Robustness and Linear Contracts

Gabriel Carroll (Working Paper, Dec 2012) summary by N. Antić

A principal contracts with an agent, who affects the principal's payoff by taking a costly, private action. Both paries are risk-neutral.

Basic Model

- Set of output values $Y \subset \mathbb{R}$ is compact and $\min(Y) = 0$
- An *action* is $(F, c) \in \Delta(Y) \times \mathbb{R}_+$; a *technology* is a set of actions available to the agent, $\mathcal{A} \subset \Delta(Y) \times \mathbb{R}_+$, assumed to be compact
  - The agent knows $\mathcal{A}$, but the principal knows some $\mathcal{A}_0 \subset \mathcal{A}$
  - (A1) Assume principal knows there are benefits from contracting, i.e., $\exists (F, c) \in \mathcal{A}_0$ for which $\mathbb{E}_F[y] - c > 0$
  - The principal believes $\mathcal{A}$ can be any superset of $\mathcal{A}_0$
  - Assume $(\delta_0, 0) \in \mathcal{A}_0$, i.e., agent can always do nothing
  - $\mathcal{A}_0$ has *full support* if $\forall (F, c) \in \mathcal{A}_0 \setminus (\delta_0, 0)$, $supp(F) = Y$
- A *contract* is $w : Y \to \mathbb{R}_+$, w cts (limited liability assumption)
  - Let w_α denote linear contracts, i.e., $w_\alpha(y) = \alpha y$, $\forall y \in Y$
- Timing of the game and payoffs are as follows:
  1. Principal offers contract w knowing $\mathcal{A}_0$
  2. Agent chooses action $(F, c) \in \mathcal{A}$
  3. Nature determines output $y \sim F$
  4. Payoffs are $y - w(y)$ for principal and $w(y) - c$ for agent
- Note that the agent chooses an action in the non-empty set:

$$A^*(w \mid \mathcal{A}) = \underset{(F,c)\in\mathcal{A}}{\arg\max}\left(\mathbb{E}_F[w(y)] - c\right)$$

  - Let $V_A(w \mid \mathcal{A})$ be the value function of the above
  - If $|A^*(w \mid \mathcal{A})| \neq 1$, agent maximizes principal's utility
- The principal is extremely ambiguity averse and maximizes:

$$V_P(w) = \inf_{\mathcal{A}\supset\mathcal{A}_0} V_P(w \mid \mathcal{A}) = \inf_{\mathcal{A}\supset\mathcal{A}_0}\left(\max_{(F,c)\in A^*(w\mid\mathcal{A})} \mathbb{E}_F[y - w(y)]\right)$$

- Note that the principal can guarantee himself a strictly positive payoff by using a linear contract w_α for some $\alpha \in [0, 1)$
  - Let $(\underline{F}, \underline{c})$ solve $\max_{(F',c')\in\mathcal{A}_0} \mathbb{E}_{F'}[\alpha y] - c' = V_A(w_\alpha \mid \mathcal{A}_0)$
  - Note that $V_A(w_\alpha \mid \mathcal{A}_0) > 0$ for α close to 1 by A1
  - For any $\mathcal{A} \supset \mathcal{A}_0$, for any (F, c) the agent chooses we have $\alpha\mathbb{E}_F[y] \geq \alpha\mathbb{E}_F[y] - c \geq \alpha\mathbb{E}_{\underline{F}}[y] - \underline{c} = V_A(w_\alpha \mid \mathcal{A}_0)$; thus:

$$V_p(w_\alpha) = (1-\alpha)\mathbb{E}_F[y] \geq \frac{1-\alpha}{\alpha} V_A(w_\alpha \mid \mathcal{A}_0) > 0 \qquad (1)$$

- If $\exists (F, 0) \in \mathcal{A}_0 \setminus (\delta_0, 0)$, w_0 can attain positive profits since $V_p(w_0) = \max_{(F,0)\in\mathcal{A}_0} \mathbb{E}_F[y] > 0$, if not then $V_p(w_0) = 0$
- Focus on contracts which perform better than w_0

Lemma. *For any $w \neq w_0$ for which $V_P(w) \geq V_P(w_0)$ we have:*

$$\begin{aligned} V_P(w) &= \min_{\{F\in\Delta(Y)\}} \mathbb{E}_F[y - w(y)] \qquad (2) \\ & \quad \textit{subject to } \mathbb{E}_F[w(y)] \geq V_A(w \mid \mathcal{A}_0). \end{aligned}$$

If $V_P(w) > 0$, the constraint binds for F attaining the minimum.

