Lecture Notes for Economics 200Cb: Games and Information Segment

Vincent Crawford **Spring 2009**

Revised 26 May 2009

1. Adverse Selection (MWG 436-450; Kreps 625-629; Varian 466-469)

Competitive labor market with many identical, expected profit-maximizing firms.

Many workers with privately observed ability θ (measured as output) distributed on a compact interval with c.d.f. $F(\theta)$ and reservation ("home") wage $r(\theta)$.

Full-information efficient benchmark outcome has each worker working iff $r(\theta) \leq \theta$, each paid his θ.

Inefficient equilibrium with $r(\theta) \equiv r$ (home wage independent of ability):

If $w \geq r$ all accept employment; if $w < r$ none do; either way equilibrium wage $w^* = E\theta$.

Whether $w^* >$ or $< r$ is determined by proportions of high- and low-ability workers:

If too many lows, firms unwilling to pay wage any will accept, too little employment (some workers with $\theta > r$ unemployed). If too many highs, firms pay wage that all will accept, too much employment (some workers with $\theta < r$ employed in equilibrium).

Asymmetric information prevents firms from paying each worker his marginal product, and thus prevents market from allocating workers efficiently between work and home.

When $r(\theta)$ varies with θ, *adverse selection* can cause market failure.

Figure 13.B.1-2 at MWG 441-442.

Suppose that $r(\theta) < \theta$ for all θ, so efficiency requires all workers to work; that $r(\theta)$ is strictly increasing; and that there is a density of abilities θ, privately known.

Then at any given wage only the less able ($r(\theta) \leq w$) will work, so that lower wage rates lower (via adverse selection) the average productivity of those who accept employment.

Equilibrium wage $w^* = E[\theta | r(\theta) \leq w^*]$, the average productivity of those who work.

To induce the best workers to work, w^* would have to equal $r(\theta)$ for the highest possible θ; but in Figure 13.B.1 firms can't break even at this level.

Thus the best workers don't work in equilibrium; the equilibrium is inefficient.

In cases like Figure 13.B.2 (left), adverse selection causes complete market failure: no one works, even though efficiency requires all to work.

Equilibrium can be unique as in Figure 13.B.1, or multiple and Pareto-ranked as in Figure 13.B.2 (right).

The equilibrium with highest wage is better for all workers and no worse for firms, who earn zero profits in any equilibrium (there’s a possibility of "coordination failure").

In a two-stage game where firms first simultaneously choose wages and workers then choose among firms, with a density of abilities θ, the highest-wage competitive equilibrium is the unique subgame-perfect equilibrium.

(Exception: If $w^* = r(\theta)$ for the lowest θ, there can be multiple subgame-perfect equilibria, but all pure-strategy subgame-perfect equilibria yield workers the same payoffs as the highest-wage competitive equilibrium.)

Proof (Proposition 13.B.1 at MWG 443-444): Firms can break lower-wage equilibria by raising wage and attracting higher-productivity workers.

We can use the notion of *(incentive-)constrained Pareto-efficient allocation* (*incentive-efficient allocation* for short) to think about the welfare effects of market intervention by a planner who faces the same informational limitations as agents in the market (their common knowledge, not their private information).

Does the market do as well as possible, given this limited information? Incentive-efficiency is the relevant criterion for mortal planners, not full-information efficiency.

A planner who can only observe whether a worker works, not his ability, must pay same wage to employed workers, and same (possibly different) wage to unemployed workers.

In this model compensation is only possible by adjusting the wages.

A planner can still implement the highest-wage competitive equilibrium by setting employed wage $w_e = w^*$ and unemployed wage $w_u = 0$ and respecting workers' wishes to work or not.

In this model, the highest-wage competitive equilibrium happens to be incentive-efficient.

Proof (Prop. 13.B.2, MWG 447-448):

Any given employed and unemployed wages w_e and w_u will make a range of low-end ability types accept employment (those for whom $w_u + r(\theta) \leq w_e$)—low-end because the pecuniary benefits of working are independent of θ, but r(θ) increases with θ.

Call the highest type that accepts employment for particular values of w_e and w_u θ^, and call the highest type that works in the highest-wage competitive equilibrium θ^*.

Given w_e and w_u, incentive-compatibility (IC) requires $w_u + r(\theta\hat{}) = w_e$ (< (>) for $\theta < (>) \theta\hat{}$).

Budget balance (BB) with ability distribution F requires $w_e F(\theta\hat{}) + w_u(1-F(\theta\hat{})) = E[\theta|\theta < \theta\hat{}]$.

Given w_e and w_u, incentive-compatibility (IC) requires $w_u + r(\theta\hat{}) = w_e$ (< (>) for θ < (>) $\theta\hat{}$).

Budget balance (BB) with ability distribution F requires $w_e F(\theta\hat{}) + w_u(1-F(\theta\hat{})) = E[\theta|\theta < \theta\hat{}]$.

Incentive-compatibility plus budget balance means that a planner who chooses w_e and w_u to make $\theta\hat{} = \theta^*$ enforces same outcome as the highest-wage competitive equilibrium.

(If he raised w_u he'd also have to raise w_e to preserve incentive-compatibility, but raising both is infeasible with $\theta\hat{} = \theta^*$.)

Can a planner Pareto-improve highest-wage competitive equilibrium by making $\theta\hat{} < \theta^*$?

Then $r(\theta\hat{}) < r(\theta^*)$, so IC implies that the gap between w_e and w_u must be smaller; and because fewer and on average less productive people work, total per capita output is smaller. Thus w_e must be smaller, so low-ability workers are worse off than in equilibrium.

Can a planner Pareto-improve highest-wage competitive equilibrium by making $\theta\hat{} > \theta^*$?

Then $r(\theta\hat{}) > r(\theta^*)$, so IC implies that the gap between w_e and w_u must be larger; and because more and on average more productive people work, total per capita output is larger, but not enough larger to cover the higher wages if $w_u \geq 0$.

Thus BB implies that w_u must be smaller, and high-ability workers are worse off than in the highest-wage competitive equilibrium.

(The detailed algebra is at MWG 447-448.)

2. Signaling and Screening (MWG 450-467; Kreps 629-652; Varian 469-471)

The inefficiency of competitive outcomes with asymmetric information leaves room for various tactics by agents to affect the outcome.

(Such tactics may or may not improve the outcome; agents don’t care about efficiency per se, but there is a weak tendency for individually beneficial tactics to enhance efficiency.)

Consider two commonly observed tactics:

Signaling (actions intended to distinguish their own types taken by informed agents)

Screening (actions to distinguish others’ types taken by uninformed agents)

“Signaling” and “screening” (usually) refer to agents' actions.

“Sorting”, “separating”, and “pooling” usually refer to kinds of equilibrium outcomes.

However, “screening” is sometimes used to describe a separating equilibrium.

Consider first signaling, then screening, in a labor market example, following MWG 450-467 (and Kreps 629-637 and 645-649). (The original screening analyses of Stiglitz (monopoly) and Rothschild-Stiglitz (competition) were set in insurance markets.)

Spence's signaling model: Two firms, one worker (can easily generalize to pool of workers).

Market structure as in Section 1, Adverse Selection: Competitive labor market with identical, expected profit-maximizing firms; here two is enough for competition.

However, the worker now has only two ability "types," with productivities $\theta_H > \theta_L > 0$, where $0 < \text{Prob}\{\theta = \theta_H\} = \lambda < 1$.

Only workers observe their types, but everyone knows λ, as common knowledge.

$r(\theta_H) = r(\theta_L) = 0$, so the unique equilibrium if workers couldn't signal would have all workers employed at wage $w^* = E\theta$, the full-information efficient outcome.

However, workers can now choose education level e, continuously variable within a bounded interval, with differentiable and uniformly higher marginal ("single crossing property") and total costs for θ_L.

Figures 13.C.2-3 at MWG 453.

Education has no effect on productivity (can easily relax; Kreps assumes education is productive, but the difference is inessential for this purpose; I follow MWG).

Although firms cannot directly observe ability, they can observe education levels; because of the different costs of education for high- and low-ability workers, chosen levels might indirectly signal workers' abilities to firms.

Modern treatments depart from Spence's by modeling the market explicitly as a game. The "rules" are as follows (extensive form in Figure 13.C.1 at MWG 451):

(i) nature chooses the worker's type θ

(ii) worker observes type and chooses education level e

(iii) firms observe e and simultaneously make wage offers w_i

(iv) worker observes w_i and chooses between firms.

Assume firms and worker play a *weak perfect Bayesian equilibrium*, with one added condition:

For all e (not just those chosen in equilibrium), firms use a common posterior μ(e) to update their beliefs about the worker's ability and to predict each other's equilibrium w_i from the equilibrium offer functions.

This consistency of beliefs and strategies off the equilibrium path yields *perfect Bayesian equilibrium* (*PBE*), here equivalent to *sequential equilibrium*.

A set of strategies and beliefs μ(e) is a PBE iff:

(i) the worker's strategy is optimal given the firms' strategies

(ii) μ(e) is derived from the worker's strategy using Bayes' Rule whenever possible

(iii) the firms' wage offers following each possible e are in Nash equilibrium in the simultaneous-move wage offer game when the probability that the worker is of high ability is μ(e)

Characterization of PBE:

As in Bertrand duopoly, if the firms observe e and have beliefs $\mu(e)$, their unique equilibrium offers are identical and equal to the worker's expected marginal product given their beliefs: $\mu(e)\theta_H + (1- \mu(e))\theta_L$.

The worker then picks either firm; it doesn't matter which.

A *separating* PBE is one in which different worker types choose different e's, so that in equilibrium, observing a worker's e perfectly reveals his type.

Lemma 13.C.1 at MWG 453: In any separating PBE, each worker type is paid its marginal product.

Proof: A straightforward modification of the standard characterization of Bertrand pricing.

Lemma 13.C.2 at MWG 454: In any separating PBE, a low-ability worker sets $e = 0$.

(Given that education is unproductive; otherwise e would maximize productivity net of cost.)

Proof: $e > 0$ costs more, but it can't help the worker because in the equilibrium he is separated, hence paid his marginal product.

Lemma: In any separating PBE, A high-ability worker chooses e~, the lowest e that a low-ability worker won't wish to imitate.

(Again, given that education is unproductive.)

Proof: Raising e costs more, so a high-ability worker, if he chooses to separate, will choose the lowest e that brings it about, given a low-ability worker's equilibrium choice.

Figures 13.C.5-7 at MWG 454-455.

Figures 13.C.5-6 show separating PBEs with different supporting wage functions w*(e), each derived from common beliefs (hence between dotted lines).

In each case firms and workers behave optimally on and off the equilibrium path: Firms bid correctly and consistently for each e, and each worker type has a generalized tangency between its indifference curve and (w,e) opportunity locus.

The PBE In Figure 13.C.6 has beliefs that satisfy a plausible *monotonicity* requirement, in that higher e never lowers the firms' estimate of the worker's expected productivity.

Figures 13.C.7-8 at MWG 455.

Figure 13.C.7 at MWG 455 shows a separating PBE in which a high-ability worker chooses e > e~, the minimum needed to separate from low-ability worker.

The e in this kind of equilibrium could be as high as the e_1 that makes a high-ability worker willing to imitate a low-ability worker.

The separating equilibria are Pareto-ranked: The one with the lowest e is best.

Firms get zero profits in all of them. Low-ability workers get the utility of obtaining zero education, being identified, and so being paid their marginal product.

High-ability workers are also identified and get paid their marginal product, and so do best when e maximizes productivity net of cost, in this case lowering it as much as is consistent with separation.

In a separating equilibrium, low-ability workers do worse than when education is impossible.

High-ability workers can do better or worse (Figure 13.C.8 at MWG 456; worse is possible because they can't duplicate the no-education outcome, in which they're pooled with low-ability workers, and the education needed to separate is costly).

Perhaps somewhat surprisingly, the set of separating equilibria is independent of λ.

A *pooling* PBE is one in which different worker types choose the same e's, so that in equilibrium, observing a worker's e reveals nothing about his type.

(In some models there can be *partial pooling* equilibria in which one type randomizes and the other doesn't, so that in equilibrium, observing e reveals noisy information about type. I ignore this possibility here.)

Figures 13.C.9-10 at MWG 457.

Figures 13.C.9-10 show the limits of pooling PBEs: e can range from 0 to e', the e that makes a low-ability worker indifferent between being identified as low-ability at e = 0 and being pooled with a high-ability worker at e = e'.

Firms and both worker types all behave optimally on and off the equilibrium path:

Firms because they bid correctly and consistently for each e.

