B_P since $B_P \leq B$.

Corollary 7.2. *The principal's payoff from contract* $B(\pi) = \min\{\pi, z + \alpha\pi\}$ *is:*

$$R = \int_0^{\frac{z}{1-\alpha}} \pi \, \mathrm{d}G + \frac{\alpha V_A(B \mid F_0) + z\left(1 - G\left(\frac{z}{1-\alpha}\right)\right)}{(1-\alpha)}.$$

This corollary gives a relation between α and z, in terms of the known parameters of the model—the lower-bound CDF G and the utility afforded to the agent under technology F_0, $V_A(B \mid F_0)$.

If $G = \delta_0$, then for any $z \geq 0$ we have that:

$$\begin{aligned} R &= \int_0^{\frac{z}{1-\alpha}} \pi \, \mathrm{d}G + \frac{\alpha V_A(B \mid F_0) + z\left(1 - G\left(\frac{z}{1-\alpha}\right)\right)}{(1-\alpha)} \\ &= \frac{\alpha V_A(B \mid F_0)}{(1-\alpha)}, \end{aligned}$$

thus the agent is (at least weakly) better off by setting $z = 0$, since any $z > 0$ is dominated. If $G = F_0$, and since worst-case scenario for Principal is G:

$$\begin{aligned} R &= \int_0^{\frac{z}{1-\alpha}} \pi \, \mathrm{d}G + \frac{\alpha V_A(B \mid F_0) + z\left(1 - G\left(\frac{z}{1-\alpha}\right)\right)}{(1-\alpha)}, \\ &= \int_0^{\frac{z}{1-\alpha}} \pi \, \mathrm{d}G + \int_{\frac{z}{1-\alpha}}^{\overline{\pi}} (\alpha\pi + z) \, \mathrm{d}G, \end{aligned}$$

which implies that we must have:

$$\int_{\frac{z}{1-\alpha}}^{\overline{\pi}} (\alpha\pi + z) \, \mathrm{d}G = \frac{\alpha V_A(B \mid F_0) + z\left(1 - G\left(\frac{z}{1-\alpha}\right)\right)}{(1-\alpha)}.$$

The above holds when $\alpha = 0$ and, as we argued earlier, such a debt contract provides the best incentives when the agent is choosing from an MLRP set in the smallest ambiguity case.

The characterization in corollary 7.2 implies that as G improves towards F_0 the set of z and a pairs which are undominated increases continuously (as the expression is continuous in G). For G sufficiently close to F_0 debt contracts become possible, however they may not be chosen for every realization of the technology set since the repayment level, z, may be too high. When G gets even closer to F_0, the repayment level decreases and debt contracts are certainly efficient when $G = F_0$ for any MLRP realization of $\mathcal{A}$. For sufficiently good realized technology sets, debt becomes Pareto optimal for G where $F_0 < G$.

8. Discussion

This paper has considered the question of Pareto optimal contracts, motivated by a central planner who cares about efficiency. Is there a way to decentralize the problem? To do so, we need to be careful to avoid any possibility of signaling and, as such, consider the following timing:

(1) The agent, knowing his realized technology set $\mathcal{A}$, as well as the information available to the principal (the baseline technology F_0, and lower-bound G) proposes a set of contracts $\mathcal{B}$;
(2) The principal accepts or rejects the set of contracts $\mathcal{B}$, based on the understanding that the agent will be able to select any $B \in \mathcal{B}$. The principal has an ex-post utility constraint, so that he will accept the set of contracts if for any $B \in \mathcal{B}$, $V(B \mid F_0) \geq R$;
(3) The agent chooses some $B \in \mathcal{B}$ and some $(e, F) \in \mathcal{A}$;
(4) Nature realizes profits and they are shared: the principal gets $B(\pi)$ and the agent gets $\pi - B(\pi)$.

In this game it is a weakly dominant strategy for the agent to propose the largest set of contracts that will get accepted, which is any contract B for which $V(B \mid F_0) \geq R$. This is related to Myerson's (1983)**?** *principle of inscrutability* in informed principal models: the informed party (the principal in Myerson's model, the agent here) should not want to reveal their private information if they can help it.[22]

The maxmin preference of the principal actually helps us here. If the principal was Bayesian and had an ex-ante participation constraint, signaling could clearly be helpful as it could indicate to the principal that "certain technologies are unlikely" and can therefore relax the participation constraint in favor of the agent. With a maxmin principal this signaling benefit is less obvious, since for any contract and any technology realization the principal has to get the required return R.

[22]Alternatively, we could have assumed that the agent is not yet aware of the set $\mathcal{A}$ in stage 1.

In the financial contracting interpretation, this decentralization looks potentially like a reverse convertible bond. With such a contract the issuer has the right to convert between a pre-agreed set of contracts, e.g., the choice that the agent gets from the set $\mathcal{B}$.

8.1. Conclusion

This paper has studied a general model of robust contracting when the principal does not know ex-ante all of the profit technologies available to the agent. The relaxation of this assumption of the textbook financial contracting model gives us a lot of traction. Firstly, it provides a complete theory of debt contracts which was the goal of this classic literature. Secondly, it shows that other, readily observable, contracts such as equity are Pareto efficient. More generally, these are examples of contracts in the class of participating preferred equity which we find to be Pareto improving for specific realized technology sets. While debt and equity are clearly common contracts, empirical work on venture capital also suggests participating preferred equity is not uncommon in practice.[23]

The key empirical implication of the results is that we should see firms in 'new' industries (such as social networking or biotech startups), where investors have little prior experience and face a lot of ambiguity about how the firm is going to generate profits, funded by equity contracts. Conversely, firms in 'established' industries (such as restaurants or accounting offices), where investors have a lot of experience and face less Knightian uncertainty, should be financed by debt.

In a very different set of models, focusing on costly information acquisition instead of moral hazard, a similar type of empirical prediction results. Dang, Gorton and Holmstrom (2012)**?** and Yang (2013)**?** find that in cases where information acquisition by the investor is not socially optimal (e.g., if the project is in a well-established industry), debt contracts should be observed as they provide the worst incentives for costly information acquisition. Yang and Zeng (2014)**?** generalize this model and find that if there are enough benefits from information acquisition by the investor (e.g., if the project is in a new industry), the class of participating (convertible) preferred equity contracts might be optimal.

Relative to the theoretical literature on robust contracting this paper shows that the robust environment leads to natural contracts, such as debt, even in environments where the pessimism of the principal is restricted, i.e., where the richness condition of Walton and Carroll (2019) fails. It also makes several novel technical observations concerning the general robustness of monotone contract and the ability to restrict attention, without

[23]Kaplan and Strömberg (2003)**?** find that 40 percent of venture capital funding rounds in their data set involve participating preferred equity.

loss of generality, from all Borel measurable contracts to lower semicontinuous ones. This latter observation is relevant as it justifies the use of minima as opposed to infima in the literature.

The analysis in this paper makes headway using two key simplifying assumptions. Firstly, in obtaining general results, we are restricting the analysis to cases where the agent is choosing from MLRP technology sets. This is to be expected, as the classic contracting literature also requires MLRP technologies to provide general conclusions, but it does leave open the question of what we can say without any restriction on the technology sets. Secondly, in the principal's minimization problem we are allowing for a relatively rich set of distributions for the principal to minimize over.[24] While we reduce the ambiguity of the principal by decreasing the size of the minimizing set (by increasing G), the techniques employed require this set to be rich (i.e., all CDFs that dominate G). It is natural to consider what happens when the richness of the minimizing sets is further restricted. These are two possible directions for future research.

[24]We also do not impose an upper-bound technology, but this is of far less improtance.

Appendix A. Proofs

A.1. Proof of Lemma 4.1

Lemma 4.1. *(i) For any Borel contract B, we have*

$$V_P(B \mid F_0) = \inf_{F \leq G} \left\{ \int_0^{\bar{\pi}} B(\pi)\, dF(\pi) \,:\, \int_0^{\bar{\pi}} (\pi - B(\pi)) dF(\pi) \geq V_A\left(B \mid F_0\right) \right\}.$$

(ii) If B is lower semicontinuous,

$$V_P(B \mid F_0) = \min_{F \leq G} \left\{ \int_0^{\bar{\pi}} B(\pi)\, dF(\pi) \,:\, \int_0^{\bar{\pi}} (\pi - B(\pi)) dF(\pi) \geq V_A\left(B \mid F_0\right) \right\}.$$

Proof. (i) Since a technology set $\mathcal{A}$ involving the baseline technology F_0 and one more constant technology F is a strict subset of $\mathcal{D}_G$, we have that

$$V_P(B \mid F_0) \leq \inf_{F \leq G} \left\{ \int_0^{\bar{\pi}} B(\pi)\, dF(\pi) \,:\, \int_0^{\bar{\pi}} (\pi - B(\pi)) dF(\pi) \geq V_A\left(B \mid F_0\right) \right\}.$$

So to prove the first claim, we are left to show that for all $\varepsilon > 0$

$$\inf_{F \leq G} \left\{ \int_0^{\bar{\pi}} B(\pi)\, dF(\pi) \,:\, \int_0^{\bar{\pi}} (\pi - B(\pi)) dF(\pi) \geq V_A\left(B \mid F_0\right) \right\} \leq V_P(B \mid F_0) + \varepsilon.$$

