# 7. ASYMMETRIC INFORMATION

**OUTLINE**



## 7.1 INTRODUCTION

*Asymmetric information* describes an economic environment in which one agent in a transaction has different information from the other agent in that transaction. There are two classes of asymmetric information problems: *adverse selection* and *moral hazard.*

### Adverse Selection

Adverse selection relates to asymmetric information about **characteristics**.

As an example, consider a market for goods of variable quality where the seller of a particular good knows its quality but the buyer does not. (The characteristic here is "quality"). The buyer will base her initial valuation of the good on the market-wide expected quality.

The seller of a high quality good, not being able to credibly convince the buyer that it is high quality and thereby charge a high price, may decide to retain the good rather than sell it at an average-quality price. Conversely, the seller of a low quality good will be happy to sell it at an average-quality price.

Thus, the market *adversely selects* the lowest quality goods for sale even though there may be buyers and sellers who would mutually benefit from the sale of the high quality

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright.

goods. Adverse selection therefore leads to a loss of social surplus relative to a setting with symmetric information.

Adverse selection can potentially lead to the collapse of the market: buyers know that only low quality sellers will be willing to sell, so when they see a good for sale they revise downward their beliefs about the quality; this drives out still more sellers whose quality is above the "revised" average, and the downward spiral potentially continues until no goods are traded.

## Moral Hazard

Moral hazard relates to asymmetric information about **actions**.

As an example, consider an insurance market where a risk averse agent, faced with some uncertainty (such as the possibility of a house fire), buys insurance from a firm.

If the agent buys full insurance (to completely cover all loss), and her actions are unobservable to the insurer, then she has no incentive to take precautionary action to prevent the loss, even if such action is not very costly. Full insurance can therefore lead to an inefficiently low level of precautionary action.

In response to this problem the insurer offers only partial insurance and so the agent is exposed to some risk; she must therefore take precautionary action anyway, and so she incurs both the cost of the action and the cost of the remaining uninsured risk.

The agent would be better-off by taking more precautionary action and getting full insurance, but the moral hazard makes this impossible. Consequently, there is foregone social surplus relative to a setting with symmetric information.

The same problem arises more generally in any **principal-agent problem**, where the payoff to the principal depends in an uncertain way on the action of the agent contracted

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright. 

to perform that action, and the principal can only base payment for the agent's services on the observed outcome (because the action itself is not observable).

## 7.2 ADVERSE SELECTION: AN EXAMPLE (THE MARKET FOR LEMONS)

Consider a product of quality *q*. Suppose the seller values the product at $\theta_S q$ and the potential buyer values it at $\theta_B q$. Assume $\theta_B > \theta_S$. Thus, Pareto efficiency requires trade because the potential buyer values the product more highly than does the seller (regardless of quality).

The seller knows *q*, but the buyer does not. Thus, there is asymmetric information. To simplify the analysis, we will assume that the buyer knows $\theta_S$.

Suppose the buyer has prior beliefs about *q* represented by a uniform distribution over the interval $[q_L, q_H]$, where $q_L \geq 0$. Thus, the **prior expectation** on quality is

$$\mu = \frac{q_L + q_H}{2}$$

Now suppose the product is offered for sale at price *p*. If $p > \theta_S q_H$ then the buyer knows that the seller would definitely accept a lower price (because $\theta_S q_H$ is the highest possible seller valuation), so we can rule out that case and focus on the case where $p \leq \theta_S q_H$.

What should the buyer infer about *q* when the product is offered for sale at $p \leq \theta_S q_H$?

If the seller is willing to sell at price *p*, the buyer knows that $\theta_S q \leq p$. Thus, the buyer can infer that the quality of the product is $q \in [q_L, p/\theta_S]$. That is, the buyer revises her beliefs about quality in response to the observation that the product is being offered for sale at price *p*.

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright.

The **conditional expected quality** (that is, expected quality conditional on the product being offered for sale at price *p*) is

$$\hat{q}(p) = \frac{q_L + p/\theta_S}{2}$$

The expected surplus for the buyer is

$$\begin{aligned} \hat{S}_B &= \theta_B \hat{q}(p) - p \\ &= \theta_B \left( \frac{q_L + p/\theta_S}{2} \right) - p \\ &= \frac{\theta_B q_L}{2} + p\left( \frac{\theta_B - 2\theta_S}{2\theta_S} \right) \end{aligned}$$

If $\theta_B > 2\theta_S$ then this is always positive. That is, there exists a price at which trade can occur, and so the information asymmetry is not an obstacle to trade.

