Game theory IV

Stefano Caria

Michaelmas Term 2025

Game theory

- Discrete games
- Continuous games
- Hidden action
- Hidden information

Roadmap

Introduction

Perfect Bayesian Equilibrium

Welfare

Productive education

Signalling

- In transactions, it is often the case that one party knows more than the other
- Signalling: idea that an agent may want to persuade a principal of something that only the agent observes
  - E.g.: A worker is looking for a job; their ability is unobserved to a firm; the worker uses education to signal ability
- Today: a model where a `Worker` (agent) tries to signal her quality to a single `Firm`/`Employer` (principal)
- Assume `Firm` is drawn from a large pool of potential firms $\Rightarrow$ (i) `Firm` optimizes and, in doing so, (ii) earns zero profit.
- Who are the players? `Nature`, `Worker`, and `Firm`

Actions

1. `Nature` draws $n \in \{H, L\}$, such that $\Pr(n = H) = q$. The type is revealed to `Worker`, but not to `Firm`.

   $\Pr(n = H) = q$ is the prior belief
2. `Worker` chooses $e \geq 0$. `Firm` observes this.
3. `Firm` offers w to `Worker`.

Payoffs: `Worker`

$$U(w, e, n) = w - C(n, e); \tag{1}$$

$$\frac{\partial C(n, e)}{\partial e} > 0; \tag{2}$$

$$\frac{\partial^2 C(n, e)}{\partial e^2} > 0; \tag{3}$$

$$C(H, e) < C(L, e); \tag{4}$$

$$\frac{\partial C(H, e)}{\partial e} < \frac{\partial C(L, e)}{\partial e}. \tag{5}$$

(5) is sometimes called the ‘single-crossing condition’.

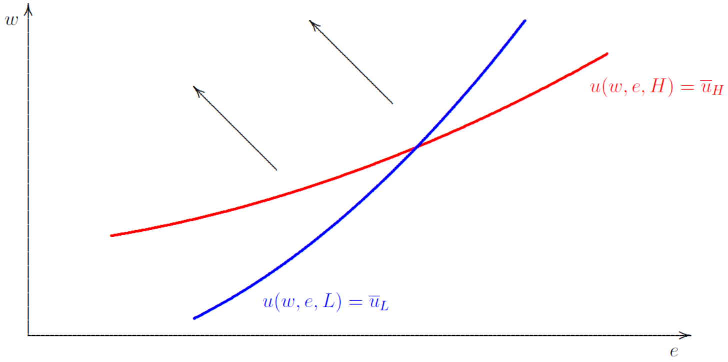
w
$u(w, e, H) = \overline{u}_H$
$u(w, e, L) = \overline{u}_L$
e

