Existence of Optimal Mechanisms in Principal-Agent Problems

Ohad Kadan, Phil Reny, Jeroen Swinkels

June 2012

A General Principal-Agent Setup

- A principal employs an agent.
- The agent may have private information (type $t \in T$, prior H).
- Agent reports t' to the principal.
- Principal recommends an action $a' \in A$ (potentially stochastic) and announces a reward scheme.
- The agent takes an action $a \in A$.
- The true action and the type of the agent affect a signal $s \in S$.
  - $P(s|a,t)$
- The principal observes the signal and gives a reward $r \in R$ to the agent (also potentially stochastic).
- Need to satisfy IR, IC.
  - The agent wants to play.
  - The agent has the incentive to report his true type, and follow the recommended action.

Preview of Our Results

- We provide conditions under which principal-agent problems, potentially with both moral-hazard and adverse selection, admit an optimal contract/mechanism.
- We allow for multi-dimensional actions and signals.
- We impose no MLRP or even an order structure on signals or actions.
- Actions, Types, Signals, Rewards can be discrete or continuous.
- Supports of signals can shift with action and type.
- Key assumptions are both limited and each have a natural economic interpretation.

Existence. Who Cares?

- Just technical wonkery?
- Not so fast.
- First, note that the Mirlees counter-example is *economically* sensible, not just something lifted from a "counterexamples in real analysis" book.
- Second, once one departs the FOA, not much known about optimal contracts.
- But then without an existence result, cannot even reason about implications of necessary conditions for optimality.

For Example

- Kadan-Swinkels 2012 “On the Moral Hazard Problem without the First Order Approach”
- Dispenses with the FOA.
- Derives a simple necessary condition for an optimal contract.
    - Essentially, a generalized shadow value as one changes both the minimum wage and the outside option simultaneously.

KS use this condition to derive new comparative statics results

- How does the utility of the principal depend on the wealth of the agent?
- How does optimally induced effort depend on the outside option of the agent?
- How does optimally induced effort depend on the level of a minimum wage, or on the degree of limited liability?

But,

- ... if optimal contracts need not exist, who knows whether any of this is relevant?
- Do "pretty good" contracts "nearly" satisfy our necessary conditions and so have roughly the right comparative statics?
- Results of this paper imply existence (in deterministic contracts) for the KS setting.
- The techniques introduced here are also critical in establishing the necessary differentiability properties of the principal's program, and in particular, in thinking about what it means for the principal's optimum to move continuously in the underlying parameters.

An Example

- Two types, t_1 and t_2.
- Actions $\{out\} \cup [0,1]$.
- Signals $S = \{out_1\} \cup \{out_2\} \cup [0,1]$.
- When t_i chooses *out*, signal is out_i.
- When $a \in (0,1]$,
  - $f(s|a,t_1) = \frac{1}{a}$
  - $f(s|a,t_2) = \frac{1}{a} + 1 - \frac{2}{a}s$,
  - each with support $[0,a]$.
- When $a = 0$, $s = 0$ with probability 1.

An Example

- $a = \frac{1}{4}$

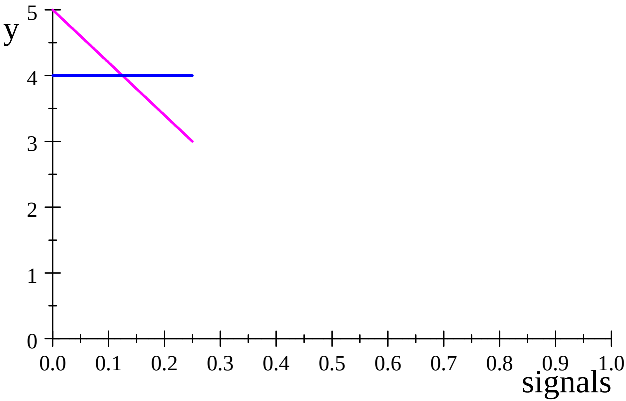

An Example

- t_1 has utility r if he chooses $a \in [0, 1]$ and receives payment r, and utility $1 + r$ from *out*.
- t_2 has utility $\sqrt{r}$ from any action.
- The set of feasible rewards is $[0, 2]$.

