Notes on Contract Theory

From Dr. Vikram Maheshri's Lectures, Compiled by Brian Murphy*†

March 17, 2025

Contents

*These notes are from my time as a student in the University of Houston PhD Economics program.
†Typos may exist in these notes. If any are found, please contact me.

1 Introduction

These notes cover key concepts covered in a Ph.D. Economics Contract Theory Course, which was my Ph.D. Microeconomics II Course. The topics including adverse selection, screening, signaling, hidden action, moral hazard, as well as a variety of related literature that discusses problems of matching, stability, search problems, learning, and communication.

2 Averse Selection & Screening

2.1 Optimal Employment Contracts

– No uncertainty or hidden information.

– Employer-employee (principal-agent) relationship.

– Constraints on how the contract is set up.

2.2 Two Goods: Labor and Transfer/Money

– **Employer:** Utility function $U(l,t)$ with an initial endowment $(\hat{l}_1, \hat{t}_1) = (0,1)$ (i.e., no labor but 1 unit of money).

– **Employee:** Utility function $u(l,t)$ with an initial endowment $(\hat{l}_2, \hat{t}_2) = (1,0)$ (i.e., all labor but no money).

Here, $U(l,t)$ and $u(l,t)$ are the **representation utilities**. Assuming that both functions are increasing and strictly concave guarantees the presence of **gains from trade**.

2.3 Joint Surplus Maximization

The joint surplus maximization problem can be expressed as:

$$\max\ U(l_1, t_1) + u(l_2, t_2)$$

subject to the constraints:

$$l_1 + l_2 = \hat{l}_1 + \hat{l}_2 = 1,$$
$$t_1 + t_2 = \hat{t}_1 + \hat{t}_2 = 1.$$

A parameter can be introduced here to reflect bargaining power between the parties.

2.4 Lagrangian Formulation

We write the Lagrangian as:

$$\mathcal{L} = U(l_1, t_1) + u(l_2, t_2) - \lambda\Big(1 - l_1 - l_2\Big) - \lambda_t\Big(1 - t_1 - t_2\Big).$$

The first-order conditions (FOCs) are:

$$\frac{\partial \mathcal{L}}{\partial l_1}: \quad U_l(l_1, t_1) - \lambda = 0 \quad \Rightarrow \quad U_l = \lambda,$$

$$\frac{\partial \mathcal{L}}{\partial l_2}: \quad u_l(l_2, t_2) - \lambda = 0 \quad \Rightarrow \quad u_l = \lambda,$$

$$\frac{\partial \mathcal{L}}{\partial t_1}: \quad U_t(l_1, t_1) - \lambda_t = 0 \quad \Rightarrow \quad U_t = \lambda_t.$$

Thus, joint surplus maximization is achieved when the marginal rates of substitution (MRS) between money and leisure are equal:

$$\frac{U_l}{U_t} = \frac{u_l}{u_t},$$

confirming the presence of **gains from trade**.

2.5 Adding Uncertainty

Introduce uncertainty by considering two states: a low state (Θ_L) and a high state (Θ_H). In the high state, suppose the individual endowments change to:

$$(\hat{l}_{1H}, \hat{t}_{1H}) = (2, 1), \quad (\hat{l}_{2H}, \hat{t}_{2H}) = (2, 1).$$

The goal now is to maximize the **expected joint employment**, which leads to the condition:

$$\frac{U_H}{U_L} = \frac{u_H}{u_L},$$

ensuring that the marginal rate of substitution is consistent across states.

2.6 Coinsurance

This section considers the optimal coinsurance contract offered by the employer. The **Employer's Problem** is formulated as:

$$\max_{\ell_{ij}, t_{ij}} \quad P_L\, U(\ell_{1L}, t_{1L}) + P_H\, U(\ell_{1H}, t_{1H}), \tag{1}$$

$$\text{s.t.} \quad \ell_{1L}, \ell_{2L} \le \ell_1, \ell_{2L}, \quad t_{1L}, t_{2L} \le t_{1L}, t_{2L}, \tag{2}$$

$$P_L\, u(\ell_{2L}, t_{2L}) + P_H\, u(\ell_{2H}, t_{2H}) \ge \bar{u}, \tag{3}$$

$$\text{where } \bar{u} = P_L\, u(\ell_{2L}, t_{2L}) + P_H\, u(\ell_{2H}, t_{2H}). \tag{4}$$

The first-order conditions include:

$$\text{FOC}(\ell_{1L}): \quad P_L\, U_l(\ell_{1L}, t_{1L}) + \lambda\, \ell_L = 0, \tag{5}$$

$$\text{FOC}(\ell_{2L}): \quad P_L\, u_\ell(\ell_{2L}, t_{2L}) + \lambda\, \ell_L = 0. \tag{6}$$

Together with similar conditions for t_{1j} and t_{2j}, these imply:

$$\frac{U_\ell(\ell_{1L}, t_{1L})}{u_\ell(\ell_{2L}, t_{2L})} = \frac{U_t(\ell_{2L}, t_{2L})}{u_t(\ell_{2L}, t_{2L})}.$$

In essence, the contract is designed to minimize risk exposure and allocates risk to the more risk-loving party, thereby ensuring **ex post efficiency**.

2.7 Hidden Information in Contract Design

Employer's Utility:

$$U\Big(\ell \cdot \theta \cdot (1 - \ell) - t\Big)$$

where:

– The constant reflects the baseline level,

– θ represents productivity,

– $1 - \ell$ represents time worked,

– t is the payment.

Employee's Utility:

$$u\Big(\theta \cdot \ell + t\Big)$$

where:

– ℓ represents leisure,

– t represents income,

– θ relates to productivity information.

Assume that the employee knows θ (with $\theta \in \{\theta_L, \theta_H\}$) while the employer only knows the probabilities P_L and P_H.

2.8 Revelation Principle

The revelation principle allows us to restrict attention to contracts that specify actions for each state. In this framework, there are two contracts:

$$(\ell_L, t_L) \quad \text{and} \quad (\ell_H, t_H).$$

2.9 Effort Test Question

The **Employer's Problem** under an effort test can be formulated as:

$$\begin{aligned}
\max_{\ell_i, t_i} \quad & P_L\, U\Big(\ell \cdot \theta_L \cdot (1-\ell_L) - t_L\Big) + P_H\, U\Big(\ell \cdot \theta_H \cdot (1-\ell_H) - t_H\Big) && (7)\\
\text{s.t.} \quad & u(\theta_L, \ell_L + t_L) \geq u(\theta_L) \quad \text{(Individual Rationality for Low type)}, && (8)\\
& u(\theta_H, \ell_H + t_H) \geq u(\theta_H) \quad \text{(Individual Rationality for High type)}, && (9)\\
& u(\ell_H, \theta_H + t_H) \geq u(\ell_L, \theta_H + t_L) \quad \text{(Incentive Compatibility for High type)}, && (10)\\
& u(\ell_L, \theta_L + t_L) \geq u(\ell_H, \theta_L + t_H) \quad \text{(Incentive Compatibility for Low type)}. && (11)
\end{aligned}$$

In this setting, high-type employees receive information rents while low-type employees do not.

2.10 Hidden Action

When effort (i.e., $1-\ell$) is unobservable, the timing of events is as follows:

Step 1: The contract is chosen.

Step 2: The state θ is realized.

Step 3: The employer observes the output.

Step 4: The employer pays based on the output.

Employee's Problem: Given a contract, the employee maximizes:

$$\max_{\ell}\; P_L\,(1-\ell)\,u\Big(t(\theta_L \ell)\Big) + P_H\, u\Big(t(\theta_H \ell)\Big).$$

Employer's Problem: The employer designs the payment scheme by solving:

$$\max_{t(\theta_i)}\; P_L\,(1-\ell)\,U\Big(\theta_L \cdot t(\theta_L)\Big) + P_H\,(1-\ell)\,U\Big(\theta_H - t(\theta_H)\Big),$$

subject to the participation constraint:

$$P_L\,(1-\ell)\,u\Big(t(\theta_L)+\ell\Big) + P_H\,(1-\ell)\,u\Big(t(\theta_H)+\ell\Big) \geq \bar{u} = u(\ell),$$

and ensuring that the employee's effort choice satisfies the incentive compatibility condition.

2.11 Adverse Selection in a Simple Model of Exchange

1/23

"Often seen with Health Insurance" happens b/c of hidden info

Agent — Buyer

$$u(q, T, \theta) = \int_0^q p(x, \theta)dx - T$$

↑ ↑ ↑

units purchased   total cost/amount paid   type/preference

Simplify

$$u(q, T, \theta) = \theta v(q) - T,$$

with $v(0) = 0$, $v' > 0$ and $v'' < 0$.
· the higher the θ, the higher the marginal utility.

Assume $\theta \in \{\theta_L, \theta_H\}$ with $\theta_H > \theta_L$.

Principal — Seller

$$\Pi = T - cq,$$

↑

unit production costs

First Degree (Perfect) Price Discrimination

Assume Seller observes θ. You can tell if θ is high or low type. Offer contracts (T_i, q_i), $i \in \{L, H\}$. Assume buyer's reservation utility is $\bar{u}$.
Seller's Problem:

$$\max_{T_i, q_i} T_i - cq_i \quad \text{s.t.} \quad \theta_i v(q_i) - T_i \geq \bar{u}$$

(IR constraint)

The Lagrangian is

$$\mathcal{L} = T_i - cq_i + \lambda_i\Big(\theta_i v(q_i) - T_i - \bar{u}\Big).$$

Differentiating with respect to T_i gives:

$$1 - \lambda_i = 0 \quad \Longrightarrow \quad \lambda_i = 1.$$

Differentiating with respect to q_i gives:

$$-c + \lambda_i \theta_i v'(q_i) = 0 \quad \Longrightarrow \quad v'(q_i^*) = \frac{c}{\theta_i}.$$

Then,

$$T_i^* = \theta_i v(q_i^*) - \bar{u}.$$

Thus, the high type buys more and pays more. There is no surplus since the constraint is set equal to $\bar{u}$.

Adverse Selection

Now, assume the Seller does **NOT** observe the buyer's type. The set of possible contracts is $T(q)$; any contract can be offered.

Linear Pricing

Seller only picks p. Starting with the Buyer:

$$\max_q \quad \theta_i v(q) - pq, \tag{12}$$

$$\text{FOC:} \quad \theta_i v'(q) = p, \tag{13}$$

so,

$$q_i^* = v'^{-1}(p/\theta_i) = D_i(p) \quad \text{(Demand, which is increasing).} \tag{14}$$

Seller's Problem becomes:

$$\max_p \quad (p-c)[\beta D_L(p) + (1-\beta)D_H(p)] \quad \text{(expected demand } = D(p)). \tag{15}$$

The FOC is:

$$D(p) + (p-c)D'(p) = 0,$$

which implies:

$$p^* = c - \frac{D(p^*)}{D'(p^*)} > c.$$

This is the monopoly price. Buyer surplus is given by:

$$S_i(p) = \theta_i v(D_i(p)) - pD_i(p).$$

Thus, the Seller always sells to high types and may or may not sell to low types, depending on the values of θ and β.

2-Part Tariff (Single)

The Seller charges a fee plus a per unit price. Let the fee be z and the per unit price be p, provided no arbitrage exists. For any p,

$$z = S_L(p),$$

i.e. the surplus of the low-type is taken as the fee so that the low type is indifferent to buying the product. If

$$T(q) = z + pq = S_L(p) + pq,$$

then the high-type buyer purchases a positive quantity and obtains surplus.

Seller's Problem

$$\max_p \quad S_L(p) + (p-c)D(p).$$

The FOC is:

$$S_L'(p) + D(p) + (p-c)D'(p) = 0, \quad \text{so that} \quad p^* = c - \frac{D(p) + S_L'(p)}{D'(p)} > c.$$

Since $S_L'(p) = -D_L(p)$, for the low type the numerator is zero while for the high type it is positive. This leads to a lower price so that consumption is closer to the full consumption level. However, there is under-consumption relative to the perfect price discrimination case.

Question: Can we do better if we draw in a non-linear pricing mechanism?

Optimal Non-Linear Pricing Problem

Seller's Problem:

$$\max_{T(q)} \quad \beta\Big(T(q_L) - cq_L\Big) + (1-\beta)\Big(T(q_H) - cq_H\Big)$$

subject to

$$q_i \in \arg\max_q \Big\{\theta_i v(q) - T(q)\Big\} \quad \text{for } i = L, H \quad \text{(Incentive Compatibility)}$$

and

$$\theta_i v(q_i) - T(q_i) \geq 0 \quad \text{(Individual Rationality)}.$$

Solve in 5 Steps:

1. Apply the revelation principle: Without loss of generality, we can restrict $T(q)$ to the set $\{(T(q_L), q_L), (T(q_H), q_H)\}$. That is, it suffices to consider contracts for the two types as long as they are incentive compatible.