The more interesting case is where $\theta_B < 2\theta_S$. In this case, the maximum price at which trade can occur (that is, at which $\hat{S}_B \geq 0$) is

$$p_{\max} = \left( \frac{\theta_B \theta_S}{2\theta_S - \theta_B} \right) q_L$$

This in turn means that the highest quality product that would ever be offered for sale is

$$\begin{aligned} q_{\max} &= \frac{p_{\max}}{\theta_S} \\ &= \left( \frac{\theta_B}{2\theta_S - \theta_B} \right) q_L \end{aligned}$$

If $q_{\max} < q_H$ then there will be some products of quality higher than $q_{\max}$ that will not be offered for sale despite the fact that are there potential buyers and sellers who could

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright.

mutually benefit from that trade (since $\theta_B > \theta_S$). In that case, the asymmetry of information causes Pareto inefficiency.

Numerical Example

Suppose $\theta_B = \frac{3}{2}$, $\theta_S = \frac{4}{3}$, $q_L = 7$ and $q_H = 12$.

The average quality of products is

$$\mu = \frac{7+12}{2} = 9.5$$

The maximum price that a buyer will pay is

$$p_{\max} = 12$$

and the maximum quality of a product offered for sale is

$$q_{\max} = 9$$

That is, only those products with quality less than or equal to $q = 9$ are offered for sale. Products with quality $q \in [9, 12]$ are not offered for sale. The market has adversely selected the lower quality products; the highest quality products are not offered for sale.

## 7.3 THE SPENCE SIGNALING MODEL: A LABOUR MARKET EXAMPLE

There are three main mechanisms through which the market can potentially deal with problems of adverse selection: warranties, reputation effects (in a repeated interaction context such as repeat sales or word-of-mouth communication), and **signaling**. In this section we focus on signaling.

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright.

**The Basic Model**

Consider a situation where a worker obtains education level *e* and demands wage *w* from the employer. The firm accepts or refuses the demanded wage. Assume that education has no productivity effect (unlike a degree in economics).

The worker is one of two types: high productivity (*H*) or low productivity (*L*). The firm cannot observe productivity prior to hiring. The firm knows the true population distribution of workers. In particular, a fraction $\alpha$ of workers are of type *L*, and a fraction $1-\alpha$ are of type *H*.

For the worker, the cost of obtaining education level *e* is correlated with her productivity. In particular, the effort-cost of education level *e* is

$$c = \frac{e}{\lambda + t}$$

where $t = L$ or $t = H$, and $\lambda \geq 0$ is a parameter common to both types.

The net payoff to a worker of type *t* who obtains education level *e* and receives wage *w* is

$$u = w - \frac{e}{\lambda + t}$$

We will see that the asymmetry in effort-cost between the *H* type and the *L* type creates the potential for the *H* type worker to signal her type via education level.

If the firm cannot distinguish between the two types, it accepts the wage demanded if and only if the wage does not exceed expected productivity; that is, if and only if $w \leq \mathbf{E}[t]$, where

$$\mathbf{E}[t] = \alpha L + (1-\alpha)H$$

This means that any $w \leq L$ is always accepted, and any $w > H$ is always refused.

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright. 

**Pooling Equilibrium**

In a pooling equilibrium (PE), both types choose $e = 0$ and so the employer cannot distinguish between them. Thus, in the PE both types are paid a wage equal to expected productivity, and the payoff to both types is

$$u^P = \alpha L + (1-\alpha)H$$

Can the *H* type do better than this by somehow convincing the employer that she is *H* type and so obtain $w = H$ ?

She may be able to do so through her choice of *e*. That is, there may be an education level $\hat{e}$ that only the *H* type would be willing to undertake, which thereby signals that the worker must be of type *H*. To put this differently, there may be an education level $\hat{e}$ that allows the *H* type to **separate** herself from *L* types.

**Separating Equilibria**

In a separating equilibrium (SE), the *H* type chooses an education level $\hat{e} > 0$ that convinces the employer that she is the *H* type because the employer knows that only the *H* type would choose this education level. If there exists such a signal, then any worker who does not choose $\hat{e}$ will be viewed by the employer as the *L* type.

