Econ 200C - Second Half

Joel Sobel

May 29, 2010

Incomplete Information

The applications all involve models of incomplete information.

We can study these games using the techniques developed using a "trick" of Harsanyi.

Assume that there is an underlying "type" space that summarizes all of the uncertainty.

Player i's type determines his payoff function and possibly other characteristics.

Assume now that the game is one of imperfect information. It begins with a move of nature – selecting players' types. Nature then tells each player his or her type (but not the types of others). Solve using "standard" techniques.

Formality

Bayesian Game: $\{I, \{S_i\}, \{u_i\}, \{T_i\}, F\}$, $S = \Pi_{i=1}^{I} S_i$, $T = \Pi_{i=1}^{I} T_i$

- I – player set.
- S_i – strategy set of player i.
- $u_i : S \times T \to \mathbb{R}$ – payoffs
- T type space (summarizes private information).
- F distribution on T.

Equilibrium is a profile of $s_i^* : T_i \to S_i$ such that for each i and $t_i \in T_i$

$$s_i^*(t_i) \text{ solves } \max_{s_i(t_i) \in S_i} Eu_i(s_i, s_{-i}^*(t_{-i}), t) dF(t \mid t_i).$$

Note: You need a specification of the strategy for all types (even though only one type is realized). It does make sense for u_i to depend on the entire vector t.

Adverse Selection

Standard Example: You are trying to buy a used car. The owner of the car knows more about the car's quality than you do. If there is a "market price" then only "bad" cars will be offered for sale.

Competitive Market

1. Given set of types, $t \in T$ ($T = [0, 1]$ will be standard).
2. $F(\cdot)$ prior on types.
3. $r(t)$ reservation wage; non-decreasing.
4. t market value of type t's product.
5. w wage.

Market outcome: $\{t : w \geq r(t)\}$ transact at w.
Demand: Infinite, any, zero depending on whether expected productivity is greater than, equal to, or less than wage.

Interpretation

t is marginal product of a worker (the value of worker to employer)
$r(t)$ is reservation wage (what worker can get if not employed)
w is market wage of employed worker.
One imagines that a market equilibrium involves a price with the property that (a) wage is equal to average quality of employed worker; (b) workers are employed if and only if wage exceeds reservation wage; (c) markets clear.

Competitive Equilibrium

A competitive equilibrium is a wage w^* and a set of types T^* such that:

1. $T^* = \{t : r(t) \leq w^*\}$
2. $w^* = E[t \mid t \in T^*]$

T^* is the set of active workers. The first condition says that workers behave rationally: They are active if and only if their reservation wage is less than the market wage w^*.
The second condition says that firms earn zero profits (so it is a market-clearing condition).
Notice that the second condition does not make sense if T^* is empty.
If T^* is empty, then (2) places no constraints on w^*.
(So there will always be an "inactive" equilibrium.)

Market Failure

There may be no competitive equilibrium with active traders even when there are always gains from trade. This is called adverse selection.

Suppose that $r(t) = \alpha t$ for $\alpha \in (0, 1)$. In this case, every worker is more valuable in the market than outside. Suppose that t is uniformly distributed on $[0, 1]$.

What happens if the wage is w?

All workers with $\alpha t \leq w$ enter the market. So the quality of a worker is uniformly distributed on $[0, w/\alpha]$ and the average quality is $w/2\alpha$.

An interior equilibrium requires that $w^* = w^*/2\alpha$ or $\alpha = .5$. If $\alpha > .5$ the market cannot be active since the average quality of worker is less than the wage.

(When $\alpha < .5$ all workers are active.)

Strategic Model of Adverse Selection

The competitive model does not allow firms to compete – they take wages as given.
Does this matter?

- Nature picks types according to distribution.
- Workers learn type, but firms do not.
- Firms (assume that there are two) simultaneously make wage offers.
- Workers select a firm to work for (or opt out of the market).
- Preferences as before.

Analysis

Informally: The equilibrium outcome is the highest competitive market wage.
More formally: Let $\bar{w}$ be the highest competitive market wage and let $\underline{t}$ be the lowest type.
Either: $\bar{w} = r(\underline{t})$, in which case the market shuts down (wages are not determined in equilibrium, but they must be no higher than $\bar{w}$)
or: There is a unique subgame perfect equilibrium in which firms offer $\bar{w}$ provided a regularity condition holds. A sufficient condition for the result is that $E[t \mid r(t) \leq w] > w$ for w slightly less than $\bar{w}$.
The regularity condition says that you cannot have an equilibrium with wages slightly below $\bar{w}$.
The condition will be satisfied if $E[t \mid r(t) \leq w] - w$ is strictly decreasing at $\bar{w}$. (We know that it is weakly decreasing.)

Interpretation

Adding strategic power rules out low activity (low wage) equilibria, but does not eliminate adverse selection.