2. Write the constraints:

$$\theta_H v(q_H) - T_H \geq \theta_H v(q_L) - T_L \geq \theta_L v(q_L) \geq 0.$$

   Thus, the IR constraint for the high type is not needed because the combination of the high-type incentive constraint (ICH) and the low-type individual rationality (IRL) imply it.

3. Delete one of the IC constraints and optimize the problem; then check if the deleted constraint is satisfied. (The high type may want to mimic the low type, so we ignore the low-type IC constraint initially.)

4. At the optimum, both ICH and IRL bind. If $\theta_H v(q_H) - T_H$ is strictly greater than required, one could raise the price T_H to capture more surplus. Similarly, T_L is increased until IRL binds.

5. With these constraints binding, substitute the equality constraints into the objective:

$$\max_{q_L, q_H} \quad \beta\Big(\theta_L v(q_L) - cq_L\Big) + (1-\beta)\Big(\theta_H v(q_H) - cq_H - (\theta_H - \theta_L)v(q_L)\Big).$$

   The term $\theta_L v(q_L) - cq_L$ represents the full surplus of the low type, which is fully appropriated by the seller, while the high type's surplus is maximized minus the information rent. If too much surplus is taken from the high type, the high type will mimic the low type.

FOC for q_L:

$$\beta\theta_L v'(q_L) - c - (1-\beta)(\theta_H - \theta_L)v'(q_L) = 0,$$

or equivalently,

$$\theta_L v'(q_L^*) = \frac{c}{1 - \frac{1-\beta}{\beta} \cdot \frac{\theta_H - \theta_L}{\theta_L}} > c.$$

FOC for q_H:

$$(1-\beta)\Big(\theta_H v'(q_H) - c\Big) = 0 \quad \Rightarrow \quad \theta_H v'(q_H^*) = c.$$

This implies that $q_L^* < q_H^*$. With the ICH binding and $q_L^* < q_H^*$, the ICL holds. In summary, while there is a distortion at the lower end, the allocation for the high type is efficient.

2.12 Optimal Non-Linear Pricing

2.12.1 Two Types

We will not be asked to solve these problems on tests or comprehensive exams; most questions describe the story rather than require detailed calculations.

2.12.2 Step-by-Step Approach

Step 0) Write out the seller's problem (maximizing the objective over contracts, subject to buyer individual rationality constraints).

Step 1) Apply the revelation principle. This simplifies the contracts and, consequently, the buyer's incentive compatibility (IC) constraints.

Step 2) The IC_H constraint can be eliminated because IC_H combined with IR_L implies IC_H.

Step 3) Delete the IC_L constraint initially (since the high type may mimic the low type) and check it later.

Step 4) Note that IC_H and IR_L bind at the optimum.

Step 5) Solve by unconstrained maximization.
This approach can be extended to models with three or four types.

2.12.3 Optimal Non-Linear Pricing with a Continuum of Types

Assume that $\Theta \sim F(\theta)$ with probability density function $f(\theta)$ on $[\underline{\theta}, \bar{\theta}]$. The distribution $F(\theta)$ is known to the seller, but the individual type θ is private information.

Seller's Problem:

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

subject to the incentive compatibility condition:

$$q(\theta) = \arg\max_{q} \Big\{\theta v(q) - T(q)\Big\} \quad \forall\theta,$$

and the individual rationality condition:

$$\theta v(q(\theta)) - T(q(\theta)) \geq 0.$$

1) Apply the Revelation Principle

$$\max_{T(\theta), q(\theta)} \int_{\underline{\theta}}^{\bar{\theta}} \Big[T(q(\theta)) - cq(\theta)\Big] f(\theta) d\theta,$$

subject to the constraints:

$$\theta v(q(\theta)) - T(\theta) \geq \theta v(q(\hat{\theta})) - T(\hat{\theta}) \quad \forall\theta, \hat{\theta} \in \Theta \quad (\mathrm{IC}_\theta),$$

and

$$\theta v(q(\theta)) - T(q(\theta)) \geq 0 \quad \forall\theta \quad (\mathrm{IR}_\theta).$$

2) Simplify the IR Constraints

$$\theta v(q(\theta)) - T(\theta) \geq \theta v(q(\underline{\theta})) - T(\underline{\theta}) \geq \underline{\theta} v(q(\underline{\theta})) - T(\underline{\theta}) \geq 0.$$

Thus, it suffices to impose the constraint:

$$\underline{\theta} v(q(\underline{\theta})) - T(\underline{\theta}) \geq 0.$$

3) Strategy: First, determine which contracts are implementable (i.e., which satisfy incentive compatibility). Then, choose the best contract among those.

Lemma 1: If

$$\frac{\partial}{\partial \theta}\left[\frac{v'(q)}{v''(q)}\right] > 0,$$

then the incentive compatibility constraint IC_θ holds if and only if

$$\frac{dq(\theta)}{d\theta} \geq 0 \quad \text{(monotonicity)},$$

and the local incentive compatibility (LIC) condition holds:

$$T'(\theta) = \theta v'(q(\theta)) \frac{dq(\theta)}{d\theta}.$$

In other words, the marginal rate of substitution must be increasing in type.

Monotonicity: This requires that the quantity $q(\theta)$ increases in type.

Local Incentive Compatibility: The first-order condition for the buyer's problem yields

$$\theta v'(q(\hat{\theta})) \frac{dq(\hat{\theta})}{d\hat{\theta}} - T'(\hat{\theta}) = 0 \quad \text{at } \hat{\theta} = \theta,$$

which confirms local incentive compatibility.

Envelope Theorem: Since the buyer maximizes utility, the envelope theorem implies that

$$w'(\theta) = v(q(\theta)),$$

where $w(\theta) = \theta v(q(\theta)) - T(\theta)$ represents the buyer's welfare in equilibrium. Integrating yields

$$w(\theta) = \int_{\underline{\theta}}^{\theta} v(q(x))dx + w(\underline{\theta}),$$

and because the IR constraint binds at $\underline{\theta}$, we have $w(\underline{\theta}) = 0$. Since $T(\theta) = \theta v(q(\theta)) - w(\theta)$, the seller's profit can be rewritten as:

$$\Pi = \int_{\underline{\theta}}^{\bar{\theta}} \Big[\theta v(q(\theta)) - w(\theta) - cq(\theta)\Big] f(\theta) d\theta \tag{16}$$

$$= \int_{\underline{\theta}}^{\bar{\theta}} \Big[\theta v(q(\theta)) - \int_{\underline{\theta}}^{\theta} v(q(x))dx - cq(\theta)\Big] f(\theta) d\theta. \tag{17}$$

Integration by Parts: Let

$$u(\theta) = \int_{\underline{\theta}}^{\theta} v(q(x))dx \quad \text{and} \quad dv = f(\theta)d\theta.$$

Then, by integration by parts,

$$\int_{\underline{\theta}}^{\bar{\theta}} v(q(\theta))(1 - F(\theta))d\theta$$

appears in the expression for Π. Thus, the seller's problem becomes pointwise:

$$\text{FOC: } v'(q(\theta))\Big[\theta - \frac{1 - F(\theta)}{f(\theta)}\Big] = c,$$

which implies that only the highest types receive the full surplus.

Note: The term $\frac{1-F(\theta)}{f(\theta)}$ represents the inverse hazard rate. An increasing hazard rate guarantees the monotonicity of $q(\theta)$.

3 Signaling

3.1 Signaling/Informed Principal Problem

Model Setup

- Informed principal, uninformed agent.
- Spence (1973, 1974) Model of Education:
  - Two types of workers: L and H, with productivity r_L for low and r_H for high types, where $r_H > r_L > 0$.
  - Workers are willing to work at any wage $w > 0$.
  - The firm holds a prior belief $\beta_i \in [0, 1]$ that $r = r_L$.
  - The firm is willing to hire workers at any wage below expected productivity.
  - A worker of type i can obtain e years of education at a cost

    $$C(e) = \theta_i \cdot e,$$

    where $\theta_H < \theta_L$ (i.e. the marginal cost for an extra year is lower for high types).
  - Years of education is observable (*single crossing condition*).
  - Education has no effect on productivity.
- **Stage 1:** The worker chooses e.
- **Stage 2:** The wage is chosen by bargaining. (Assume the worker holds all bargaining power.)

Benchmark (Full Info)

- $e_H = e_L = 0$, $w_L = r_L$, $w_H = r_H$ (the firm pays a wage equal to productivity).

Hidden Info

Let $\beta(\theta_i \mid e)$ be the agent's posterior belief about productivity. Then the wage function is given by:

$$w(e) = \beta(\theta_H \mid e)r_H + \beta(\theta_L \mid e)r_L \quad (18)$$

$$= \beta(\theta_H \mid e)r_H + \Big(1 - \beta(\theta_H \mid e)\Big)r_L. \quad (19)$$

Game Theory Concepts

– A game is an interaction between two parties.

– For example, in chess there are players, payoffs (e.g., 1 for a win, 0 for a loss), actions, and strategies (mappings from states to actions).

– **Nash Equilibrium:** A situation where no player has a unilateral incentive to deviate.

Posterior

Let $P_i(e)$ denote the probability that worker i obtains e years of education.

Since games can be dynamic, strategies might be mappings from histories to the action space. This necessitates the use of uncertainty and beliefs; hence, Bayesian Equilibria are applied.

Perfect Bayesian Equilibrium in Signaling Games

Definition 1 *A* ***Perfect Bayesian Equilibrium*** *in a signaling game is a set of (possibly mixed) strategies $P_i(e)$ for the principal's types and conditional beliefs $\beta(\theta_i \mid e)$ for the agent such that:*

1. *All levels of e observed with positive probability in equilibrium maximize the worker's expected payoff, i.e.,*

$$e \in \arg\max_{e'} \Big\{ \beta(\theta_H \mid e') r_H + \beta(\theta_L \mid e') r_L - \theta_i e' \Big\},$$

   which represents wage minus the cost of education.

2. *Firms pay workers their expected productivity:*

$$w(e) = \beta(\theta_H \mid e) r_H + \beta(\theta_L \mid e) r_L.$$

3. *Firms' posterior beliefs satisfy Bayes' Rule whenever $P_i(e) > 0$ for some i.*

4. *Posteriors are otherwise unrestricted.*

Reasonable Guess

A reasonable guess is to set:

$$\beta(\theta_H \mid e) = 1 \quad \text{if } e \geq \hat{e}, \quad 0 \text{ otherwise.}$$

Thus, one might have:

$$P_H(e) = 1 \quad \text{if } e = \hat{e}, \quad 0 \text{ otherwise,}$$

and

$$P_L(e) = 0.$$

Types of Equilibria

Generally, there are three types of equilibria:

1. **Separating:** The signal e_i identifies the type exactly.

2. **Pooling:** The signal e_i reveals no information about type.

3. **Semi-Separating:** The signal e_i reveals some, but not all, information.

Separating Equilibrium

Definition 2 *A **separating equilibrium** is a Perfect Bayesian Equilibrium in which $e_H \neq e_L$ and the beliefs satisfy*

$$\beta(\theta_H \mid e_H) = \beta(\theta_L \mid e_L) = 1,$$

so that the observed action reveals the worker's type and wages are set as $w_i = r_i$.

For the low type, the constraint implies:

$$e_L = 0.$$

For the high type, the incentive compatibility conditions are:

– $r_H - \theta_H e \geq r_L \quad (IC_H)$.

– $r_H - \theta_L e \leq r_L \quad (IC_L)$.

Solving these gives:

$$e \leq \frac{r_H - r_L}{\theta_H} \quad \text{and} \quad e \geq \frac{r_H - r_L}{\theta_L}.$$

Thus, a separating equilibrium requires

$$e_H^* \in \left[\frac{r_H - r_L}{\theta_L}, \frac{r_H - r_L}{\theta_H}\right].$$

Note that high types incur an education cost relative to the benchmark where no education is provided.

Equilibrium Continuum

There is a continuum of equilibria depending on beliefs:

$$\beta(\theta_H \mid e) = 1 \quad \forall e \geq e_H,$$

all of which are consistent with Bayes' Law.

Pooling Equilibrium

Definition 3 *A **pooling equilibrium** is a PBE where*

$$e_H = e_L = e_P,$$

with beliefs $\beta(\theta_H \mid e_P) = \beta_H$ and $\beta(\theta_L \mid e_P) = \beta_L$.

Then the wage is given by:

$$w(e_P) = \beta_H r_H + \beta_L r_L = \bar{r}.$$

Incentive compatibility requires, for example,

$$\bar{r} - \theta_L e_P \geq r_L \quad (IC_L).$$

Beliefs are assumed to be

$$\beta(\theta_H \mid e) = \beta_H \quad \forall e \geq e_P, \quad 0 \text{ otherwise},$$

and similarly for $\beta(\theta_L \mid e)$. Thus, there exists a pooling equilibrium with

$$e_P \in \left[0, \frac{\beta_H r_H + \beta_L r_L - r_L}{\theta_L}\right].$$

A selection mechanism is needed to choose among multiple equilibria.