In a separating equilibrium, the *H* type will choose $\hat{e}$ and receive wage $w = H$ , and obtain net payoff

$$u^H = H - \frac{\hat{e}}{\lambda + H}$$

and the L type will choose $e = 0$ and receive wage $w = L$ , and obtain net payoff

$$u^L = L$$

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright. 

If $\hat{e}$ is a separating-equilibrium level of $e$ then it must be **incentive compatible** for both types:

- the *L* type must prefer her equilibrium strategy to any alternative strategy, including one where she mimics the *H* type; and
- the *H* type must prefer her equilibrium strategy to any alternative strategy, including one where she mimics the *L* type.

These **incentive compatibility conditions** are

(7.1) $L \geq H - \dfrac{\hat{e}}{\lambda + L}$ for the L type

(7.2) $H - \dfrac{\hat{e}}{\lambda + H} \geq L$ for the H type

Equation (7.1) requires $\hat{e} \geq (\lambda + L)(H - L)$, and equation (7.2) requires $\hat{e} \leq (\lambda + H)(H - L)$. Since $H > L$, there does exist an $\hat{e}$ at which both of these conditions are satisfied. That is, there exists an $\hat{e}$ such that

$$(\lambda + L)(H - L) \leq \hat{e} \leq (\lambda + H)(H - L)$$

Thus, there exists an education level that can signal productivity. In particular, if the *H* type chooses an education level just slightly higher than

$$e^S = (\lambda + L)(H - L)$$

then the employer will know that she is indeed a *H* type. Note that this separating education level is increasing in $\lambda$. Why?

A higher value of $\lambda$ means that the productivity difference (*H* vs. *L*) is less important in the determination of effort-cost. This means that it is easier for the *L* type to mimic the education level of the *H* type, and so the *H* type must do a lot more to separate herself from the *L* type.

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright. 

The $H$ type will choose the separating education level only if it is **individually rational** for her to so. Her payoff to separating by choosing $e^S$ is

$$u^H = H - \frac{e^S}{1+\lambda H} = \frac{(H-L)(1+\lambda L)}{1+\lambda H}$$

If instead the $H$ type chooses $e = 0$ then she does not distinguish herself from the $L$ type and we remain in a pooling equilibrium, in which case the payoff to the $H$ type (and the $L$ type) is

$$u^P = \alpha L + (1-\alpha)H$$

The $H$ type will choose to signal if and only if $u^H > u^P$. This in turn requires

$$\alpha > \hat{\alpha} = \frac{\lambda + L}{\lambda + H}$$

This threshold is plotted against $\lambda$ in **Figure 7-1** for given values of $L$ and $H$.

The $H$ type signals her type only in the region above the threshold. In the region above the threshold, a large fraction of workers are $L$ type and so average expected productivity (on which the pooling wage is based) is very low. The $H$ type has a strong incentive to signal in that case. However, doing so becomes increasingly costly at higher values of $\lambda$ because the $H$ type must undertake a lot more education to separate herself from the L type. Thus, at any given value of $\alpha$, we transition from a SE to a PE as $\lambda$ grows.

Note from **Figure 7-1** that there exists a critical value of $\alpha$, denoted

$$\overline{\alpha} = \frac{L}{H}$$

below which the $H$ type would never signal, even when $\lambda = 0$. At values of $\alpha$ below $\overline{\alpha}$ the average-productivity wage is high enough to make signaling not worthwhile because a large fraction of workers are $H$ type. In the limit where $\alpha = 0$ all workers are $H$ type, so signaling via education serves no purpose.

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright. 

Note too that in the limit as $\lambda \to \infty$, education requires no effort for anyone, regardless of type, and so signaling via education is impossible. Thus, the region above the threshold in **Figure 7-1** vanishes as $\lambda \to \infty$.

## 7.4 MORAL HAZARD IN INSURANCE: AN EXAMPLE

Suppose an agent has wealth $m$ in the good state, and wealth $m - L$ if an accident occurs (the bad state). Let $\pi(e)$ denote the probability of an accident, as a function of precautionary effort *e*, where $\pi'(e) < 0$.

