4100/6001 Advanced Economic Analysis Part 2 Notes

Richard Holden
UNSW and NBER
richard.holden@unsw.edu.au

May 21, 2013

Contents

1 Adverse Selection (Hidden Information)

1.1 Asymmetric Information and Market Breakdown

1.1.1 Overview

- An embedded assumption of the First Welfare Theorem is that all participants have symmetric information
- Obvious not true in reality
  - Workers know their own ability better than employers
  - Car owners know the quality better than prospective buyers
  - Insurees know their risk better than insurers
  - ...
- More important than the realism of the assumptions is the implication of it–fundamentally changes the analysis and conclusions
- Witht asymmetric information the complete markets assumption of the FWT fails to hold
- No longer an Arrow-Debreu security for each state of the world
- Competitive markets are not Pareto efficient!
- Raises a whole host of questions
  - How severe is the market breakdown?
  - How do we characterize the set of equilibria in such an environment?
  - Can interventions lead to a Pareto improvement?
- Core idea of adverse selection due to Akerlof (QJE, 1970)–used the metaphor of used cars: "lemons"

1.1.2 Basic setup

- Labor market setting
- Many identical (risk neutral) firms that can employ workers
- CRS production technology with labor the only input
- Normalize price of output to 1
- Workers differ in their quality–the number of units of output they produce if hired
- Represent this by the scalar θ
- Routinely refer to this as the agent's *type* in hidden information models
- Let $\theta \in [\underline{\theta}, \overline{\theta}] \subset \mathbb{R}$ and assume that $0 \leq \underline{\theta} < \overline{\theta} < \infty$.
- Let CDF be $F(.)$

- Each worker's outside option is $r(\theta)$
- What would happen with symmetric information (i.e. what is the *first-best*)?
  - There will be a different equilm wage for each type, $w^*(\theta)$
  - Given perfect competition and CRS it must be that $w^*(\theta) = \theta$ for all types
  - This is clearly Pareto optimal (i.e. a necessary and sufficient condition being that it maximizes total surplus)
  - To see this, not that TS is $\int_{\underline{\theta}}^{\overline{\theta}} N[I(\theta)\theta + (1 - I(\theta))r(\theta)]dF(\theta)$, where $I(\theta)$ is an indicator function for type θ, taking value 1 if she works for an employer
- Now consider the *second best*
- There will have to be a single wage w for all types due to the non-obervability
- Supply side: type θ only willing to work if $w \geq r(\theta)$
- Set of types who work denoted $\Theta(w) = \{\theta : w \geq r(\theta)\}$
- Demand side: if an employer believes that the average productivity of workers accepting is μ then demand is $z(w) = 0$ if $\mu < w$, $z(w) \in [0, \infty]$ if $\mu = w$, $z(w) = \infty$ if $\mu > w$
- In a *rational expectations equilibrium* beliefs must be correct so it must be that $\mu = E[\theta|\theta \in \Theta^*]$ and labor demand can only equal labor supply at a positive level if $w = E[\theta|\theta \in \Theta^*]$

1.1.3 Pareto inefficiency

- Suppose $r(\theta) = r$ for all θ and $F(r) \in (0, 1)$
- Pareto optimality requires all types $\theta \geq r$ getting employed and all types $\theta < r$ exercising their outside option
- In the competitive equilm either all workers accept or non do
- So $E[\theta|\theta \in \Theta^*] = E[theta]$ for all w and hence $w^* = E[\theta]$
- If $E[\theta] \geq r$ then all workers accept, if not then none do
- Which one occurs depends on the relative proportion of good and bad types
- Inefficiency comes from not being able to distinguish between different types

1.1.4 Unravelling

- Suppose $r(\theta)$ is not constant
- Now the market can completely break down
- Suppose that $R(\theta) \leq \theta$ for all θ and that $r(.)$ is strictly increasing
- Expected value of worker productivity depends on the wage

- Higher wage means more workers willing to accept so the average productivity (i.e. quality of the pool) increases
- Technically important to assume that F has an associated PDF f with full support (i.e. $f(\theta) > 0$ for all θ)–so $E[\theta|r(\theta) \leq w]$ varies continuously with w
- Putting this all together it must be that $w^* = E[\theta|r(\theta) \leq w*]$
- Market equilm need not be efficient
- To get high types to accept the wage must be high
- Start at the top, $\theta = \overline{\theta}$
- But lower productivity workers will participate, so $E[\theta] < r(\overline{\theta})$
- So wage has to be less than $r(\overline{\theta})$ for employer to break even
- But then highest type opt out
- And then the wage has to be lower still (in REE)
- And then...
- How far can the unravelling go?
- *All the way!*
- For example: Suppose $r(\theta) = \alpha\theta$, for $\alpha < 1$ and let $\theta \subset U[0, 2]$. So $r(\underline{\theta}) = \underline{\theta}$, and $r(\theta) < 0$ for $\theta > 0$. So we have $E[\theta|r(\theta) \leq w] = w/2\theta$. For $\theta > 1/2, E[\theta|r(\theta \leq 0] = 0$ and $E[\theta|r(\theta \leq w] < w$ for all $w > 0$.
- There can be multiple equila
- Basic idea is that the slope of the conditional expectation function depends entirely on the density
- But equilibria can be Pareto ranked
- Firms always earn zero profits and workers are better off with higher equilm wages
- Low-wage Pareto inferior arise because of a coordination failure
- Wage is low because employers expect that the workers who accept will be low quality and this is self-reinforcing

1.1.5 Game Theoretic Approach

- What if employers *could* change the wage offer, but *choose not to* in equilm
- "RH: Just because it doesn't happen in equilm doesn't mean it's not important"

1.1.6 Market Intervention

2 Screening

2.1 Static Screening

2.1.1 Introduction

- A good reference for further reading is Fudenberg & Tirole chapter 7
- Different to "normal" Adverse Selection because 1 on 1, not a market setting
- 2 players: Principal and the Agent
- Payoff: Agent $G\left(u\left(q,\theta\right)-T\right)$, Principal $H\left(v\left(q,\theta\right)+T\right)$ where $G\left(\cdot\right),H\left(\cdot\right)$ are concave functions and q is some verifiable outcome (eg. output), T is a transfer, θ is the Agent's private information
- Don't use the concave transforms for now
- Say Principal is a monopolistic seller and the Agent is a consumer
- Let $v(q,\theta)=-cq$
- Principal's payoff is $T-cq$ where T is total payment (pq)
- $u(q,\theta)=\theta V(q)$
- Agent's payoff is $\theta V(q)-T$ where $V(\cdot)$ is strictly concave
- θ is type (higher $\theta\rightarrow$ more benefit from consumption)
- $\theta=\theta_1,...,\theta_n$ with probabilities $p_1,...,p_n$
- Principal only knows the distribution of types
- Note: relationship to non-linear pricing literature
- Assume that the Principal has all the bargaining power
- Start by looking at the first-best outcome (ie. under symmetric information)

First Best Case I: Ex ante no-one knows θ, ex post θ is verifiable

- Principal solves

$$\max_{(q_i,T_i)_{i=1}^{n}} p_i(T_i-cq_i)$$
$$s.t.{}_{i=1}^{n}p_i(\theta_iV(q_i)-T_i)\geq\overline{U} \qquad \text{(PC)}$$

First Best Case II: Ex ante both know θ

- Normalize $\overline{U}$ to 0
- Principal solves

$$\max_{(q_i,T_i)} \{T_i - cq_i\}$$
$$s.t. \theta_i V(q_i) - T \geq 0 \tag{PC}$$

- The PC will bind, so $T_i = \theta_i V(q_i)$
- So they just solve $\max_{q_i} \{\theta_i V(q_i) - cq_i\}$
- FOC $\theta_i V'(q_i) = c$
- This is just perfect price discrimination – efficient but the consumer does badly
- Case I folds into II by offering a contingent contract

Second-Best

- Agent knows θ_i but the Principal doesn't
- First ask if we can achieve/sustain the first best outcome
- ie. will they naturally reveal their type
- say the type is θ_2
- if they reveal themselves their payoff is $\theta_2 V(q_2^*) - T_2^* = 0$
- if they pretend to be θ_1 their payoff is $\theta_2 V(q_2^*) - T_1^* = \theta_2 V(q_1^*) - \theta_1 V(q_1^*) = (\theta_2 - \theta_1)V(q_1^*) > 0$ since $\theta_2 > \theta_1$
- can't get the first-best

Second-best with n types

- First to really look at this was Mirrlees in his 1971 optimal income tax paper – normative
- Positive work by Akerlof, Spence, Stiglitz
- Revelation Principle very useful: can look at / restrict attention to contracts where people reveal their true type *in equilibrium*

- Without the revelation principle we would have the following problem for the principal

$$\max_{T(q)} \{{}_{i=1}^{n} p_i(T(q_i) - cq_i)\}$$
$$\text{subject to}$$
$$\theta_i V(q_i) - T(q_i)) \geq 0, \forall i \quad \text{(PC)}$$
$$q_i = \arg\max_{q} \{\theta_i V(q) - T(q))\}, \forall i \quad \text{(IC)}$$

- But the revelation principle means that there is no loss of generality in restricting attention to optimal equilibrium choices by the buyers
- We can thus write the Principal's Problem as

$$\max_{(q_i, T_i)} \{{}_{i=1}^{n} p_i(T_i - cq_i)\}$$
$$\text{subject to}$$
$$\theta_i V(q_i) - T_i) \geq 0, \forall i \quad \text{(PC)}$$
$$\theta_i V(q_i) - T_i \geq \theta_i V(q_j) - T_j, \forall i, j \quad \text{(IC)}$$

- Incentive compatibility means the Agent truthfully reveals herself
- This helps a lot because searching over a schedule $T(q)$ is hard
- Before proceeding with the n types case return to a two type situation

Second-best with* 2 *types

- Too many constraints to be tractable (there are $n(n-1)$ constraints of who could pretend to be whom)
- 2 types with $\theta_H > \theta_L$
- Problem is the following:

$$\max \{p_H(T_H - cq_H) + p_L(T_L - cq_L)\}$$
$$s.t.(i)\ \theta_H V(q_H) - T_H \geq \theta_H V(q_L) - T_L \quad \text{(IC)}$$
$$(ii)\ \theta_L V(q_L) - T_L \geq 0 \quad \text{(PC)}$$

- We have eliminated two constraints: the IC constraint for the low type and the PC constraint for the high type
- Why was this ok?
- The low type constraint must be the only binding PC (high types can "hide behind" low types)
- And the low type won't pretend to be the high type

- PC must bind otherwise we could raise T_L and the Principal will always be happy to do that
- IC must always bind otherwise the Principal could raise T_H (without equality the high type's PC would not bind) – also good for the Principal
- So $\theta_H V(q_H) - T_H = \theta_H V(q_L) - T_L$ and $\theta_L V(q_L) - T_L = 0$
- Now substitute to get an unconstrained problem:

$$\max_{q_L, q_H} \{ p_H (\theta_H V(q_H) - \theta_H V(q_L) + \theta_L V(q_L) - c q_H) + p_L (\theta_L V(q_L) - c q_L) \}$$

- The FOCs are

$$p_H \theta_H V'(q_H) - p_H c = 0$$

and

$$p_L \theta_L V'(q_L) - p_L c + p_H \theta_L V'(q_L) - p_H \theta_H V'(q_L) = 0$$

- The first of these simplifies to $\theta_H V'(q_H) = c$ (so the high type chooses the socially efficient amount)
- The second of these simplifies to the following:

$$\begin{aligned} \theta_L V'(q_L) &= \frac{c}{1 - \frac{1-p_L}{p_L} \frac{\theta_H - \theta_L}{\theta_L}} \\ &> c \end{aligned}$$

(so the low type chooses too little)

- $q_H = q_H^*$ and $q_L < q_L^*$
- No incentive reason for distorting q_H because the low type isn't pretending to be the high type
- But you do want to discourage the high type from pretending to be the low type – and hence you distort q_L
- We can check the IC constraint is satisfied for the low type

$$\theta_H V(q_H) - T_H = \theta_H V(q_L) - T_L \text{ (high type's IC is binding)}$$

now recall that (recalling that $\theta_H > \theta_L, q_H > q_L$), so we have

$$\theta_L V(q_L) - T_L \geq \theta_L V(q_H) - T_H$$

- So the low type's IC is satisfied
- High type earns rents – PC does not bind
- Lots of applications: optimal taxation, banking, credit rationing, implicit labor contracts, insurance, regulation (see Bolton-Dewatripont for exposition)

2.1.2 Optimal Income Tax

- Mirrlees (Restud, 1971)
- Production function $q = \mu e$ (for each individual), where q is output, μ is ability and e is effort
- Individual knows μ and e but society does not
- Distribution of μs in the population, μ_L and μ_H in proportions π and $1-\pi$ respectively
- Utility function $U(q - T - \psi(e))$ where T is tax (subsidy if negative) and $\psi(e)$ is cost of effort (presumably increasing and convex)
- The government's budget constraint is $\pi T_L + (1-\pi)T_H \geq 0$
- Veil of Ignorance – rules are set up before the individuals know their type
- So the first-best problem is:

$$\max_{e_L, e_H, T_L, T_H} \{\pi U(\mu_L e_L - T_L - \psi(e_L)) + (1-\pi)U(\mu_H e_H - T_H - \psi(e_H))\}$$
$$\text{subject to}$$
$$\pi T_L + (1-\pi)T_H \geq 0$$

- But the budget constraint obviously binds and hence $\pi T_L + (1-\pi)T_H = 0$
- Then we have $T_H = -\pi T_L / (1-\pi)$
- The maximization problem can be rewritten as

$$\max_{e_L, e_H, T_L} \{\pi U(\mu_L e_L - T_L - \psi(e_L)) + (1-\pi)U(\mu_H e_H + (\pi T_L / 1 - \pi) - \psi(e_H))\}$$

- The FOCs are

  (i) $-U'(\mu_L e_L - T_L - \psi(e_L)) = U'(\mu_H e_H + (\pi T_L / 1 - \pi) - \psi(e_H))$

  (ii) $\mu_L = \psi'(e_L)$

  (iii) $\mu_H = \psi'(e_H)$

- Choose e_L, e_H efficiently in the first-best
- Everyone has same marginal cost of effort so the higher marginal product types work harder
- (i) just says the marginal utilities are equated
- Hence $\mu_L e_L - T_L - \psi(e_L) = \mu_H e_H + T_H - \psi(e_H)$
- The net payoffs are identical so you are indifferent between which type you are

- Consistent with Veil of Ignorance setup
- There is no DWL because of the lump sum aspect of the transfer

Second-Best

- Could we sustain the first-best?
- No because the high type will pretend to be the low type, $\mu_H e = q_L$ so $q_L - T_L - \psi(q_L/\mu_H) > q_L - T_L - \psi(e_L)$ since $q_L/\mu_H < e_L$
- Basically the high type can afford to slack because they are more productive - hence no self sustaining first-best
- The Second-Best problem is

$$\max_{e_L,e_H,T_L,T_H} \{\pi U(\mu_L e_L - T_L - \psi(e_L)) + (1-\pi)U(\mu_H e_H - T_H - \psi(e_H))\}$$
$$s.t.(i)\mu_H e_H - T_H - \psi(e_H) \geq \mu_L e_L - T_L - \psi(\mu_L e_L/\mu_H)$$
$$(ii)\pi T_L + (1-\pi)T_H \geq 0$$

- Solving yields $e_H = e_H^*$
- and $\mu_L = \psi'(e_L) + \beta(1-\pi)(\mu_L - \mu_L/\mu_H \psi'(\mu_L e_L/\mu_H))$
- where $\beta = \frac{U_L' - U_H'}{U_L'}$ (marginal utilities evaluated at their consumptions levels)
- but $U_L < U_H$ so $U_L' > U_H'$ (by concavity) and hence $0 < \beta < 1$
- Since $\psi(\cdot)$ is convex we have $\psi'\left(\frac{\mu_L e_L}{\mu_H}\right) < \psi'(e_L)$
- $\mu_L > \psi'(e_L) + \beta(1-\pi)(\mu_L - \mu_L/\mu_H \psi'(e_L))$
- and hence:
$$\psi'(e_L) < \frac{\mu_L - \beta(1-\pi)\mu_L}{1 - \beta(1-\pi)\mu_L/\mu_H} < \mu_L$$
- (the low type works too little)
- To stop the high type from misrepresenting themselves we have to lower the low type's required effort and therefore subsidy
- High type is better off $\rightarrow$ lose the egalitarianism we had before for incentive reasons
- Can offer a menu (q_L, T_L), (q_H, T_H) and people self select
- If you have a continuum of types there would be a tax schedule $T(q)$
- Marginal tax rate of the high type is zero (because they work efficiently) so $T'(q) = 0$ at the very top and $T'(q) > 0$ elsewhere with a continuum of types

2.1.3 Regulation

- Baron & Myerson (Ecta, 1982)
- The regulator/government is ignorant but the firm knows its type
- Firm's characteristic is $\beta \in \{\underline{\beta}, \overline{\beta}\}$ with probabilities ν_1 and $1 - \nu_1$
- Cost is $c = \beta - e$
- Cost is verifiable
- Cost of effort is $\psi(e) = e^2/2$
- Let $\Delta\beta = \overline{\beta} - \underline{\beta}$ and assume $\Delta\beta < 1$
- Government wants a good produced with the lowest possible subsidy - wants to minimize expected payments to the firm
- The First-Best is simply
$$\min_e \{\beta - e + e^2/2\}$$
- The FOC is $e^* = 1$ and the firm gets paid $\beta - 1/2$
- Can we sustain the FB?
- No because $p_L = \beta_L - 1/2$ and $p_H = \beta_H - 1/2$

Second-Best

- Two cost levels $\underline{c}$ and $\overline{c}$
- Two price levels $\underline{p}$ and $\overline{p}$ (payments)
- Government solves

$$\begin{gathered}\min\{\nu_1 \underline{p} + (1-\nu_1)\overline{p}\} \\ s.t.(i)\ \underline{p} - \underline{c} - e^2/2 \geq \overline{p} - \overline{c} - (\overline{e} - \Delta\beta)^2/2 \\ (ii)\ \overline{p} - \overline{c} - \overline{e}^2/2 \geq 0\end{gathered}$$

- noting that $\underline{e} = \overline{e} - \Delta\beta$ (from cost equation and low type pretending to be high type)
- Define $\underline{s} = \underline{p} - \underline{c} = \underline{p} - \underline{\beta} + \underline{e}$ and $\overline{s} = \overline{p} - \overline{c} = \overline{p} - \overline{\beta} + \overline{e}$ (these are the "subsidies")
- The government's problem is now

$$\begin{gathered}\min\{\nu_1(\underline{s} + \underline{\beta} - \underline{e}) + (1-\nu_1)\overline{s} + \overline{\beta} - \overline{e}\} \\ s.t.(i)\ \underline{s} - e^2/2 \geq \overline{s} - (\overline{e} - \Delta\beta)^2/2 \\ (ii)\ \overline{s} - \overline{e}^2/2 \geq 0\end{gathered}$$

- Since the constraints must hold with equality we can substitute and write this as an unconstrained problem

$$\min_{\underline{e},\overline{e}} \left\{ \nu_1 \left(\frac{\overline{e}^2}{2} + \underline{e}^2/2 - \frac{(\overline{e} - \Delta\beta)^2}{2} \right) + (1-\nu_1) \left(\frac{\overline{e}^2}{2} - \overline{e} \right) \right\}$$

- The FOCs are

(1) $\underline{e} = 1$

(2) $\nu_1 \overline{e} - \nu_1 (\overline{e} - \Delta\beta) + (1-\nu_1)\overline{e} - (1-\nu_1) = 0$

- (2) implies that:

$$\overline{e} = \frac{1 - \nu_1 - \nu_1 \Delta\beta}{1-\nu_1} = 1 - \frac{\nu_1 \Delta\beta}{1-\nu_1}$$

- The low cost ("efficient") type chooses $\underline{e} = 1$
- The high cost ("bad") types chooses $\overline{e} = 1 - \frac{\nu_1 \Delta\beta}{1-\nu_1}$
- Offer a menu of contracts: fixed price or a cost-sharing arrangement
- The low cost firm takes the fixed price contract, becomes the residual claimant and then chooses the efficient amount of effort
- See also Laffont & Tirole (JPE, 1986) – costs observable

2.1.4 The General Case – n types and a continnum of types

- Problem of all the incentive compatibility constraints
- It turns out that we can replace the IC constraints with downward adjacent types
- The constraints are then just:

(i) $\theta_i V(q_i) - T_i \geq \theta_i V(q_{i-1}) - T_{i-1} \quad \forall i = 2, ..., n$

(ii) $q_i \geq q_{i-1} \quad \forall i = 2, ..., n$

(iii) $\theta V(q_1) - T_1 \geq 0$

- (ii) is a monotonicity condition
- It is mathematically convenient to work with a continuum of types – and we will
- Let $F(\theta)$ be a cdf and $f(\theta)$ the associated density function on the support $[\underline{\theta}, \overline{\theta}]$
- The menu being offered is $T(\theta), q(\theta)$

- The problem is

$$\max_{T(\cdot),q(\cdot)} \left\{ \int_{\underline{\theta}}^{\overline{\theta}} [T(\theta) - cq(\theta)] f(\theta) d\theta \right\}$$

$$s.t.(i)\ \theta V(q(\theta) - T(\theta) \geq \theta V\left(q\left(\widehat{\theta}\right)\right) - T\left(\widehat{\theta}\right) \forall \theta, \widehat{\theta} \quad \text{(IC)}$$

$$(ii)\ \theta V(q(\theta) - T(\theta) \geq 0, \forall \theta \quad \text{(PC)}$$

- We will be able to replace all the IC constraints with a Local Adjacency condition and a Monotonoicity condition

Definition 1. *An allocation $T(\theta), q(\theta)$ is implementable if and only if it satisfies IC $\forall \theta, \widehat{\theta}$*

Proposition 1. *An allocation $T(\theta), q(\theta)$ is implementable if and only if $\theta V'(q(\theta)) \frac{dq(\theta)}{d\theta} - T'(\theta) = 0$ (the local adjacency condition) and $\frac{dq(\theta)}{d\theta} \geq 0$ (the monotonicity condition).*

Proof. $\Rightarrow$ direction:

Let $\widehat{\theta} = \arg\max_{\theta} \left\{ \theta V\left(q\left(\widehat{\theta}\right)\right) - T\left(\widehat{\theta}\right) \right\}$. Now $\frac{d}{d\widehat{\theta}} = \theta V'\left(q\left(\widehat{\theta}\right)\right) - \frac{dq(\widehat{\theta})}{d\widehat{\theta}} - T'\left(\widehat{\theta}\right)$

so $\theta V'(q(\theta)) - \frac{dq(\theta)}{d\theta} - T'(\theta) = 0, \forall \theta$

Now, by revealed preference:

$$\theta V(q(\theta)) - T(\theta) \geq \theta V(q(\theta')) - T(\theta')$$

and

$$\theta' V(q(\theta')) - T(\theta') \geq \theta' V(q(\theta)) - T(\theta)$$

combining these yields:

$$\theta [V(q(\theta)) - V(q(\theta'))] \geq T(\theta) - T'(\theta) \geq \theta' [V(q(\theta)) - V(q(\theta'))]$$

the far RHS can be expressed as $(\theta - \theta')(V(q(\theta)) - V(q(\theta'))) \geq 0$

hence if $\theta > \theta'$ then $q(\theta) \geq q(\theta')$ □

- This really just stems from the **Single-Crossing Property** (or **Spence-Mirrlees Condition**), namely $\frac{\partial U}{\partial q}$ is increasing in θ
- Note that this is satisfied with the separable functional form we have been using–but need not be satisfied in general
- Higher types are "even more prepared" to buy some increment than a lower type

Proof. $\Leftarrow$ direction

Let $W\left(\theta, \widehat{\theta}\right) = \theta V\left(q\left(\widehat{\theta}\right)\right) - T\left(\widehat{\theta}\right)$. Fix θ and suppose the contrary. This implies that $\exists \widehat{\theta}$ such that $W\left(\theta, \widehat{\theta}\right) > W(\theta, \theta)$.

