DegreeLive

Bachelor of Commerce (B2001) | Level 2 • Part B Major | Monash University

Intermediate microeconomics

Course Code: ECC2000 • Module 4 Notes

Module 4: General Equilibrium, Welfare Economics & Asymmetric Information

Chapter 13: Pure Exchange General Equilibrium and the Edgeworth Box

13.1 Feasible Allocations, the Edgeworth Box, and Pareto Efficiency

Whereas partial equilibrium analyses a single market holding prices elsewhere fixed, **General Equilibrium (GE)** analyses the simultaneous determination of prices and allocations across all interconnected markets in an economy. Consider a two-consumer (A, B), two-good $(1, 2)$ pure exchange economy with initial endowments $\omega^A = (\omega_1^A, \omega_2^A)$ and $\omega^B = (\omega_1^B, \omega_2^B)$. Aggregate endowments are $\Omega_1 = \omega_1^A + \omega_1^B$ and $\Omega_2 = \omega_2^A + \omega_2^B$. An allocation $(x^A, x^B) = ((x_1^A, x_2^A), (x_1^B, x_2^B))$ is **feasible** if:

$$x_1^A + x_1^B = \Omega_1, \qquad x_2^A + x_2^B = \Omega_2, \qquad x_i^h \geq 0$$

Where:

- x_i^h denotes consumer $h \in \{A, B\}$'s consumption of good $i \in \{1, 2\}$.
- Ω_1 and Ω_2 denote total economy-wide endowments of good 1 and good 2.

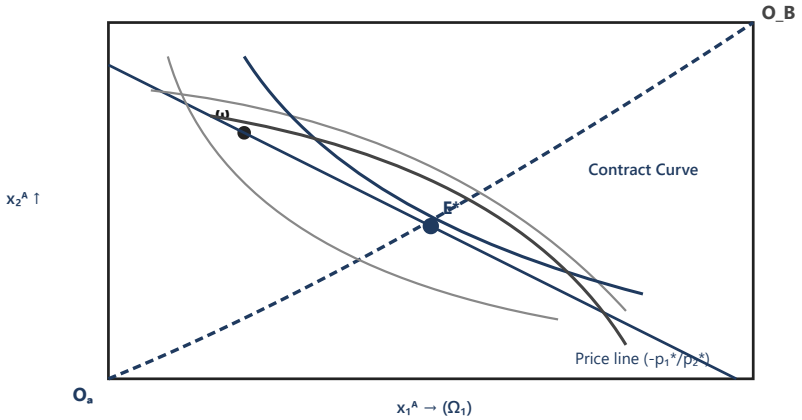
Every feasible allocation can be plotted inside an **Edgeworth Box** of width Ω_1 and height Ω_2, measuring A's bundle (x_1^A, x_2^A) from the bottom-left origin O_A and B's bundle $(x_1^B, x_2^B) = (\Omega_1 - x_1^A, \Omega_2 - x_2^A)$ from the inverted top-right origin O_B. Consumer A's indifference curves bow toward O_A and increase north-east; consumer B's indifference curves bow toward O_B and increase south-west.

A feasible allocation (x^A, x^B) is **Pareto efficient (Pareto optimal)** if there exists no other feasible allocation $(\tilde{x}^A, \tilde{x}^B)$ that makes at least one consumer strictly better off without making the other strictly worse off. From an initial endowment ω, the shaded lens formed by the two indifference curves passing through ω represents all mutually advantageous trades. At any interior Pareto-efficient allocation, the two consumers' indifference curves must be tangent—otherwise a mutual lens of gainful trade remains unexploited:

$$MRS_{1,2}^A(x_1^A, x_2^A) = MRS_{1,2}^B(x_1^B, x_2^B)$$

Where $MRS_{1,2}^h = MU_1^h / MU_2^h$ is consumer h's marginal rate of substitution. The locus of all Pareto-efficient allocations inside the Edgeworth Box is the **Pareto set**, and the subset that is weakly preferred by both agents to the initial endowment ω is the **contract curve (core)**.

DegreeLive



13.2 Walrasian Equilibrium and Walras' Law

In a decentralised market system with price vector $p = (p_1, p_2) \gg 0$, consumer h's nominal wealth is the market value of their endowment: $m^h(p) = p_1\omega_1^h + p_2\omega_2^h$. Each consumer solves their utility maximisation problem, generating Marshallian demands $x_i^h(p, p \cdot \omega^h)$ and **excess demand functions** for each good $i = 1, 2$:

$$z_i(p_1, p_2) = \sum_{h \in \{A,B\}} \left(x_i^h(p, p \cdot \omega^h) - \omega_i^h \right)$$

Where:

- $z_i(p)$ denotes aggregate excess demand for good i ($z_i > 0$ indicates shortage; $z_i < 0$ indicates surplus).
- Because Marshallian demands are homogeneous of degree zero in (p, m), excess demands depend only on the relative price ratio p_1/p_2, permitting one good to be normalised as the **numeraire** ($p_2 = 1$).