Take any $\varepsilon > 0$ and let $\{\mathcal{A}_n\}_{n=1}^{\infty} \subseteq \mathcal{D}_G$ be a minimizing sequence of $V_P(B \mid F_0) = \inf_{\mathcal{A}} \inf_{(e,F) \in A^*(B|\mathcal{A})} \int_0^{\bar{\pi}} B\left(\pi\right)\ \mathrm{d}F\left(\pi \mid e\right)$. By definition, there exists $N \in \mathbb{N}$ such that, for all $n \geq N$,

$$V_P(B \mid F_0) + \frac{\varepsilon}{2} > \lim_{k \to \infty} \int_0^{\bar{\pi}} B(\pi)\ dF_k^n(\pi),$$

where $\{F_k^n\}_{k=1}^{\infty} \subseteq \mathcal{A}_n$ is the maximizing sequence of distributions for the agent would solving $\sup\limits_{(e,F) \in \Gamma(\mathcal{A}_n)} \int_0^{\bar{\pi}} (\pi - B(\pi))\, dF(\pi \mid e) - c(e)$. Now even when $\lim\limits_{k \to \infty} F_k^n \notin \mathcal{A}_n$ there exists K such that, for all $k \geq K$,

$$\left| \int_0^{\bar{\pi}} B(\pi) dF_k^n(\pi) - \lim_{j \to \infty} \int_0^{\bar{\pi}} B(\pi) dF_j^n(\pi) \right| < \frac{\varepsilon}{2}$$

and $\int_0^{\bar{\pi}}(\pi - B(\pi))dF_k^n(\pi) \geq \int_0^{\bar{\pi}}(\pi - B(\pi))dF_0(\pi)$, as otherwise the agent would pick $F_0 \in \Gamma(\mathcal{A}_n)$. Now, take this F_K^n and observe that[25]

$$\begin{aligned}
&\inf_{F\leq G}\left\{\int_0^{\bar{\pi}} B(\pi)\,dF(\pi)\,:\,\int_0^{\bar{\pi}}(\pi - B(\pi))dF(\pi) \geq V_P(B \mid F_0)\right\} - V_P(B \mid F_0)\\
&\quad\leq \int_0^{\bar{\pi}} dF_K^N(\pi) + \left(-\lim_{k\to\infty}\int_0^{\bar{\pi}} B(\pi)dF_k^N(\pi) + \lim_{k\to\infty}\int_0^{\bar{\pi}} B(\pi)dF_k^N(\pi)\right) - V_P(B \mid F_0)\\
&\quad\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.
\end{aligned}$$

(ii) **?**We need to prove that the minimization problem is well defined when B is lower semicontinuous.[26]

Note that, by the portmanteau theorem,[27] $F \to \int_0^{\bar{\pi}} B(\pi)\,dF(\pi)$ is lower semicontinuous in the weak* topology, denoted $\sigma(\Delta(\Pi), \mathcal{C}_b(\Pi))$, since B is lower semicontinuous.[28] The result then follows from the theorem of the maximum provided that the set

$$\mathcal{D} = \left\{F \in \Delta(\Pi) :\ F \leq G \text{ and } \int_0^{\bar{\pi}}(\pi - B(\pi))dF(\pi) \geq \int_0^{\bar{\pi}}(\pi - B(\pi))dF_0(\pi)\right\}$$

is nonempty (which is true by assumption, since $F_0 \in \mathcal{D}$) and compact in $\sigma(\Delta(\Pi), \mathcal{C}_b(\Pi))$.

To see the latter, observe that $\Delta(\Pi)$ is compact in the weak* topology as a consequence of the Banach-Alaoglu theorem, since $\Delta(\Pi)$ is a weak* closed subset of $(\mathcal{C}(\Pi))$ which is a separable Banach lattice for Π compact metrizable.[29] We are thus left to show that $\mathcal{D}$ is closed.

From the portmanteau theorem, $F \to \int_0^{\bar{\pi}}(\pi - B(\pi))dF(\pi)$ is upper semicontinuous in the weak* topology (since B is lower semicontinuous). It follows that

$$\left\{F \in \Delta(\Pi)\,:\,\int_0^{\bar{\pi}}(\pi - B(\pi))dF(\pi) \geq \int_0^{\bar{\pi}}(\pi - B(\pi))dF_0(\pi)\right\}$$

is closed in the weak* topology.

Furthermore let $\bar{F} \in \overline{\{F \in \Delta(\Pi) : F \leq G\}}$ and $\{F_n\}_{n=1}^{\infty} \subseteq \{F \in \Delta(\Pi) : F \leq G\}$ be a sequence converging to $\bar{F}$ in the weak* topology.[30] Then $\lim_n F_n(\pi) = \bar{F}(\pi)$ for all π such that $\bar{F}$ is continuous at π. But $\bar{F}$ is monotone, hence has at most countably many

[25]We can use the limiting distribution if $\lim_{k\to\infty} F_k^n \in \mathcal{A}_n$.

[26]Special thanks to Theó Durandard for pointing out a significant generalization to the previous version of this argument that relied on G being differentiable.

[27]Specifically, see Aliprantis and Border (2006) theorem 15.5.

[28]Here $\mathcal{C}_b(\Pi)$ denotes the set of bounded continuous real functions on Π, see Aliprantis and Border (2006) chapter 15.

[29]See Stone-Weierstrass theorem and theorem 9.14 in Aliprantis and Border (2006).

[30]The weak* topology is metrizable, which allows us to work with sequences.

discontinuity points by Froda's theorem, and, therefore, $\bar{F}(\pi) = \lim_n F_n(\pi)$ on a comeager set $C(\bar{F})$ of Π. Clearly, on $C(\bar{F}) \subseteq \Pi$, $\bar{F}(\pi) \leq G(\pi)$, since, for all $n \in \mathbb{N}$, $F_n(\pi) \leq G(\pi)$.

We now show that $\bar{F}(\pi) \leq G(\pi)$ on the whole set Π. Recall that $\bar{F}(\pi)$ is right continuous, since it is a CDF, and hence $\forall \{\pi_n\}_{n=1}^{\infty} \subseteq \Pi$ such that $\pi_n \downarrow \pi$, $\lim_{n\to\infty} \bar{F}(\pi_n) = \bar{F}(\pi)$. But, for all $\pi < \bar{\pi}$, since $C(\bar{F})$ is dense in Π, there exists a sequence $\{\pi_n\}_{n=1}^{\infty} \subseteq C(\bar{F})$ such that $\pi_n \downarrow \pi$. Therefore $\bar{F}(\pi) = \lim_{n\to\infty} \bar{F}(\pi_n) \leq G(\pi)$, since $\forall n \in \mathbb{N}$, $\bar{F}(\pi_n) \leq G(\pi_n)$. Finally, $G(\bar{\pi}) = 1 \geq 1 = \bar{F}(\bar{\pi})$.

So, $\forall \pi \in \Pi$, $\bar{F}(\pi) \leq G(\pi)$, and therefore $\overline{\{F \in \Delta(\Pi) \,:\, F \leq G\}} = \{F \in \Delta(\Pi) \,:\, F \leq G\}$. Thus $\mathcal{D}$ is the intersection of closed sets, hence closed. □

A.2. Proof of Theorem 4.3

Theorem 4.3. *For any $G \in \Delta(\Pi)$ and any non-monotonic lower semicontinuous contract $B(\pi)$ there exists a monotonic contract $B_m(\pi)$ such that*

$$V_P(B \mid F_0) \leq V_P(B_m \mid F_0)$$

i.e., the principal (weakly) prefers the monotone contract, and since $B(\pi) \geq B_m(\pi)$, we have that for all $\mathcal{A}$, $V_A(B \mid \mathcal{A}) = V_A(B_m \mid \mathcal{A})$, i.e., the agent prefers the monotonic contract.

Define the feasible set of distributions that the principal minimizes over for a given contract B as follows:

$$\mathcal{F}(B) = \left\{ F \leq G : \int_0^{\bar{\pi}} (\pi - B(\pi))dF(\pi) \geq \int_0^{\bar{\pi}} (\pi - B(\pi))dF_0(\pi) \right\}. \tag{A.1}$$

Lemma A.1. *Take a lower semicontinuous contract B and let $F^B = \arg\min_{F\in\mathcal{F}(B)} \int_0^{\bar{\pi}} B(\pi)\ \mathrm{d}F(\pi)$. For any interval (π_1, π_2) such that there exists $\pi^* > \pi_2$ for which $B(\pi^*) < B(\pi')$ for all $\pi' \in (\pi_1, \pi_2)$, F^B is constant on $[\pi_1, \pi_2)$.*

Proof. Assume by way of contradiction $\lim_{\pi'\uparrow\pi_2} F^B(\pi') - F^B(\pi_1) = \gamma > 0$. Consider:

$$F^*(\pi) = \begin{cases} F^B(\pi) & \text{if } \pi < \pi_1 \\ F^B(\pi_1) & \text{if } \pi \in [\pi_1, \pi_2) \\ F^B(\pi) - \gamma & \text{if } \pi \in [\pi_2, \pi^*) \\ F^B(\pi) & \text{if } \pi \geq \pi^* \end{cases},$$

and note that $F^*(\pi) \leq F^B(\pi) \leq G(\pi)$ for all π and thus satisfies the constraint. Furthermore, by construction we have that:

$$\int_0^{\bar{\pi}} B(\pi)\ \mathrm{d}F^B(\pi) > \int_0^{\bar{\pi}} B(\pi)\ \mathrm{d}F^*(\pi),$$

and thus F^B could not have solved the minimization problem, if $F^* \in \mathcal{F}(B)$. We note that F^* first-order stochastically dominates F^B and thus $\int_0^{\bar{\pi}} \pi \, dF^*(\pi) \geq \int_0^{\bar{\pi}} \pi \, dF^B(\pi)$, so that

$$\int_0^{\bar{\pi}} (\pi - B(\pi)) \; dF^*(\pi) \geq \int_0^{\bar{\pi}} (\pi - B(\pi)) \; dF^B(\pi) \geq \int_0^{\bar{\pi}} (\pi - B(\pi)) dF_0(\pi),$$

so that $F^* \in \mathcal{F}(B)$. □

We are now ready to prove the theorem.

Proof of Theorem 4.3. Consider the following contract:

$$B_m(\pi) = \min_{\pi' \in [\pi, \bar{\pi}]} B(\pi').$$

Clearly $B_m(\pi) \leq B(\pi)$ for all π and hence the agent weakly prefers B_m for any technology set $\mathcal{A}$. By lemma 4.2 we can represent the principal's utility by expression 4.2, so that we need to show that:

$$\min_{F \in \mathcal{F}(B_m)} \int_0^{\bar{\pi}} B_m(\pi) dF(\pi) \geq \min_{F \in \mathcal{F}(B)} \int_0^{\bar{\pi}} B(\pi) dF(\pi),$$

where the feasible sets of distributions for the contract B and B_m are as defined in equation A.1.