Case 1: $\widehat{\theta} > \theta$

$$W\left(\theta,\widehat{\theta}\right)-W\left(\theta,\theta\right)=\int_{\theta}^{\widehat{\theta}}\frac{\partial W}{\partial\tau}\left(\theta,\tau\right)d\tau=\int_{\theta}^{\widehat{\theta}}\theta V'\left(q\left(\tau\right)\right)\frac{dq}{d\tau}-T'\left(\tau\right)d\tau$$

But $\tau>\theta$ implies that:

$$\begin{aligned}&\int_{\theta}^{\widehat{\theta}}\theta V'\left(q\left(\tau\right)\right)\frac{dq}{d\tau}-T'\left(\tau\right)d\tau\\ \leq\ &\int_{\theta}^{\widehat{\theta}}\left(\tau V'\left(q\left(\tau\right)\right)\frac{dq}{d\tau}-T'\left(\tau\right)\right)d\tau=0\end{aligned}$$

because the integrand is zero. Contradiction. Case 2 is analogous. □

- This proves that the IC constraints are satisfied globally, not just the SOCs (the common error)
- Now we write the problem as:

$$\max_{T(\cdot),q(\cdot)}\left\{\int_{\underline{\theta}}^{\overline{\theta}}\left[T\left(\theta\right)-cq\left(\theta\right)\right]f(\theta)d\theta\right\}$$

$$s.t.(i)\ \theta V'(q\left(\theta\right))\frac{dq\left(\theta\right)}{d\theta}-T'\left(\theta\right)\geq 0\quad\forall\theta \qquad \text{(Local Adjacency)}$$

$$(ii)\ \frac{dq\left(\theta\right)}{d\theta}\geq 0\quad\forall\theta \qquad \text{(Monotonicity)}$$

$$(iii)\underline{\theta}V\left(q\left(\underline{\theta}\right)\right)-T\left(\underline{\theta}\right)=0 \qquad \text{(PC-L)}$$

- Let $W\left(\theta\right)\equiv W\left(\theta,\theta\right)=\theta V\left(q\left(\theta\right)\right)-T\left(\theta\right)=\max_{\hat{\theta}}\left\{\theta V(q(\hat{\theta}))-T\left(\hat{\theta}\right)\right\}$
- Recall that in the 2 type case we used the PC for the lowest type and the IC for the other type
- We could have kept on going for higher and higher types
- Now, from the FOCs:

$$\frac{dW\left(\theta\right)}{d\theta}=\theta V'\left(q\left(\theta\right)\right)\frac{dq}{d\theta}-\frac{dT}{d\theta}+V\left(q\left(\theta\right)\right)=V\left(q\left(\theta\right)\right)$$

  (by adding $V(q(\theta))$ to both sides)

$$W\left(\theta\right)-W\left(\underline{\theta}\right)=\int_{\underline{\theta}}^{\theta}\frac{dW\left(\tau\right)}{d\tau}d\tau=\int_{\underline{\theta}}^{\theta}V\left(q\left(\tau\right)\right)d\tau$$

(change of measure trick)

- But $W\left(\underline{\theta}\right)=0$ (PC of low type binding at the optimum)
- Now $T\left(\theta\right)=-\int_{\underline{\theta}}^{\theta}V\left(q\left(\tau\right)\right)d\tau+\theta V\left(q\left(\theta\right)\right)$ (by substitution)

- So the problem is now just

$$\max_{q(\cdot)}\left\{\int_{\underline{\theta}}^{\overline{\theta}}\left[\theta V\left(q\left(\theta\right)\right)-\int_{\underline{\theta}}^{\theta}V\left(q\left(\tau\right)\right)d\tau-cq\left(\theta\right)\right]f\left(\theta\right)d\theta\right\}$$

$$s.t.\frac{dq\left(\theta\right)}{d\theta}\geq 0\quad\forall\theta$$

- Proceed by ignoring the constraint for the moment and tackle the double integral using integration by parts
- Recall that

$$\int_{\underline{\theta}}^{\overline{\theta}}uv'=uv\Big|_{\underline{\theta}}^{\overline{\theta}}-\int_{\underline{\theta}}^{\overline{\theta}}u'v$$

- So let $v'=f\left(\theta\right)$ and $u=\int V\left(q\left(\tau\right)\right)d\tau$, and we then have

$$\begin{aligned}\int_{\underline{\theta}}^{\overline{\theta}}\left[\int_{\underline{\theta}}^{\theta}V\left(q\left(\tau\right)\right)d\tau\right]f\left(\theta\right)d\theta &= \int_{\underline{\theta}}^{\theta}V\left(q\left(\tau\right)\right)d\tau F\left(\theta\right)\Big|_{\underline{\theta}}^{\overline{\theta}}-\int_{\underline{\theta}}^{\overline{\theta}}V\left(q\left(\theta\right)\right)F\left(\theta\right)d\theta\\ &= \int_{\underline{\theta}}^{\overline{\theta}}V\left(q\left(\tau\right)\right)d\tau-\int_{\underline{\theta}}^{\overline{\theta}}V\left(q\left(\theta\right)\right)F\left(\theta\right)d\theta\\ &= \int_{\underline{\theta}}^{\overline{\theta}}V\left(q\left(\theta\right)\right)\left[1-F\left(\theta\right)\right]d\theta\end{aligned}$$

- So we can write the problem as:

$$\max_{q(\cdot)}\left\{\int_{\underline{\theta}}^{\overline{\theta}}\left(\left(\theta V(q\left(\theta\right)-cq\left(\theta\right)\right)f\left(\theta\right)-V(q\left(\theta\right)\left[1-F\left(\theta\right)\right]\right)d\theta\right\}$$

- Now we can just do pointwise maximization (maximize under the integral for all values of θ)

$$\theta V'(q\left(\theta\right))=V'(q\left(\theta\right))\left(\frac{1-F\left(\theta\right)}{f\left(\theta\right)}\right)+c,\quad\forall\theta \tag{1}$$

- From 1 we can say the following:
(1)

$$\theta=\overline{\theta}\rightarrow\overline{\theta}V\left(q\left(\overline{\theta}\right)\right)=c$$

(2)

$$\theta<\overline{\theta}\rightarrow\overline{\theta}V\left(q\left(\overline{\theta}\right)\right)>c$$

($q\left(\theta\right)$ is too low)

- Since efficiency requires $\theta V'(q\left(\theta\right))=c$
- Now differentiate (1) and solve for $\frac{dq}{d\theta}\geq 0$

- This implies that $\frac{f(\theta)}{1-F(\theta)}$ is increasing in θ (*this is a sufficient condition in general, but is a necessary and sufficient condition in this buyer-seller problem*)
- This property is known as the **Monotone Hazard Rate Property**
- It is satisfied for all log-concave distributions
- We've been considering the circumstance where θ announces their type, θ^a and gets a quantity $q(\theta^a)$ and pays a tariff of $T(\theta^a)$
- This can be reinterpreted as: given $\widehat{T}(q)$, pick q
- For each q there can only be one $T(q)$ by incentive compatibility
- $\widehat{T}(q) = T(\theta^{-1}(q))$
- The optimization problem becomes

$$\max_q \left\{\theta V(q) - \widehat{T}(q)\right\}$$

- The FOC is $\theta V'(q) = \widehat{T}'(q) \equiv p(q)$

$$p(q) = \frac{p(q(\theta))}{\theta}\left(\frac{1-F(\theta)}{f(\theta)}\right) + c$$

$$\frac{p-c}{p} = \frac{1-F(\theta)}{\theta f(\theta)}$$

- Recall that we ignored the constraint $\frac{dq}{d\theta} \geq 0$
- The FOC implies

$$\left(\theta - \frac{1-F(\theta)}{f(\theta)}\right) V'\left(q\left(\theta\right)\right) = c$$

- Differentiating this wrt θ yields

$$\frac{dq}{d\theta} = -\frac{g'\left(\theta\right) v'\left(q\left(\theta\right)\right)}{v''\left(q\left(\theta\right)\right) g\left(\theta\right)},$$

  where $g\left(\theta\right) = \theta - \frac{1-F(\theta)}{f(\theta)}$
- Since the following holds

$$\theta V'(q(\theta)) = V'(q(\theta))\left(\frac{1-F(\theta)}{f(\theta)}\right) + c$$

  we have

$$V'(q(\theta)) = \frac{c}{\theta - \left[\left(1-F(\theta)\right)/f(\theta)\right]}$$

- We require that $V'(q(\theta))$ be falling in θ and hence require that $\theta - \frac{1-F(\theta)}{f(\theta)}$ be increasing in θ
- That is, that the hazard rate be increasing
- Now turn attention to $T(q)$
- $\widehat{T'}(q) > c$ except for at the very top where $\widehat{T'} = c$
- Therefore it can't be convex
- Note that

$$1 - \frac{c}{p} = \frac{1 - F(\theta)}{\theta f(\theta)}$$

$$\frac{\theta f(\theta)}{1 - F(\theta)} \uparrow \theta \Leftrightarrow \frac{dp}{dq} < 0$$

- And note that $\frac{dp}{dq} = \widehat{T''}(q)$
- So the IHRC $\Rightarrow \frac{dp}{dq} < 0$
- If the IHRC does not hold the Monotonicity Constraint binds and we need to applying "Ironing" (See Bolton & Dewatripont)
- Use Pontryagin's Principle to find the optimal cutoff points
- Require $\lambda(\theta_1) = \lambda(\theta_2) = 0$, where λ is the Lagrange multiplier
- Still get optimality and the top and sub-optimality elsewhere

2.1.5 Random Schemes

- Key paper is Maskin & Riley (RAND, 1984)
- A deterministic scheme is always optimal if the seller's program is convex
- But if the ICs are such that the set of incentive feasible allocations is not convex then random schemes may be superior

dtbpF3.3529in2.0678in0ptFigure

- Both types are risk-averse
- So S loses money on the low type, but may be able to charge enough more to the high type to avoid the randomness if the high type is more risk-averse
- If they are sufficiently more risk-averse (ie. the types are far enough apart), then the random scheme dominates
- Say: announce $\theta = \theta^a$ and get a draw from a distribution, so get $(\widetilde{q}, \widetilde{T})$
- If the high type is less risk-averse than the low type then the deterministic contract dominates

- The only incentive constraints that matter are the downward ones
- So if the high type is less risk-averse then S loses money on that type from introducing randomness
- And doesn't gain anything on the low type, because her IR constraint is already binding and so can't extract more rents from her

2.1.6 Extensions and Applications

- Jullien (2000) and Rochet & Stole (2002) consider more general PCs (egs. type dependent or random)
- Classic credit rationing application: Stiglitz & Weiss (1981)

Multi-Dimensional Types

- So far we have assumed that a single parameter θ captures all relevant information
- Laffont-Maskin-Rochet (1987) were the first to look at this
- They show that "bunching" is more likely to occur in a two-type case than a one-type case (ie. Monotone Hazard Rate condition violated)
- Armstrong (Ecta, 1996) provides a complete characterization
  - Shows that some agents are always excluded from the market at the optimum (unlike the one-dimensional case where there is no exclusion)
  - In one dimension if the seller increases the tariff uniformly by ε then profits go up by ε on all types whose IR was slack enough (so that they still participate), but lose on all the others
  - With multi-dimensional types the probability that an agent had a surplus lower than ε is a higher order term in ε – so the loss is lower from the increase even if there is exclusion
- Rochet-Chone (1997) shows that
  - Upward incentive constraints can be binding at the optimum
  - Stochastic contracts can be optimal
  - There is no generalization of the MHRC which can rule out bunching
- Armstrong (1997) shows that with a large number of independently valued dimensions the the optimal contract can be approximated by a two-part tariff

Aside: Multi-Dimensional Optimal Income Taxation

- Mirrlees (JPubE, 1976) considered the problem of multi-dimensional optimal income taxation
- Strictly harder than the above problems because he doesn't assume quasi-linear utility functions only

- He shows how, when $m < n$ (i.e. the number of characteristics is smaller than the number of commodities), the problem can be reduced to a single eliptic equation which can be solved by well-known method
- When $m \geq n$ (i.e. the number of characteristics is at least at large than the number of commodities) the above approach does not lead to a single second-order partial differential equation, but a system of m second-order partial differential equations for the m functions M_j
- Numerical evidence has shown recently that a lot of the conclusions from the one-dimensional case go away in multiple dimensions (eg. the no distortion at the top result)
- But the system of second-order PDEs seem very hard to solve

3 Signaling and Perfect Bayesian Equilibrium

3.1 Introduction

- What about the informed side of the market? Can they improve their outcome in equilm?
- Under certain circumstances they, indeed, can
- Seminal paper is Spence (1973)

3.2 Setup and Basic Analysis

- Labor market example again
- Suppose that there are two types of worker with types $\theta_H > \theta_L$ and that $pr(\theta_H) = \lambda$ which is interior
- Worker can obtain education and that is observable to all market participants
- Assume that education has no effect on productivity
- Let the cost of education be $c(e, \theta)$ with $c(0, \theta) = 0, c_e(e, \theta) >), c_e e(e, \theta) > 0, c_\theta(e, \theta) < 0$ for all $e > 0$ and let $c_{e,\theta}(e, \theta) < 0$ (notice single-crossing)
- Utility is $u = w - c(e, \theta)$ and outside option is $r(\theta)$
- Can education, though useless, serve as a signaling device?
- For simplicity assume that $r(\theta) = 0$, so that in the FB we have $w^* = E[\theta]$
- Extensive form: (i) nature chooses type, (ii) worker chooses education, (iii) firm makes job offer, (iv) worker accepts or rejects
- Need to formalize the *solution concept*
- Here we will use *Perfect Bayesian Equilibrium*–remark about Sequential Equilibrium and Fudenberg-Tirole (1992)

- PBE involves there existing a $\mu(e) \in [0,1]$ such that the firm's belief that the work is the high type after observing e is $\mu(e)$ and that after the worker has chose e any other firm believes that the worker's type is high and the other firm(s) has played according to equilm.
  - Worker's streategy is optimal given the firm's strategy
  - Worker's belief function is derived using Bayes Rule where possible
  - If there are multiple firms their offers are Nash equila in the subgame
- Proceed by backward induction
-
-

3.3 Refinements and Cho-Kreps

4 Persuason and Information Transmission

4.1 Cheap Talk

- Crawford-Sobel (Ecta, 1982)
- Main question: how much information can be transmitted when communication is costless, but interests are not necessarily aligned?
- cf. signalling models: the key ingredient there is that communication is *costly*
  - eg. Spence job market signalling: to get separation need education to more costly for certain types than for others
- Two parties: a decision maker who is uninformed, and an informed expert
- DM is to make a decision $d \in [0,1]$
- State of nature is $\theta \in [0,1]$
- DM has a uniform prior about θ
- DM's payoff is $U(d,\theta) = -(d-\theta)^2$
- The E(xpert) knows the value of θ and her payoff is $V(b,\theta,d) = -(d-(\theta+b))^2$
- $b \geq 0$ is a measure of the bias of the expert
- E may send a message $m \in [0,1]$
- Timing:

1. 1. E observes θ
   2. E sends m to DM
   3. DM chooses d

- Solution concept: PBE

Proposition 2. *For all b there exists a "babbling equilbrium" in which E sends a random message ("babbles") and hence no information is conveyed.*

- Intuition: in a babbling equilibrium DM believes there is no information content in the message. E then has no incentive to send an informative message, so is happy to babble
- Bigger question: are there informative equilibria?
- Preliminary question: are there equilibria in which information is truthfully conveyed?

Proposition 3. *There exists an equilibrium in which information is fully revealed if and only if $b = 0$.*

- Proof sketch: suppose $b > 0$ and E truthfully revealed θ. In this equilibrium she is believed, but her payoff could be increased in some states by deviating to a message $\theta + b$–a contradiction.
- Now we constuct an equilibrium in which *some* information is conveyed
- Let DM's posterior distribution about the value of θ given m be $G(\theta|m)$
- Given quadratic preferences

$$\begin{aligned} d^*(m) &\equiv \max_{d\in[0,1]} \left\{ \int U(d,\theta)\, G(\theta|m)\, d\theta \right\} \\ &= E[\theta|m]. \end{aligned}$$

- E knows this, of course, and could be faced with the following problem
- Suppose message m leads to action d and message m' leads to action $d' > d$
- Also, suppose that in state $\theta'' > \theta'$ E prefers d' to d but in state θ' prefers d to d'
- Noting that V satisfies single crossing, $d^2V/d\theta dd > 0$ and hence E prefers d' to d for all $\theta > \theta''$
- Therefore, by the Intermediate Value Theorem, there exists a state $\hat{\theta}$ such that $\theta' < \hat{\theta} < \theta''$ in which E is indifferent between d and d'
- This is the same as saying that the distance between E's bliss point and d in state $\hat{\theta}$ is the same as the distance between the bliss point and d'
- ie. $\hat{\theta} + b - d = d' - \left(\hat{\theta} + b\right)$
- So E sends message m for all $\theta < \hat{\theta}$ and message m' for all $\theta > \hat{\theta}$
- For this to be an equilbrium we need to find d, d' and $\hat{\theta}$ such that

$$\begin{aligned} \hat{\theta} + b - d &= d' - \left(\hat{\theta} + b\right), \text{ and} \\ d(m) &= E[\theta|m]. \end{aligned}$$

- Solving we have

$$\begin{aligned} d &= \frac{\hat{\theta}}{2}, \\ d' &= \frac{1+\hat{\theta}}{2}. \end{aligned}$$

- Substituting into $\hat{\theta} + b - d = d' - \left(\hat{\theta} + b\right)$ we have

$$\hat{\theta} = \frac{1}{2} - 2b.$$

- Clearly such an equilibrium exists
- Moreover the cutoff $\hat{\theta}$ is uniquely determined by b
- If bias is too big then non-existence (ie. $b > 1/4$)
- This is a particular equilm with just two partitions
- But when bias is "small" there exist equilibria with more than two partitions

Theorem 1 (Crawford-Sobel)**.** *There exists a partition equilibrium of every size (ie. number of partitions) from 1 (completely uninformative) to $N(b)$ (the most informative).*

- Many equilibria!
- One thing to focus on is the impossibility of perfectly informative communication
- Another is the following quite general message: when preferences are somewhat aligned cheap talk "can" improve both party's payoff
- Cheap talk is just that: not an announcement in a mechanism, not a costly signal, just an unverifiable utterance
- One might think it could never help much (eg. Yogi Berra: "a verbal contract isn't worth the paper it's written on"), but the CS theorem shows that it *could*
- A large literature explores concrete settings in which it *does*
- Basic idea: cheap talk (by construction) does not directly affect payoffs, but it can affect them indirectly

4.2 Improved Communication

4.2.1 Conversation

- Krishna-Morgan (JET, 2004)
- Mulitiple messages from E can be subsumed as one message
- What about two-sided communication

- DM knows "nothing" and hence has no new information to reveal, but can act as a randomization device
- Illustration: suppose $b = 1/12$
- If only E gets to talk then the most informative equilm reveals whether θ is above or below 1/3
- Timing:

  1. 1. E observes θ
     2. E and DM meet face to face
     3. E delivers a "report"
     4. DM chooses d

- Consider the following equim:
  - In the meeting E reveals whether $\theta > 1/6$ or not and sends some other message (this determines whether the meeting is a "success" or not
  - If E sends the message that $\theta \leq 1/6$ then the meeting is deemed a failure and DM chooses $d = 1/12$
  - If E sends $\theta > 1/6$ then the report is conditional on the success or failure of the meeting
  - If the meeting was a failure then $d = 7/12$ (the optimal action conditional on $\theta > 1/6$
  - But if the meeting was a "success" then the report further partitions the interval $[1/6, 1]$ into $[1/6, 5/12]$ and $[5/12, 1]$
  - In the first subinterval $d^* = 7/24$ and in the second $d^* = 17/24$
  - If $\theta = 1/6$ then E prefers $d = 1/12$ to $d = 7/12$
  - So we need "uncertainty" about the outcome of the meeting–otherwise E would not be willing to reveal whether the state was above or below 1/6
  - If $\theta < 1/6$ then E would say $\theta \in [1/6, 5/12]$ and induce $d = 7/24$ and if $\theta > 1/6$ then E would announce $\theta < 1/6$ and induce $d = 1/12$ rather than $d = 7/12$
  - It turns out that when $\theta = 1/6$ then with probability $p = 16/21$ E is indifferent between $d = 1/12$ and the lottery where she gets $d = 7/12$ with probability p and $d = 7/24$ with probability $1 - p$ gets $d = 7/12$
  - When $\theta < 1/6$ E prefers $d = 1/12$ to the lottery and when $\theta > 1/6$ E prefers the lottery
  - So can we get the meeting to be successful with probability $p = 16/21$?
  - KM show that we can, as follows
  - Suppose E sends message (low, A_i) or $(high, A_i)$ and DM sends a message A_j with $i, j \in \{1, ..., 21\}$
  - Low means $\theta \leq 1/6$ and high means $\theta > 1/6$
  - The A_i and A_j parts of the message serve as a coordination device about the success of the meeting

- E chooses A_i randomly (i.e. from a uniform distribution)
- DM does similarly for A_j
- The meeting is deemed a success if $0 \leq i - j < 16$ or of $j - 1 > 5$ and a failure otherwise
- With this structure the probability that the meeting is a success is exactly $p = 16/21$
- So more information is conveyed than in any CS equilm
- Striking thing: having the DM participate in the conversation helps even though she is completely uninformed
- Aumann and Hart (2003) show that even with unlimited communication full revelation is impossible (cf. Geanakoplos-Polemarcharkis)

4.2.2 Delegation

- Can we do better by delegating to E?
- Tradeoff: E has her own preferences and is thus biased, but she is also informed
- Suppose $b = 1/12$ then direct computation yields a payoff of $-1/36$ in the most informative partition equilm of the CS model, but under delegation the action which is chosen is $d = \theta + b$, by construction, and the payoff is $-b^2 = -1/144$, so delegation is optimal
- This conclusion is more general than this example (see Dessein, 2002)
- DM can do even better by combining the amount of delegation/discretion
- Here the optimal thing to do is limit E's discretion to $d \in [0, 1 - b]$

4.2.3 Compensation

- An obvious ommision in what we did is to preclude the possibility of compensating E for her advice
- Can we do better with an optimal contract?
- Now add a transfer such that the payoffs are
  - DM's payoff is $U(d, \theta) = -(d - \theta)^2 - t$
  - E's payoff is $V(b, \theta, d) = -(d - (\theta + b))^2 + t$
- Again use mechanism design to find the optimal contract
- Can apply the revelation principle here and restrict attention to mechanisms/contracts whereby E announces d and θ truthfully in equilm
- Aside: this isn't cheap talk any more–talk affects payoffs directly here

- Suppose $t\left(\hat{\theta}\right) = 2b\left(1 - \hat{\theta}\right)$ and $d\left(\hat{\theta}\right) = \hat{\theta}$, then the FB decision is acheived and there is full revelation
- But this is costly for DM
- eg. when $b = 1/12$ her payoff is $-1/12$, whereas it is $-1/36$ in the best CS equilm
- General result: Krishna-Morgan (2004): Full revelation is in general feasible, but never optimal

4.2.4 Multiple Senders and Multiple Dimensions

- Battaglini (2002): two sender cheap talk with a one dimensional state space
  - Also showed that with a multi-dimensional Euclidean state space a perfectly revealing PBE can be constructed
  - Moreover, there are no out of equilm messages and so these equilbria survive any refinements which place restrictions on out of equilm beliefs
  - Construction: each sender conveys information only along directions in which her interests coincide with DM (ie. directions which are orthogonal to the bias of E)
  - Since these generically span the whole state space DM can extract all the information and perfectly identify the true state
- Ambrus-Takahashi (2007) consider restricted state spaces
  - eg. some policies may not be feasible
  - or some may never be chosen by DM (and so they are not rationalizable)
- AT provide the following example

dtbpF3.378in2.2857in0ptFigure

- Suppose DM needs to allocate a fixed budget to "education," "military spending," and "healthcare"
- Suppose there are two perfectly informed experts, a left-wing E and a right-wing E
- Left-wing E has a bias towards spending more on education, while right-wing E has a bias towards spending more on the military, but both of them are unbiased with respect to healthcare
- The state space in this example is represented by triangle ABC
- At B it is optimal for DM to spend the whole budget on the military
- At C it is optimal to spend all money on education
- At A it is optimal to spend no money on either education or military

- Left-wing E's bias is orthogonal to AB in the direction of C and right-wing E's bias is orthogonal to AC in the direction of B
- Battaglini's solution would have left-wing E report along a line parallel to AC (like asking how much money to spend on the military), right-wing E to report along a line parallel to AB
- But here it is not true that any pair of such reports identifies a point in the state space!
- Look at state θ
- If left-wing E sends a truthful report, then the right-wing analyst can send reports that put you outside the state space
- ie. they say that expenditure should be larger than the budget
- Doesn't happen in equilm, but have to specify out of equilm beliefs and this can cause problems for the construction
- Key points:
  - With multiple senders, the amount of information that can be transmitted in equilm depends on fine details such as: the shape of the boundary of the state space, how similar preferences of the senders are,...
  - Also properties of the state space and sender preferences cannot be investigated independently if one allows state-dependent preferences

5 Moral Hazard

5.1 Introduction

- Many applications of principal-agent problems
  - Owner / Manager
  - Manager / Worker
  - Patient / Doctor
  - Client / Lawyer
  - Customer / Firm
  - Insurer / Insured
- History:
  - Arrow ('60s)
  - Pauly (68), Spence-Zeckhauser
  - Ross (early '70s)
  - Mirrlees (mid '70s)
  - Holmström ('79)
  - Grossman-Hart ('83)

5.2 The Basic Principal-Agent Problem

5.2.1 A Fairly General Model

- $a \in A$ (Action Set)
- This leads to q (verifiable revenue)
- Stochastic relationship $F(q; a)$
- Incentive scheme $I(q)$
- The Principal solves the following problem:

$$\max_{\widehat{I}(\cdot),\widehat{a}} \left\{ \int \left(q - \widehat{I}(q) \right) dF(q; \widehat{a}) \right\}$$

$$s.t.(i)\ \widehat{a} \text{ solves } \max_{a \in A} \left\{ \int u(a, \widehat{I}(q)) dF(q; a) \right\} \tag{ICC}$$

$$(ii) \int u(\widehat{a}, \widehat{I}(a)) dF(q; \widehat{a}) \geq \overline{U} \tag{PC}$$

- Use the deterministic problem of the Principal inducing the Agent to choose the action because there may be multiple actions which are equivalent for the Agent but the Principal might prefer one of them
- The Principal is really just a risk-sharing device

5.2.2 The First-Order Approach

- Suppose $A \subseteq \mathbb{R}$
- The problem is now

$$\max_{a,I(\cdot)} \left\{ \int_{\underline{q}}^{\bar{q}} (q - I(q)) f(q|a) dq \right\}$$

subject to

$$a \in \arg\max_{\hat{a} \in A} \left\{ \int_{\underline{q}}^{\bar{q}} u(I(q)) f(q\,|a|) dq - G(a) \right\} \tag{ICC}$$

$$\int_{\underline{q}}^{\bar{q}} u(I(q)) f(q\,|a|) dq - G(a) > \overline{U} \tag{PC}$$

- IC looks like a tricky object
- Maybe we can just use the FOC of the agent's problem
- That's what Spence-Zeckhauser, Ross, Harris-Raviv did

- FOC is
$$\int_{\underline{q}}^{\bar{q}} u(I(q))f_a(q|a)dq = G'(a)$$

- SOC is
$$\int_{\underline{q}}^{\bar{q}} u(I(q))f_{aa}(q|a)dq = G''(a)$$

- If we use the first-order condition approach:

$$\begin{aligned}\frac{\partial}{\partial I} &= 0 \Rightarrow -f(q;a) + \mu u'(I(q))f_a(q|a) + \lambda u'(I(q))f(q|a)) = 0 \\ &\Rightarrow \frac{1}{u'(I(q))} = \lambda + \mu\frac{f_a(q;a)}{f(q;a)}\end{aligned}$$

- f_a/f is the likelihood ratio
- $I \uparrow q \Leftrightarrow \frac{f_a}{f} \uparrow q$
- But the FOC approach is not always valid – you are throwing away all the global constraints
- The $I(q)$ in the agent's problem is endogenous!
- MLRP $\Rightarrow$ "the higher the income the more likely it was generated by high effort"

Condition 1 (Monotonic Likelihood Ratio Property ("MLRP")). *(Strict) MLRP holds if, given $a, a' \in A$, $a' \preceq a \Rightarrow \pi_i(a')/\pi_i(a)$ is decreasing in i.*

Remark 1. *It is well known that MLRP is a stronger condition than FOSD (in that MLRP $\Rightarrow$ FOSD, but FOSD $\nRightarrow$ MLRP).*

Condition 2 (Covexity of the Distribution Function Condition). $F_{aa} \geq 0$.

