lecture notes si680

Hidden Characteristics I: Base Model

Updated 28-March-08

The problem

- Risk-neutral principal contracts with agent (risk-averse, or neutral) to carry out task
- Two types of agents: Good and Bad
  - Good has disutility of effort $v(e)$, and with the usual convexity assumption: $v'(e) > 0$, $v''(e) > 0$.
  - Bad has higher disutility of effort (will work less): $kv(e)$, with $k > 1$
  - Agent utility:

$$U^G(t, e) = u(t) - v(e)$$
$$U^B(t, e) = u(t) - kv(e)$$

  - Since agent's *effort* now verifiable and principal risk-neutral, the effect of effort on different probabilities does not need to be explicitly modeled (outcomes not *informative* about effort because we already know effort), so let expected profit as a function of effort be $\Pi(e) = \sum_{i=1}^{n} p_i(e)x_i$. As usual, we'll assume decreasing marginal profit returns to the effort input: $\Pi'(e) > 0, \Pi''(e) < 0$.
  - Game timing shown in figure:

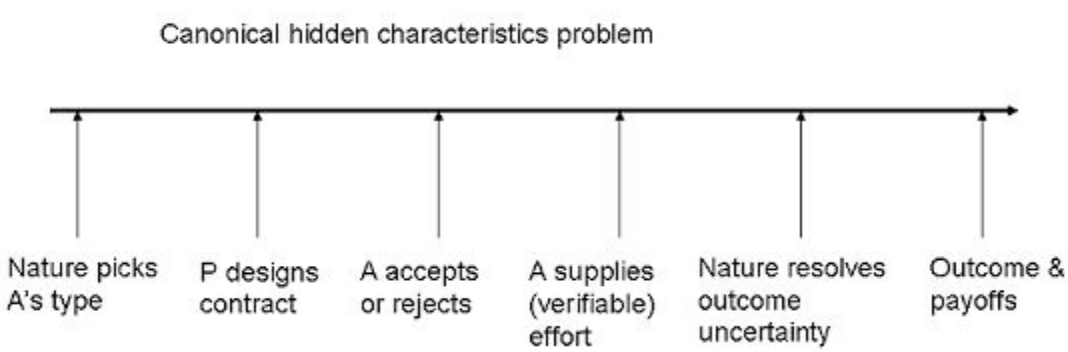

Symmetric information contract

If principle contracting with G agent:

$$\max_{e,t} \quad \Pi(e) - t$$
$$s.t. \quad u(t) - v(e) \geq U_0 \qquad \text{(PC)}$$

Lagrangean:

$$\mathcal{L} = \Pi(e) - t + \lambda(u(t) - v(e) - U_0)$$

First-order conditions:

$$\frac{\partial \mathcal{L}}{\partial e} : \Pi'(e^G) - \lambda v'(e) = 0 \qquad \text{(FOC-e)}$$
$$\frac{\partial \mathcal{L}}{\partial t} : -1 + \lambda u'(t) = 0 \qquad \text{(FOC-t)}$$
$$KT : \lambda(u(t) - v(e) - U_0) = 0$$

As is often the case for the participation constraint, it is simple to show that it must be binding (that is, an equality): Suppose not. Then the principal could subtract a constant from t^G to take away expected utility from the agent until the (PC) was binding, and the principal would be better off (smaller transfer).

If the (PC) is binding, then by the (KT) condition $\lambda \neq 0$ (the constraint is costly to the principal).

The best contract the principal can offer to each type of agent is the one that maximizes the principal's profits, subject to the agent being willing to sign (the PC constraint is satisfied with equality). This is isomorphic to the problem of a consumer maximizing utility subject to a budget constraint. An intuitive way to solve it is to graph the constraint, and then choose the highest utility indifference curve (we call it the iso-profit curve when utility is simply expected profit) possible give the constraint. That's what the picture below shows for the good type agent.

