GRADUATE MICROECONOMICS

UC IRVINE

2015-2016

(Partial and Incomplete)

"A barbarian is not aware that he is a barbarian."
– Jack Vance, *Big Planet*

Quarter I

"The person who, let us say, expects generosity from a bank, efficient flexibility from a government agency, open-mindedness from a religious institution will be disappointed ... The poor fool might as quickly discover love among the mantises."
– Jack Vance, *The Book of Dreams*

This section will ask the question: how does the world work? The microeconomic picture consists of individual actors, with self–interest, coherency of utility, and rationality (defined by tautological assumptions), who may be professionals, producers, or consumers, who may care about status, hierarchy, and/or happiness, and who have demand for inputs, demand for outputs, and who can supply outputs and inputs. In this quarter, the focus will be on the basic neoclassical theory of production and consumption, and the introductory ideas from general equilibrium. The main purpose is to equip you with the basic tools that will allow you to analyze a broad range of economic problems on your own using the neoclassical approach. Acquiring the necessary technical competence and developing an intuitive grasp of each topic is usually a slow process that requires repetition and viewing of the same subject from different angles. Problem–solving is very important for the assimilation of the material covered.

1.1. Neoclassical Production Theory

1.1.1. Methodology

Reading: Friedman, Milton, "The Methodology of Positive Economics," in Essays in Positive Economics, 1953, Chicago: University of Chicago Press.

1.1.2. Alternative Views

Reading: Ronald H. Coase, "The Institutional Structure of Production" (University of Chicago Law Occasional Paper No. 28, 1992).

1.1.3. Production Functions

There is 1 output and n inputs. The output, y, is produced with a Neoclassical production function, $F : \mathbb{R}^n \to \mathbb{R}$, such that

$$y = F(x_1, \ldots, x_n).$$

The production function is increasing in all inputs. Some example of productions functions are

- Cobb–Douglas Production: $F(x_1, x_2) = x_1^\alpha x_2^\beta$, where $\alpha, \beta > 0$.
- Substitutes in Production: $F(x_1, x_2) = Ax_1 + Bx_2$, where $A, B > 0$.
- Leontiff (Fixed-Coefficients) Production: $F(x_1, x_2) = \min\{cx_1, dx_2\}$, s. t. $c, d > 0$.

In the short–run, the production function is fixed. Note that input and output analysis is often used.

Definition: The marginal product is the change in output that results from a change in an input

$$\text{MP} \equiv \frac{\Delta y}{\Delta x_i},$$

or more formally for a continuous function

$$\text{MP} \equiv \frac{\partial F(x)}{\partial x_i}.$$

Example: Marginal Productivity with Leontiff Production
Note that, given a Leontiff production function, $F(x_1, x_2) = \min\{cx_1, dx_2\}$, then

$$\frac{\partial F(x_1,x_2)}{\partial x_2} = \begin{cases} C & \text{if } x_2 < x_1, \\ 0 & \text{if } x_2 \geq x_1. \end{cases}$$

Now, consider the change in marginal product from a change in an input.

Example: Marginal Change in MP with Leontiff Production
Given a Leontiff production function, then

$$\frac{\partial^2 F(x_1,x_2)}{\partial x_1 \partial x_2} = 0,$$

and

$$\frac{\partial(\frac{\partial F(x_1,x_2)}{\partial x_1})}{\partial x_1} = 0.$$

Example: Marginal Change in MP with Cobb–Douglas Production
Given a Cobb–Douglas production function, then

$$\frac{\partial^2 F(x_1,x_2)}{\partial x_1 \partial x_2} = \frac{\partial(\alpha x_1^{\alpha-1} x_2^{\beta})}{\partial x_2} = \alpha\beta x_1^{\alpha-1} x_2^{\beta-1} > 0,$$

and

$$\frac{\partial(\frac{\partial F(x_1,x_2)}{\partial x_1})}{\partial x_1} = \frac{\partial(\alpha x_1^{\alpha-1} x_2^{\beta})}{\partial x_1} = (\alpha-1)\alpha x_1^{\alpha-2} x_2^{\beta} \begin{cases} > 0 & \text{if and only if } \alpha > 1, \\ = 0 & \text{if and only if } \alpha = 1, \\ < 0 & \text{if and only if } \alpha < 1. \end{cases}$$

Definition: When the change in the marginal product of an input from a change in the input is negative, the production function exhibits decreasing marginal returns in that input.

Definition: An isoquant is the set of input combinations such that any combination of inputs yields a set level of output

$$\{(x_1, x_2) | \overline{y} = F(x_1, x_2)\}.$$

Example: An Isoquant with Cobb–Douglas Production
Given a Cobb–Douglas production function an isoquant is given by

$$\overline{y} = x_1^{\alpha} x_2^{\beta}$$

$$x_2^{\beta} = \frac{\overline{y}}{x_1^{\alpha}}$$

$$x_2 = \left(\frac{\overline{y}}{x_1^{\alpha}}\right)^{\frac{1}{\beta}}.$$

Example: An Isoquant with Substitutes in Production
Given substitutes in production, then any given isoquant is linear

$$\overline{y} = Ax_1 + Bx_2$$

$$x_2 = \frac{\overline{y}}{B} - \frac{A}{B} x_1.$$

Considering the shape of an isoquant given a particular production function is important.

Example: The Shape of an Isoquant with Cobb–Douglas Production
Given a Cobb–Douglas production function, $F(x_1, x_2) = x_1^{\alpha} x_2^{\beta}$, an isoquant is given by

$$x_2 = \overline{y}^{\frac{1}{\beta}} x_1^{-\frac{\alpha}{\beta}}.$$

Notice that the the isoquant is decreasing

$$\frac{\partial x_2}{\partial x_1} = \Big(-\frac{\alpha}{\beta} \Big) x_1^{-\frac{\alpha}{\beta}-1} \overline{y}^{\frac{1}{\beta}} < 0,$$

and decreasing at an increasing rate

$$\frac{\partial^2 x_2}{\partial x_1^2} = \Big(-\frac{\alpha}{\beta} \Big)\Big(-\frac{\alpha}{\beta} - 1 \Big) x_1^{-\frac{\alpha}{\beta}-2} \overline{y}^{\frac{1}{\beta}} > 0.$$

Definition: Two inputs are complementary if

$$\frac{\partial^2 F(x)}{\partial x_i \partial x_j} > 0.$$

Theorem: The Implicit Function Theorem For a continuously differentiable function, $h(z_1, \ldots, z_n)$, with non–zero partial derivatives, consider a point $(z_1', \ldots, z_n')$ such that $h(z_1', \ldots, z_n') = 0$. There is an open neighborhood around $(z_1', \ldots, z_n')$ such that for every point, $(z_1, \ldots, z_n)$, in that neighborhood there exists a function $g(z_2, \ldots, z_n) = z_1$. Furthermore, $g(z_2, \ldots, z_n) = z_1$ has the property that

$$\frac{\partial g(z_2, \ldots, z_n)}{\partial z_i} = -\frac{\partial h(z_1, \ldots, z_n)/\partial z_i}{\partial h(z_1, \ldots, z_n)/\partial z_1},$$

for $i = 2, \ldots, n$.

It is important to know how to perform total differentiation on a productions function. Given $y = F(x_1, x_2)$, then it follows that

$$\begin{aligned}
F(x_1, x_2) - y &= 0 \\
\mathrm{d}F(x_1, x_2) &= 0 \\
\frac{\partial F(x_1, x_2)}{\partial x_1}\,\mathrm{d}x_1 + \frac{\partial F(x_1, x_2)}{\partial x_2}\,\mathrm{d}x_2 &= 0 \\
\frac{\partial F(x_1, x_2)}{\partial x_1}\frac{\mathrm{d}x_1}{\mathrm{d}x_2} + \frac{\partial F(x_1, x_2)}{\partial x_2}\frac{\mathrm{d}x_2}{\mathrm{d}x_2} &= 0 \\
\frac{\partial F(x_1, x_2)}{\partial x_1}\frac{\mathrm{d}x_1}{\mathrm{d}x_2} &= -\frac{\partial F(x_1, x_2)}{\partial x_2} \\
\frac{\mathrm{d}x_1}{\mathrm{d}x_2} &= -\frac{\partial F(x_1, x_2)/\partial x_2}{\partial F(x_1, x_2)/\partial x_1}.
\end{aligned}$$

Definition: The marginal rate of technical substitution (MRTS) is the rate that one input can be substituted for another while remaining at the same level of output. It is the slope of an isoquant of a production function and is given by

$$\text{MRTS} = -\frac{\partial F(x_1, x_2)/\partial x_2}{\partial F(x_1, x_2)/\partial x_1}.$$

Notice that from the Implicit Function Theorem that the MRTS is negative and it can be shown that the MRTS is decreasing (i.e. $\text{MRTS}' < 0$.).

Definition: A production function exhibits constant returns to scale if $\forall\theta > 0$

$$F(\theta x) = \theta F(x).$$

Definition: A production function exhibits increasing returns to scale if $\forall\theta > 1$

$$F(\theta x) > \theta F(x).$$

Definition: A production function exhibits decreasing returns to scale if $\forall\theta > 1$

$$F(\theta x) < \theta F(x).$$

Example: Given a Cobb–Douglas production function, $F(x_1, x_2) = x_1^\alpha x_2^\beta$, then

$$F(\theta x_1, \theta x_2) = \theta^{\alpha+\beta} x_1^\alpha x_2^\beta.$$

$$\text{If } \theta^{\alpha+\beta} \begin{cases} > \theta & \text{then there are increasing returns to scale,} \\ = \theta & \text{then there are constant returns to scale,} \\ < \theta & \text{then there are decreasing returns to scale.} \end{cases}$$

Thus, the returns to scale of a Cobb–Douglas production function depend on if $\alpha + \beta$ is greater than, less than, or equal to 1.

Definition: A function is homogeneous of degree r if $\forall\theta > 0$

$$f(\theta x) = \theta^r f(x).$$

Theorem: Euler's Law
If $f(x)$ is homogeneous of degree r, then its derivatives are homogeneous of degree $r - 1$ and

$$rf(x_1, \ldots, x_n) = \sum_{i=1}^{n} \frac{\partial f(x)}{\partial x_i} x_i.$$

1.1.4. Profit Maximization

A firm may wish to maximize its profits. If so, the firm's objective is

$$\max_{y,x_1,\ldots,x_n} \pi = py - \sum_{i=1}^{n} w_i x_i \;\text{ s. t. }\; y = f(x_1,\ldots,x_n).$$

Substituting output for the production function yields

$$\max_{y,x_1,\ldots,x_n} \pi = pf(x_1,\ldots,x_n) - \sum_{i=1}^{n} w_i x_i.$$

The first order conditions are

$$p\frac{\partial f(x)}{\partial x_i} - w_i = 0 \;\text{ for all } i = 1,\ldots,n.$$

Definition: The input demand function is

$$x_i^* = x_i(p,w) \;\text{ for all } i = 1,\ldots,n.$$

Definition: The factor demand function is

$$x(p,w) = \arg\max_x pf(x) - wx.$$

Definition: The change in the input demand function with respect to price is positive

$$\frac{\partial x_i(p,w)}{\partial p} > 0 \;\text{ for all } i = 1,\ldots,n..$$

Next look at the change in an input choice, x_i, with its input price. The input demand function has the property of negativity

$$\frac{\partial x_i}{\partial w_i} < 0,$$

and

$$\frac{\partial x_j}{\partial w_i} < 0.$$

Furthermore, the input demand function has the property of symmetry

$$\frac{\partial x_i}{\partial w_j} = \frac{\partial x_j}{\partial w_i} \;\text{ for all } i \text{ and } j.$$

Definition: The output supply function is given by

$$y(p,w) \equiv f(x_1^*,\ldots,x_n^*) = f(x_1(p,w),\ldots,x_n(p,w)) = f(x(p,w)).$$

The change in output with respect to the price is positive

$$\frac{\partial y(p,w)}{\partial p} > 0.$$

Proof.

$$\frac{\partial y(p,w)}{\partial p} = \underset{(+)}{\frac{\partial f(\cdot)}{\partial x_1}}\underset{(+)}{\frac{\partial x_1}{\partial p}}, + \cdots + \underset{(+)}{\frac{\partial f(\cdot)}{\partial x_n}}\underset{(+)}{\frac{\partial x_n}{\partial p}}$$

□

The profit function is given by

$$\pi(p,w) \equiv py(p,w) - \sum_{i=1}^{n} w_i x_i(p,w).$$

Theorem: The profit function is non–decreasing in output price, p,

$$p_1 \geq p^2 \Rightarrow \pi(p^1,w) \geq \pi(p^2,w).$$

Proof. Suppose

$$p^1 \geq p^2,$$

then by the profit maximization assumption, the profit function must satisfy

$$\pi(p^1,w) = p^1 y(p,w) - wx(p^1,w) \geq p^1 y - wx \quad \text{for all } y \text{ and } x.$$

It follows that

$$\pi(p^1,w) = p^1 y(p,w) - wx(p^1,w) \geq p^1 y(p^2,w) - wx(p^2,w)$$

$$\pi(p^1,w) \geq p^1 y(p^2,w) - wx(p^2,w) \geq p^2 y(p^2,w) - wx(p^2,w) \equiv \pi(p^2,w)$$

$$\therefore \pi(p^1,w) \geq \pi(p^2,w).$$

□

Theorem: The profit function is non–increasing in input prices, w,

$$w^1 \geq w^2 \Rightarrow \pi(p,w^1) \leq \pi(p,w^2).$$

Proof. Left to the reader as a problem. □

Theorem: The profit function is homogeneous of degree 1 in prices, (p,w),

$$\pi(\theta p, \theta w) = \theta \pi(p,w) \quad \text{for all } \theta > 0.$$

Proof.

$$\pi(\theta p, \theta w) = \theta p y(\theta p, \theta w) - \theta w x(\theta p, \theta w)$$

$$\pi(\theta p, \theta w) = \theta[py(\theta p, \theta w) - wx(\theta p, \theta w)] = \theta \pi(p,w),$$

where

$$py(\theta p, \theta w) - wx(\theta p, \theta w) = py(p,w) - wx(p,w) = \pi(p,w),$$

because input choices and output level remain the same

$$\max_x \theta(pf(x) - wx)$$

$$\theta\left(p\frac{\partial f(x)}{\partial x_i} - w_i\right) = 0.$$

$$\therefore \pi(\theta p, \theta w) = \theta \pi(p,w).$$

□

The intuition to the above proof is that the profit function, $\pi(p,w)$ is convex.

Theorem: The profit function is convex in prices. Let (p^1, w^1) and (p^2, w^2) be price vectors. Let

$$(p^0, w^0) = \theta(p^1, w^1) + (1-\theta)(p^2, w^2),$$

then, for any $\theta \in [0, 1]$,

$$\pi(p^0, w^0) \leq \theta\pi(p^1, w^1) + (1-\theta)\pi(p^2, w^2).$$

Proof.

$$\begin{aligned}
\pi(p^0, w^0) &= p^0 y^0 - w^0 x^0 \\
\pi(p^0, w^0) &= (\theta p^1 + (1-\theta)p^2)y^0 - (\theta w^1 + (1-\theta)w^2)\theta^0 \\
\pi(p^0, w^0) &= \theta p^1 y^0 + (1-\theta)p^2 y^0 - \theta w^1 x^0 - (1-\theta)w^2 x^0 \\
\pi(p^0, w^0) &= \theta(p^1 y^0 - w^1 x^0) + (1-\theta)(p^2 y^0 - w^2 y^0),
\end{aligned}$$

and by profit maximization

$$p^1 y^0 - w^1 x^0 \leq p^1 y^1 - w^1 x^1.$$

Thus the profit function is convex in prices

$$\pi(p^0, w^0) \leq \theta\pi(p^1, w^1) + (1-\theta)\pi(p^2, w^2).$$

□

Theorem: Hotelling's Lemma

Hotelling's Lemma is that the change in profit with respect to output price is proportional to the output level

$$\frac{\partial \pi(p, w)}{\partial p} = y(p, w),$$

and the change in profit with respect to a change in an input price is negatively proportional to the input level

$$\frac{\partial \pi(p, w)}{\partial w_i} = -x_i(p, w).$$

Proof. First,

$$\frac{\partial \pi(p, w)}{\partial p} = \frac{\partial[py(p, w) - wx(p, w)]}{\partial p}.$$

Utilizing the Envelope Theorem

$$\begin{aligned}
\frac{\partial \pi(p, w)}{\partial p} &= y(p, w) + \left[\frac{\partial f(x)}{\partial x_1}\frac{\partial x_1}{\partial} + \cdots + \frac{\partial f(x)}{\partial x_n}\frac{\partial x_n}{\partial}\right] - w_1\frac{\partial x_1}{\partial p} - \cdots - w_n\frac{\partial x_n}{\partial p} \\
\frac{\partial \pi(p, w)}{\partial p} &= y(p, w) + \frac{\partial x_1}{\partial p}\left[p\frac{f(x)}{\partial x_1} - w1\right] + \cdots + \frac{\partial x_n}{\partial p}\left[p\frac{f(x)}{\partial x_n} - wn\right]
\end{aligned}$$

$$\therefore \frac{\partial \pi(p, w)}{\partial p} = y(p, w).$$

Second, note that

$$\frac{\partial \pi(p, w)}{\partial w_i} = \frac{\partial(pf(x(p, w)) - wx(p, w))}{\partial w_i}$$

$$\therefore \frac{\partial \pi(p, w)}{\partial w_i} = -x_i(p, w).$$

□

Example: A Monopolist
A monopolist sets the price, $p(y)$, and faces the profit maximization problem

$$\max_{y,x} p(y)y - wx \quad \text{s. t. } f(x) = y \text{ and } p'(y) < 0.$$

First, substitute output for the production function

$$\max_{x} p(f(x))f(x) - wx.$$

The first order condition with respect to x_i is

$$p'(f(x))\frac{\partial f(x)}{\partial x_i}f(x) + p(f(x))\frac{\partial f(x)}{\partial x_i} - w_i = 0.$$

Let $p = p(f(x))$. Then

$$[p'(f(x))f(x) + p]\frac{\partial f(x)}{\partial x_i} - w_i = 0.$$

Equating the monopolist's first–order condition to the first–order condition in perfect competition

$$[p'(f(x^m))f(x^{\mathrm{M}}) + p]\frac{\partial f(x^{\mathrm{M}})}{\partial x_i} = p\frac{\partial f(x^{\mathrm{PC}})}{\partial x_i}.$$

Thus

$$p\frac{\partial f(x^{\mathrm{M}})}{\partial x_i} + p'(f(x^m))f(x^{\mathrm{M}})\frac{\partial f(x^{\mathrm{M}})}{\partial x_i} = p\frac{\partial f(x^{\mathrm{PC}})}{\partial x_i},$$

and it can be concluded that

$$\frac{\partial f(x^{\mathrm{M}})}{\partial x_i} > \frac{\partial f(x^{\mathrm{PC}})}{\partial x_i}.$$

The marginal product of the monopolist will be higher than the perfectly competitive firm. Furthermore, the marginal product of an input is higher than its real price, and there is unexploited welfare. This is known as a dead–weight loss.

1.1.5. Cost Minimization

A firm may wish to minimize its costs, c. If so, the firm's objective is

$$\min c = \sum_{i=1}^{n} w_i x_i \text{ s. t. } y = f(x).$$

The corresponding Lagrangian is

$$\mathscr{L}(x, \lambda) = \sum_{i=1}^{n} w_i x_i - \lambda(f(x) - y),$$

where $\lambda^* > 0$ and x_i^* for all $i = 1, \dots, n$. The corresponding first order conditions are

$$\begin{aligned} \frac{\partial \mathscr{L}}{\partial x_i} &= w_i - \lambda \frac{\partial f(x_i^*}{\partial x_i} = 0, \\ \frac{\partial \mathscr{L}}{\partial \lambda} &= -f(x^*) + y \quad = 0. \end{aligned}$$

Definition: The conditional factor demand function is

$$h_i(w, y) \equiv x_i^*.$$

Definition: The technical rate of substitution (TRS) is

$$\text{TRS} \equiv -\frac{w_i}{w_j} = -\frac{\partial f(x^*)/\partial x_i}{\partial f(x^*)/\partial x_j}.$$

Definition: An isocost curve is the set of input combinations such that any combination of inputs yields a set level of cost.

Let cost be

$$c = w_1 x_1 + w_2 x_2.$$

Then an isocost curve is given by

$$x_2 = \frac{c}{w_2} - \frac{w_1}{w_2} x_1.$$

Example: Cost Minimization with Cobb–Douglas Production
Given a Cobb–Douglas production function,

$$y = f(x_1, x_2) = x_1^{\alpha} x_2^{\beta},$$

the cost minimization problem for the firm is

$$\min_{x_1, x_2} w_1 x_1 + w_2 x_2 \text{ s. t. } y = x_1^{\alpha} x_2^{\beta}.$$

The Lagrangian can be written

$$\mathscr{L}(x, \lambda) = w_1 x_1 + w_2 x_2 - \lambda(x_1^{\alpha} x_2^{\beta} - y).$$

The first–order conditions are

$$\begin{aligned} \frac{\partial \mathscr{L}}{\partial x_1} &= w_1 - \lambda \alpha x_1^{\alpha-1} x_2^{\beta} = 0 \\ \frac{\partial \mathscr{L}}{\partial x_2} &= w_2 - \lambda \beta x_1^{\alpha} x_2^{\beta-1} = 0 \\ \frac{\partial \mathscr{L}}{\partial \lambda} &= \quad -x_1^{\alpha} x_2^{\beta} + y \quad = 0. \end{aligned}$$

It follows that

$$\frac{\alpha x_1^{\alpha-1} x_2^{\beta}}{\beta x_1^{\alpha} x_2^{\beta-1}} = \frac{w_1}{w_2}$$
$$\frac{\alpha x_2^*}{\beta x_1^*} = \frac{w_1}{w_2},$$

then

$$y = x_1^* \left(\frac{w_1}{w_2} \frac{\beta}{\alpha} x_1^* \right)^{\beta}$$
$$y = x_1^{*\alpha+\beta} \left(\frac{w_1}{w_2} \frac{\beta}{\alpha} \right)^{\beta}.$$

Thus, the conditional factor demand functions given Cobb–Douglas production are

$$x_1^* = y^{\frac{1}{\alpha+\beta}} \left(\frac{w_2}{w_1} \frac{\alpha}{\beta} \right)^{\frac{\beta}{\alpha+\beta}}$$
$$x_2^* = y^{\frac{1}{\alpha+\beta}} \left(\frac{w_1}{w_2} \frac{\beta}{\alpha} \right)^{\frac{\alpha}{\alpha+\beta}}.$$

Definition: A cost function is defined as

$$c(w, y) \equiv \sum_{i=1}^{n} w_i h_i(w, y).$$

Example: The cost function given the two input Cobb–Douglas production function is

$$w_1 x_1^* + w_2 x_2^* = w_1 \left[y^{\frac{1}{\alpha+\beta}} \left(\frac{w_2}{w_1} \frac{\alpha}{\beta} \right)^{\frac{\beta}{\alpha+\beta}} \right] + w_2 \left[y^{\frac{1}{\alpha+\beta}} \left(\frac{w_1}{w_2} \frac{\beta}{\alpha} \right)^{\frac{\alpha}{\alpha+\beta}} \right].$$

Theorem: Cost is non–decreasing in input prices, w, that is, if

$$w_i^1 \geq w_i^2 \Rightarrow c(w^1, y) \geq c(w^2, y).$$

Proof. Assume that $w_j^1 \geq w_j^2$. Then

$$c(w^1, y) = \sum_{j=1}^{n} w_j^1 h_j(w^1, y) \geq \sum_{j=1}^{n} w_j^2 h_j(w^1, y),$$

is not cost minimizing. So,

$$c(w^1, y) = \sum_{j=1}^{n} w_j^1 h_j(w^1, y) \geq \sum_{j=1}^{n} w_j^2 h_j(w^2, y) = c(w^2, y),$$

is also not cost minimizing. Thus,

$$c(w^1, y) \geq c(w^2, y).$$

□

Proof. Alternative proof of non–decreasing cost with respect to input prices.
Given that $w_j^1 \geq w_j^2$, let

$$c(w^2, y) = \sum_{i=1}^{n} w_i^2 x_i^2,$$

where

$$x_i^1 \equiv h_i(w^1, y)$$
$$x_i^2 \equiv h_i(w^2, y).$$

Then by cost minimization

$$\sum_{i=1}^{n} w_i^2 x_i^2 \leq \sum_{i=1}^{n} w_i^2 x_i,$$

for all x_i such that $f(x) = y$. It follows that

$$c(w^2, y) \leq \sum_{i=1}^{n} w_i^2 x_i^1 \leq \sum_{i=1}^{n} w_i^1 x_i^1 \equiv c(w^1, y).$$

Thus,

$$c(w^2, y) \leq c(w^1, y).$$

□

Theorem:

- A conditional factor demand function is homogeneous of degree 0

$$h_i(\theta w, y) = h_i(w, y).$$

- A cost function is homogeneous of degree 1

$$c(\theta w, y) = \theta c(w, y).$$

Proof. When all input prices rise proportionally, then the combination of conditional factor demands will remain the same to minimize cost at a given level of output. Thus, the conditional factor demand functions must be homogeneous of degree 0 (formally prove if you wish). It follows that

$$c(\theta w, y) = \sum_{i=1}^{n} \theta w_i h_i(\theta w, y) = \sum_{i=1}^{n} \theta w_i h_i(w, y) = \theta c(w, y)$$

$$\therefore c(\theta w, y) = \theta c(w, y).$$

□

Example: Assume that there are inputs, x, of a logarithmic production function

$$f(x) = \begin{cases} 0 & \text{if } x \leq 1. \\ \ln(x) & \text{if } x > 1. \end{cases}$$

for output, y. The cost minimization problem is

$$\min_x wx \quad \text{s. t. } f(x) = y.$$

Assuming that $x > 1$, then

$$\begin{aligned} f(x) &= \ln(x) \\ y &= \ln(x) \\ x &= \mathrm{e}^y. \end{aligned}$$

The cost function is

$$c(w, y) = w\mathrm{e}^y.$$

The cost function is increasing in input prices

$$\frac{\partial c}{\partial w} = \mathrm{e}^y = x,$$

where $x \geq 0$.

The cost function is increasing in output

$$\frac{\partial c}{\partial y} = w\mathrm{e}^y = wx,$$

where $x \geq 0$.

Example: Assume that there are two sets of two inputs, x_1 and x_2, x_3 and x_4, that are substitutes in a Leontiff production function

$$f(x) = \min\{x_1, x_2\} + 2\min\{x_3, x_4\},$$

for output, y. The cost minimization problem is

$$\min_x wx \quad \text{s. t. } f(x) \geq y.$$

If all inputs are used, then $x_1^* = x_2^*$ and $x_3^* = x_4^*$. It follows that

$$y = x_1 + 2x_3.$$

The three cases are as follows.

(1) If $(w_1 + w_2)y > (w_3 + w_4)2y$, then $x_1^* = x_2^* = 0$ and $x_3^* = x_4^* = y$,

(2) If $(w_1 + w_2)y < (w_3 + w_4)2y$, then $x_1^* = x_2^* = y$ and $x_3^* = x_4^* = 0$

(3) If $(w_1 + w_2)y = (w_3 + w_4)2y$, then $x_1^* = x_2^* = \alpha y$ and $x_3^* = x_4^* = (1 - \alpha)y$,

where $\alpha \in [0, 1]$. The cost function is

$$c(w, y) = \begin{cases} \frac{1}{2}(w_3 + w_4)y & \text{if } x_1^* = x_2^* = 0 \text{ and } x_3^* = x_4^* = y, \\ (w_1 + w_2)y & \text{otherwise.} \end{cases}$$

Theorem: The cost function, $c(w, y)$, is concave in input prices, w.

Proof. First,

$$c(\theta w^1 + (1-\theta)w^2, y) = (\theta w^1 + (1-\theta)w^2)x^0 = \theta x^1 x^0 + (1-\theta)w^2 x^0,$$

where

$$\begin{aligned} x^0 &\equiv h(\theta w^1 + (1-\theta)w^2, y) \\ x^1 &\equiv h(w^1, y) \\ x^2 &\equiv h(w^2, y). \end{aligned}$$

Then by cost minimization

$$\theta x^1 x^0 + (1-\theta)w^2 x^0 \geq \theta^1 x^1 + (1-\theta)w^2 x^2 = \theta c(w^1, y) + (1-\theta)c(w^2, y).$$

Thus,

$$c(\theta w^1 + (1-\theta)w^2, y) \geq \theta c(w^1, y) + (1-\theta)c(w^2, y).$$

□

Theorem: Shepard's Lemma
Shepard's Lemma states that the change in cost with respect to an input price is proportional to the level of the input's conditional demand

$$\frac{\partial c(w, y)}{\partial w_i} = h_i(w, y) \quad \text{for all } i = 1, \ldots, n.$$

Proof. Let x^* be the cost minimizing at w^*. Define the function

$$z(w) \equiv \sum_{i=1}^{n} w_i x_i^* - c(w, y).$$

Since x^* is not necessarily optimal at w_i, then $z(w) \geq 0$ and $z(w^*) = 0$. Therefore,

$$\frac{\partial z(w^*)}{\partial w_i} = x_i^* - \frac{\partial c(w, y)}{\partial w_i} = 0.$$

It follows that

$$\frac{\partial c(w, y)}{\partial w_i} = x_i^* = h_i(w, y).$$

□

It can then be shown that the conditional factor of demand has the property of negativity

$$\frac{\partial h_i(w, y)}{\partial w_i} = \frac{\partial(\frac{\partial c(w,y)}{\partial w_i})}{\partial w_i} = \frac{\partial^2 c(w, y)}{\partial w_i^2} \leq 0,$$

because the cost function is concave in prices.

Furthermore, the cost function has the property of symmetry

$$\frac{\partial h_i(w, y)}{\partial w_j} = \frac{\partial(\frac{\partial c(w,y)}{\partial w_i})}{\partial w_j} = \frac{\partial^2 c(w, y)}{\partial w_i \partial w_j} = \frac{\partial(\frac{\partial c(w,y)}{\partial w_j})}{\partial w_i} = \frac{\partial h_j(w, y)}{\partial w_i}.$$

Example: Given two inputs, x_1 and x_2, that are substitutes in the production of output, y, the cost minimization problem is

$$\min_{x_1,x_2} w_1x_1 + w_2x_2 \text{ s. t. } x_1 + x_2 = y,$$

a Lagrangian could be written

$$\mathscr{L}(x_1, x_2, \lambda) = w_1x_1 + w_2x_2 - \lambda(x_1 + x_2 - y).$$

The first–order conditions are

$$\begin{aligned} w_1 - \lambda^* &= 0 \\ w_2 - \lambda^* &= 0. \end{aligned}$$

It follows that

$$w_1 = w_2.$$

Notice that the isocost curve is

$$x_2 = \frac{c}{w_2} - \frac{w_1}{w_2}x_1.$$

In this circumstance,

$$h_1(w_1, w_2, y) = \begin{cases} 0 & \text{if } w_1 > w_2, \\ k \in [0, y] & \text{if } w_1 = w_2, \\ y & \text{if } w_1 < w_2. \end{cases} \qquad h_2(w_1, w_2, y) = \begin{cases} y & \text{if } w_1 > w_2, \\ k \in [0, y] & \text{if } w_1 = w_2, \\ 0 & \text{if } w_1 < w_2. \end{cases}$$

Theorem: The Kuhn–Tucker Conditions
Given a minimization problem with an inequality constraint,

$$\min_x \phi(x) \text{ s. t. } g(x) \geq 0,$$

and $x_i \geq 0$ for all $i = 1, \ldots, n$, a Lagrangian can be written

$$\mathscr{L}(x, \lambda) = \phi(x) - \lambda(g(x)).$$

The first–order conditions are

$$\frac{\partial \mathscr{L}}{\partial x_i} = \frac{\partial \phi(x^*)}{\partial x_i} - \lambda \frac{\partial g(x^*)}{\partial x_i} \geq 0,$$

where

$$\begin{aligned} \frac{\partial \mathscr{L}}{\partial x_i} x_i^* &= 0 \\ x_i^* &\geq 0. \end{aligned}$$

and

$$\frac{\partial \mathscr{L}}{\partial \lambda} = -g(x^*) \leq 0,$$

where

$$\begin{aligned} \frac{\partial \mathscr{L}}{\partial \lambda} \lambda^* &= 0 \\ \lambda^* &\geq 0. \end{aligned}$$

The first–order conditions must bind or else the optimal quantity is a corner solution (i.e. its optimal choice is 0).