A **Walrasian (Competitive) Equilibrium** is a price vector $p^* = (p_1^*, p_2^*)$ and allocation (x^{A*}, x^{B*}) such that every agent maximises utility on their budget set and all markets clear ($z_1(p^*) = 0$ and $z_2(p^*) = 0$). Summing both consumers' binding budget constraints ($p_1 x_1^h + p_2 x_2^h = p_1\omega_1^h + p_2\omega_2^h$) establishes **Walras' Law**, an algebraic identity holding at *every* positive price vector p (not merely at equilibrium):

$$p_1 z_1(p_1, p_2) + p_2 z_2(p_1, p_2) \equiv 0$$

Where the aggregate value of excess demands across all L markets equals zero identically. An immediate corollary is that if $L - 1$ markets clear ($z_1(p^*) = 0$) and $p_2 > 0$, then the L-th market must clear automatically ($z_2(p^*) = 0$). Thus, finding a competitive equilibrium in an L-good economy reduces to solving $L - 1$ independent excess-demand equations for $L - 1$ relative prices.

Worked Example 13.1: Contract Curve and Walrasian Equilibrium in an Edgeworth Box

Consider a pure exchange economy with two traders, A and B, over two commodities: water rights (x_1, megalitres) and composite agricultural goods (x_2, dollars). Preferences are $u^A(x_1^A, x_2^A) = (x_1^A)^{1/3}(x_2^A)^{2/3}$ and $u^B(x_1^B, x_2^B) = (x_1^B)^{2/3}(x_2^B)^{1/3}$. Endowments are $\omega^A = (10, 80)$ and $\omega^B = (50, 40)$, so total endowments are $\Omega_1 = 60$ and $\Omega_2 = 120$.

DegreeLive

1. **Step 1 — Derive the Pareto contract curve:** Computing marginal rates of substitution:

$$MRS^A_{1,2} = \frac{\frac{1}{3}x^A_2}{\frac{2}{3}x^A_1} = \frac{x^A_2}{2x^A_1}, \qquad MRS^B_{1,2} = \frac{\frac{2}{3}x^B_2}{\frac{1}{3}x^B_1} = \frac{2x^B_2}{x^B_1}$$

Equating $MRS^A_{1,2} = MRS^B_{1,2}$ and substituting $x^B_1 = 60 - x^A_1$ and $x^B_2 = 120 - x^A_2$:

$$\frac{x^A_2}{2x^A_1} = \frac{2(120 - x^A_2)}{60 - x^A_1} \implies x^A_2(60 - x^A_1) = 4x^A_1(120 - x^A_2)$$

Rearranging gives the exact contract curve equation:

$$60x^A_2 + 3x^A_1x^A_2 = 480x^A_1 \implies x^A_2 = \frac{160x^A_1}{20 + x^A_1}$$

Notice at $x^A_1 = 0, x^A_2 = 0$ (O_A), and at $x^A_1 = 60, x^A_2 = 9600/80 = 120$ (O_B).

2. **Step 2 — Walrasian equilibrium relative price:** Normalise $p_2 = 1$ and let $p = p_1/p_2$. Endowment incomes are $m^A = 10p + 80$ and $m^B = 50p + 40$. Using Cobb-Douglas budget shares:

$$x^A_1(p) = \frac{1}{3}\frac{10p + 80}{p}, \qquad x^B_1(p) = \frac{2}{3}\frac{50p + 40}{p}$$

Setting market clearing in good 1 ($x^A_1 + x^B_1 = 60$):

$$\frac{10p + 80}{3p} + \frac{100p + 80}{3p} = 60 \implies 110p + 160 = 180p \implies 70p = 160 \implies p^* = \frac{16}{7} \approx 2.286$$

3. **Step 3 — Equilibrium allocation and verification:** At $p^* = 16/7$:

$$m^A = 10\left(\frac{16}{7}\right) + 80 = \frac{720}{7}, \qquad x^{A*}_1 = \frac{1}{3} \cdot \frac{720/7}{16/7} = 15, \qquad x^{A*}_2 = \frac{2}{3}\left(\frac{720}{7}\right) = \frac{480}{7} \approx 68.57$$

By feasibility, $x^{B*}_1 = 45$ and $x^{B*}_2 = 120 - 480/7 = 360/7 \approx 51.43$. Checking good 2 market clearing: $x^B_2 = \frac{1}{3}m^B = \frac{1}{3}[50(16/7) + 40] = \frac{360}{7}$, confirming Walras' Law. Checking the contract curve formula at $x^A_1 = 15$: $x^A_2 = \frac{160(15)}{20+15} = \frac{2400}{35} = \frac{480}{7}$, confirming the Walrasian equilibrium lies directly on the Pareto-efficient contract curve!