Remark 2. *This is an awkward and somewhat unnatural condition–and it has little or no economic interpretation. The CDFC holds for no known family of distributions*

- MLRP and CDFC ensure that it will be valid (see Mirrlees 1975, Grossman and Hart 1983, Rogerson 1985)
- FOC approach valid when FOC≡ICC
- In general they will be equivalent when the Agent has a convex problem
- To see why (roughly) they do the trick suppose that $I(q)$ is almost everywhere differentiable (although since it's endogenous there's no good reason to believe that)
  - The agent maximizes
$$\int_{\underline{q}}^{\bar{q}} u(I(q))f(q|a)dq - G(a)$$

- Integrate by parts to obtain

$$u(I(\bar{q})) - \int_{\underline{q}}^{\bar{q}} u'(I(q))\, I'(q)\, F(q|a)dq - G(a)$$

- Now differentiate twice w.r.t. a to obtain

$$-\int_{\underline{q}}^{\bar{q}} u'(I(q))\, I'(q)\, F_{aa}(q|a)dq - G''(a) \quad (*)$$

- MLRP implies that $I'(q) \geq 0$
- CDFC says that $F_{aa}(q|a) \geq 0$
- $G''(a)$ is convex by assumption
- So (*) is negative

- Jewitt's (Ecta, 1988) assumptions also ensure this by restricting the Agent's utility function such that this is the case

- Grossman and Hart (Ecta, 1983), proposed the LDFC, (initially referred to as the Spanning Condition).

- Mirrlees and Grossman-Hart conditions focus on the Agent controlling a family of distributions and utilize the fact that the ICC is equivalent to the FOC when the family of distributions controlled by the Agent is one-dimensional in the distribution space (which the LDFC ensures), or where the solution is equivalent to a problem with a one-dimensional family (which the CDFC plus MLRP ensure)

Remark 3. *Single-dimensionality in the distribution space is not equivalent to the Agent having a single control variable – because it gets convexified*

- It is easy to see why the LDFC works because it ensures that the integral in the IC constraint is linear in e.

5.2.3 Beyond the First-Order Approach I: Grossman-Hart

Grossman-Hart with 2 Actions

- Grossman-Hart (Ecta, 1983)

- Main idea of GH approach: split the problem into two step

  - Step 1: figure out the lowest cost way to implement a given action
  - Step 2: pick the action which maximizes the difference between the benefits and costs

- $A = \{a_L, a_H\}$ where $a_L < a_H$ (in general we use the FB cost to order actions–this induces a complete order over A if A is compact)

- Assume $q = q_1 < ... < q_n$

- Note: a finite number of states

- Payment from principal to agent is I_i in state i
- $a_H \to (\pi_1(a_H), ..., \pi_n(a_H))$
- $a_L \to (\pi_1(a_L), ..., \pi_n(a_L))$
- Agent has a v-NM utility function $U(a, I) = V(I) - G(a)$
- $G(a_H) > G(a_L)$
- Reservation utility of $\overline{U}$
- Assume V defined on $(\underline{I}, \infty)$
- $V' > 0, V'' < 0, \lim_{I \to \underline{I}} V(I) = -\infty$ (avoid corner solutions, like $\ln(I)$ instead of $I^{1/2}$)
- Of course, a legitimate v-NM utility function has to be bounded above and below (a result due to Arrow), but...

First Best (a verifiable):

- Define $h \equiv V^{-1}$
- $V(h(V)) = V$
- Pick a
- Let $C_{FB}(a) = h(\overline{U} + G(a))$
- since $V(I) - G(a) = \overline{U}$, $V(I) = G(a) + \overline{U}, I = h(\overline{U} + G(a))$
- Can write the problem as

$$\max_{a \in A} \{_{i=1}^{n} \pi_i(a) q_i - C_{FB}(a)\}$$

Second Best:

- $a = a_L$ then pay you $C_{FB}(a_L)$ regardless of the outcome
- $a = a_H$

$$\min_{I,\dots,I_n} \left\{ \sum_{i=1}^{n} \pi_i(a_H) I_i \right\}$$

$$s.t.(i) \sum_{i=1}^{n} \pi_i(a_H) V(I_i) - G(a_H) \geq \sum_{i=1}^{n} \pi_i(a_L) V(I_i) - G(a_L) \qquad \text{(ICC)}$$

$$(ii) \sum_{i=1}^{n} \pi_i(a_H) V(I_i) - G(a_H) \geq \overline{U} \qquad \text{(PC)}$$

- We use the Vs as control variables (which is OK since V is strictly increasing in I)

- $v_i = V(I_i)$

$$\min_{v_1,\ldots,v_n} \left\{ \sum_{i=1}^{n} \pi_i(a_H) h(v_i) \right\} \tag{*}$$

$$s.t.(i) \sum_{i=1}^{n} \pi_i(a_H) v_i - G(a_H) \geq \sum_{i=1}^{n} \pi_i(a_L) v_i - G(a_L) \tag{ICC}$$

$$(ii) \sum_{i=1}^{n} \pi_i(a_H) v_i - G(a_H) \geq \overline{U} \tag{PC}$$

- Now this is just a convex programming problem
- Note, however, that the constraint set is unbounded – need to be careful about the existence of a solution

Claim 1. *Assume $\pi_i(a_H) > 0, \forall i$. Then $\exists$ a unique solution to (*)*

Proof. (sketch): The only way there could not be a solution would be if there was an unbounded sequence $(v_1', ..., v_n') \Rightarrow I$s are unbounded above $\Rightarrow VarI \to \infty$, where $I_i = h(v_i)$. $\underset{\sim}{V}$ unbounded $\Rightarrow \underset{\sim}{I}$ unbounded above (if not Is$\to \underline{I}$ and vs$\to -\infty \Rightarrow$PC violated. With $V(\cdot)$ strictly concave $E[\underset{\sim}{I}] \to \infty$ as $\underset{\sim}{I} \to \infty$ if $\underline{I} \neq -\infty$. If $\underline{I} = -\infty$ the PC will be violated because of risk-aversion. □

- Solution must be unique because of strict convexity with linear constraints
- π_is are all positive
- Let the minimized value be $C(a_H)$
- Compare $\sum_{i=1}^{n} \pi_i(a_H)q_i - C(a_H)$ to $\sum_{i=1}^{n} \pi_i(a_L)q_i - C_{FB}(a_L)$
- This determines whether you want a_H or a_L in the second-best

Claim 2. $C(a_H) > C_{FB}(a_H)$ *if* $G(a_H) > G(a_L)$. *The second-best is strictly worse than the first-best if you want them to take the harder action.*

Proof. (sketch): Otherwise the ICC would be violated because all of the π_is are positive and so all the vs would have to be equal - which implies perfect insurance. □

Claim 3. *The PC is binding*

Proof. (sketch): If $\sum_{i=1}^{n} \pi_i(a_H)v_i - G(a_H) > \overline{U}$ then we can reduce all the v_is by ε and the Principal is better off without disrupting the ICC. □

- FB=SB if:

1. Shirking is optimal
2. V is linear and the agent is wealthy (risk neutrality) – make the Agent the residual claimant (but need to avoid the wealth constraint)

3. $\exists i$ sth $\pi_i(a_H) = 0, \pi_i(a_L) > 0$ (MOVING SUPPORT). If the Agent works hard they are perfectly insured, if not they get killed.

- Now form the Lagrangian:

$$\begin{aligned} = & \sum_{i=1}^{n} \pi_i(a_H) h(v_i) \\ & -\mu \left(\sum_{i=1}^{n} \pi_i(a_H) v_i - G(a_H) - \sum_{i=1}^{n} \pi_i(a_L) v_i + G(a_L) \right) \\ & -\lambda \left(\sum_{i=1}^{n} \pi_i(a_H) v_i - G(a_H) \right) \end{aligned}$$

- The FOCs are:

$$\frac{\partial}{\partial v_i} = 0, \forall i$$

$$\pi_i(a_H) h'(v_i) - \mu \pi_i(a_H) + \mu \pi_i(a_L) - \lambda \pi_i(a_H) = 0$$

$$\frac{1}{V(I_i)} = h'(v_i) = \lambda + \mu - \mu \frac{\pi_i(a_L)}{\pi_i(a_H)} \quad \forall i = 1, ..., n$$

- Note that $\mu > 0$ since if it was not then $h'(v_i) = \lambda$ which would imply that the v_is are all the same, thus violating the ICC
- Implication: Payments to the Agent depend on the likelihood ratio $\frac{\pi_i(a_L)}{\pi_i(a_H)}$

Theorem 2. *In the Two Action Case, Necessary and Sufficient conditions for a monotonic incentive scheme is the MLRP*

- This is because the FOC approach is valid the in the 2 action case even w/out the CDFC
- *This behaves like a statistical inference problem even though it is not one (because the actions are endogenous)*
- Linearity would be a very fortuitous outcome
- Note: in equilm the Principal knows exactly how much effort is exerted and the deviations of performance from expectation are stochastic – but this is optimal *ex ante*

5.3 The Value of Information in Agency Problems

5.4 Motivating Questions

1. How valuable is better information in an agency relationship?
2. What do we mean by "better" information?

3. Are there differences between the classic model, the multi-task model, and the career concerns model?

4. Can it be good for the principal to add noise?

- Other motivation: there are a number of very useful tools that are used to answer these questions, and it's useful as either a consumer or producer of applied theory to be familiar with them.

5.5 Information in the Linear Model

- Recall the Holmstrom-Milgrom linear model
- Say $w = t + vq$
- Assume normally distributed performance and CARA (exponential) utility
- Let $q = a + \varepsilon$ with $\varepsilon \sim N(0, \sigma^2)$
- Assume the Principal is risk-neutral
- The Agent is risk-averse with:

$$U(w, a) = -e^{-r(w-\psi(a))}$$

- Let $\psi(a) = \frac{ca^2}{2}$
- Note that r is the coefficient of absolute risk-aversion $-u''/u'$
- The Principal solves:

$$\max_{a,t,v} E[q - w]$$

$$s.t.(i)E[-e^{-r(w-\psi(a))}] \geq -e^{-r\overline{w}} \quad \text{(IR)}$$

$$(ii)a \in \arg\max_{a} E[-e^{-r(w-\psi(a))}] \quad \text{(IC)}$$

- Let $x \sim N(0, \sigma_x^2)$
- $E[e^{\gamma x}] = e^{\gamma^2 \sigma_x^2 / 2}$ (this is essentially the calculation done to yield the moment generating function of the normal distribution – see Varian for a more detailed derivation)

$$\begin{aligned}
& E[-e^{-r(w-\psi(a))}] \\
= \quad & -E[-e^{-r(t+va+v\varepsilon-\psi(a))}] \\
= \quad & -e^{-r(t+va-\psi(a))} E[e^{-rv\varepsilon}] \\
= \quad & e^{-r\widehat{w}(a)}
\end{aligned}$$

- $\widehat{w}(a) = t + va - \frac{r}{2}v^2\sigma^2 - \frac{1}{2}ca^2$
- Now $\max_{a} \{\widehat{w}(a)\}$

- FOC is $v - ca = 0 \Rightarrow a = v/c$
- Replace a with v/c in the Principal's Problem and they solve:

$$\max_{v,t} \left\{ \frac{v}{c} - (t + \frac{v^2}{c}) \right\} \tag{2}$$

$$s.t.\widehat{w}(a) = \widehat{w}(\frac{v}{c}) = \overline{w} \tag{3}$$

- The IR constraint is, written more fully:

$$t + \frac{v^2}{c} - \frac{r}{2}v^2\sigma^2 - \frac{v^2}{2c}$$

- ie.

$$t + \frac{v^2}{2c} - \frac{r}{2}v^2\sigma^2 = \overline{w}$$

- Substituting for t:

$$\max_{v} \left\{ \frac{v}{c} - \frac{v^2}{c} + \frac{v^2}{2c} - \frac{r}{2}v^2\sigma^2 - \overline{w} \right\}$$

- The FOC is:

$$\frac{1}{c} - \frac{v}{c} - rv\sigma^2 = 0$$

- Hence:

$$v = \frac{1}{1 + rc\sigma^2}$$

- As the amount of noise increases, the intensity of incentives goes down, and so does the principal's payoff.
- This conclusion is extremely general in single agent principal-agent models–even when the first-order approach isn't valid (Holden, 2006).
  - Grossman-Hart (1983) show that a Blackwell garbling increases agency costs
  - Kim (1995) shows the same thing for a mean preserving spread of the likelihood ratio (when the first-order approach is valid)

5.6 The Sufficient Statistic Theorem

- Recall the definition of a sufficient statistic

Definition 2. *A statistic $T(x)$ is* ***sufficient*** *for a parameter θ if the conditional distribution of the data X, given the statistic $T(x)$, does not depend on the parameter θ. i.e. $pr(X = x|T(x) = t, \theta) = pr(X = x|T(x) = t)$.*

- Setup
  - Agent takes an action $a \in A$, where A is some (possibly high dimensional) compact set.

- Possible outcomes (revenues to P) $\{q_1, ..., q_n\}$.
- Action induces a probability distribution over outcomes, so that the probabilities of the n states are $\pi_1(a), ..., \pi_n(a)$.
- Agent's cost of effort is $\psi(a)$.
- Payment from P to agent in state i is w_i.
- Agent's utility function is u and P's payoff function is V.

- Say there is a signal which is realized after effort is chosen by the Agent but before the realization of the outcome such that :

$$\pi_{ij}(a) = \pi(i, j \mid a)$$

- ie. probability of outcome i, signal j conditional on action a
- Signal does not enter directly into objective functions – only though the probabilities
- Now, letting $\psi(a)$ be the cost of effort, the Principal solves:

$$\max \left\{ \sum_{i,j} \pi_{ij}(a) V(q_i - w_{ij}) \right\}$$

$$s.t.(i) \sum_{i,j} \pi_{ij}(a) u(w_{ij}) - \psi(a) \geq \overline{U} \tag{IR}$$

$$(ii) a \in \arg\max \left\{ \sum_{i,j} \pi_{ij}(a) u(w_{ij}) - \psi(a) \right\} \tag{IC}$$

- Put the Lagrange multiplier λ on the IR
- The IC FOC is $\sum \pi'_{ij}(a) u(w_{ij}) = 1$
- Forming the Lagrangian and finding $\frac{\partial L}{\partial w_i} = 0, \forall i, \forall j$ yields:

$$\frac{V'(q_i - w_{ij})}{u'(w_{ij})} = \lambda + \mu \frac{\pi'_{ij}(a)}{\pi_{ij}(a)} \tag{4}$$

- When is the optimal w_{ij} independent of j ?
- Same as before if

$$\frac{\pi'_{ij}(a)}{\pi_{ij}(a)} = \frac{\pi'_i(a)}{\pi_i(a)}$$

- In the continuum case (denoting the additional signal s this is just:

$$\frac{g_a\left(q, s|a\right)}{g\left(q, s|a\right)} = \frac{f_a\left(q|a\right)}{f\left(q|a\right)}$$

- Integrating this object with respect to a means that it is equivalent to the existence of two functions $m(q|a)$ and $n(q|s)$ such that:

$$g(q,s|a) = m(q|a)\, n(q|s).$$

- That is, that q is a *sufficient statistic* for the pair (q,s) with respect to a
- This representation is known as the Halmos-Savage factorization criterion (or theorem) – see DeGroot (1971) for further details
- So, the optimal incentive scheme is conditioned on s if and only if s is informative about a, given that q is already available

5.6.1 Random Schemes

- Can one do better with random schemes? Do you want to add noise ?
- Suppose the Principal decided to "flip a coin", $j \in \{1, ..., m\} \to pr(j) = q(j)$
- $\pi_{ij}(a) = q_j \pi_i(a)$
- Suppose w_i was the optimal scheme and let $\widetilde{w}_i$ be the certainty equivalent:

$$u(\widetilde{w}_i) = \sum_j q_j u(w_{ij}) \quad , \forall i$$

- But we haven't changed the IC or IR
- However, the Principal has cost $\widetilde{w}_i$ and $\widetilde{w}_i < \sum_j q_j w_{ij}$ due to the concavity of $u(\cdot)$. So the Principal is better off. Contradiction
- Therefore random schemes cannot be better
- They put more risk onto the risk-averse Agent and that requires the Agent to be compensated for bearing that risk
- Can also use the sufficient statistic result - the random scheme adds no information about the likelihood ratio (and generalizes to the case where the Principal is risk-averse)

5.7 Linear Contracts

- Very little that you can say in a general moral hazard model (Grossman and Hart 83)
- Say $w = t + vq$
- Assume normally distributed performance and CARA (exponential) utility
- Let $q = a + \varepsilon$ with $\varepsilon \sim N(0, \sigma^2)$
- Assume the Principal is risk-neutral

- The Agent is risk-averse with:

$$U(w, a) = -e^{-r(w-\psi(a))}$$

- Let $\psi(a) = \frac{ca^2}{2}$
- Note that r is the coefficient of absolute risk-aversion $-u''/u'$
- The Principal solves:

$$\max_{a,t,v} E[q - w]$$
$$s.t.(i)E[-e^{-r(w-\psi(a))}] \geq -e^{-r\overline{w}} \quad \text{(PC)}$$
$$(ii)a \in \arg\max E[-e^{-r(w-\psi(a))}] \quad \text{(ICC)}$$

- Let $x \sim N(0, \sigma_x^2)$
- $E[e^{\gamma x}] = e^{\gamma^2 \sigma_x^2/2}$ (this is essentially the calculation done to yield the moment generating function of the normal distribution - see Varian for a more detailed derivation)

$$\begin{aligned} & E[-e^{-r(w-\psi(a))}] \\ = & -E[-e^{-r(t+va+v\varepsilon-\psi(a))}] \\ = & -e^{-r(t+va-\psi(a))}E[e^{-rv\varepsilon}] \\ = & e^{-r\widehat{w}(a)} \end{aligned}$$

- $\widehat{w}(a) = t + va - \frac{r}{2}v^2\sigma^2 - \frac{1}{2}ca^2$
- Now $\max_a \{\widehat{w}(a)\}$
- FOC is $v - ca = 0 \Rightarrow a = v/c$
- Replace a with v/c in the Principal's Problem and they solve:

$$\max_{v,t} \left\{ \frac{v}{c} - (t + \frac{v^2}{c}) \right\} \quad (5)$$
$$s.t.\widehat{w}(a) = \widehat{w}(\frac{v}{c}) = \overline{w}PC \quad (6)$$

- The PC is, written more fully:

$$t + \frac{v^2}{c} - \frac{r}{2}v^2\sigma^2 - \frac{v^2}{2c}$$

- ie.

$$t + \frac{v^2}{2c} - \frac{r}{2}v^2\sigma^2 = \overline{w}$$

- Substituting for t:

$$\max_v \left\{ \frac{v}{c} - \frac{v^2}{c} + \frac{v^2}{2c} - \frac{r}{2} v^2 \sigma^2 - \overline{w} \right\}$$

- The FOC is:

$$\frac{1}{c} - \frac{v}{c} - rv\sigma^2 = 0$$

- Hence:

$$v = \frac{1}{1 + rc\sigma^2}$$

- Which is a nice, simple, closed form solution
- But the linearity restriction is not at all innocuous
- In fact, linear contracts are not optimal in this setting!
- Without the restriction one may approximate the first-best

Example 1: Moving Support

- $q = a + \varepsilon$ and ε is uniformly distributed on $[-k, k]$ with $k > 0$
- So the Agent's action moves the support of q

Claim 4. *The first-best can be implemented by a non-linear contract*

Proof. Let a^* be the first-best level of effort. q will take values in $[a^* - k, a^* + k]$. Scheme: pay w^* whenever $q \in [a^* - k, a^* + k]$ and pay $-\infty$ otherwise. Just a Mirrlees Scheme (which is certainly not linear) □

- With bounded support the Principal can rule out certain outcomes *provided the Agent chooses the FB action.*

Example 2:

- $q = a + \varepsilon$ and $\varepsilon \sim N[0, \sigma^2]$

$$\Rightarrow f(q, a) = \frac{1}{(2\pi\sigma)^{1/2}} e^{-(q-a)^2/2\sigma^2}$$

- Calculate the likelihood ratio:

$$f_a(q, a) = -\frac{1}{(2\pi\sigma)^{1/2}} e^{-(q-a)^2/2\sigma^2} \times \frac{-(q-a)}{\sigma^2}$$

- $\frac{f_a}{f} = \frac{q-a}{\sigma^2}$
- as $q \to \infty^+$, $\frac{f_a}{f} \to \infty$
- So the likelihood ratio can take on values on $(-\infty, \infty)$

- For extreme values (ie. in the tails of the distn) the Principal gets almost perfect information

Claim 5. *FB a^* can be arbitrarily approximated*

Proof. Suppose the Principal chooses an incentive scheme as follows: if $q < \underline{q} \rightarrow$ low transfer k, if $q \geq \underline{q} \rightarrow$ transfer w^*. Suppose the Agent has a utility function $u(y), u'(y) > 0, u''(y) < 0$ and cost of effort $\psi(a)$. To implement a^* under the above scheme we need that:

$$IC : \int_{-\infty}^{\underline{q}} u(k) f_a(q, a^*) dq + \int_{\underline{q}}^{\infty} u(w^*(q)) f_a(q, a^*) dq = \psi'(a^*)$$

But this violates the PC by:

$$l = \int_{-\infty}^{\underline{q}} \left[u(w^*(q)) - u(k) \right] f(q^*) dq$$

□

Claim 6. *One can choose $\underline{q}$ and k to make l arbitrarily small.*

Proof. Given $-M, \exists \underline{q}$ such that:

$$\frac{f_a(q,a)}{f(q,a)} \leq -M \quad \text{for } q \leq \underline{q}$$

$\Rightarrow \frac{f_a}{f} \left(\frac{-1}{M} \right) \geq 1 \Leftrightarrow f \leq f_a \left(\frac{-1}{M} \right)$

$$\begin{aligned} \Rightarrow \quad & l \leq \int_{-\infty}^{\underline{q}} \left[u(w^*(q)) - u(k) \right] f_a(q^*, a) \left(\frac{-1}{M} \right) dq \\ = \quad & \frac{-1}{M} (\cdot) \end{aligned}$$

Therefore one can make l arbitrarily small by making M arbitrarily large □

- The *expected* punishment is bounded away from ∞
- Mirrlees's (1974) idea again - this time without the moving support
- Although the size of the punishment grows, its frequency falls at a faster rate