Payoffs: Firm

$$y(H, e) > y(L, e); \tag{6}$$

$$\frac{\partial y(n, e)}{\partial e} \geq 0. \tag{7}$$

Firm profit under unproductive education

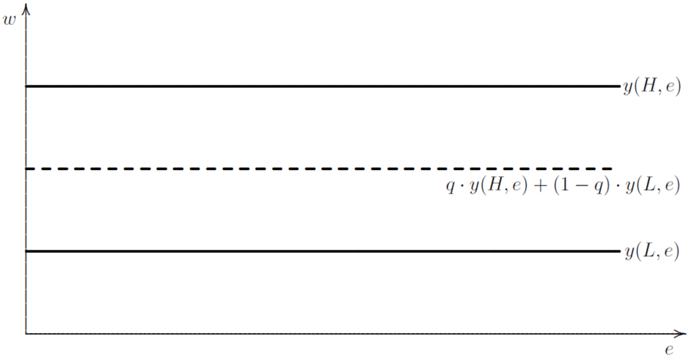

Benchmark case 1: education is illegal

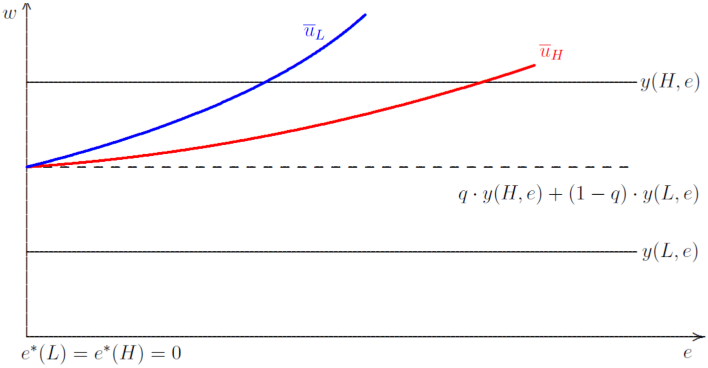

Benchmark case 2: full information

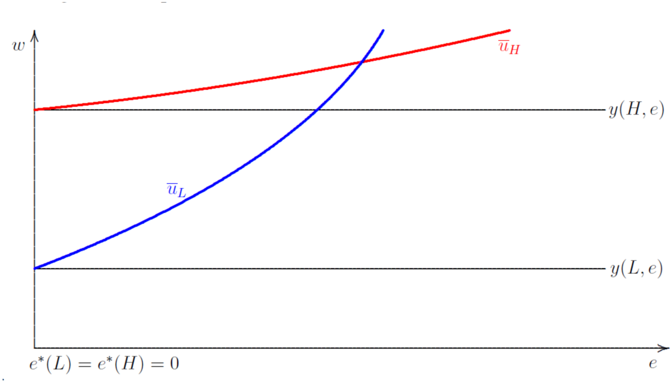

Roadmap

Solution concept: Perfect Bayesian Equilibrium

A PBE is a combination of

(i) a strategy for `Worker`, $e^*(n)$,

(ii) a strategy for `Firm` $w^*(e)$, and

(iii) a (posterior) belief for `Firm` $\mu(e) \in [0, 1]$,

such that

1. `Firm` acts as if $\Pr(n = H \mid e) = \mu(e)$.
2. For each $e \geq 0$, $w^*(e)$ maximises the payoff of `Firm`, given its belief $\mu(e)$.
3. For each $n \in \{H, L\}$, $e^*(n)$ maximises the utility of `Worker`, given `Firm`'s strategy $w^*(e)$.
4. For each $e \geq 0$ and each $n \in \{H, L\}$, if $\Pr(e^*(n) = e) > 0$, then $\mu(e)$ must be formed using Bayes' Rule and the strategy $e^*(n)$.

Solution concept: Perfect Bayesian Equilibrium

- Now the equilibrium specifies both strategies and beliefs;
- Condition (4) is key: it stipulates that beliefs have to be rational for *equilibrium actions*.

An initial implication

Given that we assume the firm makes zero profits due to the threat of competition, the optimal strategy of the firm is

$$w^*(e) = \mu(e)y(H, e) + (1 - \mu(e))y(L, e) \tag{8}$$

Types of equilibria

We consider two types of equilibria:

- **Separating equilibrium**: each type of `Worker` chooses a different action, so that upon observing `Worker`'s action, `Firm` knows `Worker` type.
- **Pooling equilibrium**: all types of `Worker` choose the same action, so that the `Worker`'s action gives `Firm` no clue about `Worker`'s type

A separating equilibrium

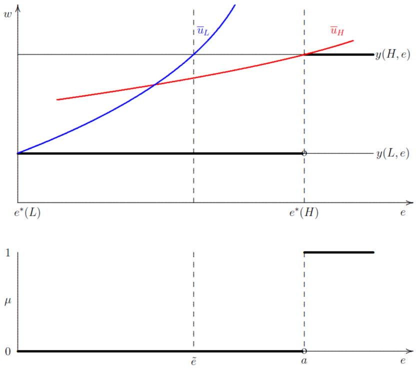

What are the strategies and beliefs in this PBE?

$$w^*(e) = \mu(e)y(H, e) + (1 - \mu(e))y(L, e) \tag{9}$$

$$e^*(L) = 0 \quad e^*(L) = a \tag{10}$$

$$\mu(e) = \begin{cases} 1 & \text{if } e \geq a, \\ 0 & \text{if } e < a. \end{cases} \tag{11}$$

To verify this is a PBE: (i) check whether strategies are best responses, (ii) check whether beliefs are correct *given the strategies used in equilibrium*.

