Learning and Evidence in Principal-Agent Environments

Kym Pram

EUI

July 13, 2018

Motivation

- Many contracting environments feature both hard evidence and learning.
  - Medical tests in insurance contracting.
  - Technical reports in procurement contracting.
  - Qualifications in labor markets.
- New technologies are making this type of information more prevalent and accurate.
  - Example: genetic testing, already a subject of policy debate (Genetic Information Nondiscrimination Act).
- It is important to understand the welfare effects. Both in terms of the positive question "will, e.g., genetic testing benefit consumers or insurance companies, and under what conditions?" and in terms of the normative question "should we allow contracts to depend on endogenous evidence, such as genetic test results?"

Approach

- I consider environments where an agent with no initial private information contracts with a principal.
  - At some cost the agent can acquire *evidence* which both acts as a signal to the agent and can be credibly and voluntarily disclosed to the principal.
  - The principal moves first and must incentivize the agent to acquire or not acquire evidence.
- Firstly, I consider a very general environment and show how the principal's and agent's payoffs depend on the cost of evidence.
- Secondly, I consider an application to the insurance market and show how payoffs of the insurer and the insured change when evidence can or cannot be contracted upon.

Results - General model

- There are two possibilities:
  1. Lower costs of evidence acquisition are best for the principal and the agent is always held to the outside option.
  2. The principal's payoff in non-monotone (U-shaped) in the cost of evidence and the agent may benefit for intermediate costs of evidence.
- In the U-shaped case, aggregate welfare is minimized at a low-intermediate cost of evidence, in the monotone case, aggregate welfare is minimized at extreme costs of evidence acquisition.
- I give conditions on primitives that distinguish the two cases.

Results - Insurance Application

- If the cost of evidence is low and the probability of a high-risk signal is high, then allowing evidence to be contracted upon is Pareto-improving.
- The agent is better off when evidence cannot be contracted upon if (i) costs are in an intermediate range or (ii) costs are low and the probability of a high-risk signal is low.
  - In these cases aggregate welfare may also be lower when evidence can be contracted upon.
- The existence of evidence reduces market efficiency and may distort contracts away from first best even when not used on-path.

General Model - Actions

- A principal faces a single agent with type $t \in T$, initially unknown to both.
- The final outcome space is two dimensional: $A \times \mathbb{R}$.
- The agent can learn, at a cost, the type. If he does so the type can be credibly disclosed.
- The principal moves first and commits to:
  - An outcome for each type, t, that might be disclosed.
  - An outcome if nothing is disclosed.
- The agent has access to an outside option, denoted $0 \in A \times \mathbb{R}$.

General Model - Information

- Initially both the principal and the agent know only the common prior, Π.
- When the agent chooses whether to acquire evidence he knows the mechanism.
  - WLOG the principal does not randomize.
- If the agent acquires evidence then he knows the type before participating in the mechanism.
- If the agent does not acquire evidence he must choose whether to particpate knowing only the prior.

General Model - Payoffs

- The principal's ex-post payoffs are given by $v(a, p, t)$, the agent's by $u(a, p, t)$.
- The principal's payoffs are increasing in p and decreasing in a, the agent's payoffs are increasing in a and decreasing in p.
- If the agent acquires evidence, then, for a function $g(t) : T \to \Delta(A \times P)$. the agent's net ex-ante payoff is given by:

$$E_t[u(g(t), t)] - c$$

  - If $g(t)$ is a lottery, then take $u(g(t), t)$ to be the expected utility. Technical Assumptions

Timing

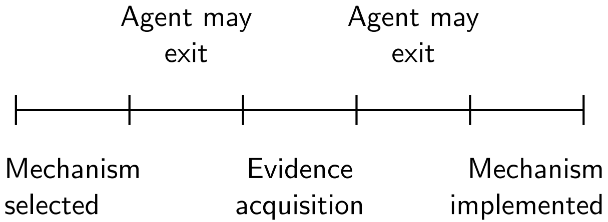
Agent may exit
Agent may exit
Mechanism selected
Evidence acquisition
Mechanism implemented

The mechanism design problem - No acquisition

$$\begin{aligned} \underset{g(N)}{\text{maximize}} \quad & E_t[v(g(N),t)] \\ \text{subject to} \quad & E_t[u(g(N),t)] \geq E_t[u(0,t)] \;\; (IR_N) \\ & E_t[u(g(N),t)] \geq E_t[max\{u(0,t),u(g(N),t)\}] - c \;\; (MH) \end{aligned}$$

- If the cost is high enough, the ex-ante optimal contract is feasible.
- If the cost is below some threshold, $\tilde{c}$, the contract is distorted to prevent learning. The MH constraint is binding.