Equilibrium Selection

To choose between PBEs, refinements such as the **Cho-Kreps Intuitive Criterion** are applied.

Definition 4 *Let*

$$U_i^* = w_i^*(e_i) - \theta_i e_i$$

denote the equilibrium payoff for worker i. The Cho-Kreps Intuitive Criterion then states that for any deviation e such that

$$r_H - \theta_j e < U_j^* \quad \textit{and} \quad r_H - \theta_i e \geq U_i^*$$

for some $j \neq i$, the posterior belief should satisfy

$$\beta(\theta_j \mid e) = 0.$$

In other words, if a deviation is strictly dominated for one type but not for another, the dominated type should not be assigned positive probability off the equilibrium path. Applying this to pooling equilibria typically eliminates them because there always exists a profitable deviation for the high type.

Educational Productivity

What if education *is* productive? In that case, r also reflects productivity, so higher education implies higher productivity. Then there exists a maximum education level in the interval

$$\left[\frac{r_H - r_L}{\theta_L}, \frac{r_H - r_L}{\theta_H}\right].$$

Applying the Cho-Kreps incentive constraint to separating equilibria, we have:

$$e_L = 0 \quad \text{and} \quad e_H \in \left[\frac{r_H - r_L}{\theta_L}, \frac{r_H - r_L}{\theta_H}\right].$$

A calculation shows that only

$$e_H^* = \frac{r_H - r_L}{\theta_L}$$

survives, meaning high types acquire just enough education to signal their type. Even if the probability of being low type, β_L, is very small, signaling remains wasteful but important.

3.2 The Market for Lemons — Akerlof (1970)

"You get this thing but you don't know if it sucks or not. A lemon is a piece of poor quality."

1. Product quality may be heterogeneous.
2. Sellers are better judges of quality than buyers.

Stylized Model

Preliminarily, assume there are four types of cars:

1. Good, with probability q.
2. Bad, with probability $1 - q$.
3. New.
4. Old.

Here, q is common knowledge and, after owning a car, its quality (good or bad) is revealed.

Other Assumptions

– Both good and bad cars must sell at some price (buyers cannot distinguish quality).

– Used cars cannot have the same valuation as new cars.

– If they were valued equally, trading would continue until only good cars remain.

– This implies that most traded used cars are lemons.

– Consequently, the market may unravel.

Generalized Model

Define:

$$p \quad \text{– price of a car,} \tag{20}$$

$$\mu \quad \text{– average quality of a car} \quad \Rightarrow \quad Q^D = D(p, \mu), \quad \mu \leq \mu(p). \tag{21}$$

There are two groups of traders who maximize expected utility:

$$U_1 = M + \sum x_i, \tag{22}$$

$$U_2 = M + \sum \frac{3}{2} x_i, \tag{23}$$

where M denotes all other consumption and x_i is the quality of the ith car. Assume Group 1 owns N cars and Group 2 owns none, with car quality $x \sim U[0, 2]$. Let the price of M be 1 and Y_i be the income of a type i trader.

The demand by Group 1 is:

$$D_1 = \begin{cases} \frac{Y_1}{p} & \text{if } \mu \geq p, \\ 0 & \text{if } \mu < p. \end{cases}$$

The supply offered by Group 1 to Group 2 is:

$$S_1 = \frac{p \cdot N}{2} \quad \text{if } p \leq 2,$$

with

$$\mu = \frac{p}{2}.$$

The demand by Group 2 is:

$$D_2 = \begin{cases} \frac{Y_2}{p} & \text{if } \frac{3}{2}\mu \geq p, \\ 0 & \text{if } \frac{3}{2}\mu < p. \end{cases}$$

Supply by Group 2 is zero. Hence, total demand is:

$$D(p) = \begin{cases} \frac{Y_1 + Y_2}{p} & \text{if } p \leq \mu, \\ \frac{Y_2}{p} & \text{if } \mu < p \leq \frac{3}{2}\mu, \\ 0 & \text{if } \frac{3}{2}\mu < p. \end{cases}$$

At price p, since $\mu = \frac{p}{2}$, we have

$$Q^D = \begin{cases} \frac{Y_1 + Y_2}{p} & \text{if } p \leq \frac{p}{2}, \\ \frac{Y_2}{p} & \text{if } \frac{p}{2} < p \leq \frac{3p}{4}, \\ 0 & \text{if } \frac{3p}{4} < p. \end{cases}$$

Because $\frac{p}{2} \leq p$ and $\frac{3p}{4} \leq p$ always hold, the model suggests we are always in a no-trade zone. Nevertheless, for any $p \in [0, 2]$ there exists some seller-buyer combination that would like to trade. Continuous updating of beliefs may eventually lead to a situation where the only car offered is of zero quality. This is a manifestation of adverse selection and is analogous to issues observed in health insurance markets—one reason for mandates or subsidies.

3.3 Milgrom & Roberts — Advertisements

"What if expenditure on ads were observed?"
Then, ad expenditures serve as a signal of quality. One might ask, *"Can we learn about quality from prices?"*

Model:

- A firm produces a product of quality $q \in \{H, L\}$.
- The firm knows its own quality, while the consumer does not.
- The firm chooses both a price p and advertising expenditures A.
- Only the firm's initial choices and the resulting consumer beliefs (not the dynamic aspects) are analyzed.
- Let $\ell(p, A)$ denote the consumer's belief that $q = H$.
- The profit function $\pi(p, q, \ell)$ is such that ads affect gross profits only indirectly via ℓ (i.e., ads do not affect demand directly).

Let p^q denote the profit-maximizing price for a firm producing quality q when consumers believe the product is of quality q. Define

$$\pi(p, q, L) = \pi(p, q, 0) \quad \text{and} \quad \pi(p, q, H) = \pi(p, q, 1).$$

Under full information, the firm chooses $p = p^q$ and $A = 0$.

Proposition 1: A separating equilibrium exists if and only if there exists $(p, A) \geq 0$ such that the low-quality firm does not advertise while the high-quality firm advertises. Specifically, the incentive constraints are:

$$\pi(p, H, H) - A \geq \pi(p_L^H, H, L) \quad (IC_H)$$

and

$$\pi(p, L, H) - A \leq \pi(p_L^L, L, L) \quad (IC_L).$$

Together, this implies:

$$\pi(p, H, H) - \pi(p_L^H, H, L) \geq A \geq \pi(p, L, H) - \pi(p_L^L, L, L).$$

In any separating equilibrium, the low-quality firm chooses $(p_L^L, 0)$ while the high-quality firm selects some (p, A) so that $\ell(p, A) = 1$ and $\ell(p_L^L, 0) = 0$, thereby satisfying incentive compatibility.

Proposition 2: Considering Kreps' strategies, there exists a separating equilibrium if and only if there exists (p, A) satisfying Proposition 1. In any separating equilibrium, the pair (p_H, A_H) must solve:

$$\max_{p,A} \quad \pi(p, H, H) - A, \tag{24}$$

$$\text{s.t.} \quad \pi(p, L, H) - A \leq \pi(p_L^L, L, L), \quad p, A \geq 0. \tag{25}$$

If $A^* > 0$, then the optimal price p^* solves:

$$\max_p \quad \Big\{\pi(p, H, H) - \pi(p, L, H)\Big\},$$

which reflects the idea that advertising induces consumers to believe the firm is of high quality. The constraint

$$\pi(p, L, H) - \pi(p_L^L, L, L) > 0$$

ensures incentive compatibility for the low type. Notice that in this formulation the variable A drops out because if a positive expenditure is optimal then the marginal benefit equals the marginal cost.

3.4 Screening Paper — Torture & the Commitment Problem

Consider a planned terrorist attack. A suspect, who may possess valuable information, is interrogated. The suspect could be either uninformed (or innocent) or the planner (informed). Note that any rational initiation of torture can be used to justify its continuation.

Commitment Problems:

1. The interrogator (P) cannot commit to stopping torture, reducing the suspect's incentive to reveal information.
2. Even if, after extensive torture, it becomes clear that the suspect is innocent, the incentive to stop torture may be insufficient, thereby encouraging the suspect to remain silent.

Thus, an informed suspect will generally not reveal information in order to prompt a cessation of torture.

What is the Value of Torture? It has an upper bound; in equilibrium the value of torture can be zero.

Model:

- **Torturer (P)**
- **Suspect (A):** Informed with probability μ_0 (possessing information quantity x) and uninformed with probability $1 - \mu_0$.
- In continuous time, torture imposes a flow cost of Δ on the suspect and a cost $c > 0$ on the torturer.
- **Ticking Time Bomb:** If the suspect reveals $z \leq x$ and is tortured for $\tau \leq T$, then A's payoff is

  $$-z - \Delta\tau,$$

  while P's payoff is

  $$z - c\tau.$$

Infinite Horizon:

P demands y units of information and commits to torturing for a duration $\tau(y)$ if the suspect does not confess. One must account for the possibility that an informed suspect may act as if they were uninformed. P's payoff becomes:

$$\mu_0 y - (1 - \mu_0)c\tau(y).$$

For an informed agent, the incentive constraint (ICA) requires that:

$$y \leq A\tau(y).$$

When this binds, we have $y = A\tau(y)$. Assuming that $\tau(y) = \frac{y}{\Delta}$, P's payoff simplifies to:

$$\mu_0 y - (1-\mu_0)\frac{cy}{\Delta} = y\left(\mu_0 - (1-\mu_0)\frac{c}{\Delta}\right). \tag{26}$$

If the term in parentheses is positive, then the optimal is $y = x$ (i.e., the suspect reveals all information). Hence, P can commit to a menu of torture durations:

Theorem 1 — Full Commitment (Infinite Horizon): If $\mu_0\Delta - (1-\mu_0)c \geq 0$, then P demands $y = \min\{x, T\Delta\}$. (27)

and tortures for

$$\min\left\{\frac{x}{\Delta}, T\right\}.$$

In this case, (1) innocent or uninformed suspects are always tortured (since they cannot provide the demanded information), and (2) informed suspects will reveal all their information if the time horizon is long enough. Without a mechanism to exempt the innocent, even the guilty might have no incentive to disclose information.

Limited Commitment:
If P is not bound to a precommitted torture duration (i.e., P can extend torture at will):

– If the suspect provides some information, P may try to extract more.

– If the suspect does not provide information, P may intensify the torture.

– P will only stop when the cost becomes prohibitive.

Time is divided into discrete periods. Suppose P demands y_t in period t and commits to suspending torture if the suspect complies:

$$\text{If the suspect complies:} \quad \text{P's payoff } u_t = y_t, \quad \text{A's payoff } v_t = -y_t; \tag{28}$$

$$\text{If the suspect does not comply:} \quad u_t = -c, \quad v_t = -\Delta. \tag{29}$$

With a common discount factor δ, the total payoffs are:

$$U = (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} u_t, \tag{30}$$

$$V = (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} v_t. \tag{31}$$

Multiple equilibria may arise. In a worst-case equilibrium, regardless of how high μ_0 or the demanded information z is, there exists an equilibrium in which the value of torture is zero. This outcome is driven by the commitment problem: actions taken today are not credible for the future. If A confesses, the continuation game becomes one of complete information and P, knowing that A is informed, may renegotiate to extract the remaining information $x - y$. Consequently, A might never confess.

When x is uncertain, the problem worsens. One can establish a lower bound for the value of torture (which is zero, independent of μ_0) and an upper bound by recasting the problem as a finite game. In finite games, backward induction reveals that in the last period the suspect can reduce the value of torture, leading to negative payoffs. Thus, in settings involving screening, commitment problems inherently arise since actions today may not align with future actions.

4 Hidden Action & Moral Hazard

4.1 Overview

In settings with hidden action, the principal (P) must design a contract that induces the agent (A) to behave as if his effort were perfectly observable. Key points include:

– P hires A to perform a task.

– P cares only about performance.

– P chooses the agent's effort, which is both unobservable and costly.

– Performance increases with effort.

Optimal Contract Considerations:

– Employees desire a guaranteed baseline. When incentives are added, the principal must "sacrifice" part of the marginal product to provide the guarantee.

– Classic examples include insurance and sharecropping.

– Often a copay or fixed component is introduced.

4.2 Two Outcome Case

Assume that the performance measure q takes on only two values: $q = 0$ (failure) and $q = 1$ (success). The agent chooses an action a (effort) such that performance increases in a. We suppose that

$$P(q = 1 \mid a) = p(a), \quad \text{with } p'(a) > 0,\ p''(a) < 0, \tag{32}$$

$$p(0) = 0, \quad p(\infty) = 1, \quad \text{and} \quad p'(0) = 1. \tag{33}$$

The utilities are given by:

– **Principal's Utility:** $V(q - w)$.

– **Agent's Utility:** $u(w) - \Psi(a)$, with $u'(w) > 0$, $\Psi'(a) > 0$, and $\Psi''(a) \geq 0$.

