First Year Microeconomics

Charles Christopher Hyland

July 13, 2026

Contents

1 Consumer Theory

1.1 Preference Relations

We are interested in analysing the idea of rational choice. There are two approaches to analysing individual choice:

1. Preference relation: Preferences are primitive concept
2. Revealed preferences: Choices are primitive concept

In this course, we only look at preference relations. Throughout this section, we shall let X denote the set of mutually exclusive alternatives to choose from.

Definition 1 (Preference Relation)
Let X be the set of mutually exclusive alternatives. A relation $\succsim$ denotes a **preference relation** if for $x, y \in X$,

$$x \succsim y \tag{1}$$

means that x is **weakly preferred** to y.

Rational choice theory starts with the idea that individuals have preferences and choose according to those. From any preference relation $\succsim$, we can define two more relations.

Definition 2 (Strict Preference Relation)
Let X be the set of mutually exclusive alternatives. The **strict** preference relation $\succ$ is defined by $x \succ y$ if

$$x \succsim y$$

but **not**

$$y \succsim x.$$

We say that x is **strictly preferred** to y.

Definition 3 (Indifference Relation)
Let X be the set of mutually exclusive alternatives. The **indifference relation** $\sim$ is defined by

$$x \sim y$$

if and only if

$$x \succsim y \quad \text{and} \quad y \succsim x.$$

We say that **x is indifferent to y.**

We now state two fundamental assumptions that pin down what we mean by *rational choice*.

Definition 4 (Complete Preference Relation)
$\succsim$ is **complete** if for every $x, y \in X$, we have either $x \succsim y$ or $y \succsim x$ (or both).

A complete preference relation can be summarised as that we can compare every pair of choices in our set X. The agent will always have an opinion on which choice they prefer. We now state the second assumption.

Definition 5 (Transitive Preference Relation)
The preference relation $\succsim$ is **transitive** if for $x, y, z \in X$, if $x \succsim y$ and $y \succsim z$ then $x \succsim z$.

We can combine these two assumptions to now formally state rational preferences.

Definition 6 (Rational Preference Relation)
A preference relation $\succsim$ is **rational** if it is both:

1. Complete;
2. Transitive.

We have two important implications of **rational preferences** on **strict preferences**:

1. $\succ$ is **irreflexive** (we cannot have $x \succ x$);
2. $\succ$ is **transitive** ($x \succ y$ and $y \succ z$ implies $x \succ z$).

We have three important implications of **rational preferences** on **indifference relation**

1. $\sim$ is **reflexive**;
2. $\sim$ is **transitive**;
3. $\sim$ is **symmetric** (if $x \sim y$ then $y \sim x$).

We also get the following transitive relation between $\succ$ and $\succsim$:

$$x \succ y \text{ and } y \succsim z \text{ implies that } x \succ z.$$

A big issue can arise when transitivity is violated as we can have preference cycles where there are no best options. Two causes for why transitivity may fail are:

1. Imperceptible differences: People can't tell options apart;
2. Endogenous taste change: Habit formation or temptation preferences.

From this point onwards, we shall assume that

1. The choice set is over L commodities and defined over positive values $X = \mathbb{R}^L_+$;
2. The preference relation $\succsim$ is defined on X;
3. The preference relation $\succsim$ is a **rational** preference relation.

We now introduce other mathematical objects which we will need.

Definition 7 (Indifference Curve)
The indifference curve for x is

$$\{y \in X : y \sim x\}. \tag{2}$$

Definition 8 (Upper Contour Set)
The upper contour curve for x is

$$\{y \in X : y \succsim x\}. \tag{3}$$

Definition 9 (Lower Contour Set)
The lower contour curve for x is

$$\{y \in X : x \succsim y\}. \tag{4}$$

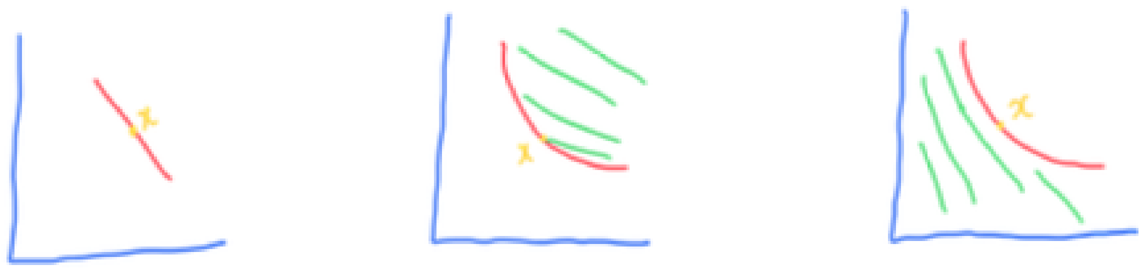

Figure 1: Indifference curve, upper contour set, and lower contour set for a point $x \in X$.

To proceed further, we need to make further restrictions on preferences. We introduce our first restriction of convexity.

Definition 10 (Convex preference relations)
A preference relation $\succsim$ is **convex** if the **upper contour sets** are **convex sets**.

That is, if $y, z \succsim x$, then

$$\alpha y + (1 - \alpha) z \succsim x$$

for any $\alpha \in [0, 1]$.

Definition 11 (Strictly convex preference relations)
A preference relation $\succsim$ is **strictly convex** if $y, z \succsim x$ and $y \neq z$ imply that

$$\alpha y + (1 - \alpha) z \succ x$$

for any $\alpha \in (0, 1)$.

Convexity captures the idea that agents like diversity. Convexity is very important: as we will see later, if consumer preferences are not convex, then market-clearing prices may not exist. We now state the next restrictions of monotonicity and local non-satiation.

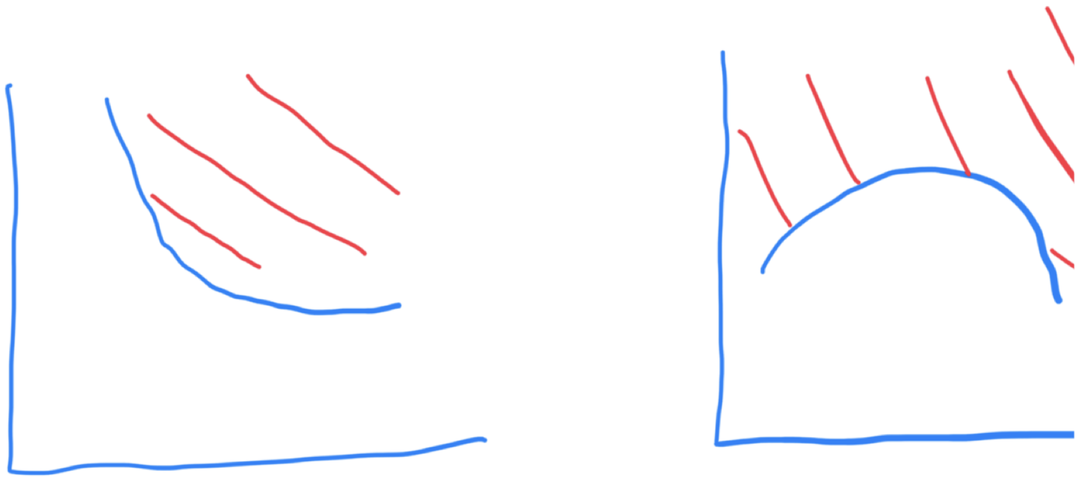

Figure 2: A strictly convex upper contour set and a non-convex upper contour set.

Definition 12 (Strongly Monotonic Preferences)
$\succsim$ is **strongly monotonic** if $x \in \mathbb{R}^L_+$ and $y \in \mathbb{R}^L_+$ are two bundles such that $y \neq x$ and $y_\ell \geq x_\ell$ for each good ℓ, then we have that $y \succ x$.

Strongly monotonic preferences can be summarised as that all goods are desirable and indifference curves cannot be *thick*. More is always better than less.

Definition 13 (Weakly Monotonic Preferences)
$\succsim$ is **weakly monotonic** if $x \in \mathbb{R}^L_+$ and $y \in \mathbb{R}^L_+$ are two bundles such that $y \neq x$ and $y_\ell \geq x_\ell$ for each good ℓ, then we have that $y \succsim x$.

Unlike strongly monotonic preferences, indifference curves under weakly monotonic preferences can be *thick*.

Remark 14. *We can see that for weakly monotonic preferences, we can pick a point, increase the quantity of one of the goods, and still be on the same indifference curve.*

In Fig. 3, the thin indifference curve represents a strongly monotonic preference whilst the thick indifference curve represents a weakly monotonic preference. We can see that for the weakly monotonic preference relation, it is possible to increase the quantity of either good 1 or 2 but still remain in the same indifference "curve" which is the area shaded in red.

Definition 15 (Locally non-satiated Preferences)
$\succsim$ is **locally non-satiated** if for every $x \in X$ and $\epsilon > 0$, there exists a $y \in X$ in the ball $B_\epsilon(x)$ such that $y \succ x$.

If an agent's preferences are locally non-satiated, this means that there is no bundle in X that is *ideal* or even locally ideal. There is always an improvement that can be made. Note that locally non-satiation depends on **both** the preferences $\succsim$ and the choice set X described.

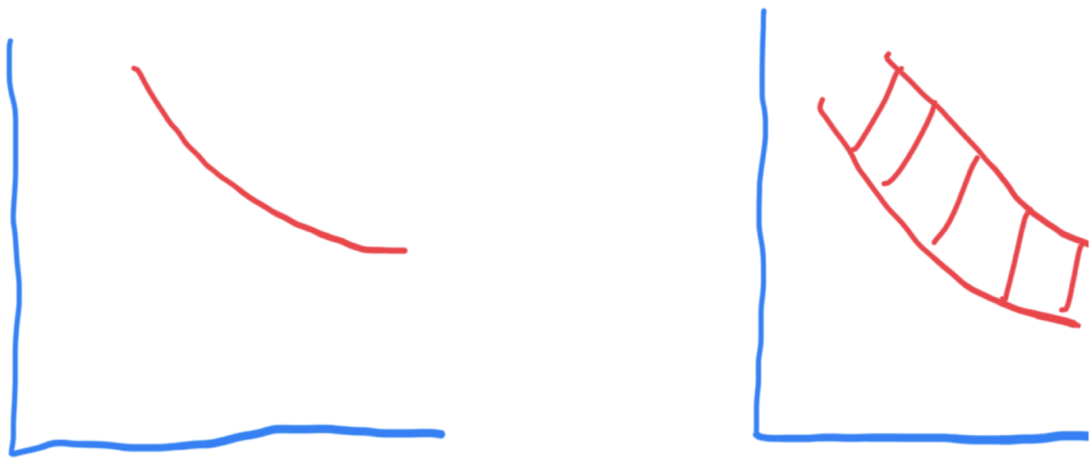

Figure 3: A thin indifference curve and a thick indifference curve.

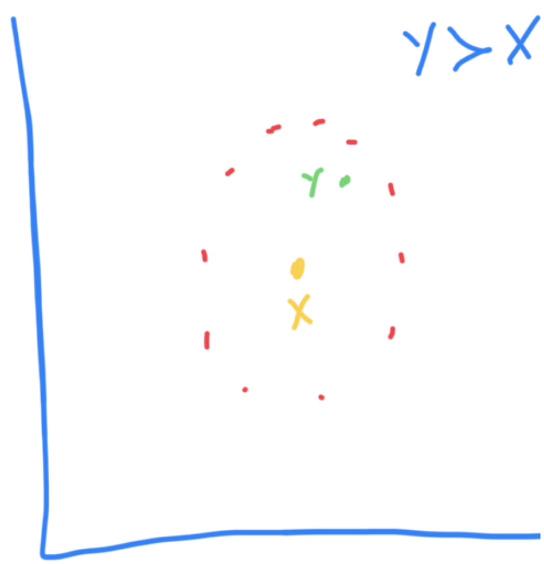

Figure 4: A locally non-satiated preference relation.

Proposition 16 (Strong monotonicity is the strongest restriction)
Strong monotonicity implies local non-satiation.

Proof. (Sketch). Pick a point $x \in \mathbb{R}^L_+$ and add the scaled unit vector $\lambda 1$ to it where 1 is a $L \times 1$ vector. For any $\epsilon > 0$, pick a $\lambda := \lambda(\epsilon)$ such that $x + \lambda 1 \in B_\epsilon(x)$. Then, we have that $x + \lambda 1 >> x$ by construction and by strong monotonicity, we have that $x + \lambda 1 \succ x$. $\square$

Remark 17. *Locally non-satiated preference relations do not imply strong monotonicity. To show this, we can pick a utility function such that the function is increasing with respect to some goods but reaches a satiation point for the remaining goods.*

Pick

$$u(x_1, x_2) = x_1 - |1 - x_2|$$

as an example. Here, when $x_2 = 1$, the consumer is satiated with respect to good 2. Therefore, locally non-satiated preferences do not imply strong monotonicity.

We now state an important property for preference relations we will need for when we define utility functions.

Definition 18 (Continuous Preference Relations)
The preference relation $\succsim$ on X is continuous if for any sequence of bundles $\{(x_n, y_n)\}_{n=1}^{\infty}$ with $x_n \succsim y_n$ for all n, $x = \lim_{n\to\infty} x_n$, $y = \lim_{n\to\infty} y_n$, then we have that $x \succsim y$.

Continuity can be summarised as saying that there are no jumps in consumers' preferences.

Proposition 19 (Alternative Characterisation of Continuous Preference Relations)
For all x, the upper contour set $\{y \in X : y \succsim x\}$ and lower contour set $\{y \in X : x \succsim y\}$ are closed sets.

Example 20 (Example of preference relation that is **not** continuous)
The lexicographic preferences are $\boldsymbol{x} \succsim \boldsymbol{y}$ if either $x_1 > y_1$ or $\{x_1 = y_1$ and $x_2 \geq y_2\}$. The lexicographic preferences are **not** continuous. Use the sequence $x_n = (\frac{1}{n}, 0)$ and $y_n = (0, 1)$ to verify this. $x_n \succ y_n$ for all n but $y \succ x$.

1.2 Utility Functions

In practice, it is hard to work with preference relations $\succsim$ and sets. Instead, it is much easier to work with utility functions that can **represent** these preferences.

Definition 21 (Utility Function)
A function $u : X \to \mathbb{R}$ is a **utility function** representing the preference relation $\succsim$ if, for all $x, y \in X$, $x \succsim y$ if and only if $u(x) \geq u(y)$.

Remark 22. *Utility functions are* **not unique**. *They are* **ordinal**, *and hence any strictly increasing transformation represents the same preferences.*

Example 23
Squaring is not a valid increasing transformation on all of $\mathbb{R}$. If a utility representation $u(x)$ can take negative values, the transformation $u(x)^2$ may reverse its ordering and therefore need not represent the same preferences.

We now ask whether every preference relation has an associated utility function and vice versa.

Proposition 24 (Utility Functions Represent Rational Preferences)
If a preference relation $\succsim$ can be represented by a utility function $u(\cdot)$, then $\succsim$ is **rational**.

Proof. Idea: We take an arbitrary utility function u, and by definition of utility functions, we know it has a preference relation associated with it.

Completeness: As u is a **real-valued** function, we know that for any $x, y \in X$, we have that either $u(x) \geq u(y)$ or $u(y) \geq u(x)$. By definition of utility function, we know that it has an associated preference relation $\succsim$ and therefore either $x \succsim y$ or $y \succsim x$. As x, y were arbitrary elements, we can conclude that $\succsim$ is complete.

Transitive: Suppose that $x \succsim y$ and $y \succsim z$. Then, $u(x) \geq u(y)$ and $u(y) \geq u(z)$. By the transitivity property of utility function, then $u(x) \geq u(z)$. Therefore, this implies that $x \succsim z$.

□

However, the converse does not hold as not every preference relation can have a utility function. In fact, this does not hold even if the preference relation is a rational preference relation.

Example 25 (Lexicographic preferences are rational preferences that do not admit a utility function representation)
A useful example is the lexicographic preference relation on $\mathbb{R}^2$, whereby $x \succsim y$ if either $x_1 > y_1$ or, if $x_1 = y_1$, then $x_2 \geq y_2$. It is rational but does not have a utility function representation. To see why, suppose that there exists a utility function that represents lexicographic preferences. Then, for every real number $x_1 \in \mathbb{R}$, we can pick a rational number $r(x_1)$ such that $u(x_1, 2) > r(x_1) > u(x_1, 1)$. Then, for $x_1 > x_0$, this implies that $r(x_1) > r(x_0)$. We have therefore constructed a one-to-one mapping from the real numbers $\mathbb{R}$ to the rationals $\mathbb{Q}$. However, the real numbers are uncountable whilst the rationals are countable. This is impossible and therefore a contradiction. Hence, there is no such utility function.

Therefore, only certain preference relations can have utility function representations.

Proposition 26 (Sufficient Conditions for Utility Function Representations)
If the preference relation $\succsim$ satisfies the following properties:

1. rational,
2. strongly monotonic,
3. continuous,

then there exists a **continuous** function $u(\cdot)$ such that $u(y) \geq u(x)$ if and only if $y \succsim x$.

The numbers that are assigned by utility functions are not economically significant. However, the geometry of the indifference curves tells us a lot of information. For example, in Section 1.2, the absolute value of the slope of the indifference curve tells us the **marginal rate of substitution** between products 1 and 2.

Definition 27 (Marginal Rate of Substitution)
The marginal rate of substitution between goods A and B is:

$$MRS = \frac{\partial u(x)/\partial x_A}{\partial u(x)/\partial x_B}. \tag{5}$$

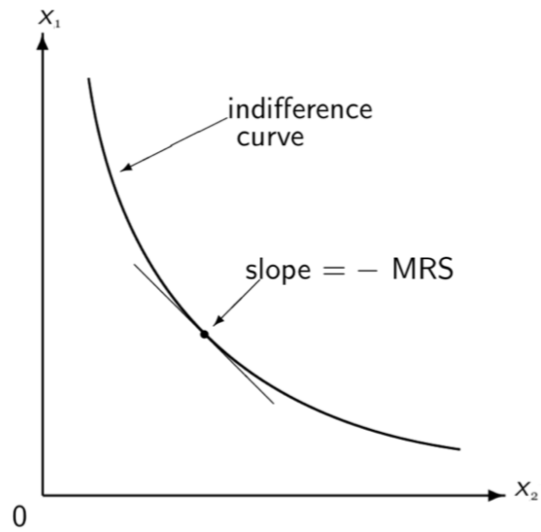

Figure 5: Marginal rate of substitution on an indifference curve.

We can give another derivation of the marginal rate of substitution as being the slope of the utility indifference curve. Define the following identity:

$$u(x_A, x_B(x_A)) \equiv \overline{u}$$

where $\overline{u}$ is an arbitrary utility level and $x_B(x_A)$ is the level of x_B needed to achieve utility level $\overline{u}$ when the level of x_A consumed is x_A. We can then differentiate this identity with respect to x_A to get:

$$\frac{\partial u}{\partial x_A} + \frac{\partial u}{\partial x_B}\frac{dx_B}{dx_A} = 0$$

We then re-arrange to get:

$$\frac{dx_B}{dx_A} = -\frac{\frac{\partial u}{\partial x_A}}{\frac{\partial u}{\partial x_B}}$$

which is again the marginal rate of substitution!

We can see that in fact, this definition of the marginal rate of substitution does not depend on the particular choice of the utility function.

Theorem 28 (MRS is Invariant to Utility Function Form)
The definition of the marginal rate of substitution is invariant to the choice of the utility function $u(\cdot)$:

$$MRS = \frac{\partial u(x)/\partial x_A}{\partial u(x)/\partial x_B}. \tag{6}$$

Proof. Define $f(\cdot)$ to be a monotonically increasing transformation of $u(x)$. This preserves the ordering. Then, let $U(x_A, x_B) = f[u(x_A, x_B)]$ and $f' > 0$. Then the MRS is

$$MRS = \frac{f' u_A}{f' u_B} = \frac{u_A}{u_B}$$

□

Indifference curves, because they depend on underlying preferences and not on their representation, are the bedrock upon which classical demand theory is built.

Proposition 29 (Convexity Implications for Marginal Rate of Substitution)
For differentiable, convex preferences over two goods, the marginal rate of substitution is weakly diminishing as one moves along an indifference curve towards more of the good in the numerator.

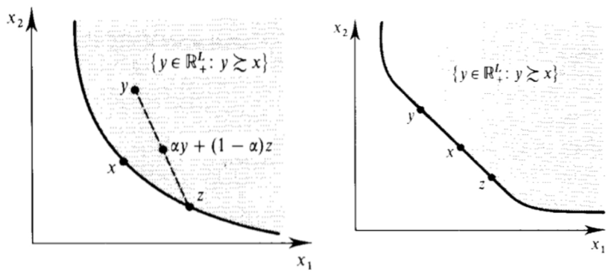

Figure 6: Convexity implies diminishing marginal rate of substitution on an indifference curve.

Remark 30. *This captures the consumer's taste for variety. Under strict convexity, a consumer will strictly prefer a non-trivial mixture of two distinct bundles between which they are indifferent.*

We will often require utility functions to be quasiconcave.

Definition 31 (Quasiconcave Utility)
A utility function $u(\cdot)$ defined on a set S is quasiconcave if every upper level set of $u(\cdot)$ is a convex set. That is, the upper level set:

$$P(a) = \{x \in S : u(x) \geq a\} \tag{7}$$

is a convex set for every value a.

We relate quasiconcavity to concavity which is likely more familiar.

Proposition 32 (Concavity Implies Quasiconcavity)
Every concave function is quasiconcave.

We may now ask how to check whether a utility function is quasiconcave.

Proposition 33 (Condition for Quasiconcavity of a Utility Function)
A utility function $u(\cdot)$ is quasiconcave if and only if all of its upper contour (upper level) sets are convex.

Remark 34. *In a differentiable two-good problem with monotone preferences, one can often write an indifference curve as $x_2 = g(x_1)$ and verify convexity by checking that $g''(x_1) \geq 0$. The general criterion, however, concerns the entire upper contour set rather than the indifference curve alone.*

1.2.1 Special Types of Preferences: Homothetic Preferences

We now move onto special types of preferences. The first one we will look at are **homothetic preferences.**

Definition 35 (Homothetic Preferences)
A preference is **homothetic** if we have that for $x \sim y$ then $kx \sim ky$ for any positive scaling factor k.

This says that if we have two bundles of goods between which we are indifferent, then scaling both bundles by the same positive factor preserves indifference. Geometrically, their utility function representation can be seen in Fig. 7. The indifference curves are magnified versions of lower indifference curves. It is also notable that the marginal rate of substitution at the point y is the **same** as the marginal rate of substitution at $2y$. Hence, the marginal rate of substitution is unaffected by the scaling factor k. This gives us a sufficient condition for checking whether a utility function represents homothetic preferences.

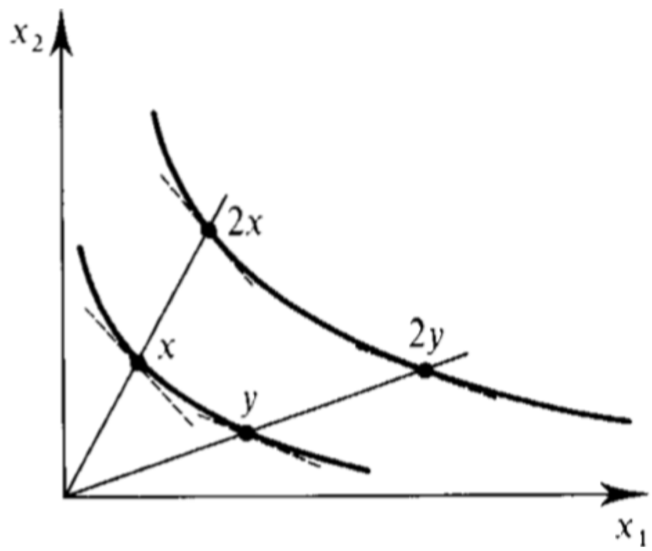

Figure 7: Indifference curves of a utility function representing homothetic preferences. If $x \sim y$, then scaling both bundles gives $2x \sim 2y$. For a homogeneous representation, the MRS is unaffected by the scaling factor k and is constant along rays through the origin.

Proposition 36 (Sufficient Conditions for Homothetic Preferences)
If marginal rates of substitution are constant along the rays through the origin, then the preferences that the utility function represents are homothetic.

Homothetic preferences are useful since if we are given one indifference curve, this gives information on all indifference curves as they are radial expansions of one another. We can extend this definition of homothetic preferences to utility functions. First, we need to recall what *homogeneity* is.

Definition 37 (Homogeneity of Degree r)
A function $f(x)$ is **homogeneous of degree r** if for any $\lambda > 0$, we have that:

$$f(\lambda x) = \lambda^r f(x). \tag{8}$$

We can now characterise homothetic preferences in terms of utility functions.

Proposition 38 (Alternative Characterisation of Homothetic Preferences)
A continuous preference relation $\succsim$ is **homothetic** if and only if it can be represented by a utility function $u(\cdot)$ that is **homogeneous of degree 1**, i.e.

$$u(\lambda x) = \lambda u(x)$$

for all $\lambda > 0$ and x.

Remark 39. *Note that there is a difference between* **homothetic preferences** *and* **homothetic functions***. A* **homothetic function** *is a function $h(\cdot)$ that can be written as $h(x) = g(f(x))$, where $f(\cdot)$ is* **homogeneous of degree** *r and $g(\cdot)$ is strictly increasing. A utility function represents homothetic preferences precisely when it is an increasing transformation of a homogeneous utility representation; arbitrary utility functions need not be homothetic.*

An implication of homothetic preferences is that, at fixed prices, budget shares are independent of income and income elasticities are equal to one. This is restrictive and is often rejected by the data. For example, *Engel's law* states that the budget share of food falls with income.

Example 40 (Cobb-Douglas Utility Function)
The Cobb-Douglas utility function is:

$$u(x_1, x_2) = x_1^{\alpha} x_2^{1-\alpha}$$

for $0 < \alpha < 1$. The Marshallian demand functions for this utility function are given by:

$$x_1(p, w) = \frac{\alpha w}{p_1},$$
$$x_2(p, w) = \frac{(1-\alpha)w}{p_2}.$$

The ratio of Marshallian demands x_1 and x_2 does not depend on wealth w. Therefore, the Engel curves are straight lines. The indirect utility function is given by:

$$v(p, w) = w\left(\frac{\alpha}{p_1}\right)^{\alpha}\left(\frac{1-\alpha}{p_2}\right)^{1-\alpha}. \tag{9}$$

1.2.2 Special Types of Preferences: Quasilinear Preferences

We now move onto another type of special preferences: **quasilinear preferences**.

Definition 41 (Quasi-linear preferences)
Let the set of choices be $X = (-\infty, \infty) \times \mathbb{R}_+^{L-1}$. Let $x \sim y$ be two bundles of goods whereby $x, y \in X$. Then, for any k, we have that:

$$(x_1 + k, x_2, \cdots, x_L) \sim (y_1 + k, y_2, \cdots, y_L). \tag{10}$$

Such a preference relation $\sim$ is called a **quasi-linear preference**.

Remark 42. *Suppose we have two bundles $x, y \in X$ between which we are indifferent. Adding the same amount of one particular good (in this case the first good) preserves indifference. Notice that the first good can also take negative quantities.*

Like homothetic preferences, we are able to draw inferences from all indifference curves if we are only given one indifference curve. We can also state the utility function representation of quasi-linear preferences.

Definition 43 (Utility Function representation of quasi-linear preferences)
The utility function representation of quasi-linear preferences is given by:

$$x_1 + u(x_2, \cdots, x_L). \tag{11}$$

Remark 44. *The linear part x_1 is interpreted as money or called the outside good. It is often set as the numeraire.*

In Fig. 8, we plot the indifference curves of a quasi-linear utility function.

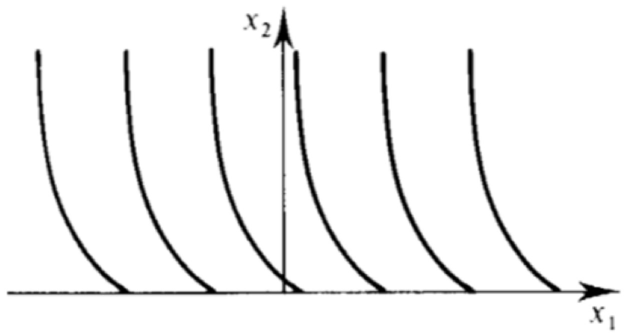

Figure 8: Quasi-linear utility function. Here, x_1 is the linear good. All indifference curves are horizontal shifts of each other. The marginal rate of substitution only depends on x_2.

Example 45 (Quasilinear Utility Function)
The quasilinear utility function is given by:

$$u(x) = x_1 + \sqrt{x_2}$$

The associated Marshallian demand curves are:

$$x_1(p, w) = -\frac{1}{4}\frac{p_1}{p_2} + \frac{w}{p_1}$$
$$x_2(p, w) = \left(\frac{p_1}{2p_2}\right)^2$$

The indirect utility function is given by:

$$v(p, w) = \frac{1}{4}\frac{p_1}{p_2} + \frac{w}{p_1}. \tag{12}$$

Proposition 46 (Income Effects for Quasi-linear utility)
Under quasi-linear utility functions, all the income effects are concentrated on the outside good.

1.2.3 Special Types of Preferences: Examples

Proposition 47 (Preferences That Are Homothetic and Quasilinear)
The preferences that are both homothetic and quasilinear are perfect-substitutes preferences, which admit the

representation

$$u(x_1, x_2) = ax_1 + bx_2.$$

The MRS is $-a/b$ and is constant and the indifference curves are parallel straight lines.

Definition 48 (Perfect Complements Utility Function)
The perfect complements utility function is given by:

$$u(x_1, x_2) = \min[ax_1, bx_2]$$

These utility curves are L-shaped with kinks lying on a ray through the origin of slope a/b.
The preferences are homothetic but not quasilinear.

Definition 49 (Cobb-Douglas Utility Curve)
The Cobb-Douglas utility curve is given by:

$$u(x_1, x_2) = a \ln(x_1) + b \ln(x_2) \tag{13}$$

for $a, b > 0$. It is homothetic, its indifference curves are smooth, and the MRS $ax_2/(bx_1)$ is diminishing.

1.2.4 Special Types of Preferences: Separability

The next preference restriction we look at is **separability**. Separability allows us to focus on a subset of the world and ignore everything else.

Definition 50 (Separability of Preferences)
Partition the bundle of goods into (x, z) and the price vector into (p, q). The preference $\succsim$ are separable in the $x-$goods if the preference $\succsim$ has the property that:

$$(x, z) \succsim (x', z) \leftrightarrow (x, z'') \succsim (x', z''). \tag{14}$$

for all x, x', z, z''.

Remark 51. *Informally, the preferences over the goods x are independent from the preferences over other goods z. If x is preferred to x' given some quantity of other goods z, then x is still preferred to x' for any quantities of other goods z''.*

Separability restricts how marginal rates of substitution depend on goods outside the group. It does not, in general, eliminate all cross-price effects, because outside prices can still affect the budget allocated to the group.

Proposition 52 (Utility Function under Separable Preferences)
Suppose that $\succsim$ is a **separable** preference in the x-goods. Suppose that $u(\cdot)$ is the utility function that represents $\succsim$. Then, we may write $u(\cdot)$ as:

$$u(v(x), z) \tag{15}$$

whereby $v(x)$ is a **sub-utility** function and $u(v, z)$ is an increasing function of v.

Under separable preferences (in the x-goods), the overall utility is written as a function of sub-utility of x and level of consumption of the z-goods. This gives us a very useful result.

Proposition 53 (Marginal Rate of Substitution of Separable Utility Function)
Let $u(v(x), z)$ be a separable utility function. Then, the marginal rate of substitution with respect to x is independent of the z-goods.

We need to be aware on distinguishing between different notions of separability.

Definition 54 (Weakly Separable in x-goods)
The utility function structure:

$$u(v(x), z) \tag{16}$$

is **weakly separable in the x-good**.

Definition 55 (Weakly Separable)
The utility function structure:

$$u(v(x), w(z)) \tag{17}$$

is **weakly separable**.

We can clearly see from this that

$$\text{Weakly Separable} \Rightarrow \text{Weakly Separable in the x-goods}$$

Definition 56 (Additively Separable)
The utility function structure:

$$u(x, z) = v(x) + w(z) \tag{18}$$

is **additively separable**.

Example 57
Quasilinear preferences are additively separable: $u(x, z) = v(x) + z$.

The separability assumption is useful because it:

1. Allows us to focus on a subset of choices;
2. Simplifies empirical analysis by only recovering part of the preference structure and not having to model everything.

We will later see in expected utility theory, additive separability plays a role whereby the expected utility is a weighted average across different states of the world. In intertemporal choice theory, additive separability is also used whereby the lifetime utility of an agent is the sum of discounted utilities.

2 Utility Maximisation and Expenditure Minimisation

2.1 Marshallian Demand Function

For this section, we take the number of commodities to be finite and equal to L. The commodity vector is $x = (x_1, \ldots, x_L) \in \mathbb{R}^L_+$, since consumers cannot consume a negative amount of a good. Correspondingly, the strictly positive price vector is $p = (p_1, \ldots, p_L) \in \mathbb{R}^L_{++}$. We denote the consumer's wealth by $w > 0$. The consumer can afford any bundle x that satisfies:

$$p \cdot x \equiv p_1 x_1 + \ldots + p_L x_L \leq w,$$

whereby the amount you spend is less than or equal to your wealth. We can then define the Walrasian budget set from this.

Definition 58 (Walrasian Budget Set)
The budget set $B_{p,w}$ for price level p and wealth level w is defined to be:

$$B_{p,w} = \{x \in \mathbb{R}^L_+ : p \cdot x \leq w\} \tag{19}$$

Remark 59. *The linear inequality $p \cdot x \leq w$ is in fact a half-space. The budget set is therefore the intersection of the half spaces: $p \cdot x \leq w$ and $x \geq 0$.*

We can state two useful properties of the budget set.

Proposition 60 (Properties of Budget Set)
The budget set $B_{p,w}$ has the following properties:

1. $B_{p,w}$ is homogeneous of degree zero in (p, w);
2. $B_{p,w}$ is increasing with respect to w and decreasing with respect to each p_ℓ;
3. $B_{p,w}$ is a compact set.

Proof. (Sketch). (1) Scaling all prices and wealth by the same positive factor does not change the budget inequality. (2) As wealth increases or a price decreases, the set of affordable bundles expands. (3) The set is closed, and strictly positive prices imply that $0 \leq x_\ell \leq w/p_\ell$ for every ℓ, so it is bounded. Hence it is compact. □

When the budget constraint binds, we call this the **budget hyperplane**:

$$\{x \in \mathbb{R}^L_+ : p \cdot x = w\}.$$

The price vector p is in fact **perpendicular** to the budget hyperplane as seen in Fig. 9.

We can now define the consumer's utility maximisation problem.

Definition 61 (Utility Maximisation Problem)
Let $x \in \mathbb{R}^L_+$ and $p \in \mathbb{R}^L_{++}$. Let $B_{p,w}$ denote the budget set. The consumer problem is to maximise $u(x)$ over $x \in B_{p,w}$:

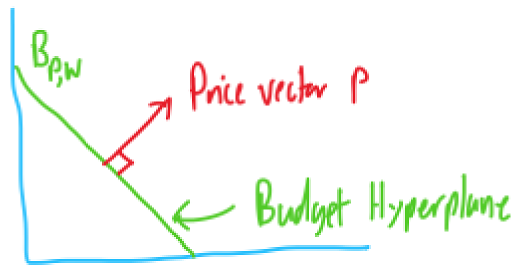

Figure 9: Budget hyperplane with orthogonal price vector p.

$$\begin{aligned} \max_x \quad & u(x) \\ \text{s.t.} \quad & p \cdot x \leq w \end{aligned} \tag{20}$$

We shall abbreviate the utility maximisation problem by **UMP**. We can show that such a problem has a solution under certain conditions.

Proposition 62 (Solution to Utility Maximisation Problem)
The consumer utility maximisation problem has a solution when the utility function $u(\cdot)$ is continuous and all prices p_ℓ are positive. Furthermore, this solution is **unique** when preferences $\succsim$ are **strictly convex**.

Proof. This holds via Weierstrass theorem. □

For an interior solution, and when local nonsatiation makes the budget constraint bind, we can set up the Lagrangian for the consumer's **UMP**:

$$\mathcal{L}(x, \lambda) = u(x) - \lambda(p \cdot x - w) \tag{21}$$

which has the associated $L + 1$ first order conditions being:

$$L_x = \nabla u(x) - \lambda p = 0 \quad \text{and} \quad L_\lambda = p \cdot x - w = 0$$

The utility-maximising vector of goods is known as the **Marshallian demand** $x(p, w)$.

Definition 63 (Marshallian Demand Function)
The **Marshallian demand correspondence** is the set of solutions to the consumer's utility maximisation problem. When the solution is unique, it is the **Marshallian demand function** $x(p, w)$.

Remark 64. *There is an alternative approach by viewing choices as primitive in order to derive the Marshallian demand function where we can view the Marshallian demand function as being* **generated**.

The Marshallian demand function has two useful properties.

Proposition 65 (Homogeneity of Marshallian Demand Function)
The Marshallian demand function $x(p, w)$ is homogeneous of degree zero:

$$x(kp, kw) \equiv x(p, w) \tag{22}$$

for any $k > 0$.

Proof. Scaling prices and wealth by the same positive factor does not change the budget set, as $B_{p,w}$ is homogeneous of degree zero. Therefore, the solution remains the same. □

Remark 66. *This implies that the units used for money do not matter. Therefore, we can normalise the price of one commodity to 1, which is known as the* **numéraire***. We can also normalise prices so that the budget is one.*

Proposition 67 (Walras' Law)
If preferences are locally nonsatiated, then the budget constraint binds:

$$p \cdot x(p, w) \equiv w. \tag{23}$$

Remark 68. *This says that consumers spend their entire wealth. Strong monotonicity is sufficient, but Walras' law also holds under the weaker assumption of* **local nonsatiation***. In a dynamic context, the budget constraint applies to consumption and saving over the consumer's entire* **lifetime***.*

We look at specific properties of the Marshallian demand function.

Definition 69 (Budget Share)
The **budget share** of product ℓ is denoted by:

$$b_\ell(p, w) \equiv \frac{p_\ell x_\ell(p, w)}{w} \tag{24}$$

Remark 70. *The budget share b_ℓ tells us what fraction of our wealth w is spent on product ℓ.*

Proposition 71 (Budget Shares Add up to One)
By Walras' law, the budget shares add up to one:

$$b_1 + \ldots + b_L \equiv 1.$$

We are now interested in asking the question of what happens to the Marshallian demand for a good as we vary the wealth of the consumer.

Definition 72 (Engel Curve)
Let $x(p, w)$ be the Marshallian demand. Fix prices $p = \overline{p}$. The **Engel curve** is the function $x(\overline{p}, w)$ as you vary w:

$$E(\overline{p}) = \{x(\overline{p}, w) : w > 0\}. \tag{25}$$

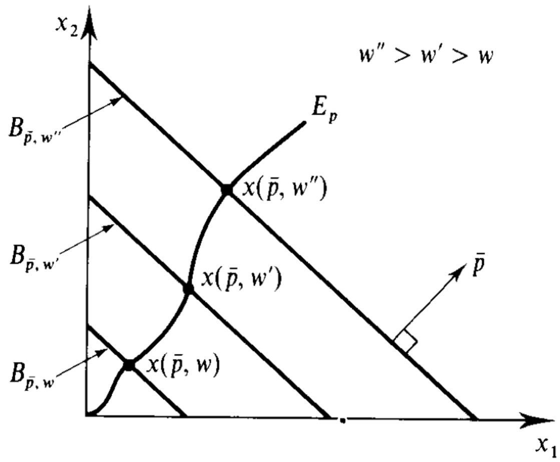

Figure 10: Engel curve.

Remark 73. *This tells us what is our Marshallian demand for goods as we vary our wealth. This is also known as* **income expansion path**

An example of an Engel curve is given in Section 2.1.

We can then define what it means for the Marshallian demand for a good to change as we vary wealth and then use this definition to classify goods.

Definition 74 (Wealth Effect for a Good)

At any price and wealth (p, w), the **wealth effect** for good ℓ is given by

$$\frac{\partial x_\ell}{\partial w}. \tag{26}$$

We can classify goods based on their wealth effects.

Definition 75 (Normal and Inferior Goods)

We can classify different goods based on their wealth effect:

1. A good is a **normal good** if $\frac{\partial x_\ell}{\partial w} \geq 0$
2. A good is a **inferior good** if $\frac{\partial x_\ell}{\partial w} < 0$

A Marshallian demand is classified as **normal** if for every good ℓ, $\frac{\partial x_\ell}{\partial w} \geq 0$.

We can also describe the same behaviour of demand for goods as a function of wealth in the form of budget shares.

Definition 76 (Luxury and Necessity Goods)

We can classify different goods based on their budget share behavior:

1. A good is a **Luxury good** if the budget share $b_\ell(p, w)$ increases with w;
2. A good is a **Necessity good** if the budget share $b_\ell(p, w)$ decreases with w.