Proof. ($\geq$) For all $\mathcal{A} \supset \mathcal{A}_0$, any $(F, c) \in A^*(w \mid \mathcal{A})$ satisfies $\mathbb{E}_F[w(y)] \geq \mathbb{E}_F[w(y)] - c \geq V_A(w \mid \mathcal{A}_0)$.

($\leq$) Let F be the argmin of problem 2 and consider two cases:

1: *supp*$(F) \not\subset \arg\max_y w(y)$. Let $F'_\varepsilon \equiv (1-\varepsilon)F \oplus \varepsilon\delta_{y^*}$ for $y^* \in \arg\max_y w(y)$, so that if $\mathcal{A} = \mathcal{A}_0 \cup (F'_\varepsilon, 0)$, $A^*(w \mid \mathcal{A}) = (F'_\varepsilon, 0)$ and $V_P(w) = (1-\varepsilon)\mathbb{E}_F[y - w(y)] + \varepsilon\mathbb{E}_F[y^* - w(y^*)]$. As $\varepsilon \to 0$, $V_P(w) \to \mathbb{E}_F[y - w(y)]$, and thus $V_P(w) \not> \mathbb{E}_F[y - w(y)]$.

2: *supp*$(F) \subset \arg\max_y w(y)$. If $\mathbb{E}_F[w(y)] > V_A(w \mid \mathcal{A}_0)$, then for $\mathcal{A} = \mathcal{A}_0 \cup (F, 0)$, $A^*(w \mid \mathcal{A}) = (F, 0)$ and $V_P(w) = \mathbb{E}_F[y - w(y)]$. If $\mathbb{E}_F[w(y)] = V_A(w \mid \mathcal{A}_0) = \max_y w(y)$, then $K \equiv \{(G, 0) \in \mathcal{A}_0 : supp(G) \subset \arg\max_y w(y)\} \neq \emptyset$ and

$$\begin{aligned} V_P(w) &\leq V_P(w \mid \mathcal{A}_0) = \max_{(G,0)\in K} \mathbb{E}_G[y] - \max_y w(y) \\ &< \max_{(G,0)\in\mathcal{A}_0} \mathbb{E}_G[y] = V_P(w_0), \quad \Rightarrow\Leftarrow. \end{aligned}$$

Now assume $V_P(w) > 0$ and let F be the argmin of 2. If $\mathbb{E}_F[w(y)] > V_A(w \mid \mathcal{A}_0)$ consider $F_\varepsilon \equiv (1-\varepsilon)F \oplus \varepsilon\delta_0$ for small ε so that $\mathbb{E}_{F_\varepsilon}[w(y)] > V_A(w \mid \mathcal{A}_0)$. Now $\mathbb{E}_{F_\varepsilon}[y - w(y)] = (1-\varepsilon)\mathbb{E}_F[y - w(y)] + \varepsilon(0 - w(0)) < \mathbb{E}_F[y - w(y)]$, $\Rightarrow\Leftarrow$. □

- $\forall \alpha > 0$, if $V_P(w_\alpha) \geq V_P(w_0)$ and $V_P(w_\alpha) > 0$ then $V_P(w_\alpha) = \frac{1-\alpha}{\alpha} V_A(w_\alpha \mid \mathcal{A}_0) = \max_{(F,c)\in\mathcal{A}_0}\left((1-\alpha)\mathbb{E}_F[y] - \frac{1-\alpha}{\alpha}c\right)$

Theorem. *A linear contract, w_α for some α, maximizes V_P. If $\mathcal{A}_0$ has full support then every contract maximizing V_P is linear.*

Proof. Take w s.t. $V_P(w) \geq V_P(w_0)$, $V_P(w) > 0$ and find w_α s.t. $V_P(w_\alpha) \geq V_P(w)$. Let $S = \text{conv}\{(w(y), y - w(y)) : y \in Y\}$, $T = \{(u, v) : u > V_A(w \mid \mathcal{A}_0), v < V_P(w)\}$ and note that by the lemma $S \cap T = \emptyset$. Thus $\exists [(\lambda, \mu) = \kappa] \equiv \{x \in \mathbb{R}^2 : x \cdot (\lambda, \mu) = \kappa\}$, a separating hyperplane which satisfies:

$$\begin{aligned} \lambda u + \mu v \leq \kappa \quad \forall (u, v) \in S \qquad (3) \\ \lambda u + \mu v \geq \kappa \quad \forall (u, v) \in T \qquad (4) \end{aligned}$$

with $\kappa \geq 0$, $\lambda > 0$, $\mu < 0$. This is illustrated below, with the line $\lambda u + \mu v = \kappa$ in green and $\{(w(y), y - w(y)) : y \in Y\}$ in red.