The mechanism design problem - On-path acquisition

$$
\begin{aligned}
\underset{(g(t))_{t\in T}}{\text{maximize}} \quad & E_t[v(g(t),t)] \\
\text{subject to} \quad & E_t[u(g(t),t)] - c \geq E_t[u(0,t)] \;\; (IR_Y) \\
& u(g(t),t) \geq u(0,t) \;\forall\; t \in T \;\; (EPIR)
\end{aligned}
$$

- Without loss of optimality the agent receives the outside option if no evidence is presented.
- The IR_Y constraint is always binding.

Principal's payoff: U-shaped case

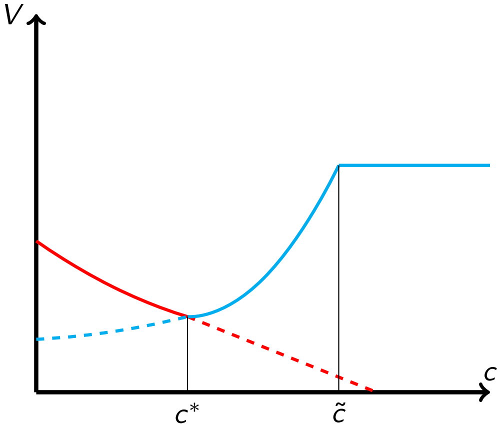

Prinicpal's payoff: Non-increasing case

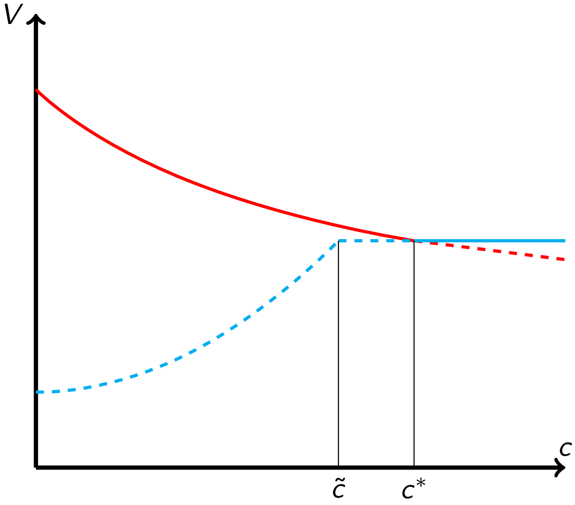

Result: Quasilinear environments

- Here the payoffs are specialized to $v(a, p, t) = V(a, t) + p$, $u(a, p, t) = U(a, t) - p$ and $u(0, 0, t) = 0$.
- Define

$$a_t^* = argmax_a(U(a, t) + V(a, t)).$$

$$a^* = argmax_a E_t[U(a, t) + V(a, t)].$$

$$VI = E_t[U(a_t^*, t) + V(a_t^*, t)] - E_t[U(a^*, t) + V(a^*, t)]$$

Proposition 1

The principal's payoff is non-increasing in c if and only if

$$VI \geq E_t[max\{0, U(a^*, t) - E_t[U(a^*, t)]\}]$$

Otherwise the principal's payoff is U-shaped in cost.

Condition with type-dependent outside option Sketch of proof

Results: General environments

- In the general environment we define the value of information as:

$$VI = max_{(g(t))_{t \in T} \text{ s.t. } EPIR} E_t[v(g(t), t)] - max_{g(N) \text{ s.t. } IR_N} E_t[v(g(N), t)]$$

- This is a generalization of the way value of information was defined in the quasilinear case.
- In general environments the value of information may be negative (Hirshleifer, 1971).
- **Result:** In general environments a modified version of Proposition 1 holds. In particular, the principal's payoff is U-shaped in c whenever $VI < 0$.