Argument

If $\bar{w} = r(\underline{t})$, then $E[t \mid r(t) \leq w] < w$ for all w, so no firm will offer more than $\bar{w}$ (else it will make negative profits).
If $\bar{w} > r(\underline{t})$,

1. Zero profits in equilibrium. (Usual argument: if positive profits, the less well off firm can raise wage by a little bit and capture essentially all profits.)
2. Highest wage must be $\bar{w}$. (Else deviate to just below $\bar{w}$. This is profitable by the regularity condition.
3. All must charge $\bar{w}$. (Otherwise high price firm can lower wage.)
4. Indicated strategies are an equilibrium. (Higher wages can't be profitable by the definition of w^*.)

Inefficiency and Intervention

We know that the equilibria in adverse selection models are inefficient: If there was complete information, it would be possible to make everyone better off.

When there is incomplete information, this is the wrong notion of efficiency (because it does not respect constraints imposed by private information).

Think of an allocation generally as $(p(t), w(t))$ as describing the probability that a type t agent is employed and the agent's wage. This allocation gives worker t utility $U(t) = (1 - p(t))r(t) + w(t)$. This allocation is incentive compatible (respects private information) if $U(t) \geq (1 - p(t'))r(t) + w(t')$, that is if it is not in t's interest to pretend to be t'.

This generality is sometimes important, but here let us restrict attention to $p(t) = 1$ or 0. In this case workers are either employed on unemployed and $w(t)$ must be constant, equal to say w_e for all employed types and constant, equal to say w_u for the others. It is consistent with the incentives to pay unemployment compensation, but not to distinguish between different employed types.

Best Competitive Outcome is Constrained Efficient

Assume that $r(\cdot)$ is strictly increasing and $r(t) \leq t$ for all t. If the density of t is positive over its support, then the highest wage competitive equilibrium is constrained efficient.

1. Some workers unemployed (else fully efficient because $r(t) \leq t$).
2. Interior type $\tilde{t}$ indifferent between work and not. Below employed; above not. (By IC, monotonicity of $r(\cdot)$.
3. $w_u + r(\tilde{t}) = w_e$ (indifference condition).
4. Budget balance: $w_e F(\tilde{t}) + w_u(1 - F(\tilde{t}) = \int_{\underline{t}}^{\tilde{t}} \theta dF(\theta)$
5. $w_u(\tilde{t}) = \int_{\underline{t}}^{\tilde{t}} \theta dF(\theta) - r(\tilde{t})F(\tilde{t})$
6. $w_e(\tilde{t}) = \int_{\underline{t}}^{\tilde{t}} \theta dF(\theta) + r(\tilde{t})(1 - F(\tilde{t}))$

Aside: so one can obtain the highest (or any) competitive equilibrium without unemployment payments.

Continued

It is not possible to improve things for everyone by picking a cutoff different from the highest competitive one. Let t^* denote the competitive cut off.

If $t^* > \tilde{t}$, then $w_e(\tilde{t}) < r(t^*)$ so that employed agents are worse off (relative to the competitive equilibrium)

If $t^* < \tilde{t}$, then $w_u(\tilde{t}) < 0$, so unemployed workers are made worse off (relative to the competitive equilibrium).

Both claims require a bit of algebra to justify.

Signaling

1. Two players, S (sender) and R (receiver)
2. Nature picks t type of sender. $p(t)$ is probability that type is t.
3. Sender observes t, selects signal, s.
4. Receiver observes s (but not t), selects action a.
5. $U_i(a, t, s)$ payoff function.

Standard application: S is worker, t is ability, s is education, R is market wage. Possible preferences: $U_R(a, t, s) = -(a - t)^2$ (market pays wage equal to expected ability)
$U_S(a, t, s) = a - \alpha s^2/t$ (workers like higher wages and lower signals, marginal cost of producing signal decreases with ability).

Basic Question

Is it possible for the signal to convey information to the receiver?
Answer: Maybe not.
Suppose that $s(t) \equiv s^*$.
The best response for the Receiver includes
$a(s^*) = arg \max EU_R(a, t, s^*)$ (prior optimal action).
If one can find $a(s)$ for $s \neq s^*$ such that
$U_S(a^*(s^*), t, s^*) \geq U_S(a(s), t, s)$ for all t and $s \neq s^*$, then there is a pooling equilibrium outcome.
Finding such an $a(\cdot)$ is not hard in leading examples (in the labor market, let $a(s) \equiv 0$, so that in a putative pooling equilibrium, agents get the average wage for signal s^* and zero otherwise.

Definition of Equilibrium

1. Sender strategy: $\sigma(t)$ mapping type to signal.
2. Receiver strategy: $\alpha(s)$ mapping signal to action.
3. Receiver belief: $\mu(t \mid s)$ updating beliefs given signal.