5.8 Moral Hazard with Multiple Tasks

5.8.1 Holmström-Milgrom

- Holmström-Milgrom (JLEO, 1991)
- Different tasks with different degrees of measurability

- Suppose the Agent can sell the Principal's product or someone else's product
- 2 tasks $i = 1, 2$
- Let $q_i = a_i + \varepsilon_i$
- $(\varepsilon_1, \varepsilon_2) \sim N(0, \Sigma)$ where

$$\Sigma = \begin{pmatrix} \sigma_1^2 & R \\ R & \sigma_2^2 \end{pmatrix}$$

- Let the Agent's utility be given by:

$$-e^{-r(w-\psi(a_1,a_2))}$$

- where $\psi(a_1, a_2) = \frac{1}{2}(c_1 a_1^2 + c_2 a_2^2) + \delta a_1 a_2$
- if $\delta > 0$ then the two tasks are technological substitutes, if $\delta < 0$ they are complements
- Assume a linear incentive scheme:

$$w = t + v_1 q_1 + v_2 q_2$$

$$\begin{aligned} \widehat{w}(a_1, a_2) &= E[w(a_1, a_2)] - \frac{r}{2} var\left(w(a_1, a_2)\right) - \psi(a_1, a_2) \\ &= E[t + v_1\left(a_1 + \varepsilon_1\right) + v_2\left(a_2 + \varepsilon_2\right)] \\ &\quad - \frac{r}{2} var(t + v_1\left(a_1 + \varepsilon_1\right) + v_2\left(a_2 + \varepsilon_2\right)) \\ &\quad - \frac{1}{2}((c_1 a_1^2 + c_2 a_2^2) + \delta a_1 a_2) \end{aligned}$$

- $E[t + v_1\left(a_1 + \varepsilon_1\right) + v_2\left(a_2 + \varepsilon_2\right)] = t + v_1 q_1 + v_2 q_2$
- $Var(\cdot) = v_1^2 \sigma_1^2 + v_2^2 \sigma_2^2 + 2R v_1 v_2$
- The Agent solves:

$$\max_{a_1, a_2} \{\widehat{w}(a_1, a_2)\}$$

- Let $R = 0$
- The FOCs are now:

$$v_1 = c_1 a_1 + \delta a_2$$

$$v_2 = c_2 a_2 + \delta a_1$$

- Using the FOC approach the Principal solves:

$$\max_{v_1,v_2,a_1,a_2} \{E[q-w] = a_1 + a_2 - t - v_1 a_1 - v_2 a_2\}$$
$$s.t.(i)\ \widehat{w}(a_1,a_2) = t + v_1 a_1 + v_2 a_2$$
$$-\frac{r}{2}\left(v_1^2\sigma_1^2 + v_2^2\sigma_2^2 + 2Rv_1v_2\right) \geq \overline{W}$$
$$(ii)\ \ v_1 = c_1 a_1 + \delta a_2$$
$$(iii)\ \ v_2 = c_2 a_2 + \delta a_1$$

- (i) must bind so we have:

$$\max_{v_1,v_2,a_1,a_2} \left\{ \begin{array}{c} a_1 + a_2 - \frac{r}{2}\left(v_1^2\sigma_1^2 + v_2^2\sigma_2^2 + 2Rv_1v_2\right) \\ -\frac{1}{2}\left(c_1 a_1^2 + c_2 a_2^2\right) - \delta a_1 a_2 \end{array} \right\}$$
$$s.t.\ v_1 = c_1 a_1 + \delta a_2$$
$$v_2 = c_2 a_2 + \delta a_1$$

- FOC1:

$$1 - r\sigma_1^2 v_1 c_1 - r\sigma_2^2 v_2 \delta - v_1 = 0$$

- $\Rightarrow$

$$v_1 = \frac{1 - r\sigma_2^2 v_2 \delta}{1 + r\sigma_1^2 v_1 c_1}$$

$$v_2 = \frac{1 - r\sigma_1^2 v_1 \delta}{1 + r\sigma_2^2 v_2 c_2}$$

- Solving simultaneously yields:

$$v_1 = \frac{1 + r\sigma_2^2(c_2 - \delta)}{1 + r\sigma_1^2 c_1 + r\sigma_2^2 c_2 + r^2\sigma_1^2\sigma_2^2(c_1c_2 - \delta^2)}$$

- and symmetrically for v_2

Results:

1. Go from $\delta = 1$ to $\delta = -1$ (ie. substitutes to compliments) and v_1, v_2 increase
2. When $\delta = 0$:

$$v_1 = \frac{1}{1 + r\sigma_1^2 c_1}$$

   which is simply the one-task case.

3. As $\sigma_2^2 \to \infty$ (task 2 is really hard to measure) then:

$$v_2 \to 0$$
$$v_1 \to \frac{r(c_2 - \delta)}{rc_2 + r^2\sigma_1^2(c_1c_2 - \delta^2)}$$

Put all the incentive on task 1.

6 Career Concerns

6.1 Overview

- Formal incentive schemes are not the only way of motivating people
- Takeovers, debt, product market competition, implicit contracts, labor market competition (ie. career concerns)
- Work hard – get a good reputation
- Fama (JPE, 1980): sort of claimed that CCs would lead to the first-best – a bit extreme

6.2 Holmstrom's Model

- Formal analysis developed by Holmstrom (Essays in Honor of Lars Wahlbeck '82, then reprinted in Restud in '99)
- 2 period version (the general case is quite impressive)
- Risk-neutral principal ("Employer") and a risk-neutral Agent ("Manager")
- $y_t = \theta + a_t + \varepsilon_t$
- $t \in \{1, 2\}$
- θ_t is the manager's ability
- a_t is her action
- ε_t is white noise
- Symmetric information other than effort observation (only M sees that) – in particular, M doesn't know her own ability so that contracting takes places under symmetric information
- $\theta \sim N(\bar{\theta}, \sigma_\theta^2)$
- $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$
- $\theta, \varepsilon_1, \varepsilon_2$ are independent
- M can move costlessly at the end of the period and there is a competitive market for M's services (same technology)
- Cost of effort $\psi(a), \psi'(a) > 0, \psi''(a)$ – and assume that $\psi(0) = 0$ and that $\psi'(0) = 0$
- Discount factor δ

- Market observes y_1 and y_2 but they are not verifiable – so can't contract on them
- Can only pay a fixed wage in each period
- With a one period model the reputation effect is absent – no incentive to work at all $\rightarrow$ get a flat wage and set $a_1 = 0 \Rightarrow y_1 = \theta + \varepsilon_1$
- Therefore $E[y_1] = E[\theta] = \bar{\theta}$
- Since there is perfect competition $w = \bar{\theta}$
- Take w_2 to be set by competition for M's services and note that $a_2 = 0$ because it is the last period

$$\begin{aligned} w_2 &= E[y_2 \mid \text{info}] \\ &= E[\theta \mid \text{info}] \\ &= E[\theta \mid y_1 = \theta + a_1 + \varepsilon_1] \end{aligned}$$

- Assume that the market has rational expectations about a_1
- Let a_1^* be the equilibrium value of a_1 (a Rational Expectations Equilibrium "REE")

$$\begin{aligned} w_2 &= E[\theta \mid \theta + a_1^* + \varepsilon] \\ &= y_1 - a_1^* \end{aligned}$$

- Update the prior such that:

$$E[\theta|(y_1 - a_1^*)] = \bar{\theta}\left(\frac{\sigma_\varepsilon^2}{\sigma_\theta^2 + \sigma_\varepsilon^2}\right) + (y_1 - a_1^*)\left(\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2}\right)$$

- Note the effect of the signal to noise ratio
- The first period problem for the Agent is:

$$\max_{a_1} \{w_1 + \delta E[w_2] - \psi(a_1)\}$$

- Which can be written as:

$$\max_{a_1} \left\{ w_1 + \delta\left(\bar{\theta}\left(\frac{\sigma_\varepsilon^2}{\sigma_\theta^2 + \sigma_\varepsilon^2}\right) + \left(\bar{\theta} + a_1 - a_1^*\right)\left(\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2}\right)\right) - \psi(a_1) \right\}$$

$$\max_{a_1} \left\{ \delta\left(a_1 - a_1^*\right)\left(\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2}\right) - \psi(a_1) \right\}$$

- The FOC is:

$$\delta\left(\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\varepsilon^2}\right) = \psi'(a_1) \tag{7}$$

- Increasing effort translates into an increased inference of agent talent

- In the FB $\psi'(a_1^{FB}) = 1$
- From (7) we know that $\psi'(a_1) < 1$ because of two things: (i) $\delta < 1$ and (ii) $\left(\frac{\sigma_\theta^2}{\sigma_\theta^2+\sigma_\varepsilon^2}\right) < 1$
- $\Rightarrow 0 < a_1^* < a_1^{FB}$
- The fact that even when the agent does nothing they are valuable in the second period prevents there being a backward induction unraveling – *but relies crucially on the additive technology*

1. $a_1^* \uparrow$ if σ_θ^2 high or σ_ε^2 low
2. Suppose that there are more periods: zero in the last period $\Rightarrow a_t \to 0$ and $t \to \infty$
3. Could also (as Holmström does) have ability getting shocked over time – need this to keep the agent working and get out of the problem in 2, above. In equilibrium the market knows how hard M is working – *disciplined with respect to the out of equilibrium moves, but no fooling in equilibrium*
4. Career concerns don't always help you - eg. in multi-tasking model the competitive labor market distorts the relative allocation of time
5. Gibbons & Murphy: looked at CEO incentive schemes - found more formal schemes later in career - empirical confirmation
6. People may work too hard early on: let $y_t = a_t + \theta + \varepsilon_t, t \in \{1,2,3\}, \varepsilon_1 \equiv 0, var(\varepsilon_2) > 0, var(\varepsilon_3) > 0$. The FOC for period 1 is $a_2 = a_3 = 0, \delta + \delta^2 = \psi'(a_1)$. The market learns about θ at the end of period 1. $\delta + \delta^2 > 1$ unless δ is smallish

6.3 Career Concerns with Multiple Tasks

- Consider an additive normal model as follows:

$$\begin{aligned} y_i &= \theta_i + a_i + \varepsilon_i \\ \theta_i &\sim N(\bar{\theta}, \sigma_\theta^2) \\ \varepsilon_i &\sim N(0, \sigma_\varepsilon^2) \end{aligned}$$

- $i \in \{1,2\}$
- Talents may be correlated, but the εs are *iid*
- Assume that the market cares about $\theta_1 + \theta_2$
- Define $\widehat{a} = a_1 + a_2$
- $(\theta_1 + \theta_2) \sim N(2\bar{\theta}, 2(1+\rho)\sigma_\theta^2)$ where ρ is the correlation coefficient between θ_1 and θ_2
- Note that $(\varepsilon_1 + \varepsilon_2) \sim N(0, 2\sigma_\varepsilon^2)$

- If the total cost of effort is $\psi(a)$ then we obtain the following FOC:

$$\psi'\left(a^{SB}\right) = \delta \frac{2(1+\rho)\sigma_\theta^2}{2(1+\rho)\sigma_\theta^2 + 2\sigma_\varepsilon^2}$$

- Note that a^{SB} increases with ρ (since an increase in ρ means that there is a higher signal to noise ratio because there is higher initial uncertainty about talent relative to pure noise)
- Implication for cluster of tasks among agents: one agent should be allocated a subset of tasks that require similar talents
- This is very different than under explicit incentives, where you increase effort by reducing uncertainty on talents and therefore uncluster tasks

6.4 Instrumental Career Concerns

- Imagine you are an advisor who has some information which is valuable to a decision maker
- Also imagine that you (A) and the decision maker (DM) have the same preferences, so no bias.
- A has an incentive to truthfully reveal her information.
- There would be an issue if DM thinks that you might be biased? How: if there is another type of A who doesn't have the same preferences as DM
- "Good" A doesn't like the idea of being thought to be biased.
- Note that this is not just shoved into the utility function! It is because it affects the degree to which her advice is listened to.
- Now you have an incentive to lie, for repetitional reasons
- So how much information ends up getting conveyed?
- Can be the case if if A is sufficiently concerned about her reputation, the no information is conveyed in equilibrium.
- Key paper here is Morris (JPE, 2001)–"Political Correctness"
- Two periods, 1 and 2.
- In period 1, DM's optimal decision depends on the state of the world which is either 0 or 1 (i.e. $\omega_1 \in \{0, 1\}$).
- Each state occurs with probability 1/2.
- This is DM's prior (and A's for that matter).
- A observes a signal about the true state, formally $s_1 \in \{0, 1\}$.)

- The signal is informative: with probability γ, the signal is equal to the true state (assume that $1/2 < \gamma < 1$.)
- With probability λ, A is "good" in the sense that she has the same preferences as DM. But with probability $1-\lambda$, A is "bad," in the sense that meaning that she is biased and just always wants DM to take the same decision (not matter what signal A got).
- After seeing the signal A can send a message $m_1 \in \{0, 1\}$.
- DM observes the message and then take an action $a_1 \in \mathbb{R}$.
- DM then sees the true state (and so does A)
- So DM will rationally update about the type of the advisor.
- If A told her the actual state then it is more likely that she is the good type.
- And the fact that A cares about how she is perceived is the whole mechanism which is interesting here.
- Period 2 is just like period 1, the state is ω_2, the message is m_2, and the signal is s_2.
- Signals are *iid* across periods.
- DM's preferences in period i are represented by the payoff function $V_{DM}^{i} = -(a_i - \omega_i)^2$, so she wants to "hit the state"
- To weight these across periods let the total payoff be $V_{DM} = -x_1(a_1 - \omega_1)^2 - x_2(a_2 - \omega_2)^2$.
- Good A has exactly these preferences.
- Bad A always wants action 1 to be taken: so that $V_{A,Bad} = y_1 a_1 + y_2 a_2$.
- Proceed by backward induction.
- Period 2: there is no more reputation building to be done, so A just focuses on her current goal–just a cheap talk game.
- There is a unique informative equilibrium in the second period of the game.
- Of course, there is also an uninformative one: a "babbling equilibrium" (Crawford and Sobel, 1982).
- Focus on the informative equilm: suppose DM learns something from the message she receives and chooses a higher action after (say) message 1
- Then bad A will want to announce 1 no matter what signal she got.
- Good A will want to truthfully report he signal, Because DM will choose a higher action if she hears 1 than if she announces 0–and this is better for DM (and hence good A).
- Can compute the optimal action of DM in period 2 (given the message and updating about the type of A), and then the value function for each type of A.

- Now work back to period 1.
- Again there is a babbling equilibrium.
- Main focus: do there exist equilibria in which good A truthfully reports.
- Suppose there is such an equilm. Does bad A want to tell the truth (ie. pool).
- No updating about types in such an equilm–no reputation building.
- But bad A wants to convince DM to take action 1, and if there's no reputational cost then bad A will lie to do that. Contradiction.
- Can show that bad A always announces 1 is she gets signal 1, and with probability v if she observes signal 0.
- And good A tells truth.
- So prob. that good A announces 1 when the true state is 1 is γ (pr. that she observed the accurate signal).
- Pr. bad A says 1 when $\omega = 1$ is $\gamma + (1-\gamma)v$, since she gets the high signal with pr γ, and with pr. $1-\gamma$ bad A observes 0 and announces 1 with pr v.
- So, Bayesian updating means that the probability that A is good, after getting message 1 and seeing observed first period state 1 is:

$$\Lambda(\lambda, 1, 1) = \frac{\lambda\gamma}{\lambda\gamma + (1-\lambda)(\gamma + (1-\lambda)v)}. \tag{8}$$

- Since $v > 0$ this probability is less than λ! ie. even though good A told the truth and was right, her reputation goes down!
- It can be shown that

$$\Lambda(\lambda, 0, 1) = \Lambda(\lambda, 0, 0) > \Lambda(\lambda, 1, 1) > \Lambda(\lambda, 1, 0). \tag{9}$$

- Both types of A have a reputational incentive to announce 0 independent of their signal (in fact, even if their signal was perfect!)
- This is true in any informative equilm.
- Good A tells the truth if the signal is zero. If the signal is 1 then there exists $\overline{x_1}$ such that for all $x_1 < \overline{x_1}$ she lies.
- Note that this leads to loss of information, but can't bias the decision–the expected value of DM's ex post belief is his ex ante belief (by the definition of conditional probability).
- If the second period is important enough (x_2 is large enough for a fixed x_1) then no information is conveyed in period 1.
- Can more repetition help? Can we get away from this conclusion that equilm is uninformative precisely when that period doesn't matter much?

7 Relational Contracts

7.1 Overview

- Even when some performance measures can be contracted on, subjective assessments of performance play an important role.
- FX traders generate fairly precise measures of performance (their books are marked to market at the end of each trading day), yet subjective bonuses are a significant part of their compensation–e.g. how well do they meet client demands, timing, etc.
- Lincoln Electric a famous example (most used HBS case ever): use a piece rate, but still about half of a typical worker's compensation is a subjective performance bonus.
- Professional services firms: put into one of several "buckets" based on subjective assessments.
- GE (under Jack Welch) formalized this with an: A,B,C,D system.
- Can think of lots of "every day" examples: promotions, compensation, etc.
- Not only within firms; also between firms
  - Alliances
  - Joint ventures
  - Supplier relationships
- Formal contracts can only be written ex ante about terms that can be verified ex post by a third party.
- A *relational contract* is a self-enforcing "contract".
- An idea which has been considered by economists and non-economists (eg. Macaulay, 1963 American Sociological Review; Klein-Leffler, 1981 JPE)
- Use the theory of repeated games and variations of the Folk Theorem to think about what relationships (relational contracts) are self-enforcing.
- Can be conditioned on non-verifiable outcomes and outcomes that may not be possible to be specified ex ante.
- Virtue: ability to adapt to new information as it becomes available, and ability to expand the range of "contractible" variables.
- Vice: (well, fragility) relies on subgame perfection
  - IR constraints need to be satisfied
  - More subtle: how much faith in rationality?
- Recall the finitely repeated prisoners' dilemma: unravels.
- Also recall the infinitely repeated prisoners' dilemma: if the discount factor is sufficiently high then (C,C) can be sustained as a SPE.
- Same principle at work here: if the future (i.e. relationship) is sufficiently valuable then the short-run temptation to cheat can be disciplined.

7.2 Observable Effort

- Start with the case where the agent's effort is observable to both P and A.
- This is basically Bull (1987) and Baker-Gibbons-Murphy (1994).
- Easy to think of this as P and A observing a noisy signal and basing a relational contract on that.
- Get's (much) trickier if want to base a relational contract based on a subjective assessment made by one player, but imperfectly observed by the other (see game theory literature on public versus private monitoring–e.g. Kandori-Matushima (ECMA, 1988)).
- Suppose one P(rincipal) and one A(gent).
- Each period A takes and action a which induces a probability distn over outcomes y.
- Let $y \in \{L, H\}$.
- Suppose $a \in [0,1]$ and $\Pr(y = H|a) = a$.
- y cannot be contracted on (maybe it's too complex to be described to a third party).
- Suppose that the compensation contract has a base salary s and a discretionary bonus b that P promises to pay A if $y = H$.
- Timing
  - P offers A a contract $w = (s, b)$
  - A accepts, or rejects and gets outside option $\bar{w}$
  - If A accepts then she chooses action $a \in [0,1]$ at private cost $c(a)$; (note that P does NOT observe a)
  - P and A observe y (in fact, it becomes common knowledge).
  - If $y = H$ then P chooses whether or not to pay b.
- P gets $y - w$
- P's discount rate is r
- A's discount rate doesn't matter, because it is only P who is trying to build a reputation
- A's payoff is $w - c(a)$, with $c(a)$ convex and satisfying $\lim_{a \to 1} \{c(a)\} = \infty$.
- In the FB we have

$$a^{FB} \in \arg\max_a \{L + a(H - L) - c(a)\}.$$

- Thus $c'\left(a^{FB}\right) = H - L$.
- If there was just one period then P won't pay a bonus (there's no point building a reputation) and so A will put in zero effort, and so $y = L$.

- If $L < \bar{w}$ then P, anticipating A's effort choice) will not pay $s > 0$, so A will reject the RC.
- Now consider the infinitely repeated game with discount rate r.
- Suppose both P and A play trigger strategies (cooperate if all previous play was cooperate, and if there was ever a defection then defect forever).
- For a given bonus b, then conditional on A believing that P will cooperate (honor the relational contract) A solves

$$\max_a \{s + ab - c(a)\}.$$

- The FOC is $c'(a^*) = b$.
- A's IR constraint is

$$s + a^*(b)\, b - c(a^*(b)) \geq \bar{w}.$$

- P rationally offers the minimum s that A will accept so that P's expected per-period payoff is

$$V(b) = L + a^*(b)(H - L) - c(a^*(b)) - \bar{w}.$$

- The key is whether A thinks P will honor the RC: if $y - H$ will P pay the discretionary bonus b?
- Fix A's strategy
  - If P does not pay b then she gets $H - s$ this period, and zero thereafter
  - If P does pay b then she gets $H - s - b$ this period, but gets the expected profit of the relationship thereafter
- So P pays the bonus IFF the FV of expected profit starting next period is bigger than the bonus, i.e.:

$$(H - s - b) + \frac{1}{r}V(b) \geq (H - s) + \frac{1}{r} \cdot 0,$$

- This is just

$$b \leq \frac{V(b)}{r}. \qquad (10)$$

- Thus, the optimal RC chooses b to maximize $V(b)$ subject to the constraint (10)–the reneging constraint.
- When r is low (P is patient/the interest rate is low) we can get the FB by setting $b = H - L$.
- For high interest rates P is not willing to pay any bonus, and so the relationship cannot be sustained.
- For intermediate interest rates we can't get the FB, but the future relationship is valuable enough to sustain some bonus and hence some effort from A.

7.3 Aside: Non-Conguent Performance Measures

- Hard to think of m(any) situation(s) where everything that the principal cares about is contractible–in the sense that it will be enforced by a court.
- Kerr (1975): "the folly of rewarding A, while hoping for B."
- A principal who failed to realize the non-congruence of what she can contract on and what she cares about would be making that error.
- Business history is littered with examples of principals who have committed this error.
  - Baker, Gibbons and Murphy (1994) point to H.J. Heinz Company rewarding division managers for increases in earnings over previous year earnings, which lead to manipulation of accounting variables through timing of booking costs and revenues.
  - At Dun & Bradstreet salespeople only received commissions if customers bought an expensive subscription to the firm's credit report services which led to salespeople fraudulently overstating their historical usage.
  - More recent examples of accounting manipulation include: Enron, Tyco and WorldCom, to name just three of the largest and most egregious cases.
- Suppose that the principal's benefit function b is given by $b = f_1 a_1 + f_2 a_2 + \epsilon$.
- But this cannot be contracted on.
- What can be contracted on is a performance measure $p = g_1 a_1 + g_2 a_2 + \phi$.
- Assume linear contracts of the form $w = t + vp$.
- Risk neutral P and A
- For simplicity assume that $E(\epsilon) = E(\phi) = 0$ and that $\psi(a_1, a_2) = \frac{1}{2}\left(a_1^2 + a_2^2\right)$.
- Agent therefore solves the following problem

$$\max_{a_1, a_2} \left\{ t + v\left(g_1 a_1 + g_2 a_2\right) - \frac{1}{2}\left(a_1^2 + a_2^2\right) \right\}.$$

- The agent's optimal actions are $a_1^*(v) = g_1 v$, and $a_2^*(v) = g_2 v$.
- What is the optimal level of v–i.e. what is the optimal intensity of incentives? Given the we have just worked out the optimal actions for the agent we can deduce that the principal's expected payoff is

$$E[b - w] = f_1 a_1^*(v) + f_2 a_2^*(v) - t - v g_1 a_1^*(v) - v g_2 a_2^*(v).$$

- A's expected payoff is

$$\begin{aligned} E[U] &= E[w] - \psi(a_1, a_2) \\ &= t + v\left(g_1 a_1^*(v) + g_2 a_2^*(v)\right) - \frac{1}{2}\left(a_1^*(v)^2 + a_2^*(v)^2\right). \end{aligned}$$

- Total surplus (the sum of the principal's and agent's payoff) is then

$$E[y] - \psi(a_1, a_2) = f_1 a_1^*(v) + f_2 a_2^*(v) - \frac{1}{2}\left(a_1^*(v)^2 + a_2^*(v)^2\right).$$

- It is easy to show that the optimal choice of v is then

$$v^* = \frac{f_1 g_1 + f_2 g_2}{g_1^2 + g_2^2}. \tag{11}$$

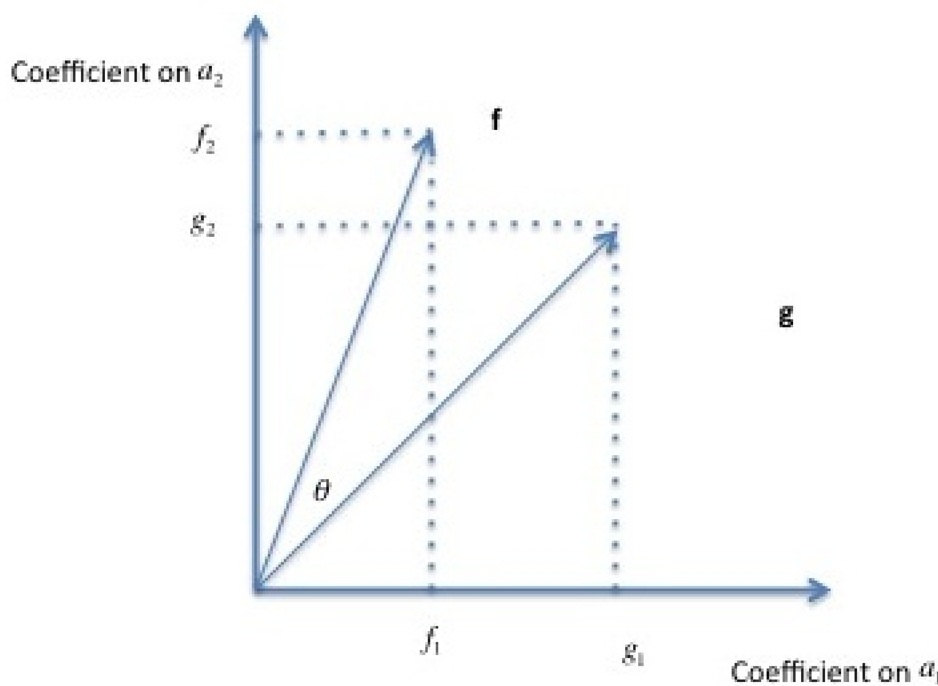

- First notice the importance of *scale*
  - Suppose both g_1 and g_2 were much larger than f_1 and f_2.
  - The agent could cause measurable performance, p, to go up by putting in lots of effort (high a_1 and a_2).
  - But this would lead to a much lower expected benefit for the principal.
  - So the efficient contract makes v–the incentive to make p high–very low.
- The second thing to note is *alignment*.
  - Suppose θ is small.