Degree Live

Chapter 14: Production General Equilibrium, Fundamental Theorems & Market Failures

14.1 General Equilibrium with Production and Overall Pareto Efficiency

Extending general equilibrium to an economy with production ($2 \; times 2 \; times 2$ model: two inputs K, L, two outputs X, Y, and two consumers A, B) introduces three simultaneous marginal conditions for overall Pareto efficiency:

1. **Exchange efficiency:** $MRS^{A}_{X,Y} = MRS^{B}_{X,Y}$ (goods are allocated efficiently across consumers).
2. **Production (technical) efficiency:** Inside the production Edgeworth Box of dimensions $(\bar{L}, \bar{K})$, inputs are allocated on the production contract curve where $MRTS^{X}_{L,K} = MRTS^{Y}_{L,K}$, placing aggregate output (X, Y) directly onto the concave **Production Possibility Frontier (PPF)**.
3. **Product-mix (top-level) efficiency:** The slope of the PPF—the Marginal Rate of Transformation $MRT_{X,Y} = MC_X / MC_Y$—must equal the common consumer marginal rate of substitution:

   $$MRS^{A}_{X,Y} = MRS^{B}_{X,Y} = MRT_{X,Y}$$

   Where $MRT_{X,Y}$ is the opportunity cost of producing one extra unit of good X in terms of forgone units of good Y.

14.2 The First and Second Fundamental Theorems of Welfare Economics

The two cornerstone results of modern welfare economics formalise Adam Smith's invisible hand and clarify the theoretical separation between efficiency and equity:

- **First Fundamental Theorem of Welfare Economics (FFTWE):** Suppose preferences satisfy local non-satiation, markets are complete and competitive (all agents act as price-takers), and there are no externalities or asymmetric information. Then *every Walrasian equilibrium allocation is Pareto efficient*. Proof intuition: if a feasible allocation $(\tilde{x}^{A}, \tilde{x}^{B})$ Pareto-dominated the equilibrium (x^{A*}, x^{B*}), then $\tilde{x}^{h}$ must cost strictly more than m^{h} for the strictly better-off agent (and at least m^{h} for the other by LNS); summing over agents yields $p^{*} \cdot (\tilde{x}^{A} + \tilde{x}^{B}) > p^{*} \cdot \Omega$, contradicting feasibility.
- **Second Fundamental Theorem of Welfare Economics (SFTWE):** Suppose all consumer preferences are convex and all firm production sets are convex (ruling out natural monopoly increasing returns). Then *any Pareto-efficient allocation can be decentralised as a Walrasian equilibrium* after a suitable redistribution of initial endowments via non-distortionary **lump-sum transfers**.

In practical Australian public policy, the SFTWE clarifies why governments do not abandon the price mechanism whenever market outcomes appear inequitable (for example, deregulating wholesale prices while compensating low-income households via direct income transfers). However, because true lump-sum taxes based on unobservable innate earning ability are infeasible, real-world progressive income taxation (such as Australian PAYG marginal brackets) introduces labour-supply distortions, creating an unavoidable **equity-efficiency trade-off**.

14.3 Externalities and Public Goods (Samuelson Condition)

When property rights are incomplete or consumption is non-rival, the FFTWE fails:

- **Externalities:** If production of Q inflicts an unpriced marginal external damage $MD(Q)$ (such as carbon emissions or Murray-Darling salinity), private equilibrium equates $P(Q) = PMC(Q)$, whereas social

DegreeLive

optimality requires $P(Q^*) = SMC(Q^*) = PMC(Q^*) + MD(Q^*)$. A **Pigouvian tax** set at $t^* = MD(Q^*)$ (or cap-and-trade emissions permits) restores efficiency. Alternatively, under the **Coase Theorem**, when property rights are clearly defined and transaction costs are zero, private bargaining achieves the efficient level Q^* regardless of initial liability assignment.