Worker types because each has a generalized tangency between its indifference curve and its (w,e) opportunity locus.

These pooling equilibria are again Pareto-ranked, with the one with e = 0 best for both worker types and the firms indifferent among them.

All pooling equilibria are weakly Pareto-dominated by the equilibrium when education is impossible.

Figure 13.C.7 at MWG 455.

Figure 13.C.7 illustrates the use of a simple equilibrium refinement to break a separating equilibrium in which a high-ability worker chooses e higher than e~.

Any e between e~ and e_1 is *equilibrium-dominated* for the low-ability type, in that it is dominated if (but only if) we assume equilibrium beliefs and bids by firms.

This kind of argument, which goes beyond sequential equilibrium (and monotonicity, etc.) to restrict out-of-equilibrium beliefs, is called in its simplest form the *intuitive criterion.*

Figures 13.C.9-10 at MWG 457.

In this model one can use this kind of argument also to rule out any pooling equilibrium.

Figures 13.C.5-6 at MWG 454-455.

The result is a unique prediction of the *outcome*, that of the separating equilibria with different beliefs in Figures 13.C.5-6.

The separating equilibria with different beliefs in Figures 13.C.5-6 may not be incentive-constrained Pareto-efficient:

If the no-signaling equilibrium Pareto-dominates the separating equilibrium, banning signaling is a Pareto-improvement.

Figure 13.C.11 at MWG 458

If the no-signaling equilibrium does not Pareto-dominate the separating equilibrium, market intervention by setting separate wages for workers with e above and below a properly chosen cutoff (Figure 13.C.11) may still allow a Pareto-improvement by making both worker types better off while allowing firms to break even by cross-subsidization (losing on low-ability workers but gaining on high-ability workers).

In more realistic models, educational signaling can improve matching between workers and jobs and/or enhance productivity. However, the desire to separate can still lead to excessive education relative to what would be optimal with observable ability.

Now consider signaling in an industrial organization example from Milgrom and Roberts' 1982 *EMT* model of entry deterrence, following Kreps 463-480.

Background: Three views of entry deterrence:

(i) Irreversible decisions that affect future interactions: Dixit and Spence entry deterrence and accommodation models (MWG 423-427)

(ii) Repeated games

(iii) Informational as in Milgrom and Roberts' analysis. (See also "crazy types" analyses of Selten's chain-store paradox in Kreps and Wilson, and Milgrom and Roberts 1982 *JET*.)

It was long assumed—Bain-Sylos/Modigliani *JPE* 1958 (one of the more famous book reviews in economics)—that firms could deter entry just by keeping prices low, signaling that they would be tough to complete with if anyone entered.

Milgrom and Roberts noted that although this sounds plausible the signaling argument should be consistent with sequential equilibrium without commitment, and it depends on the existence of post-entry-relevant private information an incumbent's behavior could transmit.

They built a model in which the incumbent has private information about its own variable cost, so its pricing behavior before entry might signal cost, hence behavior following entry.

The conclusion was surprising, as in their working paper's polemical title: "Equilibrium Limit Pricing Doesn't Limit Entry" (instead it just distorts incumbent's pricing in a dissipative way).

Their model has two expected-profit maximizing firms, Incumbent and potential Entrant.

I and E choose Quantities of goods that are perfect substitutes, I only in the first of the two periods, then I and E (if it enters) simultaneously in the second period.

I has two possible values of per-unit variable costs, constant across both periods (so there is a payoff-driven link between first-period price and second-period behavior):

\$3 or \$1 with probabilities ρ and 1- ρ.

Only I observes its cost, but ρ is common knowledge.

It is common knowledge that E's unit cost is \$3, and that both I and E have fixed costs \$3.

The rules are as follows:
(i) Nature chooses I's unit cost type c

(ii) In the first period, I observes its unit cost c and chooses Q, which determines P = 9 - Q

(iii) In the second period, E observes first-period P and decides whether or not to enter.

(iv) If E enters, I and E become Cournot competitors, taking into account whatever information is revealed in equilibrium by I's first-period choice of P. If E does not enter, I becomes a monopolist in the second period.

Figure 13.2 at Kreps 473 gives the extensive form.

Like Spence’s analysis, this one is challenging because a privately informed player, I, plays an active role. In equilibrium, I must weigh the direct effects of its first-period pricing decision against its indirect, informational effects.

First analyze Cournot subgame following entry, taking E's beliefs as given (Kreps 475).

It’s like a three-person Cournot game, with each type of I like a different player. However, the I-types are each fully informed, while E faces uncertainty about I’s quantity choice.

If E assesses that c = 3 has probability μ, the Cournot equilibrium is $Q_E = 2(2+\mu)/3$, $Q_I|(c = 1) = (10 - \mu)/3$, $Q_I|(c = 3) = (7 - \mu)/3$, with $\pi_E = 4(2+\mu)^2/9$, not including its fixed cost of 3.

Proof: E chooses Q_E to maximize $Q_E(9 - Q_E - EQ_I - 3) - 3 = Q_E(6 - Q_E - EQ_I) – 3$ (note certainty-equivalence). FOC: $6 - 2Q_E - EQ_I = 0$ so $Q_E = 3 - EQ_I/2$. (SOC – 2 < 0, okay.)

I_3 (the type of I that observes cost of 3) chooses Q_{I3} to maximize $Q_{I3}(9 - Q_E - Q_{I3} - 3) - 3 = Q_{I3}(6 - Q_E - Q_{I3}) – 3$. FOC: $6 - Q_E - 2Q_{I3} = 0$ so $Q_{I3} = 3 - Q_E/2$. (SOC okay.)

I_1 chooses Q_{I1} to maximize $Q_{I1}(9 - Q_E - Q_{I1} - 1) - 3 = Q_{I1}(8 - Q_E - Q_{I1}) – 3$. FOC: $8 - Q_E - 2Q_{I1} = 0$ so $Q_{I1} = 4 - Q_E/2$. (SOC okay.)

$EQ_I = \mu Q_{I3} + (1-\mu)Q_{I1} = \mu(3 - Q_E/2) + (1-\mu)(4 - Q_E/2) = 4 - \mu - Q_E/2$.

Thus $Q_E = 3 - EQ_I/2 = 3 - (4 - \mu - Q_E/2)/2 =$ (after simplifying) $4/3 + 2\mu/3 = 2(2+\mu)/3$. Similarly, $Q_{I1} = 10/3 - \mu/3 = (10 - \mu)/3$ and $Q_{I3} = 7/3 - \mu/3 = (7 - \mu)/3$.

E's expected profit in this equilibrium is $Q_E(6 - Q_E - EQ_I) - 3 = (4/3 + 2\mu/3)(6 - (4/3 + 2\mu/3) - 4 - \mu - (4/3 + 2\mu/3)/2) - 3 =$ (after simplifying) $4(2+\mu)^2/9 - 3$.

Thus E enters iff $4(2+\mu)^2/9 > 3$, which reduces to $\mu > 0.598$.

E.g. if E knows c = 3 (that is, if μ = 1), then I and E both set $Q_i = 2$ and get $\pi_i = 1$ (= 4 - 3), so it's profitable to enter. But if E knows c = 1 (that is, if μ = 0), then I would set $Q_I = 10/3$ and E would set $Q_E = 4/3$ and get $\pi_E = -11/9$, so it would not be profitable to enter.

Now consider I's first-period decision.

The first-period monopoly optimum is Q = 4, P = 5, and π = 13 if c = 1; but Q = 3, P = 6, π = 6 if c = 3.

Although these ideal monopoly prices are temptingly close for c = 1 and c = 3, there is no weak PBE in which each type of I chooses its monopoly optimum in the first period.

In such an equilibrium E could infer I's type by observing whether P = 5 or 6; E would then enter if P = 6, thinking that c = 3; but stay out if P = 5.

But then the high-cost type of I would get π = 6 in the first period and π = 1 in the second period, less over the two periods than the π = 5 and π = 6 it could get (in the hypothesized separating equilibrium) by switching to P = 5 and thereby preventing E from entering.

The conclusion that there is no equilibrium of this kind does not depend on zero-probability inferences, and therefore holds for any stronger notion as well as weak PBE.

Note that only one type needs to want to defect to break the equilibrium, and this is enough to invalidate it as a prediction even if that type is not realized.

Also note that the desired defection does not need to be part of an equilibrium itself: If an apparently profitable defection exists, it shows that the hypothesized equilibrium is not really an equilibrium, even if the defection does not lead to an equilibrium.

We will look later for a separating equilibrium in which one type of I deviates from its first-period monopoly optimum, and we will find one.

But first consider whether there can be a *pooling* weak PBE, in which both types of I charge the same price with probability one, and are therefore not distinguishable in equilibrium.

If $\rho < 0.598$, there is a weak PBE in which:

(i) Each type of I sets $P = 5$ in the first period

(ii) E sticks with its prior $\rho < 0.598$ and therefore stays out if $P \leq 5$ (in any weak PBE, observing P conveys no information, so E *must* stick with its prior on the equilibrium path)

(iii) E infers that I's costs are high and enters if $P > 5$

(iv) Entry leads to the Cournot equilibrium with E believing that I's costs are high

In this pooling equilibrium, the high-cost I "hides behind" the low-cost I by giving up some first-period profit to mimic the low-cost I, and both types of I successfully forestall entry.

To see that these strategies and beliefs are consistent with sequential equilibrium, note that:

(i) E's strategy is sequentially rational, given its beliefs ($\rho < 0.598$)

(ii) the beliefs are consistent with Bayes' Rule on the equilibrium path

(iii) when c = 1, I charges its favorite first-period price and prevents entry, the best of all possible worlds

(iv) when c = 3, the only way I could do better is by raising P above 5, but this would cause E to enter and thereby lower total profits

Assuming the most pessimistic conjectures about consequences of deviations from equilibrium, as in "E infers that I's costs are high and enters if P > 5," is a characteristic form of analysis ("punishing deviations"), and yields the largest possible set of weak PBE.

Here the beliefs also satisfy a natural *monotonicity* restriction, in that observing a higher P never lowers E's estimate that I's costs are high.

If $\rho > 0.598$, there is no pooling weak PBE.

In such an equilibrium $\mu > 0.598$, which would always lead to entry.

This would make a high-cost I unwilling to deviate from its first-period optimal monopoly price, but the low-cost I would set a different first-period price, even if it didn't prevent entry.

If, however, $\rho > 0.598$ (or in fact for any ρ) there is a separating weak PBE in which:

(i) the high-cost I charges its optimal monopoly price, 6, in the first period

(ii) the low-cost I charges 3.76 in the first period

(iii) E infers that costs are high if $P > 3.76$ and therefore enters

(iv) E infers that costs are low if $P \leq 3.76$ and therefore stays out

(v) both types of I charge their monopoly price in the second period if there is no entry

(vi) entry leads to the Cournot equilibrium with E believing that I's costs are high

In this separating equilibrium, the low-cost I distinguishes itself from the high-cost I by distorting its first-period price enough to prevent the high-cost I from mimicking it.

Entry occurs exactly when it would with complete information; the only effect of asymmetric information is the distortion of the low-cost I's first-period price, which benefits consumers.

To see that these strategies and beliefs are consistent with sequential equilibrium, note that:

(i) E's strategy is sequentially rational, given the hypothesized beliefs

(ii) the beliefs are (trivially) consistent with Bayes' Rule on the equilibrium path (and again monotonic)

(iii) the low-cost I would like to set $P > 3.76$ in the first-period, but that would lead to entry and reduce total profits (easy to check)

(iv) the high-cost I gets $\pi = 6$ in the first period and $\pi = 1$ following entry in the second, just above what it would get by setting $P \leq 3.76$ and forestalling entry (3.76 was chosen to make it just too costly for the high-cost I to mimic the low-cost I in this equilibrium)

These arguments don't depend on ρ, so this profile is a weak PBE for any ρ.

Now consider the Rothschild-Stiglitz competitive (insurance) screening model, transposed into the labor-market example (MWG 460-467, Kreps 638-645 and 649-650)

(Stiglitz's monopoly screening model (MWG 500-501, Kreps 661-680) is a special case of the agency models in Section 3.)

Market structure is almost the same as in Spence's signaling model:

Two firms, but many workers (inessential).

Workers have two ability "types" with productivities $\theta_H > \theta_L > 0$, where $r(\theta_H) = r(\theta_L) = 0$.

$0 < \text{Prob}\{\theta = \theta_H\} = \lambda < 1$. Only workers observe their types, but λ is common knowledge.

Workers no longer choose education, but firms can offer contracts with different "task levels" (e.g. hours) to induce workers to reveal their types by their choice of contract.