Let $F^B = \arg\min_{F \leq G} \int_0^{\bar{\pi}} B(\pi) \; \mathrm{d}F(\pi)$ and note that $\mu_{F^B}(\{\pi : B_m(\pi) < B(\pi)\}) = 0$ by lemma A.1. Thus, $V_P(B \mid F_0) = \int_0^{\bar{\pi}} B(\pi) dF^B(\pi) = \int_0^{\bar{\pi}} B_m(\pi) dF^B(\pi)$. To prove the theorem it suffices to show that for any $F \in \mathcal{F}(B_m)$,

$$\int_0^{\bar{\pi}} B_m(\pi) dF(\pi) \geq \int_0^{\bar{\pi}} B_m(\pi) dF^B(\pi).$$

Assume by way of contradiction that there exists $F \in \mathcal{F}(B_m)$ such that

$$\int_0^{\bar{\pi}} B_m(\pi) dF(\pi) < \int_0^{\bar{\pi}} B_m(\pi) dF^B(\pi),$$

for some $\varepsilon > 0$. Clearly, $\mu_F(\{\pi : B_m(\pi) < B(\pi)\}) > 0$. Since F is a regular Borel measure, there exists K compact, $K \subseteq \{\pi : B_m(\pi) < B(\pi)\}$ such that

$$\mu_F(\{\pi : B_m(\pi) < B(\pi)\} \setminus K) < \frac{\varepsilon}{3\bar{\pi}}.$$

Now, let π_* be such that $\mu_F([\pi_*, \bar{\pi}] \cap K) < \frac{\varepsilon}{3\bar{\pi}}$ and let $\pi_\# = \max\{\pi \in [\pi_*, \bar{\pi}] : B(\pi) \leq B(\pi_*)\}$. Since K is compact and B is lower semicontinuous, such π_*, $\pi_\#$ exist. Let

$$\tilde{F} = \begin{cases} F(\pi) & \text{if } \pi \in [0, \pi_*) \\ F(\pi_*) & \text{if } \pi \in [\pi_*, \pi_\#) \\ F(\pi) & \text{if } \pi \in [\pi_\#, \bar{\pi}] \end{cases},$$

and notice that

$$\begin{aligned} \int_0^{\bar{\pi}} B_m(\pi) dF(\pi) &\geq \int_0^{\bar{\pi}} B_m(\pi) d\tilde{F}(\pi) - \frac{2\epsilon}{3} \\ &= \int_0^{\bar{\pi}} B(\pi) d\tilde{F}(\pi) - \frac{2\epsilon}{3} \\ &\geq \int_0^{\bar{\pi}} B(\pi) dF^B(\pi) - \frac{2\epsilon}{3}, \end{aligned}$$

since $\tilde{F} \in \mathcal{F}(B_m)$ by construction (see argument in, and F^B is a minimizer. Putting all pieces together, we get

$$\int_0^{\bar{\pi}} B(\pi) dF^B(\pi) - \epsilon > \int_0^{\bar{\pi}} B_m(\pi) dF(\pi) \geq \int_0^{\bar{\pi}} B(\pi) dF^B(\pi) - \frac{2\epsilon}{3},$$

a contradiction. □

A.3. Proof of Theorem 5.1

Theorem 5.1. *For any $\mathcal{A} = G \cup F(\cdot \mid e) \subset \overline{\mathcal{D}}s_G$, a solution to:*

$$\max_B V_A(B \mid \mathcal{A}),$$
$$\textit{subject to} \quad V_P(B \mid F_0) \geq R,$$

is $B^D(\pi, z) := \min(\pi, z)$ for some z.

Proof. Let $B(\pi)$ be a monotonic, lower semicontinuous non-debt contract, i.e., $\{\pi : B^D(\pi, z) \neq B(\pi)\}$ is not G-null for every z. The principal's worst case in this instance is if only technology $F_0 = G$ was available. Let z_0 solve:

$$\int B^D(\pi, z_0)\ \mathrm{d}G(\pi) = \int B(\pi)\ \mathrm{d}G(\pi).$$

Note that such a z_0 exists by the intermediate value theorem since B^D is continuous in z_0.

We now show that the agent prefers the contract $B^D(\pi, z_0)$ to $B(\pi)$. Let:

$$\phi(\pi) = B^D(\pi, z_0) - B(\pi).$$

By definition:

$$\int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}G(\pi) = 0,$$

and by lemma 5.2, for any $e \geq e^*$:

$$\int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F^*(\pi \mid e) \leq \int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F^*(\pi \mid e^*) = 0,$$

so that:

$$\int_0^{\overline{\pi}} (\pi - B(\pi))\ \mathrm{d}F^*(\pi \mid e) - c(e) \leq \int_0^{\overline{\pi}} \left(\pi - B^D(\pi, z_1)\right)\ \mathrm{d}F^*(\pi \mid e) - c(e),$$

and thus the agent weakly prefers $B^D(\cdot, z_1)$ over B. Furthermore, note that because $F^* \leq G$, we have

$$\int B(\pi)\ \mathrm{d}G(\pi) \leq \int B(\pi)\ \mathrm{d}F^*(\pi \mid e^*),$$

and thus $z_0 \leq z_1$. This means that the agent also prefers the contract $B^D(\pi, z_0)$ to $B(\pi)$.

Since by definition of $B^D(\cdot, z_0)$ we have that the principal's robust constraint is satisfied, because the worst-case is still G and:

$$\int B^D(\pi, z_0)\ \mathrm{d}G(\pi) = \int B(\pi)\ \mathrm{d}G(\pi) \geq R,$$

we have that B^D is efficient. □

A.4. Proof of Lemma 5.2

Lemma 5.2. *Let $\phi(\pi)$ be a function such that $\phi(\pi) \geq 0$ for $\pi \leq \pi_B$, $\phi(\pi) \leq 0$ for $\pi \geq \pi_B$ and either:*

(1) $\int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) = 0$, *or*
(2) $\int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) \leq 0$ *and* $\phi(\pi)$ *decreasing for* $\pi \geq \pi_B$.

Then, for any $e_H > e_L$ and any MLRP family F, we have that $\int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) \geq \int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_H)$.

Proof. Note that under case 1 we have:

$$\int_0^{\pi_B} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) = -\int_{\pi_B}^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) > 0.$$

Take any $e_H > e_L$ and consider:

$$\begin{aligned}
& \left(\int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) - \int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_H)\right) \int_0^{\pi_B} \phi(\pi_L)\ \mathrm{d}F(\pi_L \mid e_L) \\
= & \int_0^{\pi_B} \phi(\pi_L)\ \mathrm{d}F(\pi_L \mid e_L) \int_0^{\overline{\pi}} \phi(\pi) \left(\frac{dF(\pi \mid e_L)}{dC(\pi)} - \frac{dF(\pi \mid e_H)}{dC(\pi)}\right) \mathrm{d}C(\pi) \\
= & \left(-\int_{\pi_B}^{\overline{\pi}} \phi(\pi_H) \frac{dF(\pi_H \mid e_L)}{dC(\pi_H)}\ \mathrm{d}C(\pi_H)\right) \int_0^{\pi_B} \phi(\pi_L) \left(\frac{dF(\pi_L \mid e_L)}{dC(\pi_L)} - \frac{dF(\pi_L \mid e_H)}{dC(\pi_L)}\right) \mathrm{d}C(\pi_L) \\
& + \left(\int_0^{\pi_B} \phi(\pi_L) \frac{dF(\pi_L \mid e_L)}{dC(\pi_L)}\ \mathrm{d}C(\pi_L)\right) \int_{\pi_B}^{\overline{\pi}} \phi(\pi_H) \left(\frac{dF(\pi_H \mid e_L)}{dC(\pi_H)} - \frac{dF(\pi_H \mid e_H)}{dC(\pi_H)}\right) \mathrm{d}C(\pi_H),
\end{aligned}$$

where we write $C(\pi)$ for $C(\pi \mid e_L, e_H)$. By Fubini's theorem (applies since the above are integrable and C is a probability measure and therefore σ-finite) the above equals:

$$\begin{aligned}
& -\int_0^{\pi_B} \int_{\pi_B}^{\overline{\pi}} \phi(\pi_H) \phi(\pi_L) \frac{dF(\pi_H \mid e_L)}{dC(\pi_H)} \left(\frac{dF(\pi_L \mid e_L)}{dC(\pi_L)} - \frac{dF(\pi_L \mid e_H)}{dC(\pi_L)}\right) \mathrm{d}C(\pi_H)\ \mathrm{d}C(\pi_L) \\
& + \int_0^{\pi_B} \int_{\pi_B}^{\overline{\pi}} \phi(\pi_H) \phi(\pi_L) \frac{dF(\pi_L \mid e_L)}{dC(\pi_L)} \left(\frac{dF(\pi_H \mid e_L)}{dC(\pi_H)} - \frac{dF(\pi_H \mid e_H)}{dC(\pi_H)}\right) \mathrm{d}C(\pi_H)\ \mathrm{d}C(\pi_L) \\
= & \int_0^{\pi_B} \int_{\pi_B}^{\overline{\pi}} \phi(\pi_H) \phi(\pi_L) \left[\frac{dF(\pi_H \mid e_L)}{dC(\pi_H)} \frac{dF(\pi_L \mid e_H)}{dC(\pi_L)} - \frac{dF(\pi_L \mid e_L)}{dC(\pi_L)} \frac{dF(\pi_H \mid e_H)}{dC(\pi_H)}\right] \mathrm{d}C(\pi_H)\ \mathrm{d}C(\pi_L) \\
\geq & \ 0,
\end{aligned}$$

where the last inequality follows since $\phi(\pi_L) \geq 0$, $\phi(\pi_H) \leq 0$ and by the generalized MLRP:

$$\frac{dF(\pi_H \mid e_L)}{dC(\pi_H)} \frac{dF(\pi_L \mid e_H)}{dC(\pi_L)} \leq \frac{dF(\pi_L \mid e_L)}{dC(\pi_L)} \frac{dF(\pi_H \mid e_H)}{dC(\pi_H)}.$$

Thus:

$$\left(\int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) - \int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_H)\right) \int_0^{\pi_B} \phi(\pi_L)\ \mathrm{d}F(\pi_L \mid e_L) \geq 0,$$

and since $\int_0^{\pi_B} \phi(\pi_L)\ \mathrm{d}F(\pi_L \mid e_L) > 0$, we have that:

$$\int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) \geq \int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_H).$$

Case 2 follows similarly (using first-order stochastic dominance of $F(\cdot \mid e_H)$ over $F(\cdot \mid e_L)$ and the fact that ϕ is decreasing for $\pi \geq \pi_B$). May need to "split up" mass at π^*.