Example: (Continued)
Let
$$\frac{\partial \mathscr{L}}{\partial x_1} = w_1 - \lambda^* A \geq 0,$$
where
$$\frac{\partial \mathscr{L}}{\partial x_1} x_1^* = 0$$
$$x_1^* \geq 0,$$
$$\frac{\partial \mathscr{L}}{\partial x_2} = w_2 - \lambda^* A \geq 0,$$
where
$$\frac{\partial \mathscr{L}}{\partial x_2} x_2^* = 0$$
$$x_2^* \geq 0.$$
and
$$\frac{\partial \mathscr{L}}{\partial \lambda} = -x_1^* - x_2^* + y = 0.$$
The cases are

(1) $$\frac{w_1}{w_2} \geq \frac{A}{B} \Leftrightarrow x_1^* = 0 \text{ and } x_2^* > 0,$$
(2) $$\frac{w_1}{w_2} \leq \frac{A}{B} \Leftrightarrow x_1^* > 0 \text{ and } x_2^* = 0,$$
(3) $$\frac{w_1}{w_2} = \frac{A}{B} \Leftrightarrow x_1^* > 0 \text{ and } x_2^* > 0.$$

Example: Cost minimization with CES Production
Assume that there are two inputs, z_1 and z_2, of a constant elasticity of substitution (CES) production function
$$f(z) = (z_1^\rho + z_2^\rho)^{\frac{1}{\rho}},$$
for output, q, where $0 < \rho \leq 1$. The cost minimization problem is
$$\min_z wz \quad \text{s. t. } q \leq f(z).$$
The Lagrangian can be written
$$\mathscr{L} = wz + \lambda(q - f(z)).$$
The first–order conditions are
$$\frac{\partial \mathscr{L}}{\partial z_1} = w_1 - \lambda(z_1^\rho + z_2^\rho)^{\frac{1-\rho}{\rho}} z_1^{\rho-1} \geq 0,$$
where
$$\frac{\partial \mathscr{L}}{\partial z_1} z_1^* = 0$$
$$z_1^* \geq 0,$$

$$\frac{\partial \mathscr{L}}{\partial z_2} = w_2 - \lambda(z_1^\rho + z_2^\rho)^{\frac{1-\rho}{\rho}} z_2^{\rho-1} \geq 0,$$

where

$$\frac{\partial \mathscr{L}}{\partial z_2} z_2^* = 0$$
$$z_2^* \geq 0,$$

and

$$\frac{\partial \mathscr{L}}{\partial \lambda} = q - f(z^*) \leq 0,$$

where

$$\frac{\partial \mathscr{L}}{\partial \lambda} \lambda^* = 0$$
$$\lambda^* \geq 0.$$

If $z_2^* > 0$, then $\lim_{z_1 \to 0} \frac{\partial f(z)}{\partial z_1} = +\infty$, therefore $z_1^* > 0$.

1.2. Neoclassical Consumption Theory

Consider a consumer who faces a consumption set $q \in X \subset \mathbb{R}^n$.

1.2.1. Preference Relations

The consumer's preferences for $x, y \in X$ are as follows.

- If $x \succeq y$, then x is at least as good as y.
- If $x \succeq y$ and $y \not\succeq x$, then $x \succ y$ and x is strongly preferred to y.

Axiom: Completeness

$$\text{For all } x, y \in X, \text{ then } x \succeq y \text{ or } y \succeq x. \tag{A.1}$$

Axiom: Transitivity

$$\text{Let } x, y, z \in X. \text{ If } x \succeq y \text{ and } y \succeq z, \text{ then } x \succeq z. \tag{A.2}$$

Axiom: Continuity

$$\text{For all } q \in X, \text{ the sets } \{x \in X | x \succeq q\} \text{ and } \{y \in Y | y \preceq q\} \text{ are closed.} \tag{A.3}$$

Axiom: Strict Convexity

$$\text{If } q_1 \succeq q_2 \text{ and } q_2 \succeq q_0, \text{ then } \theta q_1 + (1-\theta)q_2 \succ q_0, \text{ where } \theta \in (0,1). \tag{A.4}$$

Definition: A function $U : X \to \mathbb{R}$ is a utility function that represents preference relation, $\succeq$, if for all $x, y \in X$, then

$$x \succeq y \Leftrightarrow U(x) \geq U(y).$$

Theorem: If a utility function, $U : X \to \mathbb{R}$, represents preference relation, $\succeq$, then the preference relationship must satisfy completeness and transitivity.

Proof. By the definition of a utility function, $x \succeq y \Leftrightarrow U(x) \geq U(y)$. Since $U(\cdot)$ is real–valued, then it must be that for any $q_1, q_2 \in X$, then

$$U(q_1) \geq U(q_2) \text{ or } U(q_2) \geq U(q_1).$$

It follows that either $q_1 \succeq q_2$ or $q_2 \succeq q_1$. Therefore the utility function satisfies the completeness axiom (A.1).

Now, let $q_1, q_2, q_3 \in X$ and $q_1 \succeq q_2$ and $q_2 \succeq q_3$. From the definition of a utility function, then

$$U(q_1) \geq U(q_2) \quad \text{and} \quad U(q_2) \geq U(q_3),$$

and because the utility function is real–valued

$$U(q_1) \geq U(q_3).$$

It follows from the definition of a utility function that

$$q_1 \succeq q_3.$$

Thus, given $q_1 \succeq q_2$ and $q_2 \succeq q_3$, then $q_1 \succeq q_3$. Therefore, the utility function satisfies the transitivity axiom (A.2). □

Theorem: Ordinality
If a preference relation satisfies completeness, (A.1), transitivity, (A.2), and continuity, (A.3), then there exists a utility function, $U : X \to \mathbb{R}$, that represents this preference relation. Moreover, any positive monotonic transformation of $U(f : \mathbb{R} \to \mathbb{R})$ also represents the same preferences

$$\text{if } \forall x \succeq y \Leftrightarrow f\big(U(x)\big) \geq f\big(U(y)\big).$$

1.2.2. Utility Maximization

Suppose that a consumer wishes to maximize utility, u, subject to their income, m. The optimization problem is

$$\max_{q_1,\dots,q_n} u(q_1,\dots,q_n) \ \text{ s. t. } \sum_{i=1}^{n} p_i q_i = m.$$

A Lagrangian can be written

$$\mathscr{L} = u(q_1,\dots,q_n) - \lambda\left(\sum_{i=1}^{n} p_i q_i - m\right).$$

Assuming that $\lambda^*, q_i^* > 0$ for all $i = 1,\dots,n$, then the first–order conditions are

$$\frac{\partial \mathscr{L}}{\partial \lambda} = -\sum_{i=1}^{n} p_i q_i + m = 0$$
$$\frac{\partial \mathscr{L}}{\partial q_i} = \frac{\partial u(q^*)}{\partial q_i} - \lambda^* p_i = 0,$$

for all $i = 1,\dots,n$. The result is the Marshallian demand functions, $q_1^*,\dots,q_n^*$, and λ^*.

Theorem: If the second–order conditions are satisfied, and if preferences are strictly convex (A.4), then the utility function is strictly quasi–concave

$$U(\theta q_1 + (1-\theta)q_2) > \min\{U(q_1), U(q_2)\}, \text{ for all } \theta \in (0,1).$$

A strictly quasi–concave utility function implies strict convexity of preferences, this guarantees unique q_i^* for all $i = 1,\dots,n$.

Proof. Suppose that, contrary to the proposition, that there are two demands that satisfy the budget constraint and maximize utility, q^* and q'^*. Since the utility function is quasi–concave, then it must be

$$U(\theta q^* + (1-\theta)q'^*) > \min\{U(q^*), U(q'^*)\}.$$

Then from

$$\sum_{i=1}^{n} p_i q_i^* = m \qquad\qquad \sum_{i=1}^{n} p_i q_i'^* = m$$
$$\theta \sum_{i=1}^{n} p_i q_i^* = \theta m, \qquad\qquad (1-\theta)\sum_{i=1}^{n} p_i q_i'^* = (1-\theta)m,$$

it follows that

$$\sum_{i=1}^{n} p_i \theta q_i^* + \sum_{i=1}^{n} p_i (1-\theta) q_i'^* = \theta m + (1-\theta)m$$
$$\sum_{i=1}^{n} p_i(\theta q_i^* + (1-\theta)q'^*) = m.$$

Therefore, there is a bundle, $(\theta q^* + (1-\theta)q'^*)$, that satisfies the budget constraint and is strictly better than q^* and $q'*$. This contradicts the original supposition. Therefore, there must be a unique solution. □

Definition: Strict Quasi-Concavity
Say that $h(x)$ is strictly quasi–concave if, for all x_1 and x_2,

$$h(\theta x_1 + (1-\theta)x_2) > \min\{h(x_1), h(x_2)\}.$$

Definition: Marshallian Demand Function

The consumer's optimal consumption choice, given prices, p, and income, m, is the Marshallian demand

$$g(p, m) \equiv q^* = \underset{q:pq\leq m}{\operatorname{argmax}}\ u(q).$$

Example: Marshallian Demand with Cobb–Douglas Utility

Let the consumer's utility function be a Cobb–Douglas specification, $u(q) = q_1^\alpha q_2^\beta$. The arguments of the utility maximization are the Marshallian demand functions. The optimization problem is

$$\max_{q_1,q_2} q_1^\alpha q_2^\beta \quad \text{s. t. } p_1q_1 + p_2q_2 = m.$$

A Lagrangian can be written

$$\mathscr{L}(q, \lambda) = q_1^\alpha q_2^\beta - \lambda(p_1q_1 + p_2q_2 - m).$$

The first–order conditions are

$$\begin{aligned}
\frac{\partial \mathscr{L}}{\partial \lambda} &= -p_1q_1^* - p_2q_2^* + m = 0, \\
\frac{\partial \mathscr{L}}{\partial q_1} &= \alpha {q_1^*}^{\alpha-1} {q_2^*}^{\beta} - \lambda p_1 \quad = 0, \\
\frac{\partial \mathscr{L}}{\partial q_2} &= \beta {q_1^*}^{\alpha} {q_2^*}^{\beta-1} - \lambda p_2 \quad = 0.
\end{aligned}$$

It follows that

$$\begin{aligned}
\frac{\alpha}{\beta}\frac{q_2^*}{q_1^*} &= \frac{p_1}{p_2} \\
q_2^* &= \frac{p_1}{p_2}\frac{\beta}{\alpha}q_1^*
\end{aligned}$$

Then, substituting q_2^* into the budget constraint and solving yields

$$\begin{aligned}
p_1q_1^* + p_2\left(\frac{p_1}{p_2}\frac{\beta}{\alpha}q_1^*\right) &= m \\
p_1\left(q_1^* + \frac{\beta}{\alpha}q_1^*\right) &= m \\
q_1^*\left(1 + \frac{\beta}{\alpha}\right) &= \frac{m}{p_1} \\
q_1^*\left(\frac{\alpha+\beta}{\alpha}\right) &= \frac{m}{p_1} \\
q_1^* &= \frac{m}{p_1}\frac{\alpha}{\alpha+\beta}.
\end{aligned}$$

Then q_2^* can be solved for

$$\begin{aligned}
q_2^* &= \frac{p_1}{p_2}\frac{\beta}{\alpha}\frac{m}{p_1}\frac{\alpha}{\alpha+\beta} \\
q_2^* &= \frac{m}{p_2}\frac{\beta}{\alpha+\beta}.
\end{aligned}$$

Thus, the Marhsallian demand functions are

$$\begin{aligned}
g_1(p, m) = q_1^* &= \frac{m}{p_1}\frac{\alpha}{\alpha+\beta} \\
g_2(p, m) = q_2^* &= \frac{m}{p_2}\frac{\beta}{\alpha+\beta}.
\end{aligned}$$

Definition: Indirect Utility Function
The consumer's utility function as a function of the Marshallian demand is known as the indirect utility function

$$V(p, m) \equiv U(g(p, m)) = u(g_1(p, m), \ldots, u(g_n(p, m)) = u(q_1^*, \ldots, q_n^*).$$

Properties:

- Non–increasing in prices;
$$\text{if } p_i^1 \geq p_i^2, \text{ then } V(p^1, m) \leq V(p^2, m).$$
- Non–decreasing in income;
$$\text{if } m^1 \geq m^2, \text{ then } V(p, m^1) \geq V(p, m^1).$$
- Homogeneous of degree 0;
$$V(\theta p, \theta m) = V(p, m), \text{ for all } \theta > 0.$$
- Quasi-convex in prices;
$$V(\theta p^1 + (1-\theta)p^2, m) < \max\{V(p^1, m), V(p^2, m)\}.$$
- Satisfies Roy's Identity;
$$g_i(p, m) = -\frac{\partial V(p, m)/\partial p_i}{\partial V(p, m)/\partial m}.$$

Example: (Continued)
In the above example with Cobb–Douglas utility, the indirect utility function is

$$v(p_1, p_2, m) = \left(\frac{m}{p_1}\frac{\alpha}{\alpha+\beta}\right)^{\alpha}\left(\frac{m}{p_2}\frac{\beta}{\alpha+\beta}\right)^{\beta}$$
$$v(p_1, p_2, m) = \frac{m^{\alpha+\beta}}{p_1^{\alpha}p_2^{\beta}}\frac{\alpha^{\alpha}\beta^{\beta}}{(\alpha+\beta)^{\alpha+\beta}}.$$

1.2.3. Expenditure Minimization

Suppose that a consumer wishes to minimize expenditure, e, subject to their utility, u. The optimization problem is

$$\min_{q_1,\ldots,q_n} \sum_{i=1}^{n} p_i q_i \quad \text{s. t. } u(q_1, \ldots, q_n) = u.$$

A Lagrangian can be written

$$\mathscr{L}(q, \lambda) = \sum_{i=1}^{n} p_i q_i - \lambda[u(q_1, \ldots, q_n) - u].$$

Assuming that $\hat{\lambda}, \hat{q}_i > 0$ for all $i = 1, \ldots, n$, then the first–order conditions are

$$\frac{\partial \mathscr{L}}{\partial \lambda} = -u(\hat{q}) + u = 0$$
$$\frac{\partial \mathscr{L}}{\partial q_i} = p_i - \lambda\frac{\partial u(\hat{q})}{\partial q_i} = 0,$$

for all $i = 1, \ldots, n$. The result is the Hicksian demand functions, $\hat{q}_1, \ldots, \hat{q}_n$, and $\hat{\lambda}$.

Definition: Hicksian Demand Function
A Hicksian (compensated) demand function is the consumer's expenditure minimizing choice given prices, p, and utility level, u,

$$h_i(p, u) \equiv \hat{q}_i.$$

The Hicksian demand is then

$$h(p, u) \equiv \underset{q:u(q)\geq u}{\text{argmin}}\ pq.$$

Definition: Expenditure Function
The expenditure function yields the minimal expenditure given prices, p, and a desired level of utility, u,

$$e(p, u) \equiv \sum_{i=1}^{n} p_i h_i(p, u).$$

Properties:

- Non–decreasing in prices;

$$\text{if } p_i^1 \geq p_i^2, \text{ then } e(p^1, u) \geq e(p^2, u).$$

- Non–decreasing in utility;

$$\text{if } u^1 \geq u^2, \text{ then } e(p, u^1) \geq e(p, u^2).$$

- Homogeneous of degree 1;

$$e(\theta p, u) = \theta e(p, e), \text{ for all } \theta > 0.$$

- Concave in prices;

$$e(\theta p^1 + (1-\theta)p^2, u) > \min\{e(p^1, u), V(p^2, u)\}.$$

- Satisfies Shepard's Lemma

$$\frac{\partial e(p, u)}{\partial p_i} = h_i(p, u).$$

Definition: The Slutsky Equation
The relationship between the Hicksian and Marshallian demand functions is

$$h(p^\circ, u^\circ) = g_i\big(p^\circ, e(p^\circ, u^\circ)\big).$$

It follows that

$$\frac{\partial h_i(p^\circ, u^\circ)}{\partial p_j} = \frac{\partial g_i(p^\circ, e(p^\circ, u^\circ))}{\partial p_j} + \frac{\partial g_i(p^\circ, e(p^\circ, u^\circ))}{\partial m}\frac{\partial e(p^\circ, u^\circ)}{\partial p_j},$$

and from Shepard's Lemma

$$\frac{\partial e(p^\circ, u^\circ)}{\partial p_j} = h_j(p^\circ, u^\circ) = q_j^\circ.$$

Rearranging the terms provides the Slutsky Equation

$$\frac{\partial g_i(p^\circ, e(p^\circ, u^\circ))}{\partial p_j} = \underbrace{\frac{\partial h_i(p^\circ, u^\circ)}{\partial p_j}}_{\text{The Substitution Effect}} - \underbrace{\frac{\partial g_i(p^\circ, e(p^\circ, u^\circ))}{\partial m} q_j^\circ}_{\text{The Income Effect}}.$$

Also, notice that given

$$V(p^\circ, m^\circ) = V(p^\circ, e(p^\circ, u^\circ)) = u^*,$$

then

$$\frac{\partial V(p^\circ, m^\circ)}{\partial p_i} + \frac{\partial V(p^\circ, m^\circ)}{\partial m}\frac{\partial e}{\partial p_i} = 0.$$

In turn, using Shepard's Lemma, the expression can be written as Roy's Identity

$$g_i(p^\circ, m^\circ) = -\frac{\partial V(p^\circ, m^\circ)/\partial p_i}{\partial V(p^\circ, m^\circ)/\partial m}.$$

Take note the relationship between utility maximization and expenditure minimization!

$\max_q u(q)$ s. t. $pq = m$ vs. $\min_q pq$ s. t. $u(q) = u$

$\downarrow$ $\downarrow$

$g(p, m) = q^*$ $\xleftrightarrow{\text{The Slutsky Equation}}$ $h(p, u) = \hat{q}$

$\downarrow$ $\downarrow$

Roy's Identity $\uparrow$ Shepard's Lemma $\uparrow$

$V(p, m) \equiv u(g(p, m))$ $\xleftrightarrow{\textit{Inversion}}$ $e(p, u) \equiv ph(p, u)$

You can use inversion to transform the indirect utility function to the expenditure function, by noting that

$$V(p^\circ, m^\circ) = u^\circ$$
$$e(p^\circ, u^\circ) = m^\circ.$$

Example: (Continued)
Given Cobb–Douglas utility, the indirect utility function is

$$v(p_1, p_2, m) = \frac{m^{\alpha+\beta}}{p_1^\alpha p_2^\beta}\frac{\alpha^\alpha \beta^\beta}{(\alpha+\beta)^{\alpha+\beta}}.$$

By inversion,

$$u = \frac{e(p_1, p_2, u)^{\alpha+\beta}}{p_1^\alpha p_2^\beta}\frac{\alpha^\alpha \beta^\beta}{(\alpha+\beta)^{\alpha+\beta}}$$

the expenditure function is

$$e(p_1, p_2, u) = u^{\frac{1}{\alpha+\beta}} p_1^{\frac{\alpha}{\alpha+\beta}} p_2^{\frac{\beta}{\alpha+\beta}} (\alpha+\beta)(\alpha^{-\frac{\alpha}{\alpha+\beta}})(\beta^{-\frac{\beta}{\alpha+\beta}}).$$

Theorem: The Composite Commodity Theorem (Hicks,Leontiff)
If a group of prices move in parallel, then they can be treated as a composite commodity.

Proof. Suppose there are prices, p_1, p_2, and p_3, such that $p_2 = \theta p_2^\circ$ and $p_3 = \theta p_3^\circ$ for all $\theta > 0$ and some p_2°, p_3°. If $\theta p_2^\circ q_2$ and $\theta p_3^\circ q_3$ are treated as a single commodity, then

$$e(p_1, p_2, p_3, u) \Rightarrow e^*(p_1, \theta, u) = e(p_1, \theta p_2^\circ, \theta p_3^\circ, u).$$

By the derivative with respect to θ, and by Shepard's Lemma, then

$$\frac{\partial e(p_1, \theta, u)}{\partial \theta} = \frac{\partial e(p_1, \theta, u)}{\partial p_2} p_2^\circ + \frac{\partial e(p_1, \theta, u)}{\partial p_3} p_3^\circ = q_2 p_2^\circ + q_3 p_3^\circ.$$

□

Example: Indirect Utility and Expenditure Functions with CRRA Utility
Let a consumer have CRRA utility, $u(q) = \frac{c^{1-\theta}}{1-\theta}$, where $\theta = 2$. The utility maximization problem is given by

$$\max_{q_1,q_2} u(q) = -\frac{1}{q_1} - \frac{1}{q_2} \quad \text{s. t. } pq \leq m.$$

A Lagrangian can be written

$$\mathscr{L} = -\frac{1}{q_1} - \frac{1}{q_2} + \lambda(m - p_1 q_1 - p_2 q_2).$$

The first–order conditions are

$$\begin{aligned}
\frac{\partial \mathscr{L}}{\partial \lambda} &= m - p_1 q_1 - p_2 q_2 \geq 0, \\
\frac{\partial \mathscr{L}}{\partial q_1} &= \frac{1}{q_1^2} - \lambda p_1 \leq 0, \\
\frac{\partial \mathscr{L}}{\partial q_2} &= \frac{1}{q_2^2} - \lambda p_2 \leq 0.
\end{aligned}$$

You can show that marginal utility, MU_i goes to infinity for $i = 1, 2$. It follows that

$$\frac{1}{q_1^2} = \lambda p_1 \text{ and } \frac{1}{q_2^2} = \lambda p_2.$$

Then, solve for q_1

$$\begin{aligned}
p_1 q_1^2 &= p_2 q_2^2 \\
q_1 &= \left(\frac{p_2}{p_1}\right)^{\frac{1}{2}} q_2,
\end{aligned}$$

and substitute q_1 into the budget constraint

$$m - p_2 q_2 - p_1 \left(\left(\frac{p_2}{p_1}\right)^{\frac{1}{2}} q_2 \right) = 0.$$

The result is the optimal choice

$$\begin{aligned}
q_1^* &= m[p_1 + (p_1 p_2)^{\frac{1}{2}}]^{-1} \\
q_2^* &= m[p_2 + (p_1 p_2)^{\frac{1}{2}}]^{-1}.
\end{aligned}$$

It follows that the indirect utility function is

$$V(p, m) = u(q^*) = -\frac{p_1 + (p_1 p_2)^{\frac{1}{2}}}{m} - \frac{p_2 + (p_1 p_2)^{\frac{1}{2}}}{m}.$$

The expenditure function can then be found by inversion

$$\begin{aligned}
u &= -\frac{p_1 + (p_1 p_2)^{\frac{1}{2}}}{e(p, u)} - \frac{p_2 + (p_1 p_2)^{\frac{1}{2}}}{e(p, u)} \\
u &= -\frac{1}{e(p, u)} \left[p + 1 + (p_1 p_2)^{\frac{1}{2}} + p_2 + (p_1 p_2)^{\frac{1}{2}} \right] \\
e(p, u) &= -\frac{p_1 + p_2 + 2(p_1 p_2)^{\frac{1}{2}}}{u}.
\end{aligned}$$

1.2.4. The Labor Supply Decision

Consider a consumer that must choose how much to consume, c, labor, l, and leisure, q_0. The consumer earns a wage–rate, w, for each unit of labor, l, and receives non–labor income, μ. There are a total of T possible units of time. Thus,

$$q_0 = T - l.$$

The budget constraint is

$$\sum_{i=1}^{n} p_i q_i = wl + \mu,$$

or

$$\begin{aligned}
\sum_{i=1}^{n} p_i q_i &= w(T - q_0) + \mu \\
\sum_{i=1}^{n} p_i q_i &= wT - wq_0 + \mu \\
\sum_{i=1}^{n} p_i q_i + wq_0 &= wT + \mu.
\end{aligned}$$

Notice that the constraint incorporates all potential income, $wT + \mu$, and prices leisure, q_0, with the consumer's opportunity cost, w. The profit maximization problem of the consumer is

$$\max_{q_0,\dots,q_n} u(q_0, \dots, q_n) \ \text{ s. t. } \sum_{i=1}^{n} p_i q_i + wq_0 = wT + \mu.$$

The Marshallian demand functions can be found

$$g_i(p, w, x) = g_i(p, w, wT + \mu).$$

Notice from the Slutsky Equation that the income effect now includes a re-evaluation of time effect. That is, if wage, w, increases then the consumer will consume more off all goods, including leisure, q_0, but faces a higher opportunity cost of leisure. Therefore there will be the ordinary income effect, leisure, q_0 will increase and labor, l, will decrease, and there will be a negative re-evaluation of time income effect, leisure, q_o, will decrease and labor, l, will increase. Thus, the change in time spent laboring with respect to an increase in the wage–rate is ambiguous. In fact, this leads to the phenomena of a backward–bending labor supply curve. Consumers will tend to labor more from an increase in wage at low wage–rates, and will tend to labor less from an increase in wage at high wage–rates.

1.2.5. Decisions Under Risk

Consider a consumer who faces a a decision under objective risk.

Reading: Amos Tversky and Daniel Kahneman, “The Framing of Decisions and the Psychology of Choice,” Science, 1981.

Definition: Lottery
A set of possible outcomes $C = (c_1, \dots, c_n)$ with probability distribution $P = (p_1, \dots, p_n)$, where $0 \le p_i \le 1$ for $i = 1, \dots, n$ and $\sum_{i=1}^{n} p_i = 1$, is a lottery. Denoted

$$L(p_1, \dots, p_n) \in \mathcal{L}.$$

Properties:

- Completeness
$$\text{For all } L_1, L_2 \in \mathcal{L}, \text{ then } L_1 \succeq L_2 \text{ or } L_2 \succeq L_1.$$
- Transitivity
$$\text{Let } L_1, L_2, L_3 \in \mathcal{L}. \text{ If } L_1 \succeq L_2 \text{ and } L_2 \succeq L_3, \text{ then } L_1 \succeq L_3.$$
- Continuity

  For all $L \in \mathcal{L}$ and $\alpha \in [0, 1]$, the following sets must be closed;
$$\{\alpha L + (1 - \alpha)L' \succeq L''\}$$
$$\{L'' \succeq \alpha L + (1 - \alpha)L'\}.$$
- Independence Axiom

  For all $L_1, L_2, L_3 \in \mathcal{L}$ and $\alpha \in [0, 1]$, then
$$L_1 \succeq L_2 \Leftrightarrow \alpha L_1 + (1 - \alpha)L_3 \succeq \alpha L_2 + (1 - \alpha)L_3.$$

Theorem: Under completeness, transitivity, and continuity, then there is a utility function, $U : \mathcal{L} \to \mathbb{R}$, such that

$$L_1 \succeq L_2 \Leftrightarrow U(L_1) \ge U(L_2).$$

Definition: Expected Utility
The expected utility property is that utility is linear in probabilities

$$U(L) = \sum_{i=1}^{n} p_i U(c_i).$$

Under the Independence Axiom. any utility function is closed and any positive affine transformation

$$v(x) = a + bu(x),$$

where $b > 0$, represents the same preferences.

Definition: Risk Aversion
A consumer is risk averse when their utility, $u(L)$, is concave and the expected value of a lottery, $\mathbb{E}[L]$, is weakly preferred to the non–deterministic lottery, L,

$$L \preceq \mathbb{E}[L] \Leftrightarrow u(L) \leq u(\mathbb{E}[L]).$$

Definition: Risk Seeking
A consumer is risk seeking when their utility, $u(L)$, is convex and a non–deterministic lottery, L, is weakly preferred to the expected value the lottery, $\mathbb{E}[L]$,

$$L \succeq \mathbb{E}[L] \Leftrightarrow u(L) \geq u(\mathbb{E}[L]).$$

Definition: Risk Neutrality
A consumer is risk neutral when their utility, $u(L)$, is linear and they are indifferent between a non–deterministic lottery, L, and the expected value of the lottery, $\mathbb{E}[L]$,

$$L \sim \mathbb{E}[L] \Leftrightarrow u(L) = u(\mathbb{E}[L]).$$

Definition: The Arrow–Pratt Measure of Absolute Risk Aversion
The Arrow–Pratt Measure of Absolute Risk Aversion (ARA) is

$$\text{ARA} = -\frac{U''(\cdot)}{U'(\cdot)}.$$

Notice that under a monotonic affine transformation

$$\begin{aligned} v(\cdot) &= a + bu(\cdot) \\ v'(\cdot) &= \frac{\partial v(\cdot)}{\partial x} = bu'(\cdot) \\ v''(\cdot) &= \frac{\partial^2 v(\cdot)}{\partial x^2} = bu''(\cdot), \end{aligned}$$

the ARA measures will remain the same

$$-\frac{v''(\cdot)}{v'(\cdot)} = -\frac{u''(\cdot)}{u'(\cdot)}.$$

Definition: Relative Risk Aversion
The measure of relative risk aversion (RRA) is given by

$$\text{RRS} = -\frac{xU''(x)}{U'(x)}.$$

A caveat is that when income increases, then RRA decreases.

Example: Negative Exponential Utility
Let a consumer's utility be given by

$$u(x) = -\mathrm{e}^{-rx},$$

where $r > 0$. It follows from

$$u'(x) = r\mathrm{e}^{-rx}$$
$$u''(x) = -r^2\mathrm{e}^{-rx},$$

that the ARA measure is

$$-\frac{u''(x)}{u'(x)} = -\frac{-r^2\mathrm{e}^{-rx}}{r\mathrm{e}^{-rx}} = r,$$

where r is the relative risk aversion factor.

Example: An Insurance Decision
Consider a consumer with wealth, W, who faces a loss, $L > 0$, with probability $p \in [0, 1]$. The consumer's expected utility is

$$\mathbb{E}[u] = pu(W - L) + (1 - p)u(W).$$

The consumer may purchase insurance coverage, q, at a premium, πq, where $\pi > 0$. The consumer's expected utility maximization problem is then

$$\max_q \mathbb{E}[u(q)] = pu(W - L - \pi q + q) + (1 - p)u(W - \pi q).$$

The first–order condition is

$$\mathbb{E}[u'(q)] = pu'(W - L - \pi q^* + q^*)(1 - \pi) + (1 - p)u'(W - \pi q^*)(-\pi) = 0.$$

It follows that

$$(1 - \pi)pu'(W - L - \pi q^* + q^*)(1 - \pi) = (1 - p)u'(W - \pi q^*)\pi$$
$$\frac{u'(W - L - \pi q^* + q^*)}{u'(W - \pi q^*)} = \frac{1 - p}{p}\frac{\pi}{1 - \pi}.$$

Suppose that there is actuarially fair insurance: $\pi = p$. Then

$$u'(W - L - \pi q^* + q^*) = u'(W - \pi q^*).$$

From the strict convexity of the utility function

$$W - L - \pi q^* + q^* = W - \pi q^*,$$

and there is full insurance coverage: $q^* = L$.
Now, suppose that the premium exceeds the price $\pi > p$. Then

$$\frac{\pi}{1 - \pi}\frac{1 - p}{p} > 1$$
$$\frac{u'(W - L - \pi q^* + q^*)}{u'(W - \pi q^*)} > 1$$
$$u'(W - L - \pi q^* + q^*) > u'(W - \pi q^*).$$

From the strict convexity of the utility function

$$W - L - \pi q^* + q^* < W - \pi q^*,$$

and the is *less* than full insurance coverage: $q^* < L$.
However, under risk neutrality

$$u'(W - L - \pi q^* + q^*) = u'(W - \pi q^*) = k > 0,$$

where k is some constant. The first–order condition is then

$$\mathbb{E}[u'(q)] = \frac{k}{k} = \frac{\pi}{1 - \pi}\frac{1 - p}{p} = 0$$
$$\mathbb{E}[u'(q)] = pk(1 - \pi) - (1 - p)\pi k$$
$$\mathbb{E}[u'(q)] = k[p(1 - \pi) - (1 - p)\pi].$$

Note that

$$\mathbb{E}[u'(q)] = k[p(1 - \pi) - (1 - p)\pi] < k[\pi(1 - \pi) - (1 - \pi)\pi) = 0,$$

and because the marginal utility is negative

$$u'(q) < 0,$$

then there is *no* insurance coverage: $q^* = 0$.

Example: An Investment Decision with Risky Assets
A consumer with wealth, W, may invest a portion of her wealth into two assets.

- A riskless asset with return r.
- A risky asset with possible returns $\theta_1, \ldots, \theta_n$ with probabilities $p_1, \ldots, p_n$.

Let $\alpha \in [0,1]$ be the percent the consumer invests in the risky asset. The consumer's expected utility is then

$$\mathbb{E}[u(\alpha)] = \sum_{i=1}^{n} p_i u(\alpha(1+\theta_i)W + (1-\alpha)(1+r)W)).$$
$$\mathbb{E}[u(\alpha)] = \sum_{i=1}^{n} p_i u(\alpha(1+\theta_i - 1 - r)W + (1+r)W).$$
$$\mathbb{E}[u(\alpha)] = \sum_{i=1}^{n} p_i u(\alpha(\theta_i - r)W + (1+r)W).$$

Her marginal utility is

$$\mathbb{E}[u'(\alpha)] = \sum_{i=1}^{n} p_i u'(\alpha(\theta_i - r)W + (1+r)W)(\theta_i - r)W.$$

Assume that $r > \theta_i$ for all $i = 1, \ldots, n$. Then

$$\mathbb{E}[u'(\alpha)] < 0,$$

for all α and there is *no* investment in the risky asset: $\alpha^* = 0$.

If the consumer is risk neutral (i.e. $u'(\alpha) = k$ for all α), then

$$\mathbb{E}[u'(\alpha)] = \sum_{i=1}^{n} p_i k W(\theta_i - r)$$
$$\mathbb{E}[u'(\alpha)] = kw \sum_{i=1}^{n} p_i(\theta_i - r)$$
$$\mathbb{E}[u'(\alpha)] = kw \left[\sum_{i=1}^{n} p_i \theta_i - \sum_{i=1}^{n} p_i r \right]$$
$$\mathbb{E}[u'(\alpha)] = kw(\mathbb{E}[\theta_i] - r).$$

Her investment decision is then as follows.

- If $\mathbb{E}[\theta] < r$, then $\alpha^* = 0$.
- If $\mathbb{E}[\theta] = r$, then $\alpha^* \in [0,1]$.
- If $\mathbb{E}[\theta] > r$, then $\alpha^* = 1$.