Task level has no effect on productivity (inessential; easy to relax).

But type-θ worker with wage w and task level $t \geq 0$ has utility $u(w,t|\theta) = w - c(t,\theta)$, where $c(0,\theta) = 0$, $c_t(t,\theta) > 0$, $c_{tt}(t,\theta) > 0$, $c_\theta(t,\theta) < 0$ for all $t > 0$, and $c_{t\theta}(t,\theta) < 0$ ("single crossing property").

Thus tasks are purely an annoyance here, useful to the firm only because it can use them to screen workers' ability types. (In Rothschild-Stiglitz tasks are like insurance deductibles.)

Modern treatments depart from Rothschild-Stiglitz's by modeling the market as a game. The "rules" are as follows:

(i) nature chooses the workers' types θ

(ii) firms simultaneously offer sets of (any desired finite number of, but two is enough) contracts, each of which is a pair (w,t)

(iii) workers observe their types and each type chooses one of the offered contracts or no contract (assume workers who are indifferent between contracts choose the one with lower t, workers who are indifferent between a contract and no contract choose the contract, and workers whose preferred contract is offered by both firms choose each with probability ½)

Study pure-strategy subgame-perfect Nash equilibria (SPNE). Equivalent to weak PBE here because active players don't have private information, so beliefs are constrained.

Suppose first, as a benchmark, that workers' types are observable, so that firms can condition offers on a worker's type, offering a contract (w_L,t_L) restricted to low-ability workers and a contract (w_H,t_H) restricted to high-ability workers.

Proposition 13.D.1 at MWG 461-462: In any SPNE of the game with observable worker types, a worker of type θ_i accepts contract $(w_i^*,t_i^*)=(\theta_i,0)$, and firms earn zero profits.

Proof: Firms gain by replacing inefficient contracts with $t_i > 0$ by $t_i = 0$, and competition then drives w_i up to θ_i.

Now suppose that workers' types are unobservable, so that any offered contract can be accepted by a worker of either type.

Proposition 13.D.1's full-information outcome is no longer attainable, because low-ability workers prefer the high-ability contract to the low-ability contract, and they can no longer be prevented from accepting it.

Instead we will look for pooling or separating equilibria, via a series of lemmas.

Lemma 13.D.1 at MWG 462-463: In any equilibrium, pooling or separating, both firms earn zero profits.

Proof: If (w_L,t_L) and (w_H,t_H) are the contracts chosen by low- and high-ability workers, respectively, and firms have positive total profits π, at least one firm must make profit $\leq \pi/2$. Such a firm can attract all low-ability workers by offering $(w_L+\varepsilon,t_L)$ and all high-ability workers by offering $(w_H+\varepsilon,t_H)$ for some small $\varepsilon > 0$. Since ε can be as small as desired, that firm can gets profits close to π, and therefore has a profitable deviation unless $\pi \leq 0$. But $\pi < 0$ is impossible because firms aren't required to offer contracts. So in equilibrium $\pi = 0$.

Figure 13.D.3-4 at MWG 463

Lemma 13.D.2 at MWG 463: No pooling equilibria exist.

Proof: If there were a pooling equilibrium contract (w^p,t^p), by Lemma 13.D.1 it would lie on the pooled break-even (0-profit) line in Figure 13.D.3. But then either firm could gain by deviating to a contract (w~,t~) in the shaded lens with $w < \theta_H$, which attracts all high-ability workers and no low-ability workers and thus yields positive profits.

Lemma 13.D.3 at MWG 463: If (w_L,t_L) and (w_H,t_H) are contracts chosen by low- and high-ability workers in a separating equilibrium, $w_L= \theta_L$ and $w_H = \theta_H$ so both yield zero profits.

Proof: If $w_L< \theta_L$ either firm could get positive profits by offering only a contract with w a little above w_L, which all low-ability workers would accept, and which would be profitable for both low- and high-ability workers. This contradicts Lemma 13.D.1, so $w_L \geq \theta_L$ in any separating equilibrium. If $w_H < \theta_H$ as in Figure 13.D.4, then (w_L,t_L) must lie in the lens above the $w_L = \theta_L$ line as shown, by self-selection and Lemma 13.D.1 (0 profits). But then either firm could get positive profits by deviating to a contract in the upper lens below $w_H = \theta_H$ line, like (w~,t~), which would attract all high-ability workers. Thus $w_H \geq \theta_H$ in any separating equilibrium. Since firms break even in any equilibrium by Lemma 13.D.1, in fact $w_L= \theta_L$ and $w_H = \theta_H$.

Lemma 13.D.4 at MWG 464: In any separating equilibrium, low-ability workers accept $(\theta_L,0)$, the same contract they would receive in a full-information competitive equilibrium.

Proof: By Lemma 13.D.3, $w_L = \theta_L$ in any separating equilibrium. If $t_L > 0$ in such an equilibrium, a firm could do better by offering a contract with lower w_L and t_L, attracting all low-ability workers as in Figure 13.D.5 at MWG 464.

Figure 13.D.6 at MWG 464

Lemma 13.D.5 at MWG 464: In any separating equilibrium, high-ability workers accept $(\theta_H,t_H^{\wedge})$, where $t_H^{\wedge}$ is chosen so low-ability workers are indifferent between $(\theta_L,0)$ and $(\theta_H,t_H^{\wedge})$ as in Figure 13.D.6, so that $\theta_H - c(t_H^{\wedge}, \theta_L) = \theta_L - c(0, \theta_L)$.

Proof: By Lemmas 13.D.3-4, $w_H = \theta_H$ and $(w_L,t_L) = (\theta_L,0)$. For low-ability workers to accept $(\theta_L,0)$, we must have $t_H \geq t_H^{\wedge}$ in Figure 13.D.6. If the high-ability contract (θ_H,t_H) has $t_H > t_H^{\wedge}$, then either firm can get positive profits by offering an additional contract (w~,t~) with lower w_H and t_H as in Figure 13.D.6, which attracts all of the high-ability workers and does not change the choices of low-ability workers. Thus, in any separating equilibrium, the high-ability contract must be $(\theta_H,t_H^{\wedge})$.

To sum up:

Proposition 13.D.2: In any subgame-perfect equilibrium of the screening game, low-ability workers accept contract $(\theta_L,0)$ and high-ability workers accept contract $(\theta_H,t_H^{\wedge})$ in Figure 13.D.6, where $\theta_H - c(t_H^{\wedge},\theta_L) = \theta_L - c(0,\theta_L)$.

Proposition 13.D.2 tells what a separating equilibrium must look like if one exists, but does not tell us that such an equilibrium exists.

Figure 13.D.7 at MWG 465

Consider the candidates for a separating equilibrium in Figures 13.D.7.

By construction, for any λ, neither firm can gain from deviating in a way that attracts either all high- or all low-ability workers.

But varying λ allows us to move Eθ anywhere between θ_H and θ_L without affecting the candidate equilibrium. In Figure 13.D.7(b) (but not (a)), Eθ is high enough that a firm can gain by deviating to a contract (w~,t~) that attracts all workers to a *pooling* contract.

In this case, since no pooling equilibrium ever exists, no equilibrium of any kind exists. (In pure strategies; equilibrium does exist in mixed strategies, but their interpretation in this model is problematic.)

As in the signaling model's best separating equilibrium, screening equilibria are Pareto-inefficient, and low-ability workers are worse off than when screening is impossible.

However, when a screening equilibrium exists it must make high-ability workers better off. (Whenever screening would hurt high-ability workers, a pooling contract breaks the screening equilibrium candidate.)

When they exist, screening equilibria are (with a qualification) incentive-efficient.

3. Agency (MWG 477-506; Kreps 577-614 and 661-674; Varian 441-466; McMillan 91-129)

Consider a relationship between two people: a *principal* ("owner" in MWG) who could benefit from delegating a decision that affects his welfare to an *agent* ("manager" in MWG) who has relevant skills or private information.

The agent has different preferences over decisions than the principal would if he were fully informed, and the principal cannot control the agent's decisions (either because he cannot observe them, or for other, unmodeled reasons).

But the principal can design a contract or incentive scheme to influence agent's decisions.

Important distinction between:

Hidden actions/moral hazard (e.g. fire prevention, manager's effort choice that influences owner's profit, borrower's investment decisions that influence lender's return on loan)

Hidden information/adverse selection (insurer unable to observe consumer's risk class).

Distinction is logically independent of signaling-screening distinction.

Applications often have some of both. They are analytically somewhat similar in that in each case the principal cannot observe the agent's *decision rule*.

Hidden-action analysis:

Agent chooses one-dimensional effort level e from set E, which is costly to the agent.

e influences the principal's profit π. Principal wishes to maximize Eπ net of what he pays agent, but cannot observe (or directly control) e.

If relationship between e and π were deterministic, invertible, principal could infer e from π, and thereby control e; so assume π has a conditional density $f(\pi|e) > 0$ for all e ε E and all π ε $[\underline{\pi}, \pi]$, making any value of π consistent with any value of e.

Special case with two effort levels, e_H and e_L:

$f(\pi|e_H)$ first-order stochastically dominates $f(\pi|e_L)$ (that is $F(\pi|e_H) \leq F(\pi|e_L)$) for all $\pi \varepsilon [\underline{\pi}, \bar{\pi}]$, with strict inequality for a nonnegligible set of π's).

$E_{F(\pi|eH)}\pi > E_{F(\pi|eL)}\pi$, so principal prefers agent to choose e_H, other things equal.

Agent chooses $e \varepsilon \{e_H, e_L\}$ to maximize E[v(w) - g(e)], where w is wage, v(·) is strictly increasing and weakly concave so agent is risk-averse in income, and $g(e_H) > g(e_L)$ so agent dislikes effort.

Principal is risk-neutral and maximizes E[π-w].

Ultimatum model of contracting (standard in principal-agent literature): Principal proposes contract to agent, which agent can accept or reject. Acceptance yields binding contract, rejection yields agent reservation utility $\underline{u}$, a proxy for agent's best alternative in the market, assumed exogenous. Assume subgame-perfect equilibrium (SPNE) throughout.

Proposition 14.B.1 at MWG 480-481 (benchmark case): When the agent's effort is observable, an optimal (uniquely optimal when $v(\cdot)$ is strictly concave) contract for the principal specifies that the agent choose the effort e* that solves $\max_{e \varepsilon \{eH,eL\}} [E_{F(\pi|e)}\pi - v^{-1}(\underline{u}+g(e))]$ and pays the agent a fixed wage $w^* = v^{-1}(\underline{u}+g(e^*))$.

Proof: When the agent's effort is observable, a contract specifies agent's effort $e \ \varepsilon \ \{e_H,e_L\}$ and wage $w(\pi)$. The principal's problem is $\max_{e\varepsilon\{eH,eL\},w(\pi)} E_{F(\pi|e)}[\pi - w(\pi)]$ s.t. $E_{F(\pi|e)}v(w(\pi)) - g(e) \geq \underline{u}$, which last constraint is called the *participation* or *individual rationality* constraint.

First consider the best $w(\pi)$ given e, which solves $\min_{w(\pi)}E_{F(\pi|e)}\, w(\pi)$ s.t. $E_{F(\pi|e)}v(w(\pi)) - g(e) \geq \underline{u}$.

Participation constraint is always binding, with Lagrange multiplier γ and first-order condition $\gamma = 1/v'(w(\pi))$ for all π.

Given e, if $v(\cdot)$ is strictly concave this implies that $w(\pi) = w^*(e)$ for all π, and if $v(\cdot)$ is weakly concave $w(\pi) = w^*(e)$ is still one optimum. (The best way for a risk-neutral principal to get a risk-averse agent up to utility level $\underline{u}$ is for the principal to bear all the risk.)

Thus $v(w^*(e)) - g(e) = \underline{u}$, so $w^*(e) = v^{-1}(\underline{u}+g(e))$, with $w^*(e)$ increasing in e. Given $w^*(e) = v^{-1}(\underline{u}+g(e))$, the best e, e*, solves $\max_{e\varepsilon\{eH,eL\}} [E_{F(\pi|e)}\pi - v^{-1}(\underline{u}+g(e))]$.

Proposition 14.B.2 at MWG 482-483: When the agent's effort is *un*observable but the agent is risk-neutral, the optimal contract leads to the same effort and expected utilities for principal and agent as when effort is observable (Proposition 14.B.1).

Proof: Suppose the principal sets $w(\pi) \equiv \pi - \alpha$ for some constant α ("selling the project (for α) to the agent").

The agent then chooses e to solve $\max_{e \varepsilon \{eH,eL\}}[E_{F(\pi|e)}w(\pi) - g(e)] = E_{F(\pi|e)}\pi – \alpha - g(e)$.