Remark 77. *Every luxury good is a normal good but not every normal good is a luxury good. Conversely, every inferior good is a necessity.*

We saw that Engel curves describe Marshallian demand over different levels of wealth. We can define something analogous by fixing wealth and varying the price of a single good.

Definition 78 (Offer Curve)
Fix the prices of all goods except good ℓ, as well as wealth w. The **offer curve** is the Marshallian demand bundle as we vary the price p_ℓ whilst keeping all other prices $\overline{p}_{-\ell}$ and wealth $\overline{w}$ fixed:

$$O(\overline{p}_{-\ell}, \overline{w}) = \{x((p_\ell, \overline{p}_{-\ell}), \overline{w}) : p_\ell > 0\}. \tag{27}$$

An example of an offer curve is given in Fig. 11.

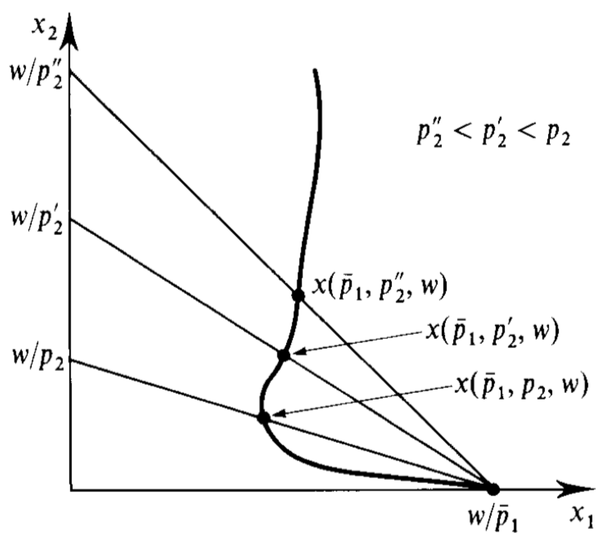

Figure 11: Offer curve.

We have defined the change in Marshallian demand as a result of a change in wealth. We can do something analogous for a change in prices.

Definition 79 (Price Effect)
The **price effect** of changing the price p_k of good k on the Marshallian demand x_ℓ for good l is given by:

$$\frac{\partial x_\ell}{\partial p_k}. \tag{28}$$

The own-price effect for good ℓ is given by:

$$\frac{\partial x_\ell}{\partial p_\ell}. \tag{29}$$

Generally, if a good's price rises, then its Marshallian demand falls. This implies that the own price effect is negative. However, there are cases when this is **not** the case.

Definition 80 (Giffen Good)
A good l is a **Giffen good** if $\frac{\partial x_l(p,w)}{\partial p_l} > 0$.

We now describe the **price effect matrix**.

Definition 81 (Price Effect Matrix)
The price effect matrix for L goods is a $L \times L$ Jacobian matrix of derivatives of the Marshallian demands for each good with respect to each price:

$$\begin{bmatrix} \frac{\partial x_1}{\partial p_1} & \frac{\partial x_1}{\partial p_2} & \cdots & \frac{\partial x_1}{\partial p_L} \\ \frac{\partial x_2}{\partial p_1} & \frac{\partial x_2}{\partial p_2} & \cdots & \frac{\partial x_2}{\partial p_L} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial x_L}{\partial p_1} & \frac{\partial x_L}{\partial p_2} & \cdots & \frac{\partial x_L}{\partial p_L} \end{bmatrix} \tag{30}$$

We can relate the price effect and wealth effect for goods.

Proposition 82 (Implications of Homogeneity of Degree 0 on Marshallian Demand)
If the Marshallian demand function $x(p,w)$ is homogeneous of degree zero, then for all p and w:

$$\sum_{k=1}^{L} \frac{\partial x_l(p,w)}{\partial p_k} p_k + \frac{\partial x_l(p,w)}{\partial w} w = 0 \tag{31}$$

for $l = 1, \ldots, L$.

Proof. As $x(p,w)$ is homogeneous of degree zero, we have

$$x(\alpha p, \alpha w) = x(p,w).$$

If we differentiate this expression with respect to α, we get the desired result (by chain rule). $\square$

Remark 83. *We see that the demand for good l depends on the price effect of the other L - 1 goods and the wealth effect of good l.*

Generally, it is easier to work with unit-free elasticities. Recall that a price elasticity is the percentage change in demand divided by the percentage change in price. Using differential notation,

$$\frac{dx_\ell(p,w)}{x_\ell(p,w)}$$

is the proportional change in Marshallian demand. Analogously, the proportional change in the price of good k is

$$\frac{dp_k}{p_k}.$$

Therefore, we can divide the two to get the price elasticity of good ℓ with respect to a price change in good k:

$$\text{elasticity} = \frac{\partial x_\ell(p,w)}{\partial p_k} \frac{p_k}{x_\ell(p,w)}.$$

We can do something analogous for the elasticity of good ℓ with respect to wealth w.

Definition 84 (Elasticities of Demand)
The **elasticity of demand for good ℓ with respect to price of good k** is:

$$\epsilon_{\ell k}(p, w) = \frac{\partial x_\ell(p, w)}{\partial p_k} \frac{p_k}{x_\ell(p, w)} \tag{32}$$

The **elasticity of demand for good ℓ with respect to wealth** is:

$$\epsilon_{\ell w}(p, w) = \frac{\partial x_\ell(p, w)}{\partial w} \frac{w}{x_\ell(p, w)} \tag{33}$$

We can classify goods again based on elasticities with respect to wealth.

1. A good ℓ is **inferior** if $\epsilon_{\ell w} < 0$;
2. A good ℓ is **normal** if $\epsilon_{\ell w} > 0$;
3. A good ℓ is **necessity** if $\epsilon_{\ell w} < 1$;
4. A good ℓ is **luxury** if $\epsilon_{\ell w} > 1$.

Hence, we can classify goods according to how their demand changes as wealth changes.

We need this result for later on.

Proposition 85 (Euler's Theorem)
Let $f(\cdot)$ be a differentiable function that is homogeneous of degree λ. Then,

$$\lambda f(\boldsymbol{x}) = \boldsymbol{x}^T \nabla f(\boldsymbol{x}). \tag{34}$$

We now move onto two important identities we will use.

Proposition 86 (Cournot Aggregation)
If the Walrasian demand function $x(p, w)$ satisfies Walras' law, then for all p and w:

$$\sum_{l=1}^{L} p_l \frac{\partial x_l(p, w)}{\partial p_k} + x_k(p, w) \equiv 0 \tag{35}$$

for $k = 1, ..., L$.

Equivalently:

$$p \cdot D_p x(p, w) + x(p, w)^T \equiv 0^T$$

Total expenditure cannot change in response to a change in prices.

Proof. Differentiate Walras' law by p:

$$p \cdot x(p, w) = w$$

and take product rule with respect to p. □

Remark 87. *The Cournot aggregation is useful since if we are able to observe prices p and the slopes of the demand curve, then we can recover the Marshallian demand $x(p, w)$ itself.*

Proposition 88 (Engel Aggregation)
If the Walrasian demand function $x(p, w)$ satisfies Walras' law, then for all p and w:

$$\sum_{l=1}^{L} p_l \frac{\partial x_l(p, w)}{\partial w} \equiv 1$$

Equivalently:

$$p \cdot D_w x(p, w) = 1$$

Total expenditure must change by an amount equal to any wealth change.

Proof. Differentiate Walras' law with respect to w:

$$p \cdot x(p, w) = w.$$

□

2.2 Utility Maximisation

We now give mathematical derivations to build intuition for Lagrange multipliers. The slope of the budget set describes the trade-offs available in the market, while the slope of an indifference curve (the marginal rate of substitution) describes the consumer's own trade-offs. The Lagrange multiplier and first-order conditions connect these objects. We now show this formally.

Suppose that we have a differentiable utility function and an interior solution $x_\ell > 0$. Say that we have two goods A and B and suppose we are currently at an optimal bundle x^*

$$p_A x_A^* + p_B x_B^* = W.$$

Then, suppose that we adjust quantities by Δx so that the bundle still remains affordable

$$p_A(x_A^* + \Delta x_A) + p_B(x_B^* + \Delta x_B) = W.$$

Subtracting the original budget equation, this can hold only if:

$$p_A \Delta x_A + p_B \Delta x_B = 0.$$

We can re-arrange this to get:

$$\Delta x_B = -\frac{p_A}{p_B} \Delta x_A.$$

Now, looking at a utility function $u(x_A, x_B)$, the impact of such an adjustment Δx can be found by taking the total derivative of the utility function:

$$\frac{\partial u(x^*)}{\partial x_A} \Delta x_A + \frac{\partial u(x^*)}{\partial x_B} \Delta x_B \equiv 0.$$

We can then plug in our expression for Δx_B to get:

$$\frac{\partial u(x^*)}{\partial x_B} \Delta x_A \left[\frac{\partial u(x^*)/\partial x_A}{\partial u(x^*)/\partial x_B} - \frac{p_A}{p_B} \right] \equiv 0.$$

At any interior solution, Δx_A can be either positive or negative. Hence, the only way such an adjustment cannot increase utility is if the terms inside the brackets add up to zero. Therefore, the marginal rate of substitution between goods A and B is:

$$MRS(A, B) = \frac{p_A}{p_B}.$$

Since this must hold for every pair of goods, we have the following.

Proposition 89 (Marginal Utility is proportional to price vector)
At the optimum, **marginal utility is proportional to price vector**:

$$\frac{\partial u(x^*)}{\partial x_\ell} = \lambda p_\ell \quad \text{for } 1 \leq \ell \leq L$$

for some positive constant λ.

If we look at our newly derived condition for good 1, we can rewrite it as follows:

$$\frac{\partial u(x^*)}{\partial x_1} = \lambda p_1 \quad \leftrightarrow \quad \frac{\frac{\partial u(x^*)}{\partial x_1}}{\lambda} = p_1.$$

Recall that λ is the marginal utility of income. Therefore, the ratio of marginal utility to the Lagrange multiplier λ gives the consumer's marginal willingness to pay for good 1, measured in units of money per unit of the good.

Hence, if $u(\cdot)$ is continuously differentiable and well-behaved, an optimal bundle $x^*(p, w)$ can be characterised using the (necessary) first order conditions. These say that if $x^*(p, w)$ is a solution, then there exists a Lagrange multiplier λ such that:

$$\nabla u(x^*) \leq \lambda p$$
$$x_\ell^* \left[\frac{\partial u(x^*)}{\partial x_\ell} - \lambda p_\ell\right] = 0, \qquad \ell = 1, \ldots, L.$$

Here, for every purchased good $x_\ell^* > 0$, the corresponding first constraint binds: $\partial u(x^*)/\partial x_\ell = \lambda p_\ell$. For a good at a corner, $x_\ell^* = 0$, the complementary-slackness condition holds even when its marginal utility is strictly below λp_ℓ.

We need to be careful at corner solutions: the marginal rate of substitution need not equal the price ratio, and the gradient need not satisfy $\nabla u(x^*) = \lambda p$ for every good. In Fig. 12, the unconstrained tangency would require a negative amount of one good. The non-negativity constraint instead forces the optimum to the boundary.

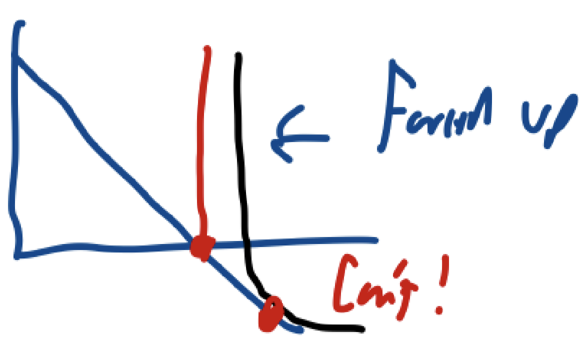

Figure 12: Corner Solution.

2.2.1 Additively Separable Preferences

We are now interested in looking at some interesting preferences. Recall that additive separability is the case where the overall utility is the sum of the sub-utilities for each goods. Therefore, under additive separability, we may write the utility function as the sum of sub-utilities:

$$U(x) = \sum_{l=1}^{L} u_l(x_l).$$

We can take the first-order conditions to get

$$u_l'(x_l) = \lambda p_l.$$

Conditional on the marginal utility of income λ, demand for good l depends only on its own price. We can then obtain the Frisch demand by inverting the marginal utility function.

Definition 90 (Frisch Demand Functions)
The **Frisch demand functions** x_l are in terms of price p and λ and given by:

$$x_l = (u_l')^{-1}(\lambda p_l). \tag{36}$$

Remark 91. *The demand for a good is a function of the price of the good and the marginal utility of income.*

Such demand functions are unrealistic according the empirical literature. We move onto another type of special preferences.

2.2.2 CES Preferences

Definition 92 (CES preferences)
The CES preferences are the sum of sub-utilities:

$$U(x) = \left[\sum_{\ell=1}^{L} \beta_\ell x_\ell^\theta\right]^{1/\theta},$$

where $\beta_\ell > 0$, $\sum_{\ell=1}^{L} \beta_\ell = 1$, and $\theta \le 1$ with $\theta \neq 0$ (the Cobb–Douglas case is obtained as a limit).

Remark 93. *The term β_ℓ is how much does good ℓ affect your utility. Furthermore, due to the ordinal nature of utility functions, we can take monotonic transformations of the CES preferences to make it easier to work with.*

CES preferences have some useful properties.

Proposition 94 (Properties of CES Preferences)
CES preferences are **homothetic**. When $0 < \theta < 1$, they also admit the additively separable representation $\sum_{\ell=1}^{L} \beta_\ell x_\ell^\theta$, which is a strictly increasing transformation of $U(x)$:

$$U(x) = \left[\sum_{\ell=1}^{L} \beta_\ell x_\ell^\theta\right]^{1/\theta}. \tag{37}$$

Proposition 95 (Frisch Demand Function for CES Preferences)
Using the additively separable representation for $0 < \theta < 1$, the Frisch demand function for good ℓ is:

$$x_\ell = \left(\frac{\lambda p_\ell}{\theta \beta_\ell}\right)^{\frac{1}{\theta-1}}. \tag{38}$$

We can then derive the Marshallian demand function from the Frisch demand functions. To do so, we need to multiply by the price p_ℓ and sum over to solve for λ.

$$x_\ell = \alpha_\ell \left(\frac{p_\ell}{P}\right)^{-\sigma} \frac{w}{P} \quad \text{where} \quad \alpha_\ell \equiv \beta_\ell^\sigma \tag{39}$$

whereby P is the **price level**, w is income, and p_ℓ/P is the relative price. We also see that w/P is the level of real income. Furthermore,

$$\sigma \equiv \frac{1}{1-\theta}$$

is the elasticity of substitution.

When $\sigma = 0$, we get back Leontief preferences. When $\sigma = 1$, we get Cobb-Douglas preferences.

The price level P is given by:

$$P(p) = \left[\sum_\ell^L \alpha_\ell p_\ell^{1-\sigma}\right]^{1/(1-\sigma)}.$$

2.3 Indirect Utility Function

We now go back to our utility maximisation problem. We are interested in what is the maximum amount of utility we can receive given a budget set.

Definition 96 (Indirect Utility Function)
The **indirect utility function** $v(p, w)$ is the maximum utility obtainable with prices p and wealth w

$$v(p, w) = \max\{u(x) : p \cdot x \leq w\}. \tag{40}$$

Remark 97. *Indirect utility functions can be derived by inserting the Marshallian demand into the direct utility function* $u(x(p, w))$.

The indirect utility function allows us to express preferences over budget sets. That is, we can assign levels of utilities to the most preferred bundle in a budget set. Then from this, we can compare the highest utility obtainable between different budget sets.

We now state some properties of the Marshallian demand function and indirect utility function.

Proposition 98 (Properties of Marshallian Demand)
The Marshallian demand function $x(p, w)$ satisfy the following properties:

1. Homogeneous of degree zero in (p, w).
2. Satisfies Walras' law.
3. Unique if preferences are **strictly** convex (this implies strictly quasi-concave utility functions).

Proposition 99 (Properties of Indirect Utility Function)
The indirect utility function $v(p, w)$ satisfy the following properties:

1. Homogeneous of degree zero in (p, w).
2. Increasing in wealth w and decreasing in price p.
3. Quasi-convex in (p, w).

Remark 100. *The first condition states that price and budget changes have no effect. Note that to show that it is hod0, you* **must** *increase both prices p and wealth w.*

$$x(\lambda p, \lambda w) = \lambda^0 x(p, w) = x(p, w).$$

The second condition says more money gets us more utility whilst more expensive things gets us less utility.

The final condition says that the **lower contour sets** *are* **convex sets***.*

We can see that the Lagrange multiplier actually has a nice interpretation with regards to the indirect utility function.

Proposition 101 (Lagrange Multiplier As Shadow Price)
The Lagrange multiplier is equal to the consumer's marginal utility of income:

$$\lambda = \frac{\partial v(p, w)}{\partial w}. \tag{41}$$

Proof. First, take derivatives of the indirect utility function:

$$\begin{aligned}
\frac{\partial v(p, w)}{\partial w} &= \frac{\partial u(x(p, w))}{\partial w} \\
&=_{(1)} \sum_{l=1}^{L} \frac{\partial u(x(p, w))}{\partial x_l} \frac{\partial x_l(p, w)}{\partial w} \\
&= \sum_{l=1}^{L} \lambda p_l \frac{\partial x_l(p, w)}{\partial w} \\
&=_{(2)} \lambda
\end{aligned}$$

whereby $=_{(1)}$ uses the chain rule and $=_{(2)}$ uses Engel's aggregation:

$$p \cdot x(p, w) = w \quad \implies \quad p \cdot \frac{\partial x(p, w)}{\partial w} = 1.$$

□

2.4 Expenditure Minimisation

The dual problem to maximising utility subject to a constraint on maximum expenditure is to minimise expenditure subject to a minimum utility constraint. That is, fix an indifference curve and find the cheapest budget set to meet it.

Definition 102 (Expenditure Minimisation Problem)
The goal is to minimise expenditure $p \cdot x$ subject to reaching a given indifference curve $u(x) \geq u$:

$$\begin{aligned} \min_{x} \quad & p \cdot x \\ \text{s.t.} \quad & u(x) \geq u \end{aligned} \tag{42}$$

The Lagrangian is then given by:

$$\mathcal{L}(x, \lambda) = p \cdot x - \lambda(u(x) - u). \tag{43}$$

The $L + 1$ first order conditions are:

$$\mathcal{L}_x = p - \lambda \nabla u(x) = 0 \quad \text{and} \quad \mathcal{L}_\lambda = u(x) - u = 0.$$

The solution is known as the **Hicksian** demand $h(p, u)$.

Definition 103 (Hicksian Demand Function)
The **Hicksian demand** $h(p, u)$ is the solution to the expenditure minimisation problem.

We can define something analogous to the indirect utility function.

Definition 104 (Expenditure Function)
The expenditure function $e(p, u)$ is the minimum amount (cost) needed to achieve at least utility level u with prices p:

$$e(p, u) = \min_{x} \{p \cdot x : u(x) \geq u\}. \tag{44}$$

Based on our definitions of the Hicksian demand and expenditure function, we can obtain the expenditure function from the Hicksian demand by insert the Hicksian demand into the budget constraint:

$$p \cdot h(p, u) = e(p, u).$$

We can derive useful properties of the Hicksian demand and expenditure function.

Proposition 105 (Homogeneity of Hicksian Demand)
The Hicksian demand $h(p, u)$ is homogeneous of degree zero in p.

Proposition 106 (Properties of Expenditure Function)
The expenditure function $e(p, u)$ is:

1. Homogeneous of degree one in p;
2. Increasing in p and in u;
3. **Concave** in p

The concavity of the expenditure function arises because suppose we purchased $h(p, u)$ at price p and prices change to p'. Then the most we will pay is $h(p, u)p'$ and we may find a cheaper alternative instead.

2.5 Duality Theory

For this section, we denote **UMP** as the **Utility Maximisation Problem** and **EMP** as the **Expenditure Minimisation Problem**.

Theorem 107 (Relationship between UMP and EMP)
If preferences are strongly monotonic, then if x^* maximizes the utility $u(\cdot)$ subject to $p \cdot x \leq w$, then setting $x = x^*$ minimises the expenditure $p \cdot x$ subject to $u(x) \geq u(x^*)$. That is,

$$h(p, v(p, w)) = x(p, w). \tag{45}$$

From this relationship, we can derive other useful properties.

Theorem 108 (Expenditure Function and Indirect Utility Function are Inverses)
The expenditure function $e(p, u)$ and the indirect utility function $v(p, w)$ are inverses of each other:

$$e(p, v(p, w)) = w$$
$$v(p, e(p, u)) = u$$

Another thing we can do is derive the indirect utility function from the expenditure function. More specifically, we can invert the expenditure function to get the indirect utility function. If we have an expenditure function $e(p, u)$ then we know that

$$e(p, v(p, w)) = w$$

Then, for a given expenditure function, we can write it as function of p and $v(p, w)$ and then solve for the indirect utility function.

Example 109
Suppose that the expenditure function was:

$$e(p, u) = (p_1 + p_2)u.$$

Then, using our expression for the expenditure function, we have that:

$$(p_1 + p_2)v(p, w) = w$$

and hence we can solve for the indirect utility function:

$$v(p, w) = \frac{w}{p_1 + p_2}.$$

We can also show a relationship between the Marshallian and Hicksian demands.

Theorem 110 (Relationship Between Marshallian and Hicksian Demand)

The Marshallian and Hicksian demands are related by:

$$h(p,u) = x(p, e(p,u))$$
$$x(p,w) = h(p, v(p,w)).$$

The first expression:

$$h(p,u) = x(p, e(p,u))$$

helps to explain why the Hicksian demand is also referred to as **compensated demand**. Here, recall that $e(p,u)$ is the minimum cost to achieve utility level u. Then, as we vary the price p to a new price p', the Hicksian demand $h(p',u)$ will be the Marshallian demand $x(p, e(p',u))$ whereby the Marshallian demand second argument $e(p',u)$ tries to find the minimum cost to maintain the same level of utility u. This cost $e(p',u)$ is then the wealth used to compute the bundle via Marshallian demand that maximises utility $x(p, e(p',u))$. Hence, the difference of this new level of cost needed $e(p',u)$ compared to original cost $e(p,u)$ for maintaining utility level u is the **compensation**.

For Hicksian demands, we have the interesting property that demand and prices move in opposite directions for price changes that are accompanied by Hicksian wealth compensation.

Proposition 111 (Compensated Law of Demand for Hicksian Demand)
For any utility level u and prices p, p', we have that:

$$(p' - p) \cdot (h(p',u) - h(p,u)) \leq 0. \tag{46}$$

Thus the price change and the compensated demand change have a weakly nonpositive inner product.

Remark 112. *Note that this result does not hold for Marshallian demands. This is because the demand for a good can decrease when its price falls in the case of Giffen goods.*

2.5.1 Duality Under Special Preferences

Empirical work often uses expenditure and indirect utility functions.

Theorem 113 (Expenditure and Indirect Utility Function under Homothetic Preferences)
Under homothetic preferences, choose a degree-one homogeneous utility representation. The expenditure and indirect utility functions then have a linear form in utility u:

$$e(p,u) = b(p)u$$
$$v(p,w) = \frac{1}{b(p)}w$$

whereby $b(p)$ is homogeneous of degree one in prices.

Remark 114. *These linear forms depend on the normalisation of utility; they apply to the chosen degree-one homogeneous representation.*

We can derive the Marshallian demand under homothetic preferences since we know the form of the indirect utility function.

Theorem 115 (Form of the Marshallian demand under Homothetic preferences)
Under homothetic preferences, the Marshallian demand is given by:

$$x_l(p, w) = \beta_l(p)w \tag{47}$$

whereby $\beta_l = \frac{1}{b(p)}\partial b(p)/\partial p_l$.

Proof. Use Roy's identity. □

Remark 116. *We have that the wealth w is a linear function of the Marshallian demand so this implies that the Engel curves will be linear rays from the origin.*

We now look at the expressions derived under **Gorman polar form preferences**. Gorman polar forms are a generalisation of **homothetic preferences**. With this in mind, we can see that the following expressions for the indirect utility function and expenditure are similar to what we saw under homothetic preferences except we now have an intercept term too.

Definition 117 (Gorman Polar Form Indirect Utility Function)
Suppose that u is a utility function representing Gorman Polar form preferences. Then, the indirect utility function is of the Gorman Polar form and it can be written as a linear function of wealth w:

$$v(p, w) = \frac{w - a(p)}{b(p)} \tag{48}$$

whereby both $a(p)$ and $b(p)$ are homogeneous of degree one in prices.

Remark 118. *Equivalently, this can be written as an affine function of wealth, $v(p, w) = A(p) + B(p)w$, where $A(p) = -a(p)/b(p)$ and $B(p) = 1/b(p)$.*

If we have the indirect utility function, $v(p, w)$, we can use the relationship between the indirect utility function and expenditure function:

$$v(p, e(p, u)) = u$$

and back out what $e(p, u)$ is. Alternatively, we could recall that $e(p, u) = w$ so then:

$$v(p, w) = \frac{e(p, u) - a(p)}{b(p)}$$

and solve for $e(p, u)$. This gives us the formula for the expenditure function under Gorman polar form.

Definition 119 (Gorman Polar Form Expenditure Function)
Suppose that u is a utility function representing Gorman Polar form preferences. Then, expenditure function is of the Gorman Polar form and it can be written as a linear function of utility u:

$$e(p, u) = a(p) + b(p)u \tag{49}$$

whereby both $a(p)$ and $b(p)$ are homogeneous of degree one in prices.

Now, recall that we have the relationship that if we have the indirect utility function, we can derive the Marshallian demand via **Roy's identity**. Therefore, we can derive the Marshallian demand under Gorman polar form preferences.

Proposition 120 (Form of Marshallian demand under Gorman Polar form)
Under Gorman polar form preferences, the Marshallian demand is given by:

$$x_l(p, w) = \alpha_l(p) + \beta_l(p)w \tag{50}$$

whereby

$$\alpha_l = \frac{\partial a(p)}{\partial p_l} - \beta_l a(p)$$
$$\beta_l(p) = \frac{1}{b(p)} \frac{\partial b(p)}{\partial p_l}.$$

Proof. Apply Roy's identity to the indirect utility function:

$$v(p, w) = \frac{w - a(p)}{b(p)}$$

which is:

$$x_i(p, w) = -\frac{\frac{\partial v(p,w)}{\partial p_i}}{\frac{\partial v(p,w)}{\partial w}}.$$

□

Now why is this helpful? If we compute the income effect of this Marshallian demand, we get:

$$\frac{\partial x(p, w)}{\partial w} = \beta_l(p).$$

Here, the income effect **does not** depend on the level of wealth w. This gives us the result that **the income expansion paths are linear just like homothetic preferences but now we don't necessarily pass through the origin due to the intercept $\alpha(p)$ term.**

Example 121 (Applications to Quasi-Linear Utility)
Recall that quasi-linear utilities are special cases of Gorman polar form. We show how the Gorman polar results reproduce our earlier results for quasi-linear utility. In the two-good case, let y be the numeraire. Then, quasi-linear utility is:

$$u(x) + y$$

whereby $u(\cdot)$ is strictly concave function. We can then compute the Marshallian demand to be:

$$x(p, w) = x(p)$$
$$y(p, w) = w - px(p)$$

whereby demand for good x does not depend on income w, while demand for good y does. The expenditure and indirect utility functions for quasi-linear utility are special cases of the Gorman polar form:

$$e(p, u) = a(p) + b(p)u$$
$$v(p, w) = \frac{w - a(p)}{b(p)}.$$

Define the consumer-surplus function $\phi(p) \equiv \max_x\{u(x) - px\}$. Matching the Marshallian demands derived for quasilinear utility, we can deduce that the expenditure and indirect utility functions are:

$$\begin{aligned} e(p, \bar{u}) &= \bar{u} - \phi(p), \\ v(p, w) &= w + \phi(p). \end{aligned}$$

This matches the Gorman polar form.

2.6 Further Duality Results

We can now show how to derive the Hicksian demand function back from the expenditure function via Shephard's lemma.

Theorem 122 (Shephard's Lemma)
Let $e(p, u)$ be an expenditure function. Then:

$$\frac{\partial e(p, u)}{\partial p_\ell} = h_\ell(p, u). \tag{51}$$

Proof. Apply the Envelope theorem:

$$e(p, u) = p \cdot h(p, u)$$

to get:

$$\frac{\partial e(p, u)}{\partial p_\ell} = h_\ell(p, u)$$

□

Likewise, we can derive the Marshallian demand from the indirect utility function via Roy's identity.

Theorem 123 (Roy's Identity)
Let $v(p, w)$ be the indirect utility function. Then,

$$x_\ell(p, w) = -\frac{\partial v(p, w)/\partial p_\ell}{\partial v(p, w)/\partial w}. \tag{52}$$

Remark 124. *Notice that Roy's identity is not the same form as Shephard's lemma. This is due to income effects that we need to consider.*

We now come to the one of the most important results. The **Slutsky equation** describes the relationship between the slopes of the Hicksian and Marshallian demand function.

Theorem 125 (Slutsky Equation)
Let $x(p, w)$ be the Marshallian demand and $h(p, u)$ be the Hicksian demand. Then, the Slutsky equation is given by:

$$\frac{\partial x_\ell(p,w)}{\partial p_k} = \underbrace{\frac{\partial h_\ell(p,u)}{\partial p_k}}_{\text{Substitution Effect}} \underbrace{- \frac{\partial x_\ell(p,w)}{\partial w} x_k(p,w)}_{\text{Wealth Effect}}. \tag{53}$$

Proof. First, recall that we showed that the Hicksian demand for good l is related to the Marshallian demand for good l by:

$$h_l(p,u) = x_l(p, e(p,u)).$$

Then, we take the partial derivative with respect to the price of good k (with the chain rule):

$$\frac{\partial h_\ell(p,u)}{\partial p_k} = \frac{\partial x_l(p,e(p,u))}{\partial p_k} + \frac{\partial x_l(p,e(p,u))}{\partial w} \underbrace{\frac{\partial e(p,u)}{\partial p_k}}_{(*)}$$

Apply Shephard's lemma to (*) to replace the derivative of the expenditure function with the Hicksian demand to get:

$$\frac{\partial h_\ell(p,u)}{\partial p_k} = \frac{\partial x_l(p,e(p,u))}{\partial p_k} + \frac{\partial x_l(p,e(p,u))}{\partial w} h_k(p,u)$$

Now, we know that as $w = e(p,u)$ then $x_l(p, e(p,u)) = x_l(p,w)$:

$$\frac{\partial h_\ell(p,u)}{\partial p_k} = \frac{\partial x_l(p,w)}{\partial p_k} + \frac{\partial x_l(p,w)}{\partial w} \underbrace{h_k(p,u)}_{(**)}$$

Finally, for (**), we use the relationship between Hicksian and Marshallian demands:

$$h_k(p,u) = x_k(p, e(p,u)) = x_k(p,w)$$

to arrive at the desired result:

$$\frac{\partial h_\ell(p,u)}{\partial p_k} = \frac{\partial x_l(p,w)}{\partial p_k} + \frac{\partial x_l(p,w)}{\partial w} x_k(p,w).$$

□

The income and wealth effects are quite important so we define them separately.

Definition 126 (Substitution Effect)

The **substitution effect** is given by:

$$\frac{\partial h_\ell(p,u)}{\partial p_k} \tag{54}$$

The substitution effect says that the Hicksian demand for good l changes as the price for good k rise, whilst staying on the same indifference curve u.

Definition 127 (Wealth Effect)

The **wealth effect** is given by:

$$\frac{\partial x_\ell(p,w)}{\partial w} x_k(p,w). \tag{55}$$

The wealth effect describes how Marshallian demand changes as the consumer's wealth changes.

Remark 128. *More specifically, for the wealth effect, when the price p_k rises by dp_k, the agent spends an extra amount $x_k(p,w) \times dp_k$ on that good and their wealth falls correspondingly, and this affects the Marshallian demand x_l via $\frac{\partial x_l}{\partial w}$.*

Now, suppose that we are looking at the Slutsky equation for good l when the price of good l itself changes:

$$\frac{\partial x_\ell(p,w)}{\partial p_l} = \frac{\partial h_\ell(p,u)}{\partial p_l} - \frac{\partial x_\ell(p,w)}{\partial w} x_l(p,w). \tag{56}$$

Here, the substitution effect $\frac{\partial h_\ell(p,u)}{\partial p_l}$ is negative whilst for a normal good, the wealth effect is positive. Therefore, since we are adding a negative term to the slope of the Hicksian slope, the Marshallian slope $\frac{\partial x_\ell(p,w)}{\partial p_l}$ is even more negative. We summarise this result in the next proposition.

Proposition 129 (Slope of Marshallian and Hicksian Demand)
The slope of the Marshallian demand is steeper than the slope of the Hicksian demand for a normal good.

We can study the relationship between price changes on the Marshallian and Hicksian demand based off Slutsky equation. First, recall

$$\frac{\partial x_\ell(p,w)}{\partial p_k} = \frac{\partial h_\ell(p,u)}{\partial p_k} - \frac{\partial x_\ell(p,w)}{\partial w} x_k(p,w)$$

Then $\frac{\partial x_\ell(p,w)}{\partial p_k} < \frac{\partial h_\ell(p,u)}{\partial p_k}$ if good l is a **normal good**
Then $\frac{\partial x_\ell(p,w)}{\partial p_k} > \frac{\partial h_\ell(p,u)}{\partial p_k}$ if good l is a inferior **good**
This gives us another useful relationship.

Proposition 130 (Hicksian and Marshallian Demand)
Hicksian demand coincides locally with Marshallian demand if there are no income effects.

Remark 131. *A general Gorman polar demand has an income effect $\partial x_l/\partial w = \beta_l(p)$, so Gorman polar form does not imply that income effects vanish. Quasilinear utility concentrates the income effect on the numeraire, and therefore the non-numeraire goods have no income effects over the range where the solution is interior. Homothetic preferences generally do have income effects because demand is proportional to wealth.*

Proposition 132 (Properties of Slutsky Matrix)
The Slutsky matrix has the following properties:

1. Symmetric;
2. Negative semi-definite.

Proof. It is symmetric due to Young's theorem:

$$\frac{\partial h_l(p,u)}{\partial p_k} = \frac{\partial h_k(p,u)}{\partial p_l}$$

We have that the effect of a price change for product k on compensated demand for product l is the same as a price change for l on product k.

It is negative semi-definite because the expenditure function $e(p,u)$ is concave in prices.

This also implies that $\frac{\partial h_l}{\partial p_l} \leq 0$ (as price increases, Hicksian demand weakly falls).

$\square$

Remark 133. *The significance of negative semi-definiteness means that own-price effects are negative.*

We define other definitions regarding the effect on demand as a result of price changes.

Definition 134 (Normal and Inferior Goods)
A **normal good** has $\partial x_l(p, w)/\partial w \geq 0$, whereas an **inferior good** has $\partial x_l(p, w)/\partial w < 0$. These definitions concern wealth changes, not own-price changes.

Definition 135 (Hicksian Substitute and Complement)
For two distinct goods $l \neq k$, good k is a **Hicksian substitute** for good l if:

$$\frac{\partial h_l(p, u)}{\partial p_k} \geq 0$$

Good k is a **Hicksian complement** for good l if:

$$\frac{\partial h_l(p, u)}{\partial p_k} \leq 0$$

We can express the fact that compensated cross-price effects are symmetric:

$$\frac{\partial h_l(p, u)}{\partial p_k} = \frac{\partial h_k(p, u)}{\partial p_l}$$

Proposition 136 (Existence of Hicksian Substitutes)
With at least two goods, a product has at least one Hicksian substitute under the weak definition above.

Proof. We know that $h_l(p, u)$ is homogeneous of degree zero in p and, by Euler's theorem,

$$\sum_{k=1}^{L} p_k \frac{\partial h_l(p, u)}{\partial p_k} = 0$$

If $\frac{\partial h_l}{\partial p_l} < 0$, this implies that at least one other product has a strictly positive compensated cross-price effect. If the own-price effect is zero, the weak definition of a substitute (a nonnegative cross-price effect) still permits zero effects. □

We define two more useful identities.

Definition 137 (Hotelling-Wold Identity)
The Hotelling-Wold identity is:

$$r_k(x) = \frac{U_k(x)}{\sum_i x_i U_i(x)} \tag{57}$$

whereby $r = \frac{p}{w}$ is the vector of normalised prices.

Definition 138 (Ville-Roy Identity)
The Ville-Roy identity is:

$$x_k(r) = \frac{V_k(r)}{\sum_i r_i V_i(r)} \tag{58}$$

whereby $V_k(r) = \partial V(r)/\partial r_k$ is the derivative of the indirect utility function with respect to normalised price r_k.

To summarise, duality allows us to study demand through expenditure and indirect utility functions, which are often more tractable than the direct utility-maximisation problem.

3 Welfare Economics

3.1 EV and CV

We now want to look at welfare analysis as we change prices. In particular, suppose that we change prices from $p^0 \to p^1$. As a result from this, the consumer's wealth also changes from $w^0 \to w^1$. Under this framework, we can then try to compare the utilities under this price and wealth change:

$$u(x(p^1, w^1)) - u(x(p^0, w^0)).$$

The issue is that, because utility is an ordinal concept, this difference can tell us only whether the consumer is better or worse off, not *how much* better off they are. We therefore work with *expenditure functions*, which are measured in monetary units. We shall also use geometry to assist us when doing welfare analysis.

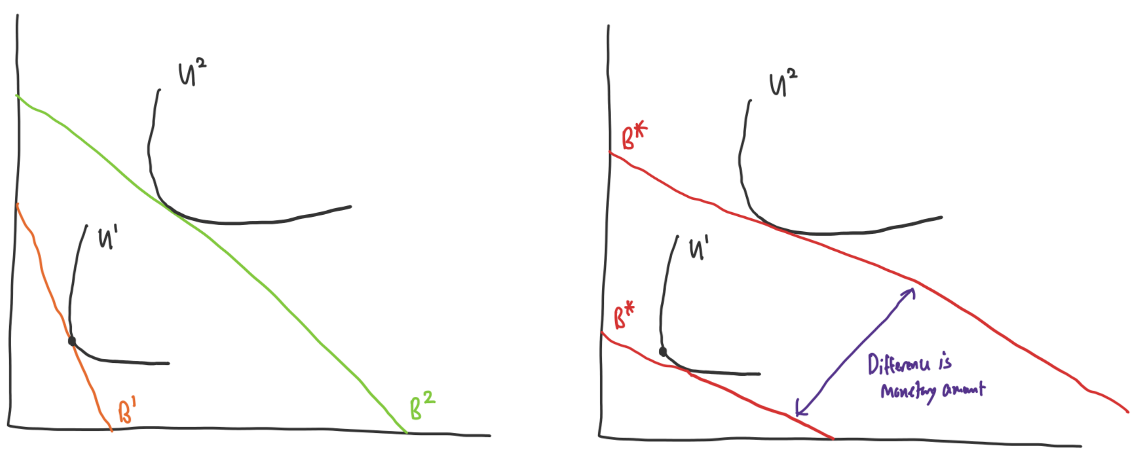

Figure 13: The graph on the left shows two indifference curves associated to the Marshallian demand with respect to two different budget sets B^1, B^2. The graph on the right shows the same indifference curves but now the same translated budget set B^* intersect the two indifference curves. The difference between the translated budget sets represent a monetary amount.

We now give the intuition used throughout this section. Refer to Fig. 13. On the left, budget set B^1 generates an initial Marshallian choice and utility level. A change in prices and wealth gives a new budget set B^2, a new Marshallian choice, and a new utility level. Because utility is ordinal, the difference between the two utility numbers is not a meaningful measure of magnitude. We therefore pick reference prices $\overline{p}$. At those prices, draw the two parallel expenditure-minimising budget lines tangent to the initial and final indifference curves, as on the right of Fig. 13. Each line records the minimum expenditure needed to attain its utility level at $\overline{p}$. The difference between these expenditures is a monetary measure of the welfare change. This comparison, based on $e(\overline{p}, v(p, w))$, will be used repeatedly throughout the section.

Therefore, we have that for a general price change $p^0 \to p^1$ and wealth change $w^0 \to w^1$, we can use the differences in expenditure functions:

$$e(\overline{p}, u^1) - e(\overline{p}, u^0) \tag{59}$$

where $u^1 = v(p^1, w^1)$ and $u^0 = v(p^0, w^0)$. This is a measure of welfare change expressed in monetary terms: it is the difference in the minimum spending needed to reach the old and new welfare levels at the reference prices $\overline{p}$.

Naturally, the reference price $\overline{p}$ we can use are either the initial prices p^0 or the new prices p^1. If we do use these prices as the reference prices, we end up with **equivalent** and **compensating** variation.

The next two concepts, *equivalent and compensating variation*, express the consumer's welfare change using the two choices of reference prices.

Definition 139 (Equivalent Variation)
Equivalent variation is the change in expenditure using **old prices** p^0 as the reference price $\overline{p}$:

$$EV = e(p^0, u^1) - e(p^0, u^0). \tag{60}$$

where $u^1 = v(p^1, w^1)$ and $u^0 = v(p^0, w^0)$

Remark 140. *EV measures how much it costs, at the original prices p^0, to move from the old indifference curve to the new one.*

Definition 141 (Compensating Variation)
Compensating variation is the change in expenditure using **new prices** p^1 as the reference price $\overline{p}$:

$$CV = e(p^1, u^1) - e(p^1, u^0) \tag{61}$$

where $u^1 = v(p^1, w^1)$ and $u^0 = v(p^0, w^0)$.