For simplicity, we assume that $\Psi(a) = a$.

4.2.1 Benchmark: Full Information

When effort is observable, the principal solves:

$$\max_{a,w_0,w_1} \; p(a)V(1 - w_1) + (1 - p(a))V(0 - w_0) \tag{34}$$

$$\text{s.t.} \quad p(a)u(w_1) + (1 - p(a))u(w_0) - a \geq \bar{u} \quad \text{(IR)}. \tag{35}$$

The corresponding Lagrangian is

$$\begin{aligned} \mathcal{L} = {} & p(a)V(1 - w_1) + (1 - p(a))V(0 - w_0) \\ & + \lambda\Big[p(a)u(w_1) + (1 - p(a))u(w_0) - a - \bar{u}\Big]. \end{aligned} \tag{36}$$

The first-order conditions (FOCs) with respect to w_0, w_1, and a are:

$$\frac{\partial \mathcal{L}}{\partial w_0}: \quad -(1-p(a))V'(0-w_0) + \lambda(1-p(a))u'(w_0) = 0, \tag{37}$$

$$\frac{\partial \mathcal{L}}{\partial w_1}: \quad -p(a)V'(1-w_1) + \lambda p(a)u'(w_1) = 0, \tag{38}$$

$$\frac{\partial \mathcal{L}}{\partial a}: \quad p'(a)\Big[V(1-w_1) - V(0-w_0)\Big] + \lambda\Big[u(w_1) - u(w_0)\Big] - \lambda = 0. \tag{39}$$

The *Borch Rule* (or optimal coinsurance condition) implies that the marginal utilities in each state are equalized. In this benchmark, one can solve for the optimal a^* once the multiplier λ is pinned down.

4.2.2 Example: Risk Neutral Principal

Assume that the principal is risk neutral so that $V(x) = x$ (and hence $V'(x) = 1$). Then, from the FOCs we obtain

$$\frac{1}{u'(w_0)} = \frac{1}{u'(w_1)} \quad \Longrightarrow \quad w_0 = w_1.$$

This implies full insurance for the agent. Solving for a then involves the condition

$$p'(a^*)\Big[1 - (w_1 - w_0) + \lambda(u(w_1) - u(w_0))\Big] - \lambda = 0.$$

With full insurance ($w_1 = w_0$), if we set $\lambda = 1/u'(w^*)$ (with $w^* = w_1 = w_0$), the marginal product (MP) of effort equals the marginal cost (MC) for the principal.

4.2.3 Example: Risk Neutral Agent

Now assume that the agent is risk neutral, so that $u(x) = x$ and $u'(x) = 1$. In this case,

$$\lambda = V'(w_0) = V'(1-w_1) \quad \Longrightarrow \quad w_0 = 1 - w_1,$$

or equivalently,

$$w_1 - w_0 = 1.$$

This differential in wages reflects the value of a successful project. In addition, the agent's incentive condition gives

$$p'(a)\big[\underbrace{V(1-w_1) - V(-w_0)}_{=0} + \lambda \underbrace{(w_1 - w_0)}_{=1}\big] - \lambda = 0,$$

so that

$$p'(a) = \frac{1}{\text{MP of effort}},$$

which in turn equals the marginal cost of effort from the agent's perspective. When the action a is unobservable, an incentive compatibility (IC) constraint must be added:

$$\max_{a, w_1, w_0} \; p(a)V(1-w_1) + (1-p(a))V(0-w_0) \tag{40}$$

$$\text{s.t. } p(a)u(w_1) + (1-p(a))u(w_0) - a \geq 0, \tag{41}$$

$$a \in \arg\max_{\hat{a}} \Big\{p(\hat{a})u(w_1) + (1-p(\hat{a}))u(w_0) - \hat{a}\Big\}. \tag{42}$$

The agent's IC condition yields

$$p'(\hat{a})\Big[u(w_1) - u(w_0)\Big] = 1.$$

4.3 Case: Bilateral Risk Neutrality

In the case where both parties are risk neutral, suppose $p'(a^*) = 1$; then from the agent's IC we have

$$u(w_1) - u(w_0) = 1,$$

implying $w_1 - w_0 = 1$. This outcome corresponds to the first-best allocation. In practice, one might implement this by selling the output up front to the agent (subject to a resource constraint). For example, if the agent is resource constrained so that $w_0 \geq 0$, one may set $w_0 = 0$ and $w_1 = 1$. However, this may leave the principal with zero expected profit, which is undesirable. To generate profit, the principal might set $w_0 = 0$ and then solve:

$$\max_{w_1} p(a)(1 - w_1) \tag{43}$$

$$\text{s.t. } p'(a)w_1 = 1 \quad \text{(IC)}. \tag{44}$$

Analysis of the FOCs shows that when $p'(a) > 1$, the optimal effort satisfies $a^* < 1$; that is, the agent exerts less than the first-best level of effort due to the resource constraint in the contract.

4.4 Incorporating Risk Aversion

4.4.1 Bilateral Risk Aversion

If the agent is risk neutral and the principal is risk averse, the optimal contract might involve selling the output up front. Conversely, if only the agent is risk averse while the principal remains risk neutral, a constant (fixed) wage results—offering no incentive for effort.

4.4.2 True Bilateral Case

When both parties are risk averse, the Lagrangian for the principal's problem becomes

$$\begin{aligned} \mathcal{L} = & p(a)V(1 - w_1) + (1 - p(a))V(-w_0) \\ & + \lambda^{IR}\Big[p(a)u(w_1) + (1 - p(a))u(w_0) - a - \bar{u}\Big] \\ & + \lambda^{IC}\Big[p'(a)(u(w_1) - u(w_0)) - 1\Big]. \end{aligned} \tag{45}$$

Taking the FOCs with respect to w_0 and w_1 gives, respectively,

$$\frac{\partial \mathcal{L}}{\partial w_0} : \quad -(1 - p(a))V'(-w_0) + \lambda^{IR}(1 - p(a))u'(w_0) - \lambda^{IC}p'(a)u'(w_0) = 0, \tag{46}$$

$$\frac{\partial \mathcal{L}}{\partial w_1} : \quad -p(a)V'(1 - w_1) + \lambda^{IR}p(a)u'(w_1) + \lambda^{IC}p'(a)u'(w_1) = 0. \tag{47}$$

These conditions can be rearranged into:

$$\frac{V'(-w_0)}{u'(w_0)} = \lambda^{IR} - \lambda^{IC}\frac{p'(a)}{1 - p(a)}, \tag{48}$$

$$\frac{V'(1 - w_1)}{u'(w_1)} = \lambda^{IR} + \lambda^{IC}\frac{p'(a)}{p(a)}. \tag{49}$$

When $\lambda^{IC} = 0$, we recover the Borch Rule. In the presence of risk aversion, however, one typically finds that $\lambda^{IC} > 0$, reflecting the additional cost of providing incentives under uncertainty.

4.5 Multidimensional Tasks

Consider now the case where the agent performs two tasks with independently distributed outputs. Suppose the agent's cost function is

$$\Psi(a_1, a_2) = \frac{1}{2}\Big(C_1 a_1^2 + C_2 a_2^2\Big) + \delta\, a_1 a_2, \tag{50}$$

with

$$0 \le \delta \le \sqrt{C_1 C_2}.$$

When $\delta = 0$, the tasks are technologically independent; if $\delta > 0$, raising effort in one task increases the marginal cost of the other (they are substitutes). A typical linear contract takes the form

$$w = t + S_1 q_1 + S_2 q_2.$$

The principal's problem is

$$\max_{q_1, q_2} E\Big[q_1 + q_2 - w\Big] \quad \text{s.t.} \quad E\Big[-\exp\Big(-\eta\Big(w - \tfrac{1}{2}(C_1 a_1^2 + C_2 a_2^2) - \delta a_1 a_2\Big)\Big)\Big] \ge \overline{w}, \tag{51}$$

where $\overline{w}$ represents the agent's certainty equivalent compensation:

$$\overline{w} = t + S_1 a_1 + S_2 a_2 - \frac{\eta}{2}\Big(S_1^2 \sigma_1^2 + S_2^2 \sigma_2^2\Big) - \frac{1}{2}\Big(C_1 a_1^2 + C_2 a_2^2\Big) - \delta a_1 a_2. \tag{52}$$

The agent chooses efforts a_1 and a_2 to maximize $\overline{w}$. The first-order conditions for the agent are:

$$S_1 - C_1 a_1 - \delta a_2 = 0, \tag{53}$$
$$S_2 - C_2 a_2 - \delta a_1 = 0. \tag{54}$$

Solving these equations yields

$$a_1 = \frac{S_1 C_2 - \delta S_2}{C_1 C_2 - \delta^2}, \quad a_2 = \frac{S_2 C_1 - \delta S_1}{C_1 C_2 - \delta^2}. \tag{55}$$

The principal then chooses piece rates S_1 and S_2 (subject to an individual rationality constraint) to maximize his expected profit. Analysis shows that if it becomes harder to monitor one task (e.g., σ_2^2 increases), then the optimal piece rate S_2 decreases, and by complementarities, S_1 may also fall.

4.6 Multidimensional Incentives: The CEO Example

Principal-Agent Framework

– **Principal:** The company's shareholders.

– **Agent:** The CEO.

Model Setup Assume that:

– Profits: $q = a + \varepsilon_q$ with $\varepsilon_q \sim N(0, \sigma_q^2)$.

– Stock price: $p = a + \varepsilon_p$ with $\varepsilon_p \sim N(0, \sigma_p^2)$.

– Covariance between q and p is given by $\sigma_{qp} = \mathrm{Cov}(q, p)$.

Common effort a generates correlation between outcomes.

Preferences and Compensation Structure The CEO's utility is assumed to be

$$u(w,a) = -\exp\Big(-\eta\Big[w - \psi(a)\Big]\Big), \quad \text{with } \psi(a) = \tfrac{1}{2}ca^2.$$

The principal is risk neutral. A typical compensation package is

$$w = t + sq + fp,$$

where:

– t is a fixed salary.

– sq is a bonus tied to profits.

– fp is the CEO's share of the firm's equity.

Principal's Problem The principal maximizes expected profit:

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

$$\text{s.t. } E\Big[-\exp\Big(-\eta\Big[w - \tfrac{1}{2}ca^2\Big]\Big)\Big] \geq -\exp\Big(-\eta\underline{w}\Big) \quad \text{(IR)} \tag{56}$$

$$a \in \arg\max_a \; E\Big[-\exp\Big(-\eta\Big[w - \tfrac{1}{2}ca^2\Big]\Big)\Big] \quad \text{(IC)}. \tag{57}$$

Solving the agent's problem yields the certainty equivalent wage:

$$\tilde{w} = t + (s+f)a - \frac{ca^2}{2} - \frac{\eta}{2}\Big[s^2\sigma_q^2 + f^2\sigma_p^2 + 2sf\sigma_{qp}\Big]. \tag{58}$$

The IR constraint then becomes

$$\tilde{w} \geq \underline{w},$$

and the IC condition implies that

$$a^* = \frac{s+f}{c}.$$

Principal's Optimization Substituting the binding IR constraint to eliminate t, the principal's problem reduces to choosing s and f to maximize

$$\max_{s,f} \; \frac{s+f}{c} - t - (s+f)\frac{s+f}{c} + \frac{c}{2}\left(\frac{s+f}{c}\right)^2 - \underline{w} \tag{59}$$

which can be rearranged into

$$\max_{s,f} \; (1-(s+f))\,\frac{s+f}{c} + \frac{\eta}{2}\Big[s^2\sigma_q^2 + f^2\sigma_p^2 + 2sf\sigma_{qp}\Big] - \underline{w}. \tag{60}$$

The first-order conditions with respect to s and f yield:

$$\text{FOC}(s): \quad \frac{1-(s+f)}{c} + \eta\Big[s\sigma_q^2 + f\sigma_{qp}\Big] = 0, \tag{61}$$

$$\text{FOC}(f): \quad \frac{1-(s+f)}{c} + \eta\Big[f\sigma_p^2 + s\sigma_{qp}\Big] = 0. \tag{62}$$

Subtracting these two conditions gives

$$f\Big(\sigma_p^2 - \sigma_{qp}\Big) = s\Big(\sigma_q^2 - \sigma_{qp}\Big) \quad \Longrightarrow \quad \frac{f}{s} = \frac{\sigma_q^2 - \sigma_{qp}}{\sigma_p^2 - \sigma_{qp}}. \tag{63}$$

After further manipulation and letting

$$K = \sigma_q^2 \sigma_p^2 - \sigma_{qp}^2,$$

one obtains the optimal shares:

$$s^* = \frac{\sigma_p^2 - \sigma_{qp}}{K\Big(1 - c\eta K\Big)}, \tag{64}$$

$$f^* = \frac{\sigma_q^2 - \sigma_{qp}}{K\Big(1 - c\eta K\Big)}. \tag{65}$$

Special Cases

– **Case 1:** If $\sigma_{qp} = 0$ (no correlation between stock price and revenues),

$$s^* = \frac{\sigma_p^2}{\sigma_q^2 \sigma_p^2 + \eta c \sigma_q^2 \sigma_p^2}, \tag{66}$$

$$f^* = \frac{\sigma_q^2}{\sigma_q^2 \sigma_p^2 + \eta c \sigma_q^2 \sigma_p^2}. \tag{67}$$

As σ_p^2 increases, s^* increases (more bonus on profits) while f^* decreases.