When $v(w) \equiv w$, $v^{-1}(w) \equiv w$, so this problem has the same solution e* that solves $\max_{e\varepsilon\{eH,eL\}}$ $[E_{F(\pi|e)}\pi - v^{-1}(\underline{u} + g(e))]$ in Proposition 14.B.1 (the maximands differ by a constant).

Setting $\alpha = \alpha^*$ where $E_{F(\pi|e^*)}\pi - \alpha^* - g(e^*) = \underline{u}$ satisfies the agent's participation constraint and yields the principal utility $\alpha^* = E_{F(\pi|e^*)}\pi - g(e^*) - \underline{u}$, the same as his utility in the optimal contract with observable effort in Proposition 14.B.1. The optimal contract with unobservable effort could not possibly improve on this.

When the agent's effort is *un*observable and the agent is risk-averse, however, there is a tension between efficient risk-sharing and providing efficient incentives for the agent that makes the problem nontrivial. E.g. perfect fire insurance dilutes incentives to take care against fire. The optimal contract is a second-best compromise between these goals.

Proposition 14.B.3 at MWG 483-488: When the agent's effort is *un*observable, the agent is risk-averse, and there are two possible effort choices, the optimal compensation scheme for implementing e_H satisfies $1/v'(w(\pi)) = \gamma + \mu[1 - f(\pi|e_L)/f(\pi|e_H)]$, gives the agent expected utility $\underline{u}$, and involves a larger Ew* than when effort is observable. The optimal scheme for implementing e_L, however, involves the same fixed w as if e were observable.

Whenever the optimal effort with observable e would be e_H, the unobservability of e causes a welfare loss:

Either it is still optimal to implement e_H, in which case the agent faces avoidable risk which the principal must compensate him for (the agent still gets $\underline{u}$).

Or it is now too expensive to implement e_H, and the principal implements e_L even though e_H would allow a Pareto-improvement.

(The fact that unobservable e makes incentive constraints bind and distorts effort downward may not be true for more than two effort levels (MWG 502-504, Exercise 14.B.4 at MWG 507).)

Proposition 14.B.3 at MWG 483-488 (copied): When the agent's effort is *un*observable, the agent is risk-averse, and there are two possible effort choices, the optimal compensation scheme for implementing e_H satisfies $1/v'(w(\pi)) = \gamma + \mu[1 - f(\pi|e_L)/f(\pi|e_H)]$, gives the agent expected utility $\underline{u}$, and involves a larger Ew* than when effort is observable. The optimal scheme for implementing e_L, however, involves the same fixed w as if e were observable.

Proof: When e is unobservable the principal's optimal contract specifies a wage w(π). The best w(π) given e solves $\min_{w(\pi)} E_{F(\pi|e)} w(\pi)$ s.t. two constraints:

(i) $E_{F(\pi|e)}v(w(\pi)) - g(e) \geq \underline{u}$ (*participation* or *individual rationality*)

(ii) e solves $\max_{e'} E_{F(\pi|e')} v(w(\pi)) - g(e')$ (*incentive compatibility*).

If it is desired to *implement* e_L, it is optimal for the principal to offer a fixed wage payment $w^*(e_L) = v^{-1}(\underline{u} + g(e_L))$, as if he were specifying e_L when effort is observable. This makes the agent choose e_L, because effort doesn't affect w and he prefers e_L, other things equal, and yields agent $\underline{u}$ just as when effort is observable. The optimal contract with unobservable effort could not possibly improve on this.

(Copied) The best $w(\pi)$ given e solves $\min_{w(\pi)}E_{F(\pi|e)}w(\pi)$ s.t. two constraints:

(i) $E_{F(\pi|e)}v(w(\pi)) - g(e) \geq \underline{u}$ (*participation* or *individual rationality*)

(ii) e solves $\max_{e'}E_{F(\pi|e')}v(w(\pi)) - g(e')$ (*incentive compatibility*).

If it is desired to implement e_H, constraint (ii) becomes

$E_{F(\pi|eH)}v(w(\pi)) - g(e_H) \geq E_{F(\pi|eL)}v(w(\pi)) - g(e_L)$.

Again letting $\gamma \geq 0$ and $\mu \geq 0$ be the Lagrange multipliers on constraints (i) and (ii) respectively, $w(\pi)$ must satisfy the following first-order condition for all π:

$-f(\pi|e_H) + \gamma v'(w(\pi))f(\pi|e_H) + \mu[f(\pi|e_H) - f(\pi|e_L)]v'(w(\pi)) = 0$,

or $1/v'(w(\pi)) = \gamma + \mu[1 - f(\pi|e_L)/f(\pi|e_H)]$.

When $e = e_H$, both constraints bind, because the agent would like to set $e = e_L$:

Lemma 14.B.1 at MWG 484: In any solution to the principal's problem with $e = e_H$, $\gamma > 0$ and $\mu > 0$.

Proof: Because $f(\pi|e_H)$ first-order stochastically dominates $f(\pi|e_L)$, there must be an open set of π throughout which $f(\pi|e_L)/f(\pi|e_H) > 1$. But if $\gamma = 0$ and $\mu \geq 0$, this contradicts the first-order condition. And if $\mu = 0$, the first-order condition implies a fixed wage payment, which implements e_L, not e_H.

Given this, this first-order condition

$1/v'(w(\pi)) = \gamma + \mu[1 - f(\pi|e_L)/f(\pi|e_H)]$

says that the agent gets a "base payment" (in utility) that is independent of π plus a "bonus" that is higher to the extent that π is evidence (in the sense of the likelihood ratio $f(\pi|e_L)/f(\pi|e_H)$) that he chose e_H.

This evidence affects the bonus not because the principal doubts that the agent chose e_H; in equilibrium, the principal knows this. Paying the agent partly according to the evidence that he chose e_H is just the cheapest way to get him to choose e_H.

We can also use the first-order condition to ask if w(π) must be increasing. Surprisingly, this is not true without further restrictions on f(·), because even when $f(\pi|e_H)$ first-order stochastically dominates $f(\pi|e_L)$, $f(\pi|e_L)/f(\pi|e_H)$ need not be decreasing in π.

The *monotone likelihood ratio property (MLRP)* says that $f(\pi|e_L)/f(\pi|e_H)$ is decreasing in π, so that higher π is evidence in favor of e_H.

(Fig. 14.B.1 at MWG 485-486 and Kreps 494-495 give examples to show why MLRP is necessary.)

One can also use the first-order condition to prove the Mirrlees-Holmstrom *Sufficient Statistic Theorem* (MWG 487-488): If (and only if) π is a sufficient statistic for the agent's choice of e, there is no gain to allowing w to depend on any other available indirect measure of e.

The first-order condition shows that the optimal incentive scheme is generally highly nonlinear and sensitive to the details of the environment, including the distribution f(·).

By contrast, real-world incentive schemes (e.g. sharecropping), tend to be simple and *robust* to environmental details.

Why this is true is still largely an open question; MWG 488 discuss a possible explanation.

Inefficiency makes devices like monitoring and cross-checking useful.

MWG 488 discuss extensions to multiple agents with relative performance evaluation (*tournaments*), long-term relationships, competition for agents among multiple principals, and multidimensional effort.

Refer to MWG 504 and Kreps 604-608 on the "first-order approach" (different from above use of first-order condition) as an imperfect alternative to min $w(\pi)$ given e when e is continuously variable.

Two problems: failure of second-order conditions and discontinuities of e^* in $w(\pi)$.

Hidden-information analysis (MWG 488-501):

Model is almost the same as for hidden actions, but now the agent chooses $e \in [0,\infty)$, e is observable, and the agent's cost of effort is unobservable.

The principal's gross profit (net of wage payments) is $\pi(e)$, with $\pi(0) = 0$, $\pi'(e) > 0$, and $\pi''(e) < 0$ for all e.

The agent's reservation utility is $\underline{u}$, and the agent's vN-M utility function is $u(w,e,\theta) \equiv v(w - g(e,\theta))$ where $v''(\cdot) < 0$, so the agent is risk averse in income.

$g(\cdot)$ measures the cost of effort, with $g(0,\theta) \equiv 0$ for all θ and, for all $e > 0$, $g_e(\cdot) > 0$, $g_{ee}(\cdot) > 0$, $g_\theta(\cdot) < 0$, and $g_{e\theta}(\cdot) < 0$ (the "single-crossing property"), so that e has positive and increasing marginal cost, and both are decreasing in θ.

We focus on the special case with two possible θs, θ_H and θ_L, with commonly known probabilities $\lambda \in (0,1)$ and $1-\lambda$.

Ultimatum model of contracting: Principal proposes contract to agent, which agent can accept or reject. Acceptance yields binding contract, rejection yields agent reservation utility $\underline{u}$, a proxy for agent's best alternative in the market, assumed exogenous. Assume subgame-perfect equilibrium (SPNE) throughout.

Consider first the case where θ is observable, so that the principal can specify the effort level e_i and wage w_i *contingent* on each realization of θ, θ_i.

In the two-outcome case the contract specifies two wage-effort pairs, (w_H,e_H) and (w_L,e_L), and the principal chooses these to solve

$$\max_{(wH,eH) \geq 0,\ (wL,eL) \geq 0} \lambda[\pi(e_H) - w_H] + (1 - \lambda)[\pi(e_L) - w_L]$$

$$\text{s.t. } \lambda v(w_H - g(e_H,\theta_H)) + (1 - \lambda)v(w_L - g(e_L,\theta_L)) \geq \underline{u}$$

Proposition 14.C.1 at MWG 492: When θ is observable, the optimal contract for the principal involves an effort level e_i^* in state θ_i such that $\pi'(e_i^*)=g_e(e_i^*,\theta_i)$, and fully insures the agent, setting the wage in each state θ_i at the level w_i^* such that $v(w_i^* - g(e_i^*,\theta_i)) = \underline{u}$.

Proof: The participation constraint must bind at the solution, because otherwise the principal could increase his profit by lowering wages. Letting $\gamma \geq 0$ be the multiplier on this constraint, we have the first-order conditions:

(14.C.2) $$-\lambda + \gamma\lambda v'(w_H^* - g(e_H^*,\theta_H)) = 0$$

(14.C.3) $$-(1-\lambda) + \gamma(1-\lambda)v'(w_L^* - g(e_L^*,\theta_L)) = 0$$

(14.C.4) $$\lambda\pi'(e_H^*) - \gamma\lambda v'(w_H^* - g(e_H^*,\theta_H))g_e(e_H^*,\theta_H) \leq 0, \text{ and } = 0 \text{ if } e_H^* > 0$$

(14.C.5) $$(1 - \lambda)\pi'(e_L^*) - \gamma(1 - \lambda)v'(w_L^* - g(e_L^*,\theta_L))g_e(e_L^*,\theta_L) \leq 0, \text{ and } = 0 \text{ if } e_L^* > 0$$

(14.C.2) $$-\lambda + \gamma\lambda v'(w_H^* - g(e_H^*,\theta_H)) = 0$$

(14.C.3) $$-(1-\lambda) + \gamma(1-\lambda)v'(w_L^* - g(e_L^*,\theta_L)) = 0$$

Combining (14.C.2) and (14.C.3) yields the standard condition for efficient insurance of a risk-averse party (the agent) by a risk-neutral party (the principal).

(14.C.6) $$v'(w_H^* - g(e_H^*,\theta_H)) = v'(w_L^* - g(e_L^*,\theta_L)).$$

Because $v''(\cdot) < 0$, (14.C.6) implies that $w_H^* - g(e_H^*,\theta_H)) = w_L^* - g(e_L^*,\theta_L)$ and $v(w_H^* - g(e_H^*,\theta_H)) = v(w_L^* - g(e_L^*,\theta_L))$.

Because the participation constraint is binding, the agent has utility $\underline{u}$ in each state.

Because $g_e(0,\theta) = 0$ and $\pi'(0) > 0$, (14.C.4) and (14.C.5) must both hold with equality and with $e_H^* > 0$ and $e_L^* > 0$.