Remark 142. *A way to remember which is which is that compensating variation uses the prices* **after** *the change, which in this case are the new prices p^1.*

Now, note that given optimising behaviour, what people spend **is** the minimum cost of getting onto that indifference curve. Hence, the value of the expenditure function is known:

$$e(p^0, u^0) = w^0$$
$$e(p^1, u^1) = w^1$$

As a result, we can rewrite the Equivalent and Compensating variation as:

$$EV = e(p^0, u^1) - w^0$$
$$CV = w^1 - e(p^1, u^0)$$

Remark 143. *Textbooks often specialise to the case in which wealth does not change, so that*

$$w^0 = w^1 = w.$$

As a result, since agents are optimising and minimising their cost to get to indifference curves u^0 and u^1, we have that:

$$e(p^0, u^0) = e(p^1, u^1) = w.$$

As a result, EV and CV may then be rewritten as:

$$EV = e(p^0, u^1) - e(p^0, u^0) = e(p^0, u^1) - e(p^1, u^1)$$
$$CV = e(p^1, u^1) - e(p^1, u^0) = e(p^0, u^0) - e(p^1, u^0)$$

These expressions are algebraically equivalent to the definitions above because the optimised expenditure in each observed state equals w. They provide a second interpretation of the same monetary measures: EV values the price change at final utility u^1, whereas CV values it at initial utility u^0.

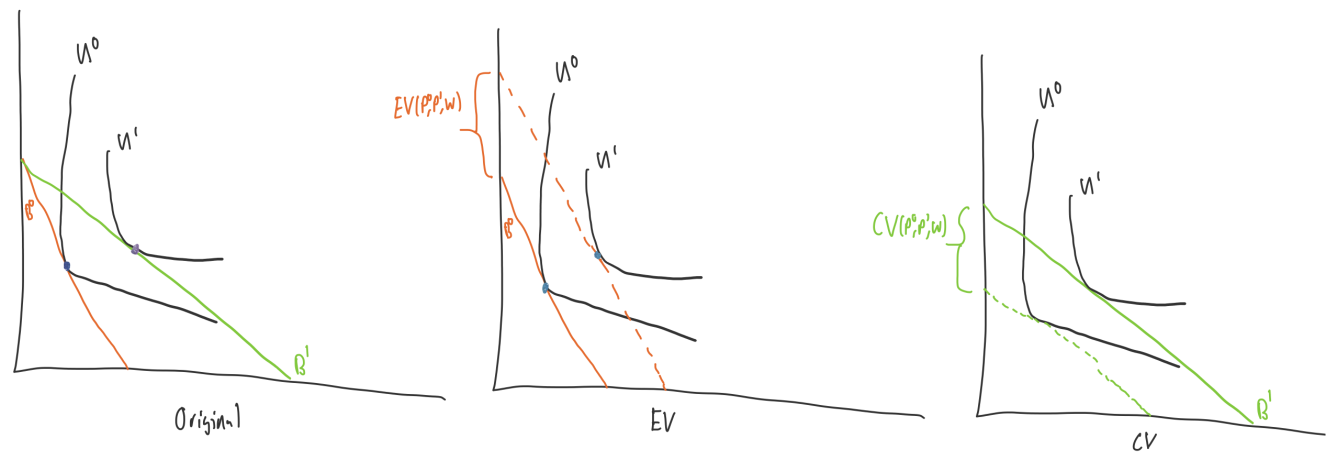

Figure 14: The left panel shows two indifference curves at two different budget sets, (p^0, w) and (p^1, w). The middle panel represents EV using the original prices p^0, while the right panel represents CV using the new prices p^1.

We now give a graphical example of EV and CV in Fig. 14. The figure on the left plots two indifference curve at two different budgets sets $(p^0, w), (p^1, w)$. We make the crucial assumption that the price of good 2 is 1 and income is fixed. Now, if we use the reference price to be the original price p^0, then, due to the crucial assumption we made, the EV is the difference on the y-axis where the budget sets intersect the y-axis. For the figure on the right, we use the new price p^1 as the reference price and again due to the crucial assumption, the difference on the y-axis is the CV. If we did not make these assumptions, then the EV and CV would be the tangential distances that we described earlier.

Proposition 144 (EV and CV Need Not Coincide)
EV and CV rank welfare changes in the same direction but need not give the same monetary amount.

We give a graphical proof of this in Fig. 15. □

Based on the previous proposition and Fig. 15, EV and CV coincide only under additional restrictions.

Proposition 145 (EV and CV Under Homothetic Preferences)
Homothetic preferences do not generally make EV and CV coincide. With a degree-one homogeneous utility representation, $e(p, u) = b(p)u$, so

$$EV = b(p^0)(u^1 - u^0), \qquad CV = b(p^1)(u^1 - u^0). \tag{62}$$

Proof.

Figure 15: The distances obtained under EV and CV are not the same.

Thus, for a non-zero welfare change, they coincide only when $b(p^0) = b(p^1)$.

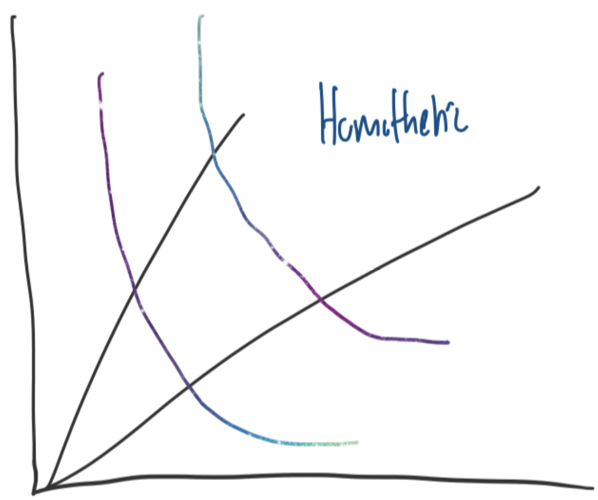

Proof.

Figure 16: Homothetic preferences are radial expansions.

Homothetic indifference curves are radial expansions, as seen in Fig. 16, but their monetary separation depends on the reference-price index $b(p)$. $\square$

Proposition 146 (Equivalent Variation and Hicksian Demand)
Suppose that wealth is fixed, $w^0 = w^1 = w$, and the price of good 1 changes $p_1^0 \to p_1^1$. Then, the equivalent variation is given by:

$$EV = \int_{p_1^1}^{p_1^0} h_1(p_1, \overline{p}_{-1}, u^1)\, dp_1 \tag{63}$$

whereby all other prices are held constant, $\overline{p}_{-1} = (\overline{p}_2, ..., \overline{p}_L)$.

Proof. We can relate equivalent variation to Hicksian demand. To do so, we can apply the Fundamental Theorem of Calculus to Shephard's lemma: □

Remark 147. *The area under the Hicksian demand before and after the price change in good 1 is the EV.*

Proposition 148 (Compensating Variation and Hicksian Demand)
Suppose that wealth is fixed, $w^0 = w^1 = w$, and the price of good 1 changes $p_1^0 \to p_1^1$. Then, the compensating variation is given by:

$$CV = \int_{p_1^1}^{p_1^0} h_1(p_1, \overline{p}_{-1}, u^0)\, dp_1 \tag{64}$$

whereby all other prices are held constant, $\overline{p}_{-1} = (\overline{p}_2, ..., \overline{p}_L)$.

Remark 149. *What is important to note is that EV is the area under Hicksian demand under u^1 whilst CV is the area under Hicksian demand under u^0.*

We now give a graphical interpretation of the relationship between EV/CV and Hicksian demand under the case where we only change the price of one good p_1.

For a **single** price change, EV and CV can be computed from Hicksian demand curves. We explain this in a few steps.

Refer to Fig. 17. Suppose we start off from the price p_1^0. We then plot a Marshallian demand curve now (in contrast to indifference curves) $x_1(p_1, \overline{p}_{-1}, w)$ whereby we have the Marshallian demand for good 1 on the x-axis and the price of good 1 on the y-axis. Now, suppose that we change the price to p_1^1 so we then move along the Marshallian demand curve to the new point.

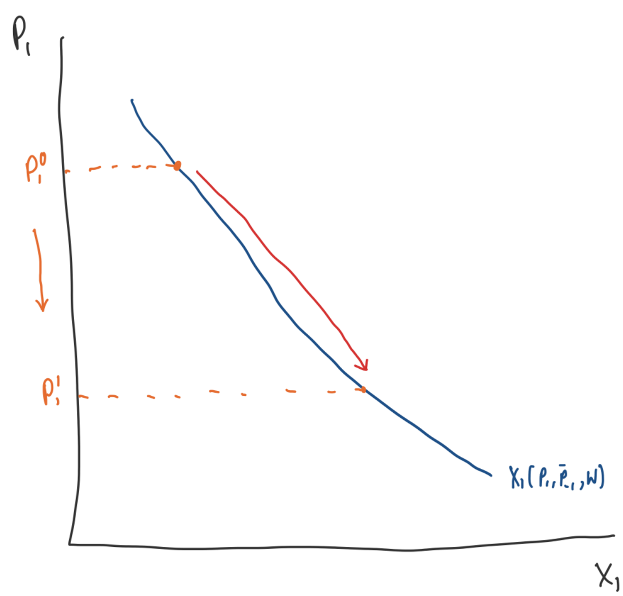

Figure 17: Marshallian demand curve as we move price from $p_1^0 \to p_1^1$.

However, the relevant welfare areas lie under Hicksian rather than Marshallian demand. When a good's price falls at fixed income, the consumer becomes better off because indirect utility is decreasing in prices. Thus, movement along a Marshallian demand curve involves movement between Hicksian demand curves indexed by different utility levels.

Refer to Fig. 18. At the initial price p_1^0, Hicksian demand at initial utility u^0 is $h_1(p_1, \overline{p}_{-1}, u^0)$. At the new price p_1^1, the price fall raises utility to u^1, giving the Hicksian demand curve $h_1(p_1, \overline{p}_{-1}, u^1)$. Hence, under this price fall, the consumer moves to a higher-utility Hicksian demand curve.

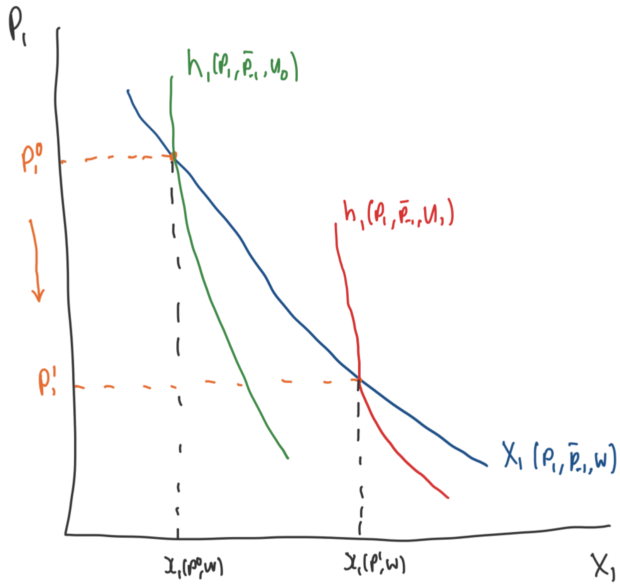

Figure 18: Marshallian demand curve as we move price from $p_1^0 \to p_1^1$.

Now, we think about why are these Hicksian demand curves not the same as the Marshallian demand curves and this is due to the **Slutsky equation**. Here, by the Slutsky equation, the slope of the Marshallian demand curve is steeper than the slope of the Hicksian demand curve:

$$\frac{\partial x_\ell(p,w)}{\partial p_k} = \underbrace{\frac{\partial h_\ell(p,u)}{\partial p_k}}_{\text{Substitution Effect}} - \underbrace{\frac{\partial x_\ell(p,w)}{\partial w} x_k(p,w)}_{\text{Wealth Effect}}.$$

or more loosely, this can be written as:

$$\underbrace{\frac{dX}{dP}}_{\text{Marshallian Slope}} = \underbrace{\frac{dH}{dP}}_{\text{Hicksian slope}} - \text{income effect}$$

Now, the Marshallian demand curve will be steeper than the Hicksian demand curve in $X - P$ (quantity-price) space as seen in the left figure of Fig. 19. However, we normally put price on the y-axis and hence we need to flip the axis around, which we did in the right figure of Fig. 19. As a result, the Marshallian demand curve will now be **flatter** than the Hicksian demand curve under the $P - X$ (price-quantity) space. Therefore, the Hicksian demand curve is steeper with respect to demand compared to the Marshallian demand curve.

We can now bring all this together to compute EV and CV. EV is the area under the integral of the Hicksian demand curve under the new utility u^1 whilst CV is the area under the integral of the Hicksian demand curve under the old utility u^0. We show this in Fig. 20.

From Fig. 20, we can see some things immediately.

Proposition 150 (EV and CV)

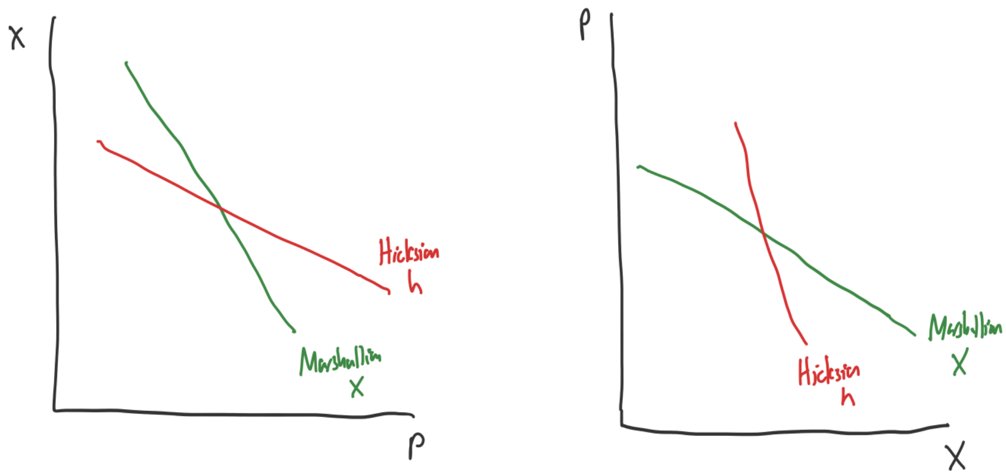

Figure 19: Marshallian demand curve and Hicksian demand curves on $X - P$ and $P - X$ space.

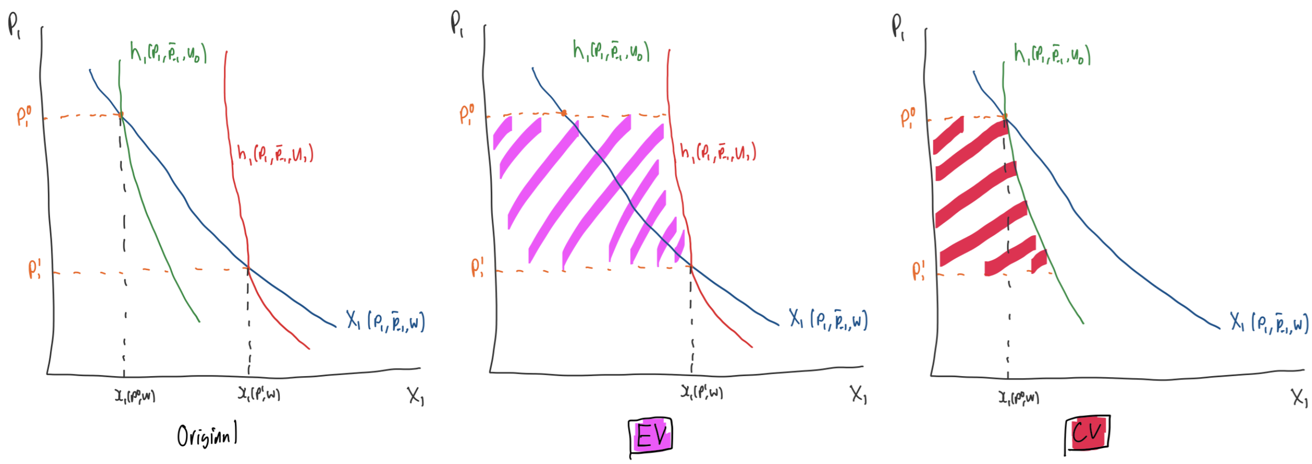

Figure 20: EV and CV are the areas under the Hicksian demand curves.

For a price fall $p_1^1 < p_1^0$ at fixed wealth, if good 1 is normal over the relevant range, then

$$\mathrm{EV} > \mathrm{CV}. \tag{65}$$

Corollary 151

The reverse is true if a good is an inferior good.

From this, we immediately get another result.

Proposition 152 (EV and CV under No Wealth Effects)

Equivalent Variation is equal to Compensating variation when there is no wealth effect.

Proof. This is due to the Slutsky equation as we now have that the derivative of the Marshallian demand is equal to the derivative of the Hicksian demand:

$$\frac{\partial x_\ell(p,w)}{\partial p_k} = \underbrace{\frac{\partial h_\ell(p,u)}{\partial p_k}}_{\text{Substitution Effect}}.$$

□

Remark 153. *This is the same as saying that Marshallian demand is equal to Hicksian demand if there are no wealth effects.*

We can also define the area under the Marshallian demand curve.

Definition 154 (Consumer Surplus)
Consumer surplus is the area under the inverse Marshallian demand curve and above the market price. For a price change, its change is measured by the corresponding integral of Marshallian demand.

Based on our earlier discussion, we can relate consumer surplus to EV and CV.

Proposition 155 (Consumer Surplus under no Wealth Effects)
Suppose that we have no wealth effects, then we have that:

$$\text{Consumer Surplus} \;=\; \text{EV} = \;\text{CV}. \tag{66}$$

Remark 156. *Consumer-surplus changes are exact money-metric welfare measures when there are no income effects. With income effects, they are generally approximations rather than exact EV or CV measures.*

Under **quasi-linear utility**, income effects are concentrated on the numeraire. Consequently, demand for the non-numeraire good has no income effect (while the solution remains interior), making quasi-linearity a popular assumption in welfare analysis.

3.2 Price Indices

Previously, under EV/CV, we compared indifference curves at fixed prices and expressed this difference in terms of money. Now, under cost-of-living indices, we will compare **prices** on a fixed **indifference curve** and express this in terms of **unit-free ratios**.

Definition 157 (Konüs Cost-of-living Index)
Suppose that we changed prices from p to p'. Then, the Konüs cost of living index is the ratio of expenditure functions:

$$P(p, p', u) = \frac{e(p', u)}{e(p, u)}. \tag{67}$$

Intuitively, this compares two different prices against a fixed indifference curve. That is, what is the **relative cost** of reaching a fixed indifference curve u under alternative prices p and p'. Notice the contrast to before where under EV and CV, we held prices constant and looked at the change in utility.

Graphically, suppose that $h(p,u) = x(p,w)$ are the initial demands. Then,

$$e(p,u) = p \cdot h(p,u).$$

The minimum cost of maintaining utility u when prices change $p \to p'$ is:

$$e(p', u) = p' \cdot h(p', u).$$

That is, we fix an indifference curve, plot two budget sets and see the expenditure functions as we rotate around the indifference curve.

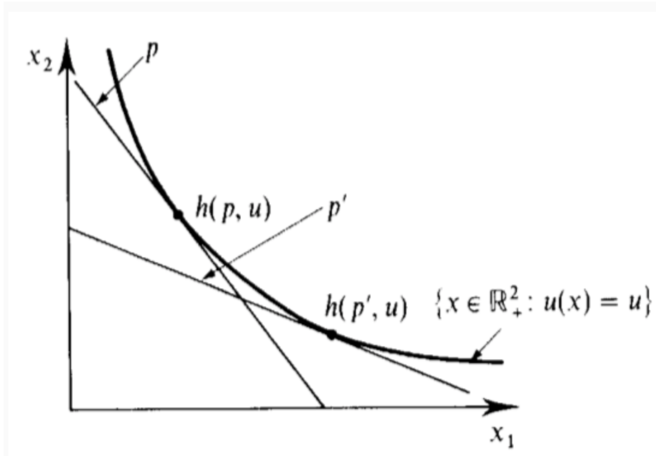

Figure 21: Konüs cost of living index.

Now, just like we needed to specify a reference prices under EV and CV, for cost of living indices, we need to specify a reference utility curve. This gives us the Laspeyres and Paasche index.

Definition 158 (Laspeyres index)
Suppose that prices changed from $p^0 \to p^1$. If we use the **initial utility** u^0 as the reference utility, the base-utility Konüs index is:

$$P_K^0(p^0, p^1, u^0) \equiv \frac{e(p^1, u^0)}{e(p^0, u^0)} = \frac{e(p^1, u^0)}{w^0}. \tag{68}$$

Definition 159 (Paasche index)
Suppose that prices changed from $p^0 \to p^1$. If we use the **final utility** u^1 as the reference utility, the current-utility Konüs index is:

$$P_K^1(p^0, p^1, u^1) \equiv \frac{e(p^1, u^1)}{e(p^0, u^1)} = \frac{w^1}{e(p^0, u^1)}. \tag{69}$$

The utility level you choose for the reference has an effect because how you value goods will affect the cost of living for you. However, we know when the reference utility does not matter.

Proposition 160 (Cost-of-Living Index under Homothetic Preferences)
Under homothetic preferences, the cost-of-living index is independent of the reference utility:

$$\frac{e(p^1, u)}{e(p^0, u)} = \frac{b(p^1)}{b(p^0)}. \tag{70}$$

Proof. When we have homothetic preferences, and therefore linear expenditure functions

$$e(p,u) = b(p)u.$$

Then, we plug this into our cost of living index:

$$\frac{e(p^1,u)}{e(p^0,u)} = \frac{b(p^1)u}{b(p^0)u} = \frac{b(p^1)}{b(p^0)}$$

$\square$

Remark 161. *Under homothetic preferences, Engel's law does not apply. As a result, price increases will be the same for everybody.*

In practice, agencies don't calculate the cost of staying on the same indifference curve but instead the cost of staying on the same consumption bundle. Hence, in practice, we compute a price index rather than a cost of living index. The idea of a price index is as follows. You fix a consumption bundle and calculate the cost of acquiring this same bundle under different prices. Hence, we have that:

$$\{p^0, x^0\} \to \{p^1, x^0\}.$$

Therefore, if we examine the cost of maintaining a fixed initial consumption bundle x^0 at different price points, we end up with a new index called the **Laspeyres price index**.

Definition 162 (Laspeyres Price Index)
Suppose we observe a price change $p^0 \to p^1$. Then, the **Laspeyres price index** is given by:

$$P_L(p^0, p^1, x^0) \equiv \frac{p^1 \cdot x^0}{p^0 \cdot x^0}. \tag{71}$$

Remark 163. *Notice that we have fixed x^0 so we are not allowing for substitution effects.*

There is an issue with the price index though. Under Slutsky's equation, recall that when we change the price of a good, substitution effects occur. However, we are no longer allowing for substitution effects as we are fixing the consumption bundle x^0 to remain the same. Hence, we can in fact show that the Laspeyres price index is a biased estimator of the cost of living index.

Proposition 164 (Laspeyres Price Index is Biased Upwards)
The Laspeyres price index P_L approximates the cost-of-living index P **from above**:

$$P_L(p^0, p^1, x^0) = \frac{p^1 \cdot x^0}{p^0 \cdot x^0} \geq \frac{e(p^1, u^0)}{e(p^0, u^0)} = P(p^0, p^1, u^0). \tag{72}$$

Proof. First, recall that for the Laspeyres price index P_L, we are fixing the Marshallian demand and hence not allowing for substitution effects. However, the *true* cost-of-living index measures the ratio of expenditure functions keeping utility constant. Therefore, as we change the price of a good, due to optimising behaviour of agents, the agents may substitute away from the good whose price has gone up and into other cheaper goods. This can cause the overall price to fall. Hence, the ratio of expenditure functions for the cost of living index can at most be the same as the Laspeyres price index.

We now mathematically prove this. The denominators of the Laspeyres price index and the cost of living index are the same due to optimising behaviour. That is,

$$p^0 \cdot x^0 = w^0 = p^0 \cdot x(p^0, w^0) = p^0 \cdot h(p^0, u^0) = e(p^0, u^0).$$

Therefore, as the denominators are the same, the comparison concerns the numerator. Expanding the cost-of-living numerator $e(p^1, u^0)$ around the original price p^0 gives:

$$e(p^1, u^0) = e(p^0, u^0) + \underbrace{(p^1 - p^0) \cdot \nabla_p e(p^0, u^0)}_{(*)} + R,$$

where concavity of the expenditure function in prices implies that the remainder $R \leq 0$. If we look at the first-order term (*), we can apply Shephard's lemma:

$$(p^1 - p^0) \cdot \nabla_p e(p^0, u^0) = (p^1 - p^0) \cdot h(p^0, u^0) = (p^1 - p^0) \cdot x(p^0, w^0)$$

We plug this into (*) and simplify:

$$\begin{aligned}
&e(p^1, u^0) = e(p^0, u^0) + \underbrace{(p^1 - p^0) \cdot \nabla_p e(p^0, u^0)}_{(*)} + R \\
\leftrightarrow\ &e(p^1, u^0) = e(p^0, u^0) + (p^1 - p^0) \cdot x(p^0, w^0) + R \\
\leftrightarrow\ &e(p^1, u^0) = p^0 \cdot x^0 + (p^1 - p^0) \cdot x(p^0, w^0) + R \\
\leftrightarrow\ &e(p^1, u^0) = p^1 \cdot x(p^0, w^0) + R
\end{aligned}$$

We can now interpret the expression:

$$e(p^1, u^0) = p^1 \cdot x(p^0, w^0) + R.$$

This says that the cost of attaining utility u^0 at the new prices p^1 is no greater than the cost of purchasing the original bundle. The remainder is nonpositive because the expenditure function is concave in prices (equivalently, its Hessian is the negative-semidefinite Slutsky matrix).

Therefore, writing out what we have derived so far, the true cost-of-living index is:

$$P(p^0, p^1, u^0) = \frac{e(p^1, u^0)}{e(p^0, u^0)} = \frac{e(p^1, u^0)}{p^0 \cdot x^0} = \frac{p^1 \cdot x^0}{p^0 \cdot x^0} + \frac{R}{p^0 \cdot x^0}$$

So we have that the true cost of living index is:

$$P(p^0, p^1, u^0) = \frac{p^1 \cdot x^0}{p^0 \cdot x^0} + \frac{R}{p^0 \cdot x^0} = P_L(p^0, p^1, x^0) + \frac{R}{p^0 \cdot x^0}$$

where the first term is the Laspeyres index and the second term is a negative term due to the Slutsky matrix being in R. Therefore, we can conclude that the Laspeyres index overestimates the true cost of living index. □

Hence, we have seen that the over-estimation by the Laspeyres index of the true cost of living index arises due to substitution effects in the Slutsky matrix S:

$$\frac{p^1 \cdot x^0}{p^0 \cdot x^0} \geq \frac{e(p^1, u^0)}{e(p^0, u^0)}.$$

Thus, if there are no substitution effects, the Slutsky matrix is zero and the indices coincide. An example is Leontief preferences. More generally, larger compensated substitution responses can create a larger gap between the Laspeyres price index and the true cost-of-living index.

Definition 165 (Paasche Price Index)
Suppose we observe a price change $p^0 \to p^1$. Then, the **Paasche price index** is given by:

$$P_P(p^0, p^1, x^1) = \frac{p^1 \cdot x^1}{p^0 \cdot x^1}. \tag{73}$$

We have a similar result whereby the Paasche index **understates** the true cost of living index.

Proposition 166 (Paasche Price Index is Biased Downwards)
The Paasche price index P_P approximates the cost-of-living index P **from below**:

$$P_P(p^0, p^1, x^1) = \frac{p^1 \cdot x^1}{p^0 \cdot x^1} \leq \frac{e(p^1, u^1)}{e(p^0, u^1)} = P(p^0, p^1, u^1). \tag{74}$$

We can use the Paasche and Laspeyres index to bound the true cost of living index.

Proposition 167 (Two-Sided Bound on Cost-of-Living Index)
Under homothetic preferences, the Paasche and Laspeyres indices give two-sided bounds on the cost-of-living index:

$$\frac{p^1 \cdot x^1}{p^0 \cdot x^1} \leq \frac{e(p^1, u)}{e(p^0, u)} \leq \frac{p^1 \cdot x^0}{p^0 \cdot x^0}. \tag{75}$$

We conclude by stating when the price indices equal the true cost-of-living index.

Proposition 168 (Equality Between Price Indices and Cost of Living Index)
With no substitution effects, the Paasche and Laspeyres indices are equal to the true cost-of-living index.

3.3 Quantity Indices

Recall that for cost-of-living indices, we expressed these as the comparison of the cost of reaching a given indifference curve as a ratio whereby we vary the prices whilst keeping the indifference curve fixed:

$$P(p^0, p^1, u) = \frac{e(p^1, u)}{e(p^0, u)}.$$

Quantity indices are also index numbers except we now hold prices fixed and compare the indifference curves:

$$Q(p, u^0, u^1) = \frac{e(p, u^1)}{e(p, u^0)}.$$

We now formally define this.

Definition 169 (Allen Quantity index/Real Income Measures)
The quantity index is given by:

$$Q(p, u^0, u^1) \equiv \frac{e(p, u^1)}{e(p, u^0)}. \tag{76}$$

Remark 170. *This is the converse of cost-of-living indices!*

Analogous to before, we can choose either the initial prices p^0 or the new prices p^1 to be the reference price. As a result, this is quite similar to EV/CV except we are now expressing the difference in terms of a ratio rather than a cash amount.

The choice of reference price defines quantity indices analogous to price indices.

Definition 171 (Laspeyres Quantity Index)

Choose the reference price to be p^0. The Laspeyres quantity index is given by:

$$Q_L(u^0, u^1, p^0) \equiv \frac{e(p^0, u^1)}{e(p^0, u^0)}. \tag{77}$$

Definition 172 (Paasche quantity index)

Choose the reference price to be p^1. The Paasche quantity index is given by:

$$Q_P(u^0, u^1, p^1) \equiv \frac{e(p^1, u^1)}{e(p^1, u^0)}. \tag{78}$$

Similar to price indices, the Laspeyres and Paasche quantity indices are also biased.

Proposition 173 (Laspeyres Quantity Index Is Upward Biased)

The Laspeyres quantity index estimates the true quantity index **from above**:

$$Q_L(u^0, u^1, p^0) \equiv \frac{e(p^0, u^1)}{e(p^0, u^0)} = \frac{e(p^0, u^1)}{w^0} \leq \frac{p^0 \cdot x^1}{p^0 \cdot x^0} \equiv Q_L(x^0, x^1, p^0). \tag{79}$$

Proposition 174 (Paasche Quantity Index Is Downward Biased)

The Paasche quantity index estimates the true quantity index **from below**:

$$Q_P(u^0, u^1, p^1) \equiv \frac{e(p^1, u^1)}{e(p^1, u^0)} = \frac{w^1}{e(p^1, u^0)} \geq \frac{p^1 \cdot x^1}{p^1 \cdot x^0} \equiv Q_P(x^0, x^1, p^1). \tag{80}$$

We can now define a general definition for a price and quantity index.

Definition 175 (General Price and Quantity Index)

A price and quantity index are a pair of functions that satisfy the function equation:

$$P(p^0, p^1, x^0, x^1)Q(p^0, p^1, x^0, x^1) = \frac{w^1}{w^0}. \tag{81}$$

Under this definition, if we knew the price index $P(p^0, p^1, x^0, x^1)$, then we can immediately work out the quantity index:

$$Q(p^0, p^1, x^0, x^1) = \frac{w^1}{w^0} \frac{1}{P(p^0, p^1, x^0, x^1)}.$$

We can then see that applying this to the cost-of-living and Allen indices, we get:

$$\underbrace{\frac{e(p^1, u^0)}{e(p^0, u^0)}}_{\text{base-utility price index}} \underbrace{\frac{e(p^1, u^1)}{e(p^1, u^0)}}_{\text{current-price quantity index}} = \frac{e(p^1, u^1)}{e(p^0, u^0)} = \frac{w^1}{w^0}.$$

The Allen index is therefore the current-weighted implicit quantity index defined using a base-weighted price index:

$$\frac{e(p^1, u^1)}{e(p^1, u^0)} = \frac{w^1}{w^0} \frac{1}{\frac{e(p^1,u^0)}{e(p^0,u^0)}}.$$

Proposition 176 (Implicit Quantity Index of the Laspeyres and Paasche Index)

The implicit quantity index defined using a Laspeyres price index is the Paasche quantity index.

The implicit quantity index defined using a Paasche price index is the Laspeyres quantity index.

Definition 177 (Fisher Index)

The Fisher price index is the geometric mean of the Laspeyres and Paasche price index:

$$P_F = \sqrt{P_L \cdot P_P}. \tag{82}$$

Remark 178. *The implicit index of the Fisher index is the direct index.*

3.4 Equivalence Scales

So far, our analysis of demand only relied on substitution and wealth effects. Clearly though, household demographics and characteristics affect demand behaviour. We now incorporate these features into our models.

Definition 179 (Equivalence Scale)

An **equivalence scale** measures the cost of living for a household of a given size and demographic composition relative to the cost of living of a reference household when **both households attain the same level of utility**.

Remark 180. *We can think of an equivalence scale as trying to compare the household with given characteristics to a baseline household.*

Definition 181 (Conditional expenditure cost function)

The conditional expenditure cost function $e(p, u, z)$ is the minimum expenditure necessary to attain utility level u, conditional on the consumer having characteristics z.

Definition 182 (Equivalence Scale)

An equivalence scale compares the conditional expenditure cost functions of different characteristics:

$$\frac{e(p, u, z^1)}{e(p, u, z^0)}. \tag{83}$$

An issue though is that equivalence scales are not identifiable. As an example, we know that due to utility functions being invariant to monotonic transformations, the equivalent scales:

$$\frac{e(p, u, z^1)}{e(p, u, z^0)} \quad \text{and} \quad \frac{e(p, \phi(u, z^1), z^1)}{e(p, \phi(u, z^0), z^0)}$$

for a monotone function ϕ can be consistent with the same demand behavior yet be quite different numerically. Hence, we need to use other methods. The solution to Equivalence scales is conditional cost of living indices.

Definition 183 (Conditional Cost of Living Indices)
The conditional cost of living index is under characteristics z^0 is:

$$P(p^0, p^1, u, z^0) = \frac{e(p^1, u, z^0)}{e(p^0, u, z^0)}. \tag{84}$$

The conditional cost of living index is under characteristics z^1 is:

$$P(p^0, p^1, u, z^1) = \frac{e(p^1, u, z^1)}{e(p^0, u, z^1)}. \tag{85}$$

Remark 184. *For example, z^0 could be a single household and z^1 could be a family of 4 household.*

We are interested in comparing households using real equivalised income which adjusts incomes between households and allows for fair comparison. To do so, we require two things:

1. A price index to remove the differences in prices facing households.
2. An equivalence scale to take out differences in characteristics of the households.

Definition 185 (Relative Equivalence Scale)
The relative equivalence scale is:

$$\frac{e(p^1, u, z^1)}{e(p^1, u, z^0)} = \frac{P(p^0, p^1, u, z^1)}{P(p^0, p^1, u, z^0)} \frac{e(p^0, u, z^1)}{e(p^0, u, z^0)}. \tag{86}$$

Remark 186. *This is the ratio of household-type-specific cost-of-living indices times the corresponding base-price-regime scale.*

Finally, we compare two households using the concept of real equivalised income. Suppose that households i and j have conditional expenditure functions:

$$\begin{aligned} w^i &= e(p^i, u^i, z^i), \\ w^j &= e(p^j, u^j, z^j). \end{aligned}$$

What we want to do is to compare the households if they faced the same prices p^0 and had the same characteristics z^0.

$$\begin{aligned} v^i &= e(p^0, u^i, z^0), \\ v^j &= e(p^0, u^j, z^0). \end{aligned}$$

where we use the fact that the expenditure function is a positive monotonic transformation of the utility function.

We can show how to turn a given conditional expenditure function to the standardised one we were interested in. To do so, we multiply and divide accordingly to get:

$$e(p^0, u^i, z^0) = \underbrace{\frac{e(p^0, u^i, z^0)}{e(p^0, u^i, z^i)}}_{\text{equivalence-scale adjustment}} \underbrace{\frac{e(p^0, u^i, z^i)}{e(p^i, u^i, z^i)}}_{\text{cost-of-living adjustment}} e(p^i, u^i, z^i).$$

Now, we use the fact that $e(p^i, u^i, z^i) = w^i$ via optimising behaviour. Hence, to obtain standardised expenditure, we apply an equivalence-scale adjustment (which removes differences in characteristics) and a cost-of-living adjustment (which removes differences in prices):

$$e(p^0, u^i, z^0) = \underbrace{\frac{e(p^0, u^i, z^0)}{e(p^0, u^i, z^i)}}_{\text{equivalence-scale adjustment}} \underbrace{\frac{e(p^0, u^i, z^i)}{e(p^i, u^i, z^i)}}_{\text{cost-of-living adjustment}} w^i.$$

We can follow a similar procedure for household j and hence compare the households.

4 2-Stage Budgeting and Aggregation

4.1 2-Stage Budgeting

We now move onto the problem of two-stage budgeting. Suppose that the agent is facing a large number of different goods in different categories. The big idea behind two-stage budgeting is that it is a utility maximising problem in two stages. In the first stage, the agent will allocate their budget across different **categories** such as food, clothes, vehicles etc. Then, in the second stage, they then choose what goods within those categories to purchase with the category budget they allocated.

Definition 187 (Two Stage Budgeting)
First, the consumer allocates her spending across broad aggregates. Second, the consumer allocates budgets across different items in the budget allotment groups.

Doing this is a form of dimension reduction whereby we have reduced the dimension of the problem. However, to do this, we require an important assumption.

Assumption (Weak Separability of Preferences)
Partition goods into two bundles (x, z) and prices associated (p, q). If preferences are **weakly separable** in the $x-$goods, then the utility function can be written as:

$$u(v(x), z) \tag{87}$$

whereby $v(x)$ is a **sub-utility function** and $u(v, z)$ is an increasing function of v.

Suppose that we did **not** carry out two-stage budgeting and instead applied our usual method of computing the Marshallian demand for the x and z goods. From this, we get the standard Marshallian demands:

$$x(p, q, w)$$
$$z(p, q, w)$$

for the x and z goods respectively.

Now, we switch back to the two-stage budgeting problem. If the overall utility function is separable in the x-goods, we can just solve for the demand over the x-goods. Suppose that we have already carried out the first stage so the allocation of the budget w that goes to the x-goods is denoted by w_x. Then, we can find the optimal choice of x-goods by solving the **sub-problem** (the second stage):

$$\begin{aligned} \max_{x} \quad & v(x) \\ \text{s.t.} \quad & p \cdot x = w_x. \end{aligned}$$

where w_x is the allocation of the overall budget to the x-goods from the first stage. Now, what is important to note is that the expenditure on x-goods from the second stage, should be equal to the expenditure on the x-goods in the original Marshallian demand $x(p, q, w)$. Mathematically, this means:

$$p \cdot x = p \cdot x(p, q, w).$$

Furthermore, we have the even stronger result that the demand for the x-goods as a function of the prices of the x-goods and the x-goods budget should be equal:

$$x(p, q, w) = x(p, w_x)$$

whereby the ordinary Marshallian demand $x(p, q, w)$ equals the second-stage conditional demand $x(p, w_x)$. Conditional on the group budget w_x, this demand no longer depends directly on the prices of the z-goods. However, those prices can still affect the first-stage allocation w_x and therefore affect unconditional demand. Empirically, if w_x is observed, only data on the x-goods are needed to estimate the conditional demand system for that group.

Definition 188 (Conditional demand function)
The conditional demand function $x(p, w_x)$ is the optimal demand given the prices of the goods in this group conditional on the level of budget allocated to the group.

4.2 Aggregation

So far we have looked at individual agents' preferences and demands. We now ask what happens when we aggregate these demands. More concretely, suppose we have I consumers with idiosyncratic preferences $\succsim_i$ and Marshallian demands $x_i(p, w_i)$. We assume that they face the same prices. Aggregate demand is the sum of individual Marshallian demands.

Definition 189 (Aggregate Demand)
Let there be I consumers each with idiosyncratic preferences $\succsim_i$ and Marshallian demands $x_i(p, w_i)$. Then, **aggregate demand** is

$$x(p, w_1, \cdots, w_I) = \sum_{i=1}^{I} x_i(p, w_i). \tag{88}$$

We can state some immediate properties of aggregate demand which arises from the properties of individual Marshallian demands.