- **Pure Public Goods:** A pure public good G (e.g., national defence, bushfire early-warning radar, or urban flood levees) is non-rival and non-excludable. Because every citizen consumes the identical aggregate quantity G, market demand is summed *vertically* rather than horizontally, yielding the **Samuelson Condition** for optimal provision:

$$\sum_{i=1}^{n} MRS_{G,x}^{i} = MRT_{G,x}$$

Where $MRS_{G,x}^{i}$ is citizen i's marginal willingness to pay for the public good in units of private numeraire x, and $MRT_{G,x} = MC_G$ is its marginal resource cost. Under voluntary private contributions, individuals ignore the positive externality on others and contribute only where their private $MRS_{G,x}^{i} = MC_G$, causing severe **free-riding** and underprovision.

Worked Example 14.1: Samuelson Condition vs Voluntary Free-Riding Equilibrium

Two Victorian shires ($i = 1, 2$) benefit from regional flood-mitigation levee height G (in metres) and private spending x_i ($p_x = 1$). Preferences are quasi-linear: $u_1(G, x_1) = 120 \ln G + x_1$ and $u_2(G, x_2) = 80 \ln G + x_2$. The marginal cost of building the levee is $MC_G = 25$ dollars per unit.

1. **Step 1 — Individual marginal willingness to pay ($MRS_{G,x}^{i}$):**

$$MRS_{G,x}^{1} = \frac{120}{G}, \qquad MRS_{G,x}^{2} = \frac{80}{G}$$

2. **Step 2 — Samuelson socially optimal provision (G^*):** Summing marginal willingness to pay vertically and equating to $MC_G = 25$:

$$\frac{120}{G^*} + \frac{80}{G^*} = 25 \implies \frac{200}{G^*} = 25 \implies G^* = 8 \quad \text{textmetres}$$

Under personalised **Lindahl pricing**, Shire 1 pays $\tau_1 = 120/8 = 15$ dollars per metre and Shire 2 pays $\tau_2 = 80/8 = 10$ dollars per metre, unanimously agreeing on $G^* = 8$.

3. **Step 3 — Voluntary contribution Nash equilibrium (G^{NE}):** If each shire contributes voluntarily ($G = g_1 + g_2$), Shire 1 alone would set $120/G = 25 \implies G_1 = 4.8$, while Shire 2 alone would set $80/G = 25 \implies G_2 = 3.2$. In Nash equilibrium, Shire 1 provides the entire $g_1 = 4.8$ metres and Shire 2 free-rides ($g_2 = 0$), yielding total levee height $G^{NE} = 4.8 < 8$ metres—an underprovision of 40%.

Degreelive

Chapter 15: Hidden Action (Moral Hazard) and Principal-Agent Contracts

15.1 The Principal-Agent Problem: Observable vs Unobservable Effort

Asymmetric information divides into two classes: **hidden action (moral hazard)**, where post-contractual effort is unobservable, and **hidden information (adverse selection)**, where pre-contractual type is private. Consider a risk-neutral Principal (a farm owner or corporate board) hiring a risk-averse Agent (a manager) whose reservation utility is $\bar{u}$. The Agent chooses binary effort $e \in \{e_L, e_H\}$ with effort disutility $c(e_H) > c(e_L) = 0$. Stochastic revenue is either high (q_H) or low ($q_L < q_H$), with probability p_H of q_H under high effort and $p_L < p_H$ under low effort. The Agent's expected utility from contingent wages (w_H, w_L) is $p_e u(w_H) + (1 - p_e)u(w_L) - c(e)$, where $u' > 0$ and $u'' < 0$.

Under the **First-Best (Observable Effort)** benchmark, the Principal observes e and forces e_H by paying a fixed wage $w_H = w_L = w^{FB}$ satisfying the binding Participation Constraint (IR): $u(w^{FB}) - c(e_H) = \bar{u}$. Because the risk-neutral Principal absorbs all output variance while giving the risk-averse Agent **full insurance**, risk-sharing is Pareto optimal.

Under **Second-Best (Moral Hazard)**, the Principal observes output $q \in \{q_L, q_H\}$ but not effort e. To induce e_H at minimum expected wage cost, the contract (w_H, w_L) must satisfy two constraints:

- **Individual Rationality (IR):** $p_H u(w_H) + (1 - p_H)u(w_L) - c(e_H) \geq \bar{u}$
- **Incentive Compatibility (IC):** $p_H u(w_H) + (1 - p_H)u(w_L) - c(e_H) \geq p_L u(w_H) + (1 - p_L)u(w_L) - c(e_L)$

Simplifying IC (with $c(e_L) = 0$) yields:

$$(p_H - p_L)[u(w_H) - u(w_L)] \geq c(e_H)$$

Where $p_H - p_L > 0$ is the incremental probability of high output generated by high effort. Because $c(e_H) > 0$, the Principal *must* set $w_H > w_L$ (performance-contingent pay), exposing the risk-averse Agent to income risk! Both IR and IC bind at the cost-minimising second-best contract when the Agent is risk-averse and unlimited liability applies. By Jensen's Inequality, the Principal must pay a strictly higher expected wage $p_H w_H + (1 - p_H)w_L > w^{FB}$ to compensate the Agent for bearing risk; this extra expected wage (or lost surplus) is the **agency cost** of moral hazard, illustrating the fundamental **incentives versus insurance trade-off**.