– **Case 2:** For a risk neutral CEO ($\eta \to 0$), one finds $f^* \to 1$ and $s^* + f^* \to 1$. In other words, the CEO becomes the sole residual claimant.

– **Case 3:** Under noisy measurement, suppose

$$\varepsilon_p = \varepsilon_q + v, \quad v \sim N(0, \sigma_v^2),$$

so that $\sigma_p^2 = \sigma_q^2 + \sigma_v^2$ and $\sigma_{qp} = \sigma_q^2$. Then one obtains $f^* = 0$, meaning no equity share is given, and all incentive is provided via the bonus s^*.

Key Takeaways:

– The optimal structure of CEO compensation depends critically on the relative variances and covariances of the performance measures.

– A higher volatility in stock price (larger σ_p^2) induces a higher profit bonus s^*.

– When performance is measured more precisely (larger σ_q^2), the optimal equity share f^* increases.

5 Moral Hazard with Adverse Selection

5.1 Main Problem

Consider a buyer (Agent A) with two possible types, θ_H and θ_L, where $\theta_L < \theta_H$. The seller (Agent P) holds the prior belief that

$$P(\theta = \theta_H) = \beta.$$

The revenue is uncertain, with the output $X \in \{0, R\}$. Revenue R is generated with probability $e \cdot \theta$, where the cost of effort is given by

$$P(e) = \frac{ce^2}{2}.$$

The seller offers contracts specified by the pair (r_i, t_i), where r_i denotes the royalty and t_i is a transfer payment.

We begin by analyzing the buyer's problem before returning to the seller's (principal's) problem.

Buyer's Optimization Problem. For a given contract (r_i, t_i), the buyer maximizes:

$$\max_{e} \quad \theta_i e(R - r_i) - t_i - \frac{ce^2}{2}. \tag{68}$$

The first-order condition (FOC) for effort is:

$$\theta_i(R - r_i) - ce = 0, \tag{69}$$

$$\Rightarrow \quad e^* = \frac{\theta_i(R - r_i)}{c}. \tag{70}$$

Substituting e^* back into the buyer's payoff yields:

$$E(\text{payoff}) = \frac{[\theta_i(R - r_i)]^2}{c} - t_i - \frac{[\theta_i(R - r_i)]^2}{2c}, \tag{71}$$

$$= \frac{[\theta_i(R - r_i)]^2}{2c} - t_i. \tag{72}$$

Seller's (Principal's) Problem. Recognizing that the seller's problem depends on the buyer's type, we have:

$$\max_{t_i, r_i} \quad \beta\left(t_H + \frac{\theta_H^2(R - r_H)r_H}{c}\right) + (1 - \beta)\left(t_L + \frac{\theta_L^2(R - r_L)r_L}{c}\right), \tag{73}$$

$$\text{s.t.} \quad \frac{[\theta_i(R - r_i)]^2}{2c} - t_i \geq 0 \quad \text{(Individual Rationality, IR);} \tag{74}$$

$$\frac{[\theta_i(R - r_i)]^2}{2c} - t_i \geq \frac{[\theta_i(R - r_j)]^2}{2c} - t_j \quad \text{(Incentive Compatibility, IC).} \tag{75}$$

How Do We Solve This?

Case: No Moral Hazard. When the buyer's type is known, the seller's problem simplifies to:

$$\max_{t_i, r_i} \quad t_i + \frac{\theta_i^2(R - r_i)r_i}{c}, \tag{76}$$

$$\text{s.t.} \quad \frac{[\theta_i(R - r_i)]^2}{2c} - t_i \geq 0 \quad \text{(IR).} \tag{77}$$

By substituting the IR constraint, we obtain an optimization problem in r_i only:

$$\max_{r_i} \quad \frac{[\theta_i(R - r_i)]^2}{2c} + \frac{\theta_i^2(R - r_i)r_i}{c} \tag{78}$$

$$= \frac{\theta_i^2(R - r_i)}{c}\left[r_i + \frac{R - r_i}{2}\right] \tag{79}$$

$$= \frac{\theta_i^2(R - r_i)}{c} \cdot \frac{r_i + R}{2} \quad \Rightarrow \quad \frac{\theta_i^2}{c}\left(R^2 - r_i^2\right). \tag{80}$$

Taking the first-order condition with respect to r_i:

$$-\frac{2\theta_i^2}{c} r_i = 0 \quad \Rightarrow \quad r_i = 0.$$

Thus, in the absence of moral hazard, there is no need to provide effort incentives; the contract is priced solely based on type. The corresponding transfer is then:

$$t_i = \frac{\theta_i^2 R^2}{2c}.$$

Fixed Effort Case (No Moral Hazard). Suppose the effort level is fixed at $\hat{e}$ and is observable. The seller's problem becomes:

$$\max_{t_i, r_i} \quad \beta \left(t_H + \theta_H \hat{e}\, r_H\right) + (1-\beta)\left(t_L + \theta_L \hat{e}\, r_L\right), \tag{81}$$

$$\text{s.t.} \quad \theta_i \hat{e}(R - r_i) - t_i - \frac{c\hat{e}^2}{2} \geq 0, \tag{82}$$

$$\theta_i \hat{e}(R - r_i) - t_i \geq \theta_i \hat{e}(R - r_j) - t_j. \tag{83}$$

Here, we may ignore the IR constraint for the high type and the IC constraint for the low type, since the IR constraint for the low type and the IC constraint for the high type are binding. Thus, we have:

$$\text{IR}_L: \quad \theta_L \hat{e}(R - r_L) - t_L - \frac{c\hat{e}^2}{2} \geq 0 \quad \Longrightarrow \quad t_L = \theta_L \hat{e}(R - r_L) - \frac{c\hat{e}^2}{2}, \tag{84}$$

$$\text{IC}_H: \quad \theta_H \hat{e}(R - r_H) - t_H \geq \theta_H \hat{e}(R - r_L) - t_L \quad \Longrightarrow \quad t_H = \theta_H \hat{e}(R - r_H) - (\theta_H - \theta_L)\hat{e}(R - r_L) - \frac{c\hat{e}^2}{2}. \tag{85}$$

The seller's objective function can then be written as:

$$\max_{r_H, r_L} \quad \beta \Big[\theta_H \hat{e}(R - r_H) - (\theta_H - \theta_L)\hat{e}(R - r_L) - \frac{c\hat{e}^2}{2} + \theta_H \hat{e}\, r_H\Big] + (1-\beta)\Big[\theta_L \hat{e}(R - r_L) - \frac{c\hat{e}^2}{2} + \theta_L \hat{e}\, r_L\Big]. \tag{86}$$

One approach is to set $r_H = R$ and $r_L = R$, with the corresponding transfers simplifying to

$$t_i = -\frac{c\hat{e}^2}{2}.$$

In this configuration, the entire return is allocated to the seller, and the firm's effort is identical regardless of ownership. The low type becomes indifferent, while the high type secures an information rent. When a royalty is imposed, t_H decreases whereas t_L remains unchanged.

General Case. Returning to the original problem and assuming that the IR constraint for the high type and the IC constraint for the low type are non-binding, we impose the binding constraints for the low type and the high type:

$$\text{IR}_L: \quad t_L = \frac{[\theta_L (R - r_L)]^2}{2c}, \tag{87}$$

$$\text{IC}_H: \quad t_H = \frac{\theta_H^2 (R - r_H)^2 - (\theta_H^2 - \theta_L^2)(R - r_L)^2}{2c}. \tag{88}$$

Substituting these expressions into the seller's objective function yields:

$$\max_{r_H, r_L} \quad \beta \left[\frac{\theta_H^2 (R - r_H)^2 - (\theta_H^2 - \theta_L^2)(R - r_L)^2}{2c} + \frac{\theta_H^2 (R - r_H) r_H}{c}\right] + (1-\beta)\left[\frac{\theta_L^2 (R - r_L)^2}{2c} + \frac{\theta_L^2 (R - r_L) r_L}{c}\right]. \tag{89}$$

The first-order condition (FOC) with respect to r_H is:

$$\frac{\beta\theta_H^2}{c}(R - r_H) + \frac{\beta\theta_H^2}{c}(R - 2r_H) = 0,$$

which implies $r_H = 0$. In this case, the business is effectively sold at a high price with the distortion arising through the royalty. Although efficiency is not completely lost, an information rent is still paid. The payoff to the high type is:

$$\frac{(\theta_H^2 - \theta_L^2)(R - r_L)^2}{2c}.$$

Similarly, the first-order condition with respect to r_L simplifies to:

$$\frac{\beta(\theta_H^2 - \theta_L^2)}{c}(R - r_L) - \frac{(1-\beta)\theta_L^2}{c}r_L = 0,$$

yielding the optimal royalty for the low type as:

$$r_L^* = \frac{\beta(\theta_H^2 - \theta_L^2)R}{\beta(\theta_H^2 - \theta_L^2) + (1-\beta)\theta_L^2}, \quad \text{with } 0 < r_L^* < R.$$

Thus, the high type captures the full surplus by charging as much as possible, while the distortion is introduced through the royalty imposed on the low type, reducing the return from effort for the latter.

5.2 Former Comp Question: Dynamic Moral Hazard

We now present a related problem involving dynamic moral hazard in a simple, two-period setting.

Problem Setup. An employee works over two periods, $t = 1, 2$, and chooses an effort level $e \in \{e_L, e_H\}$. In period 0, the cost of effort is specified as:

$$c(e_L) = 0, \quad c(e_H) = C.$$

There are two possible outputs, q_H and q_L, with:

$$P(q_t = q_H \mid e_t = e_H) = \pi_H \quad \text{and} \quad P(q_t = q_L \mid e_t = e_L) = \pi_L.$$

The contract offered is a pair of wages, $\{w_1(q_1), w_2(q_1, q_2)\}$.

Employee's Utility and Firm's Profits. The employee's utility is given by:

$$u(w_1) + u(w_2) - c(e), \quad \text{with } u'(w) > 0 \text{ and } u''(w) < 0 \quad \text{(reflecting risk aversion)}, \tag{90}$$

and the firm's profits are:

$$q_1 + q_2 - w_1 - w_2.$$

First-Best (Effort-Contingent) Case. Assume that the objective is to implement the high effort e_H. In the first-best case (where wages can be conditioned on effort), the problem is:

$$\max_{w_1, w_2} \quad E\left[q_1 + q_2 - w_1 - w_2 \,\middle|\, e_H\right], \tag{91}$$

$$\text{s.t.} \quad E\left[u(w_1) + u(w_2) - C \,\middle|\, e_H\right] \geq \bar{u}, \tag{92}$$

where $\bar{u}$ represents the employee's reservation utility. The corresponding Lagrangian is:

$$\mathcal{L} = \pi_H q_H + (1-\pi_H)q_L + \pi_H q_H + (1-\pi_H)q_L - w_1 - w_2 + \lambda\left[u(w_1) + u(w_2) - C - \bar{u}\right].$$

The first-order conditions yield:

$$u'(w_1) = u'(w_2) = \frac{1}{\lambda}.$$

Unobservable Effort Case (Moral Hazard). In the case where effort is unobservable and the seller (principal) still wishes to implement e_H, the problem becomes:

$$\max_{w_1, w_2} \quad E\left[q_1 + q_2 - w_1 - w_2 \,\middle|\, e_H\right], \tag{93}$$

$$\text{s.t.} \quad E\left[u(w_1) + u(w_2) - C \,\middle|\, e_H\right] \geq \bar{u} \quad \text{(IR)}, \tag{94}$$

$$E\left[u(w_1) + u(w_2) - C \,\middle|\, e_H\right] \geq E\left[u(w_1) + u(w_2) - C \,\middle|\, e_L\right] \quad \text{(IC)}. \tag{95}$$

Thus, the moral hazard problem is the same as the first-best problem but with an additional incentive compatibility constraint.

Wage Structure. The wage functions are specified as:

$$w_1(q_L), \quad w_1(q_H), \quad w_2(q_L, q_L), \quad w_2(q_H, q_H), \quad w_2(q_L, q_H),$$

with the property that the wages for the outcomes $w_2(q_H, q_L)$ are set equal to the corresponding wages in the symmetric outcome.

The Lagrangian for this problem is:

$$\mathcal{L} = \pi_H q_H + (1 - \pi_H) q_L - w_1 - w_2 + \lambda\left[\pi_H u\right]$$

6 Related Literature to Contract Theory

6.1 Gale & Shapley (1962): College Admissions & Stability of Marriage

Problem: How do we match two sides of a market when the two sides have incomplete information about each other?