If the consumer is risk averse, then her investment decision is then as follows.

- If $\mathbb{E}[\theta] < r$, then $\alpha^* = 0$.
- If $\mathbb{E}[\theta] = r$, then $\alpha^* = 0$.
- If $\mathbb{E}[\theta] > r$, then it depends.

1.3. Firms, Markets, and Transactions

1.3.1. Partial Equilibrium

1.4. Welfare Economics

1.4.1. Public Goods and Externalities

Definition: An externality is a benefeit or cost of an activity that is not enjoyed or borne by the agent that engages in the activity.

Definition: A negative externality is a cost of an activity that is not enjoyed or borne by the agent that engages in the activity.

Definition: A positive externality is a benefit of an activity that is not enjoyed or borne by the agent that engages in the activity.

Theorem: The Coase Theorem
If property rights are well–defined and there are zero transaction costs, bargaining will lead to an efficient outcome regardless of the initial assignment of property rights.

Definition: Private goods are rivalrous in consumption and excludable.

Definition: Public goods are non–rivalrous in consumption and non–excludable.

Definition: Club goods are non–rivalrous in consumption and excludable.

Definition: Common goods are rivalrous in consumption and non–excludable.

1.5. Introduction to Game Theory

Definition: Game Theory

- "A group of agents is said to be engaged in a game whenever the fate of an agent in the group depends not only on his own actions, but also on the actions of the rest of the agents in the group." (Binmore and Dasgupta, 1986).
- "Game theory concerns the behavior of decision makers (players) whose decisions affect each other. As in noninteractive (one–person) decision theory, the analysis is from a rational, rather than a psychological or sociological viewpoint." (Aumann, 1988).
- "Game theory can be defined as the study of mathematical models of conflict and cooperation between intelligent rational decision makers." (Myerson, 1986).

There are many types of games.

Definition: Cooperative games are games where binding agreements, commitments, promises, and threats are possible.

Definition: Non–cooperative games are games where binding agreements are not possible even if pre–play communication is possible.

Quarter II

"Human interactions, stimulated as they are by disequilibrium, never achieve balance. In even the most favorable transaction, one party—whether he realizes it or not—must always come out the worse."

– Jack Vance, *Rhialto the Marvellous*

This quarter provides a rigorous introduction to game theory and an exploration of game–theoretic applications including markets with asymmetric information, bargaining and incentive contracts.

2.1. Game Theory

Definition: Game Theory

"Briefly put, game and economic theory are concerned with the interactive behavior of Homo rationalis—rational man. . . [An] important function of game theory is the classification of interactive decision situations."

– Robert J. Aumann, 1985

"Game theory can be defined as the study of mathematical models of conflict and cooperation between intelligent rational decision-makers. Game theory provides general mathematical techniques for analyzing situations in which two or more individuals make decisions that will influence one another's welfare."

– Roger B. Myerson, 1991

"Game theory is a mathematical method for analyzing strategic interaction."

– Nobel Prize Citation, 1994

Game theory is the study of interactive decision–making or strategic interaction. It consists of a mathematical language for representing strategic interactions and a set of tools for predicting the outcome of a strategic interaction.

Definition: A prediction of behavior in a game is called a solution.

Definition: A method for making such a prediction is called a solution concept.

Example: An Electoral Competition
There are two political parties, A and B. Each party independently chooses a policy between 0 and 1 (e.g. a tax rate). A party's objective is to maximize its vote share. Each citizen has an ideal policy which is uniformly distributed between 0 and 1. Each citizen votes for the party whose policy is closest to their ideal point. What is the outcome of this game?

Example: Splitting the Bill
Five friends go out to dinner and agree in advance to split the bill. There are two items on the menu: chicken and lobster. Each person can choose one item only. They each value the chicken at \$12 and the lobster at \$20. The price is \$10 for the chicken and \$25 for the lobster. What does each person choose?

These are examples of static games with complete information.

Definition: In a static game, players move simultaneously, without knowledge of the other players' moves.

Definition: In games with complete information, players know the structure of the game and the payoff functions of the other players.

2.2. Strategic–Form Games and Mixed Strategies

Static games can be represented using a strategic–form representation (or normal–form representation). When there is complete information, a strategic–form representation has three ingredients; players, strategies, payoffs. Note that payoffs depend not only on the agent's strategy, but on the strategies of other players as well.

2.2.1. Strategic–Form Games

Definition: Strategic–Form Game
A strategic–form game (or normal–form game) consists of

- Players; A set of agents, $N = \{1, \ldots, n\}$, with typical element $i \in N$.
- Strategies; A nonempty set of strategies, S_i with typical element $s_i \in S_i$, for each $i \in N$. A strategy is a complete plan of action specifying what a player will do at every point at which she may be called upon to play.
- Payoffs; A payoff function, $u_i : S \to \mathbb{R}$ for each player i, where $S = \prod_{i=1}^{N} S_i$.

Anything with these three features can be written as a strategic–form game

$$\mathcal{G} = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle.$$

Definition: Strategy Profile
A collection of strategies, $s \in S = \prod_{i=1}^{N} S_i$, is called a strategy profile.

Games with two players and a finite number of strategies can be represented by a payoff matrix.

Example: The Prisoners' Dilemma
"Two suspects are arrested for a crime, and interviewed separately. If they both keep quiet (they cooperate with each other) they go to prison for a year. If one suspect supplies incriminating evidence (defects) then that one is freed, and the other one is imprisoned for nine years. If both defect then they are imprisoned for six years. Their preferences are solely contingent on any jail term they individually serve."

- The players are the two suspects, $N = \{1, 2\}$.
- The strategy set for player 1 is $S_1 = \{C, D\}$, and for player 2 is $S_2 = \{C, D\}$.
- The payoffs can be represented in a payoff matrix.

	C	D
C	−1, −1	−9, 0
D	0, −9	−6, −6

2.2.2. Dominance

Let $s_{-i} \equiv \{s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n\} \in S_{-i}$ be a list of strategies for all players except i.

Definition: Dominating Strategy
Strategy $s_i \in S$ strictly dominates strategy $s' \neq s_i \in S_i$ for player $i \in N$ if

$$u_i(s_i, s{-}i) > u_i(s_i', s_{-i}) \text{ for all } s_{-i} \in S_{-i} = \prod_{j=1}^{N} s_j,$$

where $j \neq i$. If the inequality is weak, $\geq$, then it is a weakly dominant strategy.

Definition: Dominated–Strategy
Strategy $s' \in S_i$ is strictly dominated if there is a $s_i \in S_i$ that strictly dominates it.

One would expect that a player would never choose a strictly dominated strategy, since she can always do better by choosing a different strategy, no matter what her opponents chooses.

Definition: Dominant–Strategy
A strategy $s_i \in S_i$ is strictly dominant for $i \in N$ if it strictly dominates all $s_i' \neq s_i \in S_i$.

If a player has a strictly dominant strategy, it is always better than her other strategies, no matter what the other players do.

Example: The Prisoners' Dilemma (Continued)
The Prisoner's Dilemma provides a simple example of a game with a strictly dominant strategy.

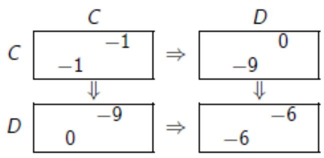

Here, D is a strictly dominant strategy for each player, and C is a strictly dominated strategy for each player. Therefore, $\{D, D\}$ is a dominant–strategy equilibrium.

Definition: Dominant–Strategy Equilibrium
A strategy $s_i^* \in S_i$ is a dominant–strategy equilibrium if

$$u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i}),$$

for all players $i \in N$ and for all strategy profiles $(s_i, s_{-i}) = s \in S$.

2.2.3. Iterated Deletion of Strictly Dominated Strategies

Consider the following game.

	L	M	R
T	4 3	2 7	0 4
B	5 5	5 −1	−4 −2

For the column player, M is strictly better than R. The game simplifies.

	L	M
T	4 3	2 7
B	5 5	5 −1

Now, for the row player, B is better than T.

	L	M
B	5 5	5 −1

$\longrightarrow$

	L
B	5 5

Thus, L beats M, leaving $\{B, L\}$ the only strategy that survives the iterated deletion of strictly dominated strategies. Notice that there are higher–order beliefs. If column player is rational, then she will not play R. If row player is rational and knows that column player is rational, then she will not play T. If column player is rational, knows that row player is rational, and knows that row player knows that column player is rational, won't play M.

The iterated deletion of strictly dominated strategies general algorithm is as follows.

- Step 1. If a strategy is strictly dominated for one player, delete it for that player and analyze the reduced game.
- Step 2. After one deletion, a strategy may become strictly dominated that was not originally so. Repeat the process until no more deletions are possible.

The remaining set of strategies for each player is nonempty and does not depend on the order of deletion, except if iteratively deleting weakly dominant strategies.

Example: The Battle of the Sexes
"Two students need to meet up to discuss their love for economics. They can meet in either the pub or the cafe. One likes coffee, and prefers the cafe. The other enjoys a pint of beer, and prefers the pub. They would both rather meet (wherever it may be) than miss each other."

- The players are the first student, (row) and the second (column).
- Row chooses $x \in \{\text{Cafe}, \text{Pub}\}$, and column chooses $y \in \{\text{Cafe}, \text{Pub}\}$.
- The payoffs can be represented in a strategic-form matrix.

	Cafe	Pub
Cafe	3 4	1 1
Pub	0 0	4 3

Notice that neither strategy is strictly (or weakly) dominated, and iterated deletion of dominated strategies does not rule out anything.

2.2.4. Nash Equilibrium

Definition: Nash Equilibrium
A Nash equilibrium is a strategy profile $s^* \in S$ such that for each $i \in N$

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*),$$

for all $s_i \in S_i$. At s^*, no i will regret playing s_i^*, that is, given all the other players' actions, i can not do better with any other strategy. Hence, a Nash equilibrium is a strategy profile from which no player has a profitable unilateral deviation.

Example: The Battle of the Sexes (Continued)

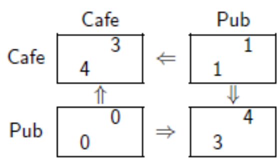

Note: The arrows represent what each player would do given the choice of their opponent.

Both $\{\text{Cafe}, \text{Cafe}\}$ and $\{\text{Pub}, \text{Pub}\}$ are Nash equilibria. Neither player has an incentive to deviate from their strategy given the strategy of their opponent.

Example: Consider the game from earlier.

	L	M	R
T	$4, 3$	$\underline{6}, \underline{7}$	$\underline{0}, 4$
B	$\underline{5}, \underline{5}$	$5, -1$	$-4, -2$

Note: Nash equilibria are strategy profiles where every payoff is underlined.

This game suggests an alternative (equivalent) definition for Nash equilibrium involving best replies, where a Nash equilibrium is a collection of mutual best replies.

Definition: Best–Reply
The best–reply correspondence for player $i \in N$ is a set–valued function B_i such that

$$B_i(s_{-i}) = \{s_i \in S_i | u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}), \forall s_i' \in S_i\}.$$

So that $B_i(s_{-i}) \subseteq S_i$ "tells" player i what to do when the other players play s_{-i}.

Definition: Nash Equilibrium (Alternative)
The strategy $s^* \in S$ is a Nash equilibrium if and only if

$$s_i^* \in B_i(s_{-i}^*),$$

for all $i \in N$. That is, a Nash equilibrium is a strategy profile of mutual best replies. Each player picks a best reply to the combination of strategies chosen by the other players.

Example: The Stag Hunt
"Two hunters simultaneously choose to hunt for rabbits, or to hunt for a stag. Successfully capturing a stag requires coordination, but there will be lots of meat. Anyone can catch a rabbit, but there will be less meat, especially when both are hunting rabbits (all that noise…)"

- The players are hunters 1 and 2 (row and column respectively).
- Each player can choose either Rabbit, R, or Stag, S.
- The payoffs can be represented in a strategic–form matrix.

	R	S
R	$\underline{3}, \underline{3}$	$4, 0$
S	$0, 4$	$\underline{5}, \underline{5}$

There are Nash equilibria at $\{R, R\}$ and $\{S, S\}$. The latter outcome is Pareto optimal, but it is not certain that it will be played.

Example: The Hawk–Dove Game
This is the classic biological game, where two players may either fight over a resource (Hawk) or yield (Dove). A Hawk beats a Dove, gaining the resource, of value v, with 0 for the Dove. Two Doves split the payoff v equally. Two Hawks have equal chance of winning the fight. The loser pays a cost c, where $v < c$. Given that $v = 4$ and $c = 6$, the game can be represented in a strategic–form game as follows.

	Hawk	Dove
Hawk	$-1, -1$	$\underline{4}, \underline{0}$
Dove	$\underline{0}, \underline{4}$	$2, 2$

There are two asymmetric Nash equilibria at $\{\text{Hawk}, \text{Dove}\}$ and $\{\text{Dove}, \text{Hawk}\}$.

Nash equilibrium requires more than rationality, it requires equilibrium knowledge where each player knows the strategy played by every other player in equilibrium.

Example: Rock–Paper–Scissors
Given two players who can play rock, R, scissors, S, or paper, P, with payoffs of 1 for winning, -1 for losing, and 0 for a tie, then the game can be represented in a strategic–form matrix.

	R	S	P
R	0, 0	1, −1	−1, 1
S	−1, 1	0, 0	1, −1
P	1, −1	−1, 1	0, 0

There are no pure–strategy Nash equilibria.

"Would row ever play R? Yes, if row thought column was playing S. Is this a rational belief? Yes, if row believes column believes row will play P. Is this a rational belief? Yes, if row believes column believes row believes column will play R. Is this a rational belief? Yes if..." Eventually, this process will return to "...believes row will play R," then R is said to be rationalizable.

2.2.5. Rationalizability

Definition: Rationalizability
Strategy $s_i \in S_i$ is rationalizable in the game $\mathcal{G} = \langle N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N}\rangle$ if for all $j \in N$, there is a set $R_j \subseteq S_j$ such that $s_i \in R_i$ and every action $s_j \in R_j$ is a best reply to a belief $\mu_j(s_j)$ of player j whose support is a subset of $R_{-j} = \prod_{k=1}^{N} R_k$, where $k \neq j$.

Note that this is circular reasoning. A strategy is rationalizable if it is a best reply to a combination of opponents' strategies that are all rationalizable. However, the formal definition does make sense! In the above examples:

- D is rationalizable in the Prisoners' Dilemma
- Cafe and Pub are rationalizable in the Battle of the Sexes.
- R and S are rationalizable in the Stag Hunt
- Hawk and Dove are rationalizable in the Hawk–Dove game.
- R, S, and P, are rationalizable in the Rock–Paper–Scissors game.

Example: Reconsider the following game.

	L	M	R
T	4, 3	2, 7	0, 4
B	5, 5	5, −1	−4, −2

Note that M is not rationalizable. For column to play M, column must believe row will play T. For this belief to be rational, column must believe that row believes that column will play R. For this belief to be rational there must be a belief for column to which R is a best reply. There is not—R is strictly dominated. In fact, only B is rationalizable for row and L rationalizable for column.

Theorem: Let $R = \prod_{i=1}^{N} R_i$, where R_i is the set of rationalizable strategies for $i \in N$, Z is the set of Nash equilibria, and A is the set that survives iterated deletion of strictly dominated strategies. Then

$$Z \subseteq R \subseteq A \subseteq S.$$

Under the current definition of rationalizability $R = A$ if the mixed extension of the game is used.

2.2.6. Mixed Strategies

Theorem: The Nash Theorem
Every finite strategic–form game has at least one Nash equilibrium.

Example: The Fashion Game
"A fashion leader and a fashion follower simultaneously choose a style of dress, A or B. The follower wants to choose the same style of dress as the leader. The leader wants to choose the opposite."

	A	B
A	-1, $\underline{1}$	$\underline{1}$, -1
B	$\underline{1}$, -1	-1, $\underline{1}$

There is no "pure–strategy" Nash equilibrium. Neither rationalizability nor iteratively eliminating strictly dominated strategies help to reach a solution.

Note that players can also mix between strategies—such a mixture is called a mixed strategy. In some contexts the most salient or even the only Nash equilibria involve mixed strategies.

Definition: Mixed–Form Games
The mixed extension of a game $\mathcal{G} = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ is the game Γ, where

- $\Gamma = \langle N, \{\Delta(S_i)\}_{i \in N}, \{U_i\}_{i \in N} \rangle$,
- $\Delta(S_i)$ is the set of probability distributions over S_i, and $\Delta(S) = \prod_{i=1}^{N} \Delta(S_i)$,
- $U_i : \Delta(S) \to \mathcal{R}$ is a von–Neumann–Morgenstern expected utility function that assigns to each $\sigma \in \Delta(S)$ the expected value under u_i of the lottery over S induced by σ.

Suppose that player i plays mixed strategy $\sigma_i \in \Delta(S_i)$ in a finite game. Denote the probability that this places on pure strategy $s_i \in S_i$ as $\sigma_i(s_i)$. Then a player's payoff is

$$U_i(\sigma) = \sum_{s \in S} u_i(s) \prod_{j \in N} \sigma_j(s_j).$$

Define $\sigma_{-i} \in \Delta(S_{-i}) \equiv \prod_{j \neq i \in N} \Delta(S_j)$ as the mixed strategies of all players except i.

Example: Matching Pennies
For matching pennies, the players are $N = \{1, 2\}$. The pure strategies are $S_i = \{H, T\}$.

	H	T
H	-1, $\underline{1}$	$\underline{1}$, -1
T	$\underline{1}$, -1	-1, $\underline{1}$

The mixed extension has the same set of players and mixed strategies

$$\sigma_1 = (p, 1-p), \text{ where } 0 \le p \le 1,$$
$$\sigma_2 = (q, 1-q), \text{ where } 0 \le q \le 1.$$

The payoff for a player is given by

$$U_i(\sigma) = \sum_{s \in S} u_i(s) \prod_{j \in \{1,2\}} \sigma_j(s_j)$$
$$U_i(\sigma) = u_i(H, H)\sigma_1(H)\sigma_2(H) + \cdots + u_i(T, T)\sigma_1(T)\sigma_2(T).$$

Definition: Mixed Strategy Best–Reply
The best–reply correspondence of the mixed extension is

$$B_i(\sigma_{-i}) = \{\sigma_i | U_i(\sigma_i, \sigma_{-i}) \geq U_i(\sigma_i', \sigma_{-i}), \forall \sigma_i'\}.$$

Example: Matching Pennies (Continued)
If $\sigma_2 = (q, 1-q)$, then

$$U_1(H, q) = (1-q) - q = 1 - 2q$$
$$U_1(T, q) = q - (1-q) = 2q - 1.$$

So, player 1 plays H if $q < \frac{1}{2}$ (i.e. $p = 1$), and plays T if $q > \frac{1}{2}$ (i.e. $p = 0$). If $q = \frac{1}{2}$, then player 1 is indifferent and any p will suffice. If $\sigma_1 = (p, 1-p)$, then

$$U_2(H, p) = p - (1-p) = 2p - 1$$
$$U_2(T, p) = (1-p) - p = 1 - 2p.$$

So, player 2 plays T if $p < \frac{1}{2}$ (i.e. $q = 0$), and plays H if $p > \frac{1}{2}$ (i.e. $q = 1$). If $p = \frac{1}{2}$, then player 2 is indifferent and any q will suffice. The best–reply correspondence is as follows.

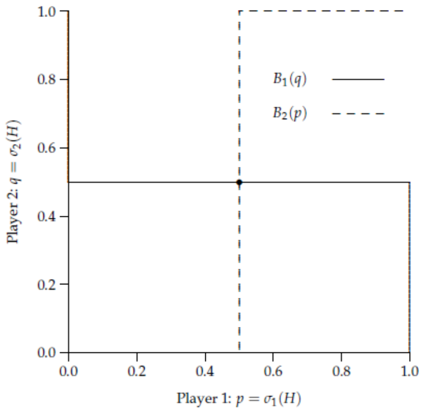

Theorem: A mixed-strategy Nash equilibrium of a game, $\mathcal{G}$, is a Nash equilibrium of its mixed extension, Γ.

Definition: Mixed–Strategy Nash Equilibrium
A mixed–strategy Nash equilibrium is a strategy profile $\sigma^* \in \Delta(S)$, such that

$$U_i(\sigma_i^*, \sigma_{-i}^*) \geq U_i(\sigma_i, \sigma_{-i}^*),$$

for all $\sigma_i \in \Delta(S_i)$ and $i \in N$.

Alternatively, $\sigma^* \in \Delta(S)$ is a Nash equilibrium if and only if

$$\sigma_i^* \in B_i(\sigma_{-i}^*),$$

for all $i \in N$.

Properties:

- For a mixed strategy, σ_i, to be a best reply to a given combination of opponents' strategies, σ_{-i}, every pure strategy in its support must also be a best reply to the opponents' strategies, σ_{-i}.
- All pure strategies in the support of the equilibrium strategy for a given player must yield the same payoff to that player.

Note that definitions for rationalizability and dominance can be extended to mixed strategies as well.

Example: Matching Pennies (Continued)
Continue consideration of the matching pennies game.
If a player mixes strategies, then they must be indifferent between the two pure strategies. The column player is indifferent only when

$$\begin{aligned} U_2(H,p) &= U_2(T,p) \\ 2p-1 &= 1-2p \\ p &= \frac{1}{2}. \end{aligned}$$

Similarly, the row player is indifferent only when

$$\begin{aligned} U_1(H,p) &= U_1(T,p) \\ 1-2q &= 2q-1 \\ q &= \frac{1}{2}. \end{aligned}$$

The mixed–strategy equilibrium is therefore $(\frac{1}{2}, \frac{1}{2})$. Notice that row player's equilibrium strategy is determined by column's payoffs and vice versa.

Example: The Battle of the Sexes Revisited
The players are student 1 (row) and student 2 (column). The strategies available to both players are Cafe and Pub. The row player chooses Cafe with probability $x \in [0, 1]$, and the column player chooses Cafe with probability $y \in [0, 1]$. The mixed extension of the game can be represented in a strategic form matrix.

	Cafe (y)	Pub ($1-y$)
Cafe (x)	4, 3	1, 1
Pub ($1-x$)	0, 0	3, 4

The expected payoffs are

$$U_1(\text{Cafe}, y) = 4y + (1-y) \qquad U_2(\text{Cafe}, x) = 3x$$
$$U_1(\text{Pub}, y) = 3(1-y), \qquad U_2(\text{Pub}, x) = x + 4(1-x).$$

Thus, the row player chooses Cafe, $x = 1$, whenever

$$\begin{aligned} U_1(\text{Cafe}, y) &> U_1(\text{Pub}, y) \\ 4y + (1-y) &> 3(1-y) \\ y &> \frac{1}{3}, \end{aligned}$$

and column player chooses Cafe, $y = 1$, whenever

$$\begin{aligned} U_2(\text{Cafe}, y) &> U_2(\text{Pub}, y) \\ 3x &> x + 4(1-x) \\ x &> \frac{2}{3}. \end{aligned}$$

The equilibria occur where $\sigma_i \in B_i(\sigma_{-i})$, for all $i = 1, 2$. There are two pure equilibria

$$\begin{aligned} x = y &= 1 \\ x = y &= 0, \end{aligned}$$

and there is one mixed equilibrium, $x = \frac{2}{3}$ and $y = \frac{1}{3}$.

You can note that, generically, there are an odd number of equilibria. Also, in mixed extensions of finite games, the best–reply correspondences will be continuous. Thus, the best–response functions must always intersect and there must exist a Nash equilibrium.

Example: Consider the following game.

	L	R
T	$1, 1$	$\underline{4}, \underline{2}$
B	$\underline{3}, \underline{3}$	$\underline{4}, 0$

By inspection, there are two pure-strategy Nash equilibria at (B, L) and $T, R)$. Suppose that player 1 (the row player) places probability p on T and probability $(1 - p)$ on B. Then, player 2's best reply is to play R if

$$\begin{aligned} U_2(L, p) &\leq U_2(R, p) \\ p + 3(1 - p) &\leq 2p \\ p &\geq \frac{3}{4}. \end{aligned}$$

If player 2 places probability q on L and $(1 - q)$ on R, then T is only ever a best reply to $q = 0$. Thus, T is weakly dominated by B. Nevertheless there is an equilibrium at (T, R). It follows that player 1's best-reply correspondence is

$$\begin{aligned} &B_1(q) = 0 \text{ if } q > 0, \\ &B_1(q) = p \text{ with } 0 \leq p \leq 1 \text{ if } q = 0, \end{aligned}$$

and player 2's best–reply correspondence is

$$\begin{aligned} &B_2(p) = 1 \text{ if } p < \frac{3}{4}, \\ &B_2(p) = 0 \text{ if } p > \frac{3}{4} \\ &B_2(p) = q \in [0, 1] \text{ if } p = \frac{3}{4}. \end{aligned}$$

The best–reply correspondances are as follows.

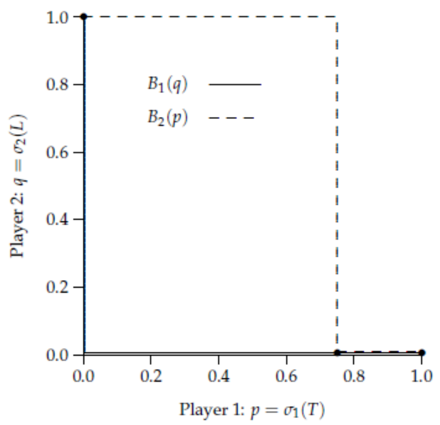

There is a continuum of mixed-strategy equilibria at $\frac{3}{4} \leq p \leq 1$, all with $q = 0$. As long as player 1 places high enough probability on T, then R is a best reply. If R is played, mixing is possible for player 1.

Example: Dominance and Mixed Strategies
In the following game, no strategy is strictly dominated by another pure strategy.

	L	M	R
T	4 10	3 0	1 3
B	0 0	2 10	10 3

Suppose that the column player plays L with probability $\frac{1}{2}$ and M with probability $\frac{1}{2}$ so that $\sigma_C = (\frac{1}{2}, \frac{1}{2}, 0)$. Then,

$$U_C(\sigma_C, \sigma_R) = 5,$$

regardless of what the row player plays for all σ_R. Thus, R is strictly dominated by a mixed strategy placing probability $\frac{1}{2}$ on L and M. Only one strategy profile survives iterated deletion of dominated strategies, $\{T, L\}$.

	L	M
T	4 10	3 0
B	0 0	2 10

$\longrightarrow$

	L	M
T	4 10	3 0

$\longrightarrow$

	L
T	4 10

Example: Never Best–Replies
Consider the above game. Below is a plot of plot of column's payoffs to each strategy. Irrespective of the value of the belief x, R is never a best–reply. In fact, a strategy is strictly dominated if and only if it is never a best–reply. Rational players would always play a best–reply given some beliefs.

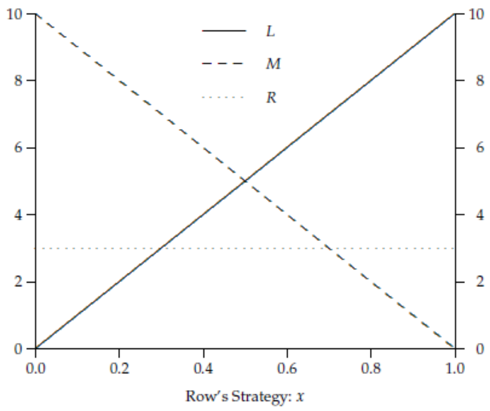

Example: A Return to Rationalizability
In the above game, only T and L are rationalizable. Consider R. It is not rationalizable as there is no belief such that R would be a best reply. Thus, B is not rationalizable, because the row player would need to believe R is to be played by the column player. Thus, M is not rationalizable, as this requires the column player to believe that the row player will play B.

If you restrict the game to pure strategies, then no strategies are strictly dominated. That is, the set of rationalizable strategies is a subset of the set of iterated deletion of dominated strategies survivors.

If you allow for mixed strategies, then R is strictly dominated by a 50 : 50 mix over L and M. Then B is strictly dominated by T, and then M is strictly dominated by L. So, now the set of rationalizable strategies is equal to the set of iterated deletion of dominated strategies survivors.

Example: Dominated Mixed Strategies
Consider the game with different payoffs.

	L	M	R
T	3 10	0 0	1 7
B	0 0	3 10	1 7

Note that neither the pure strategy L nor M is strictly dominated by R. The strategy which places probability $\frac{1}{2}$ on each of L and M earns 5. This is strictly dominated by R. The Nash equilibrium are

$$\{T, L\}, \{B, M\}, \{x \in [\tfrac{3}{10}, \tfrac{7}{10}], R\}.$$

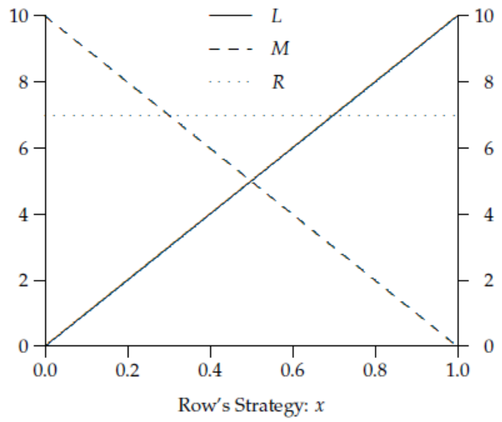

A mixed strategy with positive weight on a strictly dominated pure strategy is strictly dominated. But a mixed can be dominated by a pure even if all strategies in its support are undominated.

2.3. Static Oligopoly

2.3.1. Oligopoly versus Monopoly

Example: A Monopoly

Consider a single firm in an industry. It faces a demand curve $x(p)$, and so will choose p. The firm's profit maximization problem is

$$\max_p \pi(p) = px(p) - c\big(x(p)\big),$$

or alternatively

$$\max_q \pi(q) = p(q)q - c(q).$$

At an optimal quantity, $q^* > 0$, therefore, the first–order condition holds

$$p'(q^*)q^* + p(q^*) = c'(q^*).$$

The marginal revenue is

$$r'(q) = p'(q^*)q^* + p(q^*).$$

Thus, at the optimum marginal revenue equals marginal cost

$$r'(q^*) = c'(q^*).$$

Furthermore, note that demand is downward sloping, $p'(\cdot) < 0$. It follows from

$$p(q^*) = c'(q^*) - p'(q^*)q^*,$$

that under monopoly the equilibrium price exceeds the marginal cost of production

$$p(q^*) > c'(q^*).$$

Perfectly competitive prices are such that $p^{\mathrm{PC}} = c'(q^{\mathrm{PC}})$. Thus,

$$q^* < q^{\mathrm{PC}}$$
$$p^* > p^{\mathrm{PC}},$$

and there is a deadweight loss.

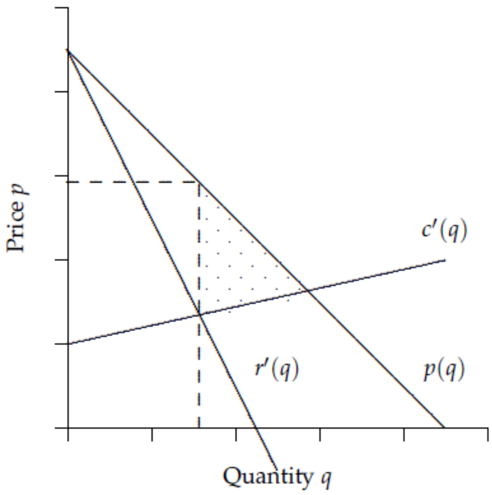

The monopoly case is essentially a decision problem. In oligopolistic industries, however, behavior is strategic where the strategy sets are continuous.

2.3.2. Cournot Competition

Two profit-maximizing firms simultaneously choose production quantities of a homogeneous good. Market price is decreasing in total quantity Q, with linear demand

$$p = a - bQ.$$

There are constant unit production costs of c for each firm.

- The players are the two firs, $i \in \{1, 2\}$.
- Player 1 chooses quantity $x \in [0, \infty)$ and player 2 chooses quantity $y \in [0, \infty)$.
- The payoffs are profits for player 1 and player 2 respectively

$$\pi_1 = x[a - b(x + y) - c],$$
$$\pi_2 = y[a - b(x + y) - c].$$

In order to find an equilibrium, first fix firm 2's strategy and calculate a best reply for firm 1, yielding a best–reply function. Then, fix firm 1's strategy and calculate a best reply for firm 2, yielding a second best–reply function. You can then solve the best–reply functions simultaneously to find a Nash equilibrium.

Fixing y, then the profits for player 1 are strictly concave in x and the first–order condition can be calculated as

$$\frac{\partial \pi_1}{\partial x} = [a - b(x + y) - c] - bx = 0$$
$$a - 2bx - by - c = 0.$$

Rearranging obtains

$$2bx = a - by - c,$$

that in turn implies

$$B_1(y) = \frac{(a - by - c)}{2b}.$$

Below is a plot of a reaction function for $a = b = 1$, and $c = 0$.