(14.C.2) $$-\lambda + \gamma\lambda v'(w_H^* - g(e_H^*,\theta_H)) = 0$$

(14.C.3) $$-(1-\lambda) + \gamma(1-\lambda)v'(w_L^* - g(e_L^*,\theta_L)) = 0$$

(14.C.4) $$\lambda\pi'(e_H^*) - \gamma\lambda v'(w_H^* - g(e_H^*,\theta_H))g_e(e_H^*,\theta_H) \leq 0, \text{ and } = 0 \text{ if } e_H^* > 0$$

(14.C.5) $$(1 - \lambda)\pi'(e_L^*) - \gamma(1 - \lambda)v'(w_L^* - g(e_L^*,\theta_L))g_e(e_L^*,\theta_L) \leq 0, \text{ and } = 0 \text{ if } e_L^* > 0$$

Combining (14.C.2) and (14.C.4) and (14.C.3) and (14.C.5) yields

(14.C.7) $$\pi'(e_i^*) = g_e(e_i^*,\theta_i), \; i = L,H,$$

the condition for efficient effort choice, the marginal benefit of effort equals its marginal (utility) cost in each state (Figure 14.C.1 at MWG 491). The principal's profit in state i is

$$\Pi_i^* = \pi(e_i^*) - v^{-1}(\underline{u}) - g(e_i^*,\theta_i), \; i = L,H.$$

From (14.C.7), $g_{e\theta}(e,\theta) < 0$, $\pi''(e) < 0$, and $g_{ee}(e,\theta) > 0$ imply $e_H^* > e_L^*$ (Figure 14.C.2 at MWG 492). This completes the proof, showing that when θ is observable, a risk-neutral principal fully insures a risk-averse agent and specifies fully efficient effort for each realization of the state θ_i.

When θ is unobservable, the principal's optimal contract must balance the provision of insurance for the agent against the need to give the agent incentives to make e vary appropriately with θ. (Because θ is unobservable, its relation to e is unobservable.)

Figure 14.C.2 at MWG 492

The first-best outcome of Proposition 14.C.1 is no longer attainable, because the agent always prefers (w_L^*, e_L^*) to (w_H^*, e_H^*) (Figure 14.C.2).

Thus if the agent is asked to report θ (directly, or indirectly by his choice of effort) he will always report $\theta = \theta_L$, and the principal will not realize the first-best outcome.

In characterizing the optimal contract in this case, we must consider the agent's incentives to misrepresent θ and how this affects the outcome.

The task is simplified by the following general result, which shows that in a sense there is no loss of generality in restricting attention to contracts in which the agent is asked to report θ (a *direct revelation mechanism*), and for which truthful reporting is consistent with equilibrium (so the mechanism is *incentive-compatible*).

Proposition 14.C.2 at MWG 493 (Revelation Principle): In determining the optimal contract, the principal can without loss of generality restrict attention to contracts in which:

(i) after the agent observes θ, he is required to report it

(ii) the contract specifies an outcome for each possible report

(iii) for every possible realization of θ, the agent finds it optimal to report θ truthfully.

Proof: Given a particular selection among any multiple equilibria that exist in the game following a set of contract proposals by the principal, one can collapse any contract that creates an incentive for the agent to lie into an equivalent contract that specifies the outcome that lying yields in equilibrium. ("Tell me the truth; I'll lie for you.")

Consider the case where θ is unobservable under the simplifying assumption that the agent is infinitely risk averse (maximin or limit of finite risk aversion).

Write the principal's problem as

$$\max_{(wH,eH) \geq 0,\ (wL\ ,eL) \geq 0} \lambda[\pi(e_H) - w_H] + (1 - \lambda)[\pi(e_L) - w_L] \text{ s.t.}$$

(i) $w_L - g(e_L,\theta_L) \geq v^{-1}(\underline{u})$ (participation constraint for θ_L)

(ii) $w_H - g(e_H,\theta_H) \geq v^{-1}(\underline{u})$ (participation constraint for θ_H)

(iii) $w_H - g(e_H,\theta_H) \geq w_L - g(e_L,\theta_H)$ (incentive-compatibility constraint for θ_H)

(iv) $w_L - g(e_L,\theta_L) \geq w_H - g(e_H,\theta_L)$ (incentive-compatibility constraint for θ_L),

where (w_L,e_L) and (w_H,e_H) are now interpreted as what happens when the agent announces θ_L or θ_H. (There is no loss of generality in doing this, by Proposition 14.C.2.)

The participation/individual rationality constraints are given for the "interim" case where the agent observes his type before contracting.

But with an infinitely risk averse agent this formulation applies equally well to the "ex ante" case where the agent signs the contract before observing his type.

The interim case is often more relevant, and in other models may have different implications than the ex ante case.

$$\max_{(wH,eH) \geq 0,\ (wL\ ,eL) \geq 0} \lambda[\pi(e_H) - w_H] + (1 - \lambda)[\pi(e_L) - w_L] \text{ s.t.}$$

(i) $w_L - g(e_L,\theta_L) \geq v^{-1}(\underline{u})$ (participation constraint for θ_L)

(ii) $w_H - g(e_H,\theta_H) \geq v^{-1}(\underline{u})$ (participation constraint for θ_H)

(iii) $w_H - g(e_H,\theta_H) \geq w_L - g(e_L,\theta_H)$ (incentive-compatibility constraint for θ_H)

(iv) $w_L - g(e_L,\theta_L) \geq w_H - g(e_H,\theta_L)$ (incentive-compatibility constraint for θ_L),

Figure 14.C.3 at MWG 495

Lemma 14.C.1 at MWG 495: Constraint (ii) is never binding, and can be ignored.

Proof: From (i) and (iii), $w_H - g(e_H,\theta_H) \geq w_L - g(e_L,\theta_H) \geq$ (because $g(e_L,\theta_H) \leq g(e_L,\theta_L)$), $w_L - g(e_L,\theta_L) \geq v^{-1}(\underline{u})$. (If Low agents are happy to sign up, High agents must be even happier.)

Lemma 14.C.2 at MWG 495-496: An optimal contract must have $w_L - g(e_L,\theta_L) = v^{-1}(\underline{u})$, so constraint (i) is binding.

Proof: Otherwise the principal could reduce both w_H and w_L by $\varepsilon > 0$, preserving incentive-compatibility and increasing profits. (Since High agents are always happier than Low agents, and the principal can screw both in a balanced way that does not interfere with incentive-compatibility, it is optimal for the principal to screw Low agents to the wall ($\underline{u}$).)

$$\max_{(wH,eH) \geq 0,\ (wL\ ,eL) \geq 0} \lambda[\pi(e_H) - w_H] + (1 - \lambda)[\pi(e_L) - w_L] \text{ s.t.}$$

(i) $w_L - g(e_L,\theta_L) \geq v^{-1}(\underline{u})$ (participation constraint for θ_L)

(ii) $w_H\text{-}g(e_H,\theta_H) \geq v^{-1}(\underline{u})$ (participation constraint for θ_H)

(iii) $w_H - g(e_H,\theta_H) \geq w_L - g(e_L,\theta_H)$ (incentive-compatibility constraint for θ_H)

(iv) $w_L - g(e_L,\theta_L) \geq w_H - g(e_H,\theta_L)$ (incentive-compatibility constraint for θ_L),

Lemma 14.C.3 at MWG 496-497: In any optimal contract:

(i) $e_L \leq e_L^*$ and (ii) $e_H = e_H^*$, where e_L^* and e_H^* are from optimal contract with observable θ.

Figures 14.C.4-6 at MWG 496-497
Proof:

(i): By Lemma 14.C.2 and incentive-compatibility, (w_L,e_L) must be on upper boundary of shaded region in Figure 14.C.4. If $e_L \geq e_L^*$, principal can increase profit by sliding (w_L,e_L) down the agent's $\underline{u}$ indifference curve to (w_L^*,e_L^*) in Figure 14.C.5, leaving Low and High agents' utilities unchanged and continuing to satisfy incentive-compatibility.

(ii): Given $(w_L\hat{},e_L\hat{})$ with $e_L \leq e_L^*$ as in Figure 14.C.6, the principal must find the (w_H,e_H) in the shaded region in Figure 14.C.6 that maximizes his profit in state θ_H. The solution occurs at a tangency like that at $(w_H\sim,e_H^*)$ in the figure, with $e_H=e_H^*$ because the only binding constraint that involves both e_H and θ_H is (iii), incentive-compatibility for the High agent.

The fact that only the incentive-compatibility constraint for the High agent is binding is common in such analyses. (With more than two types this generalizes to: only incentive-compatibility constraints for adjacent types bind, and they only in the "downward" direction.)

Figures 14.C.7-8 at MWG 497-498

Lemma 14.C.4 at MWG 497-498 (also see Appendix B at MWG 504-506): In any optimal contract, $e_L < e_L$*.

Proof: Start with (w_L,e_L)= (w_L*,e_L*) as in Figure 14.C.7 at MWG 497, which by Lemma 14.C.3 determines the state θ_H outcome, (w_H~,e_H*) in the figure. Principal's overall expected profit with (w_L,e_L) = (w_L*,e_L*) is a (λ, 1 - λ)-weighted average of his profits in states θ_H and θ_L, which can be read off the vertical axis in the figure (because π(0) = 0, the principal's profit = -w).

Sliding (w_L,e_L) a little down the Low agent's indifference curve, to (w_L,e_L^) in Figure 14.C.8(a) at MWG 498 (note typo in label of (w_L,e_L^)) yields a 0^{th}-order reduction in profit in state θ_L, because it involves a small change in (w_L,e_L) away from the first-best (w_L*,e_L*) in that state (Envelope Theorem).

But it relaxes the incentive-compatibility constraint in state θ_H and so allows the principal to lower w_H by a 1^{st}-order amount (Figure 14.C.8(b)). This increases the principal's profit.

The more likely is θ_H, the more the principal is willing to distort the θ_L outcome to get higher profits in θ_H. (Follows from the Kuhn-Tucker conditions; see MWG App. B at 504-506.)

Proposition 14.C.3 at MWG 499-500: To sum up, in the hidden-information principal-agent model with an infinitely risk-averse agent, the optimal contract sets $e_H = e_H^*$ and $e_L < e_L^*$, and the agent is inefficiently insured, getting utility > $\underline{u}$ in state θ_H and utility $\underline{u}$ in state θ_L. The principal's expected profit is lower than when θ is observable, while the infinitely risk-averse agent's utility is the same.

The conclusions would be the same if π were not publicly observable, in which case we could allow θ to affect the relationship between π and e (replacing π(e) by $\pi_L(e)$ and $\pi_H(e)$).

(We couldn't do this if π were observable, because then the principal could infer θ from π and the specified e.)

Stiglitz's (1977 *REStud*) analysis of monopolistic screening with adverse selection is just like this, except that the principal's profit depends directly on the agent's private information (MWG 500-501). Although this makes little difference, it is instructive to give Stiglitz's analysis, without assuming an infinitely risk-averse agent. Here I follow Kreps 661-674.

Figure 18.2-3 at Kreps 664-665

Basic model (Figure 18.3 at Kreps 665): Risk-neutral, expected-profit maximizing insurer (principal), risk-averse consumer (agent) with probability π_i that endowment will be Y_2 and probability 1- π_i that endowment will be Y_1, i = H,L, where $\pi_H > \pi_L$. Ultimatum contracting.

Graphing consumer's indifference curve in (y_1,y_2)-space, slope = $-(1-\pi)/\pi$ along 45° line in Figure 18.3. Risk-neutral insurer's indifference curves are linear with slope $-(1-\pi)/\pi$. Efficient insurance has tangency on the 45° line, risk-averse consumer perfectly insured.

$Y_2 < Y_1$, so Y_2 is the "accident" outcome and the consumer's endowment (Y_1,Y_2) is below the 45° line. Thus, optimal monopolistic contract when insurer knows π_i is on the 45° line where it intersects the consumer's indifference curve through his endowment (Figure 18.3).

(Competition with known π_i would also yield a contract on the 45° line, but drive insurers' expected profits to 0.)

When insurer doesn't know π_i but does know prior, ρ, that i = H, he would like to use the full-information-optimal monopolistic contracts just derived.

But because low-risk consumer cares less about y_2, his indifference curves are uniformly steeper in (y_1,y_2)-space than a high-risk consumer's (Figure 18.2). The optimal monopolistic contract for low-risk consumers will then look better for high- as well as low-risk consumers, so the insurer won't get the anticipated level of profit by offering the full-information-optimal monopolistic contracts.

When π_i is unknown, we can characterize the optimal contracts as follows. No loss of generality in restricting insurer to two contracts (two = number of consumer types), one designed for high-risk consumers and one for low-risk consumers, imposing incentive compatibility constraints to ensure that consumer types select the right contracts.