College Admissions Setting

Consider n applicants and a college with a quota of q_i available slots. The fundamental question is: What should the college do?

- **Offer admission to the top q_i applicants:**
  - *Issue:* The college does not know if the applicants applied elsewhere.
  - *Issue:* The college does not know how applicants rank the college, nor where else they have been admitted.
- One might ask whether students should be required to rank colleges and whether they might lie.
- **Waitlisting:** This mechanism can help avoid the above problems.
  - *New issue:* If a student is waitlisted at a more preferred college, will they accept the offer from the waitlisted college?

Assignment Algorithm Design

Consider n applicants and m colleges, where q_i is the quota of the i^{th} college. The inputs are as follows:

- Each applicant ranks the colleges (omitting only those they would never attend).
- Similarly, each college ranks the applicants (omitting only those it would never admit, even if the quota is not reached).

The goal is to combine these two rankings and create an assignment mechanism that satisfies certain criteria.

Notation and Stability Criteria

Let colleges be denoted by A and B, and applicants by α and β. The preferences can be expressed as:

$$A \succ_\alpha B \qquad \beta \succ_\alpha \alpha, \tag{96}$$

$$B \succ_\beta A \qquad \alpha \succ_\beta \beta. \tag{97}$$

Thus, not everyone is equally satisfied.

Definition of Unstable Assignment

An assignment is **unstable** if there exist two applicants, α and β, assigned respectively to colleges A and B, such that

$$A \succ_\beta B \quad \text{and} \quad \beta \succ_A \alpha.$$

Definition of Optimal Assignment

An assignment is **optimal** if every applicant is at least as well off as under any other stable assignment. In elementary economics, a second or third choice is always suboptimal. (For example, the first choice is optimal for men and the second for women.)

Marriage Stability Theorem

Theorem 1: There exists a stable set of marriages.

Proof (Constructive): The proof proceeds via the following algorithm:

1. Each man proposes to his favorite woman. Every woman who receives two or more proposals rejects all but her favorite proposal. (Note: Women do **not** immediately accept their favorite proposal.)

2. Rejected men propose to their second choice. Women reject all proposals except their favorite among the new proposals along with those from the first stage.

3. Continue this process until every woman has received at least one proposal. (This terminates in $\leq n^2 - 2(n+1)$ steps.)

4. Finally, match each man with a viable woman.

Proof of Stability

Suppose that an applicant α is not matched with a college A even though α prefers A to his assigned college B. Since α proposed to A before proposing to B, A must have rejected α, implying that A preferred some other applicant to α—one to whom she could have been matched. Hence, the matching is stable.

College Deferred Acceptance Algorithm

1. Students apply to their first-choice college.

2. Each college i waitlists its top q_i applicants and rejects all others.

3. Rejected students then apply to their second-choice colleges; colleges clear their waitlists.

4. The process repeats until every student is either waitlisted or rejected.

Theorem 2: Every applicant is at least as well off under the deferred acceptance mechanism as under any other stable mechanism.

Proof by Induction: Define a college as *possible* for an applicant if there exists a stable assignment that assigns the applicant to that college. Assume that, up to a given point, no applicant has been rejected from a possible college. Suppose college A receives applications from applicants $\beta_1, \ldots, \beta_q$ and rejects α. We wish to show that A is impossible for α. Since each β_i prefers A to any college that rejected them (an impossible college), considering an assignment that sends $\alpha \to A$ implies at least one β_i is assigned to a less desirable college. If such an assignment is unstable, then A is impossible for α. Thus, deferred acceptance only rejects impossible colleges.

Key Idea: One can reject offers without being forced to accept proposals immediately. In environments with limited capacity, students have no incentive to lie—they strategically distribute their applications rather than applying solely to their top choice.

6.2 Roth (1982): The Economics of Matching: Stability & Incentives

Gale and Shapley did not address incentives. Roth (1982) extends the literature by considering both stability and incentive properties of matching procedures. The following theorems are central:

Theorem 1 There exists a procedure that delivers a stable outcome (as demonstrated by Gale & Shapley).

Theorem 2 There exists a stable outcome that is efficient (i.e., optimal) from the perspective of one side of the market.

Theorem 3 (Negative Result) No stable matching procedure exists in general for which truthful revelation of preferences is a dominant strategy for all agents.

Theorem 4 There exist efficient matching procedures for which truth-telling is a dominant strategy for all agents.

Theorem 5 For any matching procedure that yields the optimal, stable outcome (e.g., deferred acceptance) for one side, truthful revelation is a dominant strategy for that side. In other words, there is no incentive for, say, men to lie; furthermore, the other side (women) has no incentive to misrepresent their top choice.

Notation: Men and Women

Let

$$M = \{m_1, m_2, \ldots, m_n\} \quad \text{and} \quad W = \{w_1, w_2, \ldots, w_n\}$$

denote the sets of men and women, respectively. Each man m_i has a strict, complete, and transitive preference relation $P(m_i)$ over W. The notation

$$w_k P(m_i) w_j$$

means that man m_i prefers woman w_k to w_j.

An outcome (or matching) is given by

$$X : M \to W, \quad X = \{(m_1, X(m_1)), (m_2, X(m_2)), \ldots\},$$

where $X(m_i) = w_j$ implies that man m_i is matched with woman w_j. Because the matching is one-to-one, it is invertible, so that

$$X^{-1}(w_j) = m_i$$

denotes the man matched with w_j.

The matching X is **stable** if there is no pair (m_k, w_l) such that:

$$w_l P(m_k) X(m_k) \quad \text{and} \quad m_k P(w_l) X^{-1}(w_l).$$

Proof Sketch of Theorem 3

The proof shows that there exists a matching problem for which no stable procedure ensures truthful revelation for all agents. Consider the following:

- Let $M = \{m_1, m_2, m_3\}$ and $W = \{w_1, w_2, w_3\}$.
- Let h be an arbitrary stable procedure such that for any preference profile P, the outcome $h(P)$ is stable (i.e., $h(P) \in C(P)$, where $C(P)$ is the set of all stable outcomes). Suppose $h(P)$ equals one of two outcomes, say X or Y, with men preferring one and women preferring the other.
- Now, consider a modified profile P' where w_1 changes her ranking from $(1, 3, 2)$ to $(1, 2, 3)$. In this case, the unique stable outcome is Y, so $h(P') = Y$.
- Next, consider another modified profile P'', where m_1 changes his ranking from $(2, 1, 3)$ to $(2, 3, 1)$, yielding $h(P'') = X$.

If $h(P) = X$, then w_1 has an incentive to report P'; if $h(P) = Y$, then m_1 has an incentive to misreport as in P''. Thus, there exists an incentive to misrepresent preferences.

Proof Sketch of Theorem 4

Consider a draft mechanism:

- Each man picks his top available partner (i.e., the partner who has not yet been selected by any previous man).
- Since men select their best available option, they cannot do better by lying. However, because women's preferences are not considered in this mechanism, the outcome is not stable.

This demonstrates that stability and efficiency on one side can be achieved. Theorem 5 extends this result by showing that in mechanisms yielding the optimal stable outcome for one side, truth-telling is a dominant strategy for that side.

Proof Sketch of Theorem 5

Consider the matching procedure introduced by Gale and Shapley:

- It suffices to show that for any agent m_i, for any deviation in reported preferences (i.e., any misrepresentation), the outcome $Y = g(P')$ is not preferred over the outcome $X = g(P)$, where P is the true preference profile.
- **Step 1:** It is sufficient to consider simple misrepresentations (i.e., deviations affecting only the top preference).
- **Step 2:** Show that no man is worse off if any man makes a simple misrepresentation.
- **Step 3:** Prove that no simple misrepresentation is successful; that is, any deviation results in the same matching.

A key lemma (Lemma 1) is used: If $Y = g(P')$ and $Z = g(P'')$, then for any man m_i, $Z(m_i) = Y(m_i)$. The proof relies on the stability and optimality of the matching and concludes that any misrepresentation does not yield a better outcome.

Big Picture: In matching theory, payoffs and preferences are fixed, and stability is determined by the "double coincidence of wants" in decentralized markets. However, centralized mechanisms and transferable utility (where one partner can compensate the other) allow for social planner optimization of total surplus.

6.3 Becker (1973)

Becker (1973) argues that analyzing marriage through the lens of transferable utility is insightful. Define:

$$u^m_{ij} = \text{utility that man } i \text{ derives from matching with } j,$$

$$u^f_{ij} = \text{utility that woman } j \text{ derives from matching with } i.$$

The focus is on the joint surplus:

$$\bar{z}_{ij} = u^m_{ij} + u^f_{ij}.$$

Question: Given $\bar{z}_{ij}$, what are the stable assignments?

Answer: A stable assignment maximizes the total output (i.e., total surplus).

Define the adjusted surplus as:

$$\bar{z}_{ij} - \bar{z}_{i0} - \bar{z}_{0j},$$

where $\bar{z}_{0j}$ is the utility for woman j when single, and $\bar{z}_{i0}$ is the utility for man i when single.

Planner's Problem

The problem is formulated as:

$$\begin{aligned}
\max_{a_{ij}} \quad & \sum_{i=0}^{M}\sum_{j=0}^{N} a_{ij}\bar{z}_{ij} \\
\text{s.t.} \quad & \sum_{j=0}^{N} a_{ij} = 1 \quad \forall\, i, \\
& \sum_{i=0}^{M} a_{ij} = 1 \quad \forall\, j.
\end{aligned}$$

Rewriting the problem (noting that $a_{0j} = 1 - \sum_{i=1}^{M} a_{ij}$), we have:

$$\begin{aligned}
\max_{a_{ij}} \quad & \sum_{i=1}^{M}\sum_{j=1}^{N} a_{ij}\bar{z}_{ij} \\
\text{s.t.} \quad & \sum_{j=1}^{N} a_{ij} \le 1 \quad \forall\, i, \\
& \sum_{i=1}^{M} a_{ij} \le 1 \quad \forall\, j.
\end{aligned}$$

This formulation is a linear program with $a^*_{ij} \in \{0, 1\}$. Note that maximizing utility and minimizing expenditure yield the same solution.

Dual Problem

The dual problem addresses how to divide the surplus:

$$\begin{aligned} \min_{u_i, v_j} \quad & \sum_{i=1}^{M} u_i + \sum_{j=1}^{N} v_j \\ \text{s.t.} \quad & u_i + v_j \geq \bar{z}_{ij} \quad \forall\, i, j. \end{aligned}$$

Here, the dual variables u_i^* and v_j^* (with $u_i \geq 0$ and $v_j \geq 0$) represent the shadow prices of the respective constraints. In effect, u_i is the share of the surplus that man i obtains in a stable match, while v_j is the share that woman j obtains.

A match occurs if and only if

$$u_i + v_j = \bar{z}_{ij},$$

which implies that any alternative match would cost more and leave the agents worse off. Consequently, u_i and v_j serve as the reservation utilities for men and women, respectively, leading naturally into models that incorporate search and expectations.

6.4 Maskin & Diamond (1979)

Maskin and Diamond (1979) address the information problem inherent in matching, particularly when agents face a decision about whether to continue searching for better partners. In their paper, *An Equilibrium Analysis of Search & Breach of Contract – Part 1*, they consider the following:

- Searching for potential partners is both costly and random.
- Negotiations occur instantaneously.
- A contract is an agreement to undertake a project with certain returns, which may be either good or bad.
- Thus, an agent can find themselves in a good partnership, a bad partnership, or remain single.

Key Observation: Agents can continue searching after partnering, and contracts can be breached. However, breach is costly and requires offering damages that are compensatory (and possibly liquidated, though the focus here is on compensatory damages).

Search Technologies: The analysis distinguishes between linear and quadratic search technologies.

Model Setup

- There are two types of agents. A partnership requires one agent of each type.
- The value of a good match is X (with $X > X'$), while a bad match has value X' (with $X' > 0$).
- New partners are found only through a search process that incurs a cost C.
- The meeting process follows a Poisson process.
- In the quadratic search model, the parameter a represents the probability that two searchers meet per unit time.

Conditional on a meeting, let Pr(Bad Match | Meeting) = P. The inflow of agents is given by $a \cdot b$ (with a higher b indicating a higher inflow). The population is segmented into:

- M: agents not in a match,
- N: agents in bad matches,
- Good matches, once formed, are stable and the agents do not leave.

When two M-type agents meet and form a good match, they split the surplus $2X$ equally. When two M-type agents form a bad match, they split $2X'$ equally or compensate the other party if a breach occurs. The compensation is given by:

$$D = V_M - V_N,$$

where V_M is the expected value of being an M-type agent and V_N is that of an N-type agent (with $V_M > V_N$). Thus, two M-type agents always match (they always obtain a positive payoff).

Consider:

1. When two N-type agents meet in a good match, breach occurs if

$$2X - 2V_N > 2D.$$

   The surplus from a new match is:

$$S = 2X - 2V_N - 2D = 2X - 4V_N + 2V_M,$$

   which is positive if the above condition holds.