15.2 Continuous Linear Contracts and the Informativeness Principle

In corporate finance and executive remuneration (such as ASX-listed CEO bonus schemes), suppose output is $q = e + \tilde{\epsilon}$ where $\tilde{\epsilon} \sim N(0, \sigma^2)$, effort cost is $c(e) = \frac{1}{2}ke^2$, and the Agent has CARA utility $-\exp(-r(w - c(e)))$ under a linear contract $w(q) = \alpha + \beta q$. The Agent's certainty equivalent is $CE = \alpha + \beta e - \frac{1}{2}ke^2 - \frac{1}{2}r\beta^2\sigma^2$. Maximising CE with respect to effort gives $e^*(\beta) = \beta/k$. Substituting e^* and binding IR into the Principal's expected profit $\mathrm{E}[q - w] = e - \frac{1}{2}ke^2 - \frac{1}{2}r\beta^2\sigma^2$ yields the optimal incentive intensity:

$$\beta^* = \frac{1}{1 + rk\sigma^2}$$

Where:

- $\beta^* \in (0, 1]$ is the optimal pay-for-performance slope.

DegreeLive

- r is the Agent's Arrow-Pratt absolute risk aversion, σ^2 is output noise variance, and $k = c''(e)$ is the marginal cost curvature of effort.

Higher risk aversion (r) or noisier environments (σ^2) reduce β^* toward zero (stronger insurance, weaker incentives). Moreover, Holmström's **Informativeness Principle** demonstrates that any costless observable signal z correlated with $\widetilde{\epsilon}$(such as industry peer benchmarks) should be included in $w(q,z)$ to filter out common shocks, lowering effective σ^2 and allowing higher-powered incentives at lower risk premia.

Worked Example 15.1: Solving a Binary Moral Hazard Contract and Agency Cost

An Australian pastoral company (risk-neutral Principal) hires a station manager with utility $u(w) - c(e) = \underline{w} - c(e)$ and reservation utility $\bar{u} = 10$. High effort (e_H) costs $c(e_H) = 2$ and yields revenue q_H = \$1,000 with probability $p_H = 0.8$ (and q_L = \$200 with probability 0.2). Low effort (e_L) costs $c(e_L) = 0$ and yields q_H with probability $p_L = 0.4$.

1. **Step 1 — First-Best benchmark (observable effort):** To induce e_H with full insurance ($w_H = w_L = w^{FB}$), binding IR gives $\underline{w^{FB}} - 2 = 10 \implies \underline{w^{FB}} = 12 \implies w^{FB} = 144$ dollars. Expected revenue is \$0.8(1,000) + 0.2(200) = 840*dollars, so first − best expected profit is* \Pi^{FB} = 840 - 144 = 696\$ dollars.
2. **Step 2 — Second-Best contract (unobservable effort):** Let $u_H = \underline{w_H}$ and $u_L = \underline{w_L}$. Binding IC and IR yield two linear equations:

$$(0.8 - 0.4)(u_H - u_L) = 2 \implies u_H - u_L = 5$$

$$0.8u_H + 0.2u_L - 2 = 10 \implies 0.8u_H + 0.2u_L = 12$$

   Substituting $u_H = u_L + 5$ into IR: \$0.8(u_L + 5) + 0.2 u_L = 12 \implies u_L + 4 = 12 \implies u_L^* = 8 *and* u_H^* = 13\$.
3. **Step 3 — Optimal wages, profit, and agency cost:** Squaring u_H^* and u_L^* gives $w_H^* = 13^2 = 169$ dollars and $w_L^* = 8^2 = 64$ dollars. The expected wage bill is:

$$\mathrm{E}[w^*] = 0.8(169) + 0.2(64) = 135.2 + 12.8 = 148$$

   Second-best profit is $\Pi^{SB} = 840 - 148 = 692$ dollars (still superior to inducing e_L, which yields \$0.4(1,000)+0.6(200) - 10^2 = 420). *The < strong > agency cost of moral hazard < /strong > is* \Pi^{FB} - \Pi^{SB} = 696 - 692 = 4\$ dollars per period.