- Then a high value of v (powerful incentives) has a great effect in increasing b, the benefit to the principal.
- Now suppose that θ is large, so that $\mathbf{f}$ and $\mathbf{g}$ are poorly aligned.
- In this case incentives for increasing the performance measure (high v) are not very helpful for increasing b.
- Indeed, in the extreme case where the f and g vectors are orthogonal, incentives are useless in increasing b.

- With a little middle-school math we can derive

$$v^* = \frac{\sqrt{f_1^2 + f_2^2}}{\sqrt{g_1^2 + g_2^2}} \cos(\theta), \tag{12}$$

where θ is the angle between the f and g vectors.

- What makes for a good performance measure?
- It is tempting to say that if p is highly correlated with b then it is a good performance measure.
- One way in which this can be the case is if the noise terms, ϵ and ϕ are highly correlated.
- Does this make p a good performance measure?
- No. Imagine that p was the stock price and b was fundamental value. It's plausible to think that shocks to both are highly correlated, but there could be a wide divergence between the two measures if markets are not efficient in the short-run (as much evidence suggests is often the case).
- Though highly correlated with true value, the stock price would not be a good performance measure.
- A stark way to see this is to suppose that $p = a_1 + \epsilon$ and $b = a_2 + \epsilon$.
- The shocks are perfectly correlated, but p is a terrible performance measure.
- Put a different way, the only way to judge what is a good performance measure is through the actions it induces. This is what I call the *Gibbons tautology*: a performance measure is valuable if and only if it induces valuable actions. In other words, alignment is more important than noise.

7.4 Mixing Relational and Explicit Contracts

- Based on Baker-Gibbons-Murphy (QJE, 1994).
- Setup basically as above, but let $H = 1$ and $L = 0$.
- y cannot be contracted on.
- A's action also affect a contractible performance measure p (that is also 0 or 1).

- Before taking a, A privately observes μ.
- $\Pr(p = 1) = \mu a$, and p and y are conditionally independent, and also assume that $E[\mu] = 1$ (on avg. p is an unbiased measure of μ).
- The implicit bonus is b and the explicit bonus is β that is paid when $p = 1$.
- Timing as above (with the objective bonus being paid after output is realized).
- Let $c(a) = \gamma a^2$.
- The FB has $a^{FB} = 1/2\gamma$.
- If A believes P will honor the implicit contract A solves

$$\max_a \left\{ s + ab + \mu a\beta - \gamma a^2 \right\}.$$

- The FOC is

$$a^* = \frac{b + \mu\beta}{2\gamma}.$$

- So we know that $a^* < a^{FB}$ if $b + \mu\beta < 1$.
- A will work for the firm if IR(A) is satisfied (before seeing μ.)
- So P's expected per-period profit , given the relational contract embodied in b and the explicit contract in β is:

$$E_\mu \left(a^*(\mu, b, \beta) - [s + a^*(\mu, b, \beta)b + \mu a^*(\mu, b, \beta)\beta] \right).$$

- P lowers s to make IR(A) binding (due to separability), so we have expected profit function for P:

$$V(b, \beta) = E_\mu \left(a^*(\mu, b, \beta) - \gamma a^*(\mu, b, \beta)^2 - \overline{w} \right)$$

- Now combine explicit and implicit
- At the end of each period P and A observe the objective performance measure p and also y.
- If $p = 1$ P pays bonus β according to the explicit contract and if $y = 1$ P chooses whether to pay A the bonus b specified in the RC.
- Before, P's expected profit if it reneged was zero in all future periods, because the trigger strategy implied that A wouldn't work for P.
- This is different when both explicit and implicit measures are available.
- One thing is that the expected per-period profit from honoring the RC is not $V(b)$ but $V(b, \beta)$–see above for the equations.
- Second thing, if P reneges, A would not accept any future RCs, but might want to take an explicit contract (if it was sufficiently remunerative).

- Without RCs, the per-period expected profit from (optimal) explicit contracting is $V(\beta^*)$ can be positive or negative.
- Depends on A's outside option $\overline{w}$ and how distorted the performance measure is $(var(\mu))$.
- If $V(\beta^*) > 0$–so explicit contracting can satisfy IR(A) and generate positive profits for P, the outside option for P if she reneges on the RC is $V(\beta^*) > 0$.
- But if $V(\beta^*) < 0$, the P's outside option is to shutdown and get zero.
- These outside options are important for understanding the optimal mix of implicit and explicit contracts.
- The reneging constraint, if $V(\beta^*) > 0$ is:

$$\frac{V(b, \beta) - V(\beta^*)}{r} \geq b \to V(b, \beta) - V(\beta^*) \geq rb.$$

- In the case where $V(\beta^*) > 0$ it can be shown (by solving for the optimal contract) that RCs cannot be (optimally) used when the discount rate is sufficiently high or the amount of distortion between the performance measure and output is sufficiently low.
  - If the performance measure is v. good then P's outside option is also v. good–so the temptation to renege is high.
  - A remarkable fact is that under parameter values that mean that an RC alone could get the FB, the very possibility of an imperfect explicit contract could render the RC infeasible.
- Also, we can get FB if the discount rate is sufficiently low.
- But even for a low discount rate FB cannot be achieved if the performance measure is nearly perfect, so the outside option in close to FB.
- Now suppose that $V(\beta^*) < 0$ so that P's outside option after reneging on an RC is to get zero profit by shutting gown the firm.
- This happens (naturally) when distortions caused by a "bad" performance measure are high.
- But could think of other ways this could happen (outside the model)
  - There are multiple implicit contracts, and destroying the reputation for one destroys all RCs
  - There are multiple recipients of RCs and ratting on one destroys the other RC
  - ...
  - Highlights the importance of beliefs about what the implicit contracting variables actually are
- Now the reneging constraint is not $V(b, \beta) - V(\beta^*) \geq rb$, but rather $V(b, \beta^*) \geq rb$.
- One can again solve for the optimal contract.

- A striking result is that as the performance measure becomes less distortionary, the implicit bonus *increases*, but the effect on the explicit bonus is ambiguous.
  - Suppose P and A have an optimal RC, but with no explicit contract.
  - Suppose also that the discount rate is high enough that the reneging constraint binds.
  - Now drop in a contractible performance measure, but one that is bad enough that it couldn't support a profitable explicit contract without an RC in the background.
  - It can still be useful.
  - A low-powered explicit contract can improve the future value of the relationship, so the reneging constraint no longer binds.
  - So the implicit bonus in the RC can be made larger.
- Here, implicit and explicit incentive are complements.
- Analogous result holds for *better* performance measures, not just the introduction of a new one.

7.5 Levin's Model

- Consider a sequence of spot contracts between a principal (P) and agent (A)
- Assume both are risk-neutral
- Assume both have common discount factor $\delta < 1$
- Let per period reservation utilities be $\bar{V}$ and $\bar{U}$ for P and A respectively and let $\bar{s} = \bar{V} + \bar{U}$
- A chooses action $a \in A$
- Output levels $q_1 < ... < q_n$
- Probability of these is $\pi_i(a)$ (just like in Grossman-Hart, where π is a mapping from A to the probability simplex)
- Denote action in period t as a_t
- Assume $\pi_i(a) > 0$ for all i and that MLRP holds
- Payment from P to A in period t is $I_t = w_t + b_t$ (interpreted as wage plus bonus)
- P's per period payoff is $q_i^t - I_t$
- A's per period payoff is $I_t - \psi(a_t, \theta_t)$, where θ_t is a cost parameter which is private information
- Let $\theta_t \in \{\theta_L, \theta_H\}$ with $\theta_L \leq \theta_H$
- Assume that these are iid over time and let $\beta = \Pr(\theta_t = \theta_H)$

- Assume ψ is convex, increasing and that $\psi(0,\theta)=0$, that $\psi_\theta(\cdot)>0$ and $\psi_{a\theta}(\cdot)>0$, where subscripts denote partial derivatives

- First best in a given period solves

$$\max_{a\in A}\left\{\sum_{i=1}^{n}\pi_i(a)\,q_i-\psi(a,\theta)\right\}$$

- Let

$$a^{FB}(\theta)=\arg\max_{a\in A}\left\{\sum_{i=1}^{n}\pi_i(a)\,q_i-\psi(a,\theta)\right\}$$

  and assume uniqueness

- Also assume

$$\sum_{i=1}^{n}\pi_i\left(a^{FB}\right)q_i-\psi\left(a^{FB},\theta\right)>\bar{s}$$

- Otherwise it would not be an interesting problem.

- Consider the game where each period the players choose whether or not to participate, A chooses an action and P chooses an output contingent bonus payment $b_t\left(q_i^t\right)$

Definition 3. *A Relational Contract is a perfect Bayesian equilibrium of the above game.*

- Let σ^A and σ^P be the strategy A and P respectively

- These are a function of observed history of play and output realizations

- Not contingent on A's action because it is not observable to P, and is sunk for A, and past actions don't affect P's continuation play.

- Assume that output realizations are observable but *not* verifiable

- Assume that past payments are observable *and* verifiable

- Let ζ^w be flow payoffs from verifiable components and ζ^b be from non-verifiable components

- ζ^b is the self-enforced part and it specifies a bonus payment $b_t(h_t)$, where h_t is the history of play and output realizations up to t

Definition 4. *We say that a relational contract is Stationary if in every period* $a_t=a(\theta_t)$, $b_t=b\left(q_i^t\right)$ *and* $w_t=w$ *on the equilibrium path.*

- Levin (2003) proves that one can restrict attention to stationary contracts wlog

  - Basic argument is that for any set of non-stationary transfers and actions one can find a stationary contract with the same payoffs

  - Can't get joint punishment with a stationary contract–but it turns out that when P's behavior is observable, optimal contracts don't involve joint punishment in equilibrium

- Fix a relational contract $\left(\sigma^A, \sigma^P, \zeta^w, \zeta^b\right)$ and let $\hat{u}$ be A's payoff under this contract and $\hat{s} - \hat{u}$ be P's payoff
- Similarly, let $\hat{w}$ be the wage (which is court enforcable), $\hat{b}\left(q_i\right)$ be the bonus payment under this contract, and $\hat{a}\left(\theta\right)$ be A's action
- Joint value is then given by the program

$$\hat{s} = \max_{a(\theta)} \left\{(1-\delta) E_{\theta,q}\left[q - \psi\left(a\left(\theta\right), \theta\right) | a\left(\theta\right)\right] + \delta E_{\theta,q}\left[\hat{s}|\hat{a}\left(\theta\right)\right]\right\}$$

subject to

$$a\left(\theta\right) \in \arg\max_{a\in A}\left\{E_q\left[\hat{w} + \hat{b}\left(q_i\right) + \frac{\delta}{1-\delta}\hat{u}|a\right] - \psi\left(a,\theta\right)\right\} \tag{IC}$$

$$\hat{b}\left(q_i\right) + \frac{\delta}{1-\delta}\hat{u} \geq \frac{\delta}{1-\delta}\bar{U} \tag{IR-A}$$

$$-\hat{b}\left(q_i\right) + \frac{\delta}{1-\delta}\left(\hat{s} - \hat{u}\right) \geq \frac{\delta}{1-\delta}\bar{V} \tag{IR-P}$$

- We are assuming that when A leaves the relationship she leaves forever (this is the strongest threat she has an gives rise to the largest set of relational contracts)–so no loss of generality.
- IR-P says P is willing to make the promised bonus payments (the RHS is what she gets if she doesn't pay and hence A quits forever; LHS is future payoff minus the bonus payment.)
- The contract which solves the program constitutes a PBE
  - If P doesn't participate at some point then P's best response is to not participate as well–and vice versa
- What about renegotiation?
  - Stationary contracts can be made renegotiation proof
    * If renegotiation precluded enforcement of inefficient punishments such that $s_\tau < \hat{s}$ after a deviation in period $t < \tau$, and RC can still be enforced with jointly efficient punishments by changing the split of the surplus.
    * P promises $\hat{w}$ and $\hat{b}(q_i)$ unless someone deviates.
    * If A deviates we move to a new contract with payments $w + b(q_i)$ such that $u = \overline{U}$ and $V = \hat{s} - u$ following the deviation.
    * If P deviates the contract has payments such that P gets $\overline{V}$.
    * If more deviations occur then do similarly.
    * Has to be renegotiation-proof because it's on the (constrained) Pareto frontier.
- What about existence?
  - It can be shown that a solution exists

- When θ and q come from finite sets one can use the standard arguments from static Moral Hazard (e.g. Grossman-Hart 1983) and Adverse Selection (e.g. Mirrlees)

- Bonus payments can be positive or negative depending on how the surplus needs to be shared
  - If P gets "a lot" of the surplus then bonuses are positive–looks like incentive pay
  - Need to give big bonuses to satisfy IR-A when $\hat{u}$ is close to $\bar{u}$
  - If A gets "a lot" of the surplus then bonuses are negative–looks like efficiency wages

- Let $\bar{b}$ and $\underline{b}$ be the highest and lowest bonuses

- Then IR-A and IR-P combine to give the "self-enforcement constraint"

$$\left(\bar{b}-\underline{b}\right) \leq \frac{\delta}{1-\delta}\left(\hat{s}-\bar{s}\right)$$

- Can now compare relational contracts to contracts contractible output in the case of moral hazard

- Moral hazard (with no adverse selection) has $\theta_L = \theta_H = \theta$ which is common knowledge

- Risk-neutral P and A so optimal contract involves making A the residual claimant

- The payment scheme is

$$I = q_i + \bar{u} - \max_{a \in A}\left\{E_q\left[q|a\right] - \psi\left(a,\theta\right)\right\}$$

- This will violate the self-enforcement constraint if

$$\left(q_n - q_1\right) > \frac{\delta}{1-\delta}\left(E_q\left[q|a^{FB}\right] - \psi\left(a^{FB},\theta\right) - \bar{s}\right)$$

- It can be shown that when this is violated the optimal relational contract is of the following form

$$\begin{aligned} b\left(q_i\right) &= \bar{b} \text{ for } q_i \geq q_k \\ b\left(q_i\right) &= \underline{b} \text{ for } q_i < q_k \end{aligned}$$

  where q_k is some interior cutoff value

- MLRP important here

- Connection to one period moral hazard with risk-neutrality and limited liability constraint (Innes, 1990)–SEC constraint acts like LL constraint.

- Can also apply the model to the case of pure adverse selection
  - That corresponds to a being observable to P and A, but θ being A's private information

- Note that Spence-Mirrlees condition is satisfied because we assumed that $\psi_{a\theta} > 0$. That condition is
$$\frac{\partial}{\partial \theta}\left(\frac{\partial[I - \psi(a,\theta)]}{\partial a}\right) < 0.$$
- Can be shown that the no distortion for the highest type no longer applies in the relational model–second-best enforceable contracts can violate SEC
  - The bonus payments in the court enforceable model can violate the self-enforcement constraint
  - So all types under-provide "effort"
  - Also get bunching with many types
- A general point–the self-enforcement constraint lowers the power of the incentives that can be provided (in either setting)
- Can also extend the model (as Levin does) to *subjective* performance measures
  - Stationary contracts now have problems
  - But the optimal contract is still quite simple
  - P pays A a base salary each period, and then a bonus if P (subjectively) judges performance to be above a threshold
  - But if below threshold then the relationship terminates
  - Inefficiency can come from the different beliefs about performance
  - So a mediator can be thought of as making the information more objective and therefore reducing the welfare loss
  - Can do better by making evaluation less frequent–can allow P to make more accurate assessments

7.6 Building Routines

- Based on Sylvain Chassang (AER, 2010).
- Basic question: suppose the details of cooperation are not common knowledge, how do players learn those details.
- Specifically: repeated game setting where one player has incomplete information about the timing and ability of the other player to affect her outcome.
- Imperfect monitoring, that at the start requires punishments that are inefficient.
- But as the common history grows there is a reduction in inefficiency–interpret as "better routines".
- Think back to Bull (1987)–we made relational contracts sound rather straightforward (e.g. subjective bonuses)
- How do parties specify what the contingencies *are* in a relational contract?!

- Link to private monitoring literature (Green-Porter 1984, Abreu-Pearce-Stachetti 1990, Fudenberg-Levine-Maskin 1994: all ECMA). See Kandori (JET, 2002) for an excellent overview: http://personal.lse.ac.uk/zapal/EC501$_2$008$_2$009/$Skrzypacz_background1.pdf$
    - 2 players in a repeated game.
    - Each period:
        - Player 1 decides whether to stay in the relationship and interact with player 2, or skip one period.
        - If player 1 stays, then player 2 gets to take an action.
        - The available actions are randomly drawn from some subset of the total action set.
        - Two types of actions: *productive actions* that give player 1 a benefit with positive probability, but also fail with some probability, and *unproductive actions* that cost player 2 zero, but do nothing for player 1.
        - At the end of each period the: set of available actions, the action P2 took, and the benefit P1 got are commonly observed.
        - Player 2 knows which actions are productive, but player 1 does not.
    - The informational asymmetry means that P1 lacks the ability to interpret what the actions that P2 has available and took mean for payoffs.
    - Wedge b/w the availability of information and the ability to interpret it.
    - Basic idea: P1 doesn't know which actions are productive and thus when she should expect P2 to take a productive action (i.e. cooperate).
    - So if P1 sees an action that has no benefit for her she doesn't know whether if P2 took a productive action and it failed, or if P2 chose and unproductive action (at zero cost to P2).
    - So P1 might use inefficient exit in equilm (c.f. Green-Porter) to support cooperation.
    - But once an action (say a_0) yields a benefit to P1, P1 knows that it's a productive action–so the monitoring problem goes away.
    - Tradeoff between new productive actions (virtue: more beneficial, vice: monitoring problem) and existing actions that are known to be productive (virtue: no monitoring problem, vice: not as productive).
    - Focus on Pareto efficient, pure strategy PBEa.
    - Main results
        - During the "specification phase" the relationship is sensitive to shocks.
        - Indeed, during this phase, a P2 action that fails is followed by punishment from P1 on the equilm path.
        - But once the learning phase is done (we're in a routine) this is no longer the case.
        - Can get path-dependence: costs of information revelation can be bigger than the efficiency gains from using a better routine.
        - Ex ante identical partnerships can have different long-run outcomes because of initial shock realizations.

8 Markets and Contracts

8.1 Overview

- A lot of what we have done thus far considers bi-lateral (or sometimes multilateral) relationships
- But in some/many contexts, contracts between agents exist in market settings
- This has been recognized for a long time–Rothschild and Stiglitz (1976) analyze screening in such a context
- But there are a number of other issues of interest
- We will only touch on a few of them here
  - Contracts as a Barrier to Entry
  - Product Market Competition
  - Equilibrium Incentive Contracts

8.2 Contracts as a Barrier to Entry

- There is a long traditional in legal scholarship/law and economics which argues that contracts can be anti-competitive in effect
- Sellers may be able to "lock up" buyers with long-term contracts which prevent or at least deter entry to some degree
- Key reference is Aghion and Bolton (1987)
- Contracts that specify penalties for early termination can be used to extract rents from future entrants who may be lower cost than the current provider
- Suppose there are two time periods $t = 1$ and $t = 2$
- At $t = 1$ there is an incumbent who can sell a product at cost $c_I \leq 1/2$ and a buyer has reservation value $v = 1$ for this widget
- At $t = 2$ a potential entrant has cost c_E which is uniformly distributed on $[0, 1]$
- Obviously $p_1 = 1$ in period 1
- Assume that if entry occurs there is Bertrand competition at $t = 2$
- So entry occurs if $c_E \leq c_I$
- If there is no contract / a spot contract then if entry occurs $p_2 = \max\{c_E, c_I\} = c_I$ and if no entry then $p_2 = 1$
- So under the spot contract the expected payoff of the buyer is

$$
\begin{aligned}
V_B &= (1 - \Pr(entry))\, 0 + \Pr(entry)(1 - c_I) \\
&= c_I\,(1 - c_I)
\end{aligned}
$$

- And the incumbent firm's payoff is

$$\begin{aligned} V_I &= p_1 - 1 + (1 - \Pr(entry))\,(1 - c_I) + \Pr(entry)\,(1 - c_I) \\ &= 1 - c_I + (1 - c_I)^2 \end{aligned}$$

- Now consider the case where the incumbent and the buyer sign a contract at $t = 1$ which specifies as price for each period and a penalty d for breach / termination
  - The contract is a triple (p_1, p_2, d)
- So the buyer will only breach the contract if the entrants price p_E is such that

$$1 - p_E \geq 1 - p_2 + d$$

i.e. surplus under the new contract compensates for the surplus under the old including damages

- The probability of entry given this contract is

$$\Pr(c_E < p_2 - d) = p_2 - d$$

- The buyer's expected payoff under the contract is

$$\begin{aligned} V_B^L &= (1 - p_1) + (1 - p_E) + d \\ &= (1 - p_1) + (1 - (p_2 - d)) + d \\ &= (1 - p_1) + (1 - p_2) \end{aligned}$$

- The incumbent's expected payoff is

$$\begin{aligned} V_I^C &= p_1 - c_I + (1 - \Pr(entry))\,(p_2 - c_I) + \Pr(entry) d \\ &= p_1 - c_I + (1 - p_2 + d)\,(p_2 - c_I) + (p_2 - d)\,d \end{aligned}$$

- The buyer will only accept the contract if

$$(1 - p_1) + (1 - p_2) \geq c_I\,(1 - c_I)$$

- So the incumbent solves

$$\max_{p_1, p_2, d} \{p_1 - c_I + (1 - p_2 + d)\,(p_2 - c_I) + (p_2 - d)\,d\}$$
$$\text{subject to}$$
$$(1 - p_1) + (1 - p_2) \geq c_I\,(1 - c_I)$$

i.e. maximize the payoff under the contract subject to the buyer being willing to accept

- The incumbent can always set $p_1 = 1$, so the problem is

$$\max_{p_2,d} \{1 - c_I + (1 - p_2 + d)(p_2 - c_I) + (p_2 - d)\, d\}$$
$$\text{subject to}$$
$$(1 - p_2) \geq c_I (1 - c_I)$$

- Noting that the constraint binds we have $1 - c_I (1 - c_I) = p_2$
- So the program is

$$\max_d \{1 - c_I + (1 - (1 - c_I (1 - c_I)) + d)((1 - c_I (1 - c_I)) - c_I) + ((1 - c_I (1 - c_I)) - d)\, d\}$$

- The solution is
$$d^* = \frac{1 + (1 - c_I)(1 - 2c_I)}{2} > 0$$
- So the probability of entry is
$$p_2 - d^* = \frac{c_I}{2}$$
- The incumbent always wants to sign the contract
- This contract is competition reducing since the probability of entry is $\frac{c_I}{2}$ instead of c_I
- Markets with contracts may not be as efficient as spot contract markets!
- Robust to certain extensions
  - Renegotiation
  - Multiple buyers

8.3 Multiple Principal-Agent Dyads

- We have focused on a number of models of a single principal and a single agent.
- Situations with several principals but one agent are quite prevalent (e.g. wholesale trade, politics).
- This is referred to as *common agency*–see Bernheim-Whinston (ECMA, 1986).
- Won't focus on that here.
- Will focus on there are multiple P-A dyads.
- Also important for understanding oligopoly–since much (most?) competition takes place between "managerial firms".
- Do contracts between principals and agents affect the nature of competition, and are there strategic effects?
- Key early paper is Fershtman-Judd (AER, 1987).
- Suppose that firm i consists of a shareholder (principal) and a manager (agent).