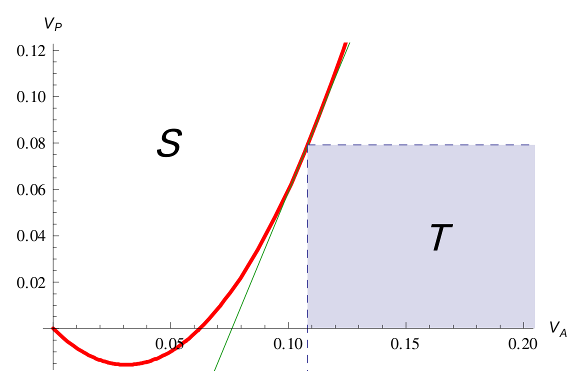

Let $\alpha \equiv \frac{-\mu}{\lambda-\mu} \in (0, 1)$. Note that w_α has the same incentives as the affine contract $w'(y) = w_\alpha(y) + \frac{\kappa}{\lambda-\mu} = \frac{\kappa-\mu y}{\lambda-\mu} \geq w(y)$ for all y, where the inequality follows by expression 3. Note that $V_P(w_\alpha) \geq V_P(w')$.

We are left to show $V_P(w') \geq V_P(w)$ and $V_P(w') > V_P(w)$ if $\mathcal{A}_0$ has full support. For any $\mathcal{A}$ and any $(F, c) \in A^*(w' \mid \mathcal{A})$, $\mathbb{E}_F[w'(y)] \geq \mathbb{E}_F[w'(y)] - c = V_A(w' \mid \mathcal{A}_0) \geq V_A(w \mid \mathcal{A}_0)$. Now:

$$\begin{aligned} V_P(w' \mid \mathcal{A}) &= \mathbb{E}_F[y - w'(y)] = \mathbb{E}_F\left[\frac{\lambda w'(y) - \kappa}{-\mu}\right] \\ &\geq \frac{\lambda V_A(w' \mid \mathcal{A}_0) - \kappa}{-\mu} \geq \frac{\lambda V_A(w \mid \mathcal{A}_0) - \kappa}{-\mu} \\ &= \frac{\lambda\mathbb{E}_{F^*}[w(y)] - \kappa}{-\mu} = \mathbb{E}_{F^*}[y - w(y)] = V_P(w), \end{aligned}$$

where F^* is the argmin of problem 2 and thus the pair $(\mathbb{E}_{F^*}[w(y)], \mathbb{E}_{F^*}[y - w(y)]) \in \overline{S}, \overline{T}$, so that $\lambda \mathbb{E}_{F^*}[w(y)] + \mu \mathbb{E}_{F^*}[y - w(y)] = \kappa$. Thus $V_P(w') \geq V_P(w)$.

Now for any $(F, c) \in A^*(w \mid \mathcal{A}_0) \not\supset \{(\delta_0, 0)\}$, if F has full support, then $\mathbb{E}_F[w'(y)] > \mathbb{E}_F[w(y)]$ unless $w = w'$, i.e., $V_A(w' \mid \mathcal{A}_0) > V_A(w \mid \mathcal{A}_0)$; thus $V_P(w_\alpha) \geq V_P(w') > V_P(w)$. □

- The optimal contract is found by solving:

$$\max_{(F,c)\in\mathcal{A}_0,\ \alpha\in[0,1]} (1-\alpha)\,\mathbb{E}_F[y] - \frac{1-\alpha}{\alpha}c$$

  - For any (F, c) choose $\alpha = \sqrt{\frac{c}{\mathbb{E}_F[y]}}$, thus solve:

$$\max_{(F,c)\in\mathcal{A}_0} \left(\sqrt{\mathbb{E}_F[y]} - \sqrt{c}\right)^2$$