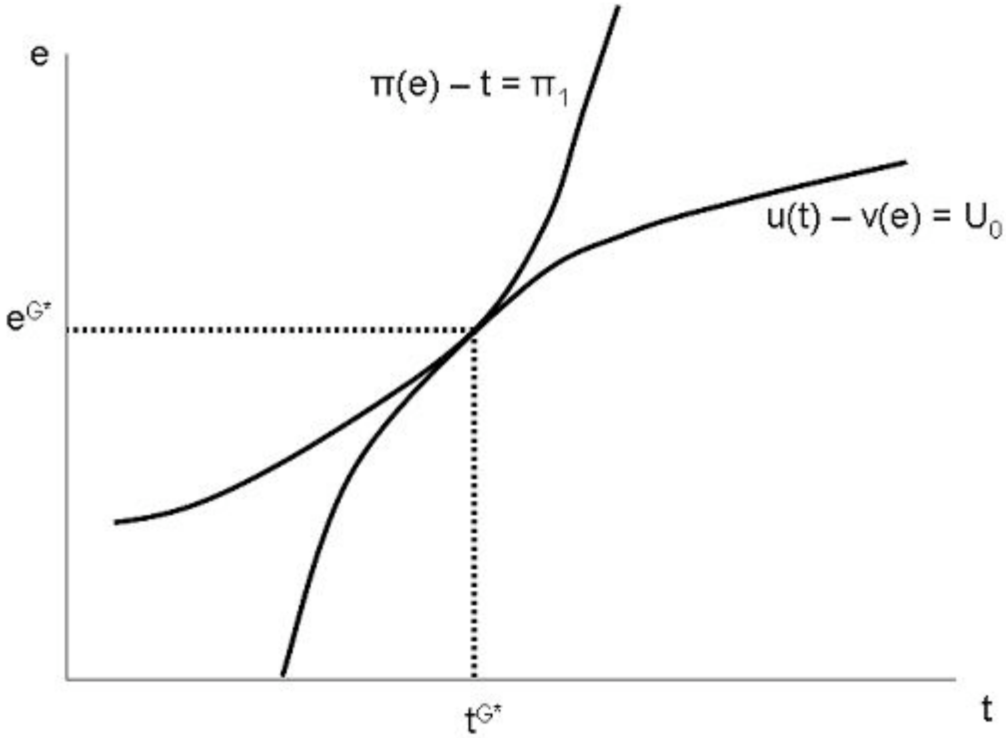

What happens with the bad type? (See figure, below.) The (PC) constraint changes to $u(t^{B*}) - kv(e^{B*}) = U_0$. What is the effect of $k > 1$? For any given t, to maintain the equality, $kv(e^{B*}) = v(e^{G*})$, so e must be lower for the bad agent, since disutility v is increasing in e. That means the "budget" constraint is *lower* for the firm contracting with a bad agent (not surprising: the opportunities available are not as good), and the result is that the firm can get only to a lower level of profit, $\Pi_2 < \Pi_1$.

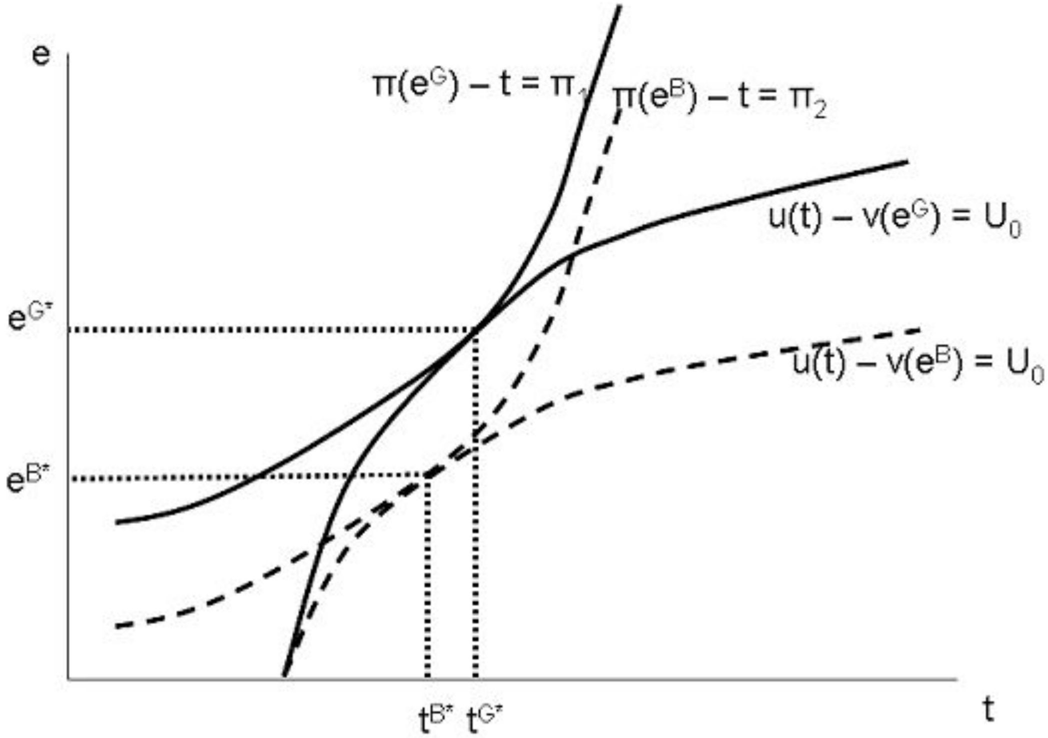

Notice that the Good agent makes a higher effort: this is a general result for the symmetric information case (the person for whom effort is less costly should make more effort). In the diagram, the Good agent also receives a higher wage — but that may not be true, because the two effects that determine the wage go in opposite directions. Indeed, Fig. 4.3 in the book is similar to my figure, except in their example, the wage is **lower** for the Good agent.