$(\sigma^*, \alpha^*, \mu^*)$ is a (weak perfect Bayesian) equilibrium if:

1. $\sigma^*(t)$ solves $\max_s U^S(\alpha^*(s), t, s)$ all t.
2. $\alpha^*(s)$ solves $\max_a EU^R(a, t, s) d\mu(t \mid s)$ all s.
3. μ^* derives from prior and σ^* using Bayes's Rule (whenever possible).

Single-Crossing Condition

If $t' > t$, and $s' > s$, then $U_S(a', t, s') = U_S(a, t, s)$ implies that $U_S(a', t', s') > U_S(a, t', s)$.

This is the fundamental sorting condition that arises in many applications of information economics.

Geometrically it states that indifference curves in (a, s) space of different types cross at most once.

Mathematically it can be thought of as a supermodularity assumption on utility.

Economically, it says that higher types are more willing to use higher signals.

When Does Single-Crossing Fail?

One Example: Cheap Talk ($U_S(\cdot)$ independent of s).

Separating Equilibrium

Let $BR(s, t)$ be the Receiver's best response to the signal s given type t:
$BR(s, t)$ solves: $\max_a U_R(a, t, s)$.

1. Lowest type as in complete information: s_0 solves $\max_s U_S(BR(s, 0), 0, s)$. Solution is $s^*(0)$.
2. Next type does as well as possible constrained by lower type does not want to imitate. That is, given $s^*(t)$, $s^*(t+1)$ is solution to $\max_s U_S(BR(s, t+1), t+1, s)$ subject to $U_S(BR(s^*(t), t), t, s^*(t)) \geq U_S(BR(s, t+1), t, s)$.

Details

1. Can one really solve the problems in the construction?
2. I described a path of play. How do you specify complete strategies?
3. Given the strategies, are they really equilibrium strategies?

The answer to (1) is, in general, no (even with single crossing). You might get a boundary solution. That is, you might have all "high" types sending the highest message. This can be ruled out with a boundary condition.

There are many answers to (2), but the most sensible is for R to assume that signals $s \in (s^*(t-1), s^*(t))$ come from $t-1$ and so $a^*(s) = BR(s, t-1)$. This "answer" raises two questions: why is $s^*(t-1) < s^*(t)$? why does this support the equilibrium? In both cases the answer is "single crossing."

The receiver best responds by definition. All we know about the sender is that the sender of type $t-1$ is indifferent between sending "her" message ($s^*(t-1)$) and the message of type t. With single-crossing, this is enough.

Properties of Separating Equilibria

- Multiplicity
- Pareto Ranked from the viewpoint of Sender.
- Good from the viewpoint of Receiver.
- Sensitive to prior.
- Excessive investment in signal

Refinements Select "Pareto dominating separating equilibrium"

The theory is out of fashion.
The theory invokes subtle dominance relationships.
The message is that one of the equilibria is "better" than the others.

Screening

Essentially the signaling problem with a different solution concept. Following a leading example, we call the informed players workers and uninformed players firms.

1. Unknown type t.
2. Action (wage) a.
3. Signal s (the interpretation of s varies with the application).
4. (Two) firms select contracts (a, s).
5. Workers pick among available contracts.

Simplifying assumptions:
Worker utility: $a - c(s, t)$
Firm utility: $t - a$.

Applications

1. Insurance: type is probability of accident; signal is the amount of risk worker accepts. Formally: t is probability of no accident.
   With accident, agent's wealth without insurance is 0; W with no accident.
   Insurance exchanges payment x for payoff I in bad state.
   Agent gets $tu(W - x) + (1 - t)u(I)$. Firm gets $tx - (1 - t)I$.
   Note: agent does not control the probability of being in accident (no moral hazard).
2. Guarantee: t is the probability that product does not fail.

Observable Types

- No signaling ($s^*(t) = 0$ all t).
- Zero profits.
- Efficiency

Zero profits come from competition. Plainly profits are non-negative on each type's contract. If they are strictly positive, one firm earns no more than half and can do better by raising its wage.

No signaling comes from competition too because a firm can do better by cutting both wages and requested signal.

Equilibrium

Assume (for simplicity) two types.

- Zero profits.
- No pooling.
- In separating contract, earn zero profits on each type (no cross subsidization).
- No distortion of low types.
- High type separates with minimum cost.
  $(\underline{t} - c(0, \underline{t}) = \bar{t} - c(\bar{s}, \underline{t})$
- Potential failure of existence.