In particular, under case 2, there exists some π^* and $\alpha \in (0,1]$ such that:

$$\int_0^{\pi_B} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) = -\lim_{\pi' \to \pi^*_-} \int_{\pi_B}^{\pi'} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) - \alpha \left[F(\pi^* \mid e_L) - F\left(\pi^*_- \mid e_L\right)\right] \phi(\pi^*).$$

We can then repeat the above, replacing $\overline{\pi}$ by π^*, with the alpha-mass adjustment. We have then shown that:

$$\left(\begin{array}{c} \int_0^{\pi^*} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) - \int_0^{\pi^*} \phi(\pi)\ \mathrm{d}F(\pi \mid e_H) \\ -\alpha\phi(\pi^*)\left[f(\pi^* \mid e_L) - f(\pi^* \mid e_H)\right] \end{array}\right) \int_0^{\pi_B} \phi(\pi_L)\ \mathrm{d}F(\pi_L \mid e_L) \geq 0, \quad \text{(A.2)}$$

where

$$f(\pi^* \mid e_L) = F(\pi^* \mid e_L) - F\left(\pi^*_- \mid e_L\right).$$

Because $F(\pi \mid e_H)$ dominates $F(\pi \mid e_L)$ with respect to the monotone likelihood ratio order, it also conditionally first-order stochastically dominates it (conditioning on any set). Conditioning on (π^*, ∞) and π^* implies that:

$$\int_{\pi^*_+}^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) + \alpha f(\pi^* \mid e_L)\,\phi(\pi^*) \geq \int_{\pi^*_+}^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_H) + \alpha f(\pi^* \mid e_H)\,\phi(\pi^*), \quad \text{(A.3)}$$

since $\phi(\pi)$ decreasing for $\pi \geq \pi_B$. Combining A.2 and A.3 we have the desired result.

Note that the above inequalities are strict, i.e., $\int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_L) > \int_0^{\overline{\pi}} \phi(\pi)\ \mathrm{d}F(\pi \mid e_H)$ if MLRP holds strictly. □

A.5. Proof of Theorem 6.1

The statement of the theorem is repeated below for convenience.

Theorem 6.1. *For any $\mathcal{A} \subset \mathcal{D}$, a solution to:*

$$\max_B V_A(B \mid \mathcal{A}),$$
$$\textit{subject to} \quad V_P(B \mid F_0) \geq R,$$

is $B_\alpha(\pi) = \alpha\pi$ for some α, i.e., a linear/equity contract.

The proof proceeds by first showing that a solution to the above must be a (weakly) convex contract B, since the principal will not put any value on concave portions of a contract and thus the lower convex hull of B is evaluated in the same way as the original contract by the principal. We then show that the appropriate linear contract is optimal within the set of convex contracts, since it is no worse for the principal and better for the agent.

Lemma A.3. *In problem* **??** *for any non-convex B, there exists a convex B_c such that $V_P(B \mid F_0) \leq V_P(B_c \mid F_0)$ and $B \geq B_c$.*

Proof. Note that it is without loss of generality to consider B which are lower semicontinuous by lemma 4.2. Let B_c be the lower convex hull of B, i.e., the largest weakly convex function majorized than B. Clearly $B \geq B_c$ and thus $V_A(B \mid F_0) \leq V_A(B_c \mid F_0)$. It suffices to consider the case where $V_A(B_c \mid F_0) = V_A(B \mid F_0) =: \upsilon$, since a larger $V_A(B_c \mid F_0)$ decreases the constraint set and thus weakly increases $V_P(B_c \mid F_0)$.

It suffices to show that:

$$\left\{\min_{F\in\Delta(\Pi)} \int_0^{\overline{\pi}} B(\pi)\ \mathrm{d}F \text{ s.t. } \int_0^{\overline{\pi}} \pi - B(\pi)\ \mathrm{d}F \geq \upsilon\right\} = \left\{\min_{\substack{F\in\Delta(\Pi)\\ supp(F)=D}} \int_0^{\overline{\pi}} B(\pi)\ \mathrm{d}F \text{ s.t. } \int_0^{\overline{\pi}} \pi - B(\pi)\ \mathrm{d}F \geq \upsilon\right\}, \text{ and}$$

$$\left\{\min_{F\in\Delta(\Pi)} \int_0^{\overline{\pi}} B_c(\pi)\ \mathrm{d}F \text{ s.t. } \int_0^{\overline{\pi}} \pi - B_c(\pi)\ \mathrm{d}F \geq \upsilon\right\} = \left\{\min_{\substack{F\in\Delta(\Pi)\\ supp(F)=D}} \int_0^{\overline{\pi}} B_c(\pi)\ \mathrm{d}F \text{ s.t. } \int_0^{\overline{\pi}} \pi - B_c(\pi)\ \mathrm{d}F \geq \upsilon\right\},$$

where $D = \{x : B(\pi) = B_c(\pi)\}$. Let F_B be the CDF which minimizes the LHS and F_c be the CDF which minimizes the RHS. Clearly, it is without loss of generality to assume that $supp(F_c) \subset D$.[31] We will show that $supp(F_B) \subset D$.

Assume by way of contradiction that there is some $\pi \in supp(F_B)$ and $\varepsilon > 0$ such that $B(\pi') > B_c(\pi')$ for all $\pi' \in \mathcal{N}_\varepsilon(\pi)$. Note that by construction there exist $\pi_L < \pi_H$ such that for all $\pi' \in \mathcal{N}_\varepsilon(\pi)$ there exists an $\beta(\pi') \in (0,1)$ such that $\pi' = \beta(\pi')\pi_L + (1-\beta(\pi'))\pi_H$ and:

$$\begin{aligned} B(\pi') > B_c(\pi') &= \beta(\pi') B_c(\pi_L) + (1-\beta(\pi')) B_c(\pi_H) \\ &= \beta(\pi') B(\pi_L) + (1-\beta(\pi')) B(\pi_H). \end{aligned}$$

Let $m = F_B(\pi+\varepsilon)_- - F_B(\pi-\varepsilon) = \int_{\mathcal{N}_\varepsilon(\pi)} \mathrm{d}F_B(\pi')$ and note that $m > 0$ since $\pi \in supp(F_B)$. Let:

$$\beta^* = \frac{1}{m}\int_{\mathcal{N}_\varepsilon(\pi)} \beta(\pi')\ \mathrm{d}F_B(\pi'),$$

so that:

$$\begin{aligned} 1-\beta^* &= 1 - \frac{1}{m}\int_{\mathcal{N}_\varepsilon(\pi)} \beta(\pi')\ \mathrm{d}F_B(\pi') \\ &= \frac{1}{m}\int_{\mathcal{N}_\varepsilon(\pi)} \mathrm{d}F_B(\pi') - \frac{1}{m}\int_{\mathcal{N}_\varepsilon(\pi)} \beta(\pi')\ \mathrm{d}F_B(\pi') \\ &= \frac{1}{m}\int_{\mathcal{N}_\varepsilon(\pi)} (1-\beta(\pi'))\ \mathrm{d}F_B(\pi'). \end{aligned}$$

[31] To see this, note that any $\pi \notin D$ is a convex combination of two elements in D and we can therefore, instead of putting mass on π, put the appropriate mass on the elements which constitute the convex combination.

Thus:

$$\begin{aligned}
&\beta^* B_c(\pi_L) + (1-\beta^*) B_c(\pi_H) \\
&\quad = \beta^* B(\pi_L) + (1-\beta^*) B(\pi_H) \\
&\quad = \frac{1}{m}\int_{\mathcal{N}_\varepsilon(\pi)} \beta(\pi') B(\pi_L)\ \mathrm{d}F_B(\pi') + \frac{1}{m}\int_{\mathcal{N}_\varepsilon(\pi)} (1-\beta(\pi')) B(\pi_H)\ \mathrm{d}F_B(\pi') \\
&\quad = \frac{1}{m}\int_{\mathcal{N}_\varepsilon(\pi)} \beta(\pi') B(\pi_L) + (1-\beta(\pi')) B(\pi_H)\ \mathrm{d}F_B(\pi') \\
&\quad < \frac{1}{m}\int_{\mathcal{N}_\varepsilon(\pi)} B(\pi')\ \mathrm{d}F_B(\pi'),
\end{aligned}$$

hence shifting mass to points π_L and π_H leads to a lower expected payoff for the principal and thus F_B could not have been a minimizer, which is a contradiction. □

We are now ready to prove the theorem. The proof goes by invoking Jensen's inequality and using a revealed preference argument to rule out technologies which the suggested replacement of contract makes worse.