Proposition 190 (Properties of Aggregate Demand)
Let $x(p, w_1, ..., w_I)$ denote aggregate demand. Then $x(p, w_1, ..., w_I)$ satisfies the following properties:

1. Aggregate demand is homogeneous of degree zero when prices and all individual wealth levels are scaled together;

2. Aggregate demand satisfies Walras' law;

3. Aggregate demand is continuous if individual demand functions are continuous.

Remark 191. *With a continuum of heterogeneous consumers, aggregation can smooth individual discontinuities under additional regularity conditions. Continuity is not automatic for a finite collection of discrete individual demands.*

We now ask whether the distribution of wealth matters for aggregate demand and, if not, under what conditions. We have defined aggregate demand as:

$$x(p, w_1, \cdots, w_I) = \sum_{i=1}^{I} x_i(p, w_i)$$

where individual wealth w_i can affect aggregate demand. We ask when it can instead be written as:

$$x(p, w_1, \cdots, w_I) = x(p, \sum_{i=1}^{I} w_i)$$

so that only aggregate wealth $\sum_{i=1}^{I} w_i$ matters. We now state the conditions under which this is the case.

Proposition 192 (When Aggregate Demand Depends Only On Aggregate Wealth)
Aggregate demand depends only on aggregate wealth when it is invariant to aggregate wealth-preserving changes in the wealth distribution

$$\sum_{i=1}^{I} \frac{\partial x_{li}(p, w_i)}{\partial w_i} dw_i = 0 \tag{89}$$

for all goods $l = 1, ..., L$ and all redistributions satisfying $\sum_{i=1}^{I} dw_i = 0$.

To interpret the equation, write:

$$\sum_{i=1}^{I} \underbrace{\frac{\partial x_{li}(p, w_i)}{\partial w_i}}_{(*)} \underbrace{dw_i}_{(**)} = 0.$$

The (*) term is individual i's income effect, while (**) is the change in that individual's wealth. Across consumers, the resulting demand changes must net to zero for every good l.

We can in fact state the necessary and sufficient condition for aggregate demand to be invariant to aggregate wealth-preserving changes in the wealth distribution as outlined in Eq. (89).

Proposition 193 (Conditions for Wealth Invariance of Aggregate Demand)
Fix any initial wealth distribution $(w_1, \ldots, w_I)$. For every redistribution $(dw_1, \ldots, dw_I)$ satisfying $\sum_{i=1}^{I} dw_i = 0$, aggregate demand is invariant to the redistribution if

$$\sum_{i=1}^{I} \frac{\partial x_{li}(p, w_i)}{\partial w_i} dw_i = 0 \tag{90}$$

for all goods $l = 1, ..., L$ if and only if

$$\frac{\partial x_{li}(p, w_i)}{\partial w_i} = \frac{\partial x_{lj}(p, w_j)}{\partial w_j} \tag{91}$$

for all goods $l = 1, ..., L$ and all consumers $i, j = 1, \ldots, I$.

Let us breakdown Eq. (91):

$$\frac{\partial x_{li}(p, w_i)}{\partial w_i} = \frac{\partial x_{lj}(p, w_j)}{\partial w_j}.$$

This says that the **wealth/income effects** for every agent, good by good, **are the same**. Recall that the wealth effect is the slope of an **Engel curve**. Thus agents' Engel curves have the same slopes, although their intercepts may differ.

The intuition is that when we take $1 from one person and give it to another, their wealth effects must cancel. This changes the distribution of wealth but not total wealth, i.e. $\sum_{i=1}^{I} dw_i = 0$. For aggregate demand to remain unchanged, the reduction in one person's demand must therefore be offset by the increase in another person's demand.

Finally, we ask when the slopes of Engel curves are the same at all levels of wealth. This occurs when everyone's preferences have the **Gorman polar form**, with a common slope in wealth:

$$v_i(p, w_i) = a_i(p) + b(p)w_i$$

This leads us to a useful result.

Proposition 194 (Conditions for Aggregate Demand to Be a Function of Aggregate Wealth)
Subject to the usual regularity conditions, aggregate demand can be written as a function of prices and aggregate wealth when everyone's preferences have a **Gorman polar form** with the same wealth coefficient:

$$v_i(p, w_i) = a_i(p) + b(p)w_i \tag{92}$$

where v_i is the indirect utility function of agent i.

Let us examine the Gorman Polar form more carefully:

$$v_i(p, w_i) = a_i(p) + b(p)w_i.$$

Here, we note that all agents i have the same slope which is a function of prices, $b(p)$. Hence, the differences in agent's preferences can only arise in the intercepts of these affine functions $a_i(p)$. Hence, all agents will have the same slope but different intercepts.

Finally, we ask whether individual-level rationality aggregates. That is, does the weak axiom of revealed preference hold in aggregate? Not necessarily. Even if individuals' actions are rational, aggregate demand need not satisfy the weak axiom; conversely, aggregate behaviour can look rational even when individual behaviour does not. Thus, **aggregate rationality alone does not identify individual agents' rationality**.

To summarise, a representative-agent demand depending only on aggregate income and prices requires strong restrictions: individual preferences must have Gorman polar form with common income effects. Without these restrictions, aggregate demand generally depends on both aggregate wealth and its distribution. We must therefore be careful when drawing inferences about individuals from aggregate data.

5 Choice Under Uncertainty

So far, we have been looking at choosing between bundles of goods. Now, we will be choosing between probability distributions. In order to do so, we need to define preferences over probability distributions. However, we also require to impose more structure on our utility functions as a result, which will culminate into what is known as the expected utility representation theorem.

5.1 Preferences Over Lotteries

We are no longer assuming that the outcome of decisions are known in advance. We need to start thinking probabilistically. Therefore, we need to define a few new concepts.

Definition 195 (Outcomes)
The set of **possible outcomes** is denoted by $\mathcal{C}$. We shall assume that $\mathcal{C}$ is a **finite** set and that the possible outcomes are indexed by $n = 1, 2, ..., N$.

Remark 196. *$\mathcal{C}$ can represent monetary payoffs or consumption bundles. It plays the role of the outcome space in probability theory.*

With these possible outcomes, we can associate probabilities to these outcomes.

Definition 197 (Simple Lottery)
A **simple lottery** L is a tuple of probabilities $L = (p_1, ..., p_N)$ whereby p_i is the probability that outcome i occurs.

Remark 198. *A simple lottery is identical to the idea of a probability mass function for a discrete random variable!*

Proposition 199 (Lotteries are Probability Measures)
A simple lottery obeys the axioms of probability:

1. $p_n \geq 0$ for all $n = 1, ..., N$;
2. $\sum_{n=1}^{N} p_n = 1$.

Remark 200. *Compound lotteries are lotteries over lotteries!*

In a simple lottery, the outcomes are certain. So we have probabilities over certain outcomes. However, we can also have lotteries where the outcomes are themselves lotteries. This has a special name.

Definition 201 (Compound Lottery)
A **compound lottery** is a lottery whose outcomes are themselves simple lotteries.

Suppose that we have K simple lotteries $L_k = (p_1^k, ..., p_N^k)$ for $k = 1, ..., K$. Suppose that we have the probabilities $\alpha_k \geq 0$ and $\sum_{k=1}^{K} \alpha_k = 1$. The compound lottery is denoted by:

$$\sum_{k=1}^{K} \alpha_k L_k \tag{93}$$

where there is a probability α_k for simple lottery L_k occurring.

Remark 202. *For example, a coin flip may determine which of two subsequent dice lotteries is played.*

Suppose we now have compound lotteries. It can be difficult to compare between compound lotteries. What can be done is to reduce compound lotteries into simple lotteries. The following theorem lets us do so.

Theorem 203 (Simple Representation of Compound Lotteries)
Every compound lottery with K simple lotteries $\sum_{k=1}^{K} \alpha_k L_k$ can be represented as a simple lottery where the probability of outcome n is the probability-weighted sum

$$p_n = \sum_{k=1}^{K} \alpha_k p_n^k.$$

Corollary 204
Certainty can be represented as a special case of a lottery with degenerate probabilities.

Definition 205 (Set of All Possible Simple Lotteries)
Given N possible outcomes $\mathcal{C}$, the set $\mathcal{L}$ denotes the set of all possible simple lotteries. It is the $(N-1)$-dimensional simplex consisting of all $p \in \mathbb{R}_+^N$ with $\sum_{i=1}^{N} p_i = 1$.

Before we can define utility functions, we need to first define continuity which is needed.

Definition 206 (Continuity Axiom)
The preference relation $\succsim$ on the space of simple lotteries $\mathcal{L}$ is **continuous** if, for any $L, L', L'' \in \mathcal{L}$ and sequence $\{\alpha_i\}_{i=1}^{\infty} \subset [0,1]$ such that $\alpha_i \to \alpha$,

$$\alpha_i L + (1-\alpha_i)L' \succsim L'' \Rightarrow \alpha L + (1-\alpha)L' \succsim L''$$

and

$$L'' \succsim \alpha_i L + (1-\alpha_i)L' \Rightarrow L'' \succsim \alpha L + (1-\alpha)L'$$

Remark 207. *Alternatively, we can say that the following sets are* **closed sets***:*

$$\{\alpha \in [0,1] : \alpha L + (1-\alpha)L' \succsim L''\} \subset [0,1]$$

$$\{\alpha \in [0,1] : L'' \succsim \alpha L + (1-\alpha)L'\} \subset [0,1]$$

We assume that agents have preferences $\succsim$ **over lotteries** in $\mathcal{L}$. That is, they can compare probability distributions.

Theorem 208 (Utility Function Representation)
Suppose that $\succsim$ represents rational and continuous preferences. Then, there is a continuous utility function $U(\cdot)$ defined on $\mathcal{L}$ such that $L \succsim L'$ if and only if

$$U(L) \geq U(L').$$

5.2 The Independence Axiom

We now see our first big departure from consumer theory, which is the independence axiom. This is quite similar to additive separability that we saw in consumer theory.

Assumption 209 (Independence Axiom)
The preference relation $\succsim$ on the space of simple lotteries $\mathcal{L}$ satisfies the **independence axiom** if for all $L, L', L^{''} \in \mathcal{L}$ and $\alpha \in (0, 1)$, we have that

$$L \succsim L' \quad \Leftrightarrow \quad \alpha L + (1-\alpha)L^{''} \succsim \alpha L' + (1-\alpha)L^{''}. \tag{94}$$

Let us break this down. On the left hand side, we have that

$$L \succsim L'$$

which is a comparison between simple lotteries. This says that we prefer simple lottery L to simple lottery L'. For the right hand side, we take our simple lotteries and construct **two compound lotteries** involving a new simple lottery $L^{''}$:

$$\alpha L + (1-\alpha)L^{''} \quad \text{and} \quad \alpha L' + (1-\alpha)L^{''}.$$

We say that we still prefer the compound lottery with the preferred simple lottery L to the compound lottery with the simple lottery L'. This assumption imposes that preferences are of a particular linear form. If we drop the independence axiom, we lose expected utility theory.

Proposition 210 (Indifference and Strict Preferences Properties of Lotteries)
Notions of indifference and strict preferences follow from the independence axiom:

$$L \sim L' \quad \Leftrightarrow \quad \alpha L + (1-\alpha)L'' \sim \alpha L' + (1-\alpha)L''$$

and

$$L \succ L' \quad \Leftrightarrow \quad \alpha L + (1-\alpha)L'' \succ \alpha L' + (1-\alpha)L''$$

Finally, if $L \succ L'$, then

$$pL + (1-p)L' \succ qL + (1-q)L' \quad \Leftrightarrow \quad p > q.$$

Remark 211. *This third property says that compound lotteries with a higher probability of a simple lottery we prefer, this compound lottery is preferred to other compound lotteries with lower weights on the simple lottery we prefer.*

For consumer theory, separability is quite a strong assumption to make as consumption of goods depend on each other. However, in choice under uncertainty, this is more reasonable as only one state of the world is ever realised at a given point in time.

Theorem 212 (Affine Utility over Lotteries)
A rational and continuous preference relation that also satisfies the independence axiom admits a utility representation that is affine in probability mixtures.

5.3 Expected Utility Theorem

We now state the most important result of this section.

Theorem 213 (Expected Utility Theorem)
Suppose that preferences $\succsim$ over lotteries are:

1. Rational;
2. Continuous;
3. Satisfy the independence axiom.

Then there exists a utility function representation for the preferences $\succsim$ which takes the expected utility form:

$$L \succsim L' \quad \Leftrightarrow \quad \sum_{n=1}^{N} p_n u_n \geq \sum_{n=1}^{N} p'_n u_n \tag{95}$$

in which the utility of the degenerate lottery yielding outcome n for certain is u_n.

The theorem says that preferences over lotteries can be represented by the probability-weighted sum of utilities of outcomes. A decision maker chooses the lottery with the highest expected utility. The entire distribution can matter through this expectation; for nonlinear u, changes in variance and other features of the distribution can change expected utility.

Proof. First, we have a preference relation $\succsim$ over all lotteries $L \in \mathcal{L}$. Rationality allows us to compare any pair of lotteries. Because $\mathcal{L}$ is compact and preferences are continuous, there is a **best** lottery B and a worst lottery W. Clearly, $B \succsim W$. Any other lottery $L \in \mathcal{L}$ lies between these two in the preference ordering:

$$W \precsim L \precsim B.$$

If $B \sim W$, all lotteries are indifferent and a constant expected-utility representation suffices. Otherwise $B \succ W$, and we proceed as follows. We assign utility levels to the best and worst lotteries by defining:

$$u(B) = 1 \quad \text{and} \quad u(W) = 0.$$

We now look at any intermediate lottery $L \in \mathcal{L}$. We **claim** that we can write this intermediate lottery as a **compound lottery** of the best and worst lottery:

$$L \sim pB + (1-p)W$$

whereby p **exists** and is **unique**. In other words, we want to be able to find a probability p such that the agent is indifferent between the lottery L and the compound lottery comprising of B and W. Our goal is now to show that p exists and is unique.

Claim 214. *The probability p exists.*

Proof. We need to show existence of the probability measure p. First, look at the two sets:

$$\{p \in [0,1] : pB + (1-p)W \succsim L\}$$

$$\{p \in [0,1] : L \succsim pB + (1-p)W\}$$

These sets are closed by continuity and nonempty: for $p = 1$, $B \succsim L$, and for $p = 0$, $L \succsim W$. Furthermore, every point in $[0,1]$ belongs to at least one of the two sets. Since the unit interval is **connected**, the two closed sets must intersect. Hence, such a p exists. □

Claim 215. *The probability p is unique.*

Proof. Suppose that p is not unique. This means that distinct p and p' both satisfy $L \sim pB + (1-p)W$ and $L \sim p'B + (1-p')W$. However, by the independence axiom, $pB + (1-p)W$ is preferred to $p'B + (1-p')W$ if and only if $p > p'$, since $B \succ W$. Thus two distinct mixtures cannot both be indifferent to L, which is a contradiction. ☐

We have shown that p exists and is unique. We now give the structure of the expected utility. The utility of the lottery L, denoted by $u(L)$ is defined to be the number (probability) p such that the agent is indifferent to having the lottery L and the compound lottery with the best outcome with probability p and worst outcome with probability 1 - p. Mathematically, we define:

$$u(L) = p \tag{96}$$

whereby p is such that

$$L \sim pB + (1-p)W$$

or equivalently under our definition for $u(L) = p$:

$$L \sim u(L)B + (1-u(L))W$$

Now, if we take any two lotteries L, L' and probability $0 < \alpha < 1$, we have that

$$\alpha L + (1-\alpha)L'.$$

We can again express both of these lotteries as compound lotteries involving the best and worst lotteries:

$$L \sim u(L)B + (1-u(L))W \tag{97}$$

$$L' \sim u(L')B + (1-u(L'))W \tag{98}$$

We can then take a compound lottery over our two lotteries L, L' (which themselves are compound lotteries) to get (after doing a bunch of algebra):

$$\alpha L + (1-\alpha)L' \sim \big[\alpha u(L) + (1-\alpha)u(L')\big]B + \big(1 - \big[\alpha u(L) + (1-\alpha)u(L')\big]\big)W \tag{99}$$

The first thing that is important is to notice that on the right hand side of Eq. (99), the coefficients on B and W satisfy:

$$\big[\alpha u(L) + (1-\alpha)u(L')\big] \in [0,1]$$

$$1 - \big[\alpha u(L) + (1-\alpha)u(L')\big] \in [0,1]$$

These are actually probability measures themselves whereby the sum of the two expressions gives us back one!

$$\Big(\big[\alpha u(L) + (1-\alpha)u(L')\big]\Big) + \Big(1 - \big[\alpha u(L) + (1-\alpha)u(L')\big]\Big) = 1$$

So if you look carefully, our expression in Eq. (99) is actually of the form:

$$\alpha L + (1-\alpha)L' \sim pB + (1-p)W \tag{100}$$

where now, we have that:

$$p = \big[\alpha u(L) + (1-\alpha)u(L')\big]$$

and

$$1 - p = 1 - \left[\alpha u(L) + (1-\alpha)u(L')\right]$$

Go back to our definition of a utility function in Eq. (96)! We said that the utility of the lottery L, denoted by $u(L)$ is defined to be the number (probability) p such that the agent is indifferent to having the lottery L and the compound lottery with the best outcome with probability p and worst outcome with probability 1 - p. Mathematically:

$$u(L) = p$$

whereby p is such that

$$L \sim pB + (1-p)W$$

or equivalently under our definition for $u(L) = p$:

$$L \sim u(L)B + (1-u(L))W$$

Apply the definition of utility to Eq. (100). The compound lottery $\alpha L + (1-\alpha)L'$ has associated probability $p = \alpha u(L) + (1-\alpha)u(L')$ of the best outcome. Therefore,

$$u(\alpha L + (1-\alpha)L') = \alpha u(L) + (1-\alpha)u(L').$$

We can therefore see that $u(\cdot)$ is linear in compound lotteries:

$$u(\alpha L + (1-\alpha)L') = \alpha u(L) + (1-\alpha)u(L').$$

This held for the case of 2 lotteries. We can easily generalise this to hold for N lotteries. Hence, we can conclude that

$$u(L) = \sum_{n=1}^{N} p_n u_n \tag{101}$$

where u_n is the utility of the degenerate lottery that gives outcome n for certain.

□

Therefore, with this expected utility structure, we can simply compute the expected utility in order to see what lottery do we prefer.

Definition 216 (von Neumann-Morgenstern Expected Utility Function)
The utility function $U : \mathcal{L} \to \mathbb{R}$ has an **expected utility form** if there is an assignment of numbers $(u_1, ..., u_N)$ to the N outcomes such that for every simple lottery $L = (p_1, ..., p_N) \in \mathcal{L}$, we have that

$$U(L) = u_1 p_1 + \cdots + u_N p_N \tag{102}$$

A utility function $U : \mathcal{L} \to \mathbb{R}$ with the expected utility form is called a **von Neumann–Morgenstern** expected utility function. The values u_i are the Bernoulli utilities of the outcomes.

Remark 217. *A von Neumann–Morgenstern utility representation is unique up to a positive affine transformation. General increasing transformations need not preserve the expected-utility form.*

5.4 Risk Aversion

We now consider lotteries over monetary outcomes x. If agents have expected utility preferences, then there is a utility function $u(x)$ such that they prefer lottery 1 to lottery 2 whenever:

$$\mathbb{E}_1[u(x)] > \mathbb{E}_2[u(x)]$$

where $\mathbb{E}_i$ refers to taking expectations with respect to the distribution corresponding to lottery $i = 1, 2$.

Definition 218 (Risk-Averse Agent)
An expected-utility maximiser is **risk averse** if they weakly prefer the sure mean of any lottery to the lottery itself:

$$\mathbb{E}[u(X)] \leq u(\mathbb{E}[X]). \tag{103}$$

This is **Jensen's inequality** for concave functions $u(\cdot)$.

Remark 219. *The expected utility is less than the utility of the average value.*

We can derive this inequality for a risk-averse agent with a concave utility function $u(\cdot)$. First, denote the mean of the distribution x by $\overline{x} = \mathbb{E}[x]$. A differentiable concave function lies below each of its tangent lines. Therefore, for all x, we have:

$$u(x) \leq u(\overline{x}) + (x - \overline{x})u'(\overline{x})$$

We then take expectations of both sides:

$$\begin{aligned}
\mathbb{E}[u(x)] &\leq \mathbb{E}[u(\overline{x})] + \mathbb{E}[(x - \overline{x})u'(\overline{x})] \\
&= u(\overline{x}) + (\mathbb{E}[x] - \overline{x})u'(\overline{x}) \\
&= u(\overline{x}).
\end{aligned}$$

which gives us:

$$\mathbb{E}[u(X)] \leq u(\mathbb{E}[X]).$$

Remark 220. *Strict concavity implies that, absent hedging or endowment effects, an agent strictly dislikes a non-degenerate fair gamble relative to its sure mean.*

In Fig. 22, we plot a risk-averse agent's utility function.

We can map what we have done so far to consumer theory as seen in Fig. 23. Suppose we have two states of the world with payoffs x_1 and x_2. Subject to a fixed expected payoff, a risk-averse agent prefers to smooth payoffs across states. The constant expected-payoff line $px_1 + (1-p)x_2 = K$ acts like a budget line. The optimum lies on the 45-degree certainty line, where wealth is the same in both states. At that point, the marginal rate of substitution is $MRS = \frac{p}{1-p}$, the odds ratio.

Definition 221 (Certainty Equivalent)
Given a particular lottery, the **certainty equivalent** of a lottery is the monetary amount C such that:

$$u(C) = \mathbb{E}[u(X)] = \int u(x)f(x)dx = \int u(x)dF(x) \tag{104}$$

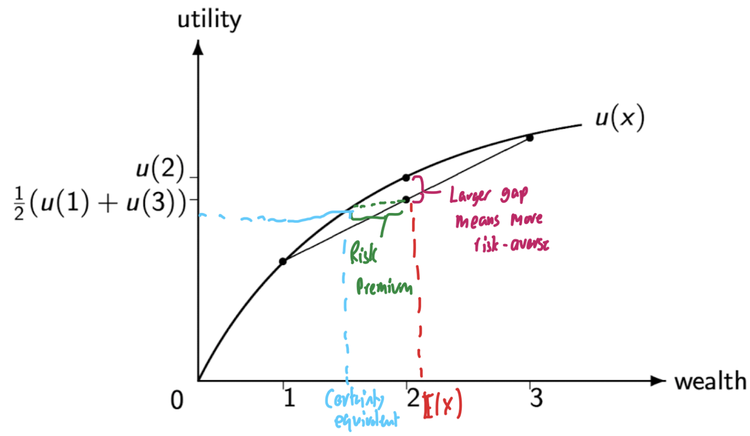

Figure 22: Risk-averse agent with concave utility function.

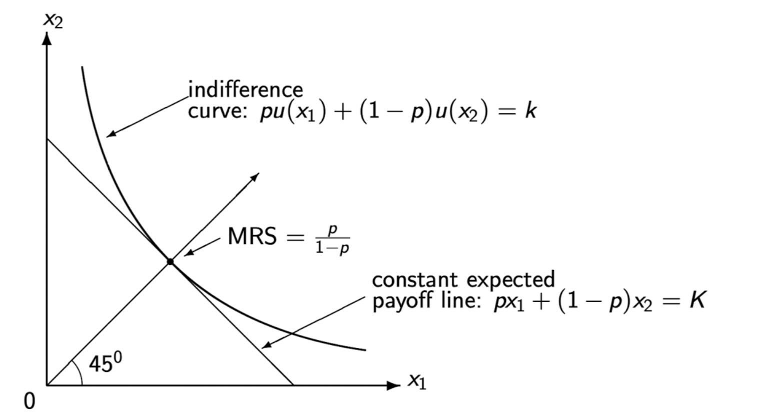

Figure 23: Expected utility indifference curve tangential to constant expected payoff line over two states of the world.

This is the guaranteed amount of money that makes the agent indifferent to the lottery.

Proposition 222 (Certainty Equivalent for Risk-Averse Agent)

For a risk-averse agent,

$$C \leq \mathbb{E}[X]$$

where C is the certainty equivalent.

We can then ask how much is the agent willing to remove all risk.

Definition 223 (Risk Premium)

The **risk premium** P is the amount of money the agent is willing to pay to eliminate all risk in the lottery so that

$$u(\mathbb{E}[X] - P) = \mathbb{E}[u(X)] \tag{105}$$

Remark 224. *P measures how much the agent is willing to pay for full insurance.*

This gives us a useful relationship.

Proposition 225 (Risk Premium and Certainty Equivalent)
The following relationship holds:

$$P + C \equiv \mathbb{E}[X] \tag{106}$$

The higher the curvature of the utility function $u(c)$, the higher the risk aversion of the agent. However, since expected utility functions are not uniquely defined (as they are defined only up to affine transformations), a measure that stays constant with respect to these transformations is needed rather than just the second derivative $u''(c)$. One measure that is used is absolute risk aversion.

We are now interested in understanding how much an agent is willing to pay to eliminate a small amount of risk.

Definition 226 (Coefficient of Absolute Risk Aversion)
The **coefficient of absolute risk aversion** is defined as:

$$r_A(x) \equiv \frac{-u''(x)}{u'(x)} > 0 \tag{107}$$

whereby recall that $u'' < 0$ due to concavity.

Proof. We can derive where this expression comes from.

Let z be a random variable such that $\mathbb{E}[z] = 0$ and $Var[z] = \sigma^2$. Let x_0 be the **risk-free wealth** and t be the risky scaling factor. Therefore, total wealth x is given by:

$$x = x_0 + tz.$$

Now note that the first moment of total wealth is:

$$\mathbb{E}[x] = \mathbb{E}[x_0 + tz] = \mathbb{E}[x_0] + t\mathbb{E}[z] = \mathbb{E}[x_0] = x_0.$$

Now, recall the risk-premium formula:

$$u[\mathbb{E}[x] - p] = \mathbb{E}[u(x)]$$

Define $p \equiv p(t)$ to be the risk-premium as a function of risk whereby $p(0) = 0$ as $t = 0$ so there is no risk.

We then apply our expression for $\mathbb{E}[x]$ to get:

$$u[\mathbb{E}[x] - p(t)] = \mathbb{E}[u(x)]$$

and using our expression for the first moment of x and the formula for x, we get:

$$u[x_0 - p(t)] = \mathbb{E}[u(x_0 + tz)]$$

We now differentiate with respect to t to get:

$$-u'(x_0 - p(t))p'(t) = \mathbb{E}[u'(x_0 + tz)z] \tag{108}$$

where we assume that we differentiate inside the expectation operator. Now, evaluate this at $t = 0$ to get:

$$\begin{aligned} -u'(x_0 - p(0))p'(0) &= \mathbb{E}[u'(x_0 + 0z)z] \\ \leftrightarrow -u'(x_0)p'(0) &= u'(x_0)\mathbb{E}[z] = 0. \end{aligned}$$

Now, since $u'(\cdot)$ is a positive value, this implies that $p'(0) = 0$.

We now differentiate Eq. (108) again with respect to t to get:

$$u''(x_0 - p(t))\big(p'(t)\big)^2 - u'(x_0 - p(t))p''(t) = \mathbb{E}[z^2 u''(x_0 + tz)]$$

We let t = 0 to get:

$$\begin{aligned} u''(x_0 - p(0))\big(p'(0)\big)^2 - u'(x_0 - p(0))p''(0) &= \mathbb{E}[z^2 u''(x_0 + 0z)] \\ \leftrightarrow u''(x_0 - p(0))\big(p'(0)\big)^2 - u'(x_0 - p(0))p''(0) &= \mathbb{E}[z^2 u''(x_0)] \\ \leftrightarrow u''(x_0 - p(0))\big(0\big)^2 - u'(x_0 - p(0))p''(0) &= \mathbb{E}[z^2 u''(x_0)] \\ \leftrightarrow -u'(x_0)p''(0) &= \mathbb{E}[z^2 u''(x_0)] \\ \leftrightarrow -u'(x_0)p''(0) &= u''(x_0)\mathbb{E}[z^2] \\ \leftrightarrow -u'(x_0)p''(0) &= u''(x_0)\sigma^2 \end{aligned}$$

which gives us the result:

$$p''(0) = -\frac{u''(x_0)}{u'(x_0)}\sigma^2$$

So now, we look at the risk-premium $p(t)$ and Taylor expand it about $t = 0$ to get:

$$\begin{aligned} p(t) &\approx p(0) + p'(0)t + \frac{1}{2}p''(0)t^2 \\ &= \frac{1}{2}\left[-\frac{u''(x_0)}{u'(x_0)}\sigma^2\right]t^2. \end{aligned}$$

Therefore, the risk-premium the agent is willing to pay is:

$$p(t) \approx -\frac{1}{2}\frac{u''(x_0)}{u'(x_0)}(t\sigma)^2$$

So the agent is willing to pay to eliminate risk as a function of their risk-aversion and the variance of the outcome $V = (t\sigma)^2$. □

We now explain why this measure is not affected by a positive affine transformation. Suppose we have a utility function $u(c_t) = \ln(c_t)$ and take the transformation:

$$g(u(c_t)) = \alpha + \beta \ln(c_t), \qquad \beta > 0.$$

Now, if we compute the absolute risk aversion of this affine transformed utility function:

$$\begin{aligned} g(u(c_t))' &= \frac{\beta}{c_t} \\ g(u(c_t))'' &= -\frac{\beta}{c_t^2}. \end{aligned}$$

The resulting measure of absolute risk aversion is:

$$r_A(x) = -\left(\frac{-\frac{\beta}{c_t^2}}{\frac{\beta}{c_t}}\right) = \frac{1}{c_t}$$

which does not depend on the affine-transformation parameters α and β.

Definition 227 (Constant Absolute Risk Aversion)
An agent has **constant absolute risk aversion** $\sigma > 0$ if their utility is of the form

$$u(x) = \frac{1 - e^{-\sigma x}}{\sigma} \tag{109}$$

Remark 228. *Under CARA, the risk premium P does not depend on the agent's initial wealth x_0.*

Proposition 229 (Wealth Effects under CARA)
In the standard portfolio problem with additive independent risks, CARA preferences imply that the optimal dollar exposure to risk is independent of initial wealth.

We can also define the notion of relative risk aversion which scales absolute risk aversion by current wealth.

Definition 230 (Coefficient of Relative Risk Aversion)
If an agent has a vNM utility function $u(x)$ defined on positive income $x > 0$, then their **coefficient of relative risk aversion** is:

$$r_R(x) = -\frac{x u''(x)}{u'(x)} = x r_A(x). \tag{110}$$

Proposition 231 (Relationship between ARA and RRA)
If r_R decreases, then so does r_A (not the converse!).
If r_R decreases, the agent invests a greater fraction of wealth in risky asset.

Definition 232 (Constant Relative Risk Aversion)
The agent has **constant relative risk aversion** $\sigma \geq 0$ if:

$$u(x) = \frac{x^{1-\sigma} - 1}{1 - \sigma} \tag{111}$$

for $\sigma \neq 1$.
As $\sigma \to 1$, we obtain $u(x) = \log(x)$ up to an additive constant.

We now give an interpretation of increasing or decreasing absolute and relative risk aversion in the context of forming a portfolio with one risky asset and one risk-free asset. If the person experiences an increase in wealth, they will choose to increase (or keep unchanged, or decrease) the number of dollars of the risky asset held in the portfolio if absolute risk aversion is decreasing (or constant, or increasing). Thus economists avoid using utility functions such as the quadratic, which exhibit increasing absolute risk aversion, because they have an unrealistic behavioral implication. Similarly, if the person experiences an increase in wealth, they will choose to increase (or keep unchanged, or decrease) the fraction of the portfolio held in the risky asset if relative risk aversion is decreasing (or constant, or increasing).

Definition 233 (Subjective Probabilities)
Suppose we have two complementary events 1 and 2. The agent's monetary payoff in event i is denoted x_i, and the agent has preferences over pairs (x_1, x_2). An agent is a **subjective expected-utility maximiser** if we can find a utility function $u(x)$ and a parameter $0 \leq p \leq 1$ such that $\succsim$ is represented by:

$$pu(x_1) + (1-p)u(x_2)$$

where p is the subjective probability that event 1 occurs.

6 Production Theory

6.1 Production Sets

We denote L commodities and index it by $\ell = 1, ..., L$.

Definition 234 (Production Vector)
A **production vector** (production plan) is a vector $y = (y_1, ..., y_L) \in \mathbb{R}^L$. $y_\ell > 0$ indicates an **output** and $y_\ell < 0$ indicates an **input**.

Definition 235 (Production Set)
The **production set** $Y \subset \mathbb{R}^L$ denotes the set of all feasible production plans.

It is hard to work with sets so we can use a transformation function instead.

Definition 236 (Transformation Function)
$F(\cdot)$ is a transformation for Y if it satisfies the property that $y \in Y$ if and only if $F(y) \leq 0$.

The boundary of Y is represented by the set of plans where $F(y) = 0$. This is called the **transformation frontier**.

Under free disposal and the usual monotonicity conditions, a plan with $F(y) < 0$ is inefficient, while efficient plans lie on the transformation frontier. We see an example in Fig. 24.

Definition 237 (Marginal Rate of Transformation)
For a production plan on the boundary, the marginal rate of transformation between two goods ℓ and k is defined by:

$$MRT = \frac{\partial F(y)/\partial y_\ell}{\partial F(y)/\partial y_k}. \tag{112}$$

The marginal rate of transformation tells us how much more (net) of a good k can be supplied if one unit less of good ℓ is supplied.

We now state commonly assumed properties of production sets.

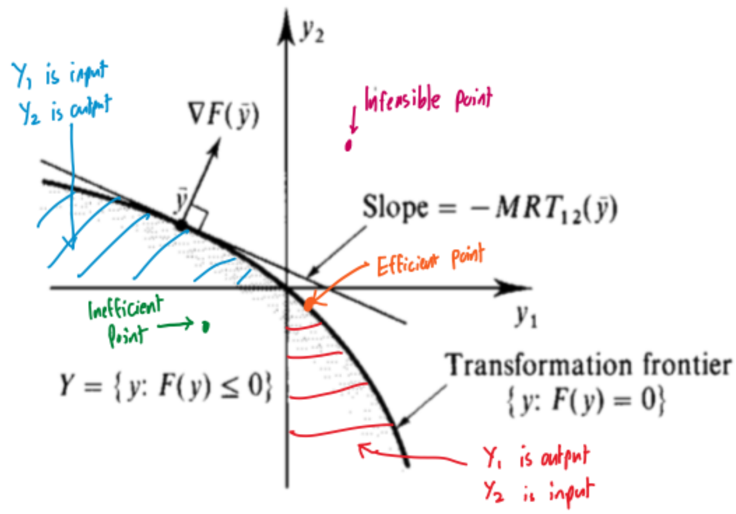

Figure 24: We see that any points below the transformation frontier to be inefficient. Any points above the frontier is infeasible. Any points that lie on it are efficient.

Definition 238 (Production Set is Closed)
The production set includes its boundary.

Remark 239. *This is a technical assumption so that when we carry out maximisation problems, we are guaranteed the existence of a solution.*

Definition 240 (No Free Lunch)
If $y \in Y$ and $y \geq 0$, then $y = 0$.

The no-free-lunch condition says that if y is a feasible production plan and **all** net-output components are weakly positive, then we **must** be at the origin. We should not be able to produce something out of nothing. An illustration can be seen in Fig. 25.

Definition 241 (Inactive)
The firm is allowed to produce nothing $0 \in Y$.

Inactivity gives the firm the option not to operate. Whether this option yields zero profit or a loss in the short run depends on whether fixed costs are avoidable or already sunk.

Definition 242 (Free Disposal)
If $y \in Y$ and $y' \leq y$ then $y' \in Y$.

Free disposal tells us that, from any feasible production plan y, we can always use weakly more inputs or dispose of some output and remain feasible. Thus the production set is downward closed. An example is shown in Fig. 26. With strictly positive prices, this property helps ensure that a profit-maximising active plan lies on the relevant frontier.

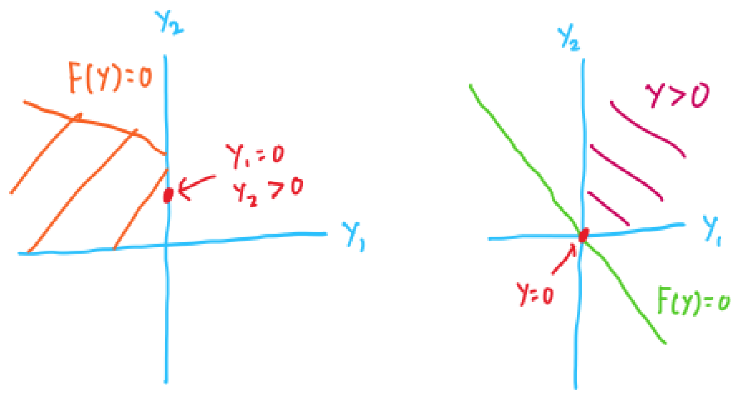

Figure 25: The figure on the left shows transformation frontiers ruled out by the no-free-lunch condition: no input is used, $y_1 = 0$, but positive output $y_2 > 0$ is produced. The figure on the right shows that every nonzero vector $y \geq 0$ must be ruled out.

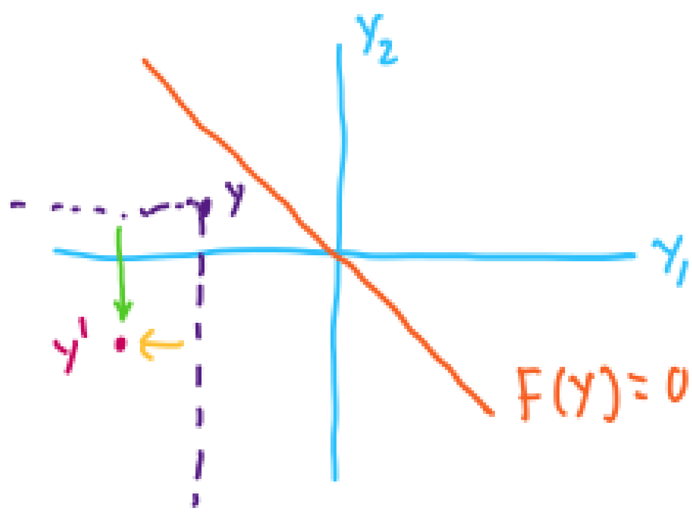

Figure 26: We are currently producing at an output of y. We can choose to use more input of y_1 and produce less output y_2.

Definition 243 (Additivity)
If $y, y' \in Y$ then

$$y + y' \in Y.$$

Definition 244 (The Production Set is Convex.)
Y is a convex set.

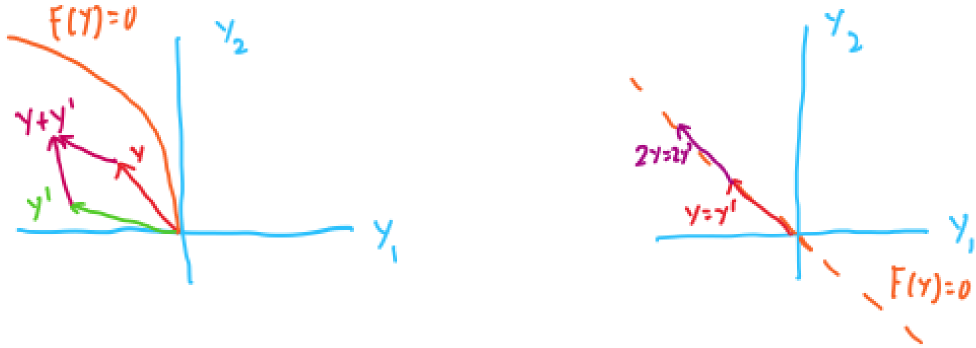

Figure 27: Additivity is related to constant returns to scale whereby if $y \in Y$ and $y' \in Y$ then $y + y' \in Y$. That is, we can merge production plans.

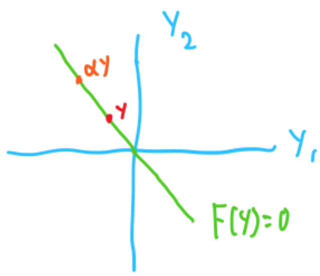

Figure 28: For every $\alpha \geq 0$, we have $\alpha y \in Y$. This means that the production set is a cone.

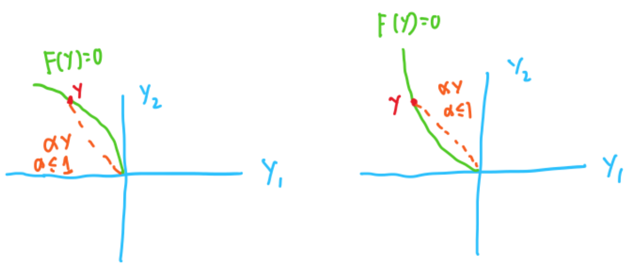

Figure 29: In the figure on the left, we have a transformation frontier that satisfies decreasing returns to scale whereby $y \in Y$ implies that $\alpha y \in Y$ for $\alpha \leq 1$. On the right, we see a transformation frontier that violates the decreasing returns to scale.

Definition 245 (Constant Returns to Scale)

Y has **constant returns to scale** (CRS) if whenever $y \in Y$, then

$$\alpha y \in Y$$

for any scale factor $0 \leq \alpha < \infty$.

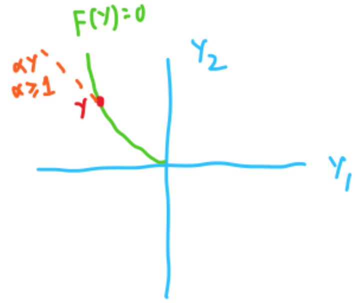

Figure 30: This transformation function satisfies increasing returns to scale where if $y \in Y$ then $\alpha y \in Y$ for $\alpha \geq 1$.