In particular, suppose

$$\pi(e) = \frac{1}{1+e}$$

Thus, if $e = 0$ then $\pi(e) = 1$; an accident is guaranteed. Conversely, if $e \to \infty$ then $\pi(e) \to 0$; an accident is impossible.

Her indirect utility function is

$$v(m,e) = m^{\frac{1}{2}} - \delta e$$

where $\delta e$ is the utility-cost of effort.

Expected utility without insurance is

$$\mathbf{E}[v^0] = \pi(e)v(m-L) + (1-\pi(e))v(m) - \delta e$$

$$= \left(\frac{1}{1+e}\right)(m-L)^{\frac{1}{2}} + \left(1 - \frac{1}{1+e}\right)m^{\frac{1}{2}} - \delta e$$

The agent chooses *e* to maximize this expected utility. The first-order condition is

$$\frac{-(m-L)^{\frac{1}{2}}}{(1+e)^2} + \frac{m^{\frac{1}{2}}}{(1+e)^2} - \delta = 0$$

and this solves for

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright. 

$$e^* = \left( \frac{m^{\frac{1}{2}} - (m-L)^{\frac{1}{2}}}{\delta} \right)^{\frac{1}{2}} - 1$$

At this level of effort, the probability of loss is

$$\pi^* = \frac{1}{1+e^*} = \left( \frac{\delta}{m^{\frac{1}{2}} - (m-L)^{\frac{1}{2}}} \right)^{\frac{1}{2}}$$

For example, suppose $m = 640000$, $L = 150000$ and $\delta = 4$. Then $e^* = 4$. At this level of precautionary action, $\pi^* = \frac{1}{5}$. Her expected loss is $\pi^* L = 30000$.

Now suppose this agent can purchase full insurance for a total premium *T*, where this premium cannot be based on *e* because *e* is not observable to the insurer. Then her expected utility is

$$\begin{aligned} \mathbf{E}[v^F] &= \pi(e)v(m-T) + (1-\pi(e))v(m-T) - \delta e \\ &= v(m-T) - \delta e \end{aligned}$$

Since *e* has a utility cost but now has no benefit, she chooses $e = 0$. That is, having obtained full insurance, thereby eliminating all risk associated with an accident, she has no incentive at all to reduce the likelihood of that accident. This is the essence of the **moral hazard problem**.

A partial solution to this problem is **co-insurance**: the insurer covers the loss in the event of an accident but the insured agent must pay an amount *d* as part of a claim. This payment is often called a **deductible**.

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright. 

In that setting, her expected utility with full coverage and deductible *d* is

$$\mathbf{E}[v^C] = \pi(e)v(m-T-d) + (1-\pi(e))v(m-T) - \delta e$$
$$= \left(\frac{1}{1+e}\right)(m-T-d)^{\frac{1}{2}} + \left(1-\frac{1}{1+e}\right)(m-T)^{\frac{1}{2}} - \delta e$$

The agent chooses *e* to maximize this expected utility. The first-order condition is

$$\frac{-(m-T-d)^{\frac{1}{2}}}{(1+e)^2} + \frac{(m-T)^{\frac{1}{2}}}{(1+e)^2} - \delta = 0$$

and this solves for

(7.3) $$\hat{e} = \left(\frac{(m-T)^{\frac{1}{2}} - (m-T-d)^{\frac{1}{2}}}{\delta}\right)^{\frac{1}{2}} - 1$$

At this level of effort, the probability of loss is

(7.4) $$\hat{\pi} = \frac{1}{1+\hat{e}} = \left(\frac{\delta}{(m-T)^{\frac{1}{2}} - (m-T-d)^{\frac{1}{2}}}\right)^{\frac{1}{2}}$$

This probability is decreasing in $d$; that is, a higher deductible reduces the probability of an accident because it creates a stronger incentive for the insured agent to undertake precautionary action.

**First-Best is Unattainable**

Now suppose the insurer attempts to set the deductible to ensure that $\hat{\pi} = \pi^*$. The actuarially fair price for insurance is then $r = \pi^*$, and the associated total premium for effective coverage $(L-d)$ is $T = \pi^*(L-d)$. If we make this substitution for *T* in $\hat{\pi}$, and then set $\hat{\pi} = \pi^*$, we can solve for $d^*$. This deductible will induce $e^*$ and associated $\pi^*$.