An Example

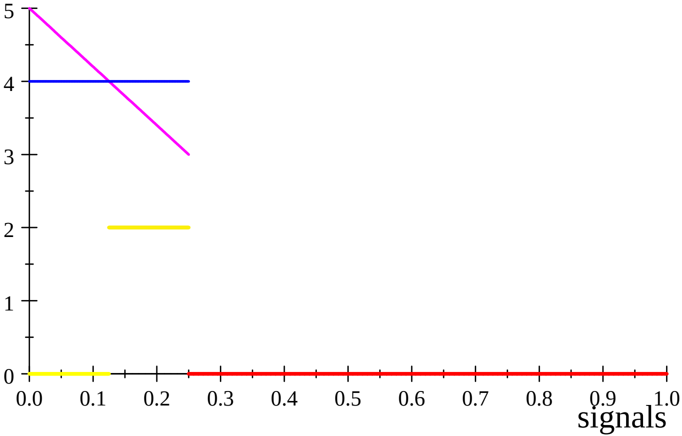

An Example

- Suppose the principal wants to implement $(a_{t_1}, a_{t_2}) = (\hat{a}, out)$.
- Let $I_{\hat{a}} = \left[\frac{\hat{a}}{2}, \hat{a}\right]$
- Compensate the agent according to the function

$$\rho_{\hat{a}}(s) = \begin{cases} 2 & s \in I_{\hat{a}} \\ 0 & s \in [0,1] \backslash I_{\hat{a}} \\ 0 & s = out_1 \\ 2(P(I_{\hat{a}} | \hat{a}, t_2))^2 & s = out_2 \end{cases} .$$

An Example

- Let $\hat{a} \longrightarrow 0$.
- Consider the resulting sequence of contracts $\{\rho_{\hat{a}}\}$.
- $\{\rho_{\hat{a}}\}$ has a pointwise limit ρ_0:
  - Agent is paid 0 if $s \in [0,1] \cup \{out_1\}$, and $1/2$ if $s = out_2$.
- But then type t_1 strictly prefers *out* to $a = 0$.
- Pointwise limit of contracts does not implement the limit action profiles.

- It is not at all clear what set of *economically motivated* assumptions would rule out these difficulties.
- We do not, in particular, see the economic motivation for the restriction that contracts, viewed as *functions* from signals to payments, come from a set that is compact in a topology under which the principal's and the agent's payoffs are continuous.
    - Holmström (1979), Page (1987), Balder (1997)

A Different Limit

- Consider for t_1 and t_2, the probability measure over $R \times S$ induced by $\rho_{\hat{a}}$.
- When t_1 chooses $\hat{a}$,
  - $\frac{1}{2}$ chance signal in $I_{\hat{a}}$ and payment 2,
  - $\frac{1}{2}$ chance signal in $[0,1]\backslash I_{\hat{a}}$ and payment 0.
- This converges (in the weak topology) to
  - $\frac{1}{2}$ chance signal 0 and payment 2,
  - $\frac{1}{2}$ chance signal 0 payment 0.
- Measure over $R \times S$ induced when t_2 chooses out_2 converges to one that assigns probability 1 to a signal equal to out_2 and a payment of $\frac{1}{2}$.

And, indeed, a contract which,

- conditional on $s = 0$ pays a fifty-fifty lottery on $r = 0$ and $r = 2$,
- pays $1/2$ when $s = out_2$,
- otherwise pays 0
  - *is* the (essentially unique) least cost way to implement $(a_{t_1}, a_{t_2}) = (0, out)$.
- Note that the limit measures do not say anything about rewards for signals outside of the support of signals given what t_1 and t_2 are supposed to do.
- Set the rewards here simply to 0.
- Also note (though it does not matter much in this simple example) that the measure we use is what happens when the right type chooses the recommended action.