Proof of Theorem 6.1. Next we show that for the principal's problem, a minimizing F for convex B puts mass on a single point. Note that for any F, where $\pi_F = \mathbb{E}_F[\pi]$, such that $\int_0^{\overline{\pi}} \pi - B(\pi)\ \mathrm{d}F \geq v = V_A(B \mid F_0)$, Jensen's inequality implies that:

$$v \leq \int_0^{\overline{\pi}} \pi - B(\pi)\ \mathrm{d}F \leq \pi_F - B(\pi_F) = \int_0^{\overline{\pi}} \pi - B(\pi)\ \mathrm{d}\delta_{\pi_F},$$

since $\pi - B(\pi)$ is concave. Furthermore, since B is convex:

$$\int_0^{\overline{\pi}} B(\pi)\ \mathrm{d}\delta_{\pi_F} = B(\pi_F) \leq \int_0^{\overline{\pi}} B(\pi)\ \mathrm{d}F.$$

Thus, the agent prefers the Dirac distribution (at the expected value) over others and the principal dislikes such distributions. Therefore, the appropriately chosen Dirac distribution (the one just attractive enough for the agent) is the principal's worst case.

Now consider replacing $B(\pi)$ by a linear contract $B_\alpha(\pi) = \alpha\pi$, where $\alpha = B(\pi^*)/\pi^*$ and:

$$\pi^* = \min \pi \text{ s.t. } \pi - B(\pi) \geq v.$$

Such a minimizer exists since $\pi - B(\pi)$ is concave and upper semicontinuous, which in particular implies $\pi - B(\pi)$ is continuous, since the domain of this function is a subset of $\mathbb{R}$.

Consider

$$\begin{aligned} V_A\left(B_\alpha \mid F_0\right) &= \int_0^{\overline{\pi}} \pi - B_\alpha\left(\pi\right)\ \mathrm{d}F_0 = \int_0^{\overline{\pi}} \frac{\pi^* - B\left(\pi^*\right)}{\pi^*}\pi\ \mathrm{d}F_0 \\ &\geq \upsilon \int_0^{\overline{\pi}} \frac{\pi}{\pi^*}\ \mathrm{d}F_0. \end{aligned}$$

Assume by way of contradiction that $\pi^* > \int_0^{\overline{\pi}} \pi\ \mathrm{d}F_0 = \pi_{F_0}$, but then by Jensen's inequality and since $\pi - B\left(\pi\right)$ is increasing for $\pi < \pi^*$ (by definition of π^*, since the subgradient at $B\left(\pi\right) \leq 1$ for $\pi < \pi^*$), we have that

$$\begin{aligned} \upsilon &= \int_0^{\overline{\pi}} \pi - B\left(\pi\right)\ \mathrm{d}F_0 \\ &\leq \pi_{F_0} - B\left(\pi_{F_0}\right) \\ &< \pi^* - B\left(\pi^*\right) \\ &= \upsilon, \end{aligned}$$

which is a contradiction. Thus

$$V_A\left(B_\alpha \mid F_0\right) = \upsilon \int_0^{\overline{\pi}} \frac{\pi}{\pi^*}\ \mathrm{d}F_0 \geq \upsilon = V_A\left(B \mid F_0\right),$$

i.e., the agent is weakly better off under the linear contract, given the baseline technology.

Consider any F such that $\int_0^{\overline{\pi}} \pi - B\left(\pi\right)\ \mathrm{d}F \geq \upsilon$ and assume by way of contradiction that $\pi^* > \int_0^{\overline{\pi}} \pi\ \mathrm{d}F = \pi_F$. We have that

$$\begin{aligned} \upsilon &\leq \int_0^{\overline{\pi}} \pi - B\left(\pi\right)\ \mathrm{d}F \\ &\leq \pi_F - B\left(\pi_F\right) \\ &< \pi^* - B\left(\pi^*\right) \\ &= \upsilon, \end{aligned}$$

where the second step follows by Jensen's inequality and the third since $\pi - B\left(\pi\right)$ is increasing for $\pi < \pi^*$ (by definition of π^*, since the subgradient at $B\left(\pi\right) \leq 1$ for $\pi < \pi^*$).

Thus, $\pi^* \leq \pi_F$ which implies that:

$$\begin{aligned} V_A(B_\alpha \mid F) &= \int_0^{\overline{\pi}} \pi - B_\alpha(\pi) \,\mathrm{d}F \\ &= \pi_F - B_\alpha(\pi_F) \\ &\geq \pi_F - B(\pi_F) \\ &\geq \int \pi - B(\pi) \,\mathrm{d}F. \end{aligned}$$

where the second line follows by linearity, the third since $\pi^* \leq \pi_F$ and the convexity of B and the final by the Jensen's inequality. Hence the agent prefers the linear contract for any F.

Note that now, the principal's payoff for any F is $\frac{\alpha}{1-\alpha} V_A(B_\alpha \mid F)$, i.e., it is perfectly aligned with the agent. When the agent chooses $F = F_0$, by construction the principal gets at least the level of utility he got with the previous contract B (since F_0 is weakly better than the principal's worst case). Other choices of F can only weakly improve the principal's payoff. □

A.6. Proof of Theorem 7.1

Theorem 7.1. *For any contract $B \neq B^P_{\alpha,z}$, there exists an MLRP technology set $\mathcal{A}$, such that the agent strictly prefers $B^P_{\alpha,z}(\pi)$ to B and $V_P\left(B^P_{\alpha,z} \mid F_0\right) \geq V_P(B \mid F_0)$.*

Proof. For an arbitrary contract B' define

$$\mathcal{F}(B') = \left\{ F \in \Delta(\Pi) \,:\, F \leq G \text{ and } \int_0^{\overline{\pi}} (\pi - B'(\pi)) dF(\pi) \geq V_A(B \mid F_0) \right\},$$

i.e., the set of distributions the principal views as feasible. Let the principal's worst-case distribution under contract B be $F_B \in \mathcal{F}(B)$. Let z_1 solve let z_1 solve:

$$\int B^D(\pi, z_1) \,\mathrm{d}F_D(\pi) = \int B(\pi) \,\mathrm{d}F_B(\pi) \geq R,$$

where $F_D = \arg\min_{F \in \mathcal{F}(B^D)} \left\{ \int B^D(\pi, z_1) \,\mathrm{d}F(\pi) \right\}$. As in the proof of theorem 5.1 the agent weakly prefers $B^D(\cdot, z_1)$ over B, if endowed with an MLRP family which dominates F_D. That is, consider technology set $\mathcal{A} = F_0 \cup F$, where F is an MLRP family (where MLRP holds strictly) such that $F(\cdot|0) = F_D(\pi)$. Since MLRP is strict, note that by the last part of the proof of lemma 5.2 the agent strictly prefers B^D to B.

Now, since F_D puts zero mass on $\pi \in (\pi_D, \overline{\pi})$, we can take $B^D(\pi, z_0)$ and replace it by its lower convex hull on the region where F_D had no support. This new contract is indeed $B^P_{\alpha,z}(\pi)$ for appropriately chosen α and z. In this case $\alpha = (1 - G(\pi_D))/(\overline{\pi} - \pi_D)$ and

$z = \pi_D(1-\alpha)$. Now one of the minimizing CDFs for the contract $B^P_{\alpha,z}(\pi)$ is F_D, and thus the principal is indifferent to this change. The agent prefers contract $B^P_{\alpha,z}(\pi)$ since $B^P_{\alpha,z}(\pi) \leq B^D(\cdot, z_0)$. Thus the agent strictly prefers $B^P_{\alpha,z}(\pi)$ to B for the defined $\mathcal{A}$. □

A.7. Proof of Corollary 7.2

Corollary 7.2. *The principal's payoff from contract* $B(\pi) = \min\{\pi, z + \alpha\pi\}$ *is:*

$$R = \int_0^{\frac{z}{1-\alpha}} \pi \, \mathrm{d}G + \frac{\alpha V_A(B \mid F_0) + z\left(1 - G\left(\frac{z}{1-\alpha}\right)\right)}{(1-\alpha)}.$$

Proof. The principal's payoff from contract $B(\pi) = \min(\pi, z + \alpha\pi)$ is:

$$\begin{aligned} R &= \int_0^{\frac{z}{1-\alpha}} \pi \, \mathrm{d}G + \min_F \int_{\frac{z}{1-\alpha}}^{\overline{\pi}} (\alpha\pi + z) \, \mathrm{d}F \\ \text{s.t.} \quad & \int_{\frac{z}{1-\alpha}}^{\overline{\pi}} (\pi - \alpha\pi - z) \, \mathrm{d}F \geq V_A(B \mid F_0). \end{aligned}$$

For any F:

$$\begin{aligned} \int_{\frac{z}{1-\alpha}}^{\overline{\pi}} (\alpha\pi + z) \, \mathrm{d}F &= \alpha \int_{\frac{z}{1-\alpha}}^{\overline{\pi}} \pi \, \mathrm{d}F + z\left(F(\overline{\pi}) - F\left(\frac{z}{1-\alpha}\right)\right) \\ &= \alpha \int_{\frac{z}{1-\alpha}}^{\overline{\pi}} \pi \, \mathrm{d}F + z\left(1 - G\left(\frac{z}{1-\alpha}\right)\right), \end{aligned}$$

and hence:

$$\begin{aligned} \int_{\frac{z}{1-\alpha}}^{\overline{\pi}} (\pi - \alpha\pi - z) \, \mathrm{d}F &= (1-\alpha) \int_{\frac{z}{1-\alpha}}^{\overline{\pi}} \pi \, \mathrm{d}F - z\left(1 - G\left(\frac{z}{1-\alpha}\right)\right) \\ &\geq V_A(B \mid F_0). \end{aligned}$$

Thus the constraint in the principal minimization problem holds as an equality.

Solving the last equation for the integral we have that:

$$\int_{\frac{z}{1-\alpha}}^{\overline{\pi}} \pi \, \mathrm{d}F = \frac{V_A(B \mid F_0) + z\left(1 - G\left(\frac{z}{1-\alpha}\right)\right)}{(1-\alpha)},$$

and hence for any F:

$$\int_{\frac{z}{1-\alpha}}^{\overline{\pi}} (\alpha\pi + z) \, \mathrm{d}F = \frac{\alpha V_A(B \mid F_0) + z\left(1 - G\left(\frac{z}{1-\alpha}\right)\right)}{(1-\alpha)},$$

which gives the characterization in the corollary. □

Appendix B. Online Appendix

This appendix contains some proofs which are more technical and involved. It starts with a discussion of the generalized notion of MLRP, due to Athey (2002), needed for the arguments.