Extensions

- The main result can be extended to more complicated settings:
  - Non-zero participation constraints for the agent
  - Somewhat more general assumptions about the principal's knowledge of $\mathcal{A}$
  - Lower-bounds on c which are functions of the expectation of F
    - Generalization to lower-bounds on c which depend on any moment of F
  - Risk-averse or risk-loving preferences, with contracts which are linear in utility
- Note that the lower bound on the principal's payoff when she knows $\mathcal{A}$ is strictly above $V_P(w_{\alpha^*})$
  - Screening by asking the agents to report $\mathcal{A}$?
- Interestingly, a menu of contracts, $\mathcal{W} = (w_\mathcal{A})$, does not beat a single (linear) contract if agent's IC needs to be satisfied:

$$V_A(w_\mathcal{A} \mid \mathcal{A}) \geq V_A(w_{\mathcal{A}'} \mid \mathcal{A}) \qquad \forall \mathcal{A}, \mathcal{A}' \supset \mathcal{A}_0 \qquad (5)$$

  - Principal's objective is $V_P(\mathcal{W}) = \inf_{\mathcal{A}\supset\mathcal{A}_0} V_P(w_\mathcal{A} \mid \mathcal{A})$

Theorem (3.3). *For any $\mathcal{W} = (w_\mathcal{A})$, $V_P(\mathcal{W}) \leq \max_w V_P(w)$.*

Proof. Let $w^0 \in \mathcal{W}$ be the contract chosen by the agent under technology $\mathcal{A}_0$ and assume by way of contradiction $V_P(w^0) < V_P(\mathcal{W})$. Then $\exists \mathcal{A}_1$ s.t. the agent chooses $(F_1, c_1) \notin \mathcal{A}_0$ given w^0 and $V_P(w^0 \mid \mathcal{A}_1) < V_P(\mathcal{W})$. WLOG let $\mathcal{A}_1 = (F_1, c_1) \cup \mathcal{A}_0$. Let $w^1 \in \mathcal{W}$ be the contract chosen by agent under $\mathcal{A}_1$. To see that $A^*(w^1 \mid \mathcal{A}_1) = \{(F_1, c_1)\}$ assume by way of contradiction $(F_0, c_0) \in \mathcal{A}_0$ is in $A^*(w^1 \mid \mathcal{A}_1)$, which by IC implies $V_A(w^1 \mid \mathcal{A}_1) \leq V_A(w^0 \mid \mathcal{A}_0)$, but $V_A(w^1 \mid \mathcal{A}_1) \geq V_A(w^0 \mid \mathcal{A}_1) > V_A(w^0 \mid \mathcal{A}_0)$, $\Rightarrow\Leftarrow$. Thus:

$$\begin{aligned} V_P(w^1 \mid \mathcal{A}_1) &= \mathbb{E}_{F_1}[y - w^1(y)] \\ &= \mathbb{E}_{F_1}[y] - c_1 - \left(\mathbb{E}_{F_1}[w^1(y)] - c_1\right) \\ &\leq \mathbb{E}_{F_1}[y] - c_1 - \left(\mathbb{E}_{F_1}[w^0(y)] - c_1\right) \\ &= \mathbb{E}_{F_1}[y - w^0(y)] = V_P(w^0 \mid \mathcal{A}_1) \\ &< V_P(\mathcal{W}), \end{aligned}$$

which contradicts the definition of $V_P(\mathcal{W})$. □

- Although the principal's lower bound when she knows $\mathcal{A}$ is higher, there are technology sets for which the linear contract w_{α^*} is optimal; this is shown in Appendix C of the paper

Jan 16, 2013 Applied Theory Working Group Discussion

Attendees: Sylvain Chassang, Stephen Morris, Ben Brooks, Konstantinos Kalfarentzos, Kai Steverson, Olivier Darmouni, Dan Zeltzer, Nemanja Antić

- Main result could have a more "constructive" proof, although the trick of using the separating hyperplane theorem is particularly useful when proving extensions
- Extension to risk-aversion is somewhat uninteresting as contracts are linear in utility, which is not typically observed
- Extension to other knowledge assumptions is not so general—principal needs to be very uncertain about at least one action; if there are many possible actions each of which the principal is not too uncertain about the result fails
- Some economic insight may be found in the optimal linear contract if more structure is imposed
- It is unclear that the limited liability assumption has its usual bite in this setting, since the two "normalizations" of $\min(Y) = 0$ and non-negative payments rely on the same zero and thus agent is unable to destroy value
- Sylvain commented that theorem 3.3 can be generalized; a statement and proof follow some definitions
  - A simple lottery over contracts is $L = \{(p_1, w^1), ..., (p_n, w^n)\}$ with $\sum_{i=1}^n p_i = 1$
  - The timing is such that the lottery is resolved only after the agent takes an action
  - A menu of simple lotteries is $\overline{\mathcal{W}} = (L_\mathcal{A})_{\mathcal{A}\supset\mathcal{A}_0}$
  - $V_P(\overline{\mathcal{W}}) = \inf_{\mathcal{A}\supset\mathcal{A}_0} \sum_{(p_i,w^i)\in L_\mathcal{A}} p_i V_P(w^i \mid \mathcal{A})$