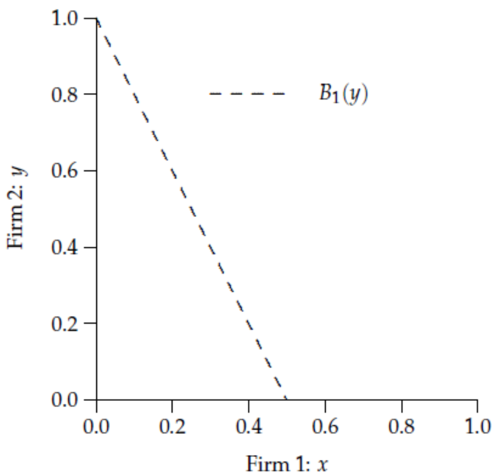

Note that the reaction function is downward sloping

$$\frac{\partial B_1(y)}{\partial y} = -\frac{1}{2} < 0.$$

Thus, the quantities, x and y, are strategic substitutes in this submodular game.

At a Nash equilibrium, players mutually best reply

$$x = B1(y),$$
$$y = B2(x).$$

So,

$$x = \frac{a - by - c}{2b} \quad \text{and} \quad y = \frac{a - bx - c}{2b}.$$

If you solve these two equations simultaneously, then the solution will be symmetric since the first-order conditions depend only on demand, Q. From symmetry, then

$$x = \frac{a - bx - c}{2b},$$

and solving yields

$$2bx = a - bx - c$$
$$3bx = a - x$$
$$x^* = \frac{a - c}{3b},$$

and by symmetry,

$$y^* = \frac{a - c}{3b}.$$

The reaction functions and equilibrium are plotted below.

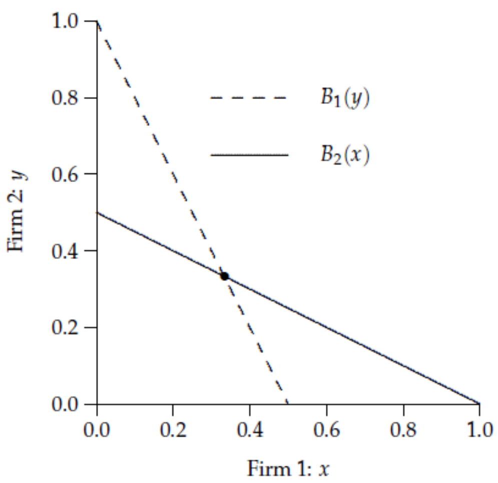

Example: Strictly Dominated Strategies in the Cournot Model
The Cournot–Nash equilibrium strategies are the only survivors from the iterated deletion of strictly dominated strategies. Assume that $a = b = 1$, and $c = 0$. Then

$$\pi_1 = x(1 - x - y)$$
$$\pi_2 = y(1 - x - y).$$

Consider the strategies $x, y \in (\frac{1}{2}, \infty)$, that are strictly dominated by $x, y = \frac{1}{2}$ with profit

$$\pi_1 = \tfrac{1}{2}(\tfrac{1}{2} - y)$$
$$\pi_2 = \tfrac{1}{2}(\tfrac{1}{2} - x).$$

Suppose to the contrary, for some y, then

$$x(1 - x - y) > \frac{(\frac{1}{2} - y)}{2} \Leftrightarrow (\tfrac{1}{2} - x)y > \tfrac{1}{4} - x + x^2.$$

Since $y \geq 0$, and $x \in (\frac{1}{2}, \infty)$, then the left–hand side is less than or equal to zero. The right-hand side is minimized at zero when $x = \frac{1}{2}$, and therefore is positive. This is a contradiction. The same is true for $y > \frac{1}{2}$.

Now consider the strategies $x \in [0, \frac{1}{4})$. These strategies are strictly dominated by $x = \frac{1}{4}$. The payoffs are $x(1 - x - y)$ and $\frac{1}{4}(\frac{3}{4} - y)$ respectively. Suppose again, to the contrary, that for some $y \leq \frac{1}{2}$, then

$$x(1 - x - y) > \frac{(\frac{3}{4} - y)}{4} \Leftrightarrow (\tfrac{1}{4} - x)y > \tfrac{3}{16} - x + x^2.$$

The inequality holds for some $y \leq \frac{1}{2}$ if and only if

$$(\tfrac{1}{4} - x)\tfrac{1}{2} > \tfrac{3}{16} - x + x^2 \Leftrightarrow 0 > \tfrac{1}{16} - \tfrac{1}{2}x + x^2,$$

which is a contradiction. This process continues until the only choice that remains is

$$x = y = \tfrac{1}{3}.$$

Example: Equilibrium in General Cournot Games
In general, let there be n firms. Firm i has a constant marginal cost, c_i, and inverse demand, $P(Q)$. The objective of the firm is to maximize profits. If demand is less than the firm's cost, $P(Q) < c_i$, then the firm does not operate, $q_i = 0$. Otherwise, the firm's maximization problem is

$$\max_{q_i} \pi_i = q_i[P(Q) - c_i].$$

The first order condition is

$$\frac{\partial \pi_i}{\partial q_i} = P(Q) - c_i + q_i P'(Q) = 0.$$

Solving yields

$$q_i^* = -\frac{P(Q) - c_i}{P'(Q)}.$$

Individual quantities are defined by industry supply Q. Thus, if $c_i = c$ for all i, then any equilibrium is symmetric.

Example: Cournot Competition with Asymmetric Cost
If the cost structures of the firms are asymmetric, $c_i \neq c$ for all i, then you can determine the equilibrium Q by summing the first–order conditions for all n firms (the total optimal production level) and dividing by the total demand, $P(Q)$. Then divide by the number of firms, n, to find a given firms production level

$$\frac{nP(Q) - \left[\sum_{i=1}^{n} c_i\right]}{P(Q)} + \frac{QP'(Q)}{P(Q)} = 0 \Leftrightarrow \frac{P(Q) - \frac{1}{n}\sum_{i=1}^{n} c_i}{P(Q)} = \frac{1}{nc}.$$

Hence outcome is determined by the industry–average of marginal cost. In games where there is a single state variable (here, Q), determining equilibria involves solving a single fixed-point equation.

Example: Cournot Competition versus Perfect Competition
Let marginal costs be constant and equal to c for every firm. In n–firm Cournot Competition then

$$\begin{aligned} P(Q) + q_i P'(Q) &= c \\ nP(Q) + P'(Q)\sum_{i=1}^{n} q_i &= nc \\ P(Q) + P'(Q)\frac{Q}{n} &= c. \end{aligned}$$

Since demand is downward sloping, $P'(\cdot) < 0$, then

$$P(Q) > c.$$

In perfect competition

$$P(Q^{\mathrm{PC}}) = c.$$

thus

$$Q < Q^{\mathrm{PC}}.$$

It can be concluded that competitive industries produce more, and at a lower price.

Example: Cournot Competition versus Monopoly
Now, suppose the monopoly optimal quantity is q^*. Suppose that $q^* > Q$. Take a particular firm i, and let firm i increase q_i so that the new industry quantity is q^*. Then joint profits must increase as they are maximized at q^* by definition. However, the aggregate quantity has risen, so price has fallen, and the other firms (who didn't alter their quantities) are worse off. As joint profits have risen, i must be better off and faces a profitable deviation. Therefore, $q^* \leq Q$. Note though, that $q^* \neq Q$ since the above equation can't be satisfied by the same Q at $n > 1$ and $n = 1$. Thus,

$$\begin{aligned} q^* &< Q, \\ P(q^*) &> P(Q), \end{aligned}$$

a monopoly produces less, and at a higher price than firms under Cournot cometition.

2.3.3. Bertrand Competition

Two firms selling identical products must simultaneously choose what price to charge. The firm that charges the lower price gains the entire market, but firms would rather charge high prices. A group of consumers will only buy if the price is less than $\overline{p}$. For simplicity, and without loss of generality, the marginal cost of production is zero, $c = 0$.

- The players are the two firms, $i \in N = \{1, 2\}$.
- The strategy of each player i is to set price $p_i \in [0, \infty]$.
- The payoff for each player is their profit

$$\pi_i = \begin{cases} p_i & \text{if } p_i < \min\{\overline{p}, p_j\}, \\ \frac{p_i}{2} & \text{if } p_i = p_j < \overline{p}, \\ 0 & \text{if } p_i \geq \overline{p} \text{ or } p_i > p_j. \end{cases}$$

There is a unique pure–strategy Nash equilibrium at $p_1 = p_2 = 0$.

Proof. If the lowest price were negative, then that firm would make a loss. If the lowest price were strictly positive, then opponent should undercut and steal the entire market. If one price is zero, e.g. $0 = p_i < p_j$, then firm i should raise its price. Hence the only possibility is $p_1 = p_2 = 0$, where there is no better reply. □

Notice that best-reply functions are not well–defined everywhere. Suppose, for example, that $0 < p_j < \overline{p}$. Then it is always a best–reply for player i to undercut player j, but she would like to do so by ε as small as possible without $\varepsilon = 0$. Mathematically, the set of feasible payoffs is open above, that is, cannot attain a maximum. The Bertrand specification is degenerate—owing to the discontinuity in payoffs.

2.3.4. The Hotelling Line

Two firms are located at either end of a unit interval $[0, 1]$. A unit mass of consumers (each with unit demand) is distributed uniformly on the interval. The firms charge p_i and p_j respectively for a homogeneous good produced with constant marginal cost, c. The cost of buying from firm i is

$$p_i + td,$$

where t is a unit transport cost and d is the distance from firm i. A particular consumer, $z \in [0, 1]$, will buy from i, who is positioned at 0, if

$$p_i + tz < p_j + t(1 - z).$$

The indifferent consumer, $\overline{z}$, satisfies

$$p_i + t\overline{z} = p_j + t(1 - \overline{z}),$$
$$\overline{z} = \frac{t + p_j - p_i}{2t}.$$

Assuming $\overline{z} \in [0, 1]$, then firm i's demand is given by

$$q_i = \overline{z} = \frac{1}{2} + \frac{p_j - p_i}{2t}.$$

Note that

$$q_i = \begin{cases} 1 & \text{if } \overline{z} > 1, \\ 0 & \text{if } \overline{z} < 0. \end{cases}$$

2.3.5. Differentiated Products

An alternative interpretation to Bertrand Competition is where there are two firms selling differentiated products who simultaneously choose prices. Total market size is a single unit mass. Each consumer is willing to pay a large amount to obtain a product. They do not necessarily buy from the cheapest firm, however, if $p_j - p_i > t$, then firm i captures the whole market, $q_i = 1$ and $q_j = 0$, and *vice versa*. If $|p_j - p_i| \leq t$, then the firms split the market depending on the price difference

$$q_i = \frac{1}{2} + \frac{p_j - p_i}{2t}.$$

- The players are the two firms, $N = \{i, j\}$.
- The strategies for each player i is to set price $p_i \in [0, \infty)$.
- The payoff for player i is

$$\pi_i = \begin{cases} 0 & \text{if } p_i - p_j > t, \\ p_i - c & \text{if } p_j - p_i > t, \\ (p_i - c)\left(\frac{1}{2} + \frac{p_j - p_i}{2t}\right) & \text{otherwise.} \end{cases}$$

Note that profit is concave in prices. Therefore, you can differentiate the profit function to obtain the first–order condition

$$\frac{\partial \pi_i}{\partial p_i} = \frac{t + p_j - 2p_i + c}{2t} = 0.$$

Solving for p_i yields the best–reply function

$$B_i(p_j) = \frac{t + c + p_j}{2}.$$

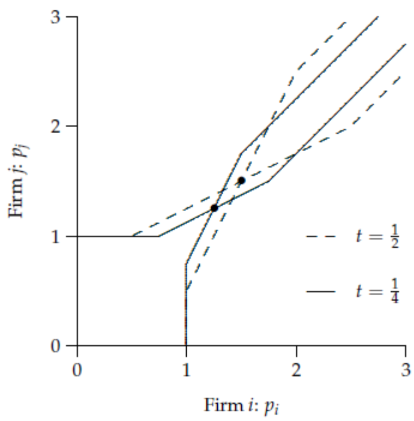

Note that the best–reply function is upward sloping $\frac{\partial B_i}{\partial p_j} > 0$. Thus, prices are strategic complements and the game is supermodular. Also, note that this solution only applies when $|p_j - p_i| \leq t$. In fact,

$$B_i(p_j) = \begin{cases} c & \text{if } p_j < c - t, \\ \frac{t + c + p_j}{2} & \text{if } c - t \leq p_j \leq 3t + c, \\ p_j - t & \text{if } 3t + c < p_j. \end{cases}$$

For an interior equilibrium

$$p_i = \frac{t + c + p_j}{2}.$$

Symmetry ensures that

$$p_i = p_j = p^*.$$

So,

$$p^* = \frac{t + c + p^*}{2}$$

$$p^* = t + c.$$

2.3.6. Submodular and Supermodular Games

Consider a game with two players, $i \in N = \{1, 2\}$, with strategies for player 1, who chooses $x \in X \subseteq \mathbb{R}$, and player 2, who chooses $y \in Y \subseteq \mathbb{R}$, and the payoffs, $u_1(x, y)$ and $u_1(x, y)$, are symmetric

$$u_1 = u_2 = u.$$

You can calculate the slope of a player's best–reply function

$$\begin{aligned} \frac{\partial u(x,y)}{\partial x} &= 0 \\ \frac{\partial^2 u(x,y)}{\partial x \partial y} + \frac{\partial^2 u(x,y)}{\partial x^2}\frac{\mathrm{d}x}{\mathrm{d}y} &= 0 \\ \frac{\mathrm{d}x}{\mathrm{d}y} &= -\frac{\partial^2 u(x,y)/\partial x \partial y}{\partial^2 u(x,y)/\partial x^2}. \end{aligned}$$

Note that from the second–order conditions that the denominator is negative. Therefore, the sign of the slope of the best–reply function is determined by numerator

$$\frac{\partial^2 u(x,y)}{\partial x^2} < 0 \Rightarrow \operatorname{sign}\left\{\frac{\mathrm{d}x}{\mathrm{d}y}\right\} = \operatorname{sign}\left\{\frac{\partial^2 u(x,y)}{\partial x \partial y}\right\}.$$

Definition: Supermodular Game (Informal)
The game $\mathcal{G} = \langle \{1, 2\}, \{X, Y\}, \{u, u\}\rangle$ is supermodular if

$$\frac{\partial^2 u(x,y)}{\partial x \partial y} > 0,$$

and so $B_i(y)$ is upward sloping, and X and Y are said to be strategic complements.

Definition: Submodular Game (Informal)
The game $\mathcal{G} = \langle \{1, 2\}, \{X, Y\}, \{u, u\}\rangle$ is submodular if

$$\frac{\partial^2 u(x,y)}{\partial x \partial y} < 0,$$

and so $B_i(y)$ is downward sloping, and X and Y are said to be strategic substitutes.

Example: Non–Monotonic Best–Replies in an Advertising Game
Two firms sell a product in a market of fixed size. Suppose that prices are fixed at 1, but that each firm must choose an advertising budget, denoted by x and y respectively. Advertising is costly, but firms want to obtain a high market share. Advertising is the sole determinant of market share, yielding sales of $\frac{x}{x+y}$ and $\frac{y}{x+y}$ respectively

- The players are the two firms, $N = \{1, 2\}$.
- The strategies are firm 1 chooses $x \in [0, \infty)$ and firm 2 chooses $y \in [0, \infty)$.
- The payoffs are the respective profits of the firms

$$\pi_1(x, y) = \frac{x}{x+y} - x \quad \text{and} \quad \pi_2(x, y) = \frac{y}{x+y} - y.$$

First, consider firm 1's maximization problem. The first–order condition implies that

$$\begin{aligned}
\frac{\partial \pi_1}{\partial x} = \frac{1}{x+y} - \frac{x}{(x+y)^2} - 1 &= 0 \\
\frac{1}{x+y} &= \frac{x}{(x+y)^2} + 1 \\
x + y &= x + (x+y)^2 \\
\sqrt{y} &= x + y, \\
x^* &= \sqrt{y} - y.
\end{aligned}$$

By symmetry the best response functions are

$$\begin{aligned}
B_1(y) &= \sqrt{y} - y \\
B_2(x) &= \sqrt{x} - x.
\end{aligned}$$

The best-reply functions for this game slope upward initially, then downward. The best–reply functions are non–monotonic. The game is neither supermodular nor submodular. Variables are both strategic complements and substitutes, depending on the region.

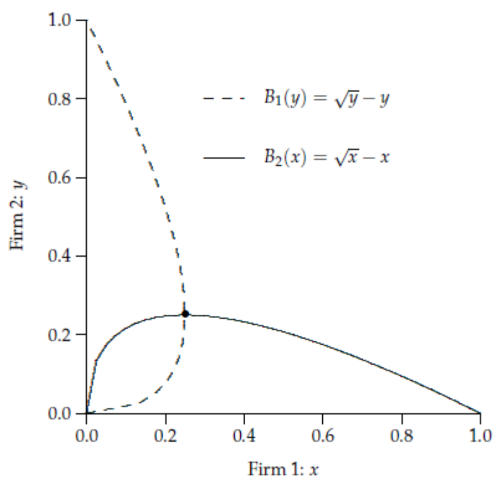

At Nash equilibrium, $x = B_1(y)$ and $y = B_2(x)$, hence

$$x^* = y^* = \tfrac{1}{4}.$$

Note that $B_1(y)$ and $B_2(x)$ only make sense for $y > 0$ and $x > 0$. If neither firm advertised, then $\pi_1(0,0) = \pi_2(0,0) = \frac{1}{2}$, is an equilibrium, although arguably 'unstable'. Whereas, $\pi_1(\frac{1}{4}, \frac{1}{4}) = \pi_2(\frac{1}{4}, \frac{1}{4}) = \frac{1}{4}$, is a stable Nash equilibrium.

Example: Mixed–Strategy Nash Equilibrium in an Investment Game
Two firms choose investment levels from the unit interval. The firm with the higher investment wins the market, which has unit value. If the same level is chosen, they split the market 50 : 50.

- The players are the two firms, $i \in N = \{1, 2\}$.
- The strategies are firm 1 chooses investment level $x \in X = [0, 1]$ and firm 2 chooses investment level $y \in Y = [0, 1]$.
- The mixed strategies for firm 1 and firm 2 are the distributions $F(x)$ and $G(x)$ on $[0, 1]$ respectively.
- The payoffs are the expected profit flows of the firms

$$\pi_1(x,y) = \begin{cases} 1-x & \text{if } x > y, \\ \frac{1}{2}-x & \text{if } x = y, \\ -x & \text{if } x < y, \end{cases} \qquad \text{and} \qquad \pi_2(x,y) = \begin{cases} -y & \text{if } x > y, \\ \frac{1}{2}-y & \text{if } x = y, \\ 1-y & \text{if } x < y. \end{cases}$$

Notice that there are no pure-strategy Nash equilibria. If $y < 1$, then player 1 would do better with $x = y + \varepsilon < 1$. If $x = y = 1$, then firm 1 would do better to choose $x = 0$. There is, however, a mixed–strategy Nash equilibrium.

Recall the indifference property of mixed equilibria. Argue that there is an equilibrium, such that in mixed equilibria a player must be indifferent across all the pure strategies they mix over. Hence all x in the support of player 1's strategy $F(\cdot)$ must yield a constant amount in expectation.

- The probability player 1 wins with x is $\Pr[x > y] = G(x)$.
- The probability player 1 loses with x is $\Pr[x < y] = 1 - G(x)$.
- The probability player 1 draws with x is $\Pr[x = y] = 0$.

For all x in the support of $F(\cdot)$, player 1's expected payoff is

$$\mathbb{E}(\pi_1) = -x[1 - G(x)] + [1 - x]G(x) = k \Leftrightarrow G(x) = x + k.$$

Suppose that $x = 0$ is in the support of $F(\cdot)$. Then $G(0) = 0$ since this is a CDF, and so $k = 0$. In this case, $G(x) = x$, and $\mathbb{E}(\pi_1) = 0$. Symmetry implies there is a Nash equilibrium where both players mix uniformly over $[0, 1]$. In fact, it can be shown that this is the only Nash equilibrium of this game[1].

Definition: Uniqueness
A mixed equilibrium strategy is unique if there are no atoms in the distribution, there are no gaps in the distribution, the distribution has full support on $[0, 1]$. There is only one such distribution.

[1] There are games that have no Nash equilibria at all (pure or mixed). This property is known as non–existence. Recall that in finite games there was always at least one (possibly mixed) equilibrium. However, the world is continuous and there are often infinite states.

2.4. Bayesian Games and Extensive–Form Games

2.4.1. Incomplete Information

Thus far, it has been assumed that the structure of the game is common knowledge. In particular, every player knows the other players' payoff functions, knows that the other players know her payoff function, knows that the other players know that she knows their payoff functions… *ad infinitum*. Is it reasonable that a player can predict with certainty how an opponent will best respond? Suppose that you face an opponent that can be one of a number of different types, each of which respond differently (e.g. an altruistic type, a spiteful type). When players do not know their opponent's type, you have a game of incomplete information. This creates a problem because a player must form beliefs about her opponent's type, about her opponent's beliefs about her type, about her opponent's beliefs about her beliefs about her opponent's types, etc. The solution proposed by Harsanyi is to assume that a player's type is a random draw from a distribution that is common knowledge. This reduces a game of incomplete information to a game of imperfect information. Such a game is commonly called a Bayesian Game.

2.4.2. Bayesian Games

Definition: Bayesian Game
A Bayesian game consists of the five following features.

- A finite set of players labeled $i \in N = \{1, \ldots, n\}$.
- For each $i \in N$, a set of types T_i, with typical member $t_i \in T_i$.
- For each $i \in N$, a mapping $s_i : T_i \to A_i$ from her types to her actions $a_i \in A_i$. The set of Bayesian strategies is $S_i = A_i \times T_i$.
- For each $i \in N$ and $t_i \in T_i$, there are beliefs, μ, with probability measure p_i over T_{-i}, written
$$\mu \equiv p_i(t_{-i}|t_i).$$
- For each $i \in N$, there are payoffs that takes the form of a von Neumann–Morgenstern utility function
$$u_i : A \times T \to \mathbb{R}$$

Anything with these five features can be written as a Bayesian game

$$\Gamma = \langle N, \{T_i\}_{i\in N}, \{A_i\}_{i\in N}, \{p_i\}_{i\in N}, \{u_i\}_{i\in N}\rangle,$$

where action profiles and type profiles are respectively defined as

$$a \in A \equiv \prod_{i=1}^{N} A_i \qquad \text{and} \qquad t \in T \equiv \prod_{i=1}^{N} T_i.$$

It is straightforward to extend the notion of a Bayesian strategy to mixed strategies. A strategy is then a mapping from a given player's type space to the set of probability distributions over their action space

$$s_i : T_i \to \Delta(A_i).$$

Definition: Bayesian–Nash Equilibrium
A Bayesian–Nash Equilibrium of a Bayesian game Γ is a combination of mutual best replies, in terms of Bayesian strategies, and a strategy profile, consisting of a Bayesian strategy for each player, from which no player has a profitable unilateral deviation[2].

[2] Equivalently, no player–type pair has a profitable deviation.

Example: Bayesian Battle of the Sexes

Two students simultaneously decide whether to go to a pub or a cafe. Student 1 prefers the cafe, while student 2 prefers the pub. Student 1 likes the second and wants to meet up with her. But the feeling may not be mutual. If student 2 likes student 1, she wants to meet. Otherwise, she prefers to avoid her. Suppose that student 2 likes student 1 with probability $\frac{1}{2}$ and hates her with probability 12.

- The players are student 1 and two types of student 2; 2_l and 2_h.
- The actions of each player can choose are Cafe and Pub.
- The payoffs are given in the below matrices, where each occurs with probability $\frac{1}{2}$.

	Cafe	Pub
Cafe	4, 3	1, 1
Pub	0, 0	3, 4

or

	Cafe	Pub
Cafe	4, 0	1, 4
Pub	0, 3	3, 1

Note that player 2 knows which matrix applies, where as player 1 does not and assigns probability $\frac{1}{2}$ to each.

This game can be written as a Bayesian game.

- The players are $N = \{1, 2\}$.
- The types for player 1 are $T_1 = \{1\}$ and for player 2 $t_2 \in T_2 = \{l, h\}$.
- For each player–type the actions available are $A_1 = A_2 = \{\text{Cafe}, \text{Pub}\}$. A Bayesian strategy for player i associates each one of her types $t_i \in T_i$ with an action $a_i \in A_i$.
- The beliefs of the players are

$$p_1(l|1) = p_1(h|1) = \tfrac{1}{2}$$
$$p_2(1|h) = p_2(1|h) = 1.$$

- The payoffs remain the same as described in the matrix above.

Notice that types are independent

$$p(t_{-i}|t_i) = p(t_{-i}).$$

Players own types do not reveal information about their opponents' types. This need not be the case (e.g. global games). Notice that types are private

$$u_i(a, t) = u_i(a, ti).$$

A player's payoffs depend upon only her own–type draws, and not directly upon the draws of opponents' types. The The Bayesian–Nash equilibrium strategy profiles are;

- Player 1 plays $s_1(1) = C$ and player 2 plays $s_2(t_2)$, where $s_2(l) = C$ and $s_2(h) = P$.
- Player 1 plays $s_1(1) = (\frac{1}{3}, \frac{2}{3})$ and player 2 plays $s_2(l) = P$ and $s_2(h) = (\frac{2}{3}, \frac{1}{3})$.
- Player 1 plays $s_1(1) = (\frac{2}{3}, \frac{1}{3})$ and player 2 plays $s_2(l) = (\frac{2}{3}, 13)$ and $s_2(h) = P$.

Example: Another Noisy Battle of the Sexes

Consider the following variation of the Battle of the Sexes. Player 1 (row) does not know what payoffs player 2 (column) receives from playing Pub. In fact, there is some 'noise' in player 2's payoff, δ. Likewise player 1's payoff to Cafe is perturbed by some 'noise', ε, from the perspective of player 2. Suppose that $\delta, \varepsilon \sim \mathcal{U}[0, a]$, where $a < 2$. The payoff matrix below illustrates this game.

	Cafe	Pub
Cafe	$4+\varepsilon$, 3	$1+\varepsilon$, $1+\delta$
Pub	0, 0	3, $4+\delta$

A Bayesian strategy for player 2 maps each of her types, δ, to an action, either Cafe or Pub. Consider if player 2 follows a 'cut–off' strategy such as to play Cafe if $\delta < \overline{\delta}$ and play Pub if $\delta \geq \overline{\delta}$. This is a pure strategy. The probability that player 2 plays Cafe, from player 1's perspective, is $\frac{\overline{\delta}}{a}$. If player 2 uses this strategy, then player 2 should play Cafe if

$$[4+\varepsilon]\tfrac{\overline{\delta}}{a} + [1+\varepsilon](1-\tfrac{\overline{\delta}}{a}) \geq 3(1-\tfrac{\overline{\delta}}{a})$$
$$\varepsilon \geq 2 - 6\tfrac{\overline{\delta}}{a}.$$

Thus, player 1's best reply to 2's cut–off strategy is itself a cutoff strategy defined by

$$\overline{\varepsilon} = 2 - 6\tfrac{\overline{\delta}}{a}.$$

If $\varepsilon \geq \overline{\varepsilon}$, then player 1 plays Cafe. If $\varepsilon < \overline{\varepsilon}$, then player 1 plays Pub. Given this strategy, an analogous argument yields

$$\overline{\delta} = 2 - 6\tfrac{\varepsilon}{a}.$$

This is a defining cut–off strategy for player 2, which is a best reply to a cut–off strategy by 1. Solving for $\overline{\delta}$ and $\overline{\varepsilon}$ results in

$$\overline{\delta} = \overline{\varepsilon} = \frac{2a}{6+a}.$$

Thus, there is a Bayesian–Nash equilibrium in cut–off strategies where player 1 plays Cafe if

$$\varepsilon \geq \frac{2a}{6+a},$$

and Pub otherwise, and player 2 plays Pub if

$$\delta \geq \frac{2a}{6+a},$$

and Cafe otherwise.

2.4.3. Purification

Definition: Purification
The probability distributions over strategies induced by the pure–strategy (Bayesian–Nash) equilibria of the perturbed game converge to the distribution of the (mixed Nash) equilibrium of the unperturbed game. This is Harsanyian purification.

The statement can be made precise, but the idea is simple.

Example: Purifying Bayesian Nash Equilibria in the Battle of the Sexes
Recall the perfect–information battle–of–the–sexes payoffs with $\varepsilon = \delta = 0$.

	Cafe	Pub
Cafe	4, 3	1, 1
Pub	0, 0	3, 4

Recall that the mixed Nash equilibrium involves player 1 playing Cafe with probability $\frac{2}{3}$ and Pub with probability $\frac{1}{3}$, and player 2 playing Cafe with probability $\frac{1}{3}$ and Pub with probability $\frac{2}{3}$. Notice that the probability with which player 1 plays Cafe in the Bayesian game is

$$\Pr\left[\varepsilon \geq \frac{2a}{6+a}\right] = 1 - \frac{2}{6+a}.$$

As $a \to 0$, then

$$1 - \frac{2}{6+a} \to \frac{2}{3}.$$

Thus, the distribution of types collapses to a point. The pure–strategy Bayesian–Nash equilibrium resembles the mixed-strategy Nash equilibrium of the unperturbed game. This process is called purification.

Notice there are other Bayesian-Nash equilibria of the perturbed game:

- Player 1 plays Cafe for all ε, and player 2 plays Cafe for all δ.
- Player 1 plays Pub for all ε, and player 2 plays Pub for all δ.

2.4.4. A First–Price Sealed Bid Auction

Two players simultaneously and independently submit sealed bids to an auctioneer who awards the object for sale to the highest bidder. The player who wins must pay the amount that she bid. In the event of a tie, the object is allocated to one player selected at random (with probability $\frac{1}{2}$), who must then pay her bid. Each player's valuation of the object v_i is an independent draw from $\mathcal{U}(0,1)$. Player 1 observes v_1, but not v_2, and *vice versa.*

- The players and $i \in \{1,2\}$.
- The strategies are $b_i(v_i)$ for each player, i, type, v_i, and bid, b_i.
- The payoff to player i of type v_i is

$$u_i(b, v_i) = \begin{cases} 0 & \text{if } b_i < b_j, \\ \frac{v_i - b_i}{2} & \text{if } b_i = b_j, \\ v_i - b_i & \text{if } b_i > b_j. \end{cases}$$

Suppose that player 2 uses a linear bidding strategy

$$b_2 = \alpha + \beta v_2.$$

Then player 1 wins if

$$b_1 > b_2 = \alpha + \beta v_2.$$

This occurs with probability

$$\begin{aligned} \Pr[b_1 > b_2] &= \Pr[b_1 > \alpha + \beta v_2] \\ \Pr[b_1 > b_2] &= \Pr[v_2 < \tfrac{b_1 - \alpha}{\beta}] \\ \Pr[b_1 > b_2] &= F[\tfrac{b_1 - \alpha}{\beta}] \\ \Pr[b_1 > b_2] &= \frac{b_1 - \alpha}{\beta}, \end{aligned}$$

as long as $\alpha \leq b_1 \leq \alpha + \beta$. Then player 1 receives a payoff of $v_1 - b_1$ if she wins and zero otherwise. Hence her expected payoff is

$$u_1(b, v_1) = (v_1 - b_1)\frac{b_1 - \alpha}{\beta}.$$

You can find player 1's optimal bid, b_1, for each one of her types, v_1, given player 2's linear bidding strategy. The first–order condition of player 1's utility maximization problem is

$$\frac{\partial u_1}{\partial b_1} = \frac{1}{\beta}(v_1 - b_1) - \frac{b_1 - \alpha}{\beta} = 0.$$

This implies that

$$\begin{aligned} \frac{2b_1}{\beta} &= \frac{v_1 + \alpha}{\beta} \\ b_1 &= \frac{v_1 + \alpha}{2} \\ b_1 &= \tfrac{1}{2}\alpha + \tfrac{1}{2}v_1. \end{aligned}$$

Hence, player 1's best reply to 2's linear strategy is itself a linear strategy.

You can rewrite

$$b_1 = \frac{1}{2}\alpha + \frac{1}{2}v_1$$
$$b_1 \equiv A + Bv_1,$$

where $A = \frac{\alpha}{2}$ and $B = \frac{1}{2}$.

You can also find player 2's best reply to player 1's linear strategy

$$b_2 = \frac{\alpha}{2} + \frac{1}{2}v_2$$
$$b_2 \equiv A + Bv_2,$$

For the linear bidding strategies to be mutual best replies, and thus constitute a Bayesian–Nash equilibrium, it must be that

$$\alpha = A = \frac{\alpha}{2} \Leftrightarrow \alpha = 0$$
$$\beta \equiv B = \frac{1}{2}.$$

Therefore, the unique BNE in linear bidding strategies is

$$b_1^*(v_1) = \frac{1}{2}v_1,$$
$$b_2^*(v_2) = \frac{1}{2}v_2.$$

Note that players bid half their valuation in equilibrium. Furthermore, these strategies meet the requirements that $a \leq b_i \leq \alpha + \beta$ for both players, $i = 1, 2$.

2.4.5. A Double Auction

A buyer has a valuation for a good $v_b \sim \mathcal{U}(0,1)$. A seller has valuation $v_s \sim \mathcal{U}(0,1)$. They each observe their own valuations, but not the other player's. They must each announce a price, p_b and p_s, simultaneously. If $p_b \geq p_s$ a sale takes place at a price half way between $\frac{pb+ps}{2}$. Otherwise there is no sale.