- Both are risk-neutral.
- Let output be q_i and profits be Π_i.
- Time 1: each dyad contracts simultaneously on a linear incentive scheme of the form $w_i = \alpha_i q_i + \beta_i \Pi_i$.
- Time 2: each dyad's contract becomes public knowledge and then the firms compete in the product market: (differentiated products) Bertand or Cournot.
- Under Cournot competition there is strategic substitutability (reaction functions slope down).
- Would like to commit to being tough because this induces competitors to act softer.
- One way to do this is to provide incentives for *output*–i.e. $\alpha_i > 0$.
- Under Bertrand competition there is strategic complementarity, so it is optimal to set $\alpha_i < 0$.
- $\alpha_i = 0$ when there is a product market in which there is no substitutability or complementarity.
- Issues:
  - Predictions depend crucially on the nature of product market competition–easy to write down in models, but harder to determine empirically. Is the world Cournot?
  - Renegotiation: each dyad would like to fool other dyads into believing that incentives for output are there, but they are not profit maximizing conditional on that belief.
  - Secret renegotiation with symmetric information implies $\alpha_i = 0$ for all i.
- Can introduce asymmetric information within each dyad to prevent perfect renegotiation (link to renegotiation in classic P-A model: Fudenberg-Tirole (ECMA, 1990), Hermalin-Katz (ECMA, 1991).
- Caillaud-Julien-Picard (ECMA, 1995) model secret renegotiation in this setting (and agents have different cost types, not known at the time of contracting).
- Two firms.
- Time 1: initial contract about R. Time 2: secret renegotiation and contract offered on output of the two firms $t(a, b)$. Time 3: agent learns type and output occurs (or agent can leave the market).
- Roughly, they find that under Cournot competition signing an initial contract with some unconditional payment R, is optimal.
- Equilm initial contract gives agent rents that limit the need to reduce output below the full-information Cournot level.
- This shifts the firm's reaction function up and hence lowers the other firm's output.

- Under differentiated products Bertrand this is not true.
- Strategic effect of initial contract R is bad–makes the other firm tougher.
- So no initial contract.
- Basic point: under Cournot the initial contract relaxes the IR constraint in an adverse selection setting.

8.4 Product Market Competition and the Principal-Agent Problem

- Classic question: does product market competition increase internal efficiency of the firm?
- Leibenstein (1967): internal firm inefficiency–"X-Inefficiency"–may be very large
- Does competition help?
- Hicks (1935): "The best of all monopoly profits is a quiet life"
- First formal model is Hart (1983)–satisfising behavior
- Scharfstein (1987) with Hart's model but different utility function obtains opposite conclusion
- Martin (1993)–Cournot competition means less effort
- Many others–see Holden (2005) for references
- Will focus on three models due to Schmidt (1997)
- Look at these through the lens of Holden (2005) framework
- There are two players, a risk-neutral principal and a risk-averse agent
- Let $\phi \in \mathbb{R}$ be a measure of product market competition which affects the profits which accrue to the principal.
- A higher value of ϕ means that, all else equal, profits are lower.
- Suppose that there are a finite number of possible gross profit levels for the firm. Denote these $q_1(\phi) < ... < q_n(\phi)$.
- These are profits before any payments to the agent.

Definition 5. *A set X is a* ***product set*** *if $\exists$ sets $X_1, ..., X_n$ such that $X = X_1 \times ... \times X_n$. X is a* ***product set in*** *$\mathbb{R}^n$ if $X_i \subseteq \mathbb{R}, i = 1, ..., n$.*

- The set of actions available to the agent, A, is assumed to be a product set in $\mathbb{R}^n$ which is closed, bounded and non-empty.
- Let S be the standard probability simplex, i.e. $S = \{y \in \mathbb{R}^n | y \geq 0, \sum_{i=1}^n y_i = 1\}$ and assume that there is a twice continuously differentiable function $\pi : A \to S$. The probabilities of outcomes $q_1(\phi), ..., q_n(\phi)$ are therefore $\pi_1(a), ..., \pi_n(a)$.

- Let the agent's von Neumann-Morgenstern utility function be of the following form:

$$U(a, I) = G(a) + K(a)V(I)$$

  where I is a payment from the principal to the agent, and $a \in A$ is the action taken by the agent.

Definition 6. *An* ***incentive scheme*** *is an n-dimensional vector* $\mathbf{I} = (I_1, ..., I_n) \in \mathcal{I}^n$.

- Given an incentive scheme the agent chooses $a \in A$ to maximize her expected utility $\sum_{i=1}^{n} \pi_i(a) U(a, I_i)$.
- Key condition for increase in product market competition to decrease agency costs is

$$\sum_{i=1}^{n} q_i'(\phi)\pi_i'(a) \geq 0, \forall a, \phi. \tag{13}$$

- When MLRP holds this become

$$\sum_{i=j+1}^{n} \pi_i'(a)q_i'(\phi) \geq \sum_{i=1}^{j} |\pi_i'(a)| \, q_i'(\phi). \tag{14}$$

Schmidt's Basic Model

- The firm goes bankrupt if realized profits are below a certain level
- Reduced form measure of product market competition, ϕ
- An increase in ϕ corresponds to a more competitive product market
- Effort by the agent affects costs
- Two possible states: high cost and low cost–states L and H
- (14) becomes:

$$\pi_L'(a) \left[q_L'(\phi) - q_H'(\phi)\right] > 0 \tag{15}$$

- By FOSD $\pi_L'(a) > 0$ (a harder action makes the low cost state more likely)
- Schmidt's result requires $q_H'(\phi) < q_L'(\phi)$
- True because loss on the agent of $\overline{L}$ if the firms goes bankrupt
  - Occurs with positive probability in the high cost state and with zero probability in the low cost state
  - He assumes that the probability of this occurring is $l(\phi)$ with $l'(\phi) > 0$
  - This loss of $\overline{L}$ is equivalent to profits being lower since it affects the agent's utility and hence the payment that the Principal must make if the participation constraint binds
  - In effect, then $q_H(\phi) \equiv \overline{q}_H(\phi) - l(\phi)\overline{L}$

- – Schmidt's main result states that the increase in agent effort is unambiguous if the PC binds
  – In such circumstances $q'_L(\phi) > q'_H(\phi)$, since the expected loss of $\mathbb{E}[L]$ occurs only in state H
  – If the PC is slack at the optimum then the effect of competition is ambiguous because the loss of L is only equivalent to profits being lower if L is sufficiently large
  – Thus, for $\overline{L}$ sufficiently small we have $q'_L(\phi) = q'_H(\phi)$ and hence the condition is not satisfied.

Schmidt's Price-Cap Model

- Now consider price-cap regulation of a monopolis
- Firm can have constant marginal cost of either c^L or $c^H > c^L$
- Regulator does not observe costs, but sets a price cap of $1/\phi$
- Larger value of ϕ interpreted as a more competitive product market.
- Denoting demand at the cap (which is assumed to be binding regardless of the cost realization) as $D(1/\phi)$, profits are:

$$q(c^j, \phi) = D\left(\frac{1}{\phi}\right)\left(\frac{1}{\phi} - c^j\right)$$

- Differentiating with respect to ϕ yields:

$$\frac{\partial q(c^j, \phi)}{\partial \phi} = -\frac{1}{\phi^2}\left[D\left(\frac{1}{\phi}\right) + D'\left(\frac{1}{\phi}\right)\left(\frac{1}{\phi} - c^j\right)\right]$$

- General condition for a harder action in this two outcome model is simply:

$$\pi'_L(a)\left[q'_L(\phi) - q'_H(\phi)\right] \geq 0$$

- Since $\pi'_L(a)$ is positive, we require $q'_L(\phi) - q'_H(\phi) \geq 0$ – i.e. $q'_L(\phi) \geq q'_H(\phi)$. This requires:

$$-\frac{1}{\phi^2}\left[D\left(\frac{1}{\phi}\right) + D'\left(\frac{1}{\phi}\right)\left(\frac{1}{\phi} - c^L\right)\right] \geq$$
$$-\frac{1}{\phi^2}\left[D\left(\frac{1}{\phi}\right) + D'\left(\frac{1}{\phi}\right)\left(\frac{1}{\phi} - c^H\right)\right]$$

- which reduces to requiring:

$$\frac{(c^L - c^H)D'\left(\frac{1}{\phi}\right)}{\phi^2} \geq 0$$

  Obviously $D'\left(\frac{1}{\phi}\right) < 0$, and, by construction, $c^H > c^L$.

- A tighter price cap leads to a harder action by the agent.

8.4.1 Equilibrium Effort Effects

Definition 7. *A Noncooperative game is a triple* $(N, \mathbf{S}, \{\mathbf{f}_i\colon \mathbf{i} \in \mathbf{N}\})$, *where* N *is a nonempty, finite set of players,* S *is a set of feasible joint strategies,* $f_i(x)$ *is the payoff function for player* i, *which is real-valued on* S, *a strategy for each player* i *is an* m_i *vector* x_i, *and a joint strategy is an* $\{x_i : i \in N\}$.

Definition 8. *A noncooperative game* $(N, \mathbf{S}, \{\mathbf{f}_i\colon \mathbf{i} \in \mathbf{N}\})$, *is a Supermodular Game if the set* S *of feasible joint strategies is a sublattice of* $\mathbb{R}^m$, *the payoff function* $f_i(\mathbf{y}_i, \mathbf{x}_{-i})$ *is supermodular in* y_i *on* S_i *for each* x_{-i} *in* S_{-i} *and each player* i, *and* $f_i(\mathbf{y}_i, \mathbf{x}_{-i})$ *has increasing differences in* $(\mathbf{y}_i, \mathbf{x}_{-i})$ *on* $S_i \times S_{-i}$ *for each* i.

Theorem 3 (Topkis 4.2.3)**.** *Suppose that* $(N, \mathbf{S}, \{\mathbf{f}_i\colon \mathbf{i} \in \mathbf{N}\})$ *is a supermodular game, the set* S *of feasible joint strategies in nonempty and compact, and the payoff function* $f_i(\mathbf{y}_i, \mathbf{x}_{-i})$ *is upper semicontinuous in* y_i *on* $S_i(\mathbf{x}_{-i})$ *for each player* i *and each* x_{-i} *in* S_{-i}. *For each* x *in* S *and each subset* N' *of* N, *let* $x_{N'} = \{x_i : i \in N'\}$. *Let* x' *be the least element of* S. *For each subset* N' *of* N, *let* $S^{N'}$ *be the section of* S *at* $x'_{N \setminus N'}$. *For each subset* N' *of* N, *each player* i *in* N', *and each* $x_{N'}$ *in* $S^{N'}$, *let* $f_i^{N'}(x_{N'}) = f_i(x_{N'}, x'_{N \setminus N'})$. *Consider the collection of supermodular games* $(N', S^{N'}, \{f_i^{N'} : i \in N'\})$ *parameterized by the nonempty subsets* N' *of* N. *Then there exists a greatest equilibrium point and a least equilibrium point for each game* N', *and for each player* i *the strategy of player* i *in the greatest (least) equilibrium point for game* N' *is increasing in* N' *where* i *is included in* N'.

- Topkis Theorem 4.2.3 provides conditions under which the strategy of *each* player in the greatest equilibrium point, and the least equilibrium point, is increasing in a parameter, t
- These two Theorems apply to a finite number of players
- But analogous results have been proved for infinitely many players–and also for quasi-supermodular games (see Milgrom and Shannon, 1996)
- Want to know conditions under which the principal of *every* firm in the market induces a harder action from her agent in the greatest and least equilibrium of the game
- Interpret a player as being a principal, and a strategy for her as being a feasible section-best action (correspondence), $a^{**} = \sup_{a \in A} \{B(a) - C(a)\}$, and a product market strategy $\mathbf{z}_i \in Z_i$, where Z_i is the set of product market strategies for player i
- If this game is a supermodular game then Topkis's theorems imply that the actions implemented by all principals are increasing in the relevant measure of product market competition
- First we need the set of feasible joint strategies be compact
- If the sets of product market strategies Z_i are nonempty and compact for all i then it follows trivially from Tychonoff's Theorem that the set S of feasible joint strategies in the Product Market with Agency Game is compact.

- e.g. if a product market strategy is a price, quantity or supply function then S will be compact.
- Second requirement: the payoff function is supermodular in $\mathbf{y}_i \in S_i$.
- The key part of this requirement is that the agent's action and the product market strategy be complements
- e.g. in a Cournot game where agent effort reduces cost this condition requires that lower costs make choosing higher quantities more desirable
- Whether or not this condition is met clearly depends on the nature of the product market and the effect of the agents' actions.
- The final important condition is that the payoff exhibit increasing differences in $(\mathbf{y}_i, \mathbf{x}_{-i})$ on $S_i \times S_{-i}$ for all i.
- Also depends on the particulars of the game.
- e.g. in Cournot, this requires that a higher effort-quantity pair from one firm makes a higher effort-quantity pair from another firm more desirable.

9 Incomplete Contracts

9.1 Introduction and History

- Coase 1937: if the market is an efficient method of resource allocation then why do so many transactions take place within the firm ?!?!
- He claimed: because markets and firms are different (markets: price and haggling, firms: authority)
- In the 1990's the value added/sales ratio was 0.397 in France and 0.337 in Germany
- The extremes seem fairly intuitive
- The challenge for economists is to explain boundaries – what determines the mix between firms & markets
- D.H.Robertson: "We find islands of conscious power in oceans of unconsciousness like lumps of butter coagulating in buttermilk"
- Neoclassical theory of the firm: there are economies of scale, and then inefficiencies beyond some point
- But why can't you get around the potential diseconomies of scale by replication (expand by hiring another manager / building another factory)
- Just introducing agency problems (a la Jensen-Meckling 1976) doesn't say much about boundaries
- What does merging even mean in a world of optimal contracting ?

- Coase: firms arise because of "transaction costs"–makes market transaction more costly
- For Coase, these were haggling costs and the cost of learning prices
- Firms economize on these costs by replacing haggling with *authority*
- But there are also costs of authority – errors. And what about delegation/agency issues?
- Alchian & Demsetz (72): where does the authority come from? Firms are just like a market mechanism
  - Grocer example: I can tell my grocer what to do but they probably won't listen to me
  - The interesting question is why authority exists within firms
- Mid 70s: Williamson (71,75,79); Klein, Crawford & Alchian (78): much more analysis of the costs of the market – "haggling" costs
- The market becomes very costly when firms have to make relationship specific investments. egs. (i) site specificity (electricity generators near coal mines), (ii) physical asset specificity, (iii) human asset specificity, (iv) dedicated assets (building new capacity)
- Williamson: The "Fundamental Transformation" (ex ante competitive, ex post bilateral monopoly)
- An obvious solution is to write a long-term contract
- Indeed, in a world of perfect contracting this would solve the problem
- But Arrow-Debreu contingent contracts don't work well with asymmetric information, hidden actions, ...
- However, perhaps one could use a revelation mechanism to get the second-best
- BUT: (i) Bounded Rationality: it's hard to think about all the possible states of the world; (ii) it's hard to negotiate these things – need a common language; (iii) still - language has to be comprehensible to a 3rd party to make the contract *enforceable*
- Actual long-term contracts tend to be highly incomplete
- Indeed, they might not be very long-term
- *Any contract is ambiguous*
- *Renegotiation is a sign of incompleteness*
- We will proceed by assuming contractual incompleteness
- Later, we will return to the issue of foundations of incomplete contracts

9.2 The Hold-Up Problem

- Renegotiation may not proceed costlessly: (i) asymmetric information, (ii) rent-seeking behavior – this is about ex post efficiency. May apply
- Even if negotiation is costless the division of the surplus may be "wrong" in the sense that it won't encourage the right ex ante investments – this is about ex ante efficiency. Always applies
- Recall the Coase Theorem
- Maybe it's more efficient to do the whole thing in one big firm
- Williamson; Klein, Crawford & Alchian then hand waive about bureaucracy costs
- Empirical work: Monteverde-Teece, Marsten, Stuckey, Joskow
- Grossman-Hart (JPE, 1986); Hart-Moore (JPE, 1990): previous work does not provide a clear description of how things change under integration. Why is there a different feasible set – and why is it sometimes better and sometimes not?!
- The firm consists of two kinds of assets: human and non-human (tangible & intangibles). Human assets can't be bought and sold
- When contracts are incomplete, not all uses of an asset will be specified – there is some discretion – "Residual Control Rights"
- The RCRs belong to the *owner*
- *This is the fundamental characteristic of asset ownership – it is the key right*

Remark 4. *Grossman and Hart introduce this in a definitional sense*

9.2.1 A Numerical Example

- Consider the relationship between a B(uyer) and a S(eller) of an intermediate good (a "widget").
- B can use the widget to produce a final good which can be sold to a consumer.
- The consumer values the final good at v.
- S can make a privately costly investment which makes the widget cheaper to produce.
- If S makes the investment, which costs \$5, then the widget can be produced for \$10, otherwise it costs \$16 to produce.
- B can make a privately costly investment which makes the final good more valuable to the consumer.
- This investment also costs \$5.
- If B makes the investment then $v = \$40$, otherwise $v = \$32$.

- Not that B and S have different human capital characteristics.
- B and S would like to write a contract which specifies that each party should make its respective investment, because that leads the total surplus in the relationship to be $40 - 10 - 5 - 5 = 20$.
- But suppose that contracts are incomplete–*observable* but not *verifiable*.
- Now B and S will have to bargain about the price that B pays to S for the widget *after* the investment stage.
- Suppose that B and S are non-integrated so that at the bargaining stage they split whatever surplus is generated 50:50.
- This split arises in a situation of Nash bargaining because B cannot produce the final good without the widget from S, and S has no use for the widget if it is not sold to be. Therefore, both B and S have zero outside options.
- In this situation B will not invest–If B does invest, she will bear a private cost of 5, but gets half of the increase in surplus of $40 - 32 = 8$, or 4.
- Similarly, S bears a cost of 5 by investing, but gets an increased payoff of $(16-10)/2 = 3$ in the bargaining. So S won't invest either.
- Thus neither B nor S invests, and total surplus is thus $32 - 16 = 16$.
- Now suppose B and S are vertically integrated, with S owning B's machine that produces a final good.
- S no longer needs to bargain with B because S owns the machine.
- So S gets *all* of the increased surplus from investing in cost reduction (that is 16-10-5), and thus will be prepared to invest.
- However B will not invest as she will get none of the benefit of making the final good more valuable.
- S cannot compel B to invest, nor contract on B making the investment.
- Total surplus is thus $32 - 10 - 5 = 17$.
- This is larger than under non-integration, so forward vertical integration is desirable.
- B ownership (backward integration), does better still.
- Now B invests, but S does not, yielding total surplus of $40 - 16 - 5 = 19$.
- This is not as good as if contracting was possible (that would yield a surplus of 20), but it is better than the other possible ownership structures.
- What makes B ownership preferable to S ownership is that B's investment is relatively more important (at the margin) than S's.
- Both cost 5, but B's has a benefit of $40 - 32 = 8$, whereas S's has a benefit of $16 - 10 = 6$.

- Since asset ownership presumably *can* be contracted upon, we would expect B ownership to emerge as the equilibrium ownership structure as it maximize joint economic surplus.

9.2.2 A Non-Numerical Example

- Consider two firms: B(uyer) and S(eller)
- Case I: RCRs shared, Case II: S has all RCRs, Case III: B has all RCRs
- Bargaining power differs under different cases
- Which is best depends on *whose investment is important*
- $t \in \{0, 1, 2\}$
- Buyer makes an investment i, revenue is $R(i), R'(i) > 0, R''(i) < 0$
- B needs some input from S (a widget) at cost c (at date 2)
- Assume $R(i) > c, \ \forall i$
- Let $c = i$
- No discounting / interest rate = 0
- Symmetric information
- FB:
$$\max \{R(i) - c - i\}$$
- FOC:
$$R'(i) = 1 \Rightarrow i = i^*$$
- Suppose no long-term contracts and standard Nash bargaining
$$p = \frac{R(i) + c}{2}$$
- Why?
  - Each player gets her threat point plus half the gains from trade
  - Gains from trade at $t = 2$ are $R(i) - c$ (if no widget then no revenues)
  - Note, i is sunk at this point
  - If $p = \frac{R(i)+c}{2}$ then S gets
$$\frac{R(i) + c}{2} - c = \frac{R(i)}{2} - \frac{c}{2}$$
    which is exactly her outside option of zero plus half the gains from trade
- $\Rightarrow$ B's payoff is $R(i) - p - i = \frac{R(i)}{2} - \frac{c}{2} - i$

- Now:
$$\max_i \left\{ \frac{R(i)}{2} - i - \frac{c}{2} \right\}$$

- FOC: $R'(i) = 2 \Rightarrow i^{SB} < i^*$

9.2.3 Solutions to the Hold-Up Problem

1. LT contract which specifies the widget price in advance – BUT contractual incompleteness – *the more incomplete the contract the more bargaining power the seller has*

2. Contract on i - stipulate that B chooses i^*, S pays $\beta\Pi$. The payoffs are:
$$\begin{aligned} B &: \frac{R(i)}{2} - \frac{c}{2} - i + \Pi \\ S &: \frac{R(i)+i}{2} - c - \Pi = \frac{R(i)}{2} - \frac{c}{2} + \Pi \\ TOTAL &: R(i) - c - i = FB \end{aligned}$$
But this crucially relies on i being *verifiable* (what if quality is uncertain, eg)

3. Allocate the bargaining power – but how would you do that?

4. Reputation - works sometimes but not always

5. Assets – give B some good outside options (a second supplier – maybe an in-house supplier). OR Vertical Integration.

- This last point is a key motivation for what we do next

9.3 Formal Model of Asset Ownership

- Hart (chapter of Clarendon Lectures)

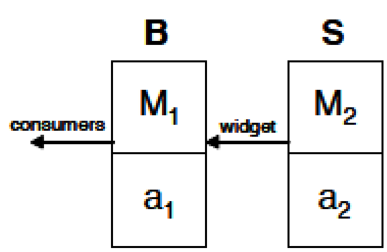

- Same time line as before
- Wealthy, risk-neutral parties
- No discounting

- No LT contracts
- ST contract at date 2
- At date 0 the parties can trade assets - this will matter at date 2 because it determines who has the Residual Control Rights (which here will just mean the right to walk away with the asset)
- 3 leading organizational forms: (i) Non Integration (M_1 owns a_1,M_2 owns a_2), (ii) Type I Integration (M_1 owns a_1 and a_2), (iii) Type II Integration (M_2 owns a_1 and a_2)
- Focus on human capital being inalienable, but physical assets being alienable
- Payoffs: M_1 invests i at cost i (think of this as market development for the final good). This leads to $R(i) - p - i$ if M_1 gets the widget from M_2 at cost p
  - But they do have an outside option – assume she can get a non-specific widget (think of it as lower quality) from a competitive market if there is no trade within the relationship, in which case the payoff is $r(i;A) - \overline{p}$ where A is the set of physical assets which M_1 owns
  - We use lower case r to indicate a lack of M_2's human capital
- $A = \{a_1\}, \{a_1, a_2\}, \varnothing$ – these correspond to Non Integration, Type I integration and Type II integration respectively
- M_2 invests e at cost e
- Production cost is $C(e)$ such that $C'(\cdot) < 0, C''(\cdot) > 0$
- If there is no trade with M_1, M_2 can supply her widget to the competitive market for general purpose widgets and receive $\overline{p} - c(e; B)$ with c decreasing in B
  - Little c indicates the lack of M_1's human capital
  - B is the set of assets. $B = \{a_2\}$ under NI, $B = \varnothing$ under type I integration and $B = \{a_1, a_2\}$ under type II integration

Formal Assumptions

1. $R(i) - C(e) > r(i;A) - c(e;B), \forall i, e, A, B,$ where $A \cup B = \{a_1, a_2\}, A \cap B = \varnothing$. ie. there are always ex post gains from trade
2. $R'(i) > r'(i; \{a_1, a_2\}) \geq r'(i; \{a_1\}) \geq r'(i; \varnothing)$, for all $0 < i < \infty$
3. $|C'(e)| > |c'(e; \{a_1, a_2\}| \geq |c'(e; \{a_2\}| \, |c'(e; \varnothing|$ for all $0 < e < \infty$

- 1 says that i and e are relationship specific – they pay off more if trade occurs
- 2 and 3 say that this relationship specificity holds in a marginal sense
- Assume that R, r, C, c, i, e are observable but not verifiable

