Buying Opinions

Mark Whitmeyer and Kun Zhang (Arizona State)

July 19, 2022

Stony Brook International Conference on Game Theory

Introduction

- In many situations, decision makers pay for advice (soft information).
  - examples: sport scouts/headhunters and consulting firms
- A bilateral contracting scenario: principal (P) pays for an agent's (A's) advice.
  - To advise P, A needs to acquire information first.
- Key features:
  - A's information acquisition is flexible, costly and private.
  - A's findings are unverifiable: after acquiring information, A sends a cheap-talk message.
  - P can condition contract on A's message and state.
  - A can take the outside option both before participating (*ex ante*) and after acquiring information (interim).
- Standard moral hazard decomposition:
  1. how to efficiently implement an information acquisition strategy
  2. what strategy to implement

Preview of Findings

- P can implement any feasible information acquisition strategy.
- A's optimal learning pins down the relative incentives (our version of IC).
- When A is risk neutral and no limited liability, any information acquisition strategy can be implemented at first-best cost.
    - Selling the project to the agent does not work!
- Characterization of optimal implementation:
    - limited liability and risk-neutral A: first-best implementation for sufficiently uninformative learning or sufficiently cheap information. Rents for A if first-best infeasible.
    - No limited liability and risk-averse A: first-best infeasible. Rents for A (generically). **(Not today)**

The Model

Model

- P (she) hires A (he) to learn about a payoff-relevant state
  - $\theta \in \Theta = \{\theta_1, \ldots, \theta_n\}$ with $n < \infty$
- P and A share common (WLOG, full support) prior $\mu \in \Delta(\Theta)$
- A can acquire information, flexibly, subject to a cost:
  - A chooses any Bayes-plausible $F \in \Delta\Delta(\Theta)$ and incurs $C(F) = \kappa \int_{\Delta(\Theta)} c \, \mathrm{d}F$
  - $\kappa > 0$ scales the cost
  - $c : \Delta(\Theta) \to \mathbb{R}_+$ is strictly convex, 2x differentiable, bounded on $\operatorname{int} \Delta(\Theta)$, and $c(\mu) = 0$
  - Class includes entropy (Sims 2003), log-likelihood (Pomatto, Strack and Tamuz 2020), and quadratic (Tsallis 1988)

Model

- After acquiring information, A sends a message to P
- True state is *ex post* observable and contractible
- Contract is a pair (M, t):
    - A compact set of messages M available to the agent, and
    - A transfer $t : M \times \Theta \to \mathbb{R}$ ($t : M \times \Theta \to \mathbb{R}_+$ if limited liability)
- This talk: A risk neutral; also consider risk averse agent in the paper
- A has outside option $v_0 \geq 0$
    - A can take this after (M, t) is proposed or after acquiring information

The Contracting Problem

First-Best Benchmark

- Write P's gross payoff as a function of the posterior $\mathbf{x} = \left(x^1, \ldots, x^{n-1}\right)$, $V(\mathbf{x})$
- Denote the set of Bayes-plausible distributions over posteriors by $\mathcal{F}(\mu)$
    - $\mathcal{F}(\mu)$ is a convex and compact subset of $\Delta\Delta(\Theta)$
- If the principal controlled the information acquisition herself, she would solve

$$\max_{F \in \mathcal{F}(\mu)} \int (V - \kappa c) \, dF \, .$$

- First-best: P can observe A's choice of F and specify transfer $t : \Delta\Delta(\Theta) \to \mathbb{R}_+$
    - Cost of acquiring information is $v_0 + C(F)$

Inducing a Distribution

- WLOG for any distribution P wants to implement, M is the support of the distribution
- A contract (M, t) induces a decision problem (μ, M, t) of the agent
  - M is the set of actions, t is the (state-dependent) utility function
  - in the decision problem, the agent acquires information and subsequently sends a message
- A distribution F is **implementable** if there exists a contract (M, t) such that
  1. $M = \text{supp}(F)$, and
  2. it is optimal for the agent to acquire F and report the realized posterior truthfully

The Agent's Decision Problem

- A chooses a distribution over posteriors to maximize her value function $W(\mathbf{x})$ Details
- A's optimal distribution is given by concavifying W: affine function $f_{\mathcal{H}}(\mathbf{x}) : \Delta(\Theta) \to \mathbb{R}$ intersects W at support of the distribution; expected payoff in the contract $f_{\mathcal{H}}(\mu)$
- Set of intersection points of $f_{\mathcal{H}}$ and W is $P_{(M,t)} \Rightarrow F$ can be implemented by (M,t) only if $\text{supp}(F) = P_{(M,t)}$
- The contract must also prevent A from walking away at any point in the interaction
- No double deviations (learn differently and walk away at some belief):

$$f_{\mathcal{H}}(\mathbf{x}) \geq v_0 - \kappa c(\mathbf{x}) \quad \text{for all} \quad \mathbf{x} \in \Delta(\Theta)\,. \tag{IR}$$