Another separating equilibrium

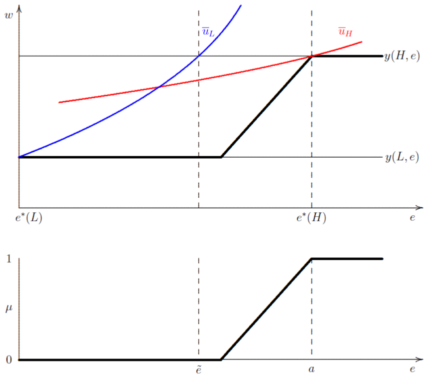

A pooling equilibrium

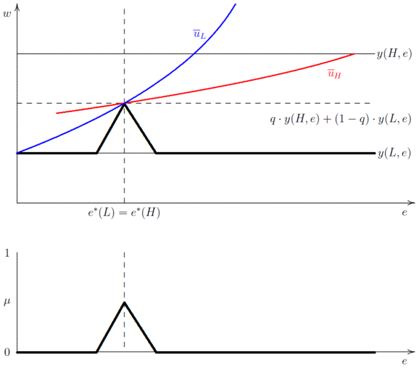

Recall the first separating equilibrium

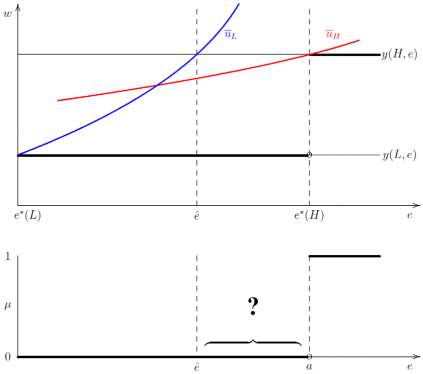

Shouldn't we put more restrictions on $\mu(e)$ for e between $\tilde{e}$ and a? These are dominated levels of e for L!

The Cho-Kreps Intuitive Criterion

Suppose that `Firm` observes a deviation from equilibrium, to education level *e*. Suppose that an action could never result in a higher payoff than the equilibrium action for type *L*, but could result in a higher payoff for type *H*, for some beliefs of `Firm`. Then:

$$\mu = 1. \tag{14.10}$$

(Note that we can flip 'H' and 'L' in this definition, in which case we have $\mu = 0$.)

One separating equilibrium remains

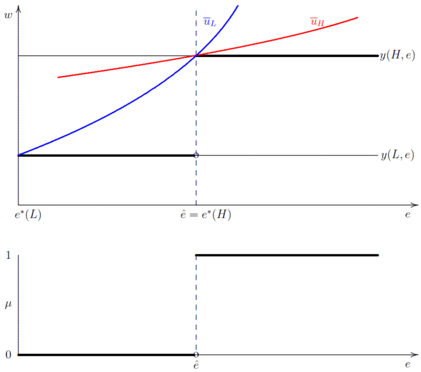

Can there still be a pooling equilibrium? (No, but why?)

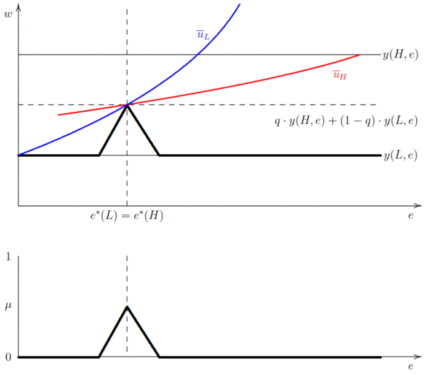

Roadmap

Is this separating equilibrium welfare enhancing?