Definition 246 (Decreasing Returns to Scale)
Y has **decreasing returns to scale** if whenever $y \in Y$, then

$$\alpha y \in Y$$

for any scale factor $\alpha \in [0, 1]$.

Definition 247 (Increasing Returns to Scale)
Y has **increasing returns to scale** if whenever $y \in Y$, then

$$\alpha y \in Y$$

for any scale factor $1 \leq \alpha < \infty$.

Examples of constant, decreasing, and increasing returns to scale production sets can be found in Fig. 28, Fig. 29, and Fig. 30.

In many examples, we can distinguish between goods that are outputs $q = (q_1, ..., q_M) \geq 0$ and inputs $z = (z_1, ..., z_{L-M}) \geq 0$ alongside corresponding prices $p = (p_1, ..., p_M) \geq 0$ and $w = (w_1, ..., w_{L-M}) \geq 0$. An important case is when we have a single output whereby we write the number of units of the output q given inputs $z = (z_1, ..., z_{L-1}) \geq 0$ as $q = f(z)$. This function $f(\cdot)$ is known as a **production function** and is used for **single output**.

Definition 248 (Marginal Rate of Technical Substitution)
The **marginal rate of technical substitution** of input ℓ for input k is defined by:

$$MRTS = \frac{\partial f(z)/\partial z_\ell}{\partial f(z)/\partial z_k}. \tag{113}$$

Remark 249. *The marginal rate of technical substitution is how much more input k is needed to keep output constant when one unit less of input ℓ is used.*

We mention an important production function.

Definition 250 (CES Production Function)
A CES production function for N inputs is:

$$\left[\sum_{i=1}^{N} \alpha_i x_i^{\theta}\right]^{\frac{\nu}{\theta}}, \tag{114}$$

where $\alpha_i > 0$, $\sum_{i=1}^{N} \alpha_i = 1$, $\theta \leq 1$ with $\theta \neq 0$, and $\nu > 0$. The elasticity of substitution between inputs is $\frac{1}{1-\theta}$, and ν is the degree of homogeneity of the production function. The Cobb–Douglas case is obtained as $\theta \to 0$.

Remark 251. *The elasticity of substitution between two inputs measures the percentage change in their input ratio relative to the percentage change in their marginal-rate-of-technical-substitution (or relative factor price at an interior cost-minimising choice). It captures the curvature of the isoquant and hence the degree of substitutability between inputs.*

The CES production function is a general case of many other well known production functions.

Proposition 252 (CES Production Cases)
The different cases of the CES production function depend on the value of θ.

1. $\theta \to 0$: Cobb-Douglas.
2. $\theta \to 1$: Perfect Substitutes.
3. $\theta \to -\infty$: Leontief.

We can immediately see then for a Cobb-Douglas production function, the elasticity of substitution is 1.

6.2 Cost Minimisation

We now focus on cost minimisation. The firm aims to produce a minimum output target by choosing inputs in the least-cost way. As before, let $q \in \mathbb{R}_+$ denote a **single output**, $z = (z_1, ..., z_{L-1}) \geq 0$ denote **inputs**, and $w = (w_1, ..., w_{L-1}) \in \mathbb{R}^{L-1}_{++}$ denote input prices. The production function is $q = f(z)$.

Definition 253 (Cost Efficient Production and Conditional Factor Demand)
The cost function is given by:

$$c(w, q) = \min_{z}\{w \cdot z : f(z) \geq q\} \tag{115}$$

which is the least-cost way to produce q units of output.

The solution to this problem, ie., the least-cost way to supply q units of output at input prices w is denoted $z(w, q)$ and is called the **conditional factor demand function**.

The cost function $c(w, q)$ is conceptually identical to the expenditure function $e(p, u)$ whilst the conditional factor demand function $z(w, q)$ plays the role of Hicksian demands $h(p, u)$.

Proposition 254 (Properties of Cost Function and Conditional Factor Demands)
The following properties hold:

1. The cost function $c(w, q)$ is homogeneous of degree one in w and concave in w.
2. The conditional factor demand function $z(w, q)$ is homogeneous of degree zero in w.

We can write the Lagrangian for our cost minimisation problem as:

$$\mathcal{L} = w \cdot z + \lambda[q - f(z)]. \tag{116}$$

Proposition 255 (Shephard's Lemma)
Shephard's lemma states that you can derive the conditional factor demand function from differentiating the cost function with respect to input prices:

$$\frac{\partial c(w, q)}{\partial w_\ell} \equiv z_\ell(w, q). \tag{117}$$

Proof. Apply the envelope theorem to our Lagrangian to get the result:

$$\frac{\partial \mathcal{L}}{\partial w_\ell} = \frac{\partial}{\partial w_\ell}\big[w \cdot z + \lambda[q - f(z)]\big] = z_\ell(w, q).$$

□

Remark 256. *Shephard's lemma says that we do not need to consider the indirect effect of input prices on the demand for inputs.*

Proposition 257 (Matrix of Derivatives of Factor Demand is Symmetric)
The matrix of derivatives of $z(w, q)$ with respect to w is symmetric and negative semi-definite.

Proof. It is symmetric due to Young's theorem. It is negative semi-definite as $c(w, q)$ is concave in w. □

Definition 258 (Marginal Cost)
The firm's **marginal cost** of supplying an additional unit of output is:

$$\frac{\partial c(w, q)}{\partial q}. \tag{118}$$

Definition 259 (Inferior Input)
An input ℓ is called inferior if $z_\ell(w, q)$ decreases with q.

Remark 260. *This says that as you scale up production q, you use less of input ℓ.*

We now distinguish between the **short run** and the **long run**. In the short run, some factors are fixed, while in the long run all factors are variable. Consequently, long-run cost is weakly below short-run cost for a given output. More formally, write $C(q)$ for the long-run cost function and $C(q|\bar{z}_2)$ for the firm's cost when the second input is fixed at $\bar{z}_2$. With two inputs, the short-run cost is:

$$C(q|\bar{z}_2) = w_1 z_1 + w_2 \bar{z}_2$$

whereby z_1 is chosen so that $q = f(z_1, \bar{z}_2)$. The short-run and long-run costs coincide when the fixed quantity $\bar{z}_2$ is also the long-run cost-minimising quantity. Therefore, long-run cost $C(q)$ is the lower envelope of the family of short-run cost curves as we vary the fixed input $\bar{z}_2$. This is known as the **Wong–Viner** envelope result.

6.3 Profit Maximisation

Let p be the price vector of both inputs and outputs. The total profit with production plan y is:

$$p \cdot y$$

as $y_i > 0$ is where the firm generates revenue and $y_i < 0$ is where the firm generates cost. This leads to the firm's problem being:

$$\max_y \left\{ p \cdot y \text{ subject to } y \in Y \right\}$$

whereby $y(p)$ is the most profitable plan. $y(p)$ is known as the **supply function**, and we denote the resulting maximised profit at price vector p by $\pi(p)$. It is important to note that the supply function $y(p)$ includes both inputs and outputs.

When the profit function is twice differentiable, Young's theorem implies that the matrix of derivatives of the supply function is symmetric:

$$\frac{\partial y_\ell(p)}{\partial p_k} \equiv \frac{\partial y_k(p)}{\partial p_\ell}.$$

The matrix of derivatives of $y(p)$ is positive semidefinite. Therefore, the own-price effect $\partial y_\ell(p)/\partial p_\ell$ is weakly nonnegative, so a price rise induces the firm to weakly increase the corresponding net output (supply more output or demand less input). This gives us an important result.

Theorem 261 (Law of Supply)
The law of supply is given by:

$$(p' - p) \cdot (y(p') - y(p)) \geq 0. \tag{119}$$

That is, prices and quantities move in the same direction.

Proof. We have the Weak Axiom of Profit Maximisation:

$$p \cdot y(p) \geq p \cdot y'$$

for all $y, y' \in Y$ whereby the profits we make with the optimal production plan $y(p)$ MUST be better than any other production plan y'. Then, this implies that:

$$p \cdot y(p) \geq p \cdot y(p')$$

$$p' \cdot y(p') \geq p' \cdot y(p)$$

as $y(p')$ is only the optimal plan at prices p'. We can then add the two inequalities together to get:

$$\big(p \cdot y(p) - p \cdot y(p')\big) + \big(p' \cdot y(p') - p' \cdot y(p)\big) \geq 0$$

We can then re-arrange to get our desired result. ☐

Supply function is nicely determined when we have **decreasing returns to scale**.

The profit-maximising plan can be shown to satisfy the first-order condition

$$\frac{\partial F(y)}{\partial y_\ell} = \lambda p_\ell \quad \text{for } 1 \leq \ell \leq L \tag{120}$$

for some positive constant λ. Here, the marginal transformation rates are proportional to prices. Intuitively, we have an isoprofit line $\{y : p \cdot y = \pi(p)\}$ and try to reach the highest isoprofit line subject to the feasible production set. We can see this graphically in Fig. 31.

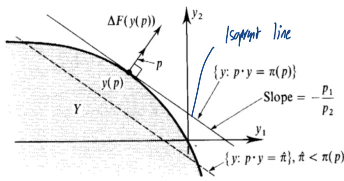

Figure 31: Profit maximisation problem represented graphically.

Proposition 262 (Properties of Profit Function)
The profit function $\pi(p)$ is homogeneous of degree one.

When the profit-maximising plan is unique, the supply function $y(p)$ is homogeneous of degree zero.

Proposition 263 (Hotelling's Lemma)
The derivative of the profit function with respect to prices gives you the supply function:

$$\frac{\partial \pi(p)}{\partial p_\ell} = y_\ell(p) \quad \text{for } 1 \leq \ell \leq L. \tag{121}$$

Remark 264. *This is the top half of Roy's identity.*

We now go back to the single-output case and have the cost vector w as a constant. We represent

$$C(q) = c(w, q)$$

as the cost function for producing q units of outputs. The **profits** of the firm is given by:

$$\pi(q) = p(q)q - C(q). \tag{122}$$

Definition 265 (Firm Profit Maximisation Problem)
The **profit maximisation problem** is given by:

$$\max_{q \geq 0} p(q)q - C(q) \tag{123}$$

where $p(q)$ allows for the possibility that if the firm changes its output decision, this can have an effect on the market price.

We can take the first-order condition with respect to q. By the product rule, we have:

$$\frac{d\pi(q)}{dq} = \frac{d}{dq}\left[p(q)q - C(q)\right] = p'(q)q + p(q) - C'(q) = 0$$

We therefore have the general result.

Proposition 266 (Marginal Revenue Equals Marginal Cost)

The first order condition for profit maximisation is:

$$\underbrace{p + p'(q)q}_{\text{Marginal Revenue}} = \underbrace{C'(q)}_{\text{Marginal Cost}} \tag{124}$$

Remark 267. *This condition that marginal revenue equals marginal cost holds for* **any** *market structure.*

To interpret the marginal revenue, we can see that there are two effects if we change the quantity q produced. The direct effect is that we will gain a revenue of p by producing an additional unit of output. However, if we produce one more unit of output, this can affect the market price for the good $p'(q)$ and can be seen as the indirect effect.

Example 268 (Perfect Competition)

Under perfect competition, this corresponds to the setting where $p(q) = p$ and $p'(q) = 0$. Here, the producer producing extra units has no impact on the market price. This means that we have:

$$p = C'(q)$$

or in words: **price equals marginal cost**. If C is a convex function, so that C' is increasing, then this first-order condition is also sufficient for profit maximisation. Then, the optimal supply function given p, which we denote by $q(p)$ satisfies $C'(q(p)) = p$.

Proposition 269 (First- and Second-Order Conditions for Profit Maximisation)

The (necessary) first-order condition is that the slopes of the revenue and cost functions are equal.

The second-order condition in the general case is:

$$\pi''(q) = 2p'(q) + qp''(q) - C''(q) \leq 0.$$

Under perfect competition, $p'(q) = p''(q) = 0$, so this reduces to $\pi''(q) = -C''(q) \leq 0$.

Therefore, to find the profit maximisation, we first need to find a level of output whereby marginal revenue equals marginal cost. Then, we must make sure that marginal costs are not decreasing in order to ensure we are not minimising profits.

Fig. 32 shows what we mean graphically. When we plot the revenue and cost functions, we look for a point where the slopes of $C(q)$ and $R(q)$ are equal. We can see that q^* in our picture is such a point. Furthermore, the revenue curve is above the cost curve, so profit is positive. However, point y is another point where the two slopes are equal but the cost curve is above the revenue curve. This is not a profit maximum. Hence, we need to check both the first- and second-order conditions, as well as compare relevant boundary points.

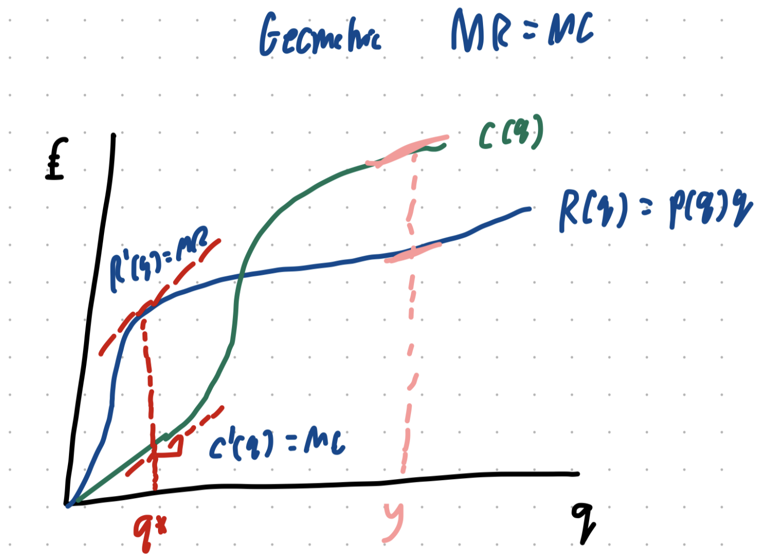

Figure 32: Profit maximisation problem represented graphically. $R(q)$ denotes the revenue function and $C(q)$ denotes the cost function.

Example 270 (Perfect Competition)
Under perfect competition, the revenue function is a straight line with slope p. We need to find the point where the slopes of the revenue and cost functions are equal and marginal cost is not decreasing. We can see this in Fig. 33.

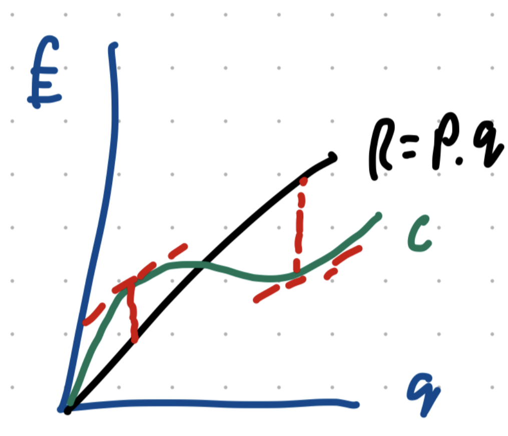

Figure 33: Profit maximisation problem represented graphically under perfect competition. $R(q)$ denotes the revenue function and $C(q)$ denotes the cost function.

Example 271 (Constant Returns to Scale)
Suppose that we have constant returns to scale and perfect competition. We can see in Fig. 34 that the two linear functions are either parallel everywhere or not parallel at all. Therefore, the firm either produces nothing, is indifferent over scale, or has unbounded desired output. Hence, profit maximisation under perfect competition

with a constant-returns Cobb–Douglas production function can fail to yield a unique finite scale.

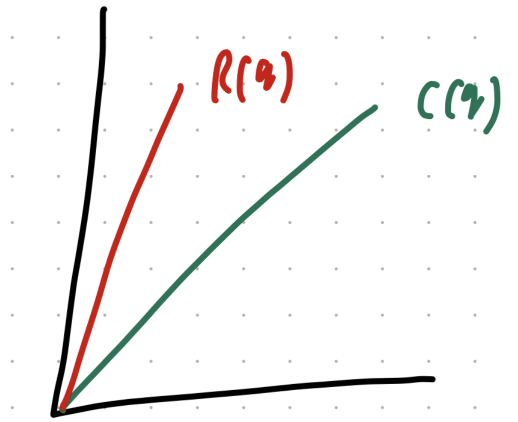

Figure 34: Profit maximisation problem represented graphically with constant returns to scale under perfect competition. $R(q)$ denotes the revenue function and $C(q)$ denotes the cost function.

We now look at what happens under the presence of **fixed costs**. Suppose that we had the cost function

$$C(q) = C_v(q) + F$$

where $C_v(q)$ is variable cost and F is fixed cost. This is a **short-run cost function**. We can then construct the profit function:

$$\pi(q) = pq - C_v(q) - F$$

If the firm produces nothing, it still needs to cover its fixed costs so then:

$$\pi(0) = -F.$$

Hence, it is not worth producing anything if losses at $q = 0$ are less than losses at $q \geq 0$:

$$\pi(0) > \pi(q)$$

Plugging in our profit functions, this is:

$$-F > pq - C_v(q) - F$$

We can re-arrange to get the rule:

$$\underbrace{p}_{\text{price}} < \underbrace{\frac{C_v(q)}{q}}_{\text{average variable cost}}.$$

Therefore, it is not worth producing a positive quantity if **price is below average variable cost**. It is worth producing if revenue covers variable costs and makes some contribution to fixed costs. Otherwise, the firm is better off shutting down. Sunk fixed costs can therefore lead firms to operate despite making negative short-run profits.

We can therefore summarise the firm's decision in the **short run**.

Definition 272 (Firm's Shutdown Condition in Short Run)

The short run supply for a competitive firm is:

$$\text{produce} = \begin{cases} q \text{ such that } p = C'(q) \text{ and } C''(q) \geq 0 & \text{if } p \geq \min_{\tilde{q}>0} \frac{C_v(\tilde{q})}{\tilde{q}} \\ q = 0 \text{ otherwise} \end{cases}$$

That is, the firm produces where price equals marginal cost, provided the second-order condition holds and price is at least the minimum average variable cost.

In the **long run**, there are no fixed costs so average variable cost is equal to average cost and therefore the shutdown condition depends on average cost.

If price p is less than average cost $C(q)/q$ at the firm's optimal output, then the firm makes negative profits $\pi(q) = pq - C(q) < 0$ and should exit the market in the long run.

If price p is greater than average cost $C(q)/q$, then the firm makes positive profits $\pi(q) = pq - C(q) > 0$, which induces entry.

The **long run equilibrium** is a state in which there is neither net exit or net entry by firms. This is when profits are equal to 0:

$$pq - C(q) = 0.$$

Re-arranging this, we can see that this occurs when:

$$\underbrace{p}_{\text{price}} = \underbrace{\frac{C(q)}{q}}_{\text{average cost}}.$$

At an interior competitive optimum, price also equals marginal cost. Thus long-run zero-profit equilibrium satisfies $p = MC(q) = AC(q)$, so the firm operates at a minimum of average cost.

7 General Equilibrium Theory

So far, we have the basic ingredients of an economy:

1. Households: Demand goods, supply capital, and own firms.
2. Firms: Demand capital, produce goods, and earn profits which are distributed to households.
3. Markets: Where goods and inputs are traded.

In this lecture, we analyse three simple examples of general equilibrium models. Each of these examples will illustrate particular components of how these ingredients interact.

7.1 Exchange Economies

We restrict our focus to exchange and consumption. We analyse a **pure exchange economy** in which no production is possible and the commodities that are ultimately consumed are those that individuals possess as **endowments**.

Assume that there are two consumers denoted by $i = 1, 2$ and two commodities denoted by $\ell = 1, 2$. We assume that these two consumers are price takers.

Definition 273 (Consumption and Endowment Vector)

Consumer i's **consumption vector** $\boldsymbol{x}_i \in \mathbb{R}^2_+$ is denoted by:

$$\boldsymbol{x}_i = (x_{1i}, x_{2i}). \tag{125}$$

Consumer i's **endowment vector** $\boldsymbol{\omega}_i \in \mathbb{R}^2_+$ is denoted by:

$$\boldsymbol{\omega}_i = (\omega_{1i}, \omega_{2i}). \tag{126}$$

Consumer i is equipped with a preference relation $\succsim_i$ over the consumption vectors.

The total endowment of good ℓ in this economy is the sum of the endowments of the two agents:

$$\overline{\omega}_\ell = \omega_{\ell 1} + \omega_{\ell 2}.$$

Definition 274 (Feasible Allocation)

An **allocation** $\boldsymbol{x} \in \mathbb{R}^4_+$ in this economy is an assignment of a non-negative consumption vector to each consumer:

$$\boldsymbol{x} = (\boldsymbol{x}_1, \boldsymbol{x}_2) = \big((x_{11}, x_{21}), (x_{12}, x_{22})\big). \tag{127}$$

An allocation $\boldsymbol{x}$ is **feasible** for the economy if:

$$x_{\ell 1} + x_{\ell 2} \leq \overline{\omega}_\ell \quad \text{for } \ell = 1, 2, \tag{128}$$

whereby total consumption of each commodity is no more than the economy's aggregate endowment of that commodity.

Allocations for which the feasibility constraints in Eq. (128) bind are called **nonwasteful**. Nonwasteful allocations can be represented by an **Edgeworth box**, as in Fig. 35. Consumer 1's quantities are measured starting from the south-west corner whilst consumer 2's quantities are measured from the north-east corner. The length of the x-axis corresponds to the total endowment of good 1 $\overline{\omega}_1$ and the length of the y-axis corresponds to the total endowment of good 2 $\overline{\omega}_2$. In Fig. 35, the endowment vector $\boldsymbol{\omega} = ((\omega_{11}, \omega_{21}), (\omega_{12}, \omega_{22}))$ is plotted. Likewise, the consumption vector $\boldsymbol{x} = ((x_{11}, x_{21}), (x_{12}, x_{22}))$ is also plotted.

We can now define what a consumer's wealth is based on prices $\boldsymbol{p}$.

Definition 275 (Consumer's Wealth)

For any prices $\boldsymbol{p} = (p_1, p_2)$, the consumer i's wealth is given by the market value of their endowments:

$$\boldsymbol{p} \cdot \boldsymbol{\omega}_i = p_1\omega_{1i} + p_2\omega_{2i}. \tag{129}$$

This implies that consumer i's budget set is a function of **only** prices. Under strongly monotone preferences, the consumer's demanded bundle exhausts wealth and therefore **lies on the budget line**.

Definition 276 (Budget Set)

The budget set for consumer i given prices $\boldsymbol{p}$ is:

$$B_i(\boldsymbol{p}) = \{\boldsymbol{x}_i \in \mathbb{R}^2_+ : \boldsymbol{p} \cdot \boldsymbol{x}_i \leq \boldsymbol{p} \cdot \boldsymbol{\omega}_i\}. \tag{130}$$

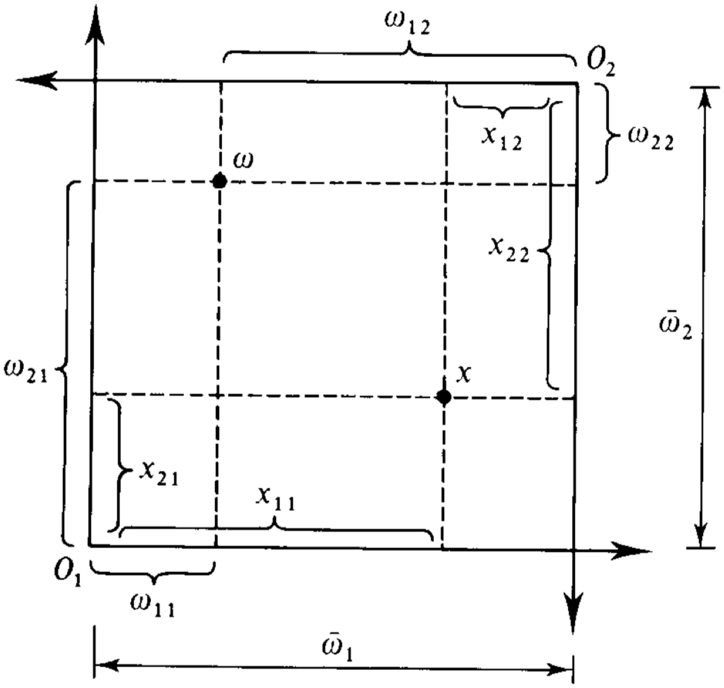

Figure 35: Edgeworth box.

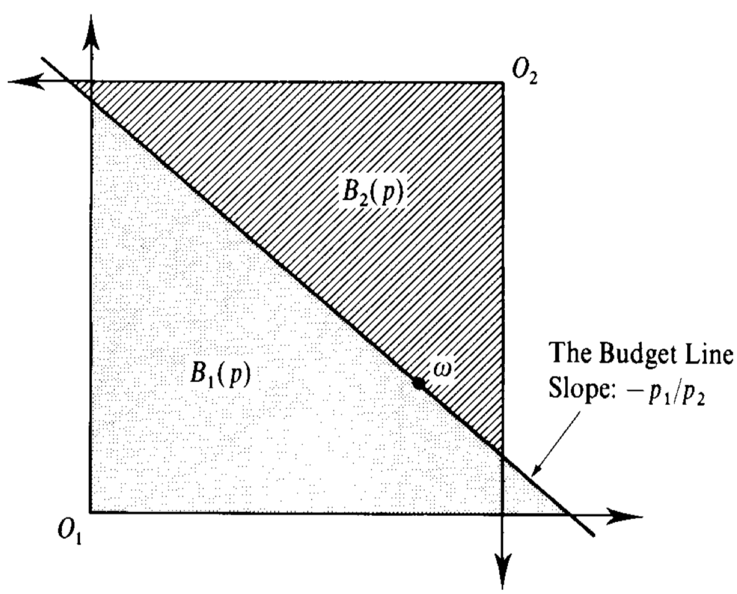

Figure 36: Budget set passing through initial endowment point $\boldsymbol{\omega}$ in Edgeworth box. Consumer 1's budget set consists of all the non-negative vectors below and to the left of the budget line. Consumer 2's budget set consists of all vectors above and to the right of the budget line. Only allocations **on** the budget line are affordable to both consumers simultaneously at prices (p_1, p_2).

We can draw budget sets of the two consumers in the Edgeworth box by drawing the **budget line** through the endowment point $\boldsymbol{\omega}$ with the slope of the negative of the price ratios $-(\frac{p_1}{p_2})$. This can be seen in Fig. 36

We can also depict each consumer i's preferences $\succsim_i$ in the Edgeworth box. We assume that these preference relations are continuous, strictly convex, and strongly monotone, so that they admit well-behaved utility representations. We can then depict their indifference curves as in Fig. 37.

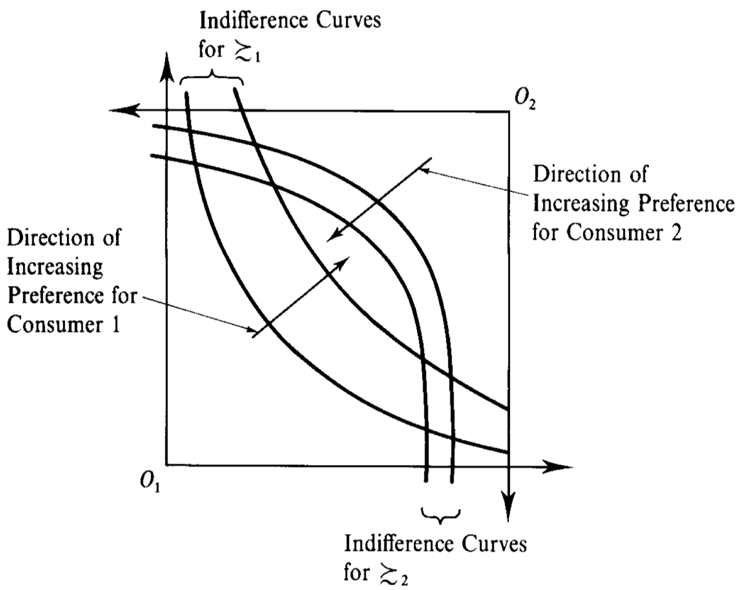

Figure 37: Indifference curves in Edgeworth box.

Using standard consumer choice theory, given a budget set and utility function, the agents can pick their most preferred point in their budget set. This corresponds to the Marshallian demand function $x_i(\boldsymbol{p}, \boldsymbol{p} \cdot \boldsymbol{\omega}_i)$ we have seen before, where wealth is now $\boldsymbol{p} \cdot \boldsymbol{\omega}_i$. An example is shown in Fig. 38.

Now, recall that the budget set passes through the endowment point $\boldsymbol{\omega}$ and has slope equal to the negative price ratio $-p_1/p_2$. As we vary the price vector $\boldsymbol{p}$, the budget line pivots **around** the endowment point $\boldsymbol{\omega}$. Associated with each pivoted line, agents have a new Marshallian demand for goods. Tracing out these Marshallian demands as prices vary gives the **offer curve**. We plot the offer curve for consumer 1 in Fig. 39. Note that the offer curve passes through the endowment point. This is because, at any prices $\boldsymbol{p}$, the consumer can **always** consume their endowment $\boldsymbol{\omega}_1 = (\omega_{11}, \omega_{21})$, so it is always affordable. Every point on the offer curve must therefore be at least as good as the endowment point.

We have a useful property characterising the offer curve.

Proposition 277 (Offer Curve Tangency)
The offer curve is tangent to the consumer's indifference curve at the endowment point ω.

We now want to characterise what an equilibrium is. We require that markets should clear at going market prices.

Definition 278 (Walrasian Equilibrium for Edgeworth Box Economy)
A **Walrasian equilibrium** for an Edgeworth box economy is a price vector $\boldsymbol{p}^*$ and a feasible allocation vector $\boldsymbol{x}^* = (\boldsymbol{x}_1^*, \boldsymbol{x}_2^*)$ such that markets clear, $\boldsymbol{x}_1^* + \boldsymbol{x}_2^* = \boldsymbol{\omega}_1 + \boldsymbol{\omega}_2$, and, for $i = 1, 2$,

$$\boldsymbol{x}_i^* \succsim_i \boldsymbol{x}_i' \tag{131}$$

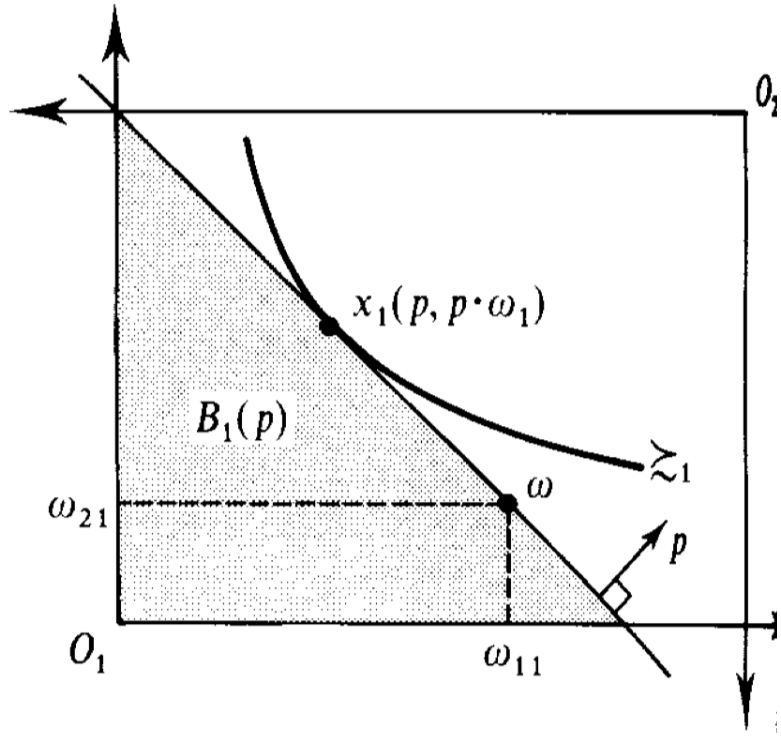

Figure 38: Marshallian demand is the point of tangency between the utility function and budget set in the Edgeworth box.

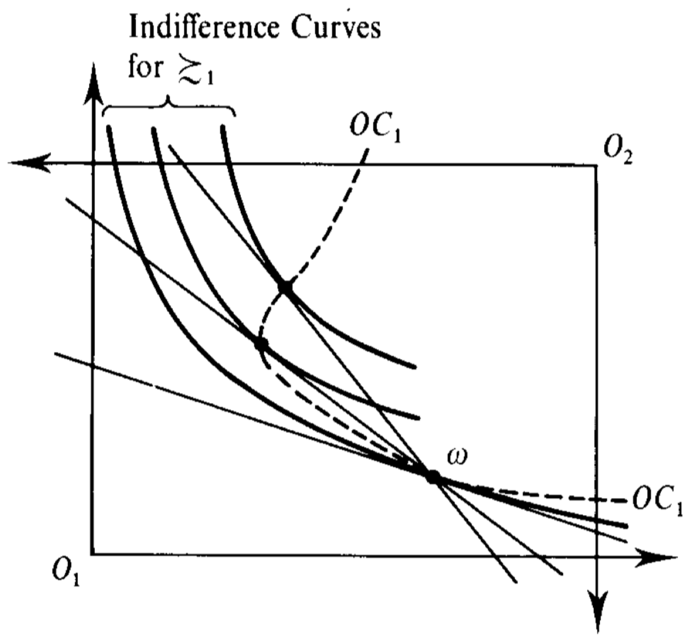

Figure 39: Offer curve OC_1 traced out for consumer 1. We see that the budget lines pivot around the endowment vector $\boldsymbol{\omega}$ as we vary the price vector $\boldsymbol{p}$.

for all $\boldsymbol{x}_i' \in B_i(\boldsymbol{p}^*)$.

A Walrasian equilibrium is given in Fig. 40. We see that the two indifference curve are tangential to each other (this only holds at **interior points**!). Furthermore, in Fig. 41, we see that the two offer curves intersect. This actually results in an important result.

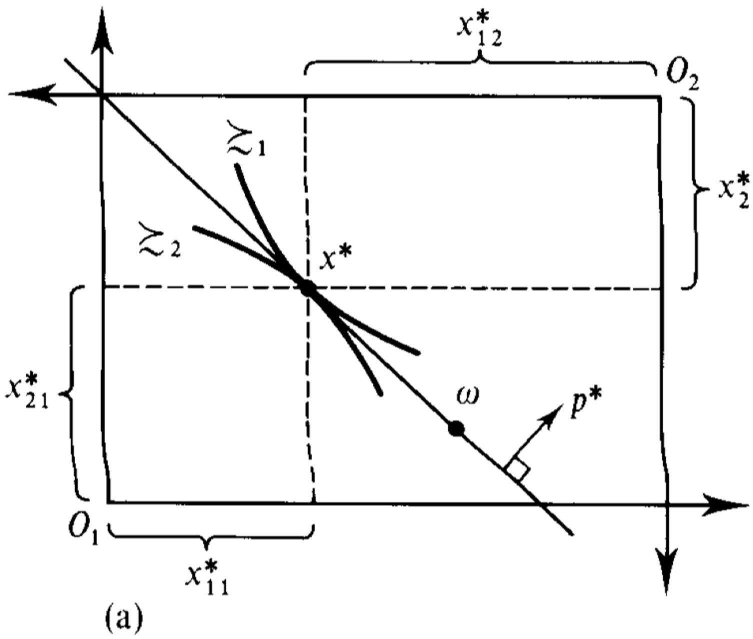

Figure 40: Walrasian equilibrium. Here, net demand is matched by net supply.

Proposition 279 (Intersections of Offer Curves are Walrasian Equilibria)
An intersection of the offer curves at an allocation different from the endowment point $\boldsymbol{\omega}$ in an Edgeworth-box economy is a Walrasian equilibrium.

Proof. Suppose that $\boldsymbol{x}^* = (\boldsymbol{x}_1^*, \boldsymbol{x}_2^*)$ is an allocation where the two offer curves intersect under prices $\boldsymbol{p}$. By definition, this implies that $\boldsymbol{x}_i^*$ is the optimal consumer bundle for both consumer i under the budget line that goes through the two points $\boldsymbol{\omega}$ and $\boldsymbol{x}^*$ under the prices $\boldsymbol{p}$. $\square$

Walrasian equilibria need not exist without suitable assumptions. For example, if one consumer has all the endowments of a particular good, there may be no strictly positive price ratio that clears the market. Equilibria may also fail to exist when consumer preferences are non-convex.

We can also investigate welfare properties of Walrasian equilibria. We define what welfare is under the Edgeworth economy.

Definition 280 (Pareto Optimal)
An allocation $\boldsymbol{x}$ in the Edgeworth box is **Pareto optimal** if there is no other allocation $\boldsymbol{x}'$ in the Edgeworth box with $\boldsymbol{x}_i' \succsim_i \boldsymbol{x}_i$ for $i = 1, 2$ and $\boldsymbol{x}_i' \succ \boldsymbol{x}_i$ for some consumer i.

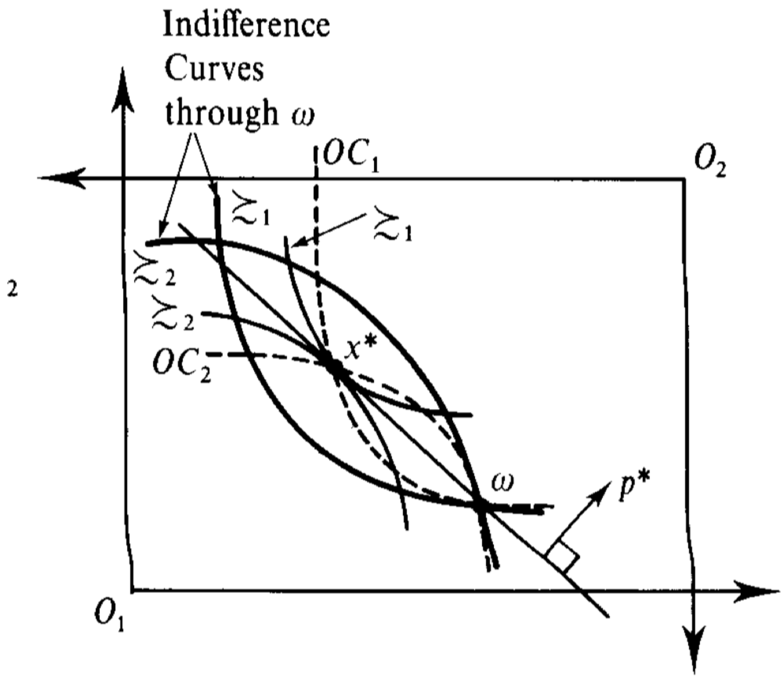

Figure 41: Intersection of offer curves result in a Walrasian equilibrium.

We show in Fig. 42 an example of a non Pareto optimal allocation and a Pareto optimal allocation. For the non Pareto optimal allocation, we see that the shaded region contains the intersection of points whereby

$$\{x_1' \in \mathbb{R}_+^2 : x_1' \succsim_1 x_1\} \cap \{x_2' \in \mathbb{R}_+^2 : x_2' \succsim_2 x_2\}$$

which are allocation points (x_1', x_2') that make both consumers weakly better off than at the current allocation $\boldsymbol{x}$. A Pareto improvement requires at least one consumer to be strictly better off. At an **interior Pareto-efficient point**, the indifference curves are tangential to one another. This does not have to hold at the boundary.

Remark 281. *At an interior allocation with differentiable preferences, tangency is a necessary condition for Pareto efficiency. Under the maintained convexity assumptions, the corresponding first-order condition is also sufficient; boundary allocations instead satisfy appropriate inequality conditions.*

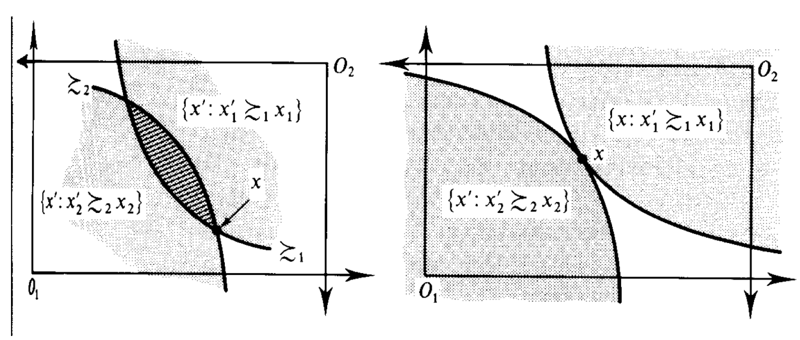

Figure 42: Image on the left depicts a non Pareto optimal allocation. The image on the right depicts a Pareto optimal allocation.

We now can collect all the set of Pareto optimal allocations.

Definition 282 (Pareto Set)
The set of all Pareto optimal allocations is known as the **Pareto set**.

In Fig. 43, we plot the Pareto set, also called the **contract curve**. The individually rational portion of the contract curve, on which both consumers do at least as well as at their initial endowment, is the **core** of this exchange economy.