Given this, the insurer solves:

$$\max_{(y1H,y2H),\,(y1L\ y2L)} \rho[(1-\pi_H)(Y_1 - y_1^H) + \pi_H (Y_2 - y_2^H)] + (1-\rho)\,[(1-\pi_L)(Y_1 - y_1^L) + \pi_L(Y_2 - y_2^L)] \text{ s.t.}$$

$$(1-\pi_H)u(y_1^H) + \pi_H\, u(y_2^H) \geq (1-\pi_H)u(Y_1) + \pi_H\, u(Y_2) \quad \text{(participation constraint for H)}$$

$$(1-\pi_H)u(y_1^H) + \pi_H\, u(y_2^H) \geq (1-\pi_H)u(y_1^L) + \pi_H\, u(y_2^L) \quad \text{(incentive compatibility constraint for H)}$$

$$(1-\pi_L)u(y_1^L) + \pi_L\, u(y_2^L) \geq (1-\pi_L)u(Y_1) + \pi_L\, u(Y_2) \quad \text{(participation constraint for L)}$$

$$(1-\pi_L)u(y_1^L) + \pi_L\, u(y_2^L) \geq (1-\pi_L)u(y_1^H) + \pi_L\, u(y_2^H) \quad \text{(incentive compatibility constraint for L)}$$

Figure 18.5-6 at Kreps 671, 674

Proposition 1 at Kreps 670: At the solution, participation for L and incentive compatibility for H are binding, and the high-risk contract (y_1^H, y_2^H) has full insurance, with $y_1^H = y_2^H$. Sketch of proof (Kreps 670-674): (i) Participation for L binds because can't have both participation constraints slack, and the one for L binds before the one for H: low-risk consumers need insurance less than high-risk consumers (Figure 18.5(b) at Kreps 671). (ii) Incentive compatibility for H binds because first-best contracts would make high-risk consumers want the low-risk contract. (iii) Thus the solution is as in Figure 18.6 at Kreps 674. Where it lies on 45° line depends on ρ: If ρ is near 1 (0), near (maybe at) ideal contract for Highs (Lows).

4. Incentives and Mechanism Design (MWG 857-910; Kreps 661-703; McMillan 133-159)

Agency theory is a leading case of the theory of incentives and mechanism design.

MWG Examples 23.B.1 (Abstract Social Choice), 23.B.2 (Pure Exchange), 23.B.3 (Public Project), and 23.B.4 (Allocation of a Single Unit of Indivisible Private Good) give four interesting examples of choice problems for which creating incentives for agents to reveal their preferences so that a Pareto-efficient outcome can be achieved is problematic.

Focus on MWG Example 23.B.4, where they assume that all agents have preferences that are quasilinear in money, so that an allocation is efficient if and only if it allocates the object to the agent whose money value for it is highest with probability 1, and no money is wasted.

Here the literature has concentrated on two leading cases:

Bilateral trade (owner and a potential buyer, who may or may not value the object more).

Auctions (a seller who values the good at 0 and I buyers, say I = 2, who may or may not value the object more).

I will further narrow the focus here to auctions, but bilateral trade and the Myerson-Satterthwaite Theorem are also of great interest (MWG 894-910, Kreps 680-703). I also focus throughout on the independent-private-values case where bidders can learn nothing about their own values from others' values. Auctions with common and/or affiliated values raise significant new difficulties.

In the auction case, suppose each buyer's value for the object is i.i.d., uniform on [0, 1].

Imagine that we try to implement a mechanism (MWG equations 23.B.3-8 at 863) that always gives the object to the buyer who states that he values it most, breaking ties in favor of buyer 1, with that buyer paying his stated value to the seller and the other buyer paying nothing.

(This is an example of a direct mechanism, in which each agent is asked to report his value (more generally, his private-information type) and then some rule is used to map the profile of reported values or types into an allocation.)

If buyers tell the truth, this mechanism ensures efficient allocation, with all of the surplus going to the seller.

But buyers are not required to tell the truth, and there is no way to check up on them.

Yet for the mechanism to work as intended, buyers most report truthfully; more precisely, telling the truth must be a Bayesian equilibrium in the game created by the mechanism.

If buyer 2 always announces his true value, will it also be optimal for buyer 1 to do so?

Given the rules of the mechanism, buyer 1's optimal report θ_1^ solves

Choose θ_1^ to solve Max (θ_1 - θ_1^)Pr{$\theta_2 \leq \theta_1$^} = (given uniformity) (θ_1 - θ_1^)θ_1^.

The optimal θ_1^ = $\theta_1/2$.

As in a first-price auction (analyzed below), it is optimal for buyer 1 to shade his report somewhat below his true value: This creates some risk of not receiving the object even though he would be willing to pay the price required to bid it away from buyer 2 (if θ_1^ ≤ θ_2), but makes up for that by lowering the price he has to pay when he does receive the object.

Buyer 2's analysis is of course the same.

Thus this apparently sensible mechanism may not work as desired.

Now imagine that we try to implement a modified version of this mechanism that still gives the object to the buyer who states that he values it most, breaking ties in favor of buyer 1, but instead of buyer i paying the seller his stated value if he wins the object, he now pays the seller the second-highest stated value (in this case, that of the only other buyer).

Again, if buyers tell the truth, this mechanism ensures efficient allocation, though now with some of the surplus going to the winning buyer as well as the seller.

Again, for the mechanism to work as intended, telling the truth must be a Bayesian equilibrium in the game created by the mechanism.

But now it is a dominant strategy for each agent to state his true value for the object. (And so truthtelling is a Bayesian equilibrium.)

Now a buyer's stated value does not affect what he has to pay if he wins, because that is completely determined by the other buyer's stated value.

The only thing his stated value affects is whether or not he wins the object.

And if he states a value other than his true value, this creates a chance that he will either win the object but have to pay more than it is worth to him (if $\theta_1 < \theta_2\text{^} \leq \theta_1\text{^}$) or lose the object when he would have been willing to pay enough to get it (if $\theta_1\text{^} < \theta_2\text{^} < \theta_1$).

Because stating his true value does not affect his payment when he wins, it is optimal.

Thus this modified mechanism can be expected to work as desired.

These examples motivate the question of what can and cannot be accomplished by designing a mechanism.

In posing this question, we should not restrict attention to direct mechanisms without further thought: Even though they are simple and natural, if we could do better (by whatever criterion is applied) with an indirect mechanism, we would presumably want to do so.

As we will see, however, there is a powerful revelation principle argument that restricting attention to direct mechanisms entails no loss of generality: Anything we could accomplish with an indirect mechanism, we can duplicate (under reasonable assumptions) with a direct mechanism.

We should also not restrict attention to incentive-compatible direct mechanisms—those that create incentives for agents to report their types truthfully (in the sense of making truth-telling a Bayesian equilibrium): If we could only do better with a mechanism that created incentives to lie, we would presumably want to do so.

However, the revelation principle also shows that restricting attention to incentive-compatible direct mechanisms entails no loss of generality.

The revelation principle yields powerful techniques for computing optimal mechanisms.

First, however, we will look more closely at commonly used mechanisms for the auction case.

Consider a first-price sealed-bid auction (MWG Example 23.B.5 at 865-6), under the same assumptions maintained before: seller's value 0, buyers' values i.i.d. uniform on [0, 1], I = 2: The highest bidder wins the object and pays his bid to the seller.

Let's look for a Bayesian equilibrium in which each bidder's strategy $b_i(\theta_i) = \alpha_i\theta_i$ for some α_i between 0 and 1. Suppose that buyer 2's strategy takes that form. Then buyer 1's optimal bid solves

Choose $b_1(\theta_1)$ to solve

Max $(\theta_1 - b_1)\Pr\{b_2(\theta_2) \leq b_1\}$ = (given uniformity and $b_2(\theta_2) = \alpha_2\theta_2$) $(\theta_1 - b_1)b_1/\alpha_2$

The optimal $b_1(\theta_1) = \theta_1/2$ if $\theta_1/2 \leq \alpha_2$ or α_2 if $\theta_1/2 > \alpha_2$.

The conclusion is of course symmetric for buyer 2.

$\alpha_1 = \alpha_2 = ½$, so that $b_1(\theta_1) = \theta_1/2$ and $b_2(\theta_2) = \theta_2/2$, is a Bayesian Nash equilibrium.

This equilibrium yields an efficient allocation, because the buyer with the highest value always wins the object and no money is wasted.

Thus, the first-price sealed-bid auction indirectly implements a mechanism that gives the object to the buyer whose bid is highest, with the implied payment to the seller.

Now consider a second-price sealed-bid auction (MWG Example 23.B.6 at 866): The highest bidder wins the object and pays the second-highest bid to the seller.

As in the example, $b_i(\theta_i) = \theta_i$ is a dominant bidding strategy for each buyer, and truthful bidding is thus a Bayesian Nash equilibrium.

Thus the second-price auction is an incentive-compatible direct mechanism, and implements a mechanism that gives the object to the buyer whose bid is highest, with the implied payment to the seller.

First- and second-price sealed-bid auctions have familiar progressive counterparts:

A second-price auction is theoretically equivalent (under certain assumptions) to an English auction, in which bids increase until all but one buyer drops out, the remaining buyer winning and paying (approximately) the second-highest bidder's last bid. Here, a strategy for the dynamic game can be represented by a cut-off above which the buyer will not bid, and this (theoretically) functions like the sealed bid in a second-price auction.

Similarly, a first-price auction is theoretically equivalent to a Dutch auction, in which the asking price decreases until one buyer says he is willing to pay it. Here, a strategy for the dynamic game can be represented by a cut-off at which the buyer will say he is willing, and this (theoretically) functions like the sealed bid in a first-price auction.

(The word auction comes from the Latin word for increase, so Roman auctions were probably "English".)

Now consider what is possible considering any mechanism, not just a direct, incentive-compatible one.

In general, a mechanism is any physically feasible way to map actions by the agents into an allocation, with no restrictions on how agents choose the actions as functions of their types, and no restrictions on how the choices interact to determine the outcome (MWG 866-867).

A mechanism, once specified, induces a Bayesian game among the agents.

A mechanism is said to implement (MWG definitions 23.B.4, 23.D.2) an outcome function if there is an equilibrium in the game that yields the specified relation between agents' true types and the outcome.

(Sometimes implementation is defined more stringently, e.g. by requiring that all equilibria yield the specified relation. I ignore this issue here.)

The notion of implementation depends on what kind of equilibria are allowed. Here I focus on Bayesian Nash equilibrium, yielding "Bayesian implementation" (MWG 883-891).

A direct revelation mechanism (MWG definition 23.B.5) is one in which agents' actions are reports of their types.

An incentive-compatible mechanism (MWG definitions 23.B.6, Example 23.B.7, definition 23.D.3) is a direct mechanism in which truthful reporting is a Bayesian Nash equilibrium.

The key result (MWG Propositions 23.C.1 at 871 for dominant-strategy equilibrium, and 23.D.1 at 884 for Bayesian Nash equilibrium) is the Revelation Principle:

If there is a mechanism that implements an outcome function (in dominant strategies, or respectively in Bayesian equilibrium), then there is a direct, incentive-compatible mechanism that also implements it (making truthtelling a dominant strategy, or respectively a Bayesian equilibrium).

Thus, in a sense there is no loss of generality in restricting attention to incentive-compatible direct mechanisms, even if other mechanisms are feasible.

Now reconsider the auction setting, under the same assumptions maintained before: seller's value 0, buyers' values i.i.d. uniform on [0, 1], I = 2.

Assume that buyers are risk-neutral.

Restrict attention to direct mechanisms, but allow the possibility of a random assignment of the object: $y_i(\theta)$ is buyer i's probability of winning the object when the vector of announced values is $\theta = (\theta_1,\ldots,\theta_I)$ and $t_i(\theta)$ is his expected transfer in the mechanism (all that matters about transfers, given that buyers are risk-neutral).

Because buyers are risk-neutral, buyer i's expected utility is $\theta_i y_i(\theta) + t_i(\theta)$.

What does it take for a mechanism to be incentive-compatible (MWG Proposition 23.D.2)?

Let $y_i^{-}(\theta_i\hat{})$ ≡ the probability that i gets the object given that he announces his type to be $\theta_i\hat{}$ and the others announce truthfully; let $t_i^{-}(\theta_i\hat{})$ ≡ i's expected transfer given that he announces his type as $\theta_i\hat{}$ and the others announce truthfully; and let $U_i(\theta_i) \equiv \theta_i y_i^{-}(\theta_i) + t_i^{-}(\theta_i)$.

Specializing here to the auction case, the mechanism is incentive-compatible if and only if:

(i) $y_i^{-}(\theta_i\hat{})$ is nondecreasing, and

(ii) $U_i(\theta_i) = U_i(0)$ + integral from 0 to θ_i of $y_i^{-}(s)ds$ for all θ_i

The proof () proceeds by turning the conditions for truthful revelation to be optimal into a condition on the derivates of $U_i(\theta_i)$, when they exist, which they do almost everywhere, integrating the derivative condition, and using the boundary condition.