Chapter 16: Hidden Information (Adverse Selection), Signalling & Screening

16.1 Adverse Selection: Akerlof's Market for Lemons

When one side of a transaction holds private pre-contractual information about product quality or personal risk type, competitive markets can suffer **adverse selection**. In George Akerlof's (1970) *Market for Lemons*, used cars (or financial assets) have quality $\theta \sim U[0,1]$, valued by sellers at $v_S(\theta) = \theta$ and by buyers at $v_B(\theta) = 1.5\theta$ (so every trade is socially efficient). If buyers cannot observe θ, they offer a uniform price P based on

average quality $\mathrm{E}[\theta \mid \theta \le P]$. Only sellers with $\theta \le P$ accept P, yielding average quality $\mathrm{E}[\theta \mid \theta \le P] = P/2$. Buyers' maximum willingness to pay is $1.5(P/2) = 0.75P < P$ for every $P > 0$! Consequently, the market unravels completely to $P^* = 0$ and zero trade occurs.

In Australia, adverse selection is a central concern in private health insurance. Under the *Private Health Insurance Act 2007 (Cth)*, insurers are legally bound by **community rating** (charging identical premiums regardless of age or pre-existing health risk). Without countervailing policy, healthy low-risk Australians would drop coverage, raising average claim costs and triggering a death spiral. Australia mitigates this adverse selection via the Medicare Levy Surcharge and **Lifetime Health Cover (LHC)** loading (a 2% premium penalty per year above age 30 that an individual delays taking hospital cover).

16.2 Spence Job Market Signalling

Michael Spence (1973) showed that informed agents can credibly communicate hidden productivity through costly, observable **signals** such as university degrees, provided the signal satisfies the **single-crossing (Spence-Mirrlees) condition**: acquiring education e is strictly less costly for high-ability workers (a_H) than for low-ability workers (a_L):

$$c(e, a_H) = c_H e, \qquad c(e, a_L) = c_L e$$
$$textwith 0 < c_H < c_L$$
$$textand a_H > a_L$$

Where:

- a_H and a_L denote the marginal productivities (and competitive wages) of high- and low-ability workers, even when education adds zero direct productivity.
- $c_H < c_L$ captures lower marginal effort cost of schooling per unit e for high-ability workers.

Competitive employers observe e and pay wage $w(e) = a_H$ if $e \ge e^*$ and $w(e) = a_L$ if $e < e^*$. In a **Separating Perfect Bayesian Equilibrium**, high-ability types choose e^* and earn a_H, while low-ability types choose $e = 0$ and earn a_L. Separation requires neither type wants to mimic the other:

- IC_L (low type prefers $e = 0$ at a_L over e^* at a_H): $a_L \ge a_H - c_L e^* \implies e^* \ge \frac{a_H - a_L}{c_L}$
- IC_H (high type prefers e^* at a_H over $e = 0$ at a_L): $a_H - c_H e^* \ge a_L \implies e^* \le \frac{a_H - a_L}{c_H}$

Combining both bounds yields the range of separating equilibrium education thresholds:

$$\frac{a_H - a_L}{c_L} \le e^* \le \frac{a_H - a_L}{c_H}$$

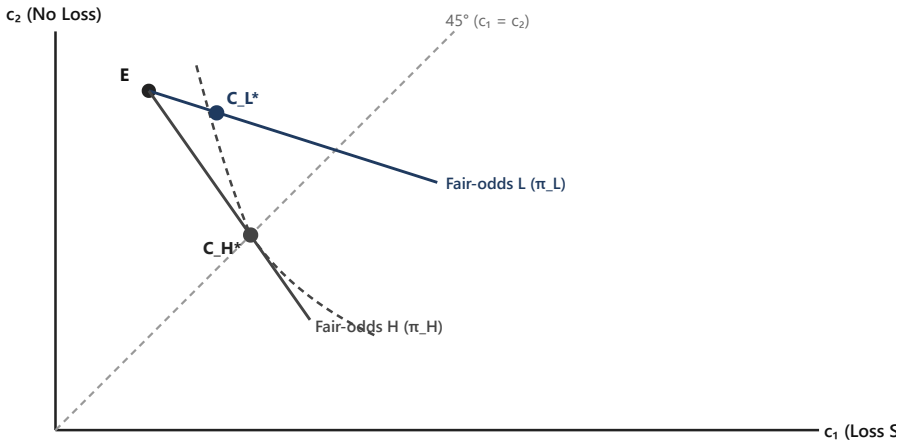
Where the least-cost separating equilibrium sets $e^* = \frac{a_H - a_L}{c_L}$ (selected uniquely by the Cho–Kreps *Intuitive Criterion*, which also eliminates pooling equilibria where both types choose the same education level). Notice that even when education has zero productivity benefit, signalling is privately rational for high-ability graduates yet socially costly ($c_H e^*$ is pure deadweight expenditure relative to full information).