Details - 1

Zero profit: Let the less profitable firm offer contracts that pay slightly higher wages than the competition. This increases the number of (profitable) workers at arbitrarily low cost.
No pooling: If one firm offers a (zero profit) pooling contract, the other firm can make a profitable deviation that attracts only good workers at a slightly higher wage with a higher signal required.
No cross subsidization: If lows are profitable, then a firm would earn positive profits by offering a contract designed for lows at a slightly higher wage. This will earn positive profits (even if it attracts highs). If highs are profitable, then a firm can design a contract that pays higher wage, requests higher signal, and attracts only highs (making positive profits on them).

Details- 2

No distortion: Otherwise a firm gains by lowering wage and required signal.
High Contract: Must separate. Must earn zero profit. Hence if not minimal separating contract, then firm can make positive profits by lowering both wage and signal in a way that attracts only high types.
Nonexistence: If there are a lot of high types, a pooling contract can be profitable.

Lessons

1. Inefficiency from incomplete information.
2. Constrained Efficiency.
3. Distortion of high types.
4. Signaling: many equilibria; Screening: no equilibria.

Agency

Central problem:
Principal has control of enterprise. Agency either has control of information or action.
The first case is called hidden (from the principal) information; the second is called hidden action.
We'll study hidden information using general language of mechanism design.

Hidden Action Model (Moral Hazard)

General Formulation:

1. π profit (observable to all).
2. $e \in E$ effort choice of agent.
3. $f(\pi \mid e)$ distribution of profits given effort.
4. w wage payment to agent (conditioned on π but generally not on e.
5. Preferences: $u(w, e)$ is agent's utility given wage w and effort e. Typically assume increasing and concave in w and decreasing and concave in e. For simplicity, $u(w, e) = v(w) - g(e)$. $\pi - w$ for principal (risk neutrality)
6. Agent has outside option that pays u_0. In order to accept job, expected utility must be at least u_0

Observable Effort

- This is the first best.
- In order to induce effort e, P sets w so that $E\{u(w(\pi), e) = u_0$ and $Ew(\pi)$ is minimized.
- Why? No need to pay strictly more than u_0. Want to find minimum expected payoff to generate effort.
- It is best to set $w(\pi)$ equal to a constant. (Mean-preserving shifts in wages lower A's utility without changes P's cost.)
- Hence to induce effort e set wage $w(\pi) \equiv v^{-1}(u_0 + g(e))$.
- Hence P solves max $\int \pi f(\pi \mid e) d\pi - v^{-1}(u_0 + g(e))$ (by picking e).

Risk Neutral Agent

Let $w = \pi - K$. This leads to first-best as agent selects the effort level that principal would select if enterprise were integrated.
Adjust K so that Agent's utility is u_0.
This is the "franchise" solution.

General Comments

- Hidden Action problem is about risk sharing.
- P doesn't care about risk.
- A does.
- To fully share risk, you'd like P to insure A fully, but in this case, A has no incentive to take effort.
- This is the fundamental tradeoff. Not a problem if A is risk neutral or P knows everything.

Goals

1. Properties of optimal contracts. Piece rates? Bonuses? Wages increasing in output?
2. Sensitivity results: Increasing risk bad? Increasing information good?

Alas, the solution to the general problem need not have any of these properties.
Simplify for some insight.

Two Effort Levels

Assume $E = \{\underline{e}, \bar{e}\}$ for $\underline{e} < \bar{e}$. Assume also that higher effort leads to stochastically higher output.

P decides first the cost of implementing each effort level, then figures out the better effort level.

The cheapest way to implement $\underline{e}$ is to fully insure the agent (there is no danger that the agent will shirk).

Implementing High Effort

$$\max \int \{\pi - w(\pi)\} f(\pi \mid \bar{e}) d\pi$$

subject to

$$\int v(w(\pi)) f(\pi \mid \bar{e}) d\pi - g(\bar{e}) \geq u_0$$

and

$$\int v(w(\pi)) f(\pi \mid \bar{e}) d\pi - g(\bar{e}) \geq \int v(w(\pi)) f(\pi \mid \underline{e}) d\pi - g(\underline{e}).$$

The first constraint is individual rationality; the second constraint in incentive compatibility.
The second constraint is new (the result of hidden actions).

First-order condition:

$$\frac{1}{v'(w(\pi))} = \gamma + \mu \left[1 - \frac{f(\pi \mid \underline{e})}{f(\pi \mid \bar{e})} \right].$$

The greek letters are non-negative multipliers.
Note:

- $\gamma > 0$ ((IR) binds):
  Otherwise wage is constant, which is inconsistent with the (IR) constraint.
- $\mu > 0$ ((IC) binds): otherwise lhs is sometimes negative, which is not possible.

What we learn

- Inefficiency
- Sensitivity of payment to likelihood ratio, $\frac{f(\pi|\underline{e})}{f(\pi|\bar{e})}$.

Optimal wage schedule is monotonic exactly when the likelihood ratio is monotone.
The monotone likelihood ratio property is therefore a key underlying assumption is agency theory.