The harder way

The book first derives this the more rigorous way, calculating the first order equations and showing how to characterize (e, t) from the simultaneous (graphical) solution of those. The authors skip a couple of steps to get to the first order conditions, so I'll present it a bit more thoroughly here. However, this is optional material at the moment: for the symmetric information (no hidden information) case the graphical approach above is sufficient.

Substitute (*FOC-t*) into (*FOC-e*) and we have the two conditions that determine optimal e and optimal t:

$$u(t^{G*}) - v(e^{G*}) = U_0$$

$$\Pi'(e^{G*}) = \frac{v'(e^{G*})}{u'(t^{G*})}$$

where the first is the binding (PC) and the second is the usual efficiency condition for a symmetric information problem: the marginal rates of substitution of transfer and effort be the same for principal and agent (the marginal utility ratios).

The same steps lead us pretty quickly to the result for a contract offered to a known bad type agent:

$$u(t^{B*}) - kv(e^{B*}) = U_0$$
$$\Pi'(e^{B*}) = \frac{kv'(e^{B*})}{u'(t^{B*})}$$

Now you can graph these two sets of simultaneous equations and find their intersections to get the optimal (e, t) for a good-agent contract and a bad-agent contract. The footnote below Fig. 4.2 in the book explains why the curves take the positions they do.

Asymmetric information: Hidden characteristics

Now suppose the principal cannot observe the agent's type direcly. This adds another constraint to the problem: If want to offer different contracts to different types (or not offer a contract to certain types), how do you know which is which? Since only the agent knows her type, you have to let her *self-select*. This imposes, as in the hidden action problem, an *incentive compatibility* constraint.

To see problem, suppose P offered these same two contracts and let agents self-select.

- Bad agent
  - Gets $u(t^B) - kv(e^B) = U_0$ if she chooses "bad" contract
  - Gets $u(t^G) - kv(e^G) < u(t^G) - v(e^G) = U_0$ if she chooses "good" contract (she has to make the verifiable effort of a Good agent)
  - So, chooses the "bad" contract
- Good agent
  - Gets $u(t^G) - v(e^G) = U_0$ if she chooses the "good" contract
  - Gets $u(t^B) - v(e^B) > u(t^B) - kv(e^B) = U_0$ if she chooses the "bad" contract
  - So, the Good agent pretends to be bad, does better by shirking

So, the efficient symmetric information contracts don't work with hidden characteristics.

What to do?

Firm has two standard options:

1. Offer a single contract to all agents, designed based on the expected shares of types who will sign up (so on average the contract profits)
2. Offer a self-selection *menu* of contracts, designed so that each type wants to correctly choose

the contract designed for its type

To make #2 work there is an *incentive-compatibility* or *self-selection* constraint on the design of the contracts.

How to design?

There is a result that greatly simplifies the design problem:

> The *Revelation Principle* states, approximately, that the results of any complex contract can be replicated by a *truthful direct revelation* contract in which the the agents are simply asked to report their types, are assigned to a contract according to their response, and the contracts are designed so that it's in the agents' best interest to tell the truth.

What does this mean for design?
Pick a menu of contracts that

1. Are contingent on the agent's announced type
2. Satisfy the (PC) for each agent if she takes the contract designed for her type
3. Satisify the (IC) that each agent is better off announcing her type truthfully than lying
4. Maximizes Principal's profits

How does this compare to what we learned about hidden action?

- Very similar
- The "action" now is agent's announcement of her hidden characteristic(s)
- As long as contracts provide incentives for her to announce truthfully, and enough expected utility for her to participate, the contract terms can depend on her announcement ("action")

$$\max_{\{(e^G,t^G),(e^B,t^B)\}} \quad q[\Pi(e^G)-t^G]+(1-q)[\Pi(e^B)-t^B]$$

$$\begin{aligned}
\text{s.t.}\quad & u(t^G)-v(e^G)\geq U_0 \qquad \text{(PC-G)}\\
& u(t^B)-kv(e^B)\geq U_0 \qquad \text{(PC-B)}\\
& u(t^G)-v(e^G)\geq u(t^B)-v(e^B) \qquad \text{(IC-G)}\\
& u(t^B)-kv(e^B)\geq u(t^G)-kv(e^G) \qquad \text{(IC-B)}
\end{aligned}$$

Results

Result 1: The only binding participation constraint is for the *least efficient agent*.