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright.

What is this critical value $d^*$? The algebra gets complicated but it is straightforward to show that $d^* = L$. That is, the only deductible that can induce the optimal amount of precautionary effort is one that effectively means no insurance coverage at all.

The key message here is that co-insurance can restore some incentive to take precautionary effort but the first-best solution, $e^*$ with full insurance, cannot be achieved in the face of the moral hazard problem.

So what is the best outcome that can be achieved?

**The Second-Best Solution**

We are looking for a contract with effective coverage $(L-d)$ and total premium *T*, that maximizes the expected utility of the insured agent subject to the insurer making zero expected profit (that is, subject to the contract being actuarially-fair).

Our first step is to characterize the schedule of contracts, comprising $\{T,d\}$ pairs, that yield zero expected profit given that the probability of a claim is $\hat{\pi}$, as identified in (7.4) above. That is, the actuarially-fair contract must satisfy

$$T - \hat{\pi}(L-d) = 0$$

This equation can be solved for $d$ as a function *T*, denoted $d_0(T)$, but it is far too complicated to be reported usefully here. However, it can be plotted for the parameter values from our earlier example, where $m = 640000$, $L = 150000$ and $\delta = 4$. This plot is depicted in **Figure 7-2**.

Note that this zero-expected-profit schedule is negatively-sloped. A lower deductible means a weaker incentive for the insured agent to take precautionary action and that raises the probability of a claim. To compensate for this higher risk, the premium must be higher.

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright. 

Our next step is to choose *T* and *d* to maximize the expected utility of the insured agent subject to the contract lying on the zero-expected-profit schedule depicted in **Figure 7-2**. That is,

$$\max_{T,d} \quad \hat{\pi}v(m-T-d)+(1-\hat{\pi})v(m-T)-\delta\hat{e}$$

$$\text{subject to} \quad T-\hat{\pi}(L-d)=0$$

where $\hat{\pi}$ is given by (7.4) and $\hat{e}$ is given by (7.3).

The solution to this problem is a tangency between the zero-expected-profit schedule and an iso-expected-utility contour for the insured agent. That tangency is depicted in **Figure 7-3** for the example parameter values.

Finding this tangency involves more mathematics than we want to undertake here, but it can be shown that the solution for our example values is $\tilde{d}=126435$ and $\tilde{T}=5150$ (as depicted in **Figure 7-3**). That is, the agent has nominal coverage equal to $150,000 but must pay a deductible of $126,435 in the event of a claim. Her total premium is $5150.

Her precautionary action under this contract is 3.58, and the associated probability of a loss is 0.2186.

In comparison, if she could somehow commit to the first-best effort level ($e^*=4$ with $\pi^*=0.2$) and then buy actuarially-fair full insurance with no deductible, her total premium would be $30000.

Would she necessarily prefer this first-best contract? Yes. In particular, it can be shown that she would be willing to pay a total premium of $31,335 for this contract and still be just better off than under the second-best contract. Thus, the moral hazard problem causes a loss of surplus (a deadweight loss) of ($31,335 – $30,000) = $1335 in this example.

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright.