Theorem. *For any $\overline{\mathcal{W}} = (L_\mathcal{A})_{\mathcal{A}\supset\mathcal{A}_0}$, $V_P(\overline{\mathcal{W}}) \leq \max_w V_P(w)$.*

Proof. Let $L \in \overline{\mathcal{W}}$ be the lottery chosen by the agent under $\mathcal{A}_0$ and $w^L = \sum_{i=1}^n p_i w^i$; w^L is clearly a contract. Assume by way of contradiction $V_P(w^L) < V_P(\overline{\mathcal{W}})$.

Note that $V_P(w^L) = V_P(\overline{\mathcal{W}} \mid \mathcal{A}_0)$, since the agent is risk-neutral and hence $V_A(L \mid \mathcal{A}_0) = V_A(w^L \mid \mathcal{A}_0)$. Thus there must be some $\mathcal{A}_1$ s.t. the agent chooses $(F_1, c_1) \notin \mathcal{A}_0$ given w^L and $V_P(w^L \mid \mathcal{A}_1) < V_P(\overline{\mathcal{W}})$. WLOG let $\mathcal{A}_1 = (F_1, c_1) \cup \mathcal{A}_0$. Note that $V_A(L \mid \mathcal{A}_1) = V_A(w^L \mid \mathcal{A}_1) > V_A(w^L \mid \mathcal{A}_0) = V_A(L \mid \mathcal{A}_0)$, since the inequality follows by the argument in theorem 3.3 in the paper.

Let $\overline{L} = \{(\overline{p}_1, \overline{w}^1), ..., (\overline{p}_n, \overline{w}^n)\}$ be chosen by the agent under $\mathcal{A}_1$. To see that $A^*(\overline{L} \mid \mathcal{A}_1) = \{(F_1, c_1)\}$ assume $(F_0, c_0) \in \mathcal{A}_0$ is in $A^*(\overline{L} \mid \mathcal{A}_1)$, which by IC implies $V_A(\overline{L} \mid \mathcal{A}_1) \leq V_A(L \mid \mathcal{A}_0)$, but $V_A(\overline{L} \mid \mathcal{A}_1) \geq V_A(L \mid \mathcal{A}_1) > V_A(L \mid \mathcal{A}_0)$, $\Rightarrow\Leftarrow$. Thus:

$$\begin{aligned} V_P(\overline{L} \mid \mathcal{A}_1) &= \mathbb{E}_{F_1}\left[y - \textstyle\sum_{i=1}^n \overline{p}_i \overline{w}^i(y)\right] \\ &= \mathbb{E}_{F_1}[y] - c_1 - \left(\mathbb{E}_{F_1}\left[\textstyle\sum_{i=1}^n \overline{p}_i \overline{w}^i(y)\right] - c_1\right) \\ &\leq \mathbb{E}_{F_1}[y] - c_1 - \left(\mathbb{E}_{F_1}\left[\textstyle\sum_{i=1}^n p_i w^i(y)\right] - c_1\right) \\ &= \mathbb{E}_{F_1}[y - w^L(y)] = V_P(w^L \mid \mathcal{A}_1) < V_P(\overline{\mathcal{W}}), \end{aligned}$$

but principal must get $V_P(\overline{\mathcal{W}})$ on any $\mathcal{A}_0$, $\Rightarrow\Leftarrow$. □

- Extension to general lotteries should be straightforward
- Timing does not matter—if lottery is drawn prior to the agent choosing an action, the principal does weakly worse:

$$\inf_{\mathcal{A}\supset\mathcal{A}_0} \sum_{(p_i,w^i)} p_i V_P(w^i \mid \mathcal{A}) \geq \sum_{(p_i,w^i)} p_i \left[\inf_{\mathcal{A}\supset\mathcal{A}_0} V_P(w^i \mid \mathcal{A})\right]$$

- The above may not be true if principal is Bayesian about $\mathcal{A}_0$