- If A cannot walk away after acquiring information, IR is just $f_{\mathcal{H}}(\mu) \geq v_0$

Illustration

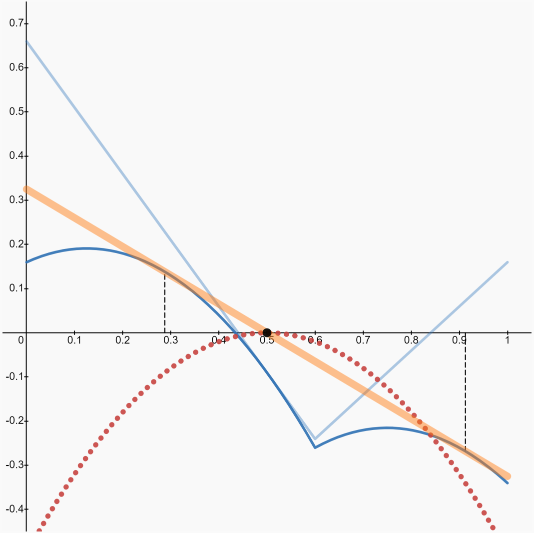

Illustration

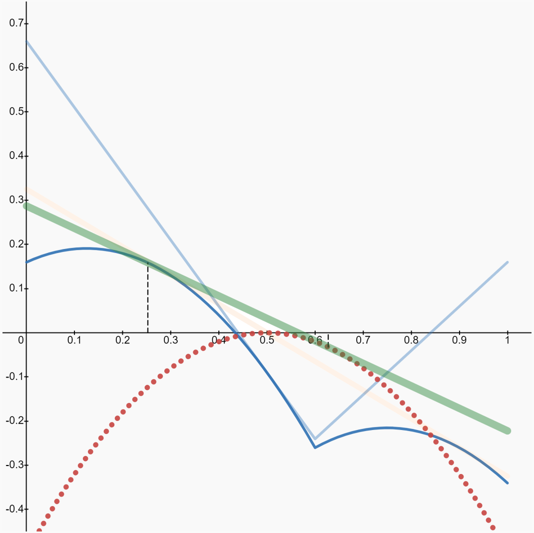

Implementation: Summary

Lemma A contract (M, t) implements a distribution F if and only if

1. $\text{supp}(F) = P_{(M,t)}$; and
2. Constraint *IR* holds; and
3. If there is limited liability, $t(m, \theta) \geq 0$ for all $\theta \in \Theta$ and $m \in M$.

Results

Two Preliminary Results

Lemma If F is a distribution over posteriors with $|\text{supp}(F)| \leq n$ and $\text{supp}(F) \subseteq \text{int}\,\Delta(\Theta)$, there exists a contract (M, t) that implements F, and the expected cost to the principal is finite.

Two Preliminary Results (Cont'd)

Corollary

(1) Every $F \in \mathcal{F}(\mu)$ with $\text{supp}(F) \subseteq \text{int}\,\Delta(\Theta)$ can be induced at a finite cost.

(2) WLOG, P only induces distributions with support on at most n points.

- Any distribution $F \in \mathcal{F}(\mu)$ can be obtained by randomizing over distributions $F_i \in \mathcal{F}(\mu)$ each with support on n or fewer points
- If P randomize first, then implement each F_i as cheaply as possible, same payoff to P gross of cost, but weakly cheaper
- Henceforth focus on distributions with $|\text{supp}(F)| \leq n$

A Big Simplification

- For each state $k = 1, \dots, n$, define $\Omega^k(i, j) := t_i^k - t_j^k$ $(i, j = 1, \dots, s)$.
- Each $\Omega^k(i, j)$ specifies the difference between the payoff to the agent from sending any message i versus message j in state k.

Proposition For an agent to learn according to a desired distribution F, the relative incentives $\left(\Omega^k(i, j)\right)_{i,j=1,\dots,s;\, k=1,\dots,n}$ are pinned down.