- The players are the buyer and the seller, $N = \{b, s\}$.
- The strategies are $p_i(v_i)$ for each player, i, type, v_i, and price p_i.
- The players both receive a payoff of 0 if $p_b < p_s$, and otherwise

$$U_b(p, v_b) = v_b - \frac{p_b + p_s}{2} \qquad \text{and} \qquad U_s(p, v_s) = \frac{p_b + p_s}{2} - v_s.$$

Suppose that the seller chooses

$$p_s(v_s) = \alpha + \beta v_s.$$

Then the buyer's expected payoff from an offer of p_b is

$$U_b(p, v_b) = \Pr[p_b \geq p_s] \left(v_b - \frac{p_b + \mathbb{E}[p_s | p_b \geq p_s]}{2} \right).$$

Note that

$$p_s = \alpha + \beta v_s \sim \mathcal{U}(\alpha, \alpha + \beta).$$

In the figure below on the left, the probability that $p_b \geq \alpha + \beta v_s$ is the shaded area, that is, $\frac{p_b - \alpha}{\beta}$. The expected value of $\alpha + \beta v_s$, given that $p_b \geq \alpha + \beta v_s$ is $\frac{\alpha + p_b}{2}$.

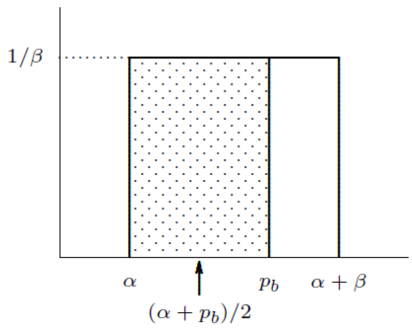

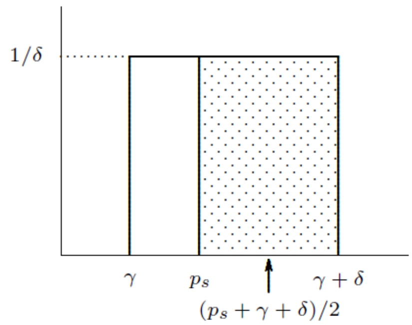

Hence,

$$U_b(p, v_b) = \frac{p_b - \alpha}{\beta} \left[v_b - \frac{1}{2} \left(p_b + \frac{\alpha + p_b}{2} \right) \right].$$

Differentiating this utility function with respect to p_b, and then setting it to zero yields

$$\frac{1}{\beta} \left[v_b - \frac{1}{2} \left(p_b + \frac{\alpha + p_b}{2} \right) \right] = \frac{3}{4} \left(\frac{p_b - \alpha}{\beta} \right.$$

$$p_b(v_b) = \tfrac{1}{3}\alpha + \tfrac{2}{3} v_b.$$

The seller's problem can be set up in an analogous way. Suppose that

$$p_b = \gamma + \delta v_b.$$

Then

$$U_s(p, v_s) = \Pr[p_b \geq p_s] \left(\frac{p_s + \mathbb{E}[p_b | p_b \geq p_s]}{2} - v_s \right).$$

In the figure above on the right, the probability term is the shaded area, given by $\frac{\gamma+\delta-p_s}{\delta}$. The conditional expectation term is $\frac{p_s+\gamma+\delta}{2}$. Differentiating this utility function with respect to p_s, and the setting it to zero yields

$$\frac{1}{\delta}\left[\frac{1}{2}\left(p_s+\frac{p_s+\gamma+\delta}{2}\right)-v_s\right]=\frac{3}{4}\left(\frac{\gamma+\delta-p_s}{\delta}\right)$$
$$p_s(v_s)=\tfrac{1}{3}(\gamma+\delta)+\tfrac{2}{3}v_s.$$

For these two strategies to be a Bayesian–Nash equilibrium, then

$$\alpha=\tfrac{1}{3}(\gamma+\delta),$$
$$\beta=\tfrac{2}{3},$$
$$\gamma=\tfrac{1}{3}\alpha,$$
$$\delta=\tfrac{2}{3}.$$

Hence,

$$\alpha=\tfrac{1}{4},$$
$$\gamma=\tfrac{1}{12}.$$

So, the linear Bayesian–Nash equilibrium strategies are

$$p_b(v_b)=\tfrac{1}{12}+\tfrac{2}{3}v_b,$$
$$p_s(v_s)=\tfrac{1}{4}+\tfrac{2}{3}v_s.$$

Trade occurs whenever $p_b \geq p_s$ or, from evaluating equilibrium strategies, whenever $v_b \geq \frac{1}{4}+v_s$. However, there is a mutually beneficial trading opportunity whenever $v_b \geq v_s$. Therefore, the Double Auction does not ensure a Pareto efficient outcome. Some mutually beneficial trades do not take place.

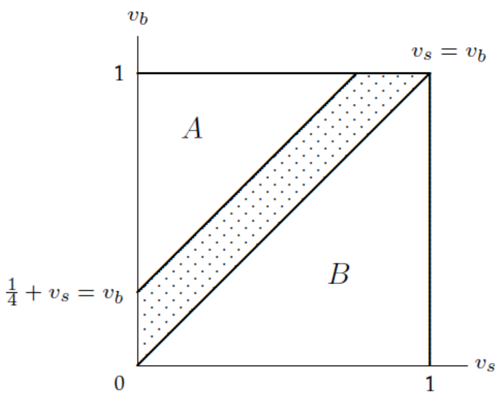

In the figure above, in area A there is efficient trade. In area B there is efficient lack of trade where the seller values the good more than the buyer. In the shaded area there is an inefficient lack of trade. There are many other equilibria, but none that are efficient. In particular whenever $v_s \leq v_b < \frac{1}{4}+v_2$, then there is an inefficient lack of trade.

2.4.6. Extensive–Form Games

2.5. Screening and Signaling

You can see that information asymmetry can lead to inefficiency–the adverse selection problem. In order to mitigate this inefficiency,'envied' types (e.g. high ability and low–risk) need to be distinguished from other types who wish to mimic them. This could be done through signaling, screening, or cheap–talk.

Definition: Screening
Screening is where the uninformed party moves first, offering a menu of contracts that separates types.

Definition: Signaling
Signaling is where the informed party moves first, taking an action with a type–dependent cost.

2.5.1. Screening

Example: Screening in Competitive Insurance Markets
Assume that there are risk–neutral competitive insurers. All individuals participating in the market have the same initial wealth, y, and potential loss, K. The probability of loss, π, is privately known. Two types, high and low–risk, $\pi_H > \pi_L$. The proportion of low–risk type L individuals is l. An insurance policy is a contingent consumption bundle, (y_{nc}, y_c). First, insurance companies simultaneously offer policies. Then each individual chooses a policy.

In subgame perfect equilibrium, given her privately–known risk p, each individual chooses the best policy offered by the insurers. Anticipating this, as well as the policies offered by the other insurers, each insurer maximizes its expected profit. The implications are that competition forces insurers to offer actuarially fair contracts (i.e. insurers break even in expectation) and at most two contracts are offered.

An equilibrium of this model can be separating or pooling. A pooling equilibrium has the firms offering one policy that is accepted by both types. In a pooling equilibrium, the price of coverage q is

$$\overline{\pi}q = [\lambda\pi_L + (1-\lambda)\pi_H]q.$$

A separating equilibrium has the firms offering distinct price contracts; one is accepted by low–risk types, the other by high–risk types. The price of coverage chosen by the high–risk and low–risk types are respectively

$$P_H(q_H) = \pi_H q_H \qquad \text{and} \qquad P_L(q_L) = \pi_L q_L.$$

The figure below depicts zero–profit lines for insuring H–types, L–types, and all types.

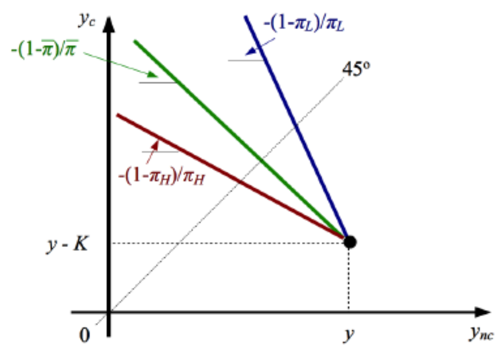

Lemma: The Spence–Mirrlees (Single-Crossing) Condition
For any point in the (y_{nc}, y_c) space, the indifference curve of the low–risk type is steeper than the indifference curve of the high–risk type running through that point.

Proof. The Expected utility of type $i \in \{L, H\}$ at (y_{nc}, y_c) is

$$U_i = (1 - \pi_i)u(y_{nc}) + \pi_i u(y_c).$$

The slope of the indifference curve, by total differentiation, is

$$\frac{\mathrm{d}y_c}{\mathrm{d}y_{nc}} = -\frac{1 - \pi_i}{\pi_i}\frac{u'(y_{nc})}{u'(y_c)}.$$

This is greater in absolute value for type L than type H at a given point, (y_{nc}, y_c), because

$$\frac{1 - \pi_L}{\pi_L} > \frac{1 - \pi_H}{\pi_H}.$$

□

In the figure below, through the same point, the L-type's indifference curve is steeper.

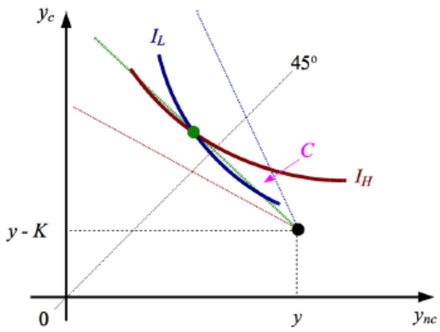

Claim: There is no pooling equilibrium in this market.

Proof. An actuarially fair pooling contract would lie on the middle budget line, as you can see in the previous figure. Through this point, the indifference curve of the L–type is steeper than that of the H–type, by the Spence–Mirrlees Lemma. Hence an insurance company would have a profitable deviation to a contract in area C—only L–types would choose it over the pooling contract, meaning that it would generate positive expected profit. Hence the putative pooling equilibrium is not an equilibrium. □

In a separating equilibrium two contracts are offered; coverage q_L with price P_L intended for the L–types, and coverage q_H with price P_H intended for the H–types. The contracts must be actuarially fair and the firms will break even

$$P_L = \pi_L q_L,$$
$$P_H = \pi_H q_H.$$

The contracts must also be incentive compatible.

Definition: Incentive Compatible
A menu is incentive compatible if each individual buys the contract meant for them.

Here, each individual will buy the insurance meant for them as a result of firm screening. The incentive compatibility constraint for the high–risk type is

$$\text{IC}_H: \ (1-\pi_H)u(y-P_H)+\pi_H u(y-P_H-K+q_H) \geq (1-\pi_H)u(y-P_L)+\pi_H u(y-P_L-K+q_L).$$

The incentive compatibility constraint for the low–risk type is:

$$\text{IC}_L: \ (1-\pi_L)u(y-P_L)+\pi_L u(y-P_L-K+q_L) \geq (1-\pi_L)u(y-P_H)+\pi_L u(y-P_H-K+q_H).$$

Note that $P_h = \pi_H q_H$ and $P_L = \pi_L q_L$ in both inequalities.

Claim: In a separating equilibrium, the high–risk type gets full coverage, $q_H = K$.

Proof. Suppose this is not the case, $q_h \neq K$. Then, a rival insurance company could deviate and offer a contract aimed at high–risk types so that more insurance is offered at a slight profit. All high–risk types prefer this contract and it makes money. If low–risk types also choose this contract then the deviating firm makes even more money. Therefore, $q_H \neq K$ cannot be part of a separating equilibrium. □

Low–risk types are offered the best contract for them that is actuarially fair, with $P_L = \pi_L q_L$, and would not be chosen by high-risk types. This occurs where IC_H binds, minimizing the distortion. Algebraically, q_L is given by

$$u(y - \pi_H K) = (1 - \pi_H)u(y - \pi_L q_L) + \pi_H u(y - \pi_L q_L - K + q_L).$$

The left–hand side is the H–type's utility from the contract aimed at her; the right-hand side is its utility from pretending to be an L–type. The other incentive compatibility constraint, IC_L, automatically holds, because of the fact that the L–type's indifference curves are steeper. The figure below on the left depicts the separating equilibrium.

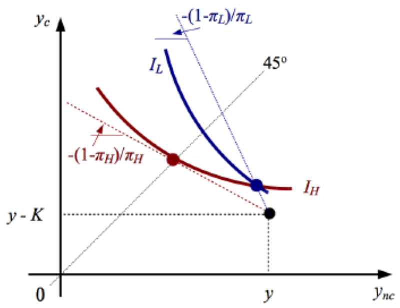

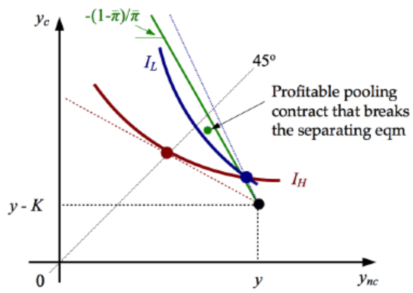

The results are; to prevent the high–risk types taking the low–risk insurance contract, it must be made just unattractive enough. Any reduction in cover hurts the high–risk types more than it hurts the low–risk types. The high–risk types are willing to pay more to avoid a reduction in cover, which is the basis for the screening. In the end, the consumption of low–risk types is distorted from the first best outcome (i.e. full coverage). The presence of high–risk types imposes an externality on the low–risk types, making it prohibitively costly for them to obtain full insurance. The figure above on the right depicts that the competitive screening model may have no equilibrium. Non–existence occurs when the proportion of low–risk types is too high[3].

[3] If an applied model does not have an equilibrium, then the model needs to be altered. In this case, it is fairly easy. Assume that firms can withdraw unprofitable contracts. With this slight modification the screening model has a unique equilibrium that is separating. In general, equilibrium non–existence can be a major problem.

Example: Monopolistic Screening in the Labor Market
A monopsonist makes a take–it–or–leave–it offer of a contract to a worker. A worker's type, here ability, which is private information, is denoted by $\theta \in \{\theta_L, \theta_H\}$. The probability of type θ_H is λ. There is a contractible action used for screening, $e \geq 0$. A contract is a pair, (w, e), composed of a wage, w and task e. When working, a type $\theta \in \{\theta_L, \theta_H\}$ worker's utility is

$$U(w, e) = w - \frac{1}{2}\frac{e^2}{\theta}.$$

Hence a low ability type incurs a higher marginal disutility from work. For simplicity, suppose that the worker's outside option is worth $r = 0$. So, she will work if

$$U(w, e) = w - \frac{1}{2}\frac{e^2}{\theta} \geq 0.$$

If employed with $e \geq 0$, a type $\theta \in \{\theta_L, \theta_H\}$ worker's marginal product is

$$\text{MP}(\theta) = \theta + \alpha e,$$

and the employer's profit is

$$\pi = \theta + \alpha e - w,$$

The individual rationality (participation) constraints of the two possible types of workers are the respective indifference curves, where

$$w = \frac{1}{2}\frac{e^2}{\theta}.$$

Below is a figure that depicts the two types of workers' productivity and indifference curves.

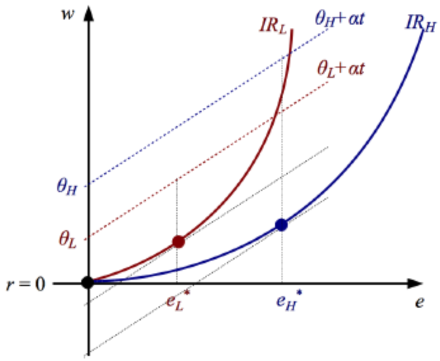

Note that in the monopsony first–best case, the firm chooses combinations, (w, e), such that the participation constraints of both types bind. Rather than firms making zero profit, as in the competitive case, workers get utility $r = 0$.

In the monopsony first–best case, the firsm sets

$$w = \frac{1}{2}\frac{e^2}{\theta},$$

for $\theta \in \{\theta_L, \theta_H\}$. This means that the firm's profit when hiring a θ type for task e is

$$\pi = \theta + \alpha e - \frac{1}{2}\frac{e^2}{\theta}.$$

The first–best task, e, maximizes profit and is such that the marginal disutility of e equals the marginal productivity of e. It follows that

$$\frac{\mathrm{d}e^2}{\mathrm{d}e2\theta} = \frac{\mathrm{d}(\theta + \alpha e)}{\mathrm{d}e}$$
$$e^*(\theta) = \alpha\theta.$$

The first–best wage for a type θ worker is

$$w(\theta) = \frac{1}{2}\frac{e^*(\theta)^2}{\theta}$$
$$w(\theta) = \frac{1}{2}\alpha^2\theta.$$

Therefore, $e_L^* < e_H^*$ and $w_L^* < w_H^*$. Furthermore, the Spence–Mirrlees single-crossing condition holds, that is, through any given point, the indifference curve the low–ability type θ_L is steeper. The high–ability type envies the deal of the low–ability type. Note that the opposite holds in the competitive screening case, where market structure can fundamentally alter the effects of asymmetric information and the nature of adverse selection.

If the worker's type is not observable to the wage–setting employer, the employer solves the following problem

$$\max_{e_L, e_H, w_L, w_H \geq 0} \pi = \lambda(\theta_H + \alpha e_H - w_H) + (1-\lambda)(\theta_L + \alpha e_L - w_L),$$
$$\text{s. t. } \mathrm{IR}_L : w_L - \frac{1}{2}\frac{e_L^2}{\theta_L} \geq 0,$$
$$\text{s. t. } \mathrm{IR}_H : w_H - \frac{1}{2}\frac{e_H^2}{\theta_H} \geq 0,$$
$$\text{s. t. } \mathrm{IC}_L : w_L - \frac{1}{2}\frac{e_L^2}{\theta_L} \geq w_H - \frac{1}{2}\frac{e_H^2}{\theta_L},$$
$$\text{s. t. } \mathrm{IC}_H : w_H - \frac{1}{2}\frac{e_H^2}{\theta_H} \geq w_L - \frac{1}{2}\frac{e_L^2}{\theta_H}.$$

The first two constraints, IR_L and IR_H, are the participation constraints, and the second two constraints, IC_L and IC_H, are incentive compatibility constraints. If IR_L holds, then

$$w_L - \frac{1}{2}\frac{e_L^2}{\theta_H} > 0,$$

because $\theta_H > \theta_L$. Hence, by IC_H, then

$$w_H - \frac{1}{2}\frac{e_H^2}{\theta_H} > 0.$$

The participation constraints for the high–ability type, IR_H, automatically holds, i.e. it is slack.

Note that it cannot be that both IR_H and IR_L hold, because then the employer could decrease both w_H and w_L by $\varepsilon > 0$, without violating any of the constraints, while increasing the employer's payoff. Therefore IR_L holds with equality, that is, the low type's participation constraint is binding. Furthermore, in the solution IC_H is binding, IC_L is slack, and $e_L < e_L^*$.

Proof. If L type's offer is at e_L (red dot) then H type's offer is in the shaded area, say at the blue dot. This implies that IC_L is slack (red dot puts L type on higher IC). If $e_L > e_L^*$, then the employer increases profit by moving e_L towards e_L^*. □

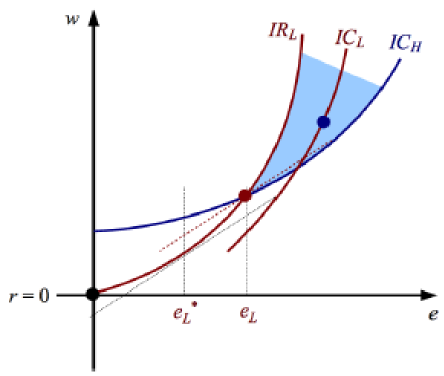

Thus, in the solution $e < e_L^*$, $e_H = e_H^*$, and the IR_L and IC_H are binding.

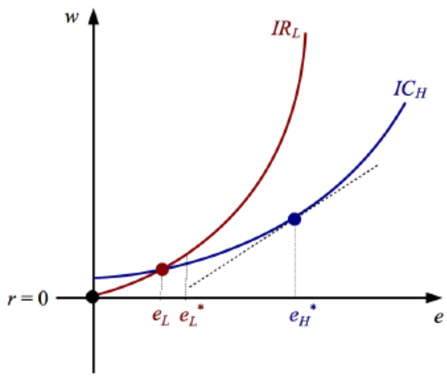

Properties:

- There is full rent extraction at the bottom: IR_L is binding (i.e. the low–ability type gets her reservation wage).
- All types but the lowest type strictly prefer the outcome of the second–best to that of the first–best. (Implied here by slack IC_H.)
- There is no distortion at the top: $e_H = e_H^*$.
- The distortion $e_L < e_L^*$ is due to the non–linearity of the indifference curves. This implies that there is an interior solution where H types enjoy some rents and L types face some distortion in e_L.

2.5.2. Signaling

Adverse selection can be resolved by the informed party moving first to signal her type. In particular, in a separating equilibrium, the 'good' type takes an action that distinguishes herself from the 'bad' type. The uninformed party then updates her beliefs regarding her opponent's type. For the signal to be credible it must be too costly for the bad type to send; of course, to be sent at all it must be sufficiently affordable for the good type.

Reading: Spence (1973), Veblen (1899), Zahavi (1975).

Example: Signaling in the Labor Market *Spence (1973)*
Consider a worker attempting to signal her ability to a competitive labor market. Nature picks the worker's type (productivity)

$$\theta \in \{\theta_L, \theta_H\},$$

where $\Pr(\theta = \theta_H) = \lambda$. The worker observes θ and chooses her education level $e \geq 0$. The worker's payoff is

$$u(w, e) = w - \frac{e}{\theta},$$

when working (the Spence–Mirrlees condition holds). The worker's outside option is $r(\theta) = 0$.

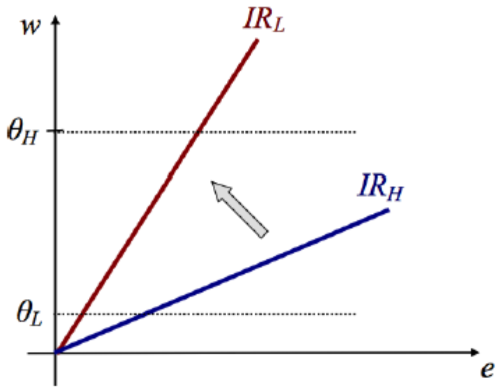

Figure II.1: Worker's Productivity and Indifference Curves

Note: Education, e, is not productive. The first–best (full information) outcome is $e = 0$ for both types and wage θ_H for H types and θ_L for L types.

The firm's payoff is

$$\pi(w) = \begin{cases} \theta - w & \text{if it employs the worker,} \\ 0 & \text{if it does not.} \end{cases}$$

The worker's strategy is $e(\theta)$ (i.e. a level of education for each $\theta \in \{\theta_L, \theta_H\}$). The firm's strategy is $w(e)$ (i.e. a wage for each level of education $e \geq 0$. The wage offered by firms depends on their belief about the worker's type given the observed education level, e. Denote the firm's belief for each education level, $e \geq 0$, as

$$\mu(e) \equiv \Pr[\theta = \theta_H | e].$$

A Perfect Bayesian Equilibrium (PBE) in this game consists of a strategy for the worker, $e(\theta)$, a strategy for each firm, $w(e)$, and each firm's beliefs, $\mu(e)$. A PBE is a triplet (e^*, w^*, μ^*) such that the follow conditions hold.

- Following any $e \geq 0$, each firm's reaction, $w^*(e)$, is optimal given its beliefs and the other firms' anticipated reactions.
- Type θ workers pick $e^*(\theta)$ optimally given their type and the firms' anticipated reactions.
- A firm's beliefs, $\mu^*(e)$, are consistent with the worker's equilibrium strategy and the prior distribution of θ.

Consistent beliefs are such that the following conditions hold.

- After an action that the worker plays in equilibrium, Bayes' rule is used to compute

$$\mu^*(e) = \Pr(\theta = \theta_H | e^*(\theta) = e).$$

- If the worker picks an out–of–equilibrium action, then the firms are free to believe anything about the worker's type.

There are pooling equilibria such that the following conditions hold.

- Both worker types choose the same $e^* \in [0, \lambda(\theta_H - \theta_L)\theta_L]$.
- Beliefs are $\mu^*(e^*) = \lambda$ and $\mu^*(e) = 0$ for all $e \neq e^*$.
- Firms set $w^*(e^*) = \lambda\theta_H + (1-\lambda)\theta_L$ and $w^*(e) = \theta_L$ otherwise.

In the pooling equilibrium, you can check that the following conditions are true.

- Firms' beliefs are consistent with the equilibrium.
- Wage equals expected productivity, as dictated by perfect competition.
- Type θ plays e^* if and only if for all e

$$\lambda\theta_H + (1-\lambda)\theta_L - \frac{e^*}{\theta} \geq \theta_L - \frac{e}{\theta}$$

$$\lambda(\theta_H - \theta_L)\theta \geq e^* - e.$$

This holds for all $e \geq 0$ for both types if and only if

$$e^* \leq \lambda(\theta_H - \theta_L)\theta_L.$$

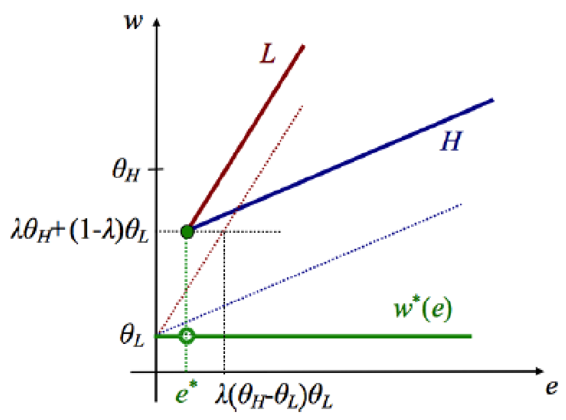

Figure II.2: Example of a Pooling Equilibrium

Note: $0 \leq e^* \leq \lambda(\theta_H - \theta_L)\theta_L$.

Next, there are separating equilibria such that the following conditions hold.

- $e^*(\theta_L) = 0$ and $e^*(\theta_H) = e_H \in [(\theta_H - \theta_L)\theta_L, (\theta_H - \theta_L)\theta_H]$.
- Beliefs are $\mu^*(e_H) = 1$ and $\mu^*(e) = 0$ for all $e \neq e_H$.
- Firms set $w^*(e_H) = \theta_H$ and $w^*(e) = \theta_L$ for all $e \neq e_H$.

You can check that the following conditions hold in equilibrium.

- Firms' beliefs are consistent and reactions are optimal given beliefs.
- Type θ_L prefers $e = 0$ to e_H if and only if

$$\theta_L \geq \theta_H - \frac{e_H}{\theta_L}$$

$$e_H \geq (\theta_H - \theta_L)\theta_L.$$

- Type θ_H prefers $e = e_H$ to 0 if and only if

$$\theta_H - \frac{e_H}{\theta_H} \geq \theta_L$$

$$e_H \leq (\theta_H - \theta_L)\theta_H.$$

Also not that

$$(\theta_H - \theta_L)\theta_L \leq e_H \leq (\theta_H - \theta_L)\theta_H.$$

An alternative equilibrium selection device involves ruling out Pareto inefficient equilibria: equilibria in which at least one player can be made better off while all other players are left no worse off. Clearly the minimal–cost separating equilibrium Pareto dominates all other separating equilibria. However, the pooling equilibrium in which $e = 0$ Pareto–dominates the minimal–cost separating equilibrium if the proportion of High types λ is sufficiently high. This is because high types waste resources distinguishing themselves without a large increase in their wage (because they are pooled mainly with other high types).

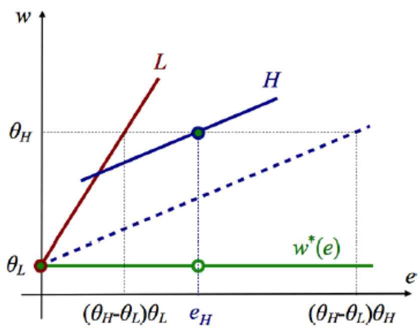

Separating Equilibrium

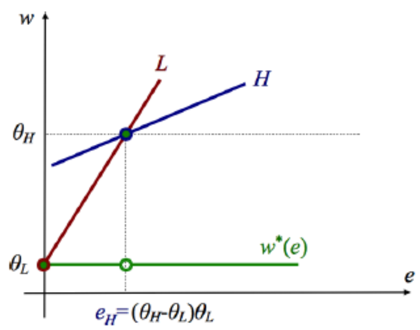

Minimal–Cost Separating Equilibrium

Is there a suitable equilibrium refinement that selects this equilibrium (otherwise known as the Riley Equilibrium)? Perform an equilibrium dominance test, that is, following a deviation by the worker, put zero probability on types whose equilibrium payoffs exceed any payoff they can get from deviation given a competitive wage setting. Recall that the Intuitive Criterion rules out every PBE that fails the equilibrium dominance test. Here, the intuitive criterion places certain restrictions on the function $w^*(e)$, by constraining out–of–equilibrium beliefs.

Theorem: Cho and Kreps 1987
In any Spencian signaling game with two sender types the Intuitive Criterion selects an equilibrium with the Riley outcome.

This is true more generally as long as types are ordered (high and low), payoffs satisfy single crossing, and there are only two types. When there are more than two types, then stronger refinements select the minimal–cost separating equilibrium outcome. The Intuitive Criterion is used for equilibrium selection in many other games of incomplete information.

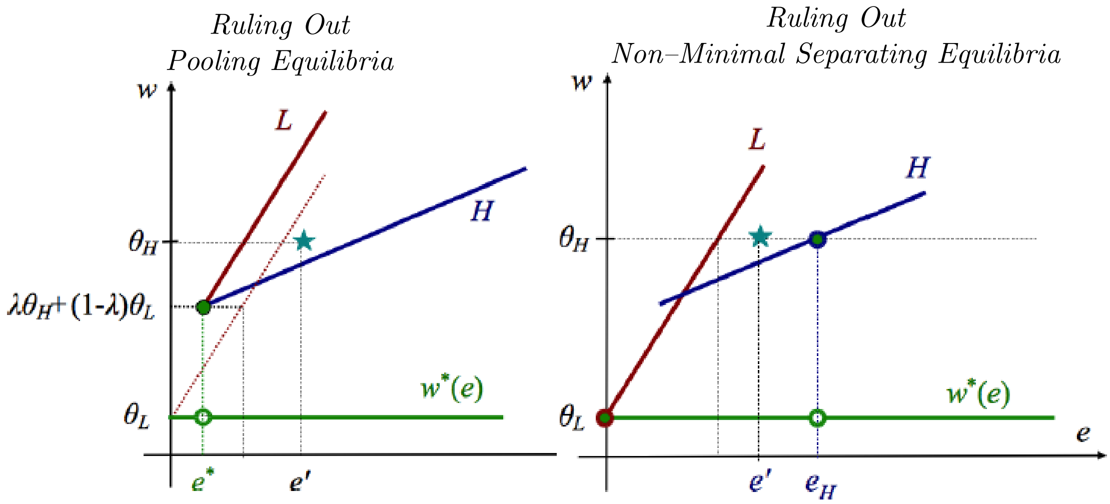

The key points to take away from Spencian Signaling are as follows. The structure; a sender wants the receiver to think that she is a high type, θ_H, and there is single–crossing, that is, the marginal cost of action is everywhere lower for θ_H. The Perfect Bayesian equilibrium and Intuitive Criterion, with two types, yields a minimal–cost separating outcome. This is the same outcome as under competitive screening. Does the signal have to be costly? Does education have to be a pure social waste? Maybe, but not necessarily, as signaling also works if education is more beneficial for the high types, in which case the high type overinvests in education relative to the first–best outcome.

2.6. Strategic Communication and Moral Hazard

2.6.1. Strategic Communication

We have seen that screening and signaling are two possible solutions to the adverse selection problem. However, could the informed party not simply disclose the hidden information? This is of course a matter of credibility. Under what conditions can an informed party credibly disclose information?

Definition: Verifiable Communication Game
In verifiable communication games, the informed party cannot report false information, but can suppress information.

Definition: Cheap Talk
Cheap–Talk is where the informed party can send a costless message regarding her type. The informed party can disclose anything she desires, and disclosure is unverifiable.

2.6.2. Verifiable Disclosure

Consider communication between an agent A and a principal P.

Quarter III

"Happiness is fugitive; dissatisfaction and boredom are real."
– Jack Vance, *Emphyrio*

This part of the microeconomic sequence starts by revising basic models of individual choice and Pareto efficiency and their extensions under uncertainty. Next, economic mechanisms—auctions, bargaining, voting, and matching—that can be used to allocate any kind of goods among any number of agents are analyzed. In contrast with the Walrasian general equilibrium, we find persistent conflicts between individual rationality and social efficiency. Finally, we discuss the more abstract part of social choice and mechanism design that establish some general tradeoffs between individual incentives and social efficiency.

3.1. Utility and Preferences

3.1.1. Preferences

The way we will model individual choice starts with an arbitrary set X and a binary relation $\succeq$ on X. The set X consists of all alternatives (payoffs, outcomes, options, etc.) that may become feasible in relevant observations.

Definition: Weak preference is $x \succeq y$. This says x is weakly preferred to y when the decision maker is willing to choose x when only two options x and y are feasible.

Definition: Strict preference is $x \succ y \Leftrightarrow x \succeq y$ and not $y \succeq x$.

Definition: Indifference is $x \sim y \Leftrightarrow x \succeq y$ and $y \succeq x$.