MWG 887-891:

Bayesian Incentive Compatibility with Linear Utility

Suppose now that each agent i's Bernoulli utility function takes the form

$$u_i(x, \theta_i) = \theta_i v_i(k) + (\bar{m}_i + t_i).$$

As before, we shall normalize $\bar{m}_i = 0$ for all i. We also suppose that each agent i's type lies in an interval $\Theta_i = [\underline{\theta}_i, \bar{\theta}_i] \subset \mathbb{R}$ with $\underline{\theta}_i \neq \bar{\theta}_i$, and that the agents' types are statistically independent. We let the distribution function of θ_i be denoted $\Phi_i(\cdot)$, and we assume that it has an associated density $\phi_i(\cdot)$ satisfying $\phi_i(\theta_i) > 0$ for all $\theta_i \in [\underline{\theta}_i, \bar{\theta}_i]$.

We begin by deriving a necessary and sufficient condition for a social choice function $f(\cdot) = (k(\cdot), t_1(\cdot), \ldots, t_I(\cdot))$ to be Bayesian incentive compatible. It is convenient to define $\bar{t}_i(\hat{\theta}_i) = E_{\theta_{-i}}[t_i(\hat{\theta}_i, \theta_{-i})]$; this is agent i's expected transfer given that he announces his type to be $\hat{\theta}_i$ and that all agents $j \neq i$ truthfully reveal their types. Likewise, we let $\bar{v}_i(\hat{\theta}_i) = E_{\theta_{-i}}[v_i(k(\hat{\theta}_i, \theta_{-i}))]$ denote agent i's expected "benefit" conditional on announcing $\hat{\theta}_i$. Because of the form of agents' utility functions, we can write agent i's expected utility when he is type θ_i and announces his type to be $\hat{\theta}_i$ (assuming that all agents $j \neq i$ tell the truth) as[32]

$$E_{\theta_{-i}}[u_i(f(\hat{\theta}_i, \theta_{-i}), \theta_i)|\theta_i] = \theta_i \bar{v}_i(\hat{\theta}_i) + \bar{t}_i(\hat{\theta}_i). \qquad (23.D.10)$$

It is also convenient to define for each i the function

$$U_i(\theta_i) = \theta_i \bar{v}_i(\theta_i) + \bar{t}_i(\theta_i),$$

giving agent i's expected utility from the mechanism conditional on his type being θ_i when he and all other agents report their true types.

Proposition 23.D.2: The social choice function $f(\cdot) = (k(\cdot), t_1(\cdot), \ldots, t_I(\cdot))$ is Bayesian incentive compatible if and only if, for all $i = 1, \ldots, I$,

(i) $\bar{v}_i(\cdot)$ is nondecreasing. (23.D.11)

(ii) $U_i(\theta_i) = U_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta_i} \bar{v}_i(s)\, ds$ for all θ_i. (23.D.12)

Proof: (i) *Necessity*. Bayesian incentive compatibility implies that for each $\hat{\theta}_i > \theta_i$ we have

$$U_i(\theta_i) \geq \theta_i \bar{v}_i(\hat{\theta}_i) + \bar{t}_i(\hat{\theta}_i) = U_i(\hat{\theta}_i) + (\theta_i - \hat{\theta}_i)\bar{v}_i(\hat{\theta}_i)$$

and

$$U_i(\hat{\theta}_i) \geq \hat{\theta}_i \bar{v}_i(\theta_i) + \bar{t}_i(\theta_i) = U_i(\theta_i) + (\hat{\theta}_i - \theta_i)\bar{v}_i(\theta_i).$$

Thus,

$$\bar{v}_i(\hat{\theta}_i) \geq \frac{U_i(\hat{\theta}_i) - U_i(\theta_i)}{\hat{\theta}_i - \theta_i} \geq \bar{v}_i(\theta_i). \qquad (23.D.13)$$

Expression (23.D.13) immediately implies that $\bar{v}_i(\cdot)$ must be nondecreasing (recall that we have taken $\hat{\theta}_i > \theta_i$). In addition, letting $\hat{\theta}_i \to \theta_i$ in (23.D.13) implies that for all θ_i we have

$$U_i'(\theta_i) = \bar{v}_i(\theta_i)$$

and so

$$U_i(\theta_i) = U_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta_i} \bar{v}_i(s)\, ds \qquad \text{for all } \theta_i.$$

(ii) *Sufficiency*. Consider any θ_i and $\hat{\theta}_i$ and suppose without loss of generality that $\theta_i > \hat{\theta}_i$. If (23.D.11) and (23.D.12) hold, then

$$\begin{aligned} U_i(\theta_i) - U_i(\hat{\theta}_i) &= \int_{\hat{\theta}_i}^{\theta_i} \bar{v}_i(s)\, ds \\ &\geq \int_{\hat{\theta}_i}^{\theta_i} \bar{v}_i(\hat{\theta}_i)\, ds \\ &= (\theta_i - \hat{\theta}_i)\bar{v}_i(\hat{\theta}_i). \end{aligned}$$

Hence,

$$U_i(\theta_i) \geq U_i(\hat{\theta}_i) + (\theta_i - \hat{\theta}_i)\bar{v}_i(\hat{\theta}_i) = \theta_i \bar{v}_i(\hat{\theta}_i) + \bar{t}_i(\hat{\theta}_i).$$

Similarly, we can derive that

$$U_i(\hat{\theta}_i) \geq U_i(\theta_i) + (\hat{\theta}_i - \theta_i)\bar{v}_i(\theta_i) = \hat{\theta}_i \bar{v}_i(\theta_i) + \bar{t}_i(\theta_i).$$

So $f(\cdot)$ is Bayesian incentive compatible. ■

Proposition 23.D.2 shows that to identify all Bayesian incentive compatible social choice functions in the linear setting, we can proceed as follows: First identify which functions $k(\cdot)$ lead every agent i's expected benefit function $\bar{v}_i(\cdot)$ to be nondecreasing. Then, for each such function, identify the expected transfer functions $\bar{t}_1(\cdot), \ldots, \bar{t}_I(\cdot)$ that satisfy condition (23.D.12) of the proposition. Substituting for $U_i(\cdot)$, these are precisely the expected transfer functions that satisfy, for $i = 1, \ldots, I$,

$$\bar{t}_i(\theta_i) = \bar{t}_i(\underline{\theta}_i) + \underline{\theta}_i v_i(\underline{\theta}_i) - \theta_i v_i(\theta_i) + \int_{\underline{\theta}_i}^{\theta_i} \bar{v}_i(s)\, ds$$

for some constant $\bar{t}_i(\underline{\theta}_i)$. Finally, choose any set of transfer functions $(t_1(\theta), \ldots, t_I(\theta))$ such that $E_{\theta_{-i}}[t_i(\theta_i, \theta_{-i})] = \bar{t}_i(\theta_i)$ for all θ_i. In general, there are many such functions $t_i(\cdot, \cdot)$; one, for example, is simply $t_i(\theta_i, \theta_{-i}) = \bar{t}_i(\theta_i)$.[33]

Auctions: the revenue equivalence theorem

Let us consider again the auction setting introduced in Example 23.B.4: Agent 0 is the seller of an indivisible object from which he derives no value, and agents $1, \ldots, I$ are potential buyers.[34] It will be convenient, however, to generalize the set of possible alternatives relative to those considered in Example 23.B.4 by allowing for a *random* assignment of the object. Thus, we now take $y_i(\theta)$ to be buyer i's *probability* of getting the object when the vector of announced types is $\theta = (\theta_1, \ldots, \theta_I)$. Buyer i's expected utility when the profile of types for the I buyers is $\theta = (\theta_1, \ldots, \theta_I)$ is then $\theta_i y_i(\theta) + t_i(\theta)$. Note that buyer i is risk neutral with respect to lotteries both over transfers and over the allocation of the good.

This setting corresponds in the framework studied in Proposition 23.D.2 to the case where we take $k = (y_1, \ldots, y_I)$, $K = \{(y_1, \ldots, y_I): y_i \in [0, 1]$ for all $i = 1, \ldots, I$ and $\sum_i y_i \leq 1\}$, and $v_i(k) = y_i$. Thus, to apply Proposition 23.D.2 we can write $\bar{v}_i(\hat{\theta}_i) = \bar{y}_i(\hat{\theta}_i)$, where $\bar{y}_i(\hat{\theta}_i) = E_{\theta_{-i}}[y_i(\hat{\theta}_i, \theta_{-i})]$ is the probability that i gets the object conditional on announcing his type to be $\hat{\theta}_i$ when agents $j \neq i$ announce their types truthfully, and $U_i(\theta_i) = \theta_i \bar{y}_i(\theta_i) + \bar{t}_i(\theta_i)$.

Proposition 23.D.3: (*The Revenue Equivalence Theorem*) Consider an auction setting with I risk-neutral buyers, in which buyer i's valuation is drawn from an interval $[\underline{\theta}_i, \bar{\theta}_i]$ with $\underline{\theta}_i \neq \bar{\theta}_i$ and a strictly positive density $\phi_i(\cdot) > 0$, and in which buyers' types are statistically independent. Suppose that a given pair of Bayesian Nash equilibria of two different auction procedures are such that for every buyer i: (i) For each possible realization of $(\theta_1, \ldots, \theta_I)$, buyer i has an identical probability of getting the good in the two auctions; and (ii) Buyer i has the same expected utility level in the two auctions when his valuation for the object is at its lowest possible level. Then these equilibria of the two auctions generate the same expected revenue for the seller.

Proof: By the revelation principle, we know that the social choice function that is (indirectly) implemented by the equilibrium of any auction procedure must be Bayesian incentive compatible. Thus, we can establish the result by showing that if two Bayesian incentive compatible social choice functions in this auction setting have the same functions $(y_1(\theta), \ldots, y_I(\theta))$ and the same values of $(U_1(\underline{\theta}_1), \ldots, U_I(\underline{\theta}_I))$ then they generate the same expected revenue for the seller.

To show this, we derive an expression for the seller's expected revenue from an arbitrary Bayesian incentive compatible mechanism. Note, first, that the seller's expected revenue is equal to $\sum_{i=1}^{I} E[-t_i(\theta)]$. Now,

$$\begin{aligned}
E[-t_i(\theta)] &= E_{\theta_i}[-\bar{t}_i(\theta_i)] \\
&= \int_{\underline{\theta}_i}^{\bar{\theta}_i} [\bar{y}_i(\theta_i)\theta_i - U_i(\theta_i)]\phi_i(\theta_i)\, d\theta_i \\
&= \int_{\underline{\theta}_i}^{\bar{\theta}_i} \left(\bar{y}_i(\theta_i)\theta_i - U_i(\underline{\theta}_i) - \int_{\underline{\theta}_i}^{\theta_i} \bar{y}_i(s)\, ds\right)\phi_i(\theta_i)\, d\theta_i \\
&= \left[\int_{\underline{\theta}_i}^{\bar{\theta}_i} \left(\bar{y}_i(\theta_i)\theta_i - \int_{\underline{\theta}_i}^{\theta_i} \bar{y}_i(s)\, ds\right)\phi_i(\theta_i)\, d\theta_i\right] - U_i(\underline{\theta}_i).
\end{aligned}$$

Moreover, integration by parts implies that

$$\begin{aligned}
\int_{\underline{\theta}_i}^{\bar{\theta}_i} \left(\int_{\underline{\theta}_i}^{\theta_i} \bar{y}_i(s)\, ds\right)\phi_i(\theta_i)\, d\theta_i &= \left(\int_{\underline{\theta}_i}^{\bar{\theta}_i} \bar{y}_i(\theta_i)\, d\theta_i\right) - \left(\int_{\underline{\theta}_i}^{\bar{\theta}_i} \bar{y}_i(\theta_i)\Phi_i(\theta_i)\, d\theta_i\right) \\
&= \int_{\underline{\theta}_i}^{\bar{\theta}_i} \bar{y}_i(\theta_i)(1 - \Phi_i(\theta_i))\, d\theta_i.
\end{aligned}$$

Substituting, we see that

$$E[-\bar{t}_i(\theta_i)] = \left[\int_{\underline{\theta}_i}^{\bar{\theta}_i} \bar{y}_i(\theta_i) \left(\theta_i - \frac{1 - \Phi_i(\theta_i)}{\phi_i(\theta_i)} \right) \phi_i(\theta_i)\, d\theta_i \right] - U_i(\underline{\theta}_i), \qquad (23.D.14)$$

or, equivalently,

$$E[-\bar{t}_i(\theta_i)] =$$

$$\left[\int_{\underline{\theta}_1}^{\bar{\theta}_1} \cdots \int_{\underline{\theta}_I}^{\bar{\theta}_I} y_i(\theta_1, \ldots, \theta_I) \left(\theta_i - \frac{1 - \Phi_i(\theta_i)}{\phi_i(\theta_i)} \right) \left(\prod_{j=1}^{I} \phi_j(\theta_j) \right) d\theta_I \cdots d\theta_1 \right] - U_i(\underline{\theta}_i). \qquad (23.D.15)$$

Thus, the seller's expected revenue is equal to

$$\left[\int_{\underline{\theta}_1}^{\bar{\theta}_1} \cdots \int_{\underline{\theta}_I}^{\bar{\theta}_I} \left[\sum_{i=1}^{I} y_i(\theta_1, \ldots, \theta_I) \left(\theta_i - \frac{1 - \Phi_i(\theta_i)}{\phi_i(\theta_i)} \right) \right] \left(\prod_{j=1}^{I} \phi_j(\theta_j) \right) d\theta_I \cdots d\theta_1 \right] - \sum_{i=1}^{I} U_i(\underline{\theta}_i). \qquad (23.D.16)$$

By inspection of (23.D.16), we see that any two Bayesian incentive compatible social choice functions that generate the same functions $(y_1(\theta), \ldots, y_I(\theta))$ and the same values of $(U_1(\underline{\theta}_1), \ldots, U_I(\underline{\theta}_I))$ generate the same expected revenue for the seller. ■

Now reconsider the equilibria of the first- and second-price sealed-bid auctions, or equivalently their progressive Dutch and English counterparts.