2. When one M-type and one N-type agent meet in a good match, breach leads to a surplus:

$$2X - V_N - V_M - D > 0,$$

   which simplifies to $2X - 2V_N > 0$ if search is costly, and then the surplus is split.

Positional Value:

- For an M-type agent: $V_M + X - V_N$.
- For an N-type agent: $V_N + X - V_N = X$.

Since an N-type agent can have an M-type partner cover half the damages, matching with an M-type is more advantageous, implying no damages and a preference for an M-type partner.

Steady-State Analysis

The steady-state population dynamics for M and N are analyzed under three cases:

Case A: N-types are searching and breach if a good match occurs, while M-types are also searching.

Case B: N-types are searching, breach only if a good match occurs, and M-types are searching.

Case C: Only M-types are searching.

Case C: Let h_m denote the number of searching M-types. The dynamics are:

$$\dot{h}_m = ab - ah_m^2.$$

At steady state, $\dot{h}_m = 0$, so

$$h_m^* = \sqrt{b}.$$

Case B: The dynamics for h_m remain:

$$\dot{h}_m = ab - ah_m^2.$$

For h_N (searching N-types),

$$\dot{h}_N = ah_m P - 2\left(ah_m h_N(1-P)\right).$$

At steady state, we have $h_m = \sqrt{b}$ and

$$h_N = \frac{P}{2(1-P)}\sqrt{b}.$$

Case A: The dynamics are given by:

$$\dot{h}_m^* = ab + ah_N^2(1-P) - ah_m^2,$$

$$\dot{h}_N^* = aPh_m^2 - 2ah_m h_N(1-P) - 2a(1-P)h_N^2.$$

In steady state, the ratio is:

$$\frac{h_m^*}{h_N^*} = \left(\frac{1-P}{P}\right)\left(1+\sqrt{\frac{1+P}{1-P}}\right),$$

with

$$h_m^A \approx \sqrt{b} \quad \text{and} \quad h_m^A + h_N^A = \sqrt{\frac{b}{1-P}}.$$

The derivation uses the fact that

$$\dot{h}_m + \dot{h}_N = ab - a(1-P)(h_m+h_N)^2.$$

It is shown that a steady state corresponding to Case B cannot exist. Particular details matter, as general assumptions are hard to justify given the model's sensitivity.

6.5 Weitzman (1979) – Optimal Search for Best Alternative

Segway from Search to Learning

Weitzman (1979) addresses a problem in which a decision maker faces several opportunities with unknown rewards:

- The rewards of opportunities are uncertain and can be resolved at a cost over time.
- Searching is sequential.
- The decision maker may stop searching at any time and choose an opportunity.

The core question is: What is the optimal search strategy?

Example: Consider two projects, A and B, with a discount rate of 10%. The following table summarizes their rewards and costs:

	A	B	C (reservation)
Reward	100 (1/2), 55 (1/2)	240 (1/5), 0 (4/5)	0
Cost	15	20	0

The research durations are 1, 2, and 0, respectively. The expected value (E.V.) computations are as follows:

$$\text{E.V. if I research A:} \quad = -15 + \left(\frac{1}{1.1}\right)\left[\frac{1}{2}(100) + \frac{1}{2}(55)\right] = 55.5, \tag{98}$$

$$\text{E.V. if I research B:} \quad = -20 + \left(\frac{1}{1.1}\right)^2\left[\frac{1}{5}(240) + \frac{4}{5}(0)\right] = 19.7. \tag{99}$$

Suppose project A is researched first and yields a payoff of 55. Should project B be developed next?

$$-20 + \left(\frac{1}{1.1}\right)^2\left[\frac{1}{5}(240) + \frac{4}{5}(55)\right] = 56 > 55. \tag{100}$$

Thus, it is optimal to develop project B after A.

The overall expected value if A is researched first is:

$$\text{E.V. if I research A first:} \quad = -15 + \left(\frac{1}{1.1}\right)\left[\frac{1}{2}(100) + \frac{1}{2}\left(-20 + \left(\frac{1}{1.1}\right)^2\left[\frac{1}{5}(240) + \frac{4}{5}(55)\right]\right)\right] = 55.9. \tag{101}$$

Conversely, if B is researched first:

$$-20 + \left(\frac{1}{1.1}\right)^2\left[\frac{1}{5}(240) + \frac{4}{5}\left(-15 + \left(\frac{1}{1.1}\right)\left[\frac{1}{2}(100) + \frac{1}{2}(55)\right]\right)\right] = 56.3. \tag{102}$$

Thus, it is optimal to research project B first, indicating that one should research opportunities with high upside early.

Pandora's Problem: n Closed Boxes

Consider n boxes, where each box i contains a reward $X_i \sim F_i(x_i)$ (the boxes are independent). Opening box i costs C_i, and the reward is revealed after a lag of t_i. Let X_0 denote an outside option available with certainty, with a discount rate r.

Dynamic Programming Formulation: Let the collection of boxes be denoted by

$$\mathcal{I} = \{1, 2, \ldots, n\},$$

and partition $\mathcal{I}$ into:

- S: the set of unopened boxes,
- $\bar{S}$: the set of opened boxes.

Define

$$y = \max_{i \in \bar{S}} X_i,$$

which is the best reward observed so far. The state of the system is $(\bar{S}, y)$, and let $\Psi(\bar{S}, y)$ denote the expected discounted value of being in this state and following the optimal policy. Then, the Bellman equation is

$$\Psi(\bar{S}, y) = \max\left\{y, \max_{i \in S}\left\{-C_i + \beta_i \int \Psi(\bar{S} \cup \{i\}, \max\{y, X_i\})\, dF_i(X_i)\right\}\right\},$$

with the boundary condition $\Psi(\emptyset, x) = x$.

Optimal Strategy and Reservation Prices

Intuition: In an optimal strategy, the marginal benefit of opening a box equals its marginal cost. Suppose there are two boxes (one open and one closed). Opening a closed box i yields a net benefit:

$$-C_i + \beta_i \left[\int_{-\infty}^{z_1} z_1 \, dF_i(x) + \int_{z_1}^{\infty} x_i \, dF_i(x) \right],$$

where z_1 is the reward from the already opened box (not box i). Box i is indifferent between being opened or left closed if

$$z_i = -C_i + \beta_i \left[\int_{-\infty}^{z_i} z_i \, dF_i(x) + \int_{z_i}^{\infty} x_i \, dF_i(x) \right], \tag{103}$$

$$C_i = \beta_i \left[\int_{z_i}^{\infty} (x_i - z_i) \, dF_i(x_i) \right] - (1 - \beta_i) z_i. \tag{104}$$

The value z_i that solves this equation is the reservation price for box i. A box is considered worth opening if its sample reward exceeds this reservation price.

Pandora's Rule:

- **Selection:** If a box is to be opened, choose the one with the highest reservation value.
- **Stop:** Stop searching when the best sample reward exceeds the reservation value of every unopened box.

Example 1: Suppose box i yields success with probability p_i and a reward R_i, and failure (reward 0) with probability $1 - p_i$, with $\beta = 1$ and $p_i R_i - C_i > 0$. Then,

$$C_i = (R_i - z_i) p_i \quad \Longrightarrow \quad z_i = \frac{p_i R_i - C_i}{p_i}.$$

For two projects with the same expected value, if one has a lower probability of success, it is optimal to research the one with the lower probability.

Example 2: Consider a mine i that contains G_i units of gold. A digging machine has a probability q_i of breaking after t units of time. Let $X_0 = 0$, $C_i = \alpha$, and assume the benefit is $-G_i$ (with $X_i > 0$). Then, with $\beta_i = 1 - q_i$, if the reservation price $z_i > 0$,

$$-G_i = (-q_i) z_i \quad \Longrightarrow \quad z_i = \frac{G_i}{q_i}.$$

Thus, the optimal strategy is to start with the mine that gives the most gold per unit probability of breakdown.

Proofs

(1) Existence and Uniqueness of Reservation Prices: Define

$$H_i(z) = \beta_i \int_{z}^{\infty} (x_i - z) \, dF_i(x_i) - (1 - \beta_i) z.$$

Since $H_i(z)$ is continuous and monotonic, with $H_i(\infty) = \infty$ and $H_i(-\infty) = -\infty$ (when $\beta_i < 1$), it follows that for any $C_i > 0$, there exists a unique z_i satisfying

$$H_i(z_i) = C_i.$$

(2) Induction on the Number of Closed Boxes: Assume Pandora's rule is optimal when there are m closed boxes, given the current best reward y. For $m + 1$ closed boxes, let box

j be the one with the highest reservation price among $\bar{S}$ (i.e., $z_j = \max_{i \in \bar{S}} z_i$). If $y \geq z_j$, it is optimal to stop; if $y < z_j$, then one must show that opening box j is optimal rather than any other box k with $z_k < z_j$. Through an inductive argument—comparing the expected values of the alternative strategies and employing auxiliary quantities such as

$$\pi_i = \Pr(X_i \geq z_i), \quad w_i = E[X_i \mid X_i \geq z_i], \tag{105}$$

$$\lambda_i = \Pr(z_h \leq X_i < z_j), \quad v_i = E[X_i \mid z_h \leq X_i < z_j], \tag{106}$$

$$\tilde{v}_i = E[\max(X_i, y) \mid z_h \leq X_i < z_j], \tag{107}$$

$$\mu_i = \Pr(z_k \leq X_i < z_h), \quad u_i = E[X_i \mid z_k \leq X_i < z_h], \tag{108}$$

one shows that the strategy of opening the box with the highest reservation price is optimal. By proving that the marginal benefit of opening any box equals its cost at the reservation price, the induction is complete.

Takeaway: Despite the complexity of general search models, Pandora's rule provides a precise prescription for the optimal search strategy regardless of specific details. The proofs illustrate how rules governing search can be rigorously converted into mathematical form, enabling clear interpretation and understanding.

6.6 Banerjee (1992)

Model Setup. Consider a population of n identical and risk–neutral agents, where assets are distributed over the interval $[0, 1]$. For the i^{th} asset, denote the asset by $a(i)$ with an associated return $z(i)$. There exists a unique index i^* such that

$$z(i^*) = Z > 0, \qquad \text{and} \qquad z(i \neq i^*) = 0.$$

All agents hold uniform priors regarding the location of i^*; that is, initially no agent possesses any additional information. Agents receive signals about the value of i^*: an agent receives a signal denoted by i' with probability α. Conditional on receiving a signal, the signal is correct (i.e., equals i^*) with probability β, and incorrect with probability $1 - \beta$. When the signal is false, i' is drawn from the uniform distribution on $[0, 1]$ (uninformative). Investment occurs sequentially: the first agent is chosen at random to invest; the second agent is chosen at random, observes the first agent's choice (but not the underlying signal), and then invests; and so on. In general, the n^{th} agent observes the choices of the previous $n - 1$ agents before making their investment decision.

Tie–Breaking Rules. The following rules are imposed to break indifferences:

(A) When an agent receives no signal and observes that everybody chose $i = 0$, the agent chooses $i = 0$.

(A) When an agent is indifferent between following their own signal and the observed choices of others, the agent follows their own signal.

(A) When indifferent between two or more previous decisions, the agent selects the higher one.

Equilibrium Behavior

First Agent.

- If the first agent receives a signal, then $a_1 = i'_1$ (i.e., the action equals the signal).
- If no signal is received, then $a_1 = 0$.

Second Agent.

- If the second agent receives a signal:
  - When $a_1 = 0$, then $a_2 = i'_2$.
  - When $a_1 \neq 0$, then $a_2 = i'_2$.
- If no signal is received:
  - If $a_1 = 0$, then $a_2 = 0$, so that $a_2 = a_1$.
  - If $a_1 \neq 0$, then $a_2 = i'_2$.

Third Agent. The decision of the third agent depends on the observed actions of the first two agents. Four cases arise:

(1) When $a_1 = a_2 = 0$:

- With no signal, set $a_3 = 0$.
- With a signal, set $a_3 = i'_3$.

(2) When $a_1 = 0$ but $a_2 \neq 0$:

- With no signal, choose $a_3 = a_2$.
- With a signal, choose $a_3 = i'_3$.

(3) When $a_1 \neq 0$ and $a_2 = 0$:

- With no signal, set $a_3 = a_1$ (using tie–breaking rule C).
- With a signal, set $a_3 = i'_3$.

(4) When $a_1 \neq 0$ and $a_2 \neq 0$:

- With no signal, set $a_3 = \max\{a_1, a_2\}$.
- With a signal, if $a_1 \neq a_2$, then $a_3 = i'_3$; if $a_1 = a_2$, then $a_3 = a_1 = a_2$.