Observing preferences can be problematic.
(i) The feasible set may be hard to control or choice avoidance can be an issue.
(ii) Indifference between x and y requires more than just one observation.
(iii) The required number of observations is at least $X * (X - 1)/2$.
(iv) Experimental observations can be costly when the payoffs are high.
(v) Path Dependence: Previous observations can interfere with the next ones.[4]

The basic properties of preferences $\succeq$ are

- reflexive if $x \succeq x$ for all $x \in X$.
- complete if $\forall\ x, y \in X$ either $x \succeq y$, or $y \succeq x$, or both.
- transitive[5] if $\forall\ x, y, z \in X$, then $x \succeq y \succeq z$ implies $x \succeq z$.
- antisymmetric if $\forall\ x, y \in X$, if $x \sim y$, then $x = y$.

Definition: Weak order preferences are complete and transitive.

Definition: Linear order preferences are complete, transitive, and antisymmetric.

[4] *The Becker–de Groot–Marschak (BDM) elicitation method can help avoid path dependence. The subject formulates a bid. The bid is compared to a random price. If the subject's bid is greater than the price, she pays the price and receives the item. If the subject's bid is lower than the price, she pays nothing and receives nothing. This method is equivalent to a Vickrey auction.*

[5] *The Condorcet Paradox* occurs when $x \succ y \succ z$ and $y \succ z \succ x$ and $z \succ x \succ y$. The voting paradox is a situation in which collective preferences can be cyclic, even if the preferences of individual voters are not cyclic. The choice of winner by a voting mechanism could be influenced by whether or not a losing candidate is available to be voted for, violating the axiom of independence of irrelevant alternatives.

3.1.2. Utility Representations

If X is finite, and $\succeq$ is complete and transitive, then there is a utility representation U. Preferences are represented by $U : X \to \mathbb{R}$ if for all $x, y \in X$, then $x \succeq y \Leftrightarrow U(x) \geq U(y)$.

- If U represents $\succeq$, then another function $U' = f \circ U$ for some strictly increasing one–variable function $f : D \to \mathbb{R}$ with a domain D that contains the range $U(X)$ of the utility U.
- U is unique up to a strictly monotonic transformation of U.
- Utility representation implies that $\succeq$ is complete and transitive.
- The converse is true if X is finite, but not tue if X is infinite[6].

3.1.3. Continuity

Let $\succeq$ be a preference on $X \subset \mathbb{R}^n$.

Definition: Call $\succeq$ upper semi–continuous if the upper contour sets $\{y \in X : y \succeq x\}$ are closed for all $x \in X$.

Definition: Call $\succeq$ lower semi–continuous if the lower contour sets $\{y \in X : y \preceq x\}$ are closed for all $x \in X$.

Definition: Call $\succeq$ continuous if all upper and lower contour sets are closed.

Theorem: Preference $\succeq$ is complete, transitive, and (upper, lower) continuous if and only if it has an (upper, lower) continuous utility representation.

[6] If X is infinite, there may not be a utility representation. Consider *lexicographic preferences* where an agent prefers any amount of one good (X) to any amount of another (Y). These preferences cannot be represented with a utility function.

3.2. Expected Utility and Pareto Efficiency

3.2.1. The Quasi–linear utility model

Let $X = Z \times \mathbb{R} = \{x = (z, m) : z \in Z, m \in \mathbb{R}\}$.

Definition: Quasi–linear utility is $U(z, m) = v(z) + m$.

The key assumptions are:

- Completeness and transitivity of preference (any utility model).
- The presence of money separability

$$(z, m) \succeq (z', m') \Rightarrow (z, m + a) \succeq (z', m' + a)$$

  for all z, z', m, m', and a[7].

Conversely, these conditions together with appropriate continuity imply the quasi–linear utility representation.

3.2.2. Expected Value and Expected Utility

Let $X = L(Z)$ be the set of all lotteries I over Z: probability distributions that deliver payoffs $z_1, \ldots, z_n \in Z$ with probabilities $p_1, \ldots, p_n$ so that $p_i \geq 0$ and $\sum_i p_i = 1$. If Z is an interval of the real line, then lotteries have monetary payoffs.

Definition: The Expected value models is

$$U(i) = \sum_{i=1}^{I} p_i z_i.$$

Definition: The Expected utility models is

$$U(I) = \sum_{i=1}^{I} p_i u(z_i),$$

where $u : Z \to \mathbb{R}$ is called a Bernoulli utility index.

In the expected value model, payoffs must be monetary. In the expected utility model, payoffs can be arbitrary (health, social status, etc.).

3.2.3. Risk Attitudes

Let $X = L(Z)$ where $Z \subset \mathbb{R}$, and $\sum_i p_i z_i$ be a sure payment.

Definition: Preference $\succeq$ is risk averse if for all I, $I \preceq \sum_i p_i z_i$.

Definition: Preference $\succeq$ is risk loving if for all I, $I \succeq \sum_i p_i z_i$.

Definition: Preference $\succeq$ is risk neutral if for all I, $I \sim \sum_i p_i z_i$.

Risk attitudes require monetary payoffs, but can be defined without assuming a particular utility representation. Within the expected utility model;

- Risk aversion is equivalent to the concavity of $u \Leftrightarrow u'$ is decreasing $\Leftrightarrow u'' \leq 0$.
- Risk loving is equivalent to the convexity of $u \Leftrightarrow u'$ is increasing $\Leftrightarrow u'' \geq 0$.
- Risk neutrality is equivalent to the linearity of u.

[7] Equivalently, $WTP(z) = WTA(z)$ where the willingness–to–pay and the willingness–to–accept are measured relative to some status quo (z^*, m^*).

3.2.4. Constant Absolute Risk Aversion

The constant absolute risk aversion (CARA) parametric model has Bernoulli utility index

$$u(z) = -e^{-\lambda z},$$

where $\lambda > 0$. If $\lambda = 0$, then $u(z) = z$.

- Utility u is defined for all z.
- CARA satisfies wealth invariance: $I \succeq I'$ implies $I \oplus a \succeq I' \oplus a$, where all payoffs are modified by $a \in \mathbb{R}$.

The downside is that risk aversion is too strong over large intervals. For further reading, see Rabin, Risk Aversion and Expected Utility Theory: A Calibration Theorem. 2000.

3.2.5. Expected Utility and Independence

Besides completeness and transitivity, expected utility model satisfies independence.

Definition: Independence is

$$I \succeq I' \Rightarrow \alpha I + (1-\alpha)I'' \succeq \alpha I' + (1-\alpha)I''.$$

The mixtures are probabilistic here, and the axiom that captures the separability principle across mutually exclusive cases that occur with probability α and $1-\alpha$ respectively. In the first case, the outcomes are determined by I and I' respectively. In the second case, the outcomes are determined by I''.

Von Neumann and Morgenstern (1944) showed that completeness, transitivity, continuity, and independence are necessary and sufficient for the expected utility representation.

The Allais Paradoxes, such as the common ratio effect, contradicts independence.

3.2.6. Uniqueness of Utility Representations

The general utility function U can be distorted by any strictly increasing function f. The composition $f(U) = f \circ U$ represents the same preference $\succeq$. If we observe $\succeq$, we cannot infer function u.

- In the quasi–linear utility model, money serves as the measuring device, and the index $v(z)$ is unique up to adding a constant

$$v(z) + \beta \text{ for } \beta \in \mathbb{R}.$$

- In the expected utility model, probabilities serve as the measuring device, and the index $u(z)$ is unique up to a positive linear transformation

$$\alpha u(z) + \beta \text{ for } \alpha > 0 \text{ and } \beta \in \mathbb{R}.$$

3.3. The Nash Bargaining Solution and Walrasian Equilibria

3.3.1. The Nash Bargaining Solution

Let I be a population of agents. Take any utility possiblity set $\mathcal{U} \subset \mathbb{R}^I$ such that

- $\mathcal{U}$ is convex and closed. For example, $\mathcal{U}$ can be the set of all utility vectors $(U_1(x), \ldots, U_I(x))$ produced by expected utility functions U_i over lotteries x.
- $(0, \ldots, 0) \in \mathcal{U}$ is a threat (status quo) point that describes what the agents obtain.
- $\mathcal{U} \cap \mathbb{R}^I_{++}$ is bounded and not empty.

The Nash Bargaining solution is

$$N(\mathcal{U}) = \underset{(u_1,\ldots,u_I)\in\mathcal{U}\cap\mathbb{R}_{++}}{\operatorname{argmax}} [\log u_1 + \cdots + \log u_I].$$

The Nash Bargaining solution is the only such function f that is defined over all suitable $\mathcal{U}$ and satisfies

- Pareto efficiency: $f(\mathcal{U}) \in \mathcal{U}$ is Pareto efficient in $\mathcal{U}$.
- Symmetry: if $\Delta = \{(u_1, \ldots, u_I) \in \mathbb{R}_{++} : \sum_i u_i \leq 1\}$, then $f(\Delta) = (\frac{1}{I}, \ldots, \frac{1}{I})$.
- Rescaling Invariance: if

$$\mathcal{V} = \{(\alpha_1 u_1, \ldots, \alpha_n u_n) : (u_1, \ldots, u_n) \in \mathcal{U}\},$$

  then $f(\mathcal{V}) = (\alpha_1 f_1(\mathcal{U}), \ldots, \alpha_n f_n(\mathcal{U}))$.
- Independence of Irrelevant Alternatives: if $f(\mathcal{U}) \in \mathcal{V} \subset \mathcal{U}$, then $f(\mathcal{V}) = f(\mathcal{U})$.

If symmetry is dropped, then Nash is generalized by

$$G(\mathcal{U}) = \underset{(u_1,\ldots,u_I)\in\mathcal{U}\cap\mathbb{R}_{++}}{\operatorname{argmax}} [\beta_1 \log u_1 + \cdots + \beta_I \log u_I],$$

for some $\beta_i > 0$ and $\sum_i \beta_i = 1$.

Example: Assignment Problems
A valuable item needs to be assigned to one of I agents with consumption values $V_i > 0$ and quasi–linear utilities. Money can be transferred. The utility possiblity set is $(u_1, \ldots, u_I) : \sum_i u_i = \max_i V_i$. The Nash bargaining solution assigns the item to the agent with the highest consumption value V_m and makes this agent pay $\frac{V_m}{I}$ to each of the other agents.

Example: Coordination Problems
I agents need to coordinate on one of the activities $(a_1, \ldots, a_k)$ with utility vectors $v_1, \ldots, v_k \in \mathbb{R}^I_{++}$ respectively. If they do not coordinate, then they receive $(0, \ldots, 0)$. Coordination can be randomized. The set $\mathcal{U}$ is the convex hull of $v_1, \ldots, v_k, 0$. The Nash bargaining solution can be found by checking all the edges (faces) of the Pareto frontier.

The Nash Bargaining solution serves as a theoretical benchmark for more practical assignment mechanisms, and is a part of other models (e.g. it can be used to assign an outcome after the utility possibility set has be determined by other means). The main problem, which will be shared by other mechanisms as well, is how does one observe $\mathcal{U}$?

Example: In the two–agent assignment problem, the agent with the higher value V will have incentives to report a lower value in order to pay less to the other individual. How can you induce truth–telling?

3.3.2. Walrasian Equilibrium

Walrasian (Competitive) Equilibrium is another theoretical solution for pure exchange economies with

- L infinitely divisible goods consumed in non–negative consumption bundles
$$(x_1, \ldots, x_L) \in \mathbb{R}^L_+.$$
- Finitely or infinitely many agents with preferences $\succeq_i$ and utility functions $u_i(x_{1i}, \ldots, x_{Li})$.
- Initial endowment vectors $\omega_i \in \mathbb{R}^L_+$.

Definition: Say that the allocation $x_1, \ldots, x_I \in \mathbb{R}^L$ is feasible if $\sum_i x_i = \sum_i \omega_i$.

If $I = 2$, then each feasible allocation $x_1, x_2 \in \mathbb{R}^L_+$ corresponds to a point in the Edgeworth box.

A Walrasian equilibrium consists of an allocation $x_i = (x_{1i} \ldots, x_{Li}) \in \mathbb{R}^L_+$ for all $i = 1, \ldots, I$ and a price vector $p \in \mathbb{R}^L_{++}$ such that

Definition: Each x_i is affordable if $p \cdot x_i \leq p \cdot \omega_i$.

Definition: Each x_i is individually optimal if $x_i \succeq_i y_i$ for all y_i such that $p \cdot y_i \leq p \cdot \omega_i$.

Definition: The allocation $x_1, \ldots, x_I \in \mathbb{R}^L_+$ is feasible if $\sum_i x_i = \sum_i \omega_i$.

Finding a Walrasian equilibria can be a hard computational problem. Note that the Nash Bargaining solution can be also defined for the threat point given by the initial endowment vectors u_i and the utility possibility set given by concave utility functions u_i or through randomization.

3.3.3. Cobb–Douglas Utility Function

If all u_i have Cobb–Douglas form utility
$$u_i(x_{1i}, \ldots, x_{Li}) = \alpha_i \log x_{1i} + \cdots + \alpha_{Li} \log x_{Li},$$
for some $\alpha_{li} > 0$ such that $\sum_l \alpha_{li} = 1$ then

- The equilibrium allocation is unique.
- The equilibrium price vector is unique if $p_1 = 1$.
- The equilibrium can be found by solving a system of linear equations.

Example: Cobb–Douglas utility functions
The problem is to
$$\max u_i(x_{1i}, \ldots, x_{Li}) \text{ s.t. } p \cdot x_i \leq M.$$
There is a unique solution
$$x_{li} = \frac{\alpha_{li} M_i}{p_l} \text{ for all } l = 1, \ldots, L \text{ and } i = 1, \ldots, I.$$
Thus
$$\underbrace{\sum_{i=1}^{I} \frac{\alpha_{li}(P \cdot w_i)}{p_l}}_{\text{“AggregateDemand}} = \underbrace{\sum_{i=1}^{I} w_{li}}_{\text{“AggregateSupply''}}$$

Problem: Cobb–Douglas utility functions
Given $I = L = 2$ and $w_1 = w_2 = (1, 1)$ and preferences

what is the Walrasian equilibrium?

Solution:

3.4. Walrasian Equilibria and its Extensions

3.4.1. Pareto Efficiency in Walrasian Equilibria

Say that preference $\succeq_i$ is locally non–satiated if for any $x \in \mathbb{R}^L_+$ and $\varepsilon > 0$, there is $y \in \mathbb{R}^L_+$ such that $\|y - x\| < \varepsilon$ and $y \succ_i x$. Any utility function that is strictly monotonic in at least one component must be locally non–satiated.

Theorem: The First Welfare Theorem
If all preferences are locally non–satiated, then any Walrasian equilibrium allocation must be Pareto efficient.

Walrasian equilibria are not the only reasonable ways to achieve Pareto efficiency in pure exchange economies with full information about preferences and initial endowments. Nash bargaining solution works as well. Both WE and NBS are susceptible to manipulation: agents have incentives to lie about their preferences (e.g. parameters of their Cobb–Douglas utilities) to get a better allocation by the mechanism rules.

3.4.2. Pareto Efficiency and Walrasian Equilibria

If all u_i are differentiable, then one can write F.O.C. for interior Pareto efficiency and WE allocations.

Definition: Let

$$\mathrm{MRS}_{kli}(x_i) = \frac{\partial u_i(x_{1i}, \ldots, x_{Li})/\partial x_{ki}}{\partial u_i(x_{1i}, \ldots, x_{Li})/\partial x_{li}}$$

be the marginal rate of substitution for agent i between goods k and l. It is how many small units of good l she is willing to sacrifice for one small unit of k.

Allocation $x_1, \ldots, x_I$ must be feasible ($\sum_i x_i = \sum_i \omega_i$) and interior ($x_i \in \mathbb{R}^I_+$ for all i). Pareto efficiency at an interior feasible allocation $x_1, \ldots, x_I$ implies that

$$\mathrm{MRS}_{kli} = \mathrm{MRS}_{klj} \text{ for all } k, l, i, j.$$

Walrasian equilibrium implies

$$\mathrm{MRS}_{kli} = \mathrm{MRS}_{klj} = \frac{p_k}{p_l} \text{ and } p \cdot \omega_i = p \cdot x_i.$$

The ratio $\frac{p_k}{p_l}$ is the relative price—how many units of good l are required to buy one unit of good k at market prices.

Problem: Pareto efficiency
Let $u_1(x) = x$ and $u_2(x) = -\alpha x$, where $x \in [0, 1]$, $\alpha > 0$. What are the Pareto efficient outcomes?

Solution: Any x is a Pareto efficient outcome.

3.4.3. The Core

Any Walrasian equilibrium allocation satisfies a stronger form of Pareto Efficiency, known as the core property. Take any non–empty coalition $C \subset I$. Say that this coalition blocks a feasible allocation $x_1, \ldots, x_I$ if there are bundles $\{y_c \in \mathbb{R}^L_+ : c \in C\}$ such that

- $y_c \succeq_c x_c$ for all $c \in C$ and $y_b \succ_b x_b$ for some $b \in C$,
- $\sum_{c \in C} y_c = \sum_{c \in C} \omega_c$.

Definition: Say that $x_1, \ldots, x_I$ has the core property if it is feasible and is not blocked by any coalition $C \subset I$.

If $C = I$, then $x_1, \ldots, x_I$ is not blocked by C if and only if it is Pareto optimal.

If $C = \{i\}$, then $x_1, \ldots, x_I$ is not blocked by C if and only if $x_i \succeq_i \omega_i$. This is called individual rationality.

The collection of all allocations with the core property is called the core of the pure exchange economy. In the 2×2 case, the core is called the contract curve.

Theorem: If all preferences are locally non–satiated, then any Walrasian equilibrium allocation must be in the core.

Note that the core is defined without money, but typically has many possible allocations.

Example: Suppose that $X^A = (1, 1)$ and $X^B = (1, 1)$. The set of Pareto efficient allocations is ... The core is ... The Walrasian equilibrium is....

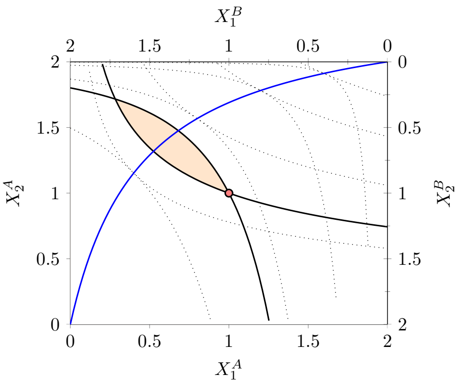

Figure III.1: The Corresponding Edgeworth Box

Note: Only type θ_H may gain from deviating to e'.
By the Intuitive Criterion $\mu(e') = 1$ and $w(e') = \theta_H$,
which rules out the equilibrium. □

3.4.4. Walrasian Equilibrium in Large Economies

Assume that the population is $I = [0, 1]$ so that no single individual $i \in I$ has any effect on the entire economy. Initial endowments and preferences are given as functions

$$\omega(t) : [0, 1] \to \mathbb{R}^L \text{ and } \succeq (t).$$

A Walrasian equilibrium consists of an allocation $x(t) : [0, 1] \to \mathbb{R}^L$ and a price vector $p \in \mathbb{R}^L_{++}$ such that for (almost all) $t \in [0, 1]$

- $x(t)$ is affordable: $p \cdot x(t) \le p \cdot \omega(t)$.
- $x(t)$ is individually optimal: $x(t) \succeq_t y$ for all y such that $p \cdot y \le p \cdot \omega(t)$.
- The allocation $x(t)$ is feasible:$\int_0^1 x(t) = \int_0^1 \omega(t0$.

$x(t)$ has the core property if there is no coalition $C \subset [0, 1]$ with positive measure that blocks $x(t)$, that is, there exists no assignment $\{y(c) \in \mathbb{R}^L_+ : c \in C\}$ such that

- $y_c \succ_c x_c$ for all $c \in C$,
- $\int_C y(c) = \int_C \omega(c)$.

Theorem: Aumann 1964
If all preferences are strictly monotonic, continuous, and measurable then the core of this pure exchange economy coincides with the set of all Walrasian equilibria allocations.

The major conclusions are

- In large economies, the Walrasian equilibria allocations can be identified even without money v.i.a. the core property.
- The core shrinks and becomes easier to find than for finite populations.

Reading: R. Aumann, Markets with a continuum of traders, Econometrica, 1964.

Problem: Pareto efficiency in large economies
Let $u_i(x) = x$ and $u_j(x) = -\alpha x$, where $x \in [0, 1]$, $\alpha > 0$, for a continuum of agents types i and j. What are the Pareto efficient outcomes?

Solution:

$$U_i = \int_0^1 x f(x)\,\mathrm{d}x = \mathbb{E}_f(x)$$

$$U_j = \int_0^1 -\alpha x f(x)\,\mathrm{d}x = -\alpha\,\mathbb{E}_f(x)$$

As $f \to g$ and $U_1(g) > U_2(f)$, then $\mathbb{E}_g(x) > \mathbb{E}_f(x)$, so $-\alpha\,\mathbb{E}_g(x) < -\alpha\,\mathbb{E}_f(x)$. Thus $U_j(g) < U_j(f)$ and any f is a Pareto efficient outcome.

Problem: Pareto efficiency in large economies with quasi–linear utility
Let $U_1(x) = x + m_1$ and $U_j(x) = -\alpha x + m_j$, where $\sum_i m_i = 0$. What is the set of Pareto efficient allocations?

Solution: First, choose $x^* \in \operatorname{argmax} \sum_i U_i(x)$.

$$\sum_i U_i(x) = x - \alpha N x = (1 - \alpha N)x$$

- If $\alpha N > 1$, then $x^* = 0$.
- If $\alpha N < 1$, then $x^* = 1$.
- If $\alpha N = 1$, then $x^* \in [0, 1]$

Proof. Suppose $\alpha N > 1$ and $x > 0$. Then $u_i = x + \overline{m}_1$ and $u_j = -\alpha x + \overline{m}_j$. So, $U_i = (\frac{x}{N})N + \overline{m}_1$ and $U_j = -\frac{x}{N} + \overline{m}_j$ is a Pareto improvement. □

3.5. Subjective Expected Utility and Arrow-Debreu Equilibria

3.5.1. Modeling Uncertainty with State Spaces

The basic model of uncertainty (as in von Neumann–Morgenstern model) assumes that the probabilities of all payoffs are known. In this case, choices are modeled as lotteries—probability distributions. It can be convenient to drop this assumption and describe all relevant uncertainty by a state space. Probabilities need not be specified in this case. A state space Ω should satisfy

Definition: Observability—for each state $\omega \in \Omega$ it should be possible to determine whether or not ω has occurred; sometimes, states can be required to verifiable so that ω can be proven to third parties.

Definition: Mutual exclusivity—no two states can occur together.

Definition: Exhaustiveness—at least one state must occur.

Furthermore, each state should describe the world in sufficient detail to determine the payoff of each relevant action.

Definition: Given a state space Ω and a consumption space X, a contingent prospect f is a function that maps Ω into X. Each f is interpreted as an action that has uncertain payoffs and delivers $f(\omega)$ if ω occurs. Let $\mathcal{F}$ be the domain of all uncertain prospects. The mismatch between real actions and $\mathcal{F}$ goes both ways. It can be difficult to specify Ω and X for real action, most uncertain prospects $f \in \mathcal{F}$ make no practical sense.

3.5.2. Contingent Commodities

- Let $\omega = \mathbb{R}$ describe a stock price. Then buying n units of the stock at price p can be viewed as a prospect $f(\omega) = n(\omega - p)$. Yet $\mathcal{F}$ includes various $g(\omega)$ such as $g(\omega) = \sin \omega$ that have little practical meaning.
- Think about a choice between taking a job in the U.S. or abroad. What are the relevant Ω and X? Too complex to make Ω exhaustive, mutually exclusive, and sufficiently detailed?
- Payoffs, even monetary ones, are often state–dependent. Take life insurance; $\Omega = \{\text{alive}, \text{dead}\}$. If insurance is bought, $f(\text{alive}) = -10K$, $f(\text{dead}) = 500K$. But is it the same kind of money?

3.5.3. Subjective Expected Utility

Given Ω and X, let $\succeq$ be the preference over the corresponding $\succeq$.

Definition: A common model for $\succeq$ is subjective expected utility

$$U(f) = \sum_{\omega \in \Omega} \pi(\omega) u(f(\omega)),$$

where $\pi(\omega) \in [0, 1]$ and $\sum_{\omega} \pi(\omega) = 1$.

The key axiom of this model is the Sure–Thing Principle, that replaces the independence axiom.

Definition: Given prospects f, event E, and payoff x, let fEx be the composite prospect that pays $f(\omega)$ if $\omega \in E$ and x otherwise. The Sure–Thing Principle states

$$fEx \succeq gEx \Rightarrow fEy \succeq gEy.$$

To make payoffs state invariant, monotonicity is required.

Definition: Monotonicity is $x \succeq y \Rightarrow xEy \succeq y$.

3.5.4. Subjective Beliefs

$$U(f) = \sum_{\omega \in \Omega} \pi(\omega) U(f(\omega))$$

The probability measure π reflects the agents subjective beliefs. The term subjective means that the probabilities $\pi(\omega)$ are attached by individuals rather than objectively through some mathematical formulas. This makes the model more general, but provides no clue about how agents should select π. Yet π is uniquely determined by $\succeq$ if payoffs are state invariant. If payoffs are not state invariant, then $\succeq$ is not sufficient to determine π uniquely.

Example: Let $\Omega = \{I, d\}$.

$$U(f) = 0.95\sqrt{f(I)} + 0.05 f(d) = 0.99(\frac{0.95}{0.99}\sqrt{f(I)} + 0.01(5f(d)).$$

There is not enough information here to figure out π.

3.5.5. A Pure Exchange Economy with Uncertainty

The pure exchange economy with uncertainty has the following components

- A population $I = \{1, \dots, I\}$ (a continuum is acceptable).
- A finite state space Ω.
- L deterministic goods consumed in non–negative bundles $(x_{1\omega}, \dots, x_{L\omega}) \in \mathbb{R}^L_+$ contingent on each $\omega \in \Omega$.
- Utility functions $U_i(x_i)$ defined over state–contingent consumption bundles

$$x_i = (x_{11i}, \dots, x_{L1i}, \dots, x_{1\omega i}, \dots, x_{L\omega i}, \dots, x_{1\Omega i}, \dots, x_{L\Omega i}) \in \mathbb{R}^{L\Omega}_+.$$

- Initial endowments $w_i \in \mathbb{R}^{L\Omega}_+$ that are state–contingent as well.

The utility functions U_i will be assumed to have the expected utility form.

Definition: Beliefs are called objective if $\pi_i(\omega) = \pi_j(\omega)$ for all i, j, and ω; and subjective otherwise.

3.5.6. The Arrow–Debreu Equilibrium

Assume that each contingent commodity $l\omega$ is tradable in the market before uncertainty is resolved. This assumption means that markets are complete. Arrow–Debreu equilibrium is a a Walrasian competitive equilibrium in the pure–exchange economy with state–contingent consumptions $x_i \in \mathbb{R}^{L\Omega}_+$ for all $i = 1, \dots, I$ and a price vector $p \in \mathbb{R}^{L\Omega}_{++}$ such that

- Each x_i is affordable: $p \cdot x_i \leq p \cdot w_i$.
- x_i is individually optimal: $x \succeq_i y_i$ for all y_i such that $p \cdot y_i \leq p \cdot w_i$.
- The allocation $x_1, \dots, x_I \in \mathbb{R}^{L\Omega}_+$ is feasible: $\sum_i x_i = \sum_i w_i$.

This is obviously a special case of Walrasian equilibrium and satisfies the first welfare theorem. The core and the Aumann Theorem can be defined as well!

3.6. Subjective Beliefs and Ambiguity Aversion

3.6.1. Implications for the 2×2 Pure–Exchange Economy

Arrow–Debreu Equilibrium can be used to make some predictions about trade with uncertainty. Let $L = 1$. Assume that all beliefs $\pi(\omega) > 0$ to avoid zero prices for some contingent commodities.

- If all agents are risk averse, beliefs are objective, and there is no aggregate uncertainty ($w = \sum_i w_i$) is constant, then Pareto efficiency implies full insurance for all agents.
- If all agents are risk averse, beliefs are objective, and there is aggregate uncertainty, then the market prices will be inversely correlated with the aggregate endowment vector w.
- If all agents are risk averse, beliefs are subjective, and there is no aggregate uncertainty, then agents are not fully insured in the Pareto efficient allocations. The agents will make bets with each other driven by the difference in their beliefs.

3.6.2. 2×2 Pure–Exchange Economies with Objective Beliefs

Let $I = 2$, $\Omega = 2$, $L = 1$. Consumption is one–dimensional.

$$U(x_{1i}, x_{2i}) = \pi u_i(x_{1i}) + (1 - \pi) u_i(x_{2i}).$$

Assume risk aversion so that $u_i'(x)$ is a strictly decreasing function. Assume $\pi \in (0, 1)$. Beliefes are the same for $i = 1$ and $i = 2$.

Case 1: There is no aggregate uncertainty ($w = \sum_i w_i$ is constant). Pareto efficiency implies full insurance for all agents. In the Arrow–Debreu equilibrium

$$\frac{\pi}{p_1} = \frac{1 - \pi}{p_2}.$$

Market prices reflect objective probabilities!

Case 2: There is aggregate uncertainty $w_{11} + w_{12} > w_{21} + w_{22}$. Pareto efficiency implies $x_{1i} > x_{21}$ for $i = 1, 2$. In the Arrow–Debreu equilibrium

$$\frac{\pi}{p_1} > \frac{1 - \pi}{p_2}.$$

The average returns of the first contingent commodity is higher than the average returns of the second contingent commodity. This is the general equilibrium version of the CAPM model. It says that assets have a higher correlation with the market (i.e. the aggregate endowment vector $w_{11} + w_{12}, w_{21} + w_{22}$) have higher expected rates of return. The reason is that the extra consumption in the good state has a smaller marginal value than in the bad state. Market prices do not reflect objective probabilities any more.

3.6.3. 2×2 Pure–Exchange Economies with Subjective Beliefs

Let $I - 2$, $\Omega = 2$, $L = 1$. Consumption is one–dimensional.

Assume risk aversion so that $u_i'(x)$ is a strictly decreasing function. Beliefs are subjective $\pi_1 > \pi_2$. There is no aggregate uncertainty; $w = \sum_i w_i$ is constant. Pareto efficiency implies betting: $x_{11} > x_{21}$ and $x_{22} > x_{12}$. The first agent gets more consumption in state 1 and the second agent gets more consumption in state 2.

Puzzle: any difference in beliefs should produce betting if markets are complete and agents maximize subjective expected utility.

3.6.4. The Ellsberg Paradox

Choose between two bets

$$\begin{cases} \$1,000 & \text{Tails} \\ \$0 & \text{Heads} \end{cases} \quad \text{or} \quad \begin{cases} \$1,000 & \text{if the oil price exceeds \$50 in one year,} \\ \$0 & \text{if the oil price is less than \$50 in one year.} \end{cases}$$

Repeat for

$$\begin{cases} \$0 & \text{Tails} \\ \$1,000 & \text{Heads} \end{cases} \quad \text{or} \quad \begin{cases} \$0 & \text{if the oil price exceeds \$50 in one year,} \\ \$1,000 & \text{if the oil price is less than \$50 in one year.} \end{cases}$$

Agents often choose to bet on the coin in each case. Within subjective EU model, or any model of choice behavior with a unique probabilistic belief,

$$\pi(T) > (\text{oil} > \$50) \text{ and } \pi(H) > \pi(\text{oil} < \$50),$$

but $\pi(T) + \pi(H) = \pi(\text{oil} > \$50) + \pi(\text{oil} < \$50) = 1$.

The tendency to bet on familiar events, such as the one in the Ellsberg Paradox, is called ambiguity aversion. Unlike risk aversion, it cannot be modeled by changing attitudes towards money. Instead, one needs to adapt the notion of subjective beliefs.

Definition: Maximum expected utility representation, or the multiple priors model, has the form

$$U(f) = \min_{\pi \in \Pi} \sum_{\omega \in \Omega} \pi(\omega) u(f(\omega)),$$

where the set Π is a convex and closed set of subjective probabilistic scenarios.

Example: Maximum Expected Utility
If $\Pi = \{\pi : \pi(T) = \frac{1}{2}, \pi(\text{oil} > \$50) \in [0.4, 0.6]\}$, then

$$\begin{cases} \$1,000 & \text{Tails} \\ \$0 & \text{Heads} \end{cases} \quad \succ \quad \begin{cases} \$1,000 & \text{if the oil price exceeds \$50 in one year,} \\ \$0 & \text{if the oil price is less than \$50 in one year.} \end{cases}$$

$$\begin{cases} \$0 & \text{Tails} \\ \$1,000 & \text{Heads} \end{cases} \quad \succ \quad \begin{cases} \$0 & \text{if the oil price exceeds \$50 in one year,} \\ \$1,000 & \text{if the oil price is less than \$50 in one year.} \end{cases}$$

Example: Let $I = 2$, $\Omega = 2$, $L = 1$, and consumption be one–dimensional.

$$U(x_{1i}, x_{2i}) = \min_{\pi \in \Pi_i} \pi u_i(x_{1i}) + (1 - \pi) u_i(x_{21})$$

Assume risk aversion so that $u_i'(x)$ is a strictly decreasing function.

- The sets of scenarios are distinct: $\Pi_1 \neq \Pi_2$.
- There is no aggregate uncertainty: $w = \sum_i w_i$ is constant.
- Pareto efficiency implies full insurance whenever $\Pi_1 \cap \Pi_2 \neq \emptyset$.