With symmetric buyers both assure that the highest valuer wins the object, and both assure that a buyer with zero value doesn’t win and gets 0 expected payoff. Thus the Revenue-Equivalence Theorem shows that both yield the same expected revenue to the seller.

In the second-price or English auction, this expected revenue is always exactly the second-highest value among the buyers.

In the first-price or Dutch auction, buyers end up paying the highest bid, so it may appear that revenue is higher on average.

However, buyers also shade their bids, and it turns out that the optimal shading (for a risk-neutral buyer) is to shade down to the expectation of the second-highest value, so the seller’s expected revenue is exactly the same as in the second-price auction.

But the realization of the seller’s revenue may be higher or lower in a first-price auction, depending on whether the second-highest value is lower or higher than its expectation).

MWG 903-905:

Example 23.F.2: *Optimal Auctions.* We consider again the auction setting introduced in Example 23.B.4. Here we determine the optimal auction for the seller of an indivisible object (agent 0) when there are I buyers, indexed by $i = 1, \ldots, I$. Each buyer has a Bernoulli utility function $\theta_i y_i(\theta) + t_i(\theta)$, where $y_i(\theta)$ is the probability that agent i gets the good when the agents' types are $\theta = (\theta_1, \ldots, \theta_I)$. In addition, each buyer i's type is independently drawn according to the distribution function $\Phi_i(\cdot)$ on $[\underline{\theta}_i, \bar{\theta}_i] \subset \mathbb{R}$ with $\underline{\theta}_i \neq \bar{\theta}_i$ and associated density $\phi_i(\cdot)$ that is strictly positive on $[\underline{\theta}_i, \bar{\theta}_i]$. We assume also that, for $i = 1, \ldots, I$,

$$\theta_i - \frac{1 - \Phi_i(\theta_i)}{\phi_i(\theta_i)}$$

is nondecreasing in θ_i.[50]

A social choice function in this environment is a function $f(\cdot) = (y_0(\cdot), \ldots, y_I(\cdot), t_0(\cdot), \ldots, t_I(\cdot))$ having the properties that, for all $\theta \in \Theta$, $y_i(\theta) \in [0, 1]$ for all i, $\sum_{i \neq 0} y_i(\theta) = 1 - y_0(\theta)$, and $t_0(\theta) = -\sum_{i \neq 0} t_i(\theta)$.[51] The seller wishes to choose the Bayesian incentive compatible social choice function that maximizes his expected revenue $E_\theta[t_0(\theta)] = -E_\theta[\sum_{i \neq 0} t_i(\theta)]$ but faces the interim individual rationality constraints that $U_i(\theta_i) = \theta_i \bar{y}_i(\theta_i) + \bar{t}_i(\theta_i) \geq 0$ for all θ_i and $i \neq 0$ [as in Section 23.D, $\bar{y}_i(\theta_i)$ and $\bar{t}_i(\theta_i)$ are agent i's probability of getting the good and expected transfer conditional on announcing his type to be θ_i] because buyers are always free not to participate. The seller's optimal choice is therefore a particular element of the set of interim incentive efficient social choice functions.

The seller's problem can be written as one of choosing functions $y_1(\cdot), \ldots, y_I(\cdot)$ and $U_1(\cdot), \ldots, U_I(\cdot)$ to solve

$$\underset{\{y_i(\cdot), U_i(\cdot)\}_{i=1}^I}{\text{Max}} \quad \sum_{i \neq 0} \int_{\underline{\theta}_i}^{\bar{\theta}_i} [\bar{y}_i(\theta_i)\theta_i - U_i(\theta_i)]\phi_i(\theta_i)\, d\theta_i \qquad (23.F.7)$$

s.t. (i) $\bar{y}_i(\cdot)$ is nondecreasing for all $i \neq 0$.

(ii) For all θ: $y_i(\theta) \in [0, 1]$ for all $i \neq 0$, $\sum_{i \neq 0} y_i(\theta) \leq 1$.

(iii) $U_i(\theta_i) = U_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta_i} \bar{y}_i(s)\, ds$ for all $i \neq 0$ and θ_i.

(iv) $U_i(\theta_i) \geq 0$ for all $i \neq 0$ and θ_i.

We note first that if constraint (iii) is satisfied then constraint (iv) will be satisfied if and only if $U_i(\underline{\theta}_i) \geq 0$ for all $i \neq 0$. As a result, we can replace constraint (iv) with

(iv′) $U_i(\underline{\theta}_i) \geq 0$ for all $i \neq 0$ and θ_i.

Next, substituting into the objective function for $U_i(\theta_i)$ using constraint (iii), and following the same steps that led to (23.D.16), the seller's problem can be written as one of choosing the $y_i(\cdot)$ functions and the values $U_1(\underline{\theta}_1), \ldots, U_I(\underline{\theta}_I)$ to maximize

$$\int_{\underline{\theta}_1}^{\bar{\theta}_1} \cdots \int_{\underline{\theta}_I}^{\bar{\theta}_I} \left[\sum_{i=1}^{I} y_i(\theta_1, \ldots, \theta_I)\left(\theta_i - \frac{1 - \Phi_i(\theta_i)}{\phi_i(\theta_i)}\right)\right]\left[\prod_{i=1}^{I} \phi_i(\theta_i)\right] d\theta_I \cdots d\theta_1 - \sum_{i=1}^{I} U_i(\underline{\theta}_i)$$

subject to constraints (i), (ii), and (iv′). It is evident that the solution must have $U_i(\underline{\theta}_i) = 0$ for all $i = 1, \ldots, I$. Hence, the seller's problem reduces to choosing functions $y_1(\cdot), \ldots, y_I(\cdot)$ to maximize

$$\int_{\underline{\theta}_1}^{\bar{\theta}_1} \cdots \int_{\underline{\theta}_I}^{\bar{\theta}_I} \left[\sum_{i=1}^{I} y_i(\theta_1, \ldots, \theta_I)\left(\theta_i - \frac{1 - \Phi_i(\theta_i)}{\phi_i(\theta_i)}\right)\right]\left[\prod_{i=1}^{I} \phi_i(\theta_i)\right] d\theta_I \cdots d\theta_1 \qquad (23.F.8)$$

subject to constraints (i) and (ii).

Let us ignore constraint (i) for the moment. Define

$$J_i(\theta_i) = \theta_i - \frac{1 - \Phi_i(\theta_i)}{\phi_i(\theta_i)}.$$

Then inspection of (23.F.8) indicates that $y_1(\cdot), \ldots, y_I(\cdot)$ is a solution to this relaxed problem if and only if for all $i = 1, \ldots, I$ we have

$$y_i(\theta) = 1 \qquad \text{if } J_i(\theta_i) > \text{Max}\,\{0, \text{Max}_{h \neq i}\, J_h(\theta_h)\}$$

and

$$y_i(\theta) = 0 \qquad \text{if } J_i(\theta_i) < \text{Max}\,\{0, \text{Max}_{h \neq i}\, J_h(\theta_h)\}. \tag{23.F.9}$$

[Note that $J_i(\theta_i) = \text{Max}\,\{0, \text{Max}_{h \neq i}\, J_h(\theta_h)\}$ is a zero probability event.] But, given our assumption that $J_i(\cdot)$ is nondecreasing in θ_i, (23.F.9) implies that $y_i(\cdot)$ is nondecreasing in θ_i, which in turn implies that $\bar{y}_i(\cdot)$ is nondecreasing. Thus the solution to this relaxed problem actually satisfies constraint (i), and so is a solution to the seller's overall problem (see Section M.K of the Mathematical Appendix). The optimal transfer functions can then be set as $t_i(\theta) = U_i(\theta_i) - \theta_i\, \bar{y}_i(\theta_i)$, where $U_i(\theta_i)$ is calculated from constraint (iii).

A few things should be noted about (23.F.9). First, observe that when the various agents have differing distribution functions $\Phi_i(\cdot)$, the agent i who has the largest value of $J_i(\theta_i)$ is *not* necessarily the same as the agent who has the highest valuation for the object. Thus, the seller's optimal auction need not be ex post (classically) efficient.

Second, in the case of *symmetric* bidders in which $\underline{\theta}_i = \underline{\theta}$ and $J_i(\cdot) = J(\cdot)$ for $i = 1, \ldots, I$, when $\underline{\theta} > 0$ is large enough so that $J(\underline{\theta}) > 0$, the optimal auction always gives the object to the bidder with the highest valuation and also leaves each bidder with an expected utility of zero when his valuation attains its lowest possible value. We can therefore conclude, using the revenue equivalence theorem (Proposition 23.D.3), that the first-price and second-price sealed-bid auctions are both optimal in this case.

Third, the optimal auction has a nice interpretation in terms of monopoly pricing. Consider, for example, the case in which $I = 1$ and $\underline{\theta}_1 = 0$. Then conditions (23.F.9) tell us that the optimal auction gives the object to the buyer (agent 1) if and only if $J_1(\theta_1) = [\theta_1 - ((1 - \Phi_1(\theta_1))/\phi_1(\theta_1))] > 0$. Suppose we think instead of the seller in this circumstance simply naming a price p and letting the buyer then decide whether to buy at this price. The seller's expected revenue from this scheme is $p(1 - \Phi_1(p))$, and so the first-order condition for his optimal posted price, say p^*, is $(1 - \Phi_1(p^*)) - p^*\phi_1(p^*) = 0$, or equivalently, $J_1(p^*) = 0$. Since $J_1(\cdot)$ is nondecreasing, we see that with this optimal posted price policy a buyer of type θ_1 gets the good with probability 1 if $J_1(\theta_1) > 0$, and with probability 0 if $J_1(\theta_1) < 0$, exactly as in the optimal auction derived above. Indeed, given the revenue equivalence theorem we can conclude that in this case this simple posted price scheme is an optimal mechanism for the seller. [For more on the monopoly interpretation of optimal auctions, see Exercise 23.F.5 and Bulow and Roberts (1989).] ■

5. Miscellany (MWG 387-400, 404-417, 423-427; covered only if time permits)

Theorem (Proposition 12.C.1 at MWG 388): Bertrand duopoly with constant returns to scale, perfectly substitutable goods: Simultaneous price choices by firms yields competitive outcome as unique equilibrium.

Theorem (Proposition 12.C.2 and Example 12.C.1 at MWG 390-392): Cournot duopoly with constant returns to scale, perfectly substitutable goods: Simultaneous quantity choices by firms yields equilibrium (not necessarily unique) with prices between competitive and monopoly prices.

Cournot outcome approaches competitive outcome as number of firms grows (MWG 391-394).

Capacity constraints and product differentiation (MWG 394-400).

Kreps-Scheinkman and importance of timing of irreversible decisions.

Simultaneous entry followed by Bertrand or Cournot competition (MWG 405-411).

Strategic precommitments, strategic complements and substitutes, direct and indirect effects of investment decisions (MWG 414-417).

Dixit and Spence entry deterrence and accommodation models (MWG 423-427).