B.1. Aside on MLRP

In this subsection, I make some basic remarks regarding a key assumption underlying most classical moral hazard models, including that of Holmström (1979) and Innes (1990)—the monotone likelihood ratio property (MLRP). The simplest version considers a family of CDFs, indexed by e, i.e., $F(\pi \mid e)$, which is twice-differentiable with respect to both π and e (as is the case in Innes (1990) and most existing models). In this case, the monotone likelihood ratio property (MLRP) states that:

$$\frac{\partial}{\partial \pi}\left(\frac{f_e(\pi \mid e)}{f(\pi \mid e)}\right) \geq 0,$$

where f is the density of F.

A slightly more general definition of the MLRP, but still requiring the existence of densities, is that the likelihood ratio:

$$\frac{f(\pi \mid e_H)}{f(\pi \mid e_L)},$$

is non-decreasing for any $e_H \geq e_L$. An equivalent way to state this is to assume that f is log-supermodular, i.e., for all $\pi_H \geq \pi_L$ and $e_H \geq e_L$:

$$\frac{f(\pi_H \mid e_H)}{f(\pi_H \mid e_L)} \geq \frac{f(\pi_L \mid e_H)}{f(\pi_L \mid e_L)}.$$

Recall that a non-negative function defined on a lattice, $h\colon X \to \mathbb{R}$ is log-supermodular if, for all $x, y \in X$, $h(x \wedge y)\, h(x \vee y) \geq h(x)\, h(y)$. Note that in this version of the definition, we can also treat f as the PMF if the measure is discrete.

However, we need to allow for general distributions in the present model—e.g., distributions which involve mixtures of continuous and discrete parts. As such, we work with general probability measures from the outset and require a general MLRP. The natural idea is to generalize the definition using Radon-Nikodym derivatives instead of densities, however one needs to take care to ensure the absolute continuity condition in the Radon-Nikodym theorem is satisfied.

This exact problem is beautifully addressed by Athey (2002)**?**, who gives the right generalization of the MLRP (see definition A1). For any $e_L < e_H \in \mathbb{R}_+$, first define a *carrying measure* as follows:

$$C(\pi \mid e_L, e_H) = \frac{1}{2}F(\pi \mid e_L) + \frac{1}{2}F(\pi \mid e_H).$$

Importantly, note that both $F(\cdot \mid e_L)$ and $F(\cdot \mid e_H)$ are absolutely continuous with respect to $C(\cdot \mid e_L, e_H)$. We say that a family of CDFs, F, satisfies the *monotone likelihood ratio property (MLRP)* if for any $e_L < e_H$, the Radon–Nikodym derivative $h(\pi, e) : (\pi, e) \mapsto \frac{dF(\pi|e)}{dC(\pi|e_L,e_H)}$ is log-supermodular for C-a.e. (π, e), where $e \in \{e_L, e_H\}$.

To give a little intuition for this, consider the special case of differentiable CDFs. We have that:

$$\begin{aligned}\frac{dF(\pi \mid e)}{dC(\pi \mid e_L, e_H)} &= \frac{dF(\pi \mid e)/d\pi}{dC(\pi \mid e_L, e_H)/d\pi} = \frac{f(\pi \mid e)}{\frac{1}{2}f(\pi \mid e_L) + \frac{1}{2}f(\pi \mid e_H)} \\ &= 2\frac{f(\pi \mid e)}{f(\pi \mid e_L) + f(\pi \mid e_H)}.\end{aligned}$$

The MLRP states that the Radon-Nykodym derivative above is log-supermodular, or:

$$\frac{dF(\pi_H \mid e_H)}{dC(\pi_H \mid e_L, e_H)}\frac{dF(\pi_L \mid e_L)}{dC(\pi_L \mid e_L, e_H)} \geq \frac{dF(\pi_H \mid e_L)}{dC(\pi_H \mid e_L, e_H)}\frac{dF(\pi_L \mid e_H)}{dC(\pi_L \mid e_L, e_H)}.$$

We write $F(\cdot \mid e_H) \overset{MLR}{\geq} F(\cdot \mid e_L)$ if the above holds. Note that in the differentiable CDF case reduces to:

$$\begin{aligned}\frac{f(\pi_H \mid e_H)}{f(\pi_H \mid e_L) + f(\pi_H \mid e_H)}\frac{f(\pi_L \mid e_L)}{f(\pi_L \mid e_L) + f(\pi_L \mid e_H)} &\geq \frac{f(\pi_H \mid e_L)}{f(\pi_H \mid e_L) + f(\pi_H \mid e_H)}\frac{f(\pi_L \mid e_H)}{f(\pi_L \mid e_L) + f(\pi_L \mid e_H)} \\ f(\pi_H \mid e_H) f(\pi_L \mid e_L) &\geq f(\pi_H \mid e_L) f(\pi_L \mid e_H) \\ \frac{f(\pi_H \mid e_H)}{f(\pi_H \mid e_L)} &\geq \frac{f(\pi_L \mid e_H)}{f(\pi_L \mid e_L)},\end{aligned}$$

which is one of the standard definitions.

The following is a useful observation about the general definition of MLRP, above.

Lemma B.1. *Take any CDF, $\widehat{F}$. Then for any $\varepsilon > 0$, there exists a family of CDFs, parametrized by $e \in [0, \overline{e}]$, which satisfies MLRP, where $F(\cdot \mid 0) = \delta_0$ and $F(\cdot \mid \varepsilon) = \widehat{F}$.*

Proof. Let $F(\cdot \mid e) = \delta_0$ if $e < \varepsilon$ and $F(\cdot \mid e) = \widehat{F}$ if $e \geq \varepsilon$. The only non-trivial case is when $e_L < \varepsilon$ and $e_H \geq \varepsilon$, so take these effort levels. This means that the carrying measure is $C(\pi) = \frac{1}{2}\delta_0 + \frac{1}{2}\widehat{F}(\pi)$, and thus that

$$\frac{dF(\pi \mid e_L)}{dC(\pi)} = \begin{cases} 2 & \text{if } \pi = 0 \\ 0 & \text{if } \pi > 0 \end{cases} \quad \text{and} \quad \frac{dF(\pi \mid e_H)}{dC(\pi)} = \begin{cases} 0 & \text{if } \pi = 0 \\ 2 & \text{if } \pi > 0 \end{cases}.$$

The Radon-Nikodym derivatives above are log-supermodular, since for any $\pi_H > \pi_L > 0$ we have

$$\frac{dF(\pi_H \mid e_H)}{dC(\pi)}\frac{dF(\pi_L \mid e_L)}{dC(\pi)} = 2 \cdot 0 = 0 \geq 0 \cdot 2 = \frac{dF(\pi_H \mid e_L)}{dC(\pi)}\frac{dF(\pi_L \mid e_H)}{dC(\pi)},$$

and if $\pi_H > \pi_L = 0$, then

$$\frac{dF(\pi_H \mid e_H)}{dC(\pi)}\frac{dF(\pi_L \mid e_L)}{dC(\pi)} = 2 \cdot 2 = 4 \geq 0 = 0 \cdot 0 = \frac{dF(\pi_H \mid e_L)}{dC(\pi)}\frac{dF(\pi_L \mid e_H)}{dC(\pi)}.$$

Thus the constructed family of distributions satisfies MLRP. □

B.2. Proof of Lemma 4.2

Lemma 4.2. *For any Borel contract B and any $G \in \Delta(\Pi)$, we have that $V_P(B \mid F_0) \leq V_P(\hat{B} \mid F_0)$, i.e., the principal prefers the lower semicontinuous contract. Furthermore, for any $\mathcal{A} \subseteq \mathcal{D}_G$, $V_A(\hat{B} \mid \mathcal{A}) \geq V_A(B \mid \mathcal{A})$, i.e., the agent also prefers the lower semicontinuous contract.*

Proof. Let $\hat{B}$ denote the lower semicontinuous envelope of B, i.e., $\forall \pi \in \Pi$,

$$\begin{aligned}\hat{B}(\pi) &= \sup\left\{\hat{b}(\pi) \,:\, \hat{b} \text{ is lower semicontinuous and } \hat{b} \leq B\right\}\\ &= \sup\left\{\tilde{b}(\pi) \,:\, \tilde{b} \text{ is Lipschitz continuous and } \hat{b} \leq B\right\},\end{aligned}$$

where the second line follows from Theorem 3.13 in Aliprantis and Border (2006).

Clearly, the agent prefers $\hat{B}$, since for all $\pi \in \Pi$, $\hat{B}(\pi) \leq B(\pi)$, and the agent (weakly) prefers the contract which gives less to the principal for any technology set.