Takeaways

- Sequences of deterministic contracts, even those with "natural" deterministic limits, may have no economically relevant limit except in mixed contracts.
- Randomized contracts may be necessary to achieve optimality.
- When contracts are interpreted in terms of the equilibrium distributions they induce on reward $\times$ signal pairs, the natural notion of convergence is weak convergence of measures.
- In this example, weak convergence provides a sensible limit, one from which one can "extract" a limit mechanism.
- In the spirit of Milgrom and Weber, 1985

The General Idea

- Starting from standard mechanisms, think instead about the distributions they induce.
- Establish a "back and forth" result.
- Establish compactness and appropriate continuity in distribution space.

Four Key Assumptions

- Utility and cost to the principal bounded below.
  - Sensible in most economic settings.
  - Rules out Mirlees.
- A (weak) continuity condition on information.
  - Will say more about this later.
- An assumption that as utility diverges, it becomes expensive to provide utility at the margin.
  - $\frac{u(w)}{w} \to 0$
- A condition that when the signal is inconsistent with (t', a') – the type the agent announced, and the action the principal recommended – there is a "simple" way to punish the agent that can depend on s, but does not require knowing the details of his true type or action (t, a).
  - trivial in standard settings - just give lowest feasible payment.

Mechanisms

- The set of mechanisms, M, is the set of all pairs (κ, α) such that

(i) $\alpha(\cdot|t) \in \Delta(A)$ for all $t \in T$ and $\kappa(\cdot|s, a, t) \in \Delta(R)$ for all $(s, a, t) \in S \times A \times T$,
(ii) rewards are set to be the worst possible given the signal when the signal is inconsistent with the type announced and action recommended.
(iii) measurability

-
    - Deterministic contracts are a special case.

Pre-Mechanisms

- For any type and action, $\kappa(\cdot|\cdot, a, t)$ generates a distribution ν on $R \times S$ which has marginal $P(\cdot|a, t)$ on signals.
  - Let W be the subset of $\Delta(R \times S) \times A \times T$ consisting of those (ν, a, t) such that ν has marginal $P(\cdot|a, t)$ on $S.W$
- The prior distribution on types, and $\alpha(\cdot|t)$ then generates a distribution on $\Delta(R \times S) \times A \times T$ with the appropriate marginal, H, on types.
  - Let Z be the set of such distributions.

Equivalence

- Say that (κ, α) and $\mu \in Z$ are equivalent if

$$\begin{aligned} & \int f(r,s,a,t)\, d\kappa(r|s,a,t)\, dP(s|a,t)\, d\alpha(a|t)\, dH(t) \\ = & \int f(r,s,a,t)\, dv(r,s)\, d\mu(v,a|t)\, dH(t) \end{aligned}$$

  for every non-negative measurable f.
- Define *IR* for a pre-mechanism in the sensible way.

How About IC for a Pre-Mechanism?

- Imagine some type t' announces t, a is recommended and the pre-mechanism spits out ν.
- Assume t' is contemplating action a'.
- Recall that ν is the distribution over $R \times S$ generated when a, t
- To evaluate the payoff to t' from a', we assume that the distribution $P(\cdot|a', t')$ is absolutely continuous wrt $P(\cdot|a, t)$ where the supports overlap, and use the Radon-Nykodym derivative.
- Off the support of $P(\cdot|a, t)$ we assume the worst reward is used.

Two Key Results

- For each mechanism (κ, α), there is an equivalent pre-mechanism $\mu \in Z$. For each pre-mechanism $\mu \in Z$, there is an equivalent mechanism (κ, α).
- Fix any given c. The subset of Z on which losses are at most c and where IC and IR hold is compact. Losses are lower semi-continuous on this subset.
- From this, existence is immediate.

Where Does Compactness Come From

- We do not assume rewards are finite, or that spaces of types or signals are bounded (A assumed compact).
- For any given (a, t), consider lotteries on $R \times S$ that have marginal $P(\cdot|a, t)$ on S, and that cost no more than some $\bar{c}$.
- These lotteries may make large payments, but can only do so rarely.
- So, under mild conditions, for every $\varepsilon > 0$, I can point to a compact subset of $R \times S$ and say "(r, s) is in this subset with probability at least $1 - \varepsilon$."
- That is, the set of lotteries is a *tight* set of measures, and so by Prohorov's theorem is compact.
- Essentially repeating the same argument, the subset of Z on which expected losses are at most c is compact.