Application: Insurance and Genetic Testing

- Question: Suppose genetic tests cannot be used in insurance. How are consumer surplus, profit and efficiency affected?
  - Note: Different benchmark to previous section
- A monopolistic insurer faces a consumer with unknown risk, $\phi \in \{\phi_L, \phi_H\}$ of loss L. The common prior puts probability λ on ϕ_H.
- At cost c the consumer can learn ϕ and can credibly disclose what he has learned to the insurer.
- Special case of the general model with a the level of insurance and p the premium.

Contract with on-path evidence acquisition

- If evidence is acquired on-path then:
  - Each type recieves full insurance, ex-post, that is $a_H(c) = a_L(c) = 1$
  - The consumer is exactly compensated ex-ante for the cost of evidence, that is:

$$\lambda u(a_H(c), p_H(c), H) + (1 - \lambda) u(a_L(c), p_L(c), L) = c$$

  - The low type's consumption level is the certainty equivalent. The high type's premium is discounted. Illustration
- Note: the pair of contracts with endogenous evidence is less responsive to the type than with exogenous evidence. Opposite of the typical finding with non-disclosable endogenous information.
- There is a threshold, $c^* > 0$ below which the principal will prefer to induce evidence acquisition.

On-path evidence acquisition with limits on contracting

- Now we consider the case where evidence cannot be contracted upon, but cost is low enough, $c < \tilde{c}$, that the principal prefers to induce information acquisition.

- Schottmuller and Lagerlof (2016): There exists a threshold probability of the high type, $\lambda^* > 0$ such that:
  - If $\lambda \geq \lambda^*$ the optimal menu features $a_H = 1$ and $(a_L, p_L) = (0, 0)$. The agent receives zero ex-ante rents.
  - If $\lambda < \lambda^*$ the optimal menu features $0 < a_L < 1$ and $a_H = 1$. The agent receives positive ex-ante rents.
- Above a threshold, $\tilde{c} < c^*$, the principal prefers to induce no on-path learning.