We next show that the principal weakly prefers $\hat{B}$ to B; i.e., $V_P(\hat{B} \mid F_0) \geq V_P(B \mid F_0)$. It is enough to consider technology sets $\mathcal{A} \subseteq \mathcal{D}_G$ of the form $\mathcal{A} = \{F_0\} \cup \{F_1\}$ for some $F_1 \in \Delta(\Pi)$ such that $F_1 \leq G$ and

$$\int_0^{\overline{\pi}} (\pi - B(\pi))dF_1(\pi) \geq \int_0^{\overline{\pi}} (\pi - B(\pi))dF_0(\pi).$$

So we want to show that

$$\begin{aligned}V_P(B \mid F_0) = \inf_{F \leq G} &\left\{\int_0^{\overline{\pi}} B(\pi)dF(\pi) \,:\, \int_0^{\overline{\pi}} (\pi - B(\pi))dF(\pi) \geq \int_0^{\overline{\pi}} (\pi - B(\pi))dF_0(\pi)\right\}\\ \leq V_P(\hat{B} \mid F_0) = \inf_{F \leq G} &\left\{\int_0^{\overline{\pi}} \hat{B}(\pi)dF(\pi) \,:\, \int_0^{\overline{\pi}} (\pi - \hat{B}(\pi))dF(\pi) \geq \int_0^{\overline{\pi}} (\pi - \hat{B}(\pi))dF_0(\pi)\right\}.\end{aligned}$$

Since $\hat{B} \leq B$,

$$V_A \equiv \int_0^{\overline{\pi}} (\pi - B(\pi))dF_0(\pi) \leq \int_0^{\overline{\pi}} (\pi - \hat{B}(\pi))dF_0(\pi),$$

so

$$V_P(\hat{B} \mid F_0) \geq \inf_{F \leq G} \left\{ \int_0^{\bar{\pi}} \hat{B}(\pi) dF(\pi) \, : \, \int_0^{\bar{\pi}} (\pi - \hat{B}(\pi)) dF(\pi) \geq V_A \right\},$$

as the feasible set is now larger. Define

$$\mathcal{F}(\hat{B}) = \left\{ F \in \Delta(\Pi) \, : \, F \leq G \text{ and } \int_0^{\bar{\pi}} (\pi - \hat{B}(\pi)) dF(\pi) \geq V_A \right\},$$

and

$$\mathcal{F}(B) = \left\{ F \in \Delta(\Pi) \, : \, F \leq G \text{ and } \int_0^{\bar{\pi}} (\pi - B(\pi)) dF(\pi) \geq V_A \right\}.$$

Now, if we are able to show that for $F^* \in \arg \min_{F \in \mathcal{F}(\hat{B})} \int_0^{\bar{\pi}} \hat{B}(\pi) dF(\pi)$,[32] and for all $\epsilon > 0$, there exists $F^\epsilon \in \mathcal{F}(B)$ such that

$$\int_0^{\bar{\pi}} B(\pi) dF^\epsilon(\pi) \leq \int_0^{\bar{\pi}} \hat{B}(\pi) dF^*(\pi) + \epsilon, \tag{B.1}$$

the conclusion follows since

$$\begin{aligned} V_P(B \mid F_0) &\leq \inf_{F \in \mathcal{F}(\mathcal{B})} \int_0^{\bar{\pi}} B(\pi) dF(\pi) \\ &\leq \int_0^{\bar{\pi}} B(\pi) dF^\epsilon(\pi) \\ &\leq V_P(\hat{B} \mid F_0) + \epsilon, \end{aligned}$$

and, since it holds for all $\epsilon > 0$, this implies that $V_P(\hat{B} \mid F_0) \geq V_P(B \mid F_0)$. The proof of the claim in equation B.1 is long and relies on various approximation arguments and is given in the subsequent lemma. □

Lemma B.2. *Let* $F^* \in \arg \min_{F \in \mathcal{F}(\hat{B})} \int_0^{\bar{\pi}} \hat{B}(\pi) dF(\pi)$. *Then for all* $\epsilon > 0$, *there exists* $F^\epsilon \in \mathcal{F}(B)$ *such that*

$$\int_0^{\bar{\pi}} B(\pi) dF^\epsilon(\pi) \leq \int_0^{\bar{\pi}} \hat{B}(\pi) dF^*(\pi) + \epsilon$$

Proof. 1. We first claim that if $B_n \to B$ uniformly, then $\hat{B}_n \to \hat{B}$ pointwise. Assume by way of contradiction that $\hat{B}_n \not\to \hat{B}$ pointwise. Then there exists $\pi \in \Pi$, $\epsilon > 0$, and a subsequence $\left\{ \hat{B}_{n_m} \right\}_{m \in \mathbb{N}}$ such that, for all $m \in \mathbb{N}$, $|\hat{B}_{n_M}(\pi) - \hat{B}(\pi)| > \epsilon$. Either $\hat{B}_{n_M}(\pi) > \hat{B}(\pi) + \epsilon$ or $\hat{B}(\pi) > \hat{B}_{n_M}(\pi) + \epsilon$. Suppose that we are in the first case, i.e., $\hat{B}_{n_M}(\pi) > \hat{B}(\pi) + \epsilon$ (the second case is similar). This means that there exists $\delta > 0$ such that, for all $\pi' \in \mathcal{B}_\delta(\pi)$,

$$B_{n_M}(\pi') > \hat{B}(\pi') + \epsilon.$$

[32] This is well defined as proven in lemma 4.1.

But then, for all $\delta' > 0$, for all $\pi' \in \mathcal{B}_\delta(\pi)$, there exists $\pi'' \in \mathcal{B}_{\delta'}(\pi')$ such that $B_{n_M}(\pi'') > B(\pi'') + \epsilon$, but this contradicts that $B_n \to B$ uniformly.

2. The above step and Egorov's theorem, imply that if $B_n \to B$ pointwise, then $\hat{B}_n \to \hat{B}$ in $L^1(F^*)$.

Let $\epsilon > 0$. Since $B_n \to B$ pointwise, by Egorov's theorem, there exists $K \subseteq \Pi$ compact such that $B_n \to B$ uniformly on K and

$$\int_{\Pi - K} dF^*(\pi) < \frac{\epsilon}{4\bar{\pi}}.$$

This means that

$$\left| \int_{\Pi - K} \left(\hat{B}_n(\pi) - B(\pi) \right) dF^*(\pi) \right| < 2\bar{\pi} \frac{\epsilon}{4\bar{\pi}} = \frac{\epsilon}{2}.$$

Furthermore, by step 1, $\hat{B}_n \to B$ pointwise on K. Therefore, by the dominated convergence theorem,[33] there exists $N \in \mathbb{N}$ such that, for all $n \geq N$,

$$\left| \int_{K} \left(\hat{B}_n(\pi) - B(\pi) \right) dF^*(\pi) \right| < \frac{\epsilon}{2}.$$

Therefore, for all $n \geq N$,

$$\left| \int_{0}^{\bar{\pi}} \left(\hat{B}_n(\pi) - B(\pi) \right) dF^*(\pi) \right| < \epsilon.$$

3. By assumption, B is Borel measurable, and by Kechris (2012), proposition 11.5,**?** it is Baire measurable. It follows that $B \in Baire^k(\Pi)$ (Baire class k) for some $k \in \mathbb{N}$.

We now show that, for all $k \in \mathbb{N}$, $Baire^{k-1}(\Pi)$ is dense in $Baire^k(\Pi)$ for the L^1-norm associated to any Borel measure $F \in \Delta(\Pi)$; i.e., for all $k \in \mathbb{N}$, for all $\tilde{B}^k \in Baire^k(\Pi)$, for all $F \in \Delta(\Pi)$, for all $\epsilon' > 0$, there exists $\tilde{B}^{k-1} \in Baire^{k-1}(\Pi)$ such that

$$\left| \int_{0}^{\bar{\pi}} \left[\tilde{B}^k(\pi) - \tilde{B}^{k-1}(\pi) \right] dF(\pi) \right| < \epsilon',$$

and

$$\left| \int_{0}^{\bar{\pi}} \left[\hat{\tilde{B}}^k(\pi) - \hat{\tilde{B}}^{k-1}(\pi) \right] dF(\pi) \right| < \epsilon'.$$

Hence let $\epsilon' > 0$. Let $\tilde{B}^k \in Baire^k(\Pi)$. By definition of Baire classes, $Baire^k(\Pi)$ is the pointwise closure of $Baire^{k-1}(\Pi)$, and therefore there exists a sequence $\left\{ \tilde{B}_n^{k-1} \right\}_{n \in \mathbb{N}} \subseteq Baire^{k-1}(\Pi)$ that converges pointwise to $\tilde{B}^k$. Let $F \in \Delta(\Pi)$. Since F is a probability measure, it is finite, and since $\tilde{B}_n^{k-1}$ converges pointwise to $\tilde{B}^k$, it is bounded (e.g., by $\bar{\pi}$). So the dominated convergence theorem applies. Similarly, by step 2., $\hat{\tilde{B}}_n^{k-1}$ converges to

[33] $\hat{B}_n$ is bounded by $\bar{\pi}$ for all n.

$\hat{\tilde{B}}^k$ in $L^1(F^*)$. Therefore, we can choose $N \in \mathbb{N}$ such that, for all $n \geq N$,

$$\left| \int_0^{\bar{\pi}} \left[\tilde{B}^k(\pi) - \tilde{B}_n^{k-1}(\pi) \right] dF(\pi) \right| < \epsilon',$$

and

$$\left| \int_0^{\bar{\pi}} \left[\hat{\tilde{B}}^k(\pi) - \hat{\tilde{B}}_n^{k-1}(\pi) \right] dF(\pi) \right| < \epsilon'.$$

This concludes the proof of the claim in step 3.

4. Observe that for all $F \in \Delta(\Pi)$,

$$\begin{aligned}
&\int_0^{\bar{\pi}} B(\pi) dF(\pi) - \int_0^{\bar{\pi}} \hat{B}(\pi) dF^*(\pi) \\
=& \int_0^{\bar{\pi}} \left[B(\pi) - B^{k-1}(\pi) + B^{k-1}(\pi) dF(\pi) - \cdots - B^0(\pi) + B^0(\pi) \right] dF(\pi) \\
&- \int_0^{\bar{\pi}} \left[\hat{B}(\pi) - \hat{B}^{k-1}(\pi) + \hat{B}^{k-1}(\pi) dF(\pi) - \cdots - \hat{B}^0(\pi) + \hat{B}^0(\pi) \right] dF^*(\pi) \\
<& \frac{\epsilon}{2} + \int_0^{\bar{\pi}} B^0(\pi) dF(\pi) - \int_0^{\bar{\pi}} \hat{B}^0(\pi) dF^*(\pi),
\end{aligned}$$

where, using step 3, we chose $B^j \in Baire^j(\Pi)$ such that, for all $j \in \{0, 1, \ldots, k-1\}$,

$$\left| \int_0^{\bar{\pi}} \left[B^{j+1}(\pi) - \tilde{B}^j(\pi) \right] dF(\pi) \right| < \frac{\epsilon}{4k},$$

and

$$\left| \int_0^{\bar{\pi}} \left[\hat{B}^{j+1}(\pi) - \hat{B}_n^j(\pi) \right] dF^*(\pi) \right| < \frac{\epsilon}{4k}.$$

5. But B^0 is continuous, so $\hat{B}^0 = B^0$, since the lower semicontinuous envelope, $\hat{f}$, of a function, f, coincides with the function exactly at the points of lower semicontinuity.