Reading: Billot, Chateauneuf, Gilboa, Tallon, Sharing Beliefs: Between Agreeing and Disagreeing, Econometrica, 2000.

Example: Another Puzzle for Subjective Beliefs
Two agents $I = 2$ can duel or not. The state space is $\Omega = \{\omega_1, \omega_2\}$. $\pi(\omega_1) = 0.9$ and $\pi_2(\omega_2) = 0.9$. The utility of winning is 1, the utility of losing is -5, and the status quo is 0.

- Not dueling is Pareto dominanted by dueling.
- But the sum of payoffs without the duel Pareto dominates the sum of payoffs with the duel.

In conclusion, when utility functions include subjective beliefs, Pareto efficiency may produce some unreasonable advice.

Reading: Gilboa, Samuelson, and Schmeidler, No–Betting Pareto Dominance, Econometrica, 2014.

Reading: Rigotti and Shannon, Sharing Risk and Ambiguity, Journal of Economic Theory, 2005.

3.7. Auctions with Private Values

Auctions can be used to sell unique items. The items for sale need not be divisible or generic. Unlike market equilibria, the definition of an auction describes a procedure rather than the precise outcome as a function of the primitives. There are many auction procedures. Some allocation is always produced.

Open–bid Auctions:

Definition: In English auctions, bidding price starts at 0 and goes up in small increments until there is one bidder. It is publicly observed when other bidders quit.

Definition: In Dutch auctions, bidding price starts at a very high price and goes down until there is one bidder who is willing to pay.

Closed–bid Auctions:

Definition: In first–price auctions, hidden bids are sent to the auctioneer. The highest bid wins and is the price.

Definition: In second–price auctions, hidden bids are sent to the auctioneer. The highest bid wins and the price is the second highest bid.

Definition: In all–pay auctions, hidden bids are sent to the auctioneer. The highest bid wins and all agents pay their bid.

3.7.1. Equilibria in Auctions with Private Values

Bidder i is assumed to have quasi–linear utility

$$U_i = \begin{cases} V_i - p(\text{win}), & \text{if she wins,} \\ -p(\text{lose}), & \text{if she does not win,} \end{cases}$$

where V_i is the reservation value and p is the price. Bidder i has non–negative utility from winning if $V_i \geq p(\text{win})$ and has negative utility from winning if $V_i < p$. The latter situation is called winner's curse. It can happen if V_i is uncertain at the time of bidding.

Reservation value is called private if it does not depend on the values or signals received by other bidders and interdependent otherwise. The assumption of private values is plausible when the auction prize has transparent quality and is meant for private consumption. Reservation values may be interdependent if bidders plan to resell the good(s) later, or believe that other bidders are better informed about the quality of the good. In the extreme case, the values are common for all bidders, but uncertain at the time of bidding.

If all bidders $i = 1, \ldots, I$ are assumed to have quasi–linear utility, and the money is paid to the seller (agent 0) with $U_o = R$ revenue, then Pareto efficiency requires requires that $\sum_{i=o}^{I} U_i = V_{\text{winner}}$. Thus, the Pareto efficient outcome requires that *the winner has the highest value of all participants*. The payments do not matter. Some sellers (government) may care about Pareto efficiency while other sellers (private) may care more about revenue.

3.7.2. English and Second–price Auctions with Private Values

Assuming that all V_i are private, then the English auction and the second–price auction have equilibria in weakly dominant strategies.

- English auction: *each bidder quits when* $\boxed{p = V_i}$

- Second–price auction: *each bidder makes a bid* $\boxed{B_i = V_i}$

  Proof. Let $M = \max_{k \neq i} B_k$. If $M < V_i$ then $U(\text{win}) = V_i - M > 0 = U(\text{lose})$, and the best outcome (win) can be obtained by $B_i = V_i$. If $M > V_i$ then $U(\text{win}) = V_i - M < 0 = U(\text{lose})$, and the best outcome (lose) can be obtained by $B_i = V_i$. If $M = V_i$ then $U(\text{win}) = U(\text{lose}) = 0$, and all strategies are equally good. Thus $B_i = V_i$ delivers the best outcome in all possible cases. □

The equilibrium outcomes are the same and Pareto efficient: *the highest value wins* and *the price is equal to the second highest value.*[8]

3.7.3. Dutch and First–price Auctions with Private Values

The Dutch auction and the first–price auction do not have weakly dominant strategies: the best bidding strategy for agent i depends on the strategies of the other agents.The two auctions are strategically equivalent: they can be modeled by the same game. There is the same pool of bidders and strategy $B_i(V_i)$ for agent i. In the Dutch auction, $B_i(V_i)$ is the price when agent i is willing to stop the auction. The outcome of the first–price auction is the same because the highest bid wins and equals the winning price. Thus, the equilibria must be identical in these two auctions. The two auctions can be analyzed through Bayesian Nash Equilibria.

Assumption: It is common knowledge that each private value V_i has an independent identical distribution with CDF F(x), where $F(x) = Prob\{V_i \leq x\}$. Consider finding a symmetric equilibrium bidding function $B(V_i)$. The function B is the same for all i by symmetry. The function B is assumed to be strictly increasing. All of these assumptions are justified: they do hold for the equilibrium we find.

If agent i "pretends" that his value is x rather than V_i, and each other j bids $B(V_j)$, then

$$U_i(x) = (V_i - B(x))\text{Prob}\{B(V_j) \leq B(x) \text{ for all } j \neq i\}$$
$$U_i(x) = (V_i - B(x))\text{Prob}\{V_j \leq x \text{ for all } j \neq i\}$$
$$U_i(x) = (V_i - B(x)) * [F(x)]^{I-1}.$$

In equilibrium, $x = V_i$ should maximize $U_i(x)$. Otherwise, agent i can deviate by bidding $B_i \neq B_i(V_i)$. The first order condition $U_i'(x) = 0$ at $x = V_i$. It follows that

$$U_i'(x) = V_i * [[F(x)]^{I-1}]' - [B(x) * [F(x)]^{I-1}]' = 0.$$

By taking $V_i = x$

$$U_i'(x) = x * [[F(x)]^{I-1}]' - [B(x) * [F(x)]^{I-1}]' = 0.$$

By integrating from 0 to any V

$$\int_0^V x * [[F(x)]^{I-1}]' \, \mathrm{d}x = B(V) * [F(V)]^{I-1} - B(0) * [F(0)]^{I-1}.$$

[8] Theory suggests that if all values are private, then each bidder will make a single bid equal to their value. However, this is not how people behave; typically, they bid more than once, and they make most bids in a short time before the auction ends. This phenomenon is called sniping. Possible explanations are that people are avoiding the sellers price manipulation, people do not like to lose, or values are not private.

Thus the equilibrium bidding function is

$$B(V) = \frac{\int_0^V x * [[F(x)]^{I-1}]' \,\mathrm{d}x}{[F(V)]^{I-1}} = \frac{\int_0^V (I-1)x * F'(x)[F(x)]^{I-2} \,\mathrm{d}x}{[F(V)]^{I-1}}.$$

The outcomes are Pareto efficient.

Given a uniform distribution on $[0, M]$, then $F(x) = \begin{cases} 0, & \text{if } x \leq 0, \\ \frac{x}{M}, & \text{if } x \in [0, M], \\ 1, & \text{if } x \geq M. \end{cases}$

The Nash Bargaining equilibria is

$$\boxed{B(V) = \frac{\int_0^V (I-1)x * x^{I-2} \,\mathrm{d}x}{V^{I-1}} = \frac{I-1}{I} V}$$

Example: Given private values $V_1 = 90$, $V_2 = 20$, and $V_3 = 10$, then the equilibrium outcome depends on the rules.

- In the English and second–price auctions, agent 1 wins and pays $p = 20$.
- In the Dutch and first–price auctions we do not know the outcome unless the distributions of V_i are specified as well. If values V_i are uniform i.i.d., then agent 1 wins and pays $p = \frac{I-1}{I} V_1 = \frac{2}{3} 90 = 60$.

Now, consider if $V_2 = 80$. In the English and second–price auctions price $p = 80$, while price $p = 60$ remains the same in the Dutch and first–price auctions.

3.7.4. All–Pay Auctions with Private Values

Consider finding a symmetric equilibrium bidding function $B(V_i)$. The function B is the same for all i by symmetry. The function B is assumed to be strictly increasing. All of these assumptions are justified: they do hold for the equilibrium we find.

If agent i "pretends" that his value is x rather than V_i, and each other j bids $B(V_j)$, then

$$U_i(x) = V_i * \text{Prob}\{B(V_j \leq B(x) \text{ for all } j \neq i\} - B(x)$$
$$U_i(x) = V_i * [F(x)]^{I-1} - B(x).$$

In equilibrium, $x = V_i$ should maximize $U_i(x)$. Otherwise, agent i can deviate by bidding $B_i \neq B_i(V_i)$. The first order condition $U_i'(x) = 0$ at $x = V_i$. It follows that

$$U_i'(x) = V_i * [[F(x)]^{I-1}]' - B'(x) = 0.$$

By taking $V_i = x$

$$B'(x) = x * [[F(x)]^{I-1}]'$$

By integrating from 0 to any V

$$B(V) = \int_0^V x * [[F(x)]^{I-1}]' \,\mathrm{d}x. = \int_0^V (I-1)x * F'(x)[F(x)]^{I-2} \,\mathrm{d}x$$

If $F(x) = \frac{x}{M}$, then the Nash Bargaining equilibria is

$$B(V) = \frac{\int_0^V (I-1)x * x^{I-2} \,\mathrm{d}x}{M^{I-1}}$$

$$\boxed{B(V) = \frac{I-1}{I}(\frac{V}{M})^{I-1}V}$$

The outcomes in the symmetric equilibria in all basic auctions are Pareto efficient, because the agent with the highest value i wins the auction. The outcome in the English and second–price auctions is Pareto efficient even if the private values are not i.i.d. The Dutch and first–price auctions are not Pareto efficient in the asymmetric case.

Example: (*Krishna pp.50–51*)
Given that V_1 is uniformly distributed on [0,100] and V_2 is uniformly distributed on [0,200].

- Then $B_1(x) > B_2(x)$ and the outcome is not necessarily Pareto efficient.

3.7.5. The Revenue Equivalence Principle

Consider the average revenues for the second–price and first–price auctions when values are i.i.d. $\mathcal{U}[0, M]$ with M=1.

- In the first–price auction, revenue is

$$R = \frac{I-1}{I} * \max_i V_i.$$

The random variable $\max_i V_i$ has CDF $F(x)^I = \text{Prob}\{\max_i V_i \leq x\} = (\frac{x}{M})^I = x^I$. Thus, average revenue is

$$\mathbb{E}(R) = \frac{I-1}{I}\int_0^1 x(F'(x))\,\mathrm{d}x = \frac{I-1}{I}\int_0^1 x(Ix^{I-1})\,\mathrm{d}x = (I-1)\int_0^1 x^I\,\mathrm{d}x = (I-1)\frac{x^{I+1}|_{x=1}}{I+1}$$

$$\boxed{\mathbb{E}(R) = \frac{I-1}{I+1}}$$

- In the second–price auction, revenue R is the second highest bid $B = V_i$ with CDF

$$F(x) = \text{Prob}\{V_i \leq V_{\max_i} \leq x\} = \text{Prob}\{V_i \leq V_{\max_i}\}\text{Prob}\{V_{\max_i} \leq x\}$$

$$F(x) = [I(1-x)]x^I = x^I + Ix^{I-1}(1-x) = Ix^{I-1} - (I-1)x^I.$$

Thus, average revnue is

$$\mathbb{E}(R) = \int_0^1 x\,\mathrm{d}(I * x^{I-1} - (I-1) * x^I) = I * (I-1)[\int_0^1 x^{I-1}\,\mathrm{d}x - \int_0^1 x^I\,\mathrm{d}x$$

$$\boxed{\mathbb{E}[R] = \frac{I-1}{I+1}}$$

These findings are a special case of the general result known as *revenue equivalence*.

Definition: Revenue equivalence principle. Assume that two auctions share population I of agents with i.i.d. private values V_i. Suppose that the Bayesian Nash Equilibria (including weakly dominant as a special case) in these auctions are such that

- the joint distribution of all V_i are the same.
- the equilibrium allocation of the prize is the same with probability 1 (Pareto efficient).
- the average payments of agents of type zero are zero.

Then the average revenues in equilibria are identical. Note that this principle does not establish the ex post equality of the revenues, only the ex ante average equality.

The average revenues do not have to be the same when the allocation of the prize is different. Pareto inefficient auctions often have higher average revenues than Pareto efficient ones (e.g. auctions with reserve prices) A small reserve price will increase the average revenue in all basic auctions. Furthermore the revenue equivalence principle does not hold well in experiments.

Reading: *Kagel(1995), Auctions: A Survey of Experimental Research*
Summary: The English auction provides the closest agreement with game theory and Pareto efficiency. Most subjects realize that they need to quit when the current price exceeds their reservation value and the rules are conducive for rational decisions. The second–price auction produces some overbidding, relative to the dominant strategy. On average, people bid about 10% above their private values. First–price and second–price generate some overbidding relative to the Dutch and English respectively. The first–price auction generates $\approx 5 - 10\%$ higher average revenues than the Dutch. Ranked by approximate revenue (lowest to highest) are; English, Second–price, Dutch, First–price, pay–to–bid, All–pay, and other exotic auctions. The English auction is the closest to Pareto efficiency; $\approx 90\%$ of auctions are won by agents with highest private value, in the Dutch and second–price $\approx 80\%$, and even fewer in the all–pay and exotic auctions. In short, rules matter even when game theory predicts the same outcomes.

3.8. Auctions with Common Values

3.8.1. Equilibria in Auctions with Interdependent and Common Values

Assume that all agents $i = 1, \ldots, I$ receive i.i.d. signals x_i and their ex post value is common: $V = x_1 + \cdots + x_I$. Then neither the Enlish nor the second–price auctions have weakly dominant strategies. The optimal bid depends on the actions of other agents. The winner's curse becomes more prominent; the winner of the auction can pay more than V. Consider a symmetric Nash equilibrium in the English auction with common value V. First, order the signals x_i so that $x_{k_1} \leq x_{k_2} \leq \cdots \leq x_{k_I}$.

- Step 0: Each agent quits at $p = Ix_i$ if nobody has quit before, and x_{k_1} is observed.
- Step 1: After one agent quits, each remaining agent quits at $p = x_{k_1} + (I-1)x_i$.
- Step n: At step n, each remaining agent quits at $p = x_{k_1} + \cdots + x_{k_n} + (I-n)x_i$.

The winner is *the agent with the highest signal*, and the price $p = x_{k1} + \cdots + x_{k_{I-1}} + x_{k_{I-1}}$

Example: A common value English auction with signals $x_1 = 10$, $x_2 = 30$, and $x_3 = 50$.

- Agent 1 quits at $p_1 = 3x_1 = 30$.
- Agent 2 quits at $p_2 = x_1 + 2x_2 = 10 + 2 * 30 = 70$.
- Agent 3 wins and pays $p = 70 < 90 = V$.
  This is a Bayesian Nash equilibrium.
  Proof. Suppose that agent 1 deviates and wins. Then agent 2 will quit when $p > 3x_2 = 90$ and agent 3 will quit when $p = x_2 + 2x_3 = 30 + 2(50) = 130 > V$. Thus it is not optimal for agent 1 to deviate. Suppose that agent 2 deviates and wins. Then agent 1 quits when $p > 3x_1 = 30$ and agent 3 quits when $p = x_1 + 2x_3 = 10 + 2(50) = 110 > V$. Therefore deviation is not optimal and agent 3 remains the winner. □

Consider the second–price auction with common value V. Let $F(\cdot)$ be the CDF of x_i. There is a symmetric Bayesian Nash Equilibrium in the second–price auction where

$$B_i(x_i) = \mathbb{E}[V : \text{second highest signal is } x_i \text{ as well}] = \mathbb{E}[V : x_k \leq x_j = x_i \text{ for all } k, j \neq i]$$

$$B_i(x_i) = 2(x_i) + (I-2)\frac{\int_0^{x_i} x_k F(x_k)' \, \mathrm{d}x_k}{F(x_i)}.$$

Example: 2$^{\text{nd}}$ Price Auction with Common Values
Let $x_1 = 100$, $x_2 = 300$, and $x_3 = 500$ be realized from $F(\cdot) \sim \mathcal{U}[0, 1000]$. Then

$$B_i(x_i) = 2x_i + (I-2)\frac{\int_0^{x_i} x_k(1/1000)\,\mathrm{d}x_k}{x_i/1000} = 2x_i + (I-2)(\frac{1}{x_i})(\frac{x_k^2|_{x_i}}{2} - 0)$$

$$\boxed{B_i(x_i) = 2x_i + (I-2)\frac{x_i}{2}}$$

- $B(x_1) = 2(100) + 50 = 250$, $B(x_2) = 2(300) + 150 = 750$, $B(x_3) = 2(500) + 250 = 1250$.
- Agent 3 wins and pays $p = 750$.

Winner's curse is possible. Consider if $x_1 = 100$, $x_2 = 420$, and $x_3 = 500$. Now

- $B(x_2) = 2(420) + 210 = 1050$.
- Agent 3 wins and pays $p = 1050 > V = 100 + 420 + 500 = 1020$.

The revenue equivalence principle does not hold if signals are not i.i.d. Pareto efficiency holds automatically in the case of common–value. However, it can be lost as well with interdependent values and affiliated signals. If signals are affiliated (a strong form of positive correlation), then $\mathbb{E}(R_{\text{English}}) \geq \mathbb{E}(R_{\text{Second–price}}) \geq \mathbb{E}(R_{\text{First–price}})$.

3.8.2. The Drainage Tract Model

Suppose some bidders are insiders (informed), others are not. Suppose agent 1 (the neighbor) knows V, all others $i = 2, \ldots, I$ just know the CDF $F(\cdot)$ of V. Take $I = 2$ and uniform distributions.

- According to Milgrom's Theorem 5.3.1. at equilibrium, the neighbor and non–neighbor bidders both bid $B(s) = \frac{1}{s}\int_0^2 v(r)\,\mathrm{d}r$. The non–neighbor receives an expected profit of zero.
- $B_n(V) = \frac{V}{2}$ and $B(s)_{nn} = \frac{s}{2}$ is optimal in the first–price auction with private values.

The Drainage tract model predicts; insiders get positive average profits, uninformed agents get average profits of zero, the number of uninformed agents does not affect the average revenue, the bids of uninformed agents are not correlated with the value. For further reading see Hendricks, Porter and Wilson, *Auctions for Oil and Gas Leases with an Informed Bidder and a Random Reservation Price*, (1994).

3.9. Auctions with Many Units

3.9.1. Multi–Unit Auctions

Let I agents compete for $N \leq I$ items. First, suppose that each agent needs just one item and has private value V_i. Multi–unit sealed–bid auctions allow each agent to make several bids of (k, b) type. Note that each bid (k, b) is equivalent to k bids of value b. So, it is convenient to assume that each bidder i can make several bids $b_{i1}, \ldots, b_{ik}$. The K bids satisfying $b^i_K \leq \ldots, \leq b^i_1$ indicate how much a bidder is willing to pay for each additional unit. If $b^i_k > b^i_{k+1}$, then at any price p lying above b^i_{k+1} and below b^i_k bidder i is willing to buy exactly k units. A bid vector (k, b) can be thought of as an "inverse demand function" and can be inverted to obtain i's demand function

$$d^i \equiv \max\{k : p \leq b^i_k\} : \mathbb{R}_+ \to \{1, 2, \ldots, K\}.$$

A total of $N \times K$ bids are collected and the K units are awarded to the K highest of these bids—that is, if bidder i has $k \leq L$ of the K highest bids, then i is awarded k units. The implicit allocation rule may be framed in conventional supply and demand terms. The aggregate demand function determines how many units are demanded in all at different prices. The auctioneer forms a demand curve

$$D(p) = \sum_{(k,b):b \geq p} k.$$

Since the number of units to be sold is fixed, the supply function is vertical. Denote c^{-i} the K–vector of highest competing bids facing bidder i. The residual supply function S^{-i} facing bidder i is

$$S^{-i}(p) = K - \max\{k : c^{-i}_k \geq p\}.$$

3.9.2. The inefficiency of equilibia in the uniform–price auctions

In a uniform–price auction all K units are sold at a "market–clearing" price p such that the total amount demanded is equal to the total amount supplied. The price is the largest price where $D(p) \geq N + 1$, and is equal to the highest losing bid

$$p = \max_i \{b^i_{k^i+1}\}.$$

The uniform–price auction reduces to a second–price sealed–bid auction when there is only a single unit $K = 1$ for sale. Similarly, consider a uniform–price auction where each agent can demand at most one item. Agent i's utility is

$$U_i = \begin{cases} V_i - p(\text{win}), & \text{if she wins one or more item,} \\ 0, & \text{if she does not win any items.} \end{cases}.$$

Let V_i be private and hence, unaffected by other V_j. Then it is a weakly dominant strategy for each i to bid $B_i = V_i$

Proof. Let M be the N^{th} highest bid across all $j \neq i$. If $V_i \geq M$, then agent i's best outcome is to win an item at price M; bidding $B_i = V_i$ will do that. If $V_i < M$, then agent i's best outcome is to win no items; bidding $B_i = V_i$ will do that. □

Now, assume that each agent can demand more than one item, but that there are no complementarities. Agent i's utility is

$$U_i = \begin{cases} V_i(q) - q \times p(\text{win}), & \text{if she wins } q, \\ 0, & \text{if she does not win any items.} \end{cases}.$$

The sequence of marginal utilities $V_{i1} = V_i(1)$, $V_{i2} = V_i(2) - V_i(1), \ldots$ is declining. The utility depends on the quantity k, but not on the combination of items. Claim: Bidding $B_{ik} = V_{ik}$ is no longer dominant. When a bidder wants to buy more than one unit and when the units have declining marginal values, a bidder generally has an incentive to reduce her demand, that is, to bid less than her value for some units.

Example: (*Milgrom pp.258–259*) Suppose there are 2 bidders and two units q for sale. Bidder 1 demands only a single unit and has value $V_1(q) = v_1$ for all $q \geq 1$, where v_1 has a uniform distribution on $\mathcal{U}[0,1]$. Bidder 2 has demand for two units, the first unit is worth v_{21} and the second is worth v_{22}, where $0 < v_{22} \leq v_{21} < 1$.

- Bidder 1 has a weakly dominant strategy to bid $b_{11} = v_1$ and $b_{12} = 0$.
- With two units for sale, bidder 2 is assured of winning at least one item if she places a positive bid. Her expected payoff is

$$U_2 = \mathbb{E}[(v_{21} + v_{22} - 2b_{11})\mathbb{1}_{\{b_{11}<x\}} + (v_{21} - x)\mathbb{1}_{\{b_{11}>x\}}$$

$$U_2 = \int_0^x (v_{21} + v_{22} - 2b_{11})\,\mathrm{d}b_{11} + \int_x^1 (v_{21} - x)\,\mathrm{d}b_{11}$$

$$U_2 = (v_{21} + v_{22})x - x^2 - 0 + (v_{21} - x) - (v_{21} - x)(x)$$

$$U_2 = v_{21} + v_{22}x - x.$$

  This function is maximized at $x = 0$; *the optimal bid for the second unit is zero.* Bidder 2 always finds it optimal to bid as if she had demand for only one unit, regardless of her actual values. This is an example of demand reduction, where Bidder 2 bids less than their value v_{22} on the second item. Note that the outcome is not Pareto efficient; suppose that $v_{22} > v_1$, then it is Pareto optimal to allocate both items to agent 2.

Other possible rules are; the auctioneer sells each item seperately at a second–price or a first–price auction. There are potential issues;

- The declining price anomaly is when price tends to decrease from one item to the next. See
- The second–price auction produces variable prices, but also reveals the discrepancy between the highest bid and the second–highest bid.
- Non–monotonicity: may occur wherein agents who bid more may pay less because of weaker competition.

These features can be problematic because participants and sellers can be put in an awkward situation in their firms/institutions. The uniform–price auction does not have this problem because all participants pay the same price.

3.9.3. Multi–Unit Auctions with Complementarities

Complementarities arise when agents attach a value $V(C)$ for each combination[9] of items and $V(C) \neq \sum_{k \in C} v_k$. When items are complementary, Pareto inefficiencies are prevalent.

Example: Two licenses are auctioned to three companies in separate auctions. Values are

$$\begin{array}{lll} V_1(x) = 100 & V_1(y) = 200 & V_1(xy) = 300 \\ V_2(x) = 200 & V_2(y) = 100 & V_2(xy) = 300 \\ V_3(x) = 0 & V_3(y) = 0 & V_3(xy) = 600 \end{array}$$

The auctions will allocate x to agent 2 and y to agent 1, unless agent 3 chooses to bid more than 200 in each auction (which is risky for her). This allocation is Pareto inefficient, because allocating x and y to agent 3 can cause a Pareto improvement.

Example: Two licenses can be auctioned to three companies only as a bundle. Values are

$$\begin{array}{lll} V_1(x) = 100 & V_1(y) = 100 & V_1(xy) = 400 \\ V_2(x) = 100 & V_2(y) = 200 & V_2(xy) = 300 \\ V_3(x) = 300 & V_3(y) = 300 & V_3(xy) = 300 \end{array}$$

Agent 1 wins and the outcome is Pareto inefficient. Allocating y to agent 2 and x to agent 3 can cause a Pareto improvement.

It appears that the optimal auction should make all agents to submit bids for all possible combinations.

[9] See *Milgrom pp.278–279* for evidence that the bids for smaller combinations are disproportionately low.

3.10. Vickrey–Clarke–Groves Mechanisms

3.10.1. Vickrey–Clarke–Groves Framework

There is a population of agents $I \cup \{0\}$ where agent 0 collects payments and pays out transfers. In applications, agent 0 is usually a government institution or a seller in auctions. $X = Z \times \mathbb{R}^I$ is the feasible consumption space for agents $i \in I$ with quasi–linear utility functions $U_i(z, p_i) = V_i(z) - p_i$. The vector $(z, p_1, \ldots, p_I)$ means that the allocation z is chosen and all agents pay $p_1, \ldots, p_I$ to agent 0. Agent 0 has utility $U_0 = \sum_{i=1}^I p_i$ with only a monetary component, and is otherwise unaffected by z. The Pareto efficient outcome requires $\sum_{i \in I} V_i(z)$ to be maximized. What happens if the auctioneer asks the agents directly about V_i? There will be incentives to misrepresent. Suppose that all agents are asked directly about their values $V_i(z)$, and then the sum $\sum_{i \in I} \hat{V}_i(z)$ is maximized based on their reports $\hat{V}_i(z)$. Then the agents will have strong incentives to misrepresent their value functions (i.e. $\hat{V}_i \neq V_i$

Example: Consider dualistic public good provision $Z = \{0, 1\}$. If not provided $V_i(0) = 0$.

- If the auctioneer asks all agents to report $V_i = V_i(1) \in [-M, M]$, then it is a *weakly dominant strategy* to report $\hat{V}_i = M$ if $V_i \geq 0$ and report $\hat{V}_i = -M$ if $V_i < 0$. The equilibrium will be $z^* = 1$ if $\#\{i : \hat{V}_i \geq 0\} > \#\{i : \hat{V}_i < 0\}$. This outcome can be far from Pareto efficient.

Vickrey–Clarke–Groves (VCG) mechanisms are designed to make it a weakly dominant strategy to report the function V_i truthfully. VCG mechanisms eliminate the incentives to lie by imposing additional payments on all agents $i \in I$. These payments are called externality taxes[10]. A special agent 0 (e.g. the government) is required.

- Step 1: Let all agents $i \in I$ report the function $\hat{V}_i(z)$. Choose $z^* \in Z$ that maximizes $\sum_{i=1}^I \hat{V}_i(z)$. Agent 0 is passive and does not report anything.
- Step 2: Let every agent k pay a tax

$$p_k = \max_{z \in Z} \sum_{j \neq k \in I} \hat{V}_j(z) - \sum_{j \neq k \in I} \hat{V}_j(z^*).$$

Agent k is called pivotal if z^* does not maximize $\sum_{j \neq k \in I} \hat{V}_j(z)$. The pivotal agent affects the social choice for the rest of the population $I \setminus k$. Pivotal agents pay $p_k > 0$.

Theorem: In the VCG mechanism, each agent $i \in I$ has a weakly dominant strategy to report $\hat{V}_i = V_i$. The outcome in this weakly dominant strategy equilibrium is Pareto efficient. Moreover, all such mechanisms have VCG form with $p_k = P(\hat{V}_{-k} - \sum_{j \neq k \in I} \hat{V}_j(z^*)$.

Proof. Agent k is free to report any $\hat{V}_k$. The mechanism selects z^* that maximizes $\sum_i \hat{V}_i(z)$. Note that z^* depends on $\hat{V}_k$. Utility

$$U_k(\hat{V}_k, \hat{V}_{-k}) = V_k(z^*) - p_k = V_k(z^*) + \sum_{j \neq k \in I} \hat{V}_j(z^*) - P(\hat{V}_{-k}),$$

is maximized if $V_k(z^*) + \sum_{j \neq k \in I} \hat{V}_j(z^*)$ is maximized. This occurs at $\hat{V}_k = V_k$. □

[10] Externality taxes may be combined with lump–sum subsidies as well.

VCG mechanisms can be used

- to choose the optimal quantities of public goods; let $V_i(Q)$ be the value function and proceed accordingly.
- to design auctions; the second–price auction is a special case for a one–unit auction.
- in the multi–unit case, VCG mechanism can be applied in the case of perfect substitutes $V_i(q) = V_{i1} + \cdots + V_{iq}$ and complements where $V_i(C)$ depends on the combination C. When bids are made for combinations of goods, such auctions are called combinatorial.

What can go wrong?

- Agent 0 can be hard to select. All payments made to agent 0 cannot be immediately redistributed because it would change the incentives to tell the truth. If agent 0 can run many VCG mechanisms then redistribute all procees across a large population through lump–sum payments, then it may work. in the auction setting, it is natural to assume that the seller can collect payments.
- Bidders may refuse to participate.
- The mechanism is open to group manipulation.
- The mechanism does not maximize the seller's revenue.
- The revenue is not monotonic with respect to the number of bidders, there can be incentives to add shill bidders.
- The function $\hat{V}_i$ can be very complex to report. If the agents fail to report some of their values, the outcome need not be Pareto efficient.

3.11. General Mechanisms

3.11.1. Mechanisms

Mechanism design starts from the following primitives

- A population of I agents.
- A set X of feasible alternatives. Each $x \in X$ can describe both the consumption of public goods and private consumption bundles. Say $x_i \in \mathbb{R}^n_+$ and $x = (x_1, \ldots, x_I)$ such that $\sum_i x_i = \sum_i w_i$.
- For each i there is a set $\Theta_i = \{\theta_i, \ldots\}$ of possible types of agent i. Each possible type θ_i includes information about preferences (utilities), initial endowments, etc.
- For each i, there is a set S_i of possible strategies (actions) that agent i can take in the mechanism. Mechanisms are called direct if $S_i = \Theta_i$, that is, each agent announces some type $\hat{\theta}_i$ as her strategy.
- A function $g : S_1 \times \cdots \times S_I \to X$ that determines the choice $g(s_1, \ldots, s_I) \in X$ for each strategy profile $(s_1, \ldots, s_I)$. If g determines some distribution in X, then X needs to be replaced by $\Delta(X)$—the space of possible distributions in X.

Problem: Given $u_1 = \hat{\alpha}^{\alpha}\hat{\beta}^{1-\alpha}$ and $u_2 = (1-\hat{\alpha})^{\beta}(1-\hat{\beta})^{1-\beta}$, find a mechanism $g(s)$ and a corresponding social choice function $f^*(s)$.

Solution: $g(s) = \{\text{Ask } \hat{\alpha}, \hat{\beta} \text{ and implement Walrasian equilibrium}\}$. It is weakly dominant to report $\hat{\alpha} = 1$, $\hat{\beta} = 0$.

$$g(s) = \{x_{11} = 1,\ x_{22} = 1\}.$$

Therefore, a SCF $f^*(s) = \{\text{Ask } \hat{\alpha} \text{ and } \hat{\beta},\ x_{11} = 1 \text{ and } x_{22} = 1 \text{ for all } \hat{\alpha}, \hat{\beta}\}$.

Problem: First–price auction with Uniform Distribution
$f(\theta) = \{\text{Highest } V_i \text{ receuves good at some price}\}$. What is a mechanism $g(s)$ and a corresponding SCF $f^*(s)$?

Solution:

$$g(s) = \{\text{Highest } s_i = b_i \text{ receives the good at price } b_i\}$$

$$f^*(s) = \{\text{Ask } \hat{V}_i \text{ and set } b_i = \frac{I-1}{I}\hat{V}_i, \text{ then follow the first–price auction with } b_i\text{'s}\}$$

3.12. Social Choice Functions

3.12.1. Social Choice Functions

Unlike a mechanism, a social choice function (SCF) describes the social choice $x \in X$ as a function of individual true types. Social choice functions are mappings $f : \Theta_1 \times \cdots \times \Theta_I \to X$. The outcome $f(\theta_1, \ldots, \theta_I)$ is obtained when the true types are θ_i.

Definition: A SCF is called Pareto efficient if $f(\theta_1, \ldots, \theta_I)$ is Paretio efficient in X when individual preferences are $\succeq_{\theta_i}$.