16.3 Rothschild-Stiglitz Competitive Insurance Screening

When the *uninformed* party moves first by offering a menu of contracts to sort informed types, the game is one of **screening**. In Rothschild and Stiglitz (1976), competitive risk-neutral insurers face a share λ of high-risk drivers (accident probability π_H) and $1 - \lambda$ of low-risk drivers ($\pi_L < \pi_H$), each with wealth W_0 and potential loss D. In

DegreeLive

state-space (c_1, c_2)—loss vs no-loss wealth—each type's indifference curve has slope $-\frac{\pi_i}{1-\pi_i}\frac{u'(c_1)}{u'(c_2)}$. Because $\pi_H > \pi_L$, at every point (c_1, c_2) high-risk types have a **steeper indifference curve** (single-crossing property): they are willing to sacrifice more no-loss wealth for an extra dollar in the loss state.



Three sharp equilibrium theorems emerge from the Rothschild-Stiglitz model:

1. **No pooling equilibrium can exist:** At any pooling contract on the pooled fair-odds line, a rival insurer can offer a slightly lower-coverage, lower-premium contract inside the wedge between the flat low-risk indifference curve and the steep high-risk indifference curve, attracting *only* profitable low-risk customers ("cream-skimming") and bankrupting the pooling contract.
2. **Unique separating candidate (C_H^*, C_L^*):** In any Nash equilibrium, high-risk types receive **full insurance** at their actuarially fair price (C_H^* on the $45°$ line), while low-risk types receive **partial insurance** (C_L^* on their own fair-odds line, carrying a deductible just large enough that high-risk types weakly prefer C_H^* over mimicking C_L^*).
3. **Possible non-existence:** If the fraction of high-risk types λ is very small, the pooled fair-odds line lies very close to the low-risk fair-odds line; a pooling contract can then offer low-risk types higher utility than the distorted separating contract C_L^*, causing even (C_H^*, C_L^*) to fail—meaning no pure-strategy Nash equilibrium exists at all.

Worked Example 16.1: Spence Job Market Signalling Equilibrium

In an Australian graduate recruitment market, 50% of applicants are high-ability with productivity a_H = \$120,000 and 50% are low-ability with productivity a_L = \$70,000. Acquiring e years of tertiary credentials costs high-ability graduates $c_H e = 10,000e$ and low-ability graduates $c_L e = 25,000e$, adding zero direct productivity.

1. **Step 1 — Separating equilibrium bounds:** Applying the incentive compatibility thresholds:

$$e^*_{min} = \frac{120,000 - 70,000}{25,000} = 2$$

$$textyears, \qquad e^*_{max} = \frac{120,000 - 70,000}{10,000} = 5$$

$$textyears$$

Any threshold $e^* \in [2, 5]$ supports a separating equilibrium where high-ability graduates acquire e^* years of study (earning net payoff $120,000 - 10,000e^*$) and low-ability graduates acquire 0 years (earning \$70,000).

2. **Step 2 — Intuitive Criterion refinement:** Under the Cho–Kreps Intuitive Criterion, high-ability graduates select the least-cost separating threshold $e^* = 2$ years, earning net payoff $120,000 - 10,000(2) = 100,000$ dollars.
3. **Step 3 — Compare with pooling and full information:** Under pooling at $e = 0$, every applicant receives the expected productivity wage $0.5(120,000) + 0.5(70,000) = 95,000$ dollars. Notice that high-ability workers earn $100,000 > 95,000$ under separating at $e^* = 2$, whereas if c_H were \$15,000 (making $120,000 - 30,000 = 90,000 < 95,000$), high-ability graduates would actually be *worse off* in the separating equilibrium than under pooling, even though pooling still fails the Intuitive Criterion!

Worked Example 16.2: Rothschild-Stiglitz Separating Insurance Menu

Suppose drivers have initial wealth $W_0 = 100$, face accident loss $D = 64$, and have Bernoulli utility $u(c) = \sqrt{c}$. High-risk drivers have $\pi_H = 0.50$ and low-risk drivers have $\pi_L = 0.25$.