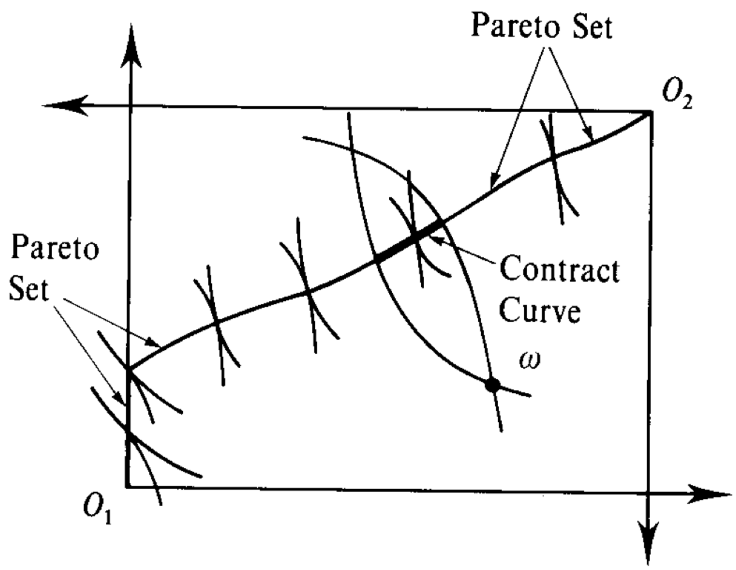

Figure 43: The contract curve and Pareto sets.

Definition 283 (Contract Curve)
The contract curve is the set of Pareto-efficient allocations in the Edgeworth box.

Remark 284. *All the points on this locus are Pareto efficient allocations, meaning that from any one of these points there is no reallocation that could make one of the people more satisfied with his or her allocation without making the other person less satisfied*

We now state a very important fact that is a special case of the **first fundamental theorem of welfare economics** that we will see later.

Proposition 285 (Walrasian Equilibria are Pareto Optimal)
Any Walrasian equilibrium allocation $\boldsymbol{x}^*$ necessarily belongs to the Pareto set.

Remark 286. *This also implies that any Walrasian equilibrium lies on the contract curve.*

We can state a special case of the **second fundamental theorem of welfare economics**.

Definition 287 (Second Fundamental Theorem of Welfare Economics)
An allocation $\boldsymbol{x}^*$ in the Edgeworth box is supportable as an equilibrium with transfers if there is a price system $\boldsymbol{p}^*$

and wealth transfers T_1 and T_2 satisfying $T_1 + T_2 = 0$ such that for each consumer i, we have that:

$$x_i^* \succsim_i x_i' \tag{132}$$

for all $x_i' \in \mathbb{R}^2_+$ such that

$$\boldsymbol{p}^* \cdot x_i' \leq \boldsymbol{p}^* \cdot \boldsymbol{\omega}_i + T_i.$$

Remark 288. *This says that under convexity assumptions, a planner can achieve any desired Pareto optimal allocation by appropriately redistributing wealth in a lump-sum fashion and then letting the market work. Note that the transfers sum to zero in the definition, so that the planner runs a balanced budget.*

We can in fact give another statement of the second welfare theorem. If the preferences of the two consumers in the Edgeworth box are continuous, convex, and strongly monotone, then **any Pareto-optimal allocation is supportable as an equilibrium with transfers**. We show this in Fig. 44. In the figure on the left, suppose that the desired outcome is the allocation $\boldsymbol{x}^*$ and the current outcome is the initial endowment $\boldsymbol{\omega}$. The tax authority can construct a wealth transfer between the two consumers that shifts the budget line, using a price vector $\boldsymbol{p}^*$ that clears both markets and supports $\boldsymbol{x}^*$. Alternatively, as shown on the right, endowments can be redistributed to $\boldsymbol{\omega}'$ or $\boldsymbol{\omega}''$ so that $\boldsymbol{p}^*$ and $\boldsymbol{x}^*$ form a Walrasian equilibrium.

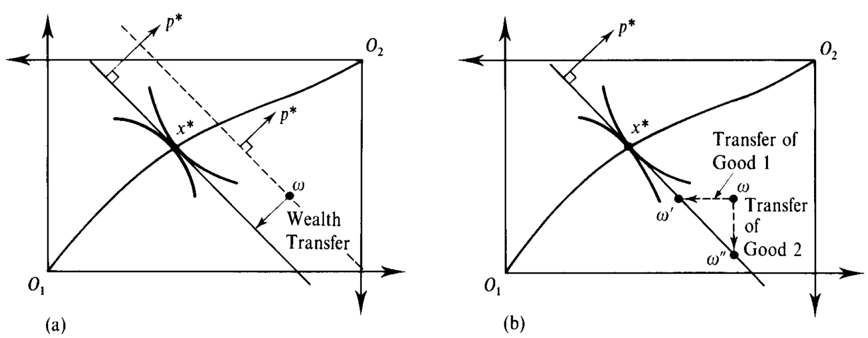

Figure 44: The second fundamental welfare theorem. On the left, wealth transfers are used. On the right, transfers of endowments are used.

The second welfare theorem may fail to hold when preferences are not convex.

7.2 Robinson Crusoe Economies

We now introduce production by studying an economy consisting of one firm and one consumer. Using this, we will explore how the production and consumption sides of the economy fit together. Suppose we have two price taking economic agents, a single household and a single firm. The two goods for the consumer to consume in the economy are leisure and the consumption good produced by the firm.

Suppose that the consumer has continuous, convex, and strongly monotone preferences $\succsim$ defined over leisure x_1 and the consumption good x_2. Suppose that they also have an endowment of $\overline{L}$ units of leisure (24 hours in a day) and **no endowment of the consumption good**.

Suppose that the firm uses labor to produce the consumption good according to an increasing and strictly concave production function $f(H)$ where H is the firm's labor input. Hence, to produce output, the firm must hire the consumer and purchase labour from them.

Definition 289 (Firm Profit Maximisation Problem)
The firm seeks to maximise profits given prices. Therefore, letting p be the price of output and w be the price of labour, the firm solves the problem:

$$\max_{H \geq 0} \quad pf(H) - wH. \tag{133}$$

The firm's output is given by $x(p, w) = f(H(p, w))$.

Firms are owned by consumers and hence profits earned by the firm $\pi(p, w)$ accrues to the consumer.

We now let $u(L, c)$ be the consumer's utility function representing the preferences $\succsim$ over leisure L and consumption of the good c. We can now state the consumer's problem.

Definition 290 (Consumer's Utility Maximisation Problem)
The consumer's utility maximisation problem is given by:

$$\begin{aligned} \max_{L,c} \quad & u(L, c) \\ \text{s.t.} \quad & pc \leq w(\overline{L} - L) + \pi(p, w) \end{aligned} \tag{134}$$

where $\overline{L}$ was the endowment of time.

Remark 291. *To make things easier, we could have endowed the agent with 1 unit of time so that their budget constraint becomes:*

$$pc + wL = w + \pi$$

where the left hand side is the expenditure of consumption and leisure whilst the right hand side is the income from the wage and firm's profits.

The budget constraint shows that the consumer can earn income from working and from the profits they receive from the firm. The consumer's optimal demands can be written as $(L(p, w), c(p, w))$.

Proposition 292 (Walrasian Equilibrium)
A Walrasian equilibrium in this economy involves a price vector (p^*, w^*) at which the consumption and labor markets clear:

$$c(p^*, w^*) = x(p^*, w^*)$$

$$H(p^*, w^*) = \overline{L} - L(p^*, w^*)$$

It can be shown that the consumer's budget line is exactly the isoprofit line associated with the solution to the firm's profit maximisation problem. A particular consumption-leisure combination can arise in a competitive equilibrium if and only if it maximises the consumer's utility subject to the economy's technological and endowment constraint.

7.3 Small Open Economies

We now analyse the production side of the economy in greater detail by discussing the allocation of resources among several firms. To do this, we look at a small open economy that takes the world prices of its outputs as fixed.

Assumption (Prices of Output Goods Are Fixed)
Prices of output goods are fixed.

Suppose that we have **J firms**. Each firm j produces a **consumer good** $q_j \in \mathbb{R}_+$ with L factors. Therefore, the factors for consumer good j is:

$$\boldsymbol{z}_j = (z_{1j}, ..., z_{Lj}) \geq 0.$$

Since there are J firms, this implies that there are J consumer goods in this economy. Firm j's production function $q_j = f_j(\boldsymbol{z}_j) \in \mathbb{R}_+$ is a concave, strictly increasing and differentiable production function $f_j(\boldsymbol{z}_j)$.

The economy has total endowments $\overline{\boldsymbol{z}} \in \mathbb{R}^L_+$ (the L factors of production) owned by the consumers:

$$\overline{\boldsymbol{z}} = (\overline{z}_1, ..., \overline{z}_L) >> 0.$$

These endowments can only be used as production inputs.

We want to focus on the factor markets of the economy. Therefore, we suppose that the prices $\boldsymbol{p} \in \mathbb{R}^J_+$ of the J produced consumption goods are fixed at:

$$\boldsymbol{p} = (p_1, ..., p_J).$$

An example we can think of is a small open economy whose trading decision in world markets for consumption goods have little effect on world prices for these goods. The output produced is sold in world markets whilst factors are immobile and must be used for production within the country.

The key question we want to know is what are the equilibrium factor prices $\boldsymbol{w} = (w_1, ..., w_L) \in \mathbb{R}^L_+$ and what are the allocations of the economy's factor endowments among the J firms. Therefore, given output prices as fixed $\boldsymbol{p} = (p_1, ..., p_J) \in \mathbb{R}^J_+$, and with input prices as $\boldsymbol{w} = (w_1, ..., w_L) \in \mathbb{R}^L$, the **profit maximising production plan** for firm j solves for:

$$\max_{\boldsymbol{z}_j \geq 0} p_j f_j(\boldsymbol{z}_j) - \boldsymbol{w} \cdot \boldsymbol{z}_j \tag{135}$$

where recall that $\boldsymbol{z}_j \in \mathbb{R}^L_+$ are the factor inputs, $\boldsymbol{w} \in \mathbb{R}^L_+$ are the costs of these factors, whilst $p_j, f_j(\boldsymbol{z}_j)$ are scalars representing the price of consumer good j p_j and the amount of good j produced $q_j = f_j(\boldsymbol{z}_j)$.

Firm j's set of **optimal input demands** given prices $(\boldsymbol{p}, \boldsymbol{w})$ is denoted by:

$$\boldsymbol{z}_j(\boldsymbol{p}, \boldsymbol{w}) \subseteq \mathbb{R}^L_+$$

whereby $\boldsymbol{z}$ is a (set of) vector for how much of each of the L factors are demanded. Note that it is a $\subset$ sign since we could have multiple demands as the solution (a correspondence!).

In this model, since the endowment that consumers receive cannot be consumed by consumers, the total factor supply is $\overline{\boldsymbol{z}} = (\overline{z}_1, ..., \overline{z}_L)$ as long as input prices $\boldsymbol{w} = (w_1, ..., w_L) >> 0$ are strictly positive.

Definition 293 (Factor Market Equilibrium)
The equilibrium for the factor markets of this economy given the fixed output prices $\boldsymbol{p} = (p_1, ..., p_J)$ consists of:

1. An input price vector $\boldsymbol{w}^* = (w_1^*, ..., w_L^*) >> 0$ and a factor allocation:

$$(\boldsymbol{z}_1^*, ..., \boldsymbol{z}_J^*) = \Big((z_{11}^*, ..., z_{L1}^*), ..., (z_{1J}^*, ..., z_{LJ}^*)\Big)$$

2. All factor markets clear. That is, each firm uses an **optimal input demand**

$$z_j^* \in z_j(\boldsymbol{p}, \boldsymbol{w}^*) \quad \text{for all } j = 1, ..., J$$

and all of the endowment goods are used in production:

$$\sum_j z_{\ell j}^* = \bar{z}_\ell \quad \text{for all } \ell = 1, ..., L.$$

Since the production functions for the firms are concave function, the first order conditions are both necessary and sufficient conditions for the characterisation of the optimal factor demands. Therefore, the $L(J+1)$ variables formed by the factor allocation $(\boldsymbol{z}_1^*, ..., \boldsymbol{z}_J^*) \in \mathbb{R}_+^{LJ}$ (the mix over L factors by the J firms) and the factor prices $\boldsymbol{w}^* = (w_1^*, ..., w_L^*)$ constitute an equilibrium if and only if they satisfy the following $L(J+1)$ equations (under the assumption of an interior solution):

$$p_j \frac{\partial f_j(\boldsymbol{z}_j^*)}{\partial z_{\ell j}} = w_\ell^* \quad \text{for } j = 1, ..., J \text{ and } \ell = 1, ..., L \tag{136}$$

and

$$\sum_j z_{\ell j}^* = \bar{z}_\ell \quad \text{for all } \ell = 1, ..., L. \tag{137}$$

based off the firm's profit maximising production plan in Eq. (135).

The **equilibrium output levels** are then:

$$q_j^* = f_j(\boldsymbol{z}_j^*)$$

for every firm j.

Alternatively, we could have written the equilibrium conditions for outputs and factors prices using the firm's cost functions $c_j(\boldsymbol{w}, q_j)$ for each firm $j = 1, ..., J$. Under this scenario, the output levels $\boldsymbol{q}^* = (q_1^*, ..., q_J^*) >> 0$ and factor prices $\boldsymbol{w}^* >> 0$ constitute an equilibrium if and only if the following first order conditions hold:

$$p_j = \frac{\partial c_j(\boldsymbol{w}^*, q_j^*)}{\partial q_j} \quad \text{for } j = 1, ..., J$$

and

$$\sum_j \frac{\partial c_j(\boldsymbol{w}^*, q_j^*)}{\partial w_\ell} = \bar{z}_\ell \quad \text{for } \ell = 1, ..., L.$$

By **Shephard's lemma**, this second condition says that firm j's optimal demand for the ℓ-th input is $z_{\ell j}^* = \partial c_j(\boldsymbol{w}, q_j^*)/\partial w_\ell$. Hence, this is the factor market clearing condition!

7.3.1 2 x 2 Production Model

We now take a specific example and look at a 2×2 production model whereby we set the number of firms $J = 2$ and the number of factors used in production $L = 2$. That is, we have two output goods being produced from two primary factors.

We assume that the production functions $f_1(z_{11}, z_{21})$ and $f_2(z_{12}, z_{22})$ are homogeneous of degree one (constant returns to scale). We interpret factor 1 as labour and factor 2 as capital.

For every vector of factor prices $\boldsymbol{w} = (w_1, w_2)$, we denote by $c_j(\boldsymbol{w})$ the **minimum cost** of producing one unit of good j and by $\boldsymbol{a}_j(\boldsymbol{w}) = (a_{1j}(w), a_{2j}(w))$ the **unique input combination** with which this minimum cost is reached.

Definition 294 (Unit Cost Function)
The unit cost function for good j is the minimum cost of producing good j:

$$c_j(\boldsymbol{w}). \tag{138}$$

By **Shephard's lemma**, we can relate the two vectors by:

$$\nabla c_j(\boldsymbol{w}) = (a_{1j}(\boldsymbol{w}), a_{2j}(\boldsymbol{w})).$$

We can plot the unit isoquant for firm j as:

$$\{(z_{1j}, z_{2j}) \in \mathbb{R}^2_+ : f_j(z_{1j}, z_{2j}) = 1\}$$

which is the mixture of factor inputs to produce one unit of output j. We can also plot the cost-minimising input combination $(a_{1j}(\boldsymbol{w}), a_{2j}(\boldsymbol{w}))$. We show this in Fig. 45.

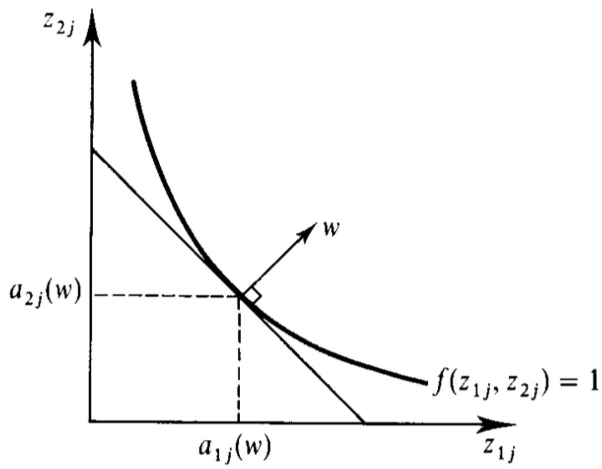

Figure 45: Unit isoquant of producing one unit of good j. We also show the cost-minimising input combination.

In Fig. 46, we plot the level curve of the unit cost function:

$$\{(w_1, w_2) : c_j(w_1, w_2) = \overline{c}\}.$$

The curve is downward sloping because, as the factor price w_1 increases, the factor price w_2 must fall to keep the minimised cost of producing one unit of good j unchanged. Moreover, the upper contour set is convex because the cost function $c_j(\boldsymbol{w})$ is concave in $\boldsymbol{w}$. The vector $\nabla c_j(\overline{\boldsymbol{w}})$ is normal to the level curve and is exactly $(a_{1j}(\overline{\boldsymbol{w}}), a_{2j}(\overline{\boldsymbol{w}}))$. As we move along the curve towards a higher factor price w_1 and lower factor price w_2, the ratio $a_{1j}(\boldsymbol{w})/a_{2j}(\boldsymbol{w})$ falls.

We now restrict our attention to **efficient factor allocations**. We can represent the allocations by an Edgeworth box in Fig. 48 where the length of the box represents the total endowment of each factor $\overline{z}_1, \overline{z}_2$. The factors used by firm 1 is measured from the southwest corner whilst the factors used by firm 2 is measured from the northeast corner. In Fig. 47, we can also draw the Pareto set of efficient factor allocations.

From all feasible factor allocations, we can represent the non-negative **output pairs** (q_1, q_2) that can be produced using the economy's available factor inputs. This is the **production possibility set**; its frontier is generated by efficient factor allocations such as those on the Pareto set in Fig. 47. We plot this convex set in Fig. 49.

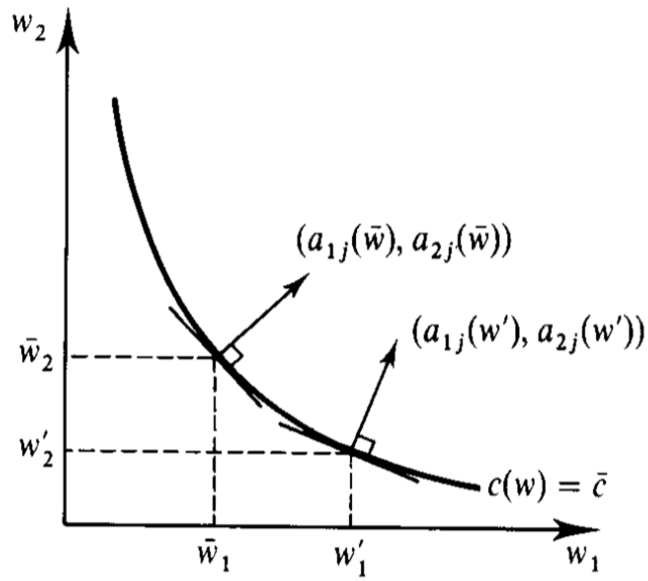

Figure 46: Unit cost function.

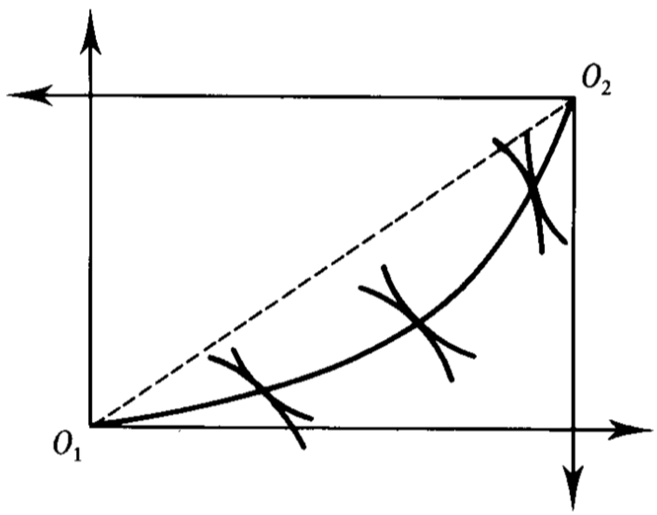

Figure 47: Edgeworth box showing Pareto efficient allocation of factors.

We are now interested in what determines the equilibrium factor allocations (z_1^*, z_2^*) and the corresponding equilibrium factor prices $\boldsymbol{w}^* = (w_1^*, w_2^*)$. Factor intensity tells us how intensively one good uses a particular factor relative to another good.

Definition 295 (Factor Intensive)
The production of good 1 is **relatively more intensive in factor 1** than is the production of good 2 if:

$$\frac{a_{11}(\boldsymbol{w})}{a_{21}(\boldsymbol{w})} > \frac{a_{12}(\boldsymbol{w})}{a_{22}(\boldsymbol{w})} \tag{139}$$

at **all factor prices** $\boldsymbol{w} = (w_1, w_2)$.

Remark 296. *Suppose that we had our output goods being computers and shoes. Furthermore, assume that our*

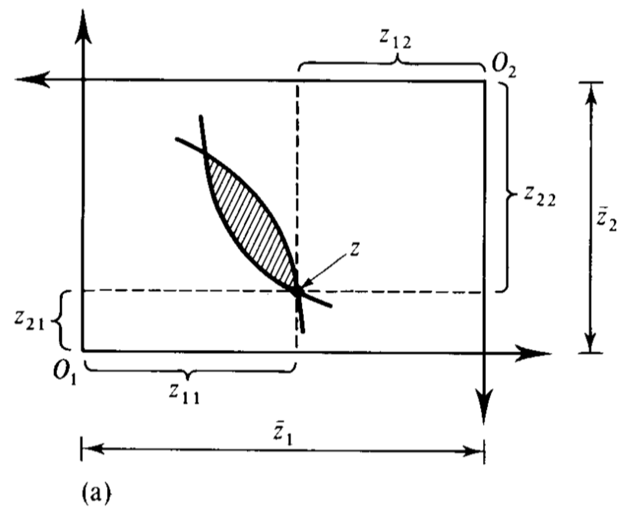

Figure 48: Edgeworth box showing allocation of factors $\boldsymbol{z}$ between two firms.

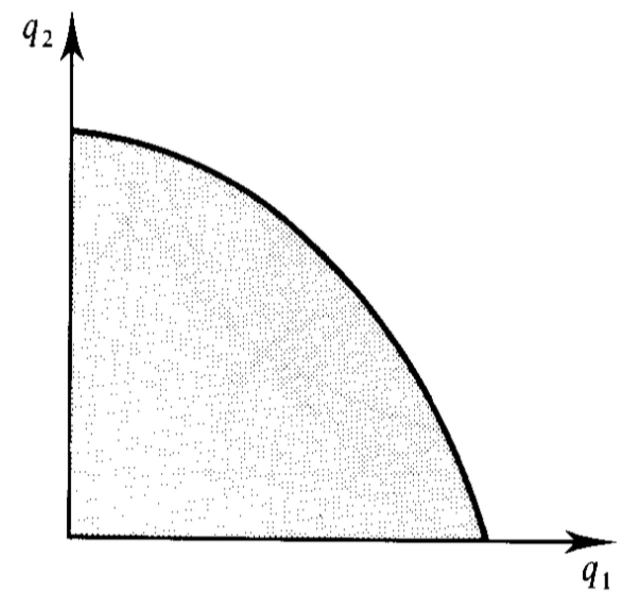

Figure 49: Production possibility set arises from the Pareto set of efficient factor allocations.

factors are capital and labour. We can say that computers are relatively more intensive in capital compared to shoes if the ratio of capital to labour used in the production of computers is higher than the ratio of capital to labour used in the production of shoes:

$$\frac{K^C}{L^C} > \frac{K^S}{L^S}$$

where K^C, L^C is the capital and labour used in producing computers and K^S, L^S is the capital and labour used in producing shoes.

We now want to determine the equilibrium factor prices. First, we assume that we have an interior equilibrium so both goods are produced. Given our constant returns to scale assumption, a necessary condition for the factor prices (w_1^*, w_2^*) to be the factor prices in an interior equilibrium requires that it satisfies the system of equations:

$$\begin{cases} c_1(w_1, w_2) = p_1 \\ c_2(w_1, w_2) = p_2 \end{cases}$$

whereby prices are equal to unit cost. This gives us two equations for the two unknown factor prices w_1, w_2. In Fig. 50, we show the equilibrium factor prices and factor intensities whereby prices are equal to unit cost. By the condition that prices equal cost, we require that the two cost curves intersect at the equilibrium factor prices (w_1^*, w_2^*). Furthermore, by the factor intensity assumption, the cost curve for firm 2 must be flatter (less negatively sloped) than that for firm 1 since the derivatives of the cost curves were given by:

$$\nabla c_j(\boldsymbol{w}) = \big(a_{1j}(\boldsymbol{w}), a_{2j}(\boldsymbol{w})\big).$$

These two cost curves intersect at most one time and hence there is **at most a single pair of factor prices that can arise as the equilibrium factor prices of an interior equilibrium**.

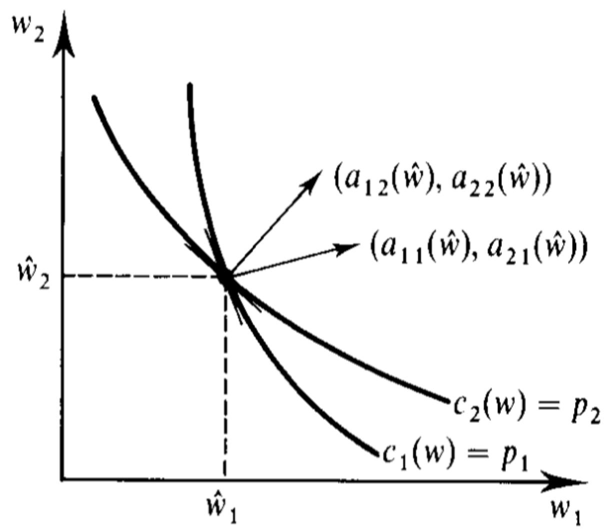

Figure 50: Equilibrium factor prices and factor intensities in an interior equilibrium. We see that prices p_j are equal to costs $c_j(w_1, w_2)$.

Once we know the equilibrium factor prices $\boldsymbol{w}^*$, we can then compute the equilibrium output levels. To do so, we go back to the Edgeworth box. We then find the unique point (z_1^*, z_2^*) in the Edgeworth box of factor allocations at which both firms have the factor intensities associated with factor prices $\boldsymbol{w}^*$:

$$\frac{z_{11}^*}{z_{21}^*} = \frac{a_{11}(\boldsymbol{w}^*)}{a_{21}(\boldsymbol{w}^*)}$$
$$\frac{z_{12}^*}{z_{22}^*} = \frac{a_{12}(\boldsymbol{w}^*)}{a_{22}(\boldsymbol{w}^*)}$$

We show this in Fig. 51.

All this gives us a key insight into our model.

Theorem 297 (Price Factor Equalisation Theorem)
In the 2×2 production model, under factor intensity and assuming that the economy does not specialise in the

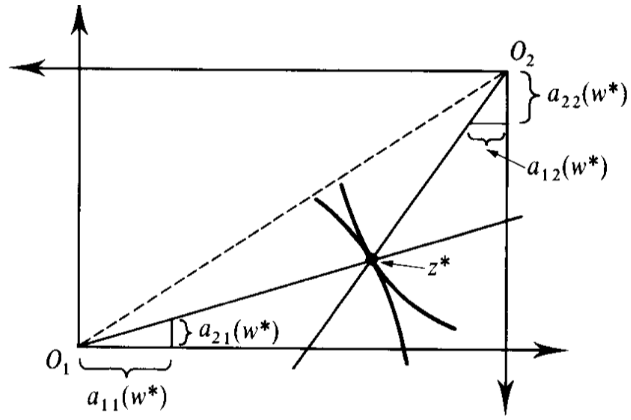

Figure 51: Edgeworth box showing the equilibrium factor allocation whereby both firms have the factor intensities associated with factor prices binding.

production of a single good, the **equilibrium factor prices depend only on the technologies of the two firms and on the output prices p.**

Remark 298. *This theorem provides the conditions under which the prices of nontradable factors are equalised across nonspecialised countries.*

We can now look at two comparative statics exercises. The first question we ask is *how does a change in the price of one of the outputs affect the equilibrium factor prices and factor allocations?* Suppose that one of the output prices, say p_1 changes. Then, the increase in the output price p_1 shifts firm 1's curve

$$\{(w_1, w_2) : c_1(w_1, w_2) = p_1\}$$

outward towards a higher factor price levels. The point of intersection of the two cost curves moves out along firm 2's curve to a higher level of w_1 and a lower level of w_2.

Proposition 299 (Stolper-Samuelson Theorem)
In the 2×2 production model with the factor intensity assumption, if p_j increases, then the equilibrium price of the factor more intensively used in the production of good j increases, while the price of the other factor decreases.

Remark 300. *If the price of computers rises, then the price of the factor input of capital rises whilst the price of the factor input of labor falls.*

Proof. Recall the conditions cost equals price conditions:

$$c_1(w_1, w_2) = p_1$$
$$c_2(w_1, w_2) = p_2$$

We employ total differentiation to get:

$$\nabla c_1(\boldsymbol{w}^*) \cdot d\boldsymbol{w} = a_{11}(\boldsymbol{w}^*)dw_1 + a_{21}(\boldsymbol{w}^*)dw_2 = dp_1$$
$$\nabla c_2(\boldsymbol{w}^*) \cdot d\boldsymbol{w} = a_{12}(\boldsymbol{w}^*)dw_1 + a_{22}(\boldsymbol{w}^*)dw_2 = dp_2$$

where we used the fact that $\nabla c_j(\boldsymbol{w}) = (a_{1j}(\boldsymbol{w}), a_{2j}(\boldsymbol{w}))$. We can write this in matrix notation to get:

$$d\boldsymbol{p} = \begin{bmatrix} a_{11}(\boldsymbol{w}^*) & a_{21}(\boldsymbol{w}^*) \\ a_{12}(\boldsymbol{w}^*) & a_{22}(\boldsymbol{w}^*) \end{bmatrix} d\boldsymbol{w}.$$

We can express this 2×2 matrix by $\boldsymbol{A}$. Then, the factor intensity assumption implies that the determinant $|\boldsymbol{A}| > 0$. Therefore, the inverse $\boldsymbol{A}^{-1}$ exists. Therefore, we can compute it to be:

$$\boldsymbol{A}^{-1} = \frac{1}{|\boldsymbol{A}|} \begin{bmatrix} a_{22}(\boldsymbol{w}^*) & -a_{21}(\boldsymbol{w}^*) \\ -a_{12}(\boldsymbol{w}^*) & a_{11}(\boldsymbol{w}^*) \end{bmatrix}.$$

Therefore, the entries of this inverse $\boldsymbol{A}^{-1}$ are positive on the diagonal entries and negative on the off-diagonals. Since $d\boldsymbol{w} = \boldsymbol{A}^{-1}d\boldsymbol{p}$, this implies that for an increase in the output price of good 1 $d\boldsymbol{p} = (1, 0)$, we have that the change in the factor price $dw_1 > 0$ and $dw_2 < 0$. That is, the factor that is used more intensively in the good that had their price changes, sees that factor's price go up!

□

We show the Stolper-Samuelson theorem in Fig. 52.

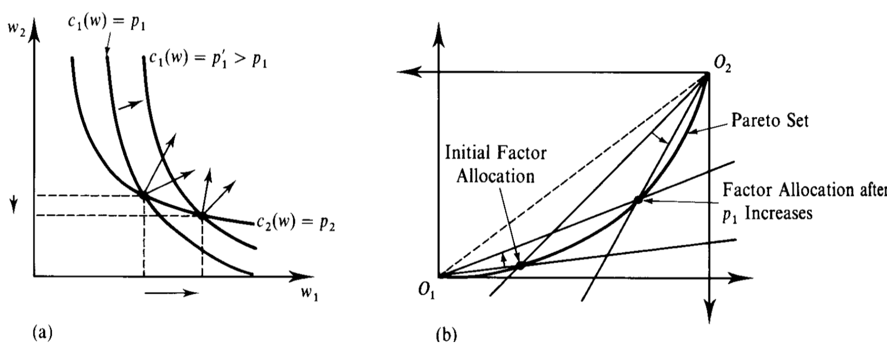

Figure 52: The image on the left shows the change in the equilibrium factor prices as the output price for good 1 increases. The image on the right shows the change in the equilibrium factor allocation in the Edgeworth box representing factors.

If p_1 increases, the Stolper–Samuelson theorem implies that w_1^*/w_2^* increases, so both firms substitute away from factor 1. Factor allocations adjust accordingly; the theorem alone does not determine the direction of output changes.

The second comparative statics question we ask is *how does the change in availability of a factor endowment affect the equilibrium factor prices and factor allocations?* Since neither output prices nor technologies have been changed, the factor input prices will **not** change either. As a result, the factor intensities do not change.

Proposition 301 (Rybczynski Theorem)
In the 2×2 production model with the factor intensity assumption, if the endowment of a factor increases, then

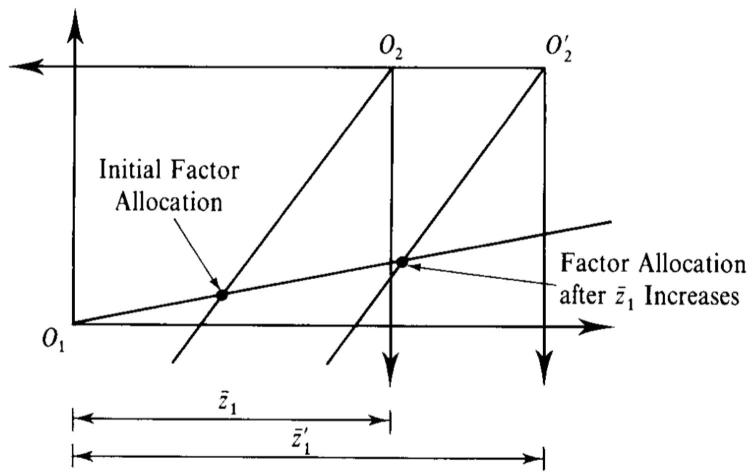

Figure 53: Rybczynski theorem. Initially, the total endowment of factor 1 is $\overline{z}_1$. After it increases to $\overline{z}_1'$, output of the good that uses factor 1 intensively rises at unchanged factor intensities.

the production of the good that uses this factor relatively more intensively increases and the production of the other good decreases.

Remark 302. *If the endowment of capital increases, then the production of computer increases and the production of shoes decreases.*

We can show this in Fig. 53.

7.4 Aggregate Demand

We assume that the economy has L commodities:

$$\mathbb{R}^L = \{x = (x_1, ..., x_L) : x_i \in \mathbb{R}\}$$

which is Euclidean space. We assume that there is a finite set of agents A. Agents behave rationally and maximise their utility functions:

$$U^a : \mathbb{R}^L \to \mathbb{R}$$

for all agents $a \in A$.

We assume that agents have endowments of goods $\omega^a \equiv (\omega_1^a, ..., \omega_L^a) \geq 0$, with $\omega^a \neq 0$, so that every agent has a positive endowment of at least one good. Furthermore, every good has a positive aggregate endowment:

$$\overline{\omega} \equiv \sum_{a \in A} \omega^a >> 0.$$

We assume that agents have perfect information about all fundamentals of the model. Goods are traded on the market and agents are price takers.

Assumption (Utility Function)

For every agent $a \in A$ her utility function U^a is:

1. **Continuous**: For any consumption sequence $x^{(n)}$ which converges to x, we have:

$$\lim_{x^{(n)} \to x} U^a(x^{(n)}) = U^a(x) \tag{140}$$

2. **Strictly increasing**: For every agent $a \in A$, for consumption bundle $x > x'$, the utility function:

$$U^a(x) > U^a(x')$$

3. **Strictly quasi-concave**: For any distinct bundles $x \neq x'$ and every $t \in (0, 1)$, we have:

$$U^a(tx + (1-t)x') > \min\{U^a(x), U^a(x')\}.$$

We now repeat ideas from consumer theory.

Definition 303 (Budget Set)
For agent a, consider a price vector $p \in \mathbb{R}^L_{++}$ and income w^a. Then, agent a's **budget set** is:

$$B(p, w^a) = \{x \in \mathbb{R}^L_+ : p \cdot x \leq w^a\}. \tag{141}$$

Definition 304 (Marshallian Demand Bundle)
A commodity bundle x^a is a **demand bundle** of agent a if it maximises the utility $U^a(\cdot)$ over the budget set $B(p, w^a)$, i.e.

$$x^a(p, w^a) \in \arg\max_{x \in B(p, w^a)} U^a(x). \tag{142}$$

We have the following properties surrounding the agent's Marshallian demand.

Proposition 305 (Marshallian Demand Properties)
Suppose that U^a is continuous, strictly increasing, and strictly quasi-concave. Then,

1. There exists a **unique** bundle $x^a(p, w^a)$ that solves the consumer problem;
2. $x^a(p, w^a)$ is continuous;
3. $x^a(p, w^a)$ is homogeneous of degree zero;
4. $x^a(p, w^a)$ exhausts the budget constraint, i.e. $p \cdot x^a(p, w^a) = w^a$.

We now want to aggregate up these demand functions.

Definition 306 (Excess Demand Function)
Agent a's **excess demand function** is defined as:

$$z^a(p) = x^a(p, p \cdot \omega^a) - \omega^a \tag{143}$$

where $x^a(p, p \cdot \omega^a)$ is the Marshallian demand and ω^a is the agent's endowment.

Proposition 307 (Properties of Individual Excess Demand)

The excess demand function $z^a(p)$ satisfies the following properties:

1. $z^a(p)$ is bounded from below by $-\omega^a$;
2. $z^a(p)$ is continuous;
3. $z^a(p)$ is homogeneous of degree zero;
4. $z^a(p)$ satisfies $p \cdot z^a(p) = 0$.

Proof. 1. Since consumption is non-negative, excess demand is bounded below by $-\omega^a$.

2. This follows from the continuity of Marshallian demand $x^a(\cdot)$.
3. This follows because Marshallian demand $x^a(\cdot)$ is homogeneous of degree zero.
4. Since Marshallian demand exhausts wealth, $p \cdot x^a(p, p \cdot \omega^a) = p \cdot \omega^a$, so $p \cdot z^a(p) = 0$.

$\square$

We can now aggregate the excess demand of all agents.

Definition 308 (Aggregate Excess Demand)

For all $p >> 0$, the **aggregate excess demand** function is defined as:

$$z(p) = \sum_{a \in A} z^a(p) = \sum_{a \in A} \left(x^a(p, p \cdot \omega^a) - \omega^a \right) \tag{144}$$

We can state some properties of the aggregate excess demand.

Proposition 309 (Properties of Aggregate Excess Demand)

The aggregate excess demand function $z(p)$ satisfies the following properties:

1. $z(p)$ is bounded from below;
2. $z(p)$ is continuous;
3. $z(p)$ is homogeneous of degree zero;

The aggregate excess demand function also satisfies another important result.

Theorem 310 (Walras Law)

For all prices $p >> 0$, the aggregate excess demand function $z(p)$ satisfies **Walras Law**:

$$p \cdot z(p) = 0. \tag{145}$$

Remark 311. *Walras' law says that the value of aggregate excess demand is zero: positive excess demands are offset, in value, by excess supplies. It does not say that every market clears at every price vector.*

7.5 Equilibrium

We now look at equilibria characterised by the aggregate excess demand function.

Definition 312 (General Equilibrium)
A **general equilibrium** in an exchange economy is a price vector p^* such that all markets clear:

$$z(p^*) = 0. \tag{146}$$

Remark 313. *Notice that markets clearing is equivalent to the excess demand function being zero.*

The question is whether such a price vector exists. Under the standard assumptions, the answer is yes.

Theorem 314 (Existence of General Equilibrium)
Suppose that the aggregate excess demand function $z(\cdot)$ satisfies the standard continuity, homogeneity, lower-bound, boundary, and Walras-law conditions. Then, there exists a price vector p^* such that

$$z(p^*) = 0. \tag{147}$$

Proof. Use fixed-point theorem. □

7.6 Two Welfare Theorems

We now work our way towards the two welfare theorems.

Definition 315 (Allocation)
An **allocation** is a map $x : A \to \mathbb{R}^L$, whereby x^a is the bundle given to agent a.

Definition 316 (Feasible Allocation)
An allocation is **feasible** if:

$$\sum_{a \in A} x^a = \sum_{a \in A} \omega^a \tag{148}$$

all allocations equal to all endowments.

Definition 317 (Walrasian Equilibrium)
A price-allocation pair (p^*, x^*) is a **Walrasian equilibrium** if the allocation is feasible, each consumer is optimising their given budget set, and

$$\sum_{a \in A} x^a(p^*, p^* \cdot \omega^a) = \sum_{a \in A} \omega^a. \tag{149}$$

Definition 318 (Pareto Improvement)
An allocation y is a **Pareto improvement** over allocation x if:

1. For every agent $a \in A$, $U^a(y^a) \geq U^a(x^a)$;
2. There exists an agent $b \in A$ such that $U^b(y^b) > U^b(x^b)$.

We can now state the first welfare theorem.

Definition 319 (Pareto Optimal)
A feasible allocation $x = \{x^a\}_{a \in A}$ is **Pareto optimal** (or **Pareto efficient**) if there is no other feasible allocation that is a Pareto improvement over it.