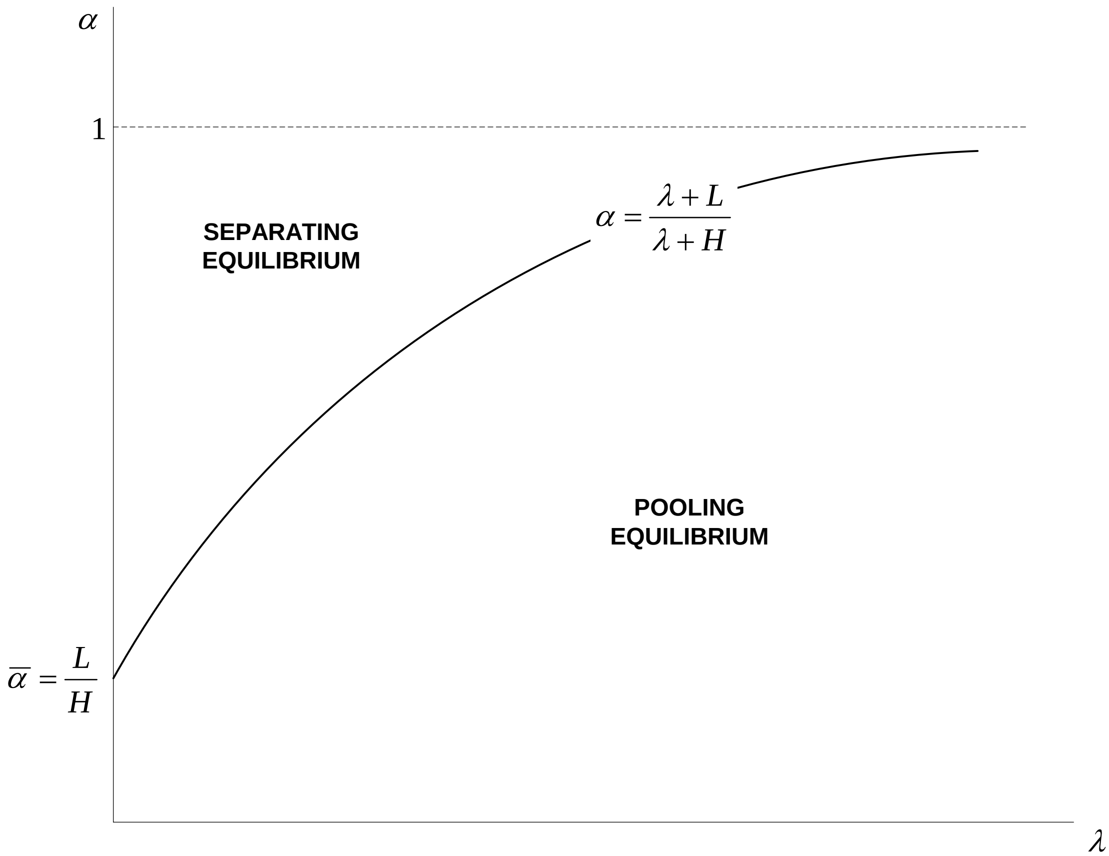


**Figure 7-1**

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright.