Copying Behavior. A calculation is provided for the conditional probabilities:

$$\text{(I)} \quad P(i^* = a_1 = a_2 \mid a_1 = a_2,\ i'_3) = \frac{\alpha^3\beta^3(1-\beta) + \alpha^2\beta^2(1-\beta)(1-\alpha)}{P(H)},$$

$$\text{(II)} \quad P(i^* = i'_3 \mid a_1 = a_2,\ i'_3) = \frac{\alpha\beta(1-\beta)(1-\alpha)}{P(H)}.$$

Since (I) is greater than (II), agents have an incentive to "copy" the previous actions. In equilibrium, whenever two agents choose the same action, subsequent agents will follow that action; thus, private information is not conveyed further through actions. In other words, while actions converge via communication, private beliefs do not.

Implications.

This model helps explain phenomena such as fads and political herding, where there is a huge value attached to being the first person in line. It also illustrates how communication can break the herd, thereby altering the flow of social learning. Even if waiting is costly or agents can change positions, the general intuition of the model remains robust.

6.7 Normal Learning Model: Observational Learning

Overview. Learning is here conceptualized as the process of encountering information and applying it. Two approaches to statistical learning are distinguished:

Classical/Frequentist Approach:
- The parameter θ is considered an unknown but fixed quantity.
- One observes a random sample (i.e., iid observations $X_1, X_2, \ldots, X_n$ drawn from a population characterized by a probability density function indexed by θ).
- Inferences about θ are then drawn using methods such as maximum likelihood estimation (MLE) or the method of moments (MoM).

Bayesian Approach:
- The parameter θ is treated as a random variable with a distribution representing our prior beliefs.
- The prior distribution, denoted by $\pi(\theta)$, is subjective and specified before data are observed.
- After observing a random sample, the prior is updated using Bayes' law, and the process can be repeated as new data become available.

Denote:

$$\pi(\theta) : \text{prior distribution}, \quad f(X|\theta) : \text{sampling distribution},$$
$$m(X) = \int f(X|\theta)\pi(\theta)\, d\theta : \text{marginal distribution}.$$

Bayes' rule then yields the posterior:

$$\pi(\theta|X) = \frac{f(X|\theta)\pi(\theta)}{m(X)}.$$

6.8 Example: Normal Learning Model

Assume that

$$X_i \overset{\text{iid}}{\sim} N(\theta, \sigma^2),$$

and that the prior for θ is

$$\theta \sim N(\mu, \tau^2),$$

with σ^2, μ, and τ^2 known. The joint probability density function of X and θ is given by

$$f(X|\theta) \cdot \pi(\theta) = \frac{1}{\sqrt{2\pi}\sigma/\sqrt{n}} \exp\left[-\frac{(X-\theta)^2}{2\sigma^2/n}\right] \cdot \frac{1}{\sqrt{2\pi}\tau} \exp\left[-\frac{(\theta-\mu)^2}{2\tau^2}\right].$$

After combining exponents, one obtains an expression of the form

$$-\frac{1}{2}\left[\frac{(\theta - K(X))^2}{V^2} + \frac{(X-\mu)^2}{\tau^2 + \sigma^2/n}\right],$$

which implies that the joint density can be viewed as the product of a normal density in θ with mean $K(X)$ and variance V^2, and a marginal normal density in X with mean μ and variance $\tau^2 + \sigma^2/n$. Consequently, the posterior expectation is

$$E[\theta|X] = K(X) = \frac{\tau^2}{\tau^2 + \sigma^2/n}\overline{X} + \frac{\sigma^2/n}{\tau^2 + \sigma^2/n}\mu,$$

where the weight on the sample mean $\overline{X}$ increases as the variance of the prior (i.e., τ^2) increases, or as the sample variance σ^2 decreases. The posterior variance is

$$V(\theta|X) = V^2 = \frac{(\sigma^2/n)\,\tau^2}{\sigma^2/n + \tau^2}.$$

Thus, if the new information is imprecise, the posterior will rely more heavily on the prior.

6.9 Jovanovic (1979) & Kane & Staiger (2005)

Application to Tenure Decisions. In these works the parameter μ represents the true effectiveness (or productivity) of a teacher. The model assumes:

- $\mu \sim N(0, \sigma_\mu^2)$,
- For each period t, the observed performance is

$$Y_t = \mu + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2),$$

  so that

$$Y_t \mid \mu \sim N(\mu, \sigma_\varepsilon^2).$$

- Observations over t periods are used to decide on granting tenure.

The posterior expectation after observing $Y_1, \ldots, Y_t$ is given by

$$E[\mu \mid Y_1, \ldots, Y_t] = \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\varepsilon^2/t}\overline{Y} = \frac{t}{t + \sigma_\varepsilon^2/\sigma_\mu^2}\overline{Y},$$

where the term $\frac{t}{t+\sigma_\varepsilon^2/\sigma_\mu^2}$ serves as a shrinkage factor, pulling the estimate toward the prior mean of zero. In practice, a cutoff C is set: if $\overline{Y} > C$, tenure is granted; if $\overline{Y} \leq C$, tenure is denied. In a real decision, tenure is granted if

$$E(\mu \mid \overline{Y}) > C, \quad \text{or equivalently} \quad \overline{Y} > \frac{t + \sigma_\mu^2/\sigma_\varepsilon^2}{t} \cdot C.$$

The probability of tenure given a cutoff C is then

$$\Pr(\text{Tenure} \mid C) = \Pr(\mu > C \mid \overline{Y}) = 1 - \Phi\left(\frac{C' - \frac{t}{t+\sigma_\varepsilon^2/\sigma_\mu^2}\overline{Y}}{\frac{\sigma_\varepsilon^2}{t}\sqrt{\sigma_\mu^2}}\right).$$

If the cutoff is set too high, experienced individuals may be erroneously dismissed. The model can be extended by incorporating an experience trend,

$$Y_t = \mu + \varepsilon_t + \beta t, \quad \text{with } Y_t \mid \mu \sim N(\mu + \beta t, \sigma_\varepsilon^2),$$

and by defining

$$Z = Y_t - \beta t,$$

so that the average $\overline{Y}$ is replaced by $\overline{Z}$. This modification allows for determining an optimal cutoff time that balances early learning (which may lock in mediocre individuals) against the risk of losing top talent.

6.10 Crawford & Sobel (1982)

Communication and Strategic Information Transmission

Crawford and Sobel (1982) develop a model of *cheap talk* where communication (via advertisement or signaling) is costless. In the model, an informed sender (S) cares about the action taken by an uninformed receiver (R). The sender observes the true state of the world, $s \in S$, and sends a message $m \in M$ to the receiver, who then chooses an action $a \in A$. The model becomes particularly interesting when the sender and receiver have misaligned objectives.

The utility functions are specified as

$$U^S(a(m), s) = -(a - (s + b))^2,$$

where b represents the distortion in the sender's objectives, and

$$U^R(a, s) = -(a - s)^2.$$

Thus, while the sender knows the state, the receiver does not. For example, consider a simple case with binary state space $S = \{0, 1\}$, message space $M = \{0, 1\}$, and action space $A = \{0, 1\}$, with each state occurring with probability $\frac{1}{2}$.

The receiver's decision problem is to choose

$$a(m) \in \arg\max_a \Big\{ U^R(a, 0)\, p(S = 0 \mid m) + U^R(a, 1)\, p(S = 1 \mid m) \Big\}.$$

Simultaneously, the sender chooses a message according to

$$q(m \mid s) \in \arg\max_m \, U^S(a(m), s) \quad \forall s,$$

where $q(m \mid s)$ is the probability that the sender sends message m when the state is s. Bayes' rule then provides the posterior

$$p(s \mid m) = \frac{q(m \mid s)\, p(s)}{\frac{1}{2}\, q(m \mid s) + \frac{1}{2}\, q(m \mid 1 - s)}.$$

Types of Perfect Bayesian Equilibrium (PBE)

Three types of PBE are considered:

(a) **Fully Separating Equilibrium:** The message fully reveals the sender's private information. A candidate strategy is

$$q(m \mid s) = \begin{cases} 1 & \text{if } m = s, \\ 0 & \text{otherwise,} \end{cases} \quad \text{with } a(m) = m \quad \forall m.$$

This equilibrium exists when b is sufficiently small.

(b) **Partially Separating Equilibrium:** Some messages perfectly reveal the state while others do not. For instance, one might have

$$q(0 \mid 0) = 1, \quad q(1 \mid 0) = 0, \quad a(0) = 0, \quad a(1) = 2b,$$

and

$$q(0 \mid 1) = 0, \quad q(1 \mid 1) = 1.$$

In this equilibrium, even when the state is 0, the sender might have an incentive to induce the receiver to choose 1 in order to obtain a higher payoff.

(c) **Pooling (Babbling) Equilibrium:** No information is conveyed by the message. A typical example is

$$q(m \mid s) = 1 \quad \forall s, \quad \text{with } a(1) = \frac{1}{2}, \quad a(0) = 0,$$

or even a randomized strategy such as

$$q(m \mid s) = \frac{1}{2} \quad \forall m, s, \quad \text{yielding } a(m) = \frac{1}{2} \quad \forall m.$$

When b is very small, the babbling equilibrium always exists. Fully separating equilibria exist only if b is small enough (specifically, when $b < \frac{1}{2}$), whereas partially separating equilibria require b to be not too small (e.g., $b > \frac{1}{4}$ may be needed to sustain non–trivial messages). In cases where b is too large, no equilibrium in which the sender truthfully reveals any information exists because the sender always prefers to push for the receiver to choose action 1.

Extension to a Continuous State Space

Now consider a continuous state space $S = [0, 1]$, with action space $A = [0, 1]$ and a finite message space M. Assume the prior distribution over states is uniform, i.e., $S \sim U[0, 1]$. Bayes' rule then gives

$$p(s \mid m) = \frac{q(m \mid s)p(s)}{\int_0^1 q(m \mid t)p(t)\,dt} = \frac{q(m \mid s)}{\int_0^1 q(m \mid t)\,dt}.$$

Several observations can be made:

1. An always babbling equilibrium exists: for example, if $q(m \mid s) = \frac{1}{|M|}$ for all s, then the receiver's optimal action is $a(m) = \frac{1}{2}$ for every message.
2. A fully separating equilibrium does not exist in this setting, since the receiver cannot invert the sender's message function.
3. There exists a state $t \in S$ such that $U^S(a, t)$ is the same for two different actions a and a', i.e., t is the sender's bliss point.
4. One can define a function $a^S(t)$ taking values in (a, a') such that the difference $a'(t) - a$ equals $a' - a^S(t)$, reflecting a symmetry condition.
5. No state t' greater than t induces the lower action a, and no state t' less than t induces the higher action a'.

These observations suggest the existence of a cutoff in the state space.

Candidate Equilibrium with Two Messages. Suppose the sender's message space is $\{0, 1\}$, and these induce two actions $a_1 < a_2$ such that $a(0) = a_1$ and $a(1) = a_2$. Then, under equilibrium, there exists a cutoff point t such that:

$$\text{If } s \leq t, \text{ send } m = 0; \qquad \text{if } s > t, \text{ send } m = 1.$$

Bayes' rule implies that

$$E(S \mid m = 0) = \frac{1}{2}t, \qquad E(S \mid m = 1) = t + \frac{1-t}{2} = \frac{1+t}{2}.$$

The receiver chooses actions to maximize expected utility:

$$a(m) \in \arg\max_a \int_0^1 U^R(a, S)p(S \mid m)\,dS,$$

leading to

$$a(0) = \frac{t}{2} \quad \text{and} \quad a(1) = \frac{1+t}{2}.$$

Equilibrium requires indifference for the sender between the two messages, which imposes

$$\frac{t}{2} + b = \frac{1-t}{2} - b \quad \Longrightarrow \quad t = \frac{1}{2} - 2b.$$

Thus, the equilibrium exists provided that $t > 0$, i.e., if $b < \frac{1}{4}$.

Generalization. In a candidate equilibrium with three cutoffs $(0, t_1, t_2, 1)$, the following condition must hold for all interior cutoffs:

$$\begin{aligned} t_i + b - \frac{1}{2}(t_{i-1} + t_i) &= \frac{1}{2}(t_i + t_{i+1}) - (t_i + b), \\ 4b &= t_{i+1} - 2t_i + t_{i-1}, \\ t_i &= \frac{t_{i-1} + t_{i+1}}{2} - 2b. \end{aligned}$$

By setting $t_N = 1$, one obtains a bound on the maximum number of partitions, and hence on the granularity of information transmission. The key insight is that while full separation is impossible when the sender has a strong incentive to misreport, partial pooling equilibria can arise in which some information is credibly transmitted. In such equilibria, extreme messages may be sent to reveal the sender's type to some extent, even though complete revelation is precluded by incentive constraints.

Takeaway. Crawford and Sobel demonstrate that even when communication is costless (cheap talk), the possibility of partially revealing information exists. Although pooling equilibria may lead to no information transmission, under certain conditions the sender can credibly convey some information. This result has broad implications for understanding strategic communication in environments where the sender and receiver have misaligned objectives.