Continuity of Payoffs? A Small Example

- Consider the pre-mechanism μ_n that no matter the announced type, recommended action or realized signal randomizes $\frac{1}{n}$ on paying n, and $\frac{n-1}{n}$ on paying 0.
- In the weak topology, converges to paying 0 always.
- For the principal, expected cost along the sequence is 1, but limit cost is 0.
- Hence, the loss to the principal need not converge.
- However, by the Portmanteau theorem, it either converges or drops at the limit.
- This drop is ok, as the principal is then minimizing something lower semi-continuous.

it would be very bad if the utility of the agent dropped...

ICIR,

- For the agent, expected utility along the sequence is

$$\frac{n-1}{n}u(0) + \frac{1}{n}u(n).$$

- But,

$$\frac{u(n)}{n} \to 0.$$

- So, utility converges nicely!
- Using the assumption that utility is expensive on the margin, this generalizes to show that utility of an agent who announces his type honestly and follows the recommended action is continuous.

Final key part of the proof

- We took convergence relative to equilibrium actions.
- But, how about the utility of a deviant?
- It is here that we need some continuity of information.

Information Continuity. The Example Revisited

- Modify the example only in that when t_2 chooses a, the distribution is the triangular density that is $\frac{2}{a}$ at 0, and 0 at a

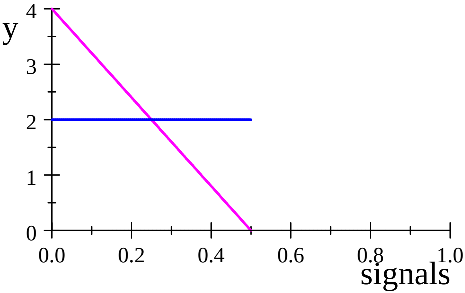

- Least cost way to implement $a_{t_1} = \hat{a} > 0$ and $a_{t_2} = out$ is again to pay 2 on $I_{\hat{a}} = [\frac{1}{2}\hat{a}, \hat{a}]$, $2P^2(I_{\hat{a}}|a, t_2)$ at $s = 2$, and 0 otherwise.
- But, $P(I_{\hat{a}}|\hat{a}, t_2) = \frac{1}{4}$, and so it is only necessary to pay $\frac{1}{8}$ for $s = 2$.

Information Continuity. The Example Revisited

- As before take the limit distribution over $R \times S$ when t_1 chooses action $\hat{a}_n$, and t_2 chooses *out* and use this limit to define a candidate contract.
- This contract, as before, pays a $\left(\frac{1}{2}, \frac{1}{2}\right)$ lottery on $(0, 2)$ when $s = 0$ is observed.
- But, it pays only $\frac{1}{8}$ to $s = 2$, while (as before), a payment of $\frac{1}{2}$ is required for t_2 not to wish to deviate to $a = 0$.
- Hence, in this example, the construction breaks down.

How do the examples differ?

- Think about

$$\frac{P(I_{\hat{a}}|\hat{a}, t_2)}{P(I_{\hat{a}}|\hat{a}, t_1)}$$

- This ratio determines how easy it is to reward t_1 for choosing $\hat{a}$ without also making $\hat{a}$ attractive to t_2.
- In the first example as $a_n \to 0$,

$$\frac{P(I|a_n, t_2)}{P(I|a_n, t_1)} \to \frac{P(\{0\}|0, t_2)}{P(\{0\}|0, t_1)} = 1,$$

and so the trade-off between rewarding t_1 and incentivizing t_2 to deviate moves continuously.

How do the examples differ?

- But, in this example, as $a_n \to 0$,

$$\frac{P(I_n|a_n, t_2)}{P(I_n|a_n, t_1)} \to \frac{1}{2} \text{ while } \frac{P(\{0\}\,|0, t_2)}{P(\{0\}\,|0, t_1)} = 1.$$

- There is a fundamental discontinuity in information.
- The trade-off between rewarding t_1 and incentivizing t_2 is strictly more difficult in the limit than along the sequence.
- Our information condition says that, at least along well chosen subsequences of actions and types, this sort of upward jump does not occur.
- This allows us to show that *IR* survives limits.