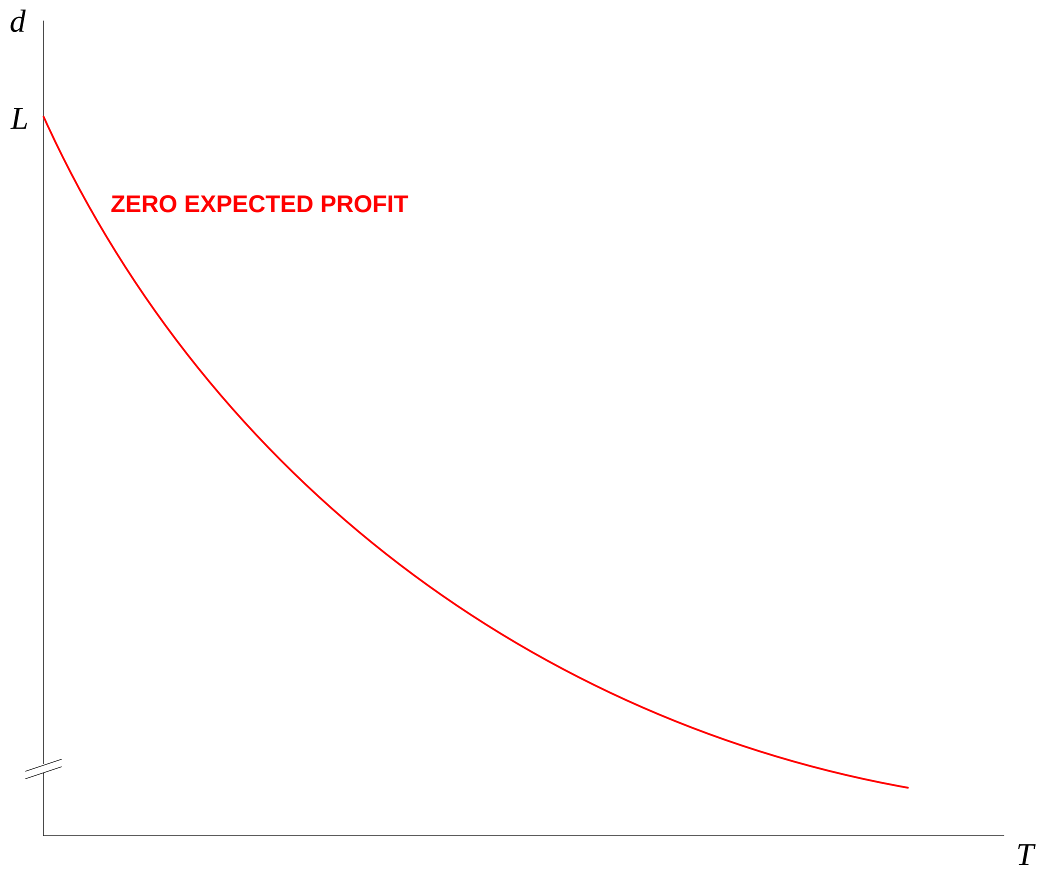


**Figure 7-2**

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright.

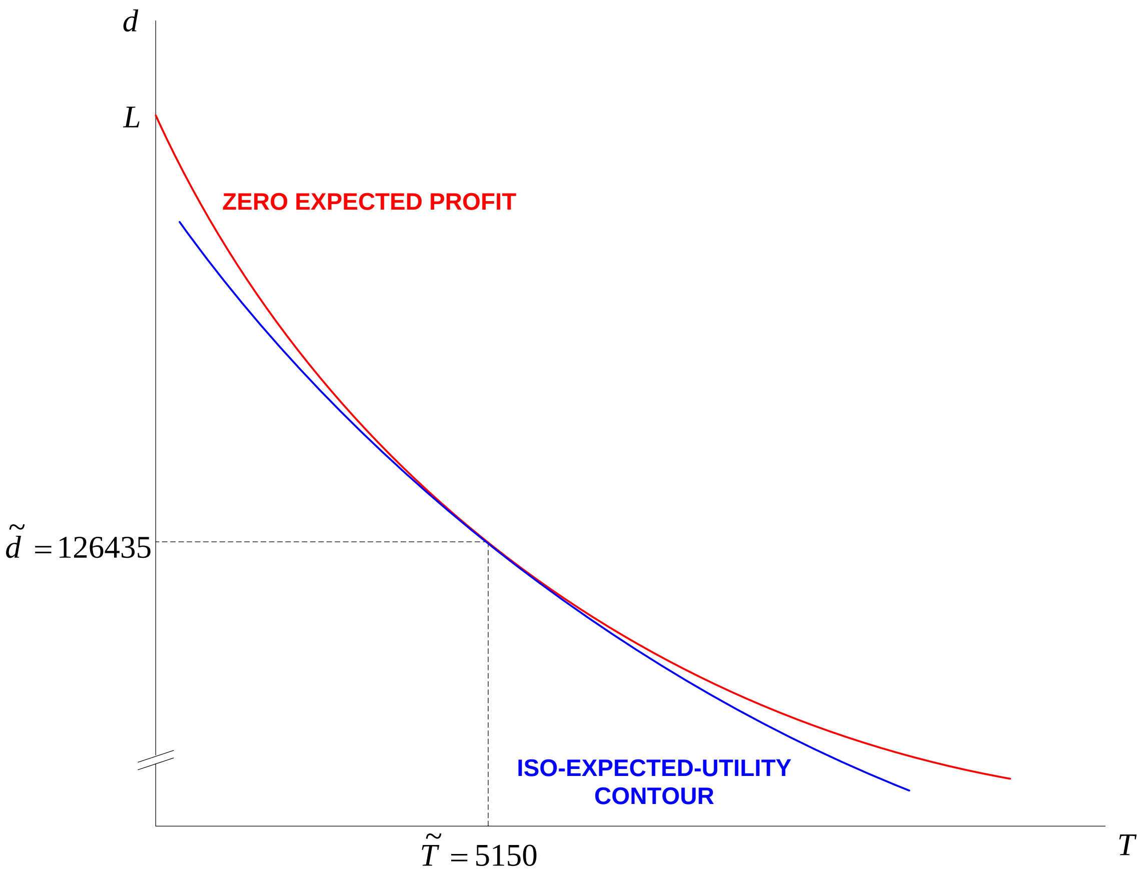


**Figure 7-3**

Posting this material to any site other than *web.uvic.ca/~pkennedy* is a violation of copyright.