Definition: A SCF is called strategy proof if $f(\theta, \theta_{-i}) \succeq_{\theta_i} f(\theta', \theta_{-i})$ for all θ_i, θ_i', and θ_{-i}.

Each SCF can be interpreted as a direct mechanism f^* where $S_i = \Theta_i$ and

$$f^*(s_1, \ldots, s_I) = f(\theta_1, \ldots, \theta_I).$$

Strategy proofness means that $s_i^*(\theta_i) = \theta_i$ is a weakly dominant strategy in this direct mechanism. If f is not strategy proof, then it is not clear why the mechanism f^* should implement f. Thus, f^* and f are not the same, though they may seem very similar.

3.12.2. Dominant Strategy Implementation

Definition: Say that $s_i^*(\theta_i)$ is a weakly dominant strategy for agent i with type θ_i in the mechanism g if

$$g(s_i^*(\theta_i), s_{-i}) \succeq_{\theta_i} g(s_i', s_{-i})$$

for all s_i' and s_{-i}.

If all i have such weakly dominant strategies for all types θ_i, then the mechanism g has an equilibrium in weakly dominant strategies. The SCF f is implemented by g, or g implements f, in weakly dominant strategies if

$$f(\theta_1, \ldots, \theta_I) = g(s_1^*(\theta_1), \ldots, s_I^*(\theta_I)).$$

3.12.3. Examples of Strategy Proofness with Type Constraints

The second–price auctions and VCG mechanisms are strategy proof when preferences are quasi–linear, and agent zero just collects money. X can be finite if all values and payments are discrete (e.g. quoted in whole dollar numbers).

The Median–peak (median–voter) social choice function with single–peaked preferences on $X \subset \mathbb{R}$, where each $\succeq_i$ has a peak p_i such that

$$x < y \leq p_i \Rightarrow y \succ_i x$$
$$p_i \geq y > x \Rightarrow y \succ_i x$$

and it is strategy proof.

Uniform rules—$X = \{(x_1, \ldots, x_I) \in \mathbb{R}^I : \sum_{i=1}^I x_i = C\}$—when preferences $\succeq_i$ are single–peaked over x_i are strategy proof.

3.13. Strategy Proofness in Practice

3.13.1. The Gibbard–Satterthwaithe Theorem

The revelation principle implies that strategy proof SCFs are exactly those that can be implemented by some mechanism in weakly dominant strategies. However, The Gibbard–Satterthwaite Theorem shows that it is not easy to combine strategy proofness with Pareto efficiency. First, say that f is dictatorial if there is a $d \in I$ such that $f(\theta_d, \theta_{-d}) \succeq_{\theta_d} y$ for any θ_d, θ_{-d}, and $y \in X$. If $\succeq_{\theta_d}$ is a linear order (complete, transitive, and antisymmetric), then there is only one such $f(\theta_d, \theta_{-d})$.

Theorem: *The Gibbard–Satterthwaithe Theorem*:
Given a population I, feasible set X, types $\Theta = \Theta_1 \times \cdots \times \Theta_I$, and a social choice function $f : \Theta :\to X$, if

- for each i, the class of possible $\succeq_{\theta_i}$ includes all linear orders on X,
- $|X| \geq 3$ is finite,
- f is Pareto efficient (if x is the best in X for all θ_i, then $f(\theta_1, \ldots, \theta_I) = x$),
- f is strategy proof,

then f is dictatorial.

Thus, the theorem states that the five conditions; that types θ_i allow all linear orders, $|X| \geq 3$, f is Pareto efficient, f is strategy proof, and f is not dictatorial, cannot hold together. Pareto efficiency and $|X| \geq 3$ can be relaxed further to the assumption that at least three alternatives can be social choices for some realizations of types. If any single condition is dropped, then examples can be found.

Example: 1. Suppose that $|X| = 2$. Majority voting satisfies that types θ_i allow all linear orders, f is Pareto efficient, f is strategy proof, and f is not dictatorial.

Example: 2. Now suppose $|X| \geq 3$; pick two alternative exogenously, then run majority voting. The outcome is Pareto inefficient, but satisfies that types θ_i allow all linear orders, $|X| \geq 3$, f is strategy proof, and f is not dictatorial.

Example: 3. The Borda count voting mechanism produces an outcome that satisfies that types θ_i allow all linear orders, $|X| \geq 3$, f is Pareto efficient, and f is not dictatorial, but the Borda count f is not strategy proof.

Example: 4. Suppose that f is dictatorial. A lexicographic dictatorship (where agent 1 chooses first, then agent 2 chooses among all options acceptable to agent 1, etc ...) produces an outcome that satisfies that types θ_i allow all linear orders, $|X| \geq 3$, f is Pareto efficient, and f is strategy proof.

Example: 5. ***(Median–Voter Mechanism)*** Suppose that all agents have single–peaked preferences on $X \subset \mathbb{R}$. This assumes that types θ_i exclude some linear orders. Let p_i be the peak of agent i and

$$f(p_1, \ldots, p_I) = \text{median}(p_1, \ldots, p_I) = p_m,$$

such that

$$\#\{i : p_i \leq p_m\} \geq \frac{I}{2} \text{ and } \#\{i : p_i \geq p_m\} \geq \frac{I}{2}.$$

Then f is strategy proof and implemented in weakly dominant strategies by the direct mechanism f^*. An interesting feature for political scientists is that p_m cannot be strictly beaten by any other alternative in majority voting. Such p_m is called a Condorcet winner. In general, Condorcet winners need not exist, as shown by the Condorcet Paradox.

3.13.2. Maskin Monotonicity

Say that f is Maskin monotonic if for all $x \in X$, $\theta_1, \ldots, \theta_I$, and θ_i' such that the lower contour set at x for type θ_i' includes the lower contour set at x for type θ_i,

$$x = f(\theta_i, \theta_{-i}) \Rightarrow x = f(\theta_i', \theta_{-i}).$$

The inclusion of the lower contour sets means that for all $y \in X$, then $x \succeq_{\theta_i} y$ implies $x \succeq_{\theta_i'} y$.

Lemma: (*Maskin*): Suppose that all preferences are linear orders. If f is strategy proof, then f is Maskin monotonic.

3.14. Bayesian Implementation

3.14.1. Bayesian Implementation

Assume that all agents have expected utility with Bernoulli index $u_i(x, \theta_i)$. Suppose that the joint distribution ϕ of types $(\theta_1, \ldots, \theta_I)$ is common knowledge.

Definition: Say that $s_i^*(\theta_i)$ are a Bayesian Nash equilibrium if

$$\mathbb{E}(u_i(g(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_{-i}) \geq \mathbb{E}(u_i(g(s_i', s_{-i}^*(\theta_{-i})), \theta_i)|\theta_{-i}),$$

for all s_i' and θ_i.

The SCF f is implemented by g, or g implements f, in BNE if

$$f(\theta_1, \ldots, \theta_I) = g(s_1^*(\theta_1), \ldots, s_I^*(\theta_I)).$$

In many contexts, economists have to use Bayesian Implementation because

- they study some exogenously given mechanism (e.g. first–price, Dutch, all–pay auctions, bargaining models, etc.).
- they want to implement some social choice function that is not strategy proof.

Example: Bargaining application
Suppose that two agents, the seller $i = 1$ and the buyer $i = 2$, need to allocate an item that is owned by the seller. Assume that the utility functions are quasi–linear.

- If there is no trade, $U_1 = \theta_1 = V_1$ and $U_2 = 0$.
- If there is a trade at price p, then $U_1 = p$ and $U_2 = \theta_2 - p = V_2 - p$.
- Pareto efficiency dictates that there should be a trade whenever $V_2 \geq V1$.

Problem: VCG Mechanism
One way to run an efficient exchange is to have agent 0 pay a fixed price P to the seller and then run a second–price auction. Let θ_1 and θ_2 be both uniformly i.i.d. distributed on $[0, M]$. What p would be sufficient for all seller types?

Solution:

- Paying M should be enough for all seller types, regardless of the buyer's behavior.
- If the buyer conforms to the equilibrium, the $\frac{M}{2}$ is enough because the seller of type $\theta_1 = M$ expects to pay that much on average at the second–price auction.
- However, the average revenue in the second–price auction is only $\frac{M}{3}$. Thus, the VCG mechanism will always run a deficit if it provides incentives to participate for all types.

Example: Split–the–Difference Mechanism
To avoid deficits, one can use a variety of mechanisms with the Bayesian equilibria. Let both agents make bids B_1 and B_2.

- If $B_1 > B_2$, then there is no trade.
- If $B_2 \geq B_1$, then there is a trade at the price $p = \frac{B1+B2}{2}$.
- Note that if both agents make bids $B_i = \theta_i$, then the outcoe is Pareto efficient. However, this is not a BNE here.

Problem: Split–the–Difference Mechanism
Let θ_1 and $\theta@$ be both uniformly i.i.d. distributed on $[0, M]$. What is a Bayesian Nash equilibrium for this mechanism? Does the equilibrium have a balanced budget? Is the outcome Pareto efficient?

Solution:

- Then

$$B_1^* = \frac{M}{4} + \frac{2}{3}\theta_1$$
$$B_2^* = \frac{M}{12} + \frac{2}{3}\theta_2$$

is a Bayesian Nash equilibrium (check!).
- This equilibrium is budget–balanced, but not Pareto efficient.

3.14.2. The Revelation Principle

Revelation principles are formulated both for dominant strategy implementation and Bayesian implementation. They characterize all social choice functions that can be implemented by some mechanism g. The revelation principle states that if f is implemented by g (in dominant strategies and BNE), then f is also implemented by its direct mechanism f^* (in dominant strategies and BNE) so that $s_i^*(\theta_i) = \theta_i$. As the equilibrium involves truth–telling, the function f is called truthfully implementable. Thus the revelation principle asserts that if f is implementable by some g, then it is also truthfully implementable by its direct mechanism f^*. In the case of dominant implementation, f must be strategy proof.

Theorem: *The Revelation Principle*: Take a social choice function $f : \Theta_i \times \cdots \times \Theta_I \to X$. The following statements are equivalent:
- f is strategy proof.
- f is implemented by the direct mechanism f^* with $\theta_i^* = \theta_i$.
- f is implemented by some mechanism g in weakly dominant strategies $s_i^*(\theta_i)$ such that $f(\theta_1, \ldots, \theta_I) = g(s_1^*(\theta_1), \ldots, s_I^*(\theta_I))$ for all θ_i.

3.15. Arrow's Impossibility Theorem

3.15.1. Arrow's Theorem

Instead of social choice functions, the Arrow Theorem establishes impossibility for social welfare aggregators $F : \Theta \to \mathcal{R}$ where $\mathcal{R}$ is the space of preferences on X.

Theorem: The following conditions cannot hold together
- $|X| \geq 3$.
- $F(\theta)$ is complete and transitive for all θ.
- F is Pareto efficient: if $x \succ_i y$ for all i, then $xP(\theta)y$.
- Independence of Irrelevant Alternatives, which states

$$xP(\theta, X)y \Leftrightarrow xP(\theta, X \,/\, \{z\})y \quad \text{for } z \neq x \text{ and } z \neq y.$$

- F is not dictatorial.

The Gibbard–Satterthwaithe and Arrow Theorem are related.

Proof. If F is possible, then f is possible as well. □

3.16. Voting

3.16.1. Voting and Single–Peaked Preferences

Definition: $D \subset I$ is decisive if there are $x, y \in X$ and types $\theta_1, \ldots, \theta_I$ such that
- x is the best choice for θ_i, $i \in D$.
- y is the best choice for θ_j, $j \notin D$.

- $x = f(\theta_1, \dots, \theta_I)$.

Pareto efficiency implies $D = I$ is decisive.

$I =$	$\underline{A} \cup$	$\underline{B} \cup$	$\neg\underline{D}$
1^{st}	x	y	z
2^{nd}	y	z	x
3^{rd}	z	x	y

Proof. If x is the best choice for $i \in D \Rightarrow x = f(\theta_1, \dots, \theta_I)$. Take any $z \in X$, then z is the best for all $i \in D$, $z = f(\theta_1, \dots, \theta_I)$. Assume D is decisive, $D = A \cup B$ for disjoint groups A, B. Then either A or B is decisive.

- Case 1: $f \neq x$, $f \neq y$.
- Case 2: $f = x$, A is decisive.
- Case 3: $f = y$, B is decisive.

□

If $D = \{d\}$, then d is a dictator.

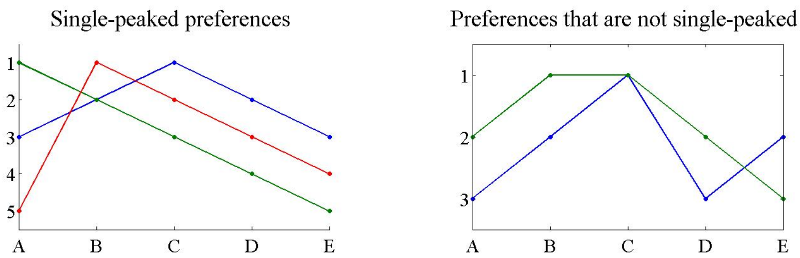

3.16.2. Voting Problems

Definition: The uniform division rule is as follows. It is assumed that X describes all possible divisions of a fixed amount C, where $X = \{(x_1, \ldots, x_I) \in \mathbb{R}^I : \sum_i x_i = C\}$ and preferences are constrained; $\succeq_i$ is single–peaked over x and there are no externalities.

- If $\sum_{i=1}^I p_i = C \Rightarrow x_i = p_i$, then everyone's demand is met.
- If $\sum_{i=1}^I p_i > C$, then we must find a ration r such that $\sum_{i=1}^I \min\{p_i, r\} = C$. There is a unique r that satisfies this condition.
- If $\sum_{i=1}^I p_i < C$, then we must find a quota q such that $\sum_{i=1}^I \max\{p_i, q\} = C$.

This is an example of constrained demand.

Problem: Uniform Division Rule with a Ration
What is the result of the uniform division rule with $C = 10$, $p_1 = 2$, $p_2 = 3$, $p_3 = 4$, $p_4 = 5$?

Solution:

- If $r \leq 2 \Rightarrow \sum_{i=1}^I \min(p_i, r) = 4r < 10$, then the ration r is too low.
- If $r \in (2, 3] \Rightarrow p_1 + 3r = 10$, then agent 1 has her demand met and the others ration

$$2 + 3r = 10$$
$$r = \frac{8}{3}.$$

- If $r \in (3, 4] \Rightarrow 5 + 6 > 10 \Rightarrow$, then the ration r is insufficient.

Thus, $r = \frac{8}{3}$ is the unique ration. The allocation is $x = (2, \frac{8}{3}, \frac{8}{3}, \frac{8}{3})$.

Problem: Uniform Division Rule with a Quota
What is the result of the uniform division rule with $C = 18$, $p_1 = 2$, $p_2 = 3$, $p_3 = 4$, $p_4 = 5$?

Solution:

- If $q \leq 2 \rightarrow \sum_{i=1}^I p_i < 18$, then the quota q is insufficient.
- If $q \in (2, 3] \rightarrow q = 6 \notin (2, 3]$, then the quota is infeasible.
- If $q \in (3, 4] \rightarrow 2q = 9$ $q = 4.5 \notin (3, 4]$, then the quota is infeasible.
- If $q \in (4, 5] \rightarrow 3q = 13$ $q = \frac{13}{3}$.

Thus, $q = \frac{13}{3}$ is the unique quota. The allocation is $x = (\frac{13}{3}, \frac{13}{3}, \frac{13}{3}, 5)$.

Note that we have Pareto efficiency and strategy proofness—no agent wants to consume less and no agent wants to consume more. If a deficit is replaced by a surplus, then $x_i \succeq x_j$ for all $j \neq i$. Also note, there is considerable interest in implementation of weakly dominant strategies—strategy proof mechanisms that reveal an agent's type.

Reading: Y. Sprumont. *The Division Problem with Single–Peaked Preferences: A Characterization of the Uniform Allocation Rule*, Econometrica, (1991)

Definition: The Borda Count is a practical mechanism used in committee work. Let $\succeq_i$ be linear orders $R_i(x) = \#\{y \in X : y \succeq_i x\}$. Rank $R_i = 1$ if x is best for i, $R_i = 2$ if x is 2^{nd} best for i, etc. Add up all the ranks, and pick the lowest one.

$$f(\succeq_i, \ldots, \succeq_I) \text{ minimizes } \sum_{i=1}^I R_i(x).$$

Example: The Borda Count

$I:$	$\underline{1}$	$\underline{2}$	$\underline{3}$
1^{st}	x	x	y
2^{nd}	y	y	x
3^{rd}	z	z	z
4^{th}	w	w	w

$$\sum_i R_i(x) = 4, \quad \sum_i R_i(y) = 5, \quad \sum_i R_i(z) = 9, \quad \sum_i R_i(w) = 17.$$

Thus, x is chosen. Note that Agent 3 could lie about preference on x to influence the outcome for a benefit. Therefore, the Borda Count is not strategy proof—it can be manipulated by lying about bad alternatives.

Example: Fake Pairwise Majority Voting
Take any $x, y \in X$ and run majority voting between x and y. This scheme is not Pareto efficient.

Example: The Median Peak
Given $p_1 = p_2 = p_3 = 0$, $p_4 = 30$, and $p_5 = 50$.

- The median is 0.

Given $p_1 = p_2 = 0$, $p_3 = 60$, $p_4 = 30$, and $p_5 = 50$.

- The median is 30.

The median always exists and does not need tie–breaking if I is odd.

3.17. Bargaining and the Myerson–Satterthwaite Theorem

Consider the two–agent bargaining problem with the seller's and buyer's private values, V_1 and V_2, that have positive density in $[0, M]$. Claim: The following conditions cannot hold together in any Bayesian Nash equilibrium in any mechanism.

- All types have *ex ante* incentives to participate (*ex ante* individual rationality).
- The average revenue of the seller is equall to the average payment of the buyer (*ex ante* balance budget).
- The *ex post* Pareto efficiency trade occurs if and only if $V_2 \geq V_1$.

3.17.1. The Myerson–SatterWaite Theorem

The Myerson–SatterWaite Theorem: Suppose that θ_1 and θ_2 are independently distributed on $[0, M]$. The distributions need not be the same, but must have positive density. Also, agents are assumed to be risk neutral. Suppose that a mechanism g implements a social choice function $f : \Theta \to \{(i, p)\}$. Then the following conditions cannot hold together.

- All types have incentives to participate—their average utility

$$U_i(\text{participating}) \geq U_i(\text{not participating}).$$

- The budget is balanced.
- The equilibrium is Pareto efficient.

The result of the theorem is an argument against the Coase Theorem, where private trade is enough to restore Pareto efficiency from any initial allocation. This is because the Coase Theorem requires perfect information!

Example: VCG mechanism
The central planner (Agent 0) pays a fixed price to the seller, and then runs the second–price auction.

- To make all types participate the price has to be at least $\frac{M}{2}$ if the distributions are uniform.
- The expected revenue is $\frac{M}{3}$.
- Thus, there is an average budget deficit of $\frac{M}{6}$.

Example: Split–the–difference mechanism
The BNE strategies are $B_1^* = \frac{M}{4} + 2\frac{V_1}{3}$ and $B_2^* = \frac{M}{12} + 2\frac{V_1}{3}$.

- The equilibrium is budget-balanced and inefficient.

Example: The direct mechanism that implements the same social choice function; agents announce values B_1 and B_2. If $\frac{M}{12} + 2\frac{B_2}{3} \geq \frac{M}{4} + 2\frac{B_1}{3}$, then there is trade and the price is the average of $\frac{M}{12} + 2\frac{B_2}{3}$ and $\frac{M}{4} + 2\frac{B_1}{3}$.

- Truth–telling is BNE, but the outcome is inefficient because of the rules.

3.18. Matching

3.18.1. Matching Problems

Matching problems are varied and complex (e.g. students and colleges, hostpitals and doctors, marriages, kidney transplants, etc.). Money is allowed to be used, sometimes, but can be restricted or totally banned. Matching problems include assignment problems (where only one–side has preferences), marriage problems (where both sides have preferences and the matching is one–to–one), and admission problems (where both sides have preferences and the matching is not one–to–one).
Take sets M and W. Each m has preferences over $W \cup \{\text{single}\}$. Each w has preference over $M \cup \{\text{single}\}$. A matching is a function $F : M \to W \cup \{\text{single}\}$ such that $F(m) = F(m')$ is possible only if $F(m) = F(m') = \text{single}$. Let $G : W \to M \cup \{\text{single}\}$ be the corresponding function for w.

A matching is stable if
- Each m prefers $F(m)$ to being single. Each w prefers $G(w)$ to being single.
- There is no pair m and w such that m prefers w to $F(m)$ and w prefers m to $G(w)$.

Stability implies Pareto efficiency. Note, stability is not gauranteed in one–sided markets.

Example: Roommate Problem

$$\begin{aligned} &m_1 : m_2 \succ m_3 \succ m_4 \\ &m_2 : m_3 \succ m_1 \succ m_4 \\ &m_3 : m_1 \succ m_2 \succ m_4 \\ &m_4 : \text{anyone.} \end{aligned}$$

3.18.2. Allocation Problems

An allocation problem is a tuple that specifies population I, the set of choices H, and preferences $\succ_i$ over H. An assignment (matching) is a function $a : I \to H$ so that $a(i) = a(j)$ implies $i = j$. Functions such as these are called injective. If $|I| = |H|$, then assignments are one–to–one. An assingment is called Pareto efficient if there is no other b such that all agents prefer $b(i)$ to $a(i)$ and some do it strictly.

Two kinds of mechanisms that achieve Pareto efficiency are
- Top trading cycles mechanisms (Gale 1962).
- Serial choice (serial dictatorship, priority) mechanisms.

3.18.3. Assignment Problems

The number of all possible assignments is $N!$ and grows very fast. An assignment $a : I \to N$ is Pareto inefficient if there is $b : I \to N$ such that
- $U_i(b(i)) \geq U_i(a(i))$ for all i and $U_i(b(i)) > U_i(a(i))$ for some i.

If no such b exists, $a : I \to N$ is efficient.

3.18.4. Top Trading Cycles

Consider a house allocation problem. Suppose that

- The number of agents is the same as the number of houses, $I = H$.
- $a(i)$ is the initial assignment.
- All ranking are strict, so that each agent i is never indifferent between two distinct options:

$$U_i(h_m) \neq U_i(h_n) \text{ whenever } m \neq n.$$

Let every agent i pick her best choice $B(i)$. Say that $i_1, i_2, \ldots, i_k$ is a top trading cycle if $B(i_n) = a(i_{n+1})$ for all $n = 1, \ldots, k-1$ and $B(i_k) = a(i_1)$. Such cycles must exist.

Proof. Assign the houses in each cycle and take the agents and houses in each TTC off the market. Repeat for the remaining agents and houses. Continue until all agents are assigned a house. □

The TCC mechanism is Pareto efficient.

Proof. Suppose that $t(i)$ is the assignment produced by the TTC and $b(i)$ is a Pareto improvement. The members of any TTC in the first round get their top choices. Thus $t(i) = b(i)$ for all i who get a house in any such cycle. By induction with respect to the number of rounds, $t = b$ for all i. □

The TCC mechanism is strategy proof.

Proof. Let agent i complete a cycle in round k. The top trading cycles that agent i can complete before round k are still available at round k. (Note that agents cannot change their minds. They report their preferences before the algorithm starts). All trading cycles that agent i can complete after round k provide no better alternatives than what she gets in round k. □

Theorem: The top trading cycle algorithm is Pareto efficient and strategy proof. It is a dominant strategy for each agent to announce her true preference (type).

This result does not contradict the Gibbard-Sattethwaite Theorem, because preferences are restricted. For example, students only care about their own university; their utility functions are assumed to be unaffected by the assignment of universities to other students.

The TTC mechanism guarantees that the final allocation has the core property, which is stronger than Pareto efficiency.

Proof. Fix the initial allocation $a : I \to H$. Let t be the assignment generated by the TTC. There is no $C \subset I$ such that members in C can all rearrange their initial assingments so that $b(c) \succ_c t(c)$ for all $c \in C$. □

Theorem: Roth and Postlewaite 1977
The matching produced by Gale's TTC algorithm is the unique core matching.

Theorem: A mechanism is strategy–proof, Pareto–efficient, and individually rational

$$t(i) \succeq_i a(i) \text{ for all } i$$

if and only if it is the TTC.

The outcome of the TCC mechanism will depend on the initial assignment.

Example: Top Trading Cycle Algorithm
Consider the TCC mechanism given the following strict preferences.

$$\begin{aligned}
1 &: h_3 \succ h_2 \succ h_4 \succ h_1 \\
2 &: h_1 \succ h_4 \succ h_2 \succ h_3 \\
3 &: h_1 \succ h_2 \succ h_3 \succ h_4 \\
4 &: h_3 \succ h_2 \succ h_1 \succ h_4
\end{aligned}$$

Assume starting assignment: $a(1) = h_1$, $a(2) = h_2$, $a(3) = h_3$, and $a(4) = h_4$. The top trading cycles are

- Step 1: $1 \to 3 \to 1$.
- Step 2: $2 \to 4 \to 2$.

The final assignment is $t(1) = h_3$, $t(2) = h_4$, $t(3) = h_1$, and $t(4) = h_2$.

Suppose, however, that the starting assignment is: $a(1) = h_3$, $a(2) = h_2$, $a(3) = h_4$, and $a(4) = h_1$. The top trading cycles are

- Step 1: $1 \to 1$.
- Step 2: $2 \to 4 \to 2$.

The final assignment is $t(1) = h_3$, $t(2) = h_1$, $t(3) = h_4$, and $t(4) = h_2$.

Suppose, however, that the starting assignment is: $a(1) = h_1$, $a(2) = h_2$, $a(3) = h_4$, and $a(4) = h_3$. The top trading cycles are

- Step 1: $4 \to 4$.
- Step 2: $1 \to 2 \to 1$.
- Step 3: $3 \to 3$.

The final assignment is $t(1) = h_2$, $t(2) = h_1$, $t(3) = h_4$, and $t(4) = h_2$.

The outcome of the TCC mechanism depends on the initial assignment.

3.18.5. Serial Choice Algorithm

In order to apply the top trading cycle algorithm, we need to start form some initial assignment. The outcome will depend on this assignment. Imagine that no such assignment is given. Then we can choose it randomly and then apply the top trading cycle algorithm. Alternatively, we can use the serial choice algorithm.

The Serial Choice Algorithm:

- Pick a random order for all agents.
- Let the first agent in the order pick her best choice.
- Let the second agent in the order pick her best choice among the remaining choices.
- Repeat until all choices have been exhausted.

The outcome is Pareto efficient and strategy proof.

Proof. Suppose that agent i completes a cycle in round k. The top trading cycles that agent i can complete before round k are still available at round k. All trading cycles that agent i can complete after round k provide no better alternatives than what she gets in round k. □

The outcome is dictatorial, because the first agent in the order always gets her most favorite choice. Moreover, it is obvious that this mechanism is not manipulable. Observe, however, that if agents can exchange money, then Pareto efficiency is not guaranteed any more.

Example: Serial Choice Algorithm
Consider the serial choice algorithm given the following strict preferences.

$$\begin{aligned} 1 &: h_3 \succ h_2 \succ h_4 \succ h_1 \\ 2 &: h_1 \succ h_4 \succ h_2 \succ h_3 \\ 3 &: h_1 \succ h_2 \succ h_3 \succ h_4 \\ 4 &: h_3 \succ h_2 \succ h_1 \succ h_4 \end{aligned}$$

Assume a random order for all agents: $1 > 2 > 3 > 4$. The serial choice assignment is $t(1) = h_3$, $t(2) = h_1$, $t(3) = h_2$, and $t(4) = h_4$.

Suppose, however, that the random order for all agents is: $4 > 3 > 2 > 1$. Then the serial choice assignment is $t(1) = h_2$, $t(2) = h_4$, $t(3) = h_1$, and $t(4) = h_3$.

The outcome of the serial choice algorithm will depend on the random order of the agents.

3.18.6. Deferred Acceptance Algorithms in Two-Sided Markets

Empirical evidence (Roth and Peranson, 1998) shows that deferred acceptance algorithms do a very good job achieving stability even for large groups (i.e. $> 20,000$ agents with $20,000$ vacancies). Note that allowing for couple matching is complex [11].

The Deferred Acceptance Algorithm:

- Step 1: Each $m \in M$ makes a proposal to their best choice $W \cup \{\text{single}\}$. Each w tentatively accepts the best proposal and rejects all others.
- Step 2: All rejected m make proposals to all w who have not rejected them yet. Each w tentatively accepts the best proposal and rejects all others.
- Repeat until there is a stable matching (as both M and W are finite, the algorithm must stop at some point).

Theorem: Gale–Shapley 1962
A stable matching exists for any two–sided market with strict preferences.

Proof. If w likes some m more than the outcome of the DAA, then m never proposed to her. Thus, m is matched to somebody that he likes better than w. □

Theorem: Knuth
When all agents have strict preferences, all m have the same ranking across all stable matches, and all women have the same ranking across all stable matchings, then the two rankings are opposite to each other.

[11] The Roth Peranson algorithm allows for couples.

3.19. Matching—Marriage Problems

3.19.1. The College Admissions Problem

There are S students and C colleges. Each student has a preference (utility) over $C \cup \{\text{single}\}$ (some colleges can be unacceptable altogether). Each college has a preference over students (some students can be unacceptable altogether).

Definition: A quota $q(c)$ determines how many agents a company c can possibly hire.

A matching $F : S \to C \cup \{\text{single}\}$ assigns a college c to each student s. The number of students assigned to college c cannot exceed $q(c)$. Note that the corresponding matching function $G : C \to S$ takes S as values, so F is more convenient.

To take account of individual incentives for both colleges and students, consider the concept of stability. Say that a matching F is stable if

- $F(s) \succ \{\text{single}\}$ for all students s.
- $F(s) = c$ implies that s is acceptable for c.
- There is no pair of a college c and a student s such that the student strictly prefers c to $F(s)$ and c prefers s to some other s' such that $F(s') = c$.
- There is no pair of a college c and a student s such that the student strictly prefers c to $F(s)$, s is acceptable for c, and the college c is matched with fewer than $q(c)$ students under F.

Under additional assumptions on the college utility functions, stability implies Pareto efficiency. Stable matchings always exist and can be obtained through the student–proposing Student optimal deferred acceptance algorithm (SODA) or the college–proposing College optimal deferred acceptance algorithm (CODA).

3.19.2. Student optimal deferred acceptance algorithm (SODA)

The SODA:

- Each student proposes (applies) to her favorite college. Then each college c provisionally accepts $q(c)$ students out of those who have applied. If less are available, then fewer are provisionally accepted. Once a student is rejected, she cannot reapply to college.
- At the next round, the rejected students apply to their next best choice. The colleges repeat with the pool of conditionally accepted and the new applicants.
- The process continues until all students are accepted by a college or rejected by all.

Example: SODA

Students A, B, C apply to colleges X, Y, Z. All quotas are 1. Preferences are

$$\begin{aligned} &A : Y \succ X \succ Z && X : B \succ C \succ A \\ &B : X \succ Z \succ Y && Y : B \succ C \succ A \\ &C : X \succ Z \succ Y && Z : A \succ B \succ C \end{aligned}$$

Solution:

- Step 1: A applies to Y; B and C apply to X. Y provisionally accepts A; X provisionally accepts B.
- Step 2: C applies to Z. Z provisionally accepts C.
  The matching (AY, BX, CZ) becomes final.

3.19.3. College optimal deferred acceptance algorithm (CODA)

The CODA:

- Each college makes $q(c)$ offers to potential students. Students provisionally accept the best of the current offers and reject all others.
- Colleges make more offers to replace rejections. Students provisionally accept a better offer if they get one.
- The process continues until all colleges fill their quota $q(c)$ or are rejected by all acceptable students.

Example: CODA

Colleges X, Y, Z make offers to students A, B, C. All quotas are 1. Preferences are

$$\begin{aligned} A &: Y \succ X \succ Z & \qquad X &: B \succ C \succ A \\ B &: X \succ Z \succ Y & \qquad Y &: B \succ C \succ A \\ C &: X \succ Z \succ Y & \qquad Z &: A \succ B \succ C \end{aligned}$$

Solution:

- Step 1: Z admits A; X and Y both admit B. A tentatively accepts Z; B accepts X.
- Step 2: Y admits C. C accepts Y.
  The matching (AZ, BX, CY) becomes final.

3.19.4. Stability

The outcomes of SODA and CODA can be different. It is always weakly better to make offers rather than to have an opportunity to reject them. Thus, SODA provides better outcomes for students and CODA provides better outcomes for colleges.

Theorem: Gale

Given that all preferences are revealed truthfully, both SODA and CODA achieve stability.

3.19.5. Manipulability

Both SODA and CODA achieve stability, but they are manipulable by the rejecting side.

Theorem: Roth 1982

There is no method that achieves stability in a non-manipulated and non-dictatorial fashion.

Theorem: Sonmez 1997

Companies (colleges) will have incentives to hide their true capacity.

Theorem: Roth JET 1985

Truth–telling is a dominant strategy for all students under the student–optimal stable mechanism.

When the market is large, it becomes unlikely that schools can profitably misrepresent their preferences (Immorlica and Mahdian 2005, Kojima and Pathak 2008).

Theorem: The Rural Hospital Theorem (Roth and Ecta 1986)

Any college that does not fill all its positions at some stable matching is assigned precisely the same set of students at every stable matching.