- For each state k, P fixes benchmark message $j(k)$, then chooses $\left(t_{j(k)}^k\right)_{k=1}^n$; the payoff to A from sending message $j(k)$ in state k

No Limited Liability

- Efficient (first-best) implementation requires $f_{\mathcal{H}}(\mu) = v_0$
- Thus, Constraint *IR* ($f_{\mathcal{H}}(\mathsf{x}) \geq v_0 - \kappa c(\mathsf{x})$ for all x) must bind at $\mathsf{x} = \mu$
- **Selling the project to the agent?**

Illustration: Optimal Contract without Limited Liability

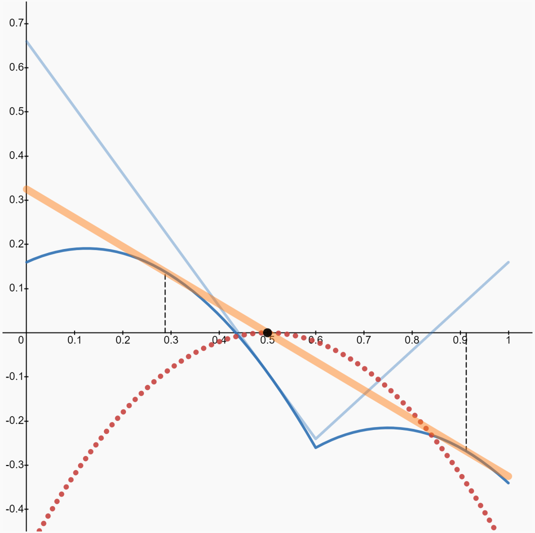

Illustration: Optimal Contract without Limited Liability

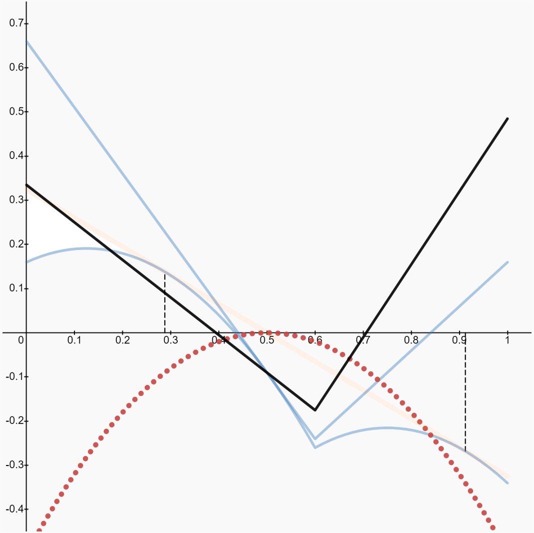

Illustration: Optimal Contract without Limited Liability

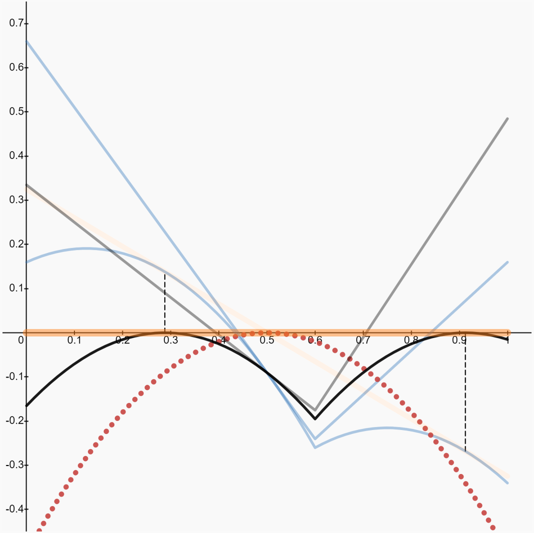

No Limited Liability

- No interim IR $\Rightarrow$ selling the project works. Key thing: $f_{\mathcal{H}}(\mu) = v_0$
- Interim IR $\Rightarrow$ selling the project doesn't work generically: now need $f_{\mathcal{H}}$ tangent to $v_0 - \kappa c$ at μ

Proposition If A is risk neutral and not protected by limited liability, every feasible F with $\text{supp}(F) \subseteq \text{int}\,\Delta(\Theta)$ can be implemented efficiently.

- Not a *shoot the agent contract*: Penalties may be mild
- If either (i) v_0 is sufficiently large, or (ii) implemented distribution sufficiently low in Blackwell order, or (iii) κ is sufficiently small, our construction works under limited liability

Limited Liability (2 States)

To ease exposition, assume $\Theta = \{\theta_1, \theta_2\}$.

Proposition Either

1. $\{x_1, x_2\}$ can be implemented efficiently (and Constraint *IR* binds); or
2. $\{x_1, x_2\}$ cannot be implemented efficiently; and either
   2.1 Constraint *IR* binds and the $t_2^1 = 0$; or
   2.2 Constraint *IR* binds and $t_1^2 = 0$; or
   2.3 Constraint *IR* does not bind and $t_2^1 = t_1^2 = 0$.