Contract with no on-path learning: Case 1

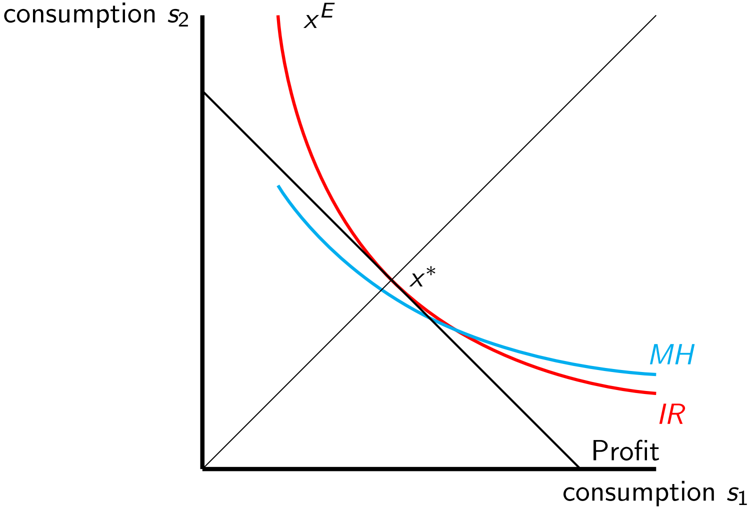

Contract with no on-path learning: Case 2

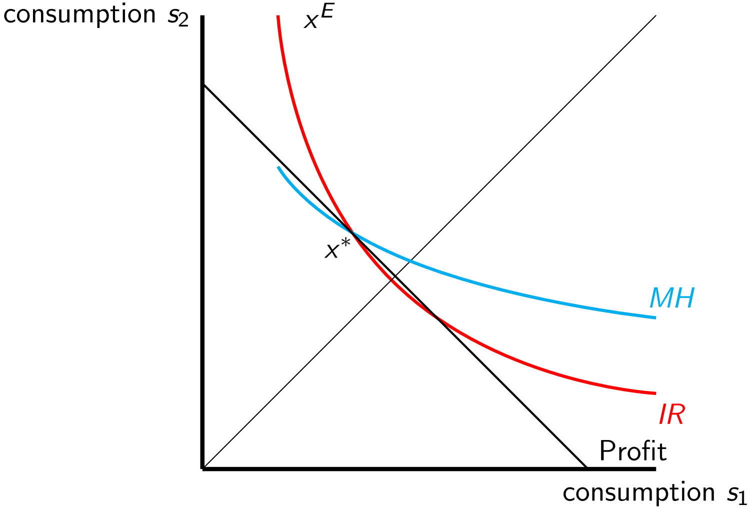

Contract with no on-path learning: Case 3

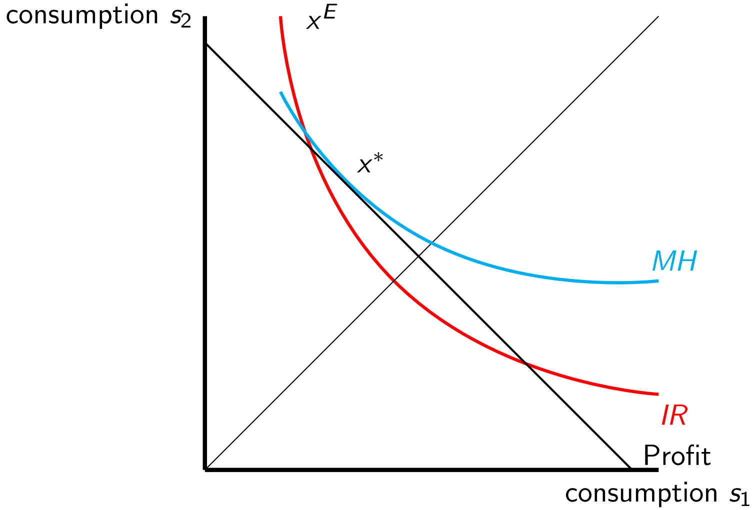

Proof of Optimality

Welfare comparisons

- There exist thresholds $c'' \geq 0$, $c^* > \tilde{c} > 0$:
  - If $c > c^*$ the the equilibrium (pooling) contract is the same whether evidence can be contracted upon or not.
  - If $c \in [c^*, c'']$ or if $c < \tilde{c}$ and λ is sufficiently low, then consumer surplus is strictly higher when evidence cannot be contracted upon.
  - If $c < \tilde{c}$ and λ if sufficiently high or if $c \in [c'', c^*]$ then the equilibrium where evidence can be contracted upon is a Pareto improvement over the equilibrium in which it cannot.
- Effect on aggregate welfare can go either way.

Contrast with endogenous soft information

- While the U-shaped and monotone payoff functions are possible with endogenous soft information, the are not the *only* cases (cf Hoppe and Schmitz, 2010).
- Agent may be best off with low cost of acquisition when endogenous information is non-verifiable.
- High-powered vs low-powered contracts:
  - High-powered contracts are a general feature of optimal contracts with endogenous soft information in several environments.
  - In contrast, I find (weakly) low-powered contracts with endogenous evidence in several examples.
  - This shows that the optimality of high powered contracts does not depend only on incentives for acquisition but also interation with screening incentives.

Conclusions

- The availability of endogenous evidence may benefit the agent and the principal's payoff may be non-monotone in the acquisition cost:
  - In the quasilinear case this occurs if and only if the value of information is less than the expected upside risk in the agent's preferences.
  - In the general case this always occurs if the value of information is negative.
- In monopoly insurance markets, the ability to contract on evidence can lead to a Pareto improvement for some parameter values.
- Conversely, the ability to contract on evidence may reduce aggregate welfare for other parameter values (cannot occur in a competitive market).
- In several examples, optimal contracts with endogenous evidence are "low-powered".