- In the First-Best:
$$\max \ \{R(i) - C(e) - i - c\}$$
- FOCs are $R'(i^*) = 1$ and $-C'(e^*) = 1 = |C'(e^*)|$
- SB: Fix the organizational form, assume no LT contract and 50/50 Nash Bargaining at date 2
- Note that the ex post gains from trade are $(R - C) - (r - c)$
- M_1 and M_2's payoffs ex post are
$$\begin{aligned}\Pi_1 &= r - \overline{p} + \frac{1}{2}\left[(R - C) - (r - c)\right] \\ \Pi_2 &= \overline{p} - c + \frac{1}{2}\left[(R - C) - (r - c)\right]\end{aligned}$$
and the price of the widget is
$$p = \bar{p} + \frac{1}{2}(R - r) - \frac{1}{2}(c - C)$$
- M_1 solves:
$$\max_i \{\Pi_1 - i\}$$
$$\max_i \left\{\frac{1}{2}R(i) + \frac{1}{2}r(i; A) - \frac{1}{2}C(e) + \frac{1}{2}c(e; B) - i\right\}$$
- The FOC is:
$$\frac{1}{2}r'(i; A) + \frac{1}{2}R'(i) = 1$$
- M_2 solves:
$$\max_e \{\Pi_2 - e\}$$
$$\max_e \left\{\bar{p} - \frac{1}{2}C(e) - \frac{1}{2}c(e; B) + \frac{1}{2}R(i) - \frac{1}{2}r(i; A) - e\right\}$$
- The FOC is:
$$\frac{1}{2}\left|C'(e)\right| + \frac{1}{2}\left|c'(e; B)\right| = 1$$
- Together, these FOCs determine a Nash equilibrium
- Recall that $R' > r' \Rightarrow i^{SB} < i^*$
- Under any ownership structure we get underinvestment since $R'' < 0$ and $C'' > 0$
- Intuition: marginal investment by M_1 increases gains from trade by $R'(i)$ but her payoff only increases by $\frac{1}{2}R'(i) + \frac{1}{2}r'(i; A) < R'(i)$
- $i^{T2} \leq i^{NI} \leq i^{T1} < i^*$ and $e^{T1} \leq e^{NI} \leq e^{T2} < e^*$

- Let $S = R(i) - C(e) - i - e$ be the total surplus given ex post bargaining
- Compute it at NI, T1, T2 and see which is larger
- Key: the Coase Theorem says we will get this outcome

Results:

1. Type 1 Integration is optimal is M_1's investment is important, Type 2 Integration is optimal if M_2's investment is important, Non Integration is optimal if both are similarly important

Definition 9. *Assets a_1 and a_2 are Independent if $r'(i; \{a_1, a_2\}) \equiv r'(i; \{a_1\})$ and $c'(e; \{a_1, a_2\}) \equiv c'(e; \{a_2\})$ (a notion of marginal incentives)*

Definition 10. *Assets a_1 and a_2 are Strictly Complimentary if either $r'(i; \{a_1\}) \equiv r'(i; \varnothing)$ or $c'(e; \{a_2\}) \equiv c'(e; \varnothing)$*

Definition 11. *M_1's human capital (respectively M_2's human capital) is Essential if $c'(e; \{a_1, a_2\}) \equiv c'(e; \varnothing)$ (respectively $r'(i; \{a_1, a_2\}) \equiv r'(i; \varnothing)$)*

2. If a_1, a_2 are Independent NI is optimal
3. If a_1, a_2 are Strictly Complimentary then some form of integration is optimal
4. If M_1's human capital is Essential then Type 1 Integration is optimal
5. If M_2's human capital is Essential then Type 2 Integration is optimal
6. If M_1 and M_2's human capital are both Essential the organizational form doesn't matter –all are equally good
7. Joint Ownership is suboptimal (one notion of joint ownership is mutual veto) – creating a veto is like turning the asset into a Strictly Complimentary Asset – creates MUTUAL Hold-Up

- All proofs follow directly from the FOCs

Investment in the asset itself:

- "Russian Roulette Agreements": 1 can name a price p to buy 2 out – 2 can accept, or reject and must buy 1 out for p (wealth constraints can be a big issue)
- Can also set up mechanisms with different percentages of the income and control rights
- Argument about joint ownership being bad relies upon investment being in human capital, not the physical asset

Comments:

- Can generalize this to many individuals and many assets (Hart-Moore (JPE, 1990))

- Robustness? (a) Human capital/physical capital thing; (b) Rajan-Zingales: 1 asset model and 1 investment with 2 people: M_1's FOC becomes $\frac{1}{2}R'(i) + \frac{1}{2}r'(i) = 1$ and M_2's FOC becomes $\frac{1}{2}R'(i) + \frac{1}{2}\underline{r}'(i) = 1$ where $r(i) \equiv r(i; \{a\}), \underline{r}(i) \equiv r(i; \varnothing)$. Suppose $\underline{r}'(i) = 0$ and $r'(i) < 0$ (eg. multi-tasking) – then one gets the opposite result from Hart-Moore. How much do you concentrate on the relationship...
- Baker, Gibbons & Murphy $r' > 0$ and $R' = \underline{r}' = 0$ (rent seeking behavior) $\rightarrow FB$: $i = 0$

9.3.1 Different Bargaining Structures

- Ex post bargaining matters
- Under Rubinstein bargaining the outside option can have a different effect
- Hart-Moore use Nash bargaining
- Binmore, Rubinstein & Wolinsky
- Suppose you can't enjoy outside options whilst bargaining
- The OUTSIDE OPTION PRINCIPLE: M_1 gets $\max\left\{\frac{1}{2}, r\right\}$
- Comes down to whether it is credible to exercise the outside option
- de Mezer-Lockwood do outside option bargaining in a similar model

9.3.2 Empirical Work

- Elfenbein-Lerner (RAND, 2003)
  - Builds on earlier work by Lerner & Merges
  - Looks at 100+ alliance contracts between internet portals and other firms
  - Material on portal sites often provided through alliances
  - Important relationship specific investments / effort: development of content, maintenance & hosting, provision of customer service, order fulfillment, billing
  - Significant alienable assets: servers, URL, customer data
  - Also specific control rights/contractual rights: eg. restrict lines of business of a party, need approval for advertising
  - Opportunism exists
  - Does allocation of asset ownership depend on the important of specific investments? Should the partner who "does a lot" own a lot of the assets?
  - Aghion-Tirole (QJE, '94) model with wealth constraints – the "logical" owner may not be able to afford them
  - EL find that relative wealth is not so important for asset ownership in their data

  - For contractual rights: depends much more on relative wealth and less on importance of investments

- Woodruff (IJIO): Mexican footware industry – relationship between producers and small retailers – integration or not?
- Quick style changes: more retailers independent ownership – consistent with downstream incentives being important in that case
- Mullainathan-Scharfstein (AER PP '01); Stein et al; Hong et al – integration does seem to matter

9.4 A General Framework

- The above model with 2 parties and 2 assets captures the key insights of PRT.
- But *a priori* unclear that it extends to large firms–many divisions, lots of employees.
- Hart-Moore (JPE, 1990) address this question.
- At time 1 I agents make investments x_i at cost $\psi(x_i)$.
- We abuse notation at let I be the number of agents and the set.
- The set of all available assets in the economy is $\overline{A}$ and the investments are made in some subset of that $A \subseteq \overline{A}$.
- At time 2 the investments of a subset $S \subseteq I$ of the agents and the assets in A generate ex post surplus of $V(S; A|x)$, where $x = (x_1, ..., x_I)$.
- Suppose that no contract can be written ex ante.
- Now that we have multiple parties involved in the ex post bargaining, Nash bargaining is problematic.
- Hart-Moore argue that (as in the case of 2 parties and Nash bargaining) that the Coase theorem should apply and hence there should be ex post efficient no matter what the asset allocation is.
- Use the Shapley value.
- If the set $\mathcal{S}$ is the power set of I, and $\mathcal{A}$ is the power set of $\overline{A}$ then the function $\omega : \mathcal{S} \to \mathcal{A}$ is the *ownership allocation*–i.e. the subset of assets owned by the subset of agents S.
- Assumption: each asset can only be controlled by, at most, one group S.
- Assumption: any asset controlled by some group is controlled by the *whole* group.
- Formally: $\omega(S) \cap \omega(I \setminus S) = \emptyset$, and $\omega(S') \subseteq \omega(S)$, so that $\omega(\emptyset) = \emptyset$.

Definition 12. *Given an ownership allocation $\omega(S)$, a vector of investments x, and associated ex post surplus for any group S, $V(S,\omega(S)|x)$, the expected ex post surplus for agent i is given by the **Shapley Value** if:*

$$B_i(\omega|x) = \sum_{S|i\in S} p(S)\left[V(S;\omega(S)|x) - V(S\setminus\{i\};\omega(S\setminus\{i\})|x)\right]$$

where $p(S) = ((s-1)!/(I-s)!)/I!$, and s is the cardinality of S.

- The Shapley value is the expected payoff taken over all possible subgroups S that agent i might join ex post–and any order of group formation is equally likely.
- It assigns each agent i in a group the difference between the surplus to the entire group and the surplus without agent i–i.e. agent i's expected contribution to surplus.
- As an example, suppose that $I = 2$ and $\overline{A} = \{a_1, a_2\}$–interpret agent 1 as a printer and agent 2 as a publisher.
- Three possible ownership structures
  - Nonintegration: $\omega(1) = \{a_1\}, \omega(2) = \{a_2\}$
  - Publisher Integration: $\omega(1) = \emptyset, \omega(2) = \{a_1, a_2\}$
  - Printer Integration: $\omega(1) = \{a_1, a_2\}, \omega(2) = \emptyset$

9.4.1 Nonintegration

- This means that no surplus can be generated without both assets.
- The ex post surplus that can be generated with only one asset is $V(\{1\};\{a_1\}|x) = V(\{2\};\{a_2\}|x) = 0$.
- If the two agents trade (i.e. form a group) then the surplus is $V(\{1,2\};\{a_1,a_2\}|x) > 0$.
- The Shapley value determines each agent's share of the surplus.
- There are only two possible orderings of group formation $\{1,2\}$ and $\{2,1\}$.
- So each agent gets $\frac{1}{2}V(x)$.

9.4.2 Printer integration

- This means that the printer owns both assets, and can thus generate surplus on her own, but the publisher cannot.
- It is plausible that the printer could do even better by hiring the publisher and thus that $V(\{1\};\{a_1,a_2\}|x)\Phi_1(x_1) < V(x)$..
- So player 1 gets
$$B_1(PI|x) = \frac{1}{2}\left(V(x) - \Phi_1(x_1)\right) + \Phi_1(x_1)$$
- And player 2 gets
$$B_2(PI|x) = \frac{1}{2}\left(V(x) - \Phi_1(x_1)\right).$$

9.4.3 Publisher Integration

- This means that the publisher owns both assets, and can thus generate surplus on her own, but the printer cannot
- Symmetric to above we have
- Player 1 gets

$$B_1(pI|x) = \frac{1}{2}\left(V(x) - \Phi_2(x_2)\right)$$

- And player 2 gets

$$B_2(pI|x) = \frac{1}{2}\left(V(x) - \Phi_2(x_2)\right) + \Phi_2(x_2).$$

9.4.4 Investments

- Suppose $V(x)$ is strictly concave and increasing in $x = (x_1, x_2)$, that $\Phi_i(x_i)$ is increasing and concave, and that $\psi_i(x_i)$ is strictly increasing and convex.
- Regardless of the ownership structure, each agent solves

$$max_{x_i}\left\{B_i(\omega(S)|x_1, x_2) - \psi_i(x_i)\right\}.$$

- Since investments can't be contracted on, they are made non-cooperatively.
- We focus on Nash equilm of the investment game.
- In general, then, the FOC for agent i is

$$\frac{\partial B_i(\omega(S)|x_1, x_2)}{\partial x_i} = \psi_i'(x_i).$$

- Under non integration we have

$$\frac{1}{2}\frac{\partial V(x_1^{NI}, x_2^{NI})}{\partial x_1} = \psi_1'(x_1^{NI}),$$

$$\frac{1}{2}\frac{\partial V(x_1^{NI}, x_2^{NI})}{\partial x_2} = \psi_2'(x_2^{NI}).$$

- Under printer integration we have

$$\frac{1}{2}\frac{\partial V(x_1^{PI}, x_2^{PI})}{\partial x_1} + \frac{1}{2}\Phi_1'(x_1^{PI}) = \psi_1'(x_1^{PI}),$$

$$\frac{1}{2}\frac{\partial V(x_1^{PI}, x_2^{PI})}{\partial x_2} = \psi_2'(x_2^{PI}).$$

- Under publisher integration we have

$$\frac{1}{2}\frac{\partial V(x_1^{pI}, x_2^{pI})}{\partial x_1} = \psi_1'(x_1^{pI}),$$

$$\frac{1}{2}\frac{\partial V(x_1^{PI}, x_2^{pI})}{\partial x_2} + \frac{1}{2}\Phi_2'(x_2^{pI}) = \psi_2'(x_2^{pI}).$$

- Printer has greater incentives to invest under printer integration than under the other two ownership structures is $\Phi_1'(x_1) > 0$.
- Same reasoning for the publisher under publisher integration.
- Note that it could be the case that $\Phi_i'(x_i) \leq 0$.
- Eg. If investments adapted to special skills of agent j, then if j is not hired ex post then the investment could be counterproductive.
- Customization toward agent j weakens agent i's bargaining position by lowering her outside option.
- Also, if $\Phi_1'(x_1)$ is too large then there could be over investment so that non-integration is optimal.
- Some form of integration optimal if $0 < \Phi_i'(x_i) \leq \partial V(x_i, x_j)/\partial x_i$, for all x_j.

9.5 Real versus Formal Authority

- Inside the firm asset ownership doesn't matter
- Authority matters inside the firm – and this is not achieved through assets
- How is authority allocated inside a firm?
- Initial model: 2 parties, P and A – what is the optimal authority between P and A
- Assumption: authority can be allocated – this can be achieved contractually (eg. shareholders allocate authority to the board)
- Boards allocate authority to management – management to different layers of management
- AT call this stuff "Formal Authority" (legal / contractual)
- Distinction between this and "Real Authority" (which is what is the case if the person with Formal authority typically "goes along" with you)
- Asymmetric information important

Model:

- $\{P, A\}$
- Each can invest in "having an idea" – only 1 can be implemented
- P chooses prob E of having an idea at cost $g_p(E)$ with $E \in [0, 1]$
- A chooses prob e of having an idea at cost $g_a(e)$ with $e \in [0, 1]$

- Assume $g_i(0) = 0, g_i'(0) = 0, g_i' > 0$ elsewhere, $g_i'' > 0, g_i'(1) = \infty \; \forall i \in \{A, P\}$, in order to ensure an interior solution
- If it exists, P's idea is implemented giving payoffs of B to P and αb to A where $\alpha \in [0, 1]$ is a congruence parameter (their preferences are "somewhat" aligned)
- If A's idea is implemented the payoffs are b to A and αB to P

Case I: P has formal authority

$$\begin{aligned} U_P &= EB + (1-E)e\alpha B - g_p(E) \quad (16) \\ U_A &= E\alpha b + (1-E)eb - g_a(e) \quad (17) \end{aligned}$$

- P maximizes (16) by choosing E and A maximizes (17) by choosing e
- The FOCs are:

$$\begin{aligned} B(1-e\alpha) &= g_p'(E) \\ b(1-E) &= g_a'(e) \end{aligned}$$

- Under a stability assumption you get a unique Nash Equilibrium
- P and A effort are substitutes – whereas in Hart-Moore they are complements
- Higher effort from P crowds-out effort from A
  - May want to "overstretch"
  - May want to find an agent with more congruent preferences

Case II: A has formal authority

- P solves:

$$\max_E \{e\alpha B + (1-e)EB - g_p(E)\}$$

- A solves:

$$\max_e \{eb + (1-e)E\alpha b - g_a(e)\}$$

- The FOCs are:

$$\begin{aligned} B(1-e) &= g_p'(E) \\ b(1-\alpha E) &= g_a'(e) \end{aligned}$$

- Which implies $E \uparrow, e \downarrow$ (effort levels are strategic substitutes)
- Comparing the FOCs with the P formal authority shows that A effort increases when A has formal authority
- If there is a P with several Agents then the P may "want to be overstretched" to give good innovation incentives to subordinates – just "puts out fires"

Comments:

1. Seems to have quite a nice flavor–sounds like the right setup
2. Ignores ex post renegotiation (since $B = b$) – imposes an ex post inefficiency.
   (i) Perhaps authority is ex post non-transferable and implementing ideas is ex post non-contractable

   (ii) But this opens another door – lead to ex post inefficiency
3. Inside a firm, what gets allocated? Formal or Real authority?

9.6 Foundations of Incomplete Contracts

9.6.1 The Maskin-Tirole Critique

- Whole premise of PRT is that there is some information that it *observable* but not *verifiable* to a third party.
- Maskin-Tirole (ReStud, 1999) argue that if it's observable then it can be *made* verifiable by using a clever mechanism.
- They use the subgame perfect implementation mechanisms introduced by Moore-Repullo (ECMA, 1988).
- The MT argument is that the parties can write a contract that specifies the *payoff* contingencies for any state of the world.
- Then, once the state has been realized, the parties can fill in the physical details.
- Of course, this requires truth-telling to be incentive compatible.
- The MR mechanism achieves this.

9.6.2 Preliminary: Implementation Literature

- Began with Maskin (WP, 1977 – reprinted ReStud, 1999)
- Observable information can be made verifiable and hence contractible through a mechanism
- Ask the parties what the state of nature was and if they don't agree then deliver a large punishment
- Can yield truth-telling as a Nash Equilm
- But: (i) There are generally other equilibria, (ii) There is an incentive to renegotiate because punishment is not in their ex post collective or individual interests, (iii) Never seen in practice
- Consider a correspondence $f(\theta)$ to be implemented
- Players announce messages $(m_1, ..., m_n)$ and the outcome is $g(m_1, ..., m_n)$

- Require: (i) Monotonocity – if $a \in f(\theta)$ then $a \in f(\widetilde{\theta})$ whenever for each individual and each outcome $b \in A$, a is weakly preferred to b by i in state θ it is also weakly preferred by i in state $\widetilde{\theta}$, and (ii) Weak No Veto Power "WNVP": $f(\theta)$ satisfies WNVP if $a \in f(\theta)$ whenever at most one agent doesn't have a as her most preferred choice, $\forall\theta$ (this is like weak non-dictatorship)

Theorem 4. *(Maskin, 1977) If $f(\theta)$ is implementable then it is Monotonic and if there are at least three agents then if $f(\theta)$ is Monotonic and satisfies WNVP then it is Nash Implementable.*

- Intuition:
  - Necessity: if an outcomes is a Nash Equilm of a mechanism in a state it will remain an equilm in another state where this outcome remains as attractive as other outcomes
  - Sufficiency: this part shows how to construct the mechanism. Get rid of equilibria we don't want by enrichening the message space of the agents. Gets rid of disagreement on the true state by allowing any individual agent to impose another outcome that she is known *not* to prefer in the true state (then monotonicity kicks in). Get rid of equilibria where agents agree on the state and $a \notin f(\theta)$ or there is no agreement by allowing agents to individually impose their favorite outcome by naming the largest integer of all the integers chosen by the agents. This works because equilibria involve pre-specified strategies, and hence integers. This unbounded strategy space ensures non-existence of such equilibria.
- Comments:
  1. Monotonicity is quite restrictive – and in particular it rules out seeking any particular distributional outcomes
  2. Integer game not at all natural

Subgame-Perfect Implementation

- Moore-Repullo (Econometrica, 1988)
- Do away with the integer game
- Main strength: get rid of the monotonicity assumption of Maskin
- The most desirable outcomes are subgame-perfect equilibria
- Following example is from Aghion, Fudenberg, Holden, Tercieux and Kunimoto (2009).
- Two parties, a Buyer B and a Seller S of a single unit of an indivisible good.
- Conditional on trade, B's payoff is $V_B = \theta - p$, where θ is the value of the good to the buyer and p is the price.
- S's payoff is $V_S = p$.
- Good can be either high or low quality.

- If H then B values it at 14, and if L then B values it at 10.
- Suppose that the quality θ is common knowledge between B and S
- Then the following mechanism induces both parties to reveal the truth as a unique equilibrium of the game induced my the mechanism

1. B announces either "high" or "low." If "high" then B pays S a price equal to 14 and the game then stops.
2. If B announces "low" and S does not "challenge" B's announcement, then B pays a price equal to 10 and the game stops.
3. If S challenges B's announcement then:
   (a) B pays a fine F to T (a third party)
   (b) B is offered the good for 6
   (c) If B accepts the good then S receives F from T (and also the 6 from B) and we stop.
   (d) If B rejects at 3b then S pays F to T
   (e) B and S Nash bargain 50:50 over the good.

- This mechanism yields truth-telling as the unique (subgame perfect) equilibrium.
- Suppose $v = 14$, and let $F = 9$.
- If B announces "high" then B pays 14 and we stop.
- If, however, B announces "low" then S will challenge because at stage 3a B pays 9 to T and, this cost being sunk, B will still accept the good for 6 at stage 3b (since it is worth 14 and $14 - 6 = 8$ is greater than $14/2 = 7$ which is what B gets if it rejects the offer at 6).
- Anticipating this, S knows that by challenging B, S receives $9 + 6 = 15$, which is greater than the 10 that S would receive if S did not challenge.
- Moving back to stage 1, if B lies and announces θ'' when the true state is θ', B gets $14 - 9 - 6 = -1$, whereas B gets $14 - 14 = 0$ if he tells the truth.

9.6.3 The Robustness Counter-Critique

- MT meant more than *observable*, they meant *common knowledge*.
- What if the state of nature is common p-belief for p close to 1.
- Monderer-Samet notion: I believe with probability p that you believe with probability p, and so on ad infinitum.
- Question: does the MT depend crucially one the CK embedded assumption?
- Answer: it does (Aghion-Fudenberg-Holden-Kunimoto-Tercieux, 2009).
- Now allow for a common p-belief perturbation from common knowledge

- Suppose that the players have a common prior that $\Pr(v = 14) = \frac{1}{2}$ and $\Pr(v = 10) = \frac{1}{2}$.
- Each player receives a conditionally independent draw from a signal structure with two possible signals: θ' or θ''.
- Let the perturbed signal structure ν_ε be as follows

	$\theta'_B\theta'_S$	$\theta'_B\theta''_S$	$\theta''_B\theta'_S$	$\theta''_B\theta''_S$
$\Pr(v = 14)$	$\frac{1}{2}\left(1-\varepsilon-\varepsilon^2\right)$	$\frac{1}{2}\varepsilon$	$\frac{1}{4}\varepsilon^2$	$\frac{1}{4}\varepsilon^2$
$\Pr(v = 10)$	$\frac{1}{4}\varepsilon^2$	$\frac{1}{4}\varepsilon^2$	$\frac{1}{2}\varepsilon$	$\frac{1}{2}\left(1-\varepsilon-\varepsilon^2\right)$

- Is there an equilibrium in pure strategies in which the buyer always reports truthfully?
- Suppose there is and suppose that B gets signal θ'_B.
  - Then she believes that, regardless of what signal player S gets, the value of the good is greater than 10 in expectation
  - So she would like to announce "low" if she expects that subsequent to such an announcement, S will not challenge
  - Now, suppose B announces low
  - In a fully revealing equilibrium, S will infer that B must have seen signal θ''_B if B announces low
  - S now believes that expected value of the good is ≤ 12 (one high signal, one low) so will not challenge because B will accept the good a stage 3 and S will pay the fine
  - But if S will not challenge then B would prefer to announce "low" when B received signal θ'_B.
  - Therefore there does not exist a truthfully revealing equilibrium in pure strategies.
- But maybe there are mixed strategy equilibria in which the mixing probability on the truthful announcement goes to 1 as ϵ goes to 0
- It turns out that this is not the case
- Mixing probabilities:

	High	Low
θ'_B	$1-\sigma'_B$	σ'_B
θ''_B	σ''_B	$1-\sigma''_B$

	Challenge	Don't Challenge
θ'_S	$1-\sigma'_S$	σ'_S
θ''_S	σ''_S	$1-\sigma''_S$

- **Proposition**: There is no sequence of equilibrium strategies $\sigma'_B, \sigma''_B, \sigma'_S$ and σ''_S which all converge to 0 as $\varepsilon \to 0$.

Proof Sketch:

- When $\varepsilon \to 0$, B ends up playing in pure strategies in stage 3 (for ε small enough, buyer B who observed $\theta"_B$ does not accept the good at price 6 while he does if he observed θ'_B)

- In stage 2, S privileges B's announcement over her own signal, thus never challenges B's announcement as $\varepsilon \to 0$
- When $\varepsilon \to 0$, B ends up playing in pure strategies in stage 1 (buyer B with signal θ''_B gets approximately $10 - 14$ if plays H, and gets zero if he plays L)
- $\to$ thus for ε small, there is no way that σ makes B indifferent between H and L
- $\to$ same reasoning for player θ'_B.
- Then, same reasoning as in the pure strategy equilibrium case shows that there is no way for σ to be an equilibrium.
- Does a bad equilm exist?
- Consider the following common p-belief perturbation ν^ε of the complete information structure.