Proof: Eliminate (*PC-G*): it is implied by (*PC-B*) and (*IC-G*). Since (PC) isn't binding for the good types, that means that more efficient agents each earn more than U_0 (*PC-G* not binding).

We call this the *information rent*. The intuition is pretty straightforward: have to make the Good agents better off than Bad, or the Good will pretend to be Bad and not work as much. By going through the math of the constrained optimization, the authors show that the amount of the information rent is $(k-1)v(e^B)$, that is,

$$u(t^G) - v(e^G) = U_0 + (k-1)v(e^B)$$

.
Intuitively: The greater is the disutility of effort for the Bad agent ($kv(e^B)$), the more he must be paid to meet the participation constraint. The more the bad type is paid, the more attractive shirking is to the Good type, and thus the higher the information rent the Good type must receive to honestly reveal her type. (The multiplier on v is $k-1$ because the term is more specifically equal to the difference between the disutility of the Bad type (which must be paid for Bad type participation) and the disutility of the Good type (thus, $kv(e^B) - v(e^B)$) which is the "net pay" the Good type receives from shirking.)

Result 2: Greater effort demanded of more efficient (Good) agent.

Proof: Implied by (*IC-G*) and (*IC-B*): Rearrange (*IC-G*) to get

$$v(e^G) - v(e^B) \leq u(t^G) - u(t^B)$$

and (*IC-B*) gives

$$u(t^G) - u(t^B) \leq k[v(e^G) - v(e^B)]$$

Combining these: $v(e^G) - v(e^B) \leq k[v(e^G) - v(e^B)]$, which, since $k > 1$ means that $v(e^G) \geq v(e^B)$ which in turn implies $e^G \geq e^B$, since disutility is an increasing function. By going through the math of the constrained optimization, the authors show that the inequality is in fact strict: two different contracts are offered, with $e^G > e^B$.

Result 3: The incentive constraint binds for the Good agent, not for the Bad agent

The proof requires going through the math of the constrained optimization. The intuition: Since the optimal contract menu demands that the Good agent work harder, costly incentives will be necessary to convince the Good agent not to pretend to be a Bad agent (so she can shirk).

Result 4: The contract is efficient for the Good agent

The proof is in the book. The result is that the ratio of marginal utilities for the Principal and Agent are equal for Good agents, so there is no inefficiency (waste). This is the same result we got for *all agents* when information was symmetric.

Why is the contract for Bad agents inefficient? To make it less attractive for Good agents; doing so means the Principal can pay less information rent (less reason for Good agents to pretend to be bad).

Comparing the asymmetric to symmetric info contracts

- Recall that if the symmetric info contracts $\{(e^{G*}, t^{G*}), (e^{B*}, t^{B*})\}$ are offered when types are hidden, the Good agent pretends to be a Bad agent, shirks, and gets higher utility than U_0. So, there is really only one contract, (e^{B*}, t^{B*}).
- How to change this contract to offer one that is more attractive to the Good agent (only)?
    - The only difference between the types is that the Bad type suffers more disutility from effort. Thus, to get more effort, the Principal can pay a wage increase large enough to reward the Good type for revelation, but not large enough to tempt the Bad type to lie and claim to be Good.
    - So, there is some contract $\{(e^{G}, t^{G}), (e^{B*}, t^{B*})\}$ with $e^G > e^{B*}$ and $t^G > t^{B*}$ that the Good agent will self-select.
    - So, the Principal can do better by demanding higher effort from the Good agent, raising the transfer just enough to give the Good agent the same utility he would get if he shirked — which then would *not* be enough to tempt the Bad type to pretend to be Good (since the Bad type needs a bigger wage increase than does the Good type)
- But, the Principal can do better still
    - The Bad contract is the Good agent's "temptation" to lie
    - By making the Bad contract a bit worse, the Good agent will be less tempted, so the Principal can lower the Good wage a bit
    - At least a bit of this trade-off is always worthwhile: the Principal saves more on the Good agent's wage than it loses from distorting the Bad agent's effort
- How much distortion?
    - Depends on the fraction of Goods in the population
        - With high q (lots of Goods), the more the Bad contract should be degraded, to

economize on rents paid to the many Goods

- If $1-q$ high (mostly Bads), degrade Bad contract less. If $1-q=1$ (*all Bads*) then Bad contract is not degraded at all: it is the efficient contract (same as in symmetric model, because type is no longer hidden)