6. Lemma B.3, proven below, shows that $F^* \in \overline{\mathcal{F}(B)}$. Then, since $\mathcal{F}(B)$ is dense in $\overline{\mathcal{F}(B)}$ (with respect to the weak* topology), there exists $F^\epsilon \in \mathcal{F}(B)$ such that, for all $\tilde{B}^0$ continuous,

$$\left| \int_0^{\bar{\pi}} \tilde{B}^0(\pi) d\left(F^\epsilon - F^*\right)(\pi) \right| < \frac{\epsilon}{2},$$

and thus, we have

$$\int_0^{\bar{\pi}} B(\pi) dF^\epsilon(\pi) - \int_0^{\bar{\pi}} \hat{B}(\pi) dF^*(\pi) < \epsilon,$$

for this $F^\epsilon \in \mathcal{F}(B)$. □

Lemma B.3. *Let $\mathcal{F}(B)$ and $\mathcal{F}(\hat{B})$ be defined as above. Then $F(\hat{B}) \subseteq \overline{\mathcal{F}(B)}$*

Proof. We first show that the lower semicontinuous envelope of $F \to \int_0^{\overline{\pi}} B(\pi)dF(\pi)$ is given by $\int_0^{\overline{\pi}} \hat{B}(\pi)dF(\pi)$. To see this, note that $\hat{B} \leq B$ implies that

$$\int_0^{\overline{\pi}} \hat{B}(\pi)dF(\pi) \leq \int_0^{\overline{\pi}} B(\pi)dF(\pi)$$

from monotonicity of the integral. Furthermore, from the Portmanteau theorem, $F \to \int_0^{\overline{\pi}} \hat{B}(\pi)dF(\pi)$ is lower semicontinuous. So, by definition of the lower semicontinuous envelope,

$$\int_0^{\overline{\pi}} \hat{B}(\pi)dF(\pi) \leq lsc\left(\int_0^{\overline{\pi}} B(\pi)d\cdot(\pi)\right)(F).$$

To prove the reverse inequality, suppose for a contradiction that there exists $F \in \Delta(\Pi)$ and $\overline{\epsilon} > 0$ such that

$$lsc\left(\int_0^{\overline{\pi}} B(\pi)d\cdot(\pi)\right)(F) - \overline{\epsilon} > \int_0^{\overline{\pi}} \hat{B}(\pi)dF(\pi).$$

Then, by lemma 1.32 (part iv) in Bauschke and Combettes (2011),**?** there exists an open neighborhood V of F such that, for all $\tilde{F} \in V$,

$$\int_0^{\overline{\pi}} B(\pi)d\tilde{F}(\pi) > \int_0^{\overline{\pi}} \hat{B}(\pi)dF(\pi) + \overline{\epsilon}.$$

Furthermore, by the same lemma 1.32 (part iv),

$$\hat{B}(\pi) = \lim_{\tilde{\epsilon}\to 0^+} \inf_{\pi' \in \{\pi' : |\pi - \pi'| < \tilde{\epsilon}\}} B(\pi').$$

Therefore, for all $\epsilon' > 0$, there exists $\tilde{\epsilon}$ such that, $\forall \overline{\epsilon} \min\{\tilde{\epsilon}, \epsilon'\}$,

$$0 \leq \inf_{\pi' \in \{\pi' : |\pi - \pi'| < \tilde{\epsilon}\}} B(\pi') - \hat{B}(\pi) < \frac{1}{2}\epsilon'.$$

Furthermore, by the definition of the infimum, for all $\epsilon' > 0$, there exists π^* such that

$$0 < B(\pi^*) - \inf_{\pi' \in \{\pi' : |\pi - \pi'| < \tilde{\epsilon}\}} B(\pi') < \frac{1}{2}\epsilon'$$

and $|\pi - \pi^*| < \tilde{\epsilon}$. Therefore

$$B(\pi^*) - \hat{B}(\pi) \leq \left(B(\pi^*) - \inf_{\pi' \in \{\pi' : |\pi - \pi'| < \tilde{\epsilon}\}} B(\pi')\right) + \left(\inf_{\pi' \in \{\pi' : |\pi - \pi'| < \tilde{\epsilon}\}} B(\pi') - \hat{B}(\pi)\right)$$
$$< \frac{\epsilon'}{2} + \frac{\epsilon'}{2} = \epsilon'.$$

Summarizing, for any $\epsilon' > 0$ and any $\pi \in \Pi$, there exists $\pi^* \in \Pi$ such that $|\pi - \pi^*| < \epsilon'$ and $B(\pi^*) - \hat{B}(\pi) < \epsilon'$.

Let $\epsilon > 0$ and consider a partition of Π, $\mathcal{P}_N = \{[\pi_0, \pi_1]\} \cup \{(\pi_{i-1} \cup \pi_i\}_{i=2}^N$, so that $N < \infty$ and, for all $i \in \{0, 1, \ldots, N-1\}$, $\pi_{i+1} - \pi_i < \frac{1}{2}\epsilon$ and $F(\pi_i) = F(\pi_i^-)$, i.e., π_i are not mass point of F, except possibly π_0.[34] Let $\pi_i' = \arg \underset{\pi \in [\pi_i, \pi_{i+1}]}{\hat{B}(\pi)}$, which is well defined by the theorem of the maximum since $\hat{B}$ is lower semicontinuous.

By the above argument, find π^* such that $|\pi'^*| < \frac{1}{2}\epsilon$ and $B(\pi_i^*) - \hat{B}(\pi_i') < \frac{\bar{\epsilon}}{2}$. Now let $f_\epsilon(\pi_i^*) = F(\pi_{i+1}) - F(\pi_i)$ and define

$$F_\epsilon(\pi) = \sum_{\pi_j^* \leq \pi} f_\epsilon(\pi_j^*).$$

Observe that, as $\epsilon \to 0$, $F_\epsilon \to F$ in the weak* topology. This implies that $F_\epsilon \in V$ for some $\epsilon > 0$ sufficiently small. Furthermore

$$\begin{aligned}
&\int_0^{\bar{\pi}} B(\pi) dF_\epsilon(\pi) - \int_0^{\bar{\pi}} \hat{B}(\pi) dF(\pi) \\
=& \sum_{i=0}^{N-1} \left(F(\pi_{i+1}) - F(\pi_i)\right) B(\pi^*) - \int_0^{\bar{\pi}} \hat{B}(\pi) dF(\pi) \\
\leq& \sum_{i=0}^{N-1} \left(F(\pi_{i+1}) - F(\pi_i)\right) \left(B(\pi^*) - \hat{B}(\pi_i')\right) \\
\leq& \frac{1}{2}\bar{\epsilon}.
\end{aligned}$$

So

$$\int_0^{\bar{\pi}} B(\pi) dF_\epsilon(\pi) \leq \int_0^{\bar{\pi}} \hat{B}(\pi) dF(\pi) + \frac{1}{2}\bar{\epsilon},$$

which is a contradiction. Therefore, this shows that

$$\int_0^{\bar{\pi}} \hat{B}(\pi) dF(\pi) \geq lsc\left(\int_0^{\bar{\pi}} B(\pi) d\cdot(\pi)\right)(F),$$

and thus that:

$$\int_0^{\bar{\pi}} \hat{B}(\pi) dF(\pi) = lsc\left(\int_0^{\bar{\pi}} B(\pi) d\cdot(\pi)\right)(F).$$

Now, by the definition of the upper semicontinuous envelope, $F \to \int_0^{\bar{\pi}} \left(\pi - \hat{B}(\pi)\right) dF(\pi)$ is the upper semicontinuous envelope of $F \to \int_0^{\bar{\pi}} (\pi - B(\pi)) dF(\pi)$. Then, by lemma 1.32 (part vi) in Bauschke and Combettes (2011),**?** the hypograph of $F \to \int_0^{\bar{\pi}} \left(\pi - \hat{B}(\pi)\right) dF(\pi)$

[34] We can indeed choose such π_i's as there can only be countably many mass point by Froda's theorem.

is the topological closure of the hypograph of $F \to \int_0^{\bar{\pi}} (\pi - B(\pi))\, dF(\pi)$; i.e.,

$$\overline{\left\{(F,\mu) \in \Delta(\Pi) \times \mathbb{R} \,:\, \mu \leq \int_0^{\bar{\pi}} \left(\pi - \hat{B}(\pi)\right) dF(\pi)\right\}} = \left\{(F,\mu) \in \Delta(\Pi) \times \mathbb{R} \,:\, \mu \leq \int_0^{\bar{\pi}} (\pi - B(\pi))\, dF(\pi)\right\}.$$

Finally, observe that

$$\begin{aligned}\mathcal{F}(\hat{B}) &= Proj_{\Delta(\Pi)}\left(\left\{(F,\mu) \in \Delta(\Pi) \times \mathbb{R} \,:\, \mu \leq \int_0^{\bar{\pi}} \left(\pi - \hat{B}(\pi)\right) dF(\pi)\right\} \cap [\{F \leq G\} \times \{V_A(B \mid F_0)\}]\right) \\ &\subseteq \overline{Proj_{\Delta(\Pi)}\left(\left\{(F,\mu) \in \Delta(\Pi) \times \mathbb{R} \,:\, \mu \leq \int_0^{\bar{\pi}} (\pi - B(\pi))\, dF(\pi)\right\} \cap [\{F \leq G\} \times \{V_A(B \mid F_0)\}]\right)} \\ &= \overline{\mathcal{F}(B)},\end{aligned}$$

since the projection operator is continuous in the product topology (by definition) and the image of the closure of a set by a continuous mapping is included in the closure of the image of this set. □