Theorem 320 (First Welfare Theorem)
If agents' utility functions are strictly increasing, then **every Walrasian allocation is Pareto efficient**.

Remark 321. *Note that the theorem does not say that a Walrasian allocation exist. But if it does, then it is Pareto efficient. Note that the welfare theorem implicitly assumes price-taking behaviour and complete markets.*

We are now interested in whether we can transfer resources between agents. An activity-independent reallocation is called a **lump-sum transfer**.

Definition 322 (Walrasian Allocation with Transfers)
An allocation $\{y^a\}_{a \in A}$ is a **Walrasian allocation with transfers** if it is feasible and there is a price vector $p^* >> 0$ and a vector of zero-sum transfers $t \in \mathbb{R}^{|A|}$, with $\sum_{a \in A} t^a = 0$, such that

$$y^a = x^a(p^*, p^* \cdot \omega^a + t^a) \tag{150}$$

where the new demand y^a is the old demand x^a under the new wealth with transfers t^a.

Here, $\{y^a\}_{a \in A}$ is an allocation that can be induced in general equilibrium by redistributing wealth among agents via lump-sum taxes and transfers. We now ask which allocations can be supported as equilibria with transfers. The next theorem tells us.

Theorem 323 (Second Welfare Theorem)
Suppose that, for all agents $a \in A$, utility U^a is continuous, strictly increasing, and strictly quasiconcave. Then, **every Pareto-efficient allocation can be supported as a Walrasian equilibrium with suitable lump-sum transfers**.

Remark 324. *The second welfare theorem requires much stronger conditions on the utility function U^a compared to the first Welfare theorem.*

7.7 Production

We are now interested in introducing production into our framework. We let $y \in \mathbb{R}^L$ be the input-output vector. Firm $j \in J$ is characterised by its **production set** and **ownership structure**. We make the same assumptions on the production set Y as we did in the production section of the course.

One new thing we introduce is that firms j are owned by the consumers a. Therefore, we denote θ_{aj} to be the fraction of firm j that is owned by agent a. Clearly, for a given firm j, the agents must own 100% of firm j:

$$\sum_{a \in A} \theta_{aj} = 1$$

The agents receive the profits from firm j proportional to their ownership of the firm. We can now define what a Walrasian equilibrium is in such a economy.

Definition 325 (Walrasian Equilibrium in Production Economy)
Given a private ownership economy specified by consumer's utility functions and endowments, $\{(U^a(\cdot), \omega^a)\}_{a \in A}$, firm's production sets $\{Y^j\}_{j \in J}$ and ownership structures $\{\theta_{aj}\}_{a \in A, j \in J}$. Then, an allocation (x_0, y_0) and a price vector p constitutes a **Walrasian equilibrium** if:

1. $y_0^j \in Y^j$ and $p \cdot y^j \leq p \cdot y_0^j$ for all firms $j \in J$ and all other production plans $y^j \in Y^j$
2. For all agents $a \in A$, $x_0^a \in B^a(p) \equiv \{x : p \cdot x \leq p \cdot \omega^a + \sum_{j \in J} \theta_{aj} p \cdot y_0^j\}$ and $U^a(x_0^a) \geq U^a(x^a)$ for all $x^a \in B^a(p)$;
3. $\sum_{a \in A} x_0^a = \overline{\omega} + \sum_{j \in J} y_0^j$.

Remark 326. *The first condition says that each firm j maximises its profits by choosing the production plan that earns the highest profits.*

The second condition says that consumers maximise utility over budget sets consisting of the market value of their endowments plus profits from the firms they own.

The third condition says that markets clear whereby what agents consume is equal to the initial endowments $\overline{\omega}$ and what firms produce.

We can also define an excess demand function like we did before.

Definition 327 (Production Included Excess Demand)
The **production inclusive excess demand function** is defined by:

$$z(p) \equiv \sum_{a \in A} x^a \left(p, p \cdot \omega^a + \sum_{j \in J} \theta_{aj} p \cdot y^j(p) \right) - \sum_{a \in A} \omega^a - \sum_{j \in J} y^j(p) \tag{151}$$

Remark 328. *Notice that the Marshallian demand depends on the wealth of an agent being the sum of their endowment plus the profits they earn from their ownership of firms. We then subtract the agent's initial endowment and what the firms produce to compute the excess demand function.*

We can now state the two Welfare theorems which are identical to what we saw before.

Theorem 329 (First Welfare Theorem)
If agents' utility functions are strictly increasing, then **every Walrasian allocation is Pareto efficient**.

Theorem 330 (Second Welfare Theorem)
Under the standard convexity assumptions on preferences and production sets, every Pareto-efficient allocation

(x^*, y^*) can be supported as a **Walrasian equilibrium** after suitable lump-sum redistribution of wealth or ownership rights.

8 Introduction to Game Theory

8.1 Introduction

We are studying decision making by individuals where the payoff of an individual depends not only on their own decision, but decisions made by others. There is **strategic interdependence**.

A **game** is a representation of a situation in which multiple individuals interact and the welfare of each individual depends not only on their own decision but on the decisions of other individuals. In games, individuals will search for the best strategy in order to win. There are four things we need to specify a game:

1. The **players** involved in the game;
2. The **rules** of the game. This includes the sequential ordering of who moves and what moves they are allowed to make. Furthermore, we need to specify what the players **know** when they are making their move;
3. The different **outcomes** of the game, which is based on the **set of actions** that could occur;
4. The **payoffs** of the game and the **players' preferences** over these outcomes.

The boldfaced terms will appear in the formal definition below. It is important to distinguish the outcome of a game from the players' payoffs. Games can be represented in two ways: **extensive form** and **normal form**.

8.2 Extensive Form Representation of a Game

The **extensive form representation** of a game specifies each of the previous elements of a game, including the timing and sequencing of the game. We can use a **tree** to represent the sequencing of decisions in a game. In Fig. 54, we give an example of a game in its extensive form representation. We now go through the components of a game in extensive form whilst making references to Fig. 54 before giving the formal definition.

Definition 331 (Finite Set of Players)
We denote the finite set of n strategic players by $N = \{1, ..., n\}$. When Nature moves, we label it player 0, so the enlarged set is $\{0\} \cup N$.

In Fig. 54, we have that $N = \{A, B\}$. Now, we need to define an important partial order.

Definition 332 (Precedence Relationship)
Let X be a finite set of nodes. Then, the **precedence relationship** is defined as:

$$p : X \longrightarrow X \cup \{\emptyset\} \tag{152}$$

where $p(x)$ is the **unique immediate predecessor** of x, with $p(x_0) = \emptyset$. The immediate-predecessor map induces the asymmetric and transitive precedence relation on the tree.

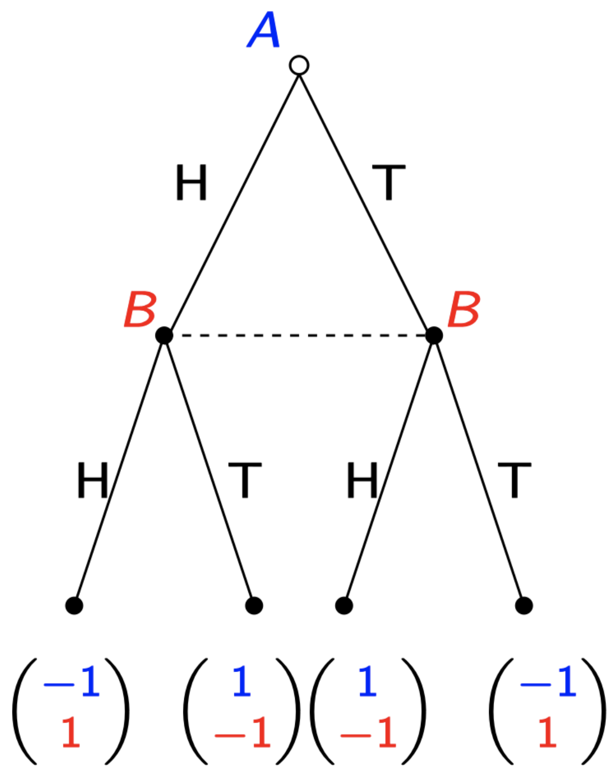

Figure 54: Extensive form representation of a game.

In Fig. 54, the precedence relationship is represented by the (directed) edges in our tree. We require that for all nodes besides the initial node (which we will define shortly), their predecessor must be in the set of nodes X. Mathematically, this is: $x \in X \setminus \{x_0\}, p(x) \in X$.

There are 3 important types of nodes in our game.

Definition 333 (Initial, Terminal, and Decision nodes.)
There are three types of nodes in our game:

- The **initial** node x_0: This is the node with **no** predecessor, so $p(x_0) = \emptyset$.
- The **terminal** nodes $t \in T \subset X$: These are the nodes with no successors. Mathematically speaking, there is no node $x \in X$ such that $t = p(x)$.
- The **decision** nodes. These are all non-terminal nodes, which includes the initial one.

Definition 334 (Set of Actions)
The set $\mathcal{A}$ is a finite set of actions and a function $\alpha : X \setminus \{x_0\} \longrightarrow \mathcal{A}$ that gives the action that leads to any non-initial node from its predecessor. α has to satisfy that if $x' \neq x''$ and $p(x') = p(x'')$ then $\alpha(x') \neq \alpha(x'')$.

We denote by $A(x) \subset \mathcal{A}$ the **set of actions available** at node $x \in X \setminus T$.

Remark 335. *We require that every decision node, and hence every information set, is assigned to a single player.*

$\mathcal{A}$ denotes the set of actions available when a player is at node x.

Definition 336 (Information Sets)
The collection of information sets $\mathcal{I}$ is a partition of the decision nodes: for every decision node $x \in X \setminus T$, there is a unique information set $I \in \mathcal{I}$ such that

$$x \in I \subseteq X \setminus T.$$

The information that a player has when they are called to act is represented by **information sets**. This is the set of decision nodes that they cannot distinguish between. This means that they do not know which **path of play** led to their current decision.

Proposition 337 (Set of Feasible Actions)
All the nodes x, x' in the same information set I have the same set of actions: for all $x, x' \in I$,

$$A(x) = A(x').$$

Remark 338. *If the set of actions at every node in the same information set was not the same, then the player can deduce where they are at, which is a contradiction! Sometimes we say $A(I)$ to denote the actions available at an information set I.*

We state an important property of the information set partition.

Proposition 339 (Information Sets are Disjoint)
Information sets are **disjoint** whereby for two information sets $I \neq I'$, we have that:

$$I \cap I' = \emptyset$$

In Fig. 54, we denote the dotted line as the two nodes being in the same information set.

Definition 340 (Assignment of Information Sets)
We denote the function

$$\iota : \mathcal{I} \longrightarrow \{0, 1, ..., n\}$$

as the function that assigns each information set to a player or to Nature (player 0) that represents chance.

Remark 341. *The function ι assigns information sets to players. It is crucial that each information set and each action should be allocated to a single player.*

Definition 342 (Probability Measure of Chance)
For each information set I at which Nature moves ($\iota(I) = 0$), $\rho(I, \cdot)$ is a probability distribution over $A(I)$: $\rho(I, a) \geq 0$ and $\sum_{a \in A(I)} \rho(I, a) = 1$.

In Fig. 55, we give an example of where there is a 50% chance from nature of where we move.

Definition 343 (Payoffs)
We denote the collection of payoff functions as

$$u = \{u_1(\cdot), ..., u_n(\cdot)\}$$

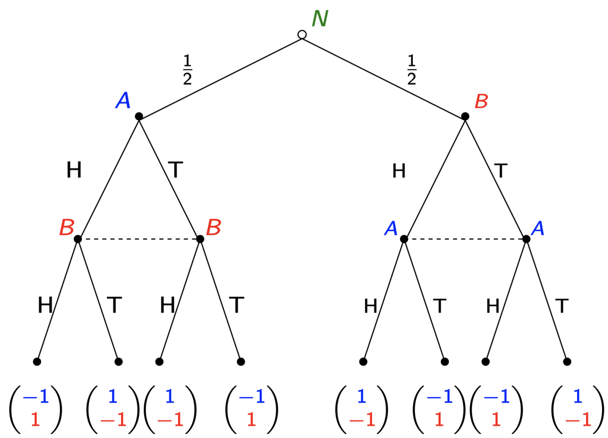

Figure 55: Game with chance.

at each terminal node whereby each payoff $u_i : T \to \mathbb{R}$ is a **Bernoulli utility**.

Remark 344. *We will assume throughout this course that players are expected utility maximisers: if there is shared uncertainty (moves by Nature) the players have the corresponding expected utility.*

We can now collect all our definitions to define a game in extensive form.

Definition 345 (Game in Extensive Form)
A game in extensive form Γ_E is specified by:

$$\Gamma_E \equiv \{N, X, p(\cdot), \mathcal{A}, \alpha(\cdot), \mathcal{I}, \iota(\cdot), \rho(\cdot), u\} \tag{153}$$

In addition, we say that the structure of the game including payoffs is known by all players. That is, the structure of the game is **common knowledge**. In our definition, we have assumed finite number of players, nodes, and actions. This is known as a **finite game**. However, this **does not** have to be the case.

Definition 346 (Perfect Information)
A game has **perfect information** if all the information sets are singleton sets.

Remark 347. *A game with perfect information implies that, at each decision node, the player who moves knows the full history of play up to that point.*

This definition of perfect information implies that a game where there is an information set with more than one node is classified as a game with **imperfect information**. We see in Fig. 56, an example of a game with perfect information and a game with imperfect information.

We make the assumption going forward that players have **perfect recall**.

Definition 348 (Perfect Recall)
Players have **perfect recall** if they do not forget the information they had before nor the action that they took.

Fig. 57 gives an example where a player does not have perfect recall.

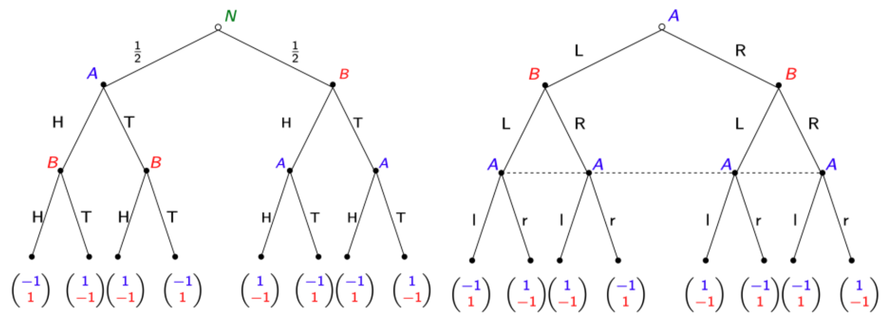

Figure 56: Game on the left has perfect information since it has only singleton information sets. The game on the right has imperfect information as it has an information set that is not a singleton set.

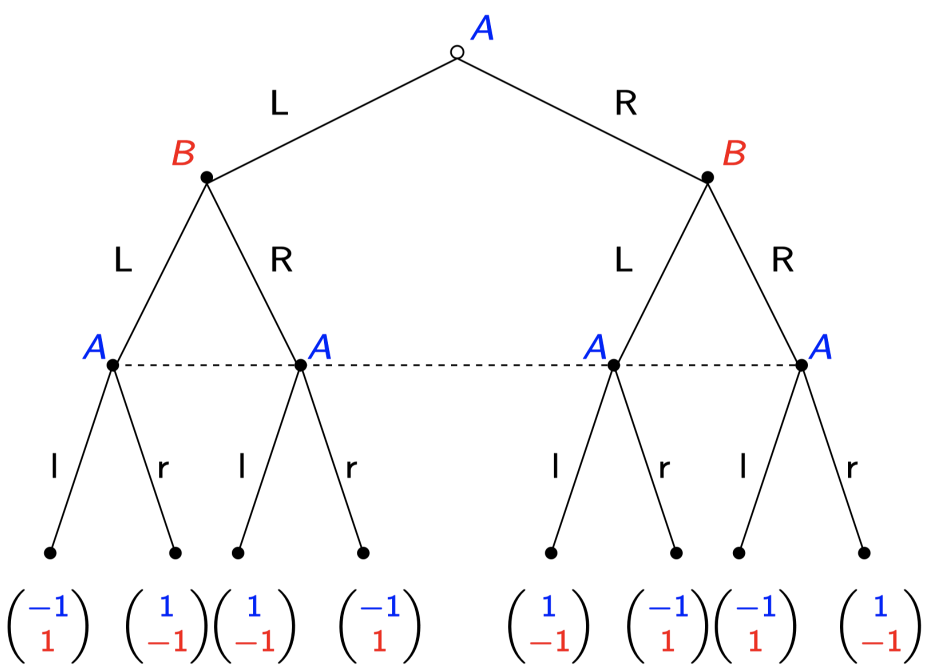

Figure 57: Player A here does not have perfect recall as they have forgotten what they have done in the first decision node. In particular, they have only one information set at the second level of the tree. If they did recall what they did in their first move, there should be at least two information sets. It is the dotted line between the two sub-trees that indicates that they have forgotten what they did in the first turn.

Definition 349 (Pure Strategy)

Let $\mathcal{I}_i$ denote the set of player i's information sets. A **pure strategy** for player i is a function

$$s_i : \mathcal{I}_i \to \mathcal{A} \tag{154}$$

such that at information set I, the action chosen by the strategy $s_i(I)$ is feasible at I, denoted by $s_i(I) \in A(I)$. We denote the **set of pure strategies** of player i by S_i.

A strategy s_i for player i tells us what action player i will take at **every** one of their information sets in the game.

Remark 350. *Strategies must always be specified even for information sets that are never reached as it allows us to* *make comparisons on what to do.*

We can collect together the strategies of all the players in our game.

Definition 351 (Profile of Strategies)
We denote by S_i the set of strategies of player i. If there are n players, and player i has a set of strategies S_i, a profile of strategies is

$$s = (s_1, ..., s_n)$$

where $s_i \in S_i$.

Every profile of strategies induces an outcome of the game.

Example 352
[Writing out strategies]

We apply our definitions to the example in Fig. 58. A has one information set at the initial node and two possible strategies, $s_A^1 = H$ and $s_A^2 = T$.

For B, they have **two** information sets. Furthermore, they have two actions available to them at each information set I_1 and I_2. Therefore, they have **four** possible strategies as a strategy requires both the information set and the action.

$$s_B^1 = \left(s_B^1(I_1) = H, s_B^1(I_2) = H\right)$$

$$s_B^2 = \left(s_B^2(I_1) = H, s_B^2(I_2) = T\right)$$

$$s_B^3 = \left(s_B^3(I_1) = T, s_B^3(I_2) = H\right)$$

$$s_B^4 = \left(s_B^4(I_1) = T, s_B^4(I_2) = T\right)$$

If we look at the strategy s_B^2, we can see that it contains a mapping to what action to take depend on what information set we are at. If we are at information set I_1, then we should play H ($s_B^2(I_1) = H$) and if we are at information set I_2, then we should play T ($s_B^2(I_2) = T$).

An example of a strategy profile would be:

$$(s_A^1, s_B^2) = \left(H, \left(s_B^2(I_1) = H, s_B^2(I_2) = T\right)\right)$$

Remark 353. *In simultaneous move games, strategies coincide with actions.*

8.3 Normal Form Representation of a Game

The **normal form representation** of a game is a reduced version of the extensive form representation of a game whereby we replace all the specific rules and timing by strategies. We can think about such representations as where agents decide their strategies simultaneously at the start of the game.

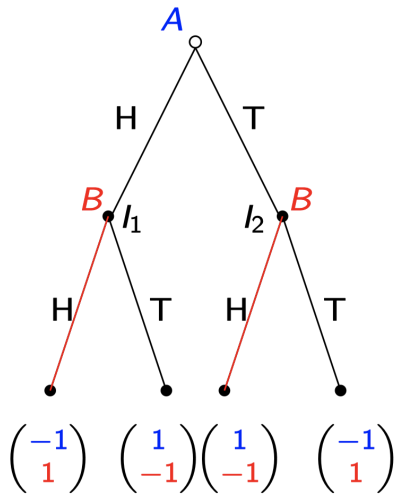

Figure 58: Hide and seek with cheating.

> **Definition 354** (Normal Form Game)
> A game in strategic form is given by
> $$\Gamma_N \equiv [N, \{S_i\}, \{u_i(\cdot)\}] \tag{155}$$
> whereby
>
> - $N = \{1, ..., n\}$ is the set of players.
> - S_i is the set of strategies of player i.
> - $u_i(s_1, ..., s_n)$ is the payoff of player i associated by the outcome arising from the strategies $(s_1, ..., s_n)$ when the terminal node $t \in T$ is reached.

Remark 355. *We will assume von Neumann-Morgenstern utilities in case Nature is a player.*

We can denote $-i$ as all other players besides player i.

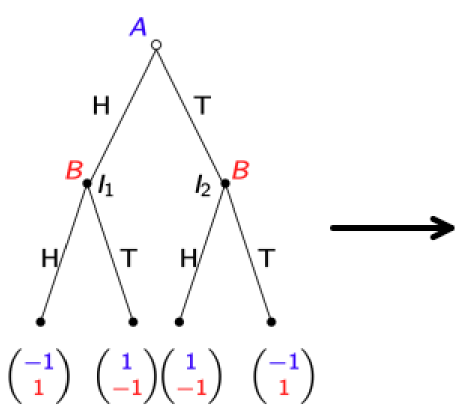

		Bob			
		HH	*TH*	*HT*	*TT*
Ana	H	-1,1	1,-1	-1,1	1,-1
	T	1,-1	1,-1	-1,1	-1,1

Figure 59: Hide and seek with cheating in normal form.

Example 356
[Normal form representation of an extensive form game]

We use Fig. 59 to illustrate how an extensive-form game is converted into normal form. Continuing from Example 352, the strategy sets for players A and B are:

$$s_A^1 = H \quad \text{and} \quad s_A^2 = T$$

and

$$s_B^1 = \left(s_B^1(I_1) = H, s_B^1(I_2) = H\right)$$
$$s_B^2 = \left(s_B^2(I_1) = H, s_B^2(I_2) = T\right)$$
$$s_B^3 = \left(s_B^3(I_1) = T, s_B^3(I_2) = H\right)$$
$$s_B^4 = \left(s_B^4(I_1) = T, s_B^4(I_2) = T\right)$$

We can then write out their contingency plans as:

$$s_A \in \{H, T\}$$
$$s_B \in \{HH, TH, HT, TT\}$$

The strategy TH for player B corresponds to s_B^3: B plays T at I_1 and H at I_2. Because B has two information sets, every strategy must specify an action at each one. We place these strategies along the rows and columns of the table in Fig. 59, then enter the payoff associated with each strategy profile.

We arrive at a particularly useful result.

Proposition 357 (Representation of Extensive Form Games)
Every extensive form representation of a game has an **unique** normal form representation.

Unfortunately, this does not hold in the reverse.

Proposition 358 (Representation of Normal Form Games)
A normal form representation of a game can have **multiple** extensive form representations.

Fig. 60 gives an example of a normal form game having multiple extensive form representations. Intuitively, this occurs because you are losing information when you are condensing a game into a normal form representation.

Definition 359 (Mixed Strategy)
Given a player i's (finite) pure strategy set S_i, a **mixed strategy** σ_i for player i is a probability measure

$$\sigma_i : S_i \to [0, 1] \tag{156}$$

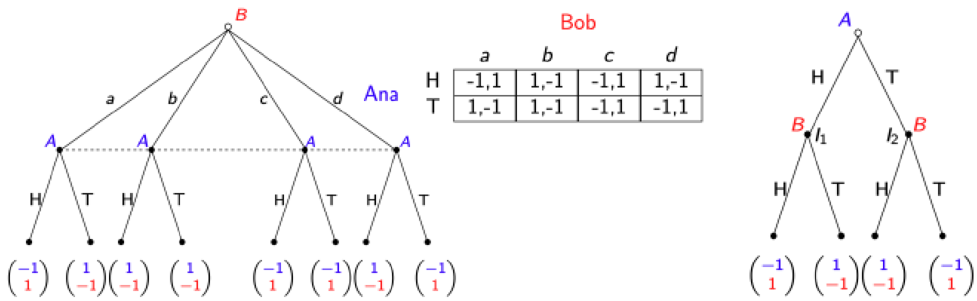

	a	b	c	d
H	-1,1	1,-1	-1,1	1,-1
T	1,-1	1,-1	-1,1	-1,1

Figure 60: The normal form game in the middle has two different extensive form game representations.

that assigns a probability $\sigma_i(s_i) \geq 0$ to each pure strategy s_i such that

$$\sum_{s_i \in S_i} \sigma_i(s_i) = 1.$$

The **set of mixed strategies** for player i is denoted by $\Delta(S_i)$.

Remark 360. *A mixed strategy is literally a probability mass function over all the strategies that is in the set of strategies S_i available to player i.*

Note that a pure strategy s_i can be seen as a mixed strategy in which $\sigma_i(s_i) = 1$. With some abuse of notation, and writing $S = \prod_i S_i$, we write

$$u_i(\sigma_1, \ldots, \sigma_n) = \sum_{s \in S} \sigma_1(s_1) \cdots \sigma_n(s_n) u_i(s)$$

and call this the **expected payoff** to player i.

Example 361

[Computing expected payoffs] We continue the example from Example 356. We have shown that the strategies available to A and B were:

$$s_A \in \{H, T\}$$
$$s_B \in \{HH, TH, HT, TT\}$$

In order to assign mixed strategies, we need to assign probabilities over these pure strategies for each player **respectively**. For example, the following is a mixed strategy for player A:

$$\sigma_A : \sigma_A(H) = \frac{1}{2} \quad \text{and} \quad \sigma_A(T) = \frac{1}{2}$$

This puts a probability of a half on both A's pure strategy of H and T. For player B, we need to put probabilities over the 4 pure strategies available to B. The following is an example of a mixed strategy for B:

$$\sigma_B : \sigma_B(HH) = \frac{1}{6}, \quad \sigma_B(TH) = \frac{1}{6}, \quad \sigma_B(HT) = \frac{1}{3}, \quad \sigma_B(TT) = \frac{1}{3}.$$

The final thing we need are the payoffs from reaching the terminal nodes. We read them off from the normal form matrix. For example:

$$u_A(H, HH) = -1 \quad \text{and} \quad u_A(T, HT) = -1$$

Therefore, we can compute the expected payoff for player A playing the mixed strategy σ_A as:

$$\begin{aligned} u_A(\sigma_A, \sigma_B) = &\tfrac{1}{2}\tfrac{1}{6}u_A(H, HH) + \tfrac{1}{2}\tfrac{1}{6}u_A(H, TH) + \tfrac{1}{2}\tfrac{1}{3}u_A(H, HT) + \tfrac{1}{2}\tfrac{1}{3}u_A(H, TT) \\ &+\tfrac{1}{2}\tfrac{1}{6}u_A(T, HH) + \tfrac{1}{2}\tfrac{1}{6}u_A(T, TH) + \tfrac{1}{2}\tfrac{1}{3}u_A(T, HT) + \tfrac{1}{2}\tfrac{1}{3}u_A(T, TT) = -\tfrac{1}{6} \end{aligned}$$

Definition 362 (Mixed Form Representation)
The mixed extension of the normal form representation of a game Γ_N is:

$$\tilde{\Gamma}_N \equiv [N, \{\Delta(S_i)\}, \{u_i(\cdot)\}] \tag{157}$$

Remark 363. *We now have a set of sets of probability measures over our pure strategies $\{\Delta(S_i)\}$.*

8.4 Randomised Choice

We have looked at mixing strategies in normal form representation of games. We can also mix our strategies in extensive form games. However, we shall in fact not randomise strategies but randomise **actions** at each information set. When we are called to play, we will mix between actions available to us at that information set.

Definition 364 (Behavioural Strategy)
A **behavioural strategy** for player i is a collection of independent probability measures:

$$\{\beta_i(I_i)\}_{I_i \in \mathcal{I}_i} \tag{158}$$

whereby $\beta_i(I_i)$ is a probability measure over the feasible actions for player i at information set I_i, denoted by $A(I_i)$.

Remark 365. *At each information set I_i, we have a probability measure $\beta_i(I_i)$ that will pick actions available to us at that information set $A(I_i)$.*

Example 366 (Behavioural Strategies in hide and seek)
We build off Example 361. We have previously used the mixed strategies:

$$\sigma_A : \sigma_A(H) = \frac{1}{2} \quad \text{and} \quad \sigma_A(T) = \frac{1}{2}$$

$$\sigma_B : \sigma_B(HH) = \frac{1}{6}, \quad \sigma_B(TH) = \frac{1}{6}, \quad \sigma_B(HT) = \frac{1}{3}, \quad \sigma_B(TT) = \frac{1}{3}.$$

We now give examples of behavioural strategies for each player. In Fig. 59, player A has one information set at the initial node and two available actions. Thus, the following probabilities define a valid behavioural strategy:

$$\beta_A : \beta_A(H) = \frac{1}{2} \quad \text{and} \quad \beta_A(T) = \frac{1}{2}$$

Player B is more complicated. Recall that player B has **two information sets** I_1 and I_2. At both information sets, B had the actions H and T available to them. A valid behavioural strategy for player B **at information set** I_1 is:

$$\beta_B(I_1, H) = \frac{1}{2} \quad \text{and} \quad \beta_B(I_1, T) = \frac{1}{2}$$

Notice that here, the sum of these probabilities over actions at information set I_1 equals one. A valid behavioural strategy for player B at information set I_2 is:

$$\beta_B(I_2, H) = \frac{1}{3} \quad \text{and} \quad \beta_B(I_2, T) = \frac{2}{3}$$

which also adds up to one at this information set.

Finally, we note the **distribution over outcomes**: $o(a_A, a_B) \in \{o(H,H), o(H,T), o(T,H), o(T,T)\}$.

We now want to compute what is the probability in reaching a given outcome based on the mixed strategies. Using the outcome (T, H) as an example, under a mixed strategy, we require that player A chooses T and player B chooses H. Under this, the probability that player A plays T is:

$$\sigma_A(T) = \frac{1}{2}.$$

The probability that player B plays H is given by:

$$\sigma_B(HH) + \sigma_B(TH) = \frac{1}{6} + \frac{1}{6}$$

Therefore, the probability that (T, H) occurs under the mixed strategy (MS) is:

$$\frac{1}{2}(\frac{1}{6} + \frac{1}{6})$$

Regarding the behavioural strategy, the probability that player A plays T is given by:

$$\beta_A(T) = \frac{1}{2}$$

and the probability that player B plays H is:

$$\beta_B(I_2, H) = \frac{1}{3}$$

where we specify the information set I_2 since we want T to be the payoff for player B. Therefore, the probability of the outcome (TH) is:

$$\frac{1}{2} \cdot \frac{1}{3}.$$

We can repeat this calculation for all the other outcomes. The resulting distribution is shown below. The two strategies induce the same probabilities over the four terminal outcomes (HH), (HT), (TH), and (TT).

$$\text{MS: } \left(\tfrac{1}{2}(\tfrac{1}{6} + \tfrac{1}{3}), \tfrac{1}{2}(\tfrac{1}{6} + \tfrac{1}{3}), \tfrac{1}{2}(\tfrac{1}{6} + \tfrac{1}{6}), \tfrac{1}{2}(\tfrac{1}{3} + \tfrac{1}{3})\right) = \left(\tfrac{1}{4}, \tfrac{1}{4}, \tfrac{1}{6}, \tfrac{1}{3}\right)$$

$$\text{BS: } \left(\tfrac{1}{2}\tfrac{1}{2}, \tfrac{1}{2}\tfrac{1}{2}, \tfrac{1}{2}\tfrac{1}{3}, \tfrac{1}{2}\tfrac{2}{3}\right) = \left(\tfrac{1}{4}, \tfrac{1}{4}, \tfrac{1}{6}, \tfrac{1}{3}\right)$$

Figure 61: Distributions over outcomes induced by the mixed and behavioural strategies.

From the previous example, we can see that the mixed and behavioural strategies induce the same distribution over the final outcomes. We can in fact show that this result holds in general.

Theorem 367 (Kuhn (1953) - Equivalency of Mixed and Behavioural Strategies)
In games of **perfect recall**, for every mixed strategy of player i, there is a corresponding behavioural strategy for that player that leads to the same distribution over outcome.

Conversely, for every behavioural strategy of player i, there is a mixed strategy for that player that leads to the same distribution over outcomes.

Remark 368. *This theorem says that there is an equivalency of mixed and behavioural strategies.*

In practice, we generally use behaviour strategies to analyse extensive form representation whilst we use mixed strategies to analyse normal form representations.

We make the following further assumptions.

Assumption (Players are Rational)
Players are **rational**. That is, players have well-behaved preferences, are aware of their alternatives, form expectations about unknowns, and choose actions that lead to their most preferred outcomes.

Assumption (Common Knowledge of Rationality)
Players are rational, they know that the other players are rational, they know that the other players know this, and so on.

9 Games of Complete Information

9.1 Dominance Strategies

We will restrict our analysis to games of **complete information** whereby every player knows every other players' payoff. Furthermore, we shall restrict our analysis to static games.

Definition 369 (Static Game)
A **static** game is a game in which players move only once and all players move at the same time or without knowing what the others have done.

Remark 370. *They are also called* **simultaneous-move** *games.*

It is important to note that even sequential games can be considered a simultaneous move game since players' strategies are complete contingent plans chosen at the beginning of the game once and for all. Therefore, the following concepts will also apply to dynamic games.

For these games it is more convenient to work with the normal-form representation:

$$\Gamma_N = [N, \{S_i\}, \{u_i(\cdot)\}]$$

for pure strategies and

$$\tilde{\Gamma}_N = [N, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$$

for mixed strategies.

Notation: Sometimes it will be convenient to isolate agent i from the rest. We denote by

$$s_{-i} \equiv (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$$

the profile of strategies for i's opponents. The whole profile can be written as $s = (s_i, s_{-i})$. We denote by

$$S = S_1 \times \ldots \times S_n$$

the set of strategy profiles, and by

$$S_{-i} = \prod_{j \neq i} S_j$$

the set of strategy profiles for i's opponents.

Analogously, when working with mixed strategies we will denote by

$$\sigma_{-i} \equiv (\sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_n)$$

for the profile of strategies for i's opponents and we will be able to write the whole profile of strategies as $\sigma = (\sigma_i, \sigma_{-i})$.

Our goal is to understand and predict how these game will be played. We start by the concept of dominance. Sometimes predicting how a game will be played is relatively simple. This is the case if each player has a strategy that is optimal regardless of what the others do. We call those strategies **dominant strategies**.

Definition 371 (Strictly Dominant Strategy)
A strategy $s_i \in S_i$ is a **strictly dominant strategy** for player i if for all $s_i' \neq s_i$, we have

$$u_i(s_i, s_{-i}) > u_i(s_i', s_{-i}) \tag{159}$$

for all $s_{-i} \in S_{-i}$.

We now define what it means for an agent to be rational with beliefs.

Definition 372 (Rational With Beliefs)
Player i is **rational with beliefs** μ if

$$s_i \in \arg\max_{s_i \in S_i} \mathbb{E}_{\mu_i(s_{-i})} u_i(s_i, s_{-i}) \tag{160}$$

that is, if s_i maximises i's expected payoff over all other players' strategies $\sum_{s_{-i}} u_i(s_i, s_{-i})\mu_i(s_{-i})$.

This allows us to state what it means for an agent to always want to play a strictly dominant strategy.

Theorem 373 (Rational Agents play Dominant Strategies)
If a player has a strictly dominant strategy and the player is rational, she should always play it.

Proof. Let s_i be our strictly dominant strategy. Then, for any other strategy s_i' and any beliefs μ_i,

$$\begin{aligned}\mathbb{E}_{\mu_i(s_{-i})} u_i(s_i, s_{-i}) &= \sum_{s_{-i}} u_i(s_i, s_{-i})\mu_i(s_{-i}) \\ &> \sum_{s_{-i}} u_i(s_i', s_{-i})\mu_i(s_{-i}) = \mathbb{E}_{\mu_i(s_{-i})} u_i(s_i', s_{-i})\end{aligned}$$

□

Remark 374. *Note that we did not require common knowledge of rationality for this.*

Proposition 375 (Non-degenerate Mixed Strategies are not Strictly Dominant)
A non-degenerate mixed strategy is never a strictly dominant strategy.

Proof. The general idea is that a mixed strategy will never dominate the pure strategies in the support of the mixed strategy. Let us denote $\sigma_i = (\sigma_i(s_1), ..., \sigma_i(s_{n_i}))$ to be the mixed strategy for player i over their n_i pure strategies. Then, the payoff for player i using this mixed strategy is:

$$u_i(\sigma_i, s_{-i}) = \sum_{k=1}^{n_i} \sigma_i(s_k) u_i(s_k, s_{-i})$$

The mixed-strategy payoff is a **convex combination of pure-strategy payoffs**. It therefore cannot be strictly greater than the payoff of every pure strategy in its own support. $\square$

Example 376 (Prisoner's Dilemma)
In the prisoner's dilemma, the dominant strategy for both players is to defect. Although both players act in their own self-interest, the resulting outcome is Pareto inefficient.

However, in most cases there are no dominant strategies. Still, we might be able to rule out some strategies by considering the rationality of agents.

Definition 377 (Strictly and Weakly Dominated Pure Strategy)
A pure strategy $s_i \in S_i$ is a **strictly dominated strategy** for player i if there exists a strategy $\sigma_i' \in \Delta(S_i)$ such that for all $s_{-i} \in S_{-i}$,

$$u_i(\sigma_i', s_{-i}) > u_i(s_i, s_{-i}). \tag{161}$$

A pure strategy $s_i \in S_i$ is a **weakly dominated strategy** for player i if there exists a strategy $\sigma_i' \in \Delta(S_i)$ such that $u_i(\sigma_i', s_{-i}) \geq u_i(s_i, s_{-i})$ for every $s_{-i} \in S_{-i}$, with strict inequality for at least one s_{-i}.

Example 378
We see in Fig. 62 that there are no dominant strategies. However, strategy B is dominated by M as player 1 can receive a payoff of 2 from playing M regardless of what player 2 does whilst player 1 will only receive a payoff of 1 from playing B. Therefore, if player 1 is rational, they will never play B.

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

Figure 62: A game in which player 1's strategy B is strictly dominated by M.

There are three comments on Definition 377.

First, note that we only need to consider profiles of **pure** strategies s_{-i} for the opponents of player i. The reason is that if these inequalities hold for all s_{-i}, then,

$$\begin{aligned} u_i(\sigma_i', \sigma_{-i}) &= \textstyle\sum_{s_{-i} \in S_{-i}} \sigma_1(s_1)...\sigma_{i-1}(s_{i-1})\sigma_{i+1}(s_{i+1})...\sigma_n(s_n)u_i(\sigma_i', s_{-i}) \\ &> \textstyle\sum_{s_{-i} \in S_{-i}} \sigma_1(s_1)...\sigma_{i-1}(s_{i-1})\sigma_{i+1}(s_{i+1})...\sigma_n(s_n)u_i(s_i, s_{-i}) \\ &= u_i(s_i, \sigma_{-i}) \quad \text{for all } \sigma_{-i} \end{aligned}$$

Therefore, if it holds for an opponent's pure strategies, then it will hold for their mixed strategies too.

Second, in the definition, we cannot replace the mixed strategy σ_i' by a pure strategy s_i'. This is because a pure strategy might be dominated by a mixed strategy even though it is not dominated by any pure strategy. We give an example in Example 379.

Example 379

We can see in Fig. 63, the pure strategy B is **not** dominated by any other pure strategies.

However, it is dominated by the **mixed strategy**:

$$\sigma_1(T) = \frac{1}{2} \quad \sigma_1(M) = \frac{1}{2}$$

since if player 2 plays L:

$$\frac{1}{2}u_1(T, L) + \frac{1}{2}u_1(M, L) = \frac{1}{2}(3) + \frac{1}{2}(0) = 1.5 > 1 = u_1(B, L)$$

and if player 2 plays R

$$\frac{1}{2}u_1(T, R) + \frac{1}{2}u_1(M, R) = \frac{1}{2}(0) + \frac{1}{2}(3) = 1.5 > 1 = u_1(B, R)$$

Therefore, B is dominated by the mixed strategy $\sigma_1 = (\sigma_1(T) = \frac{1}{2}, \sigma_1(M) = \frac{1}{2})$.

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

Figure 63: Player 1's strategy B is strictly dominated by a mixture of T and M.

We can extend our definition for strictly dominated strategies to mixed strategies.

Definition 380 (Strictly Dominated Mixed Strategy)

A mixed strategy $\sigma_i \in \Delta(S_i)$ is a **strictly dominated strategy** if there is another strategy $\sigma_i' \in \Delta(S_i)$ such that for all $s_{-i} \in S_{-i}$, $u_i(\sigma_i', s_{-i}) > u_i(\sigma_i, s_{-i})$.

We are now interested in asking the question on whether is a pure strategy s_i dominated or not. There is a three step procedure in doing so.

Proposition 381 (Procedure to check for Dominated Strategies)

1. Check if the strategy s_i is a best response to another player's strategy. That is, does there exists a s_{-i} such that $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$ for all s_i'. If so, then s_i is clearly undominated as it is a best response.

2. If not, check if s_i is dominated by another **pure strategy**.

3. If not, consider all mixed strategies with the support over the pure strategies in S_i that we know are not dominated strategies from having survived step 1 and 2. Then check is there any way for this mixed strategy to dominate our strategy of interest s_i.