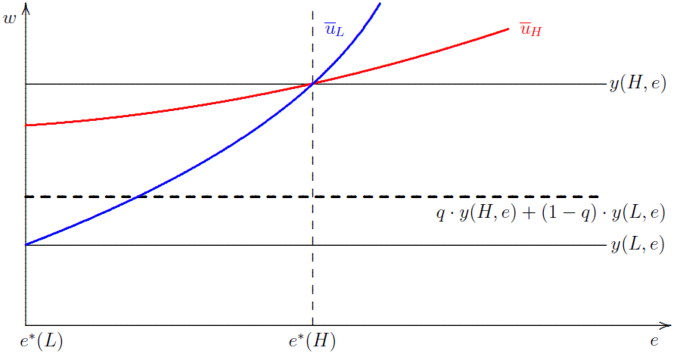

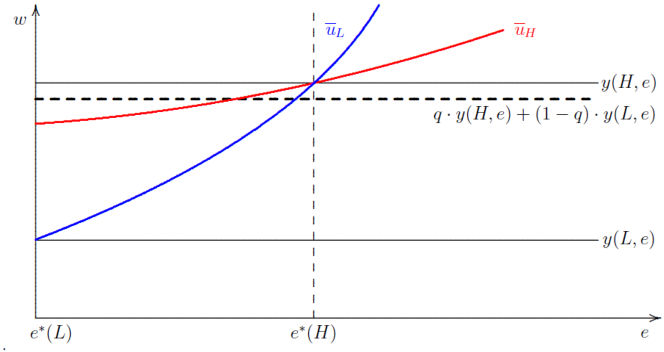

w
$\overline{u}_L$
$\overline{u}_H$
$y(H,e)$
$q \cdot y(H,e) + (1-q) \cdot y(L,e)$
$y(L,e)$
$e^*(L)$
$e^*(H)$
e

Going back to Abebe et al. 2021: Is this figure at all surprising?

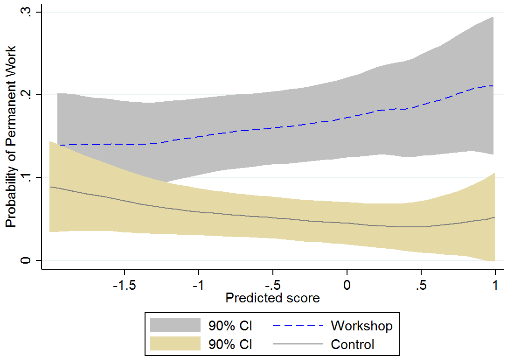

Roadmap

Productive education: no-envy cases are possible and should be ruled out for a signalling game to emerge

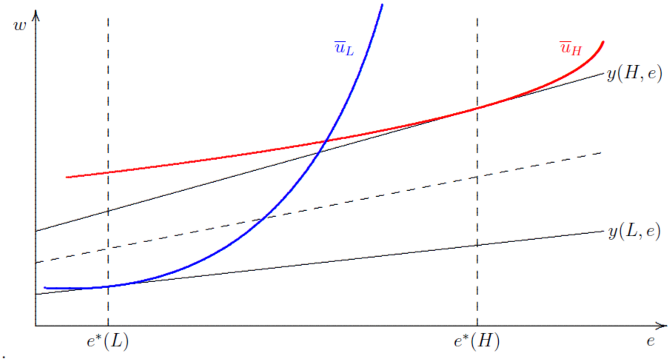

Productive education: *L* envies *H*

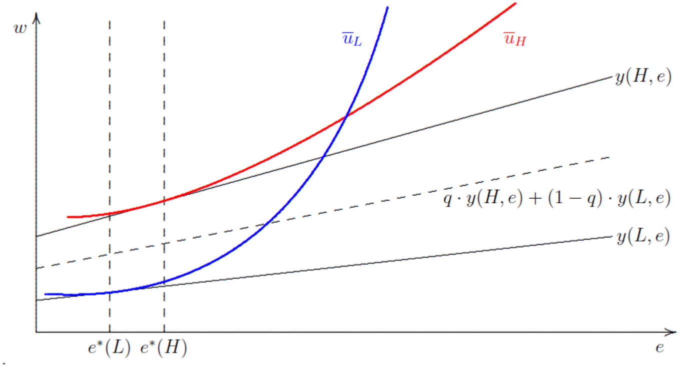

Productive education: separating equilibrium

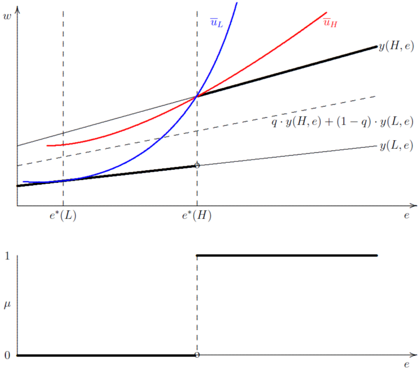

Takeaways

1. A credible signal is a costly action that separates types.
2. Education can be a signal. (But this is a ‘second best’ technology. . .)
3. Asymmetric information can be a cause of ‘missing market’ problems.
4. The private return to education need not be the public return. (Consider the econometric implications. . .)
5. Education can be a metaphor for many other forms of signalling.

Many other applications

For example:

- Leadership
- Protest and repression
- Criminal codes (e.g., tattoos.. see Diego Gambetta's Code of the Underworld)
- Governments and bond markets

References

I followed Simon Quinn's presentation of this material. You can check Simon's extensive lecture notes here.