1. **Step 1 — High-risk full-insurance contract (C_H^*):** High-risk drivers pay fair premium $\pi_H D = 0.50(64) = 32$ for full coverage ($K_H = 64$), receiving guaranteed wealth $c_{1,H}^* = c_{2,H}^* = 100 - 32 = 68$ and utility $u(C_H^*) = \sqrt{68} \approx 8.2462$.
2. **Step 2 — Low-risk separating contract (C_L^*):** Low-risk types receive indemnity $K_L < 64$ at fair premium $\gamma_L K_L = 0.25K_L$, giving state wealths $c_{1,L} = 36 + 0.75K_L$ and $c_{2,L} = 100 - 0.25K_L$. To prevent high-risk types from mimicking C_L^*, K_L must satisfy the binding IC_H constraint:

$$0.5\sqrt{36 + 0.75K_L} + 0.5\sqrt{100 - 0.25K_L} = \sqrt{68} \approx 8.2462$$

Testing values along the low-risk fair-odds line: at $K_L = 0, LHS = 0.5(6) + 0.5(10) = 8.0 < 8.2462$. At $K_L = 16, c_1 = 48, c_2 = 96, LHS = 0.5\sqrt{48} + 0.5\sqrt{96} = 3.4641 + 4.8990 = 8.3631 > 8.2462$. Solving numerically gives $K_L^* \approx 9.95$: low-risk drivers receive only \$9.95 of coverage out of a \$64 loss (retaining a self-funded deductible of \$54.05) strictly to separate themselves from high-risk drivers.

16.4 Mechanism Design, the Revelation Principle & Spectrum Auctions

In market environments characterised by asymmetric information, **mechanism design** inverts traditional game theory: rather than predicting equilibrium outcomes under fixed rules, the social planner or seller designs the game's rules (the mechanism) to induce strategic agents to reveal their private valuations truthfully. The fundamental theoretical pillar is the **Revelation Principle**: any Bayesian Nash equilibrium of an arbitrary game can be replicated by a direct, incentive-compatible mechanism in which agents find it weakly dominant to report their true private types.

In private-value auction settings where N risk-neutral bidders have valuations drawn independently from a continuous distribution $F(v)$, the **Revenue Equivalence Theorem** (Vickrey, 1961; Myerson, 1981) establishes that any auction mechanism satisfying two conditions—(1) the object is allocated to the highest bidder, and (2) the lowest valuation bidder expects zero surplus—yields the identical expected revenue to the seller and identical expected payments from each bidder type.

DegreeLive

$$E[R] = N\int_{\underline{v}}^{\bar{v}} \left(v - \frac{1-F(v)}{f(v)}\right) F(v)^{N-1} f(v), dv$$

Where $v - \frac{1-F(v)}{f(v)}$ represents the bidder's *virtual valuation*, reflecting information rent that must be conceded to induce truthful disclosure.

Worked Numerical Example: Optimal Reserve Pricing in an Australian ACMA 5G Spectrum Auction
The Australian Communications and Media Authority (ACMA) auctions regional 5G spectrum licences. Suppose $N = 3$ major telecommunications carriers participate, with independent private valuations uniformly distributed over $[0,; 100]$ million AUD ($F(v) = v/100, f(v) = 1/100$).

Determine ACMA's optimal reserve price r^* and compare expected revenue under an unconstrained auction vs an auction with the optimal reserve price.

Under Myerson's optimal auction design, the optimal reserve price r^* equates the virtual valuation to the seller's cost (here assumed zero):

$$v - \frac{1-F(v)}{f(v)} = v - \frac{1 - v/100}{1/100} = v - (100 - v) = 2v - 100 = 0 \implies r^* = 50 \text{ million AUD}$$

In an unconstrained second-price auction without a reserve price, expected revenue equals the expected second-highest order statistic from $N = 3$ uniform draws:

$$E[v_{(2)}] = \frac{N-1}{N+1} \times 100 = \frac{2}{4} \times 100 = 50.00 \text{ million AUD}$$

Introducing the optimal reserve price of $r^* = 50$ million AUD risks leaving the licence unsold if all bids fall below 50 (probability $(0.5)^3 = 0.125$), but significantly extracts revenue when at least one bidder exceeds 50, elevating expected auction revenue to:

$$E[R^*] = 100 \times \left(1 - \frac{N}{N+1}\left(\frac{r^*}{100}\right) - \frac{1}{N+1}\left(\frac{r^*}{100}\right)^{N+1}\right) = 54.6875 \text{ million AUD}$$

DegreeLive

The optimal reserve price generates an expected revenue gain of $54.6875 - 50.00 = +4.6875$ million AUD (+9.38%), illustrating the powerful public finance utility of mechanism design in Australian spectrum allocation.

DegreeLive