ν^ε	$\theta^{'}_B, \theta^{'}_S$	$\theta^{'}_B, \theta^{"}_S$	$\theta^{"}_B, \theta^{'}_S$	$\theta^{"}_B, \theta^{"}_S$
$v = 14$	α	0	0	0
$v = 10$	0	$(1-\alpha)\varepsilon/2$	$(1-\alpha)\varepsilon/2$	$(1-\alpha)(1-\varepsilon)$

- Then consider the following strategy profile of the game with prior ν^ε
  1. B announces low regardless of his signal
  2. If B has announced low, S does not challenge regardless of her signal.
  3. Off the equilibrium path, i.e. if B announced low and S subsequently challenged, then B always rejects S's offer.
- To complete the description of the candidate sequential equilibrium, we also have to assign beliefs over states and signals for each signal of each player and for any history of play.
  1. Before playing the game but after receiving their private signals, we assume that agents's beliefs are given by ν^ε conditioned on their private signals.
  2. Out of equilibrium, if B is offered the good for 6 (which requires that S will have challenged), we assume that B always believe with probability one that the state is $v = 10$ and that S has received signal $\theta^{"}_S$.
- For $\varepsilon > 0$ sufficiently small, the above strategy profile is *sequentially rational* given the beliefs we just described and conversely these beliefs are *consistent* given the above strategy profile (Kreps-Wilson).
- Proof of sequential rationality:
  1. At Stage 3, regardless of his signal, B believes with probability one that the state is $v = 10$. Accepting S's offer at 6 generates $10 - 9 - 6 = -5$ and rejecting it generates $5 - 9 = -4$. Thus, it is optimal for B to reject the offer.

2. Moving back to Stage 2

   ⟶ if S chooses "Challenges," S anticipates payoff approximately equal to $7-9 = -2$ if her signal is $\theta_S^{'}$ and to $5-9=-4$ if the signal is $\theta_S^{"}$ as ε becomes small

   ⟶ if S chooses "No Challenge," S guarantees a payoff of 10.

   ⟶ thus, regardless of her signal, it is optimal for S not to challenge.

3 Moving back to Stage 1, B "knows" that S does not challenge regardless of her signal. Suppose that B receives $\theta_B^{'}$

- 1. ⟶ Then, as ε becomes small, B believes with high probability that the state is $v = 14$ so that his expected payoff approximately results in $14 - 10 = 4$.

     ⟶ this is larger than 0, which B obtains when announcing "High."

     ⟶ therefore, it is optimal for B to announce "Low." Obviously, this reasoning also shows that when B has received signal $\theta_B^{"}$, it is optimal for him to announce "Low."

- General MR setting:
  - n states of the world; two players; quasi-linear utility
  - Player 1 announces her preferences
  - Player 2 can challenge this, or not
  - If player 2 challenges then player 1 is given a choice b/w two alternatives (cleverly constructed)
  - Repeat to elicit player 2's preference

- This mechanism can implement *any* SCF as a *unique* equilibrium.

- Two agents: 1 and 2, whose preferences over a social decision $d \in D$ are parametrized by $\omega_i \in \Omega_i$ for $i = 1, 2$.

  → let $\Omega_i = \{\omega_i^1, ..., \omega_i^n\}, \Omega = \Omega_1 \times \Omega_2$.

- The agents have utility functions as follows:

$$u_1(d, \omega_1) - t_1,$$
$$u_2(d, \omega_2) + t_2$$

  where d is a collective decision, t_1 and t_2 are monetary transfers.

- The agent's ω's are common knowledge

- Let $f = (D, T_1, T_2)$ be a social choice function, where for each $(\omega_1, \omega_2) \in \Omega_1 \times \Omega_2$:

  → the social decision is $d = D(\omega_1, \omega_2)$

  → the transfers are $(t_1, t_2) = (T_1(\omega_1, \omega_2), T_2(\omega_1, \omega_2))$.

- The following mechanism induces truth-telling as unique equilibrium outcome:

  → the mechanism has two phases, one phase for each agent

  → each phase consists of three stages

→ below we describe phase 1, where player 1 first announces ω_1

→ phase 2 is the same as phase 1 with the roles of players 1 and 2 being reversed, i.e. with player 2 first announcing ω_2.

1. *Stage 1:* Player 1 announces a preference ω_1, and we proceed to stage 2.
2. *Stage 2:* If player 2 agrees then the phase ends and we proceed to phase 2; If player 2 does not agree and "challenges" by announcing some $\phi_1 \neq \omega_1$, then we proceed to stage 3.
3. *Stage 3:* Player 1 chooses between $\{x; t_x + \Delta\}$ and $\{y; t_y + \Delta\}$ where

$$\begin{aligned} u_1(x, \omega_i) - t_x &> u_1(y, \omega_1) - t_y; \\ u_1(x, \phi_1) - t_x &< u_1(y, \phi_1) - t_y. \end{aligned}$$

4. Then:

   → if player 1 chooses $\{x; t_x + \Delta\}$, then player 2 receives $t_2 = t_x - \Delta$ (and a third party receives 2Δ).

   → if player 1 chooses $\{y; t_y + \Delta\}$, then player 2 receives $t_2 = t_y + \Delta$.

- If agent 1 lied at stage 1 then agent 2 could challenge with the truth and then at stage 3 agent 1 will find it optimal to choose $\{y; t_y + \Delta\}$.
- If Δ is sufficiently large then this will be worse for agent 1 than telling the truth and having the choice function f implemented

  → moreover, agent 2 will be happy with receiving $t_y + \Delta$.
- If agent 1 tells the truth at stage 1 then agent 2 will not challenge

  → because she knows that agent 1 will choose $\{x; t_x + \Delta\}$ at stage 3 which will cause agent 2 to pay the fine of Δ.

⟶ Now consider a common p-belief perturbation parameterized by $\varepsilon \longrightarrow 0$, where for ε small player i tends to disregard his own signal if it differs from player j's announcement with regard to ω_j.

- **Definition:** Take a SCF f which is MR implementable under common knowledge. *Truthtelling is robust* for this SCF if for any $\varepsilon-$perturbation of the complete information structure, there is a profile of mixed equilibrium strategies for the perturbed game which converge to truth-telling when $\varepsilon \to 0$.
- **Proposition:** Truth-telling is not robust for the above SCF $f = (D, T_1, T_2)$.
- Now let us go beyond MR mechanism and consider *any* mechanism
- The failure of Maskin monotonicity in contractual situations has a very general implication

  → whenever there exists a "good" sequential equilibrium in the complete information game, there always exists also a "bad" sequential equilibrium in arbitrarily small p-belief perturbations of that game.

- **Idea:** by introducing just a small amount of incomplete information one may rapidly increase the sets of (sequential) beliefs that are consistent with dynamic rationality

  $\rightarrow$ in particular one can make a bad NE become consistent with rationality and thereby turn it into a sequential equilibrium

- Suppose n players

- Each player i has a utility function $u_i(a, \omega_i)$

- Players do not observe the state directly, but are informed of the state via signals

- A mechanism Γ together with a state $\omega \in \Omega$ defines an extensive form game $\Gamma(\omega)$.

  $\rightarrow$ let $SPE(\Gamma(\omega))$ denote the set of subgame perfect equilibria of the game $\Gamma(\omega)$

- **Theorem:** Assume finite state space and finite strategy spaces. And suppose that a mechanism Γ SPE-implements a non-monotonic SCF f under common knowledge. Then there exists a sequence of p-belief perturbations parametrized by ε and a corresponding sequence of sequential equilibria of the games induced by Γ under these perturbations, whose outcomes do not converge to $f(\omega)$ in some state ω as $\varepsilon \rightarrow 0$.

- In particular this result implies that whenever a SCF cannot be implemented using static mechanisms (with NE as a solution concept), there is no hope to implement it using sequential mechanisms if we want mechanisms to be robust to small common p-belief perturbations from common knowledge

  $\rightarrow$ this defeats the purpose of using subgame perfect implementation

- Suppose SCF f is not monotonic

  $\rightarrow$ then there would exist ω and ω' such that for all players i and for all alternative b

  $$f(\omega) \succeq_{i,\omega} b \Longrightarrow f(\omega) \succeq_{i,\omega'} b \tag{I}$$

  and nevertheless $f(\omega) \neq f(\omega')$

- Since Γ implements f, there exists a SPE $m_\omega \in SPE(\Gamma(\omega))$ such that $g(m_\omega) = f(\omega)$

- But then m_ω is a Nash equilibrium in state ω' (same as in Maskin's proof above)

- ....and necessarily a "bad" NE since $f(\omega) \neq f(\omega')$

- Then, one can use a common p-belief perturbation to "rationalize" this bad NE and turn it into a sequential equilibrium of the perturbed games.

- How do we achieve this? Illustration in the HM example

- There, having B always announce a low v at stage 1 and then having S never challenge at stage 2, is a bad NE equilibrium but is not a sequential equilibrium under common knowledge

  $\longrightarrow$ if stage 3 were to be reached under common knowledge, then B would just infer that S deviated from the equilibrium, but never update his beliefs about v or about S's perception of v

- However, moving from common knowledge to common p-beliefs changes things radically:

  $\longrightarrow$ now, if stage 3 is reached, then B updates his beliefs about which signal S must have seen

  $\longrightarrow$ in particular, if B's updating puts weight on S having the low signal, then B will not take the offer at price 6 and anticipating this S will indeed not challenge in equilibrium

- Thus by moving from common knowledge to common p-belief, we have:

  1. extended the set of consistent beliefs

     $\longrightarrow$ under common knowledge it could not be a consistent belief that S saw $\underline{v}$ if B "knew" that $v = \overline{v}$

     $\longrightarrow$ under common p- belief, achieving stage 3 can be made consistent with S having observed the low signal

  2. turned a bad (non-sequential) NE of the game with common knowledge into a sequential equilibrium of the game with common p-belief.

- The above bad sequential equilibria survive the Cho-Kreps selection criterion

- Introduce an investment stage prior to the HM-MR mechanism considered above

  $\rightarrow$ S invests into increasing the good's quality (as in Che and Hausch (1999))

- Let S chooses investment i at cost $c(i)$, and let the $\Pr(v = 14) = \beta i$.

- First best benchmark:

$$\max_i \{\beta i 14 + (1-\beta i)10 - c(i)\}$$

  with first-order condition

$$4\beta = c'\left(i^{FB}\right).$$

- MR mechanisms under common knowledge implement this first-best investment by ensuring that S gets paid 14 when $v = 14$ and not more than 10 when $v = 10$.

- Is the first best investment also always approximately achieved through MR mechanisms under common p-belief?

- No

- Three steps:

  1. recall that the SCF $f(\theta) = (1, -\theta, \theta)$ that would induce ex ante first best investment by the seller, is not constant and therefore not monotonic.

  2. in common p-belief (value) perturbations ν_ε from common knowledge, ex ante the seller will choose her investment so as to

$$\max_i \left\{ \begin{array}{c} \left[\beta i(1 - \Pr(L|\theta_B = \theta'_B)) + (1-\beta i)\Pr\left(H|\theta_B = \theta''_B\right)\right] 14 \\ + \left[(1-\beta i)\left(1 - \Pr(H|\theta_B = \theta''_B)\right) + \beta i \Pr\left(L|\theta_B = \theta'_B\right)\right] 10 \\ -c(i) \end{array} \right\}.$$

3 but from Theorem 2 there exist common p-belief (value) perturbations ν_ε and a corresponding sequence of sequential equilibria where $\Pr(L|\theta_B = \theta'_B)$ remains bounded away from zero as $\varepsilon \longrightarrow 0$.

- 1. $\Longrightarrow$ this in turn implies that the equilibrium investment, defined by the first-order condition

$$4\beta \left(1 - \Pr(L|\theta_B = \theta'_B) - \Pr\left(H|\theta_B = \theta''_B\right)\right) = c'(i),$$

remains bounded away from the first-best level of investment as $\varepsilon \to 0$.

- Thus:

  $\longrightarrow$ there exist a sequence common p-belief perturbations from common knowledge that are becoming arbitrarily small as $\varepsilon \to 0$...

  $\longrightarrow$ and there exists a sequence of sequential equilibria associated to that sequence of perturbations....

  $\longrightarrow$ for which S's investment remains bounded away from the first-best level under non-integration of B and S.

  $\longrightarrow$ this restores the role for vertical integration along similar lines as in Grossman and Hart (1986).

- The same applies not only to the MR mechanism but to *any* extensive-form mechanism that implements $f(\theta) = (1, -\theta, \theta)$ under common knowledge since $f(\theta)$ is non-monotonic

9.6.4 Renegotiation Design

Aghion, Dewatripont & Rey (Econometrica, 1994)

- Suppose ω not verifiable so Arrow-Debreu contracts cannot be written – CAN STILL GET FB !
- Consider $q = \overline{q}, p = \overline{p}$ and then renegotiate at date 2
- Buyer can make an offer and if Seller accepts then trade occurs on those terms – otherwise trade takes place at $(\overline{q}, \overline{p})$
- For simplicity assume that S has all the bargaining power
- B's payoff is:

$$E_\omega \left[R(\overline{q}, \omega, e)\right] - \overline{p} - e$$

- Maximizing this w.r.t. e yields:

$$\frac{\partial E_\omega \left[R(\overline{q}, \omega, e^*)\right]}{\partial e} = 1$$

- Solve for $\overline{q}$ which exists
- Since B has all the bargaining power she will offer the ex post efficient quantity and maximize joint surplus – S will be indifferent b/w this and the default

- Anticipating getting the default S will end p choosing the optimal investment level by the construction of $\overline{q}$
- Since B has all the bargaining power she is the residual claimant on investment and therefore chooses the optimal investment conditional on S choosing the optimal investment on her side
- *Gets around the moral hazard in teams problem!*
- B is the residual claiamant
- S (more interestingly) has the right incentives because the default option gets more attractive as the cost of production goes down – which she controls
- The default introduces another instrument which allows one to target a second exogenous variable
- Key: shows that a foundation for incomplete contracts must be based on *ex post non contractibility*
- At-Will contracting is essentially a necessary condition for non-verifiability leading to incompleteness
- Frames what all the implementation literature cannot do without – *ex post non contractibility*

10 The Firm as a Subeconomy

10.1 Overview

- Based entirely on: Holmstrom, Bent (1999) "The Firm as a Subeconomy", *Journal of Law, Economics and Organizations* 15(1), pp.74-102.
- There are *lots* of important ideas in the paper.
- One way to think about it is as observing an empirical fact: asset ownership is clustered in firms–and then providing an explanation for this fact through a synthesis of a number of existing theories.
  - GHM property-rights theory: ownership of assets provides incentives for ex ante investment by the owner, and *removes* and incentive for the non-owner
  - Alchian-Demsetz (1972): the firm monitors inputs to try and solve free-rider problems.
  - Holmstrom-Milgrom (1991, 1994): incentive design in the face of imperfect measurement.
- An issue with PRT: seems like a good explanation of why *individuals* might own assets; less satisfactory account of why *firms* own assets.
- Taken to it's logical conclusion, PRT suggests that each asset should be owned separately–counterfactual

- Perhaps ownership confers contracting rights: says who can utilize a given assets, under what conditions, and on what terms.
- This leads to the "subeconomy" view: a firm is a mini-economy where HQ can regulate trade; it can
  - Assign tasks
  - Allocate authority
  - Provide explicit incentives
  - Provide implicit incentives
- Sounds a little bit like a government?!
- Firms and governments (should) do very similar things: they both structure the environment in which certain parties interact to internalize externalities
- PRT very good on incentives created (or removed) by asset ownership
- Silent on incentives within firms

10.2 Moral Hazard in Teams

- Holmstrom (Bell, 1982)
- n agents $1, ..., n$ who choose actions $a_1, ..., a_n$
- This produces revenue $q(a_1, ..., a_n)$ with $q(\cdot)$ concave
- Agent's utility function is $I_i - \psi_i(a_i)$ with $\psi_i(\cdot)$ convex
- In the first-best:
$$\max \left\{ q(a_1, ..., a_n) - \sum_{i=1}^{n} \psi_i(a_i) \right\}$$
- The FOC is:
$$\frac{\partial q}{\partial a_i} = \psi_i'(a_i) \quad , \forall i$$
- In the second-best assume that a_i is observable only to agent i but that q is observable to everyone
- A partnership consists of sharing rules $s_i(a_i), i = 1, ..., n$ such that
$$\sum_i s_i(q) \equiv q \tag{18}$$
- Might suppose that $s_i(q) \geq 0, \forall i$
- In a Nash Equilibrium each agent solves:
$$\max_{a_i} \{ s_i(q(a_i, a_{-i})) - \psi_i(a_i) \}$$

- The FOC is:
$$s_i'(q)\frac{\partial q\left(a_i,a_{-i}\right)}{\partial a_i}=\psi_i'(a_i)$$

- Need $s_i'(q)=1,\forall i$ to get the FB

- But we know from (18) that $\sum_i s_i'(q)\equiv 1$

- Can't get the FB

- Nothing to do with risk-aversion – there is no uncertainty here

- Say we introduce an $(n+1)$th party such that:

$$s_i(q)\equiv q\left(a^*\right)-F_i,\quad \forall i=1,...,n$$

$$s_{n+1}(q)=\sum_i F_i-nq\left(a^*\right)$$

- This will be profitable for the $(n+1)$th party if we pick F_i such that $\sum_{i=1}^n F_i+q\left(a^*\right)\geq nq\left(a^*\right)$

- And also profitable for the agents if $F_i\leq q\left(a^*\right)-\psi_i\left(a_i^*\right)$

- These can both be satisfied because at the FB $q\left(a^*\right)-\sum_{i=1}^n\psi_i\left(a_i^*\right)>0$

- We have made everyone the residual claimant

- However, the $(n+1)$th party wants it to fail. They might burn the factory down, ... Call them the Budget Breaker ("BB")

- They might also collude with one of the Agents

- A side contract between BB and i – this merges BB and i into one person and we are back into the n agent case

- n people could collude to "borrow" q and game the BB

- This mechanism (making everyone the residual claimant) is similar to Groves-Clarke we we saw earlier

- Alchian-Demsetz have a different solution: introduce a monitor who can observe the inputs and use a forcing contract.

- But: (i) why doesn't the monitor renneg? and (ii) what incentive does the monitor have to actually monitor?

10.3 Regulating Trade Within the Firm

- Moral hazard arises precisely because there is a hidden action–i.e. the action is imperfectly measured.
- Can happen in three ways
  1. Contractible performance depends on the agent and nature: $x = f(e, \theta)$
  2. Contractible performance measures are biased: $x = \theta e$, but output $y = e$
  3. Contractible performance measure is manipulable: $x = e + m$, where m is unobserved
- Suppose that costs of effort and manipulation are separable so that the cost to the agent is $c(e) + d(m) = \frac{1}{2}\left(e^2 + \lambda m^2\right)$, where λ measures the cost of manipulation.
- Output is $y = pe$, but x (the performance measure) is all that can be contracted on.
- Focus on linear sharing rules of the form $s(x) = \alpha x + \beta$.
- NB: in this setting that restriction is without loss of generality.
- The FOCs for the agent imply
  - $e(\alpha) = \alpha$, and
  - $m(\alpha) = \alpha/\lambda$.
- Optimal α maximizes total surplus (recall risk-neutrality).

$$\begin{aligned} TS &= p(e(\alpha)) - c(e(\alpha)) - d(m(\alpha)) \\ &= p\alpha - \frac{1}{2}\left(\alpha^2 + \alpha^2/\lambda\right). \end{aligned}$$

- This means that $\alpha^* = \alpha p/(1+\lambda)$.
- Of course, $\alpha^{FB} = p > \alpha^*$.
- Now suppose there are two tasks with efforts e_1 and e_2.
- A's costs is now $\frac{1}{2}\left(e_1^2 + e_2^2\right) + \frac{1}{2}\left(\lambda m^2\right)$
- P's benefit function is $y = p_1 R(e_1) + p_2 e_2$, where R is a strictly increasing, concave function.
- There are two performance measures: $x_1 = R(e_1)$ and $x_2 = e_2 + m$.
- P pays A according to: $s(x_1, x_2) = \alpha_1 x_1 + \alpha_2 x_2 + \beta$.
- A solves

$$\max_{e_1, e_2, m} \left\{ \alpha_1 R(e_1) + \alpha_2(e_2 + m) - \frac{1}{2}(e_1 + e_2)^2 - \frac{1}{2}\lambda m^2 \right\}.$$

- The FOCs are

  1. $\alpha_1 R'(e_1) = \alpha_2$,
  2. $e_2 = \alpha_2 - e_1$, and
  3. $m = \alpha_2/\lambda$.

- The loss from manipulation (even though the amount is known in equilm) is $\frac{1}{2}\lambda m^2 = (\alpha_2^2/2\lambda)$.
- When manipulation is easier more of it happens.
- In the FB $\alpha_1 = p_1$ and $\alpha_2 = p_2$.
- In the SB, incentives are lower-powered.
- This reduces wasteful manipulation.
- A decrease in λ makes incentive lower-powered.
- i.e. When performance measures are less informative, optimal incentives are weaker.
- Incentives affect *opportunity costs*.
- Recall DJT (1999b) and relative intensity of incentives of government agencies and private firms–though, there, for career concerns reasons.
- Now go back to the single-task version and add a parameter τ to the cost function so that A's cost is $c(e, \tau) + d(m, \tau)$.
- The FOC implies $c'(e, \tau) = \alpha = d'(m, \alpha)$.
- So as α increases, so do e and m.
- Use the same functional form for d as before .
- Since $de/d\alpha = 1/C''(\tau)$, C'' measures how responsive A's action is to the incentive coefficient α.
- Also suppose that $dC''/d\tau < 0$ and that marginal cost is increasing in τ.
- This captures the idea that A can engage in private (unobserved) activities in addition to effort e, and that total time spent on all activities affects cost of effort.
- Then τ is like to number of private tasks that are allowed–so it's a measure of freedom.
- Thus, total surplus is

$$TS(\alpha, \tau, p, \lambda) = pe(\alpha, \tau) - c(e(\alpha, \tau), \tau) - d(m(\alpha, \lambda), \lambda)$$

- Notice that

$$\frac{dTS}{d\alpha} = \frac{(p-\alpha)\partial e}{\partial \alpha} - \alpha \frac{\partial m}{\partial \alpha}$$

- Now, differentiating this wrt τ and noting that $p > \alpha$ and $\partial^2 e/\partial\alpha\partial\tau > 0$ by assumption, we obtain

$$\frac{\partial^2 TS}{\partial\alpha\partial\tau} > 0.$$

- The other cross partials are:

$$\frac{\partial^2 TS}{\partial\alpha\partial p} > 0.$$

$$\frac{\partial^2 TS}{\partial\alpha\partial\lambda} = 0.$$

$$\frac{\partial^2 TS}{\partial p\partial\tau} < 0.$$

$$\frac{\partial^2 TS}{\partial\lambda\partial\tau} = 0.$$

- TS is not supermodular in all variables.

- An increase in the cost of manipulating performance increases α and τ–better performance measurement implies more freedom and more intense incentives.

- If p increases then we need more intense incentives (higher α), a reduction in opportunity cost (lower τ), but when α goes up we need τ to go up. Ambiguity.

- Key points

  1. Contracting on separate instruments creates externalities–have to think about designing the incentive system as a whole.
  2. Indirect effects matter (because of lack of super modularity)–which exogenous variable changed matters for predictions about (say) freedom.

10.4 Asset Ownership

- When instruments are controlled by different firms it's more difficult to have coherent design

- Separate design creates contractual externalities and integration can help internalize them

- Consider 2 agents: 1 and 2

- One asset: A

- If operated by agent 1 it produces output $y_1 = R(e_1)$ and if operated by agents 2 $y_2 = e_2$.

- No manipulation is possible ($m = 0$).

- By bargaining outside the relationship agent 1 can get price p_1 for each unit y_1.
- Agent 2 is an "expert" on marketing y_2 but has no investment decision.
- Her human capital is essential for selling y_2 at price p_2 (per unit).
- Cannot contract on y_1, y_2–non verifiable.
- But can bargain.
- If agent 1 owns A then she can get $(1/2)p_2 y_2$ from agent 2 by threatening to withhold output y_2, plus she gets $p_1 R(e_1)$ by bargaining with an outside party (since producing y_1 doesn't involve agent 2).
- If agent 2 owns A she can withhold her human capital in producing y_1, y_2.
- Under 50:50 Nash Bargaining agent 1 gets $(1/2)[p_1 R(e_1) + p_2 e_2]$ (so does agent 2).
- We can now analyze things as above
- If agent 1 owns A she faces an incentive scheme with $\alpha_1 = p_1$ and $\alpha_2 = p_2/2$.
- If she doesn't own the asset then $\alpha_1 = p_1/2$ and $\alpha_2 = p_2/2$.
- Agent 2 doesn't invest.
- Agent 1's efforts when she owns A are $e_1 = p_1 R'(e_1) - e_2$ and $e_2 = (1/2)p_2 - e_1$.
- So e_1 is greater than FB and e_2 is lower than FB.
- When agent 1 doesn't own A we have $e_1 = (1/2)p_1 R'(e_1) - e_2$ and $e_2 = (1/2)p_2 - e_1$.
- Thus, $e_1^* = e_1^{FB}$, but e_2 is lower than the FB level.
- But e_2 is closer to the FB level than under the alternative ownership structure.
- So agent 1 NOT owning the asset provides superior incentives.
- Key: balanced incentives–even if low-powered–can be better than imbalanced incentives in the presence of asset ownership.
- Ownership by agent 1 can be bad for agent 1 because it diverts effort from task 2.
- Can be better for a third party to own the assets...