Technical Conditions

- The set A is a compact subset of $\mathbb{R}$.
- The functions $v(\cdot, t)$, $u(\cdot, t)$ are continuous on the product space $A \times P$ for all $t \in T$.
- Let $\phi \in \Delta(A \times P)$ and let $\bar{p} = E_{\phi}(p)$ and $\bar{a} = marg_A \phi$. Then $v(\bar{a}, \bar{p}, t) \geq E_{\phi} v((a, p), t)$ and $v(\bar{a}, \bar{p}, t) \geq E_{\phi} u((a, p), t)$ for all $t \in T$.
- For each $t \in T, a \in A, c \in \mathbb{R}$ there exists $p \in \mathbb{R}$ such that:

$$u(a, p, t) = c.$$

Back

Sketch of Proof: No on-path learning

- It is clear that case 1 is optimal when feasible.
- Similarly, if case 1 is not feasible and case 3 is then case 3 is optimal.
- Using Berge's Maximum theorem and an intermediate value argument we can establish that there is a threshold below which case 3 is optimal.
- If neither case is feasible we can establish that case 2 is optimal. Intuitively, the closest feasible point on IR to full insurance must be the most profitable.
- At the threshold above which case 1 is optimal, case 3 can be shown to be infeasible, establishing the existence of case 2.

Back

Condition with type-dependent outside option

Proposition 1b

The principal's payoff is non-increasing in cost if and only if

$$VI \geq E_t[max\{U(0,t) - E_t U(0,t), U(a^*,t) - E_t[U(a^*,t)]\}]$$

Back

Illustration of mechanism with on-path evidence

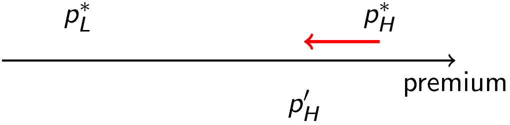

Back

Proof: Quasilinear environments

- Idea of proof: In quasilinear environments the optimal mechanism when evidence is acquired on path sets

$$g(t) = (a_t^*, U(a_t^*) - U(0,t) - c)$$

so that the principal's payoff is $E_t[V(a_t^*, t) + U(a_t^*, t)] - E_t[U(0,t)] - c$.
- The principal's payoff is $U - \textit{shaped}$ if, at $\tilde{c}$, the lowest cost at which $g(N) = (a^*, E_t[U(a^*, t)] - E_t[U(0,t)])$ is feasible, the best mechanism with on-path evidence acquisition does worse than the best mechanism with no acquisition.

Proof: Quasilinear environments

- The threshold cost, $\tilde{c}$ is defined by:

$$E_t[U(0,t)] = E_t[max\{U(0,t), U(a^*,t) - E_t(a^*,t) + E_t[U(0,t)]\}] - \tilde{c}$$

- The right-hand side in the condition comes from this expression.
- The left hand side comes from the difference in gross payoffs (ignorning cost) between the best mechansim with evidence acquisition and the best mechanism without acquisition.

Back

Related literature

- **Mechanism Design with Endogenous Information:** Cremer and Khalil (1992), Cremer et al (1998), Persico (2000), Szalay (2009), Shi (2012)
- **Mechanism Design with Evidence:** Green and Laffont (1986), Bull and Watson (2004, 2007), Ben-Porath and Lipman (2012), Kartik and Tercieux (2012), Koessler and Perez-Richet (2015), Sher and Vohra (2015), McAdams (2011), Pram (2017)
- **Third Degree Price Discrimination:** Schmalensee (1981), Varian (1985), Aguirre, Cowan and Vickers (2010) Schwartz (1990), Bergemann, Brooks and Morris (2014)
- **Welfare Effects of Learning:** Hirshleifer (1971), Roesler and Szentes (2017)

Related literature

- **Monopoly Insurance with Endogenous Information:** Schottmuller and Lagerlof (2016)
- **Competitive Insurance with Endogenous Testing:** Doherty and Thisle (1996), Hoy and Polborn (2000), Strohmenger and Wambach (2000)
- **Models with Evidence and Learning:** Eso and Wallace (2016), DeMarzo, Kremer and Skrzypczak (2017)