- If $\{x_1, x_2\}$ is in the region corresponding to 2.3, same result holds even when A is risk averse

Entropy Reduction Cost: An Example

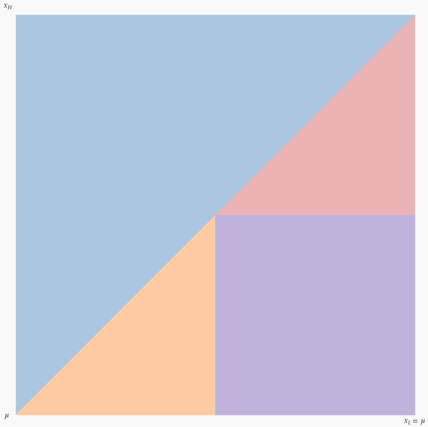

Moderate outside option (or moderate cost of info. acqui.), $\mu = 0.5$

Entropy Reduction Cost: An Example

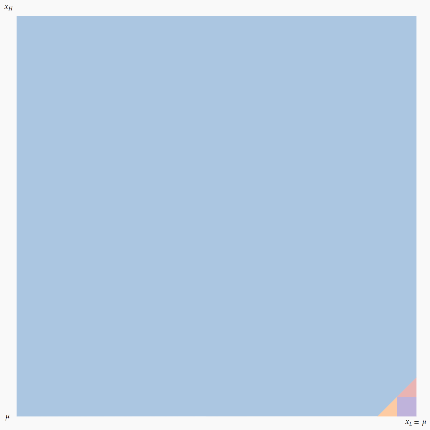

Low outside option (or expensive information), $\mu = 0.5$

Entropy Reduction Cost: An Example

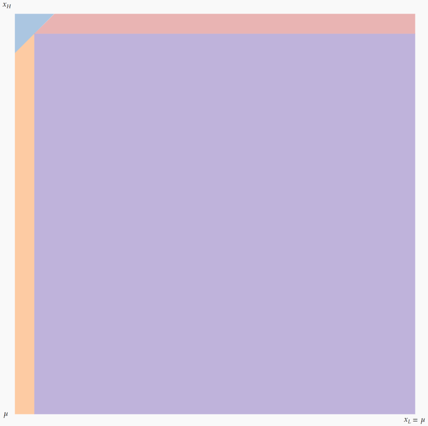

High outside option (or cheap information), $\mu = 0.5$

Related Work

Related Work

- Rappoport and Somma (2017): posteriors are contractible.
  - Hard (them) versus soft (us) information.
- Yoder (Forthcoming): posteriors are contractible, agent's marginal cost of information (κ) is private information.
  - Screening is now important;
  - Contracting on experiment versus posteriors.
- Zermeño (2011), Clark and Reggiani (2021): decision-making delegated to the agent;
  - Can payoffs depend on true state?
  - Decomposition of Pareto optimal contracts.

Thank you!

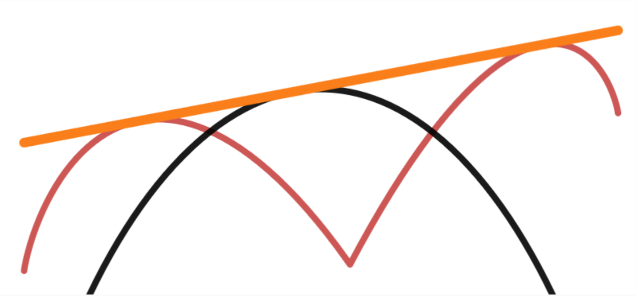

mark.whitmeyer@gmail.com kunzhang@asu.edu

The Agent's Decision Problem: Details

- For any $m \in M$, define A's *net utility* $N(\mathbf{x} \mid m)$:

$$N(\mathbf{x} \mid m) = \sum_{i=1}^{n-1} x^i t(m, \theta_i) + \left(1 - \sum_{i=1}^{n-1} x^i\right) t(m, \theta_n) - \kappa c(\mathbf{x}) \;,$$

  where x^i is the i-th entry of $\mathbf{x} = \left(x^1, \ldots, x^n\right)$.
- The agent's value function is thus $W(\mathbf{x}) = \max_{m \in M} N(\mathbf{x} \mid m)$.

Back

Picking a Point on the "Outside Option Curve"

- With interim IR, problem reduces to picking a point, $\mathbf{x}^*$, on $v_0 - \kappa c(\mathbf{x})$ where $f_{\mathcal{H}}(\mathbf{x})$ is tangent
- Generically $\mathbf{x}^* \neq \mu \Rightarrow$ Agent gets rents
  - Without interim IR, Agent gets no rents
- Efficient implementation is impossible (unless $F = \delta_\mu$)