We now give an example for using this procedure.

Example 382

Look at Fig. 63. We will apply our three steps to this example. We want to find all dominated strategies for player 1.

In our first step, we can see that if player 2 plays L, then player 1 playing T is a best response. Therefore, this rules out the strategy T as a dominated strategy. Moreover, we can see that if player 2 plays R, then player 1's best response is M. This rules out M as a best response too. This leaves the strategy B.

In our second step, we see that the strategy B is neither dominated by the pure strategies T or M. We therefore move onto the third step.

In this final step, we consider the mixed strategies over T and M (which we have shown to not be dominated) and see if we can dominate B with it. First, we form the mixed strategy:

$$\sigma_1 = (\sigma_1(T) = p, \sigma_1(M) = 1 - p)$$

where we need to now figure out what the probability p should be. Suppose that player 2 plays L, then we have that the payoff from the mixed strategy is higher than playing the pure strategy B if:

$$u_1(\sigma_1, L) = u_1(T, L)(p) + u_1(M, L)(1 - p) = (3)(p) + (0)(1 - p) > 1 = u_1(B, L)$$

which holds if and only if

$$p > \frac{1}{3}.$$

We now suppose that player 2 plays R and we see that the payoff from the mixed strategy is higher than the pure strategy B if:

$$u_1(\sigma_1, R) = u_1(T, R)(p) + u_1(M, R)(1 - p) = (0)(p) + (3)(1 - p) > 1 = u_1(B, R)$$

which holds if and only if

$$p < \frac{2}{3}.$$

Combining these two inequalities for p, the pure strategy B is strictly dominated by σ_1 for any $\frac{1}{3} < p < \frac{2}{3}$. This does not mean that σ_1 is itself a strictly dominant strategy.

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

Figure 64: The same dominance example, revisited for the three-step procedure.

A rational player will never play a strictly dominated strategy, since there is another strategy that will lead to a strictly better outcome no matter what the others do. Therefore, if players know that player i is rational and has a dominated strategy, they can reason that player i will never play such strategy. Eliminating such options might result in some of them having a dominated strategies which they will never play. This leads us to our next concept.

Definition 383 (Iterated Deletion of Strictly Dominated Strategies)
The process of **iterated deletion of strictly dominated strategies** involves removing strictly dominated strategies until there are no more dominated strategies left.

Note that in each iteration we are demanding a new layer of knowledge of rationality. We will require not only that they are rational, but that they know others are rational, and that they know that others know that they are rational and so on. If players are that sophisticated, we can argue that only those strategies that survive the IDSDS should ever be played.

Theorem 384 (Order of Deletion Does Not Matter for Strictly Dominated Strategies)
The set of strategies surviving the process of iterated deletion of strictly dominated strategies is independent of the order of deletion.

Remark 385. *The order* **may** *matter if we have* **weakly dominated strategies***.*

Definition 386 (Dominance Solvable)
A game is **dominance solvable** if the iterated deletion of strictly dominated strategies leads to a **unique** strategy profile.

Using common knowledge of rationality (CKR) we can predict the outcome of the game. Note that if there is a unique strategy surviving the IDSDS, it necessarily is a pure strategy.

We now give two famous examples of dominance solvable games.

Example 387
(Second Price Auction / Vickrey Auction). Suppose there is one seller of an indivisible good with n potential buyers with valuations $0 \leq v_1 \leq ... \leq v_n$. Since this is a complete information game, these valuations are common knowledge.

The strategy for buyers is to simultaneously submit bids $b_i \in [0, +\infty)$. This means we are now in an infinite game.

The highest bidder wins the good and pays the **second highest bid**. The other bids get nothing. We ignore ties in this setup. From this, the payoff for player i is given by:

$$u_i(b_i, b_{-i}) = \begin{cases} v_i - \max_{j\neq i} b_j & \text{if } b_i > \max_{j\neq i} b_j \\ 0 & \text{otherwise} \end{cases}$$

whereby $\max_{j\neq i} b_j$ denotes the highest bid amongst all other bids that are not ours, which is the highest bid in the case of $b_i > \max_{j\neq i} b_j$. We now make the following claim.

Claim 388. *Bidding our own valuation $b_i = v_i$ is a* **weakly dominant strategy***.*

Proof. To show this claim, we need to show that bidding our valuation $b_i = v_i$ does weakly better than any other bid we could have made b_i' against any other opponents' bids b_{-i}. For notational purposes, denote $r_i = \max_{j\neq i} b_j$

to be the highest bid that is **not** our bid.

Suppose that we are in the case that we bid higher than our valuation $b_i' > v_i$. There are then 3 cases to consider based on what the highest bid that was not our own bid did r_i.

1. If we did not make the highest bid, $r_i > b_i'$, then we did not win the auction and received a payoff of 0: $u_i(b_i', b_{-i}) = 0$. This is also the payoff if we bid our valuation, $u_i(v_i, b_{-i}) = 0$.

2. If we made the highest bid, $r_i < b_i'$, **and** the second-highest bid was below our valuation, $r_i \leq v_i$, then we receive a weakly positive payoff: $u_i(b_i', b_{-i}) = v_i - r_i \geq 0$. This payoff is also equal to the payoff we would receive if we bid our valuation, $u_i(v_i, b_{-i})$.

3. If we made the highest bid, $r_i < b_i'$, but the second-highest bid was also above our valuation, $r_i > v_i$, then we receive a strictly negative payoff because we pay more than our valuation: $u_i(b_i', b_{-i}) = v_i - r_i < 0$. However, if we bid our valuation, $b_i = v_i$, then we would not win the auction and would receive a payoff of 0, $u_i(v_i, b_{-i}) = 0$, which is strictly better than bidding above our valuation. Therefore, bidding $b_i = v_i$ does better in this case.

We can conclude that bidding our valuation $b_i = v_i$ weakly dominates any other bid greater than our valuation $b_i' > v_i$.

We can repeat the argument to show that bidding our valuation v_i does better than bidding below our valuation, $b_i' < v_i$. The only additional case is where the second-highest bid r_i is above our bid b_i' but below our valuation v_i. We would then be better off bidding our valuation and receiving $v_i - r_i \geq 0$, rather than losing the auction and receiving 0.

We can therefore conclude that $b_i = v_i$ is a **weakly dominant strategy**. Thus, truthful bidding is a weakly dominant strategy for every bidder. □

Example 389

(Perfect Competition and Monopoly Pricing).

Recall that the firm's maximisation problem was to choose the quantity to maximise profits:

$$\max_q \pi(q) = \max_q \Big\{ p(q)q - c(q) \Big\} \tag{162}$$

where $p(q)$ is the inverse demand function. We further assume that $p(\cdot)$ and $c(\cdot)$ are twice differentiable and that $p(0) > c'(0)$.

The first-order condition (differentiating with respect to q) is:

$$\underbrace{p'(q)q + p(q)}_{\text{Marginal Revenue}} = \underbrace{c'(q)}_{\text{Marginal Cost}} \tag{163}$$

We now make the assumptions of **linear demands**

$$p(q) = a - bq$$

and **constant returns to scale cost function**:

$$c(q) = cq$$

whereby $a > c \geq 0$ and $b > 0$.

Under constant returns, this implies that marginal cost is:

$$c'(q) = \frac{d}{dq}\big[c(q)\big] = \frac{d}{dq}\big[cq\big] = c.$$

Under **perfect competition**, we assume that the firm's production has no impact on the market price, so $p(q) = p$ and $p'(q) = 0$. This gives the familiar price-equals-marginal-cost relationship:

$$\underbrace{p^{\circ}}_{\text{Price}} = \underbrace{c}_{\text{Marginal Cost}}$$

We can then compute the quantity by plugging the price into the linear demand:

$$p^{\circ} = a - bq^{\circ} \quad \leftrightarrow \quad q^{\circ} = \frac{a - p^{\circ}}{b}$$

$$\Rightarrow q^{\circ} = \frac{a-c}{b}$$

where we used $p^{\circ} = c$. Therefore, under perfect competition, price and quantity are characterised by:

$$p^{\circ} = c \quad \text{and} \quad q^{\circ} = \frac{a-c}{b}.$$

We now want to examine the scenario under a **monopoly**. Using our marginal revenue equals marginal cost first order condition Eq. (163) and linear demand and constant returns:

$$(-b)(q^m) + (a - bq^m) = c$$

We can re-arrange to get the monopoly quantity:

$$q^m = \frac{a-c}{2b}$$

which is less than the output under perfect competition!

To retrieve the monopoly price, we plug the monopoly quantity into the linear demand:

$$p^m(q^m) = a - b\Big[\frac{a-c}{2b}\Big] = a - \frac{a-c}{2}$$

We therefore have that the monopoly price is:

$$p^m = \frac{a+c}{2}$$

which is higher than the price under perfect competition.

This principle holds more generally: monopoly output is lower than the perfectly competitive outcome, $q^m < q^{\circ}$, and monopoly price is higher, $p^m > p^{\circ}$. The associated loss of total surplus is the **deadweight loss**, illustrated in Fig. 65.

9.1.1 Cournot Duopoly with Linear Demand

Suppose we have two firms $i = 1, 2$ that sell an identical product. They simultaneously choose which **quantity** to produce. Both firms have constant returns to scale marginal cost per unit c. They both face linear inverse demand function:

$$p = a - bQ$$

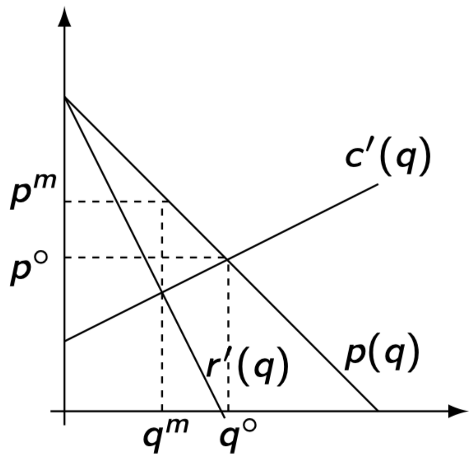

Figure 65: Perfect-competition and monopoly outcomes under linear demand and constant marginal cost.

where $Q = q_1 + q_2$ is the total quantity produced in the market. Hence, the amount that a firm produces will affect the market price. Furthermore, we assume that $a > c$ so that it is always worthwhile to produce.

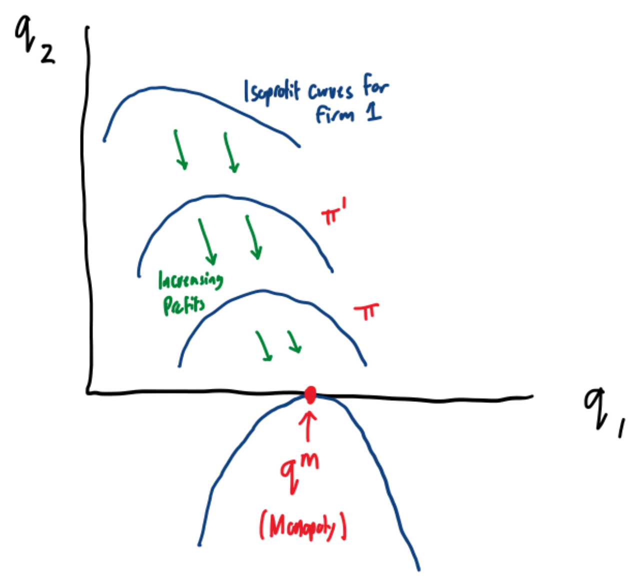

Figure 66: Isoprofit curves for firm 1. Lower curves represent higher profit, and the maximum attainable profit occurs at the monopoly quantity q^m when firm 2 produces zero.

Each firm's payoff is its revenue less the cost of producing q_i:

$$\pi_1 = q_1(p(Q) - c) = q_1[a - b(q_1 + q_2) - c]$$
$$\pi_2 = q_2(p(Q) - c) = q_2[a - b(q_1 + q_2) - c]$$

We now consider the **iso-profit curve for firm 1**, corresponding to the profit level π. Recall that the iso-profit curve is the locus of quantities produced that has the same profit level (similar to an indifference curve)

$$\{(q_1, q_2) : q_1[a - b(q_1 + q_2) - c] = \pi\}$$

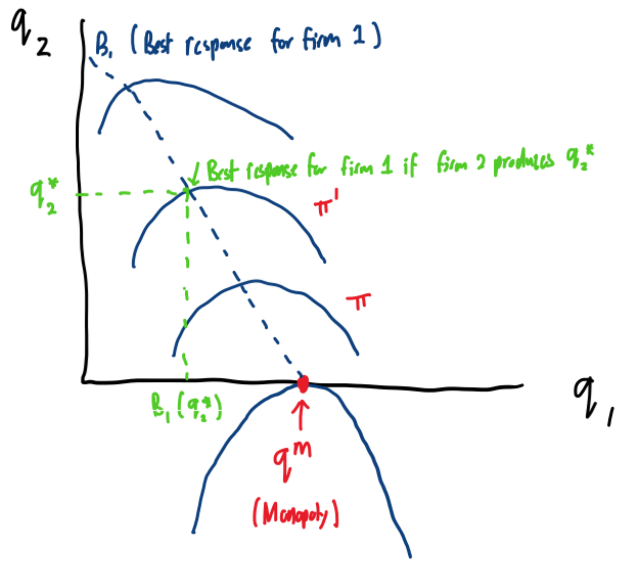

Figure 67: Best response for firm 1.

We want to plot this graphically so we can solve for q_2 as a function of q_1 whilst keeping profits fixed:

$$q_2^{\pi}(q_1) = \frac{a-c}{b} - q_1 - \frac{\pi}{bq_1} \tag{164}$$

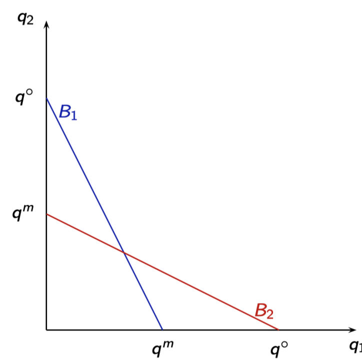

Figure 68: Best response for firm 1 and firm 2.

We now interpret and graph **firm 1's** isoprofit curves using Eq. (164). For low levels of q_1, q_2^{π} is increasing, and for high levels it is decreasing, so the isoprofit curves are **concave**. For a higher profit level $\pi' > \pi$, the corresponding curve lies lower: $q_2^{\pi'}(q_1) < q_2^{\pi}(q_1)$. Thus, higher-profit isoprofit curves lie closer to the x-axis. Firm 1 achieves its highest possible profit when $q_2 = 0$, in which case its optimal quantity is:

$$q_1 = \frac{a-c}{2b} = q^m$$

which is actually the monopoly quantity! We graphically show this in Fig. 66.

For any given level of output that firm 2 produces q_2, the peaks of the iso-profit curves correspond to the best response for what firm 1 should produce in response q_1 as this is the highest profit that firm 1 can achieve in response to firm 2's production. Then, drawing a curve through the peaks of the iso-profit traces out the best response function

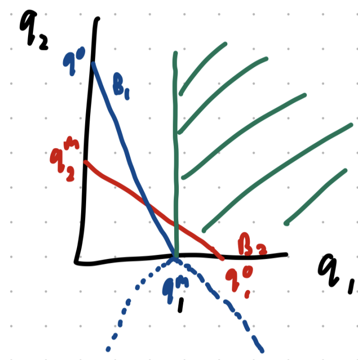

Figure 69: Any production q_1 in the green area can be shown to be dominated by firm 1's decision to produce the monopoly output q^m.

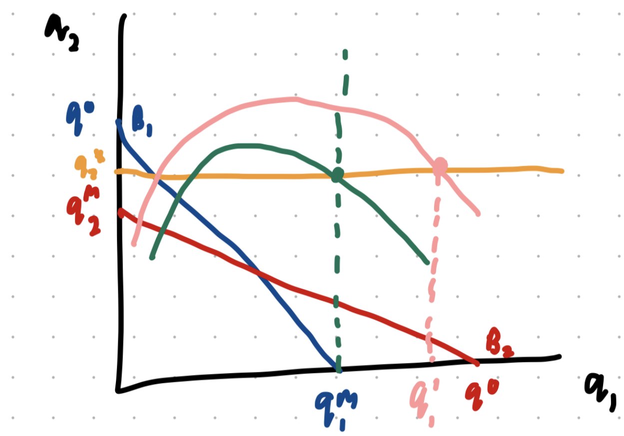

Figure 70: Firm 2 picks q_2^*. Firm one can either respond with playing q_1^1 or q_1^m. We can clearly see that playing the monopoly output is on a higher isoprofit line (higher profit means a curve closer to x-axis!) and therefore it is more desirable to pick that one.

for firm 1, which we denote by B_1.

Fig. 67 shows the best response for firm 1.

We can repeat this process for firm 2 in order to trace out its best-response function, denoted by B_2. This is shown in Fig. 68. When firm 1 produces the perfectly competitive quantity $q_1 = q^\circ$, firm 2's best response is to produce nothing, because any positive output would generate a loss.

We shall now apply iterated deletion of strictly dominated strategies in order to derive the unique strategies in this model.

Claim 390. *The unique strategies surviving IDSDS are $\{q^C, q^C\}$, where $q^C = (a - c)/(3b)$, corresponding to the intersection of the two best-response functions.*

Proof. First, we show that any quantity above the monopolistic quantity is strictly dominated. That is, for any $q_1 > q^m$, the strategy $q^m \succ q_1$ is preferred. In Fig. 69, we want to show that any point $q_1 > q^m$ is dominated by firm 1 playing q^m.

We use Fig. 70 to show that it is always better for firm 1 to produce the monopoly amount rather than any larger amount. Without loss of generality, suppose that firm 2 produces q_2^*. Firm 1's monopoly output q_1^m lies on a higher

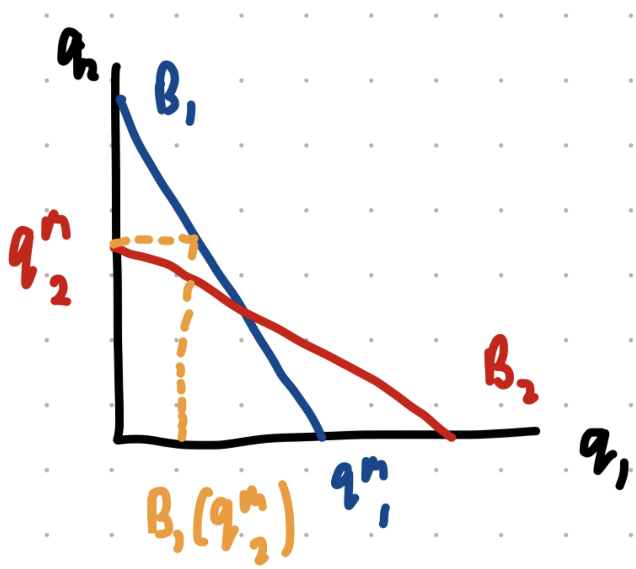

Figure 71: Firm 2 plays the monopoly output q_2^m and firm 1 can respond with the best response $B_1(q_2^m)$.

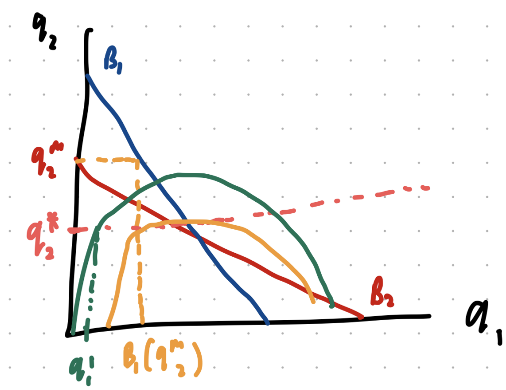

Figure 72: Firm 2 plays the output q_2^* and firm 1 can respond with q_1'. However, this is on a lower iso-profit line compared to if firm 1 instead responded with the best response $B_1(q_2^m)$.

isoprofit curve than any $q_1^1 > q_1^m$. This argument holds for every q_2 that firm 2 produces. Therefore, we can delete any strategy $q_1 > q_1^m$ for firm 1. By symmetry, firm 2 can delete any strategy $q_2 > q_2^m$.

We now show that, after eliminating strategies above q_1^m, firm 1 will never choose a quantity below the best response to firm 2's monopoly output: $q_1 < B_1(q_2^m)$. This is shown in Fig. 71.

We now show why firm 1 will never play q_1 below its best response. In Fig. 72, suppose that firm 2 picks q_2^* and firm 1 responds with $q_1' < B_1(q_2^m)$. The isoprofit curve associated with $B_1(q_2^m)$ (the yellow curve) gives higher profit than the curve associated with q_1' (the green curve). Thus q_1' is strictly dominated by $B_1(q_2^m)$. By symmetry, we can delete all $q_2' < B_2(q_1^m)$. Repeating these upper- and lower-bound deletions yields nested intervals converging to the unique fixed point $\{q^C, q^C\}$ shown in Fig. 73.

□

To summarise, we have shown that a game is **dominance-solvable** if the IDSDS predicts a unique outcome of the game. However, there are often many strategies that survive IDSDS. Therefore, if we want to have more narrow predictions, we need to make additional assumptions on top of common knowledge of rationality.

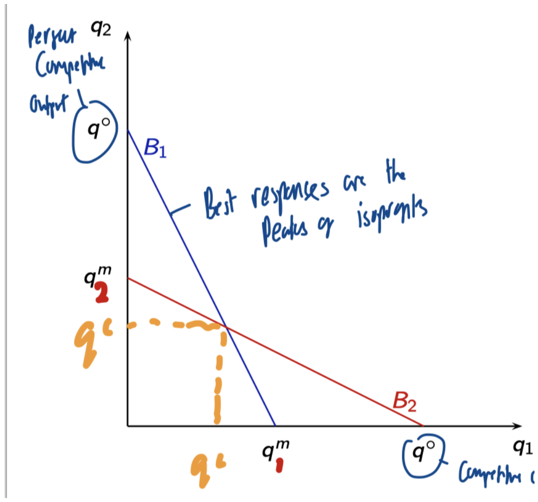

Figure 73: Surviving strategies are the intersection of the two best response functions.

9.2 Nash Equilibrium

We have seen that trying to find a unique strategy that survives iterated deletion of strictly dominated strategies does not always work. Therefore, in addition to common knowledge of rationality, we now introduce a stronger concept to narrow our predictions.

A Nash equilibrium is a profile of strategies such that each player's strategy is an optimal response to the other player's strategies.

Definition 391 (Pure Nash Equilibrium)
A strategy profile $s = (s_1, ..., s_n)$ is a **pure Nash Equilibrium** (NE) if for every $i = 1, ..., n$,

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \tag{165}$$

for all $s_i' \in S_i$.

Remark 392. *A Nash equilibrium is* **strict** *if each player has a unique best response to the other players' strategies. Furthermore, a Nash equilibrium refers to a* **profile of strategies**, *whereas a dominated strategy refers to a* **single strategy**.

The concept of NE requires players to be **rational**: they best respond to beliefs about how the others will play. Furthermore, it requires those beliefs to be **correct**. Common knowledge of rationality alone is not enough to justify NE; rational players may still disagree about how the game will be played.

Proposition 393 (Dominance Solvable Implies a Nash Equilibrium)
If a game is dominance-solvable, then the solution is the **unique Nash Equilibrium**.

If a strategy forms part of a NE, it cannot be strictly dominated, because it is a best response to the opponents' equilibrium strategies s_{-i}. Thus, strategies used in a NE necessarily survive IDSDS. In other words, the set of NE is a subset of the profiles that can be formed from strategies surviving IDSDS.

We can also define Nash equilibria using best-response correspondences.

Definition 394 (Best Response Correspondence)

Player i's **best response correspondence** $B_i : S_{-i} \rightrightarrows S_i$ assigns to each $s_{-i} \in S_{-i}$, the set

$$B_i(s_{-i}) = \{s_i \in S_i : u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \text{ for all } s_i' \in S_i\}.$$

Remark 395. *The notation $\rightrightarrows$ refers to a correspondence.*

Definition 396 (Pure Nash Equilibrium)

A strategy profile $s = (s_1, ..., s_N)$ is a Nash Equilibrium if and only if $s_i \in B_i(s_{-i})$ for all $i = 1, ..., N$.

Remark 397. *This says that the strategy profile is everyone's best response to everyone else.*

Therefore, using this definition, we can find Nash equilibria by checking whether each component of a strategy profile is a best response to the others.

Example 398

(Best Response Example). We will look for the Nash equilibrium in the game Fig. 74 using best responses. Suppose that player 2 plays L, then the best response for player 1 is $B_1(L) = B$ as it gives the highest payoff. We can therefore compute all the best responses for player 1 against player 2's strategies and vice versa:

$$B_1(L) = B; \quad B_1(C) = M; \quad B_1(R) = T$$
$$B_2(T) = L; \quad B_2(M) = C; \quad B_2(B) = R$$

	L	C	R
T	0, $\underline{7}$	2, 5	$\underline{7}$, 0
M	5, 2	$\underline{3}$, $\underline{3}$	5, 2
B	$\underline{7}$, 0	2, 5	0, $\underline{7}$

Figure 74: The NE of this game is (M,C).

The strategy profile (M, C) consists of mutual best responses, since $B_1(C) = M$ and $B_2(M) = C$. Hence, it is a Nash equilibrium.

We now give three examples of computing NEs in industrial organisation examples.

Example 399

(Bertrand Competition). Suppose we have two firms $i = 1, 2$ that sell an identical product. They simultaneously choose which **price** to charge. Both have constant marginal cost c. The firm with the lower price serves the entire market; if prices are equal, the firms split the market equally. We now compute each firm's best responses.

Claim 400. *There is a unique Nash equilibrium (p_1^*, p_2^*) in the Bertrand duopoly model. In this equilibrium, firms set prices to be equal to marginal cost: $p_1^* = p_2^* = c$.*

Proof. First, it is clear that $p_1 = p_2 = c$ is indeed a Nash equilibrium as it is the best response to each other. To see why, without loss of generality, suppose that firm 1 sets their price at $p_1 = c$. Then if firm 2 sets their price $p_2 < c$ then they will make negative profits. If firm 2 sets their price $p_2 > c$, then they will make no profits as they will not gain any market share. Therefore, their best response is to play $p_2 = c$. This holds by symmetry for firm 1's best response.

We now show that this is a **unique** Nash equilibrium. Suppose that (p_1, p_2) is another equilibrium. Without loss of generality, suppose that $p_1 \leq p_2$. We have three cases:

1. If $p_1 < c$, then firm 1 makes losses and has a profitable deviation to raise prices to c.
2. If $p_1 = c < p_2$, then firm one can raise their prices p_1 and increase profits whilst still retaining market share.
3. If $c < p_1 \leq p_2$ then firm 2 should undercut the price of firm 1 and gain full market share in order to make profits.

Therefore, we conclude that no other equilibrium could exist. □

The Bertrand model makes the (strange) claim that as long as there are two firm, then distortions due to market power disappear.

Example 401

(General Cournot Competition). Suppose that have n firms that sell an identical product. They all simultaneously choose which quantity to produce q_i. We have heterogeneity as each firm i can have constant marginal cost c_i. Suppose they all face an inverse demand function $p(Q)$ where Q is the total quantity in the market.

We make the assumptions that:

1. $p(\cdot)$ is decreasing;
2. $p(\cdot)$ is concave;
3. $p(0) > \min_i c_i$.

This last assumption says that there needs to be some production so that people are willing to pay for the first unit and that this price paid is higher than the marginal cost of the most efficient firm.

Now, the payoff for firm i is their total revenue less than their total cost:

$$\pi_i = q_i[p(Q) - c_i]. \tag{166}$$

We can then compute the first order condition by differentiating with respect to q_i and applying the product rule to Eq. (166). Doing so will give us an expression characterising the best response function for firm i: $q_i^* = B_i(q_{-i}^*)$. And so, the first order condition $\frac{\partial}{\partial q_i}\left[q_i p(Q) - q_i c_i\right]$ is equal to:

$$q_i^* p'(Q^*) + p(Q^*) = c_i \tag{167}$$

if $p(Q^*) > c_i$ and otherwise the best response for firm i is $q_i^* = 0$. The individual best responses for each firm $q_i^* = B_i(q_{-i}^*)$ satisfies the first order condition in Eq. (167) as this is a necessary condition for optimality. We can then sum over all these best responses q_i^* and divide by n to get:

$$\frac{Q^*}{n} p'(Q^*) + p(Q^*) = \frac{\sum_i c_i}{n}$$

Therefore, the Nash equilibrium $(q_1^*, q_2^*, ..., q_n^*)$ satisfies the condition:

$$\frac{Q^*}{n}p'(Q^*) + p(Q^*) = \frac{\sum_i c_i}{n}$$

which relates the total quantity in industry, denoted by Q, and its relationship to the average marginal cost $\frac{\sum_i c_i}{n}$. Therefore, we have an equation with one single variable, which is total industry output Q.

However, we were interested in deriving the **individual quantities** produced q_i^*. In order to do so, we need to impose further assumptions on individual marginal costs c_i as this will allow us to do so. We shall impose that marginal costs are common to all firms $c_i = c$. This gives us a result for the price and quantity.

Claim 402. *In any Nash equilibrium of the Cournot model with common unit cost c, the market price $p^* = p(Q^*)$ is such that:*

$$c = p^\circ < p^* < p^m$$

whereby p° is the perfect competition price, p^* is the General Cournot price, and p^m is the monopoly price. Furthermore, the quantity produced by the General Cournot model Q^* is such that:

$$q^m < Q^* < Q^\circ$$

where q^m is the amount produced under a monopoly setting and Q° is the amount produced under perfect competition.

Proof. First, recall that the FOC that characterised the NE was:

$$\frac{Q^*}{n}p'(Q^*) + p(Q^*) = \frac{\sum_i c}{n} \tag{168}$$

where we replaced $c_i = c$ by assumption of identical marginal cost for all firms.

First, we show that price is higher than the perfectly competitive price, $p^* > p^\circ$, and total output is lower than the perfectly competitive output, $Q^* < q^\circ$. The inverse demand function $p(\cdot)$ is decreasing, so $p'(Q^*) < 0$ in Eq. (168). For the FOC in Eq. (168) to hold, we therefore require:

$$p(Q^*) > \frac{\sum_i c}{n} = \frac{nc}{n} = c$$

and hence we have shown that the price in the General Cournot model $p^* = p(Q^*) > c = p^\circ$ which was the price under perfect competition. Then, as $p^* > p^\circ$, by the decreasing nature of the inverse demand function, this implies that:

$$Q^\circ > Q^*$$

so the amount produced in the general Cournot model will be less than the amount produced in the perfectly competitive model.

We now show that the amount produced in the Cournot model is more than the monopoly setting $Q^* > q^m$ and the price is lower than the monopoly setting $p^* < p^m$. We will show this via proof by contradiction.

Suppose that the amount produced in the General Cournot model is actually lower than the monopoly amount $Q^* < q^m$. Now, fix the strategies of all other firms $-i$ and examine firm i. Firm i should be playing their best response under this setting. However, we claim that firm i has a profitable deviation to raise its own quantity q_i such that total quantity produced in the Cournot model equals to the monopoly output $Q = q^m$. This causes the **joint profits** of all firms to increase since the monopoly quantity q^m is the quantity that maximises profits for a single firm and hence maximises profits for the whole industry. As quantity Q increases, this means the price falls. As the other firms $-i$ did not change their quantities, they are worse off. Hence, all the additional profits accrue to firm i who benefits from this profitable deviation. But this is a contradiction that firm i is currently playing their best response! Therefore, Q^* cannot be less than q^m.

This then implies that the quantity produced under the general Cournot model is at least as much as the monopoly output: $Q^* \geq q^M$. However, if $Q^* = q^m$, this cannot be the case as the output under q^m was only for one firm, $n = 1$, but the first order condition Eq. (168) cannot be satisfied for both $n = 1$ and $n > 1$ if $Q^* = q^m$. This implies that the amount produced in general Cournot model is greater than monopoly output $Q^* > q^m$ and by the decreasing inverse demand function, the price $p^* < p^m$. □

What we see is that the quantity that is produced in this model is higher than that in the monopoly setting but lower than that in the competitive market setting. Conversely, the price in this model will be lower than that in the monopoly setting but higher than that in the perfectly competitive market setting.

We can therefore see that the Cournot outcome is an outcome that is in between perfect competition and the monopolistic outcome.

One final thing to note is that if we look at the first order condition:

$$\frac{Q^*}{n} p'(Q^*) + p(Q^*) = c$$

as we let the number of firms $n \to \infty$, we end up with:

$$p(Q^*) = c$$

whereby price is equal to the marginal cost. This is actually the perfectly competitive outcome!

Example 403

(Hotelling Model). We have two firms located at either end of an unit interval $[0, 1]$. They produce an identical product with constant marginal cost c. They also simultaneously choose their price. There is an unit mass of consumers who are uniformly distributed on the interval. Each consumer wants at most one unit of the good and derives benefit v from it. The cost of buying the product from firm i is

$$p_i + td$$

where t is a transport cost and d is the distance of the consumer to firm i and p_i is the price of the good sold by firm i. We can see that consumers are trading off the price of firm i and the distance to firm i. This is illustrated in Fig. 75.

Figure 75: Two firms on either of the unit interval [0,1].

We now derive demand for the two firms' products. A consumer located at $z \in [0, 1]$ buys from firm 1 (at endpoint 0) if:

$$\underbrace{p_1 + tz}_{\text{Price of firm 1 plus distance to firm 1}} < \underbrace{p_2 + t(1-z)}_{\text{Price of firm 2 plus distance to firm 2}}$$

Equivalently, the consumer buys from firm 1 when

$$z < \frac{1}{2} + \frac{p_2 - p_1}{2t}.$$

In particular, firm 1 serves no consumers if $p_1 \geq p_2 + t$ and serves the whole market if $p_1 \leq p_2 - t$.

If the consumer z is indifferent between firm 1 and 2, then we have that:

$$\begin{aligned} p_1 + tz &= p_2 + t(1-z) \\ \leftrightarrow 2tz &= p_2 - p_1 + t \\ \Rightarrow \hat{z} &= \frac{1}{2} + \frac{p_2 - p_1}{2t} \end{aligned}$$

Here, $\hat{z}$ is the notation relating the distance to firm 1 and the relative prices between firm 1 and 2.

Thus, consumer z does not purchase from firm 1 if:

$$z > \hat{z}.$$

Therefore, given prices p_1, p_2, the demand for firm 1 is given by:

$$x_1(p_1, p_2) = \begin{cases} 0 & \text{if } \hat{z} \leq 0 \\ \frac{1}{2} + \frac{p_2 - p_1}{2t} & \text{if } \hat{z} \in (0, 1) \\ 1 & \text{if } \hat{z} \geq 1 \end{cases}$$

whereby if $\hat{z} < 0$, then firm 1's price is too high relative to firm 2 and does not cover the transport cost t.

From this, the profits for firm i is given by their total revenue less their cost:

$$\pi_i = x_i(p_i, p_j)(p_i - c).$$

We then plug in our expression for $x_i(p_i, p_j)$ into our profit in order to get our best response functions for what price should each firm set:

$$B_i(p_j) = \begin{cases} c & \text{if } p_j < c - t \\ (t + c + p_j)/2 & \text{if } c - t < p_j \leq 3t + c \\ p_j - t & \text{if } 3t + c < p_j \end{cases} \tag{169}$$

Intuitively, firm i should set their price to marginal cost $p_i = c$ if the other firm's price is too low. If the other firm's price is sufficiently high, we can set our price high as well $p_i = p_j - t$.

We graphically plot this in Fig. 76 whereby the Nash equilibrium is where the two best response functions intersect each other. We see that the best response functions are upward sloping function as if one firm raises their price, then the other firm can as well. The interior equilibrium is given by:

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

and under symmetry $p_i = p_j = p^*$, we can conclude that the Nash equilibrium is:

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

This can be interpreted as a Bertrand model whereby we have differentiated firms. Here, the firms have market power so that price is above marginal cost and hence differs from the perfectly competitive setting.

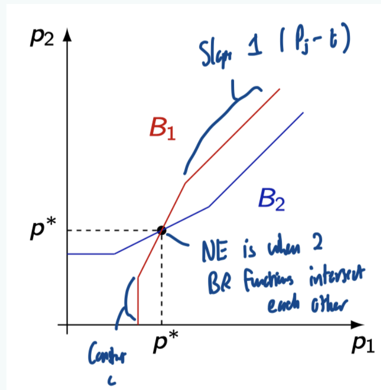

Figure 76: Best response for the Hotelling model.

One final thing we note is the relationship between dominant strategy equilibrium and Nash equilibrium.

Theorem 404 (Dominant-Strategy Equilibria are NE)
Every dominant-strategy equilibrium is a Nash equilibrium.

Remark 405. *Note that the converse does not hold. Not every Nash equilibrium is a dominant strategy equilibrium.*

9.2.1 Mixed Nash Equilibrium

We now extend our definition of Nash equilibrium to mixed strategies.

Definition 406 (Mixed Nash Equilibrium)
A mixed strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$ is a **mixed Nash Equilibrium** (NE) if for every $i = 1, ..., n$,

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(s_i, \sigma_{-i}) \tag{170}$$

for all $s_i \in S_i$.

Remark 407. *Note that it is enough to check that σ_i does better than any other pure strategy $s_i \in S_i$ rather than any other mixed strategy $\sigma'_i \in \Delta(S_i)$. Indeed, if that is the case then $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i})$ for all $\sigma'_i \in \Delta(S_i)$:*

$$\begin{aligned}
u_i(\sigma'_i, \sigma_{-i}) &= \textstyle\sum_{s_i \in S_i} \sigma'_i(s_i) u_i(s_i, \sigma_{-i}) \\
&\leq \textstyle\sum_{s_i \in S_i} \sigma'_i(s_i) u_i(\sigma_i, \sigma_{-i}) \\
&= u_i(\sigma_i, \sigma_{-i}) \textstyle\sum_{s_i \in S_i} \sigma'_i(s_i) \\
&= u_i(\sigma_i, \sigma_{-i})
\end{aligned}$$

We can also define mixed-strategy Nash equilibria in terms of best-response correspondences.

Definition 408 (Nash Equilibrium in Mixed Strategies)
Define player i's **best response correspondence** as the set

$$B_i(\sigma_{-i}) = \{\sigma_i \in \Delta(S_i) : u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i}) \text{ for all } \sigma_i' \in \Delta(S_i)\}.$$

Then, a strategy profile $\sigma = (\sigma_1, ..., \sigma_N)$ is a **Nash Equilibrium** (NE) if $\sigma_i \in B_i(\sigma_{-i})$ for all i.

Example 409
(Matching Pennies).

The matching pennies game can be seen in Fig. 77. There is no pure Nash equilibrium. However, there is a mixed Nash equilibrium.

		q	1-q
		H	T
p	H	-1, 1	1, -1
1-p	T	1, -1	-1, 1

Figure 77: Matching pennies game.

The mixed strategies are the probability measures on the pure strategies:

$$\sigma_1(H) = p; \quad \sigma_1(T) = 1 - p$$

$$\sigma_2(H) = q; \quad \sigma_2(T) = 1 - q$$

The expected payoff for player 1 by playing the mixed strategy σ_1 is given by:

$$\begin{aligned} u_1(\sigma_1, \sigma_2) &= (p)(q)u_1(H, H) + (p)(1-q)u_1(H, T) + (1-p)(q)u_1(T, H) + (1-p)(1-q)u_1(T, T) \\ &= (p)(q)(-1) + (p)(1-q)(1) + (1-p)(q)(1) + (1-p)(1-q)(-1) \\ &= -pq + p(1-q) + q(1-p) - (1-p)(1-q) \\ &= p(2-4q) - (1-2q) \end{aligned}$$

Therefore, the expected payoff for player 1 is:

$$u_1(\sigma_1, \sigma_2) = p(2-4q) - (1-2q)$$

which is a function of the probability that player 2 puts on H, which is denoted by q. Player 1 can choose a randomisation probability p such that player 1 has a best response to player 2's choice of q in order to maximise the payoff $u_1(\sigma_1, \sigma_2)$.

First, let q be $q^+ > \frac{1}{2}$, then the expected payoff for player 1 is:

$$u_1(\sigma_1, q^+) = p(2-4q^+) - (1-2q^+)$$

Since $2 - 4q^+ < 0$, the expected payoff is decreasing in p. Therefore, player 1 should pick $p = 0$ as their best response to player 2 choosing $q > \frac{1}{2}$.

Now, let q be $q^- < \frac{1}{2}$, then the expected payoff for player 1 is:

$$u_1(\sigma_1, q^-) = p(2 - 4q^-) - (1 - 2q^-)$$

Since $2 - 4q^- > 0$, the expected payoff is increasing in p. Therefore, player 1 should pick $p = 1$ as their best response to player 2 choosing $q < \frac{1}{2}$.

Finally, look at when $q = \frac{1}{2}$. Then, the expected payoff for player 1 is:

$$u_1(\sigma_1, q) = p(2 - 4q) - (1 - 2q) = p(2 - 4\frac{1}{2}) - (1 - 2\frac{1}{2}) = 0.$$

Therefore, no matter what mixing probability player 1 chooses for p, they will get the same expected payoff. Therefore, player 1 is indifferent between playing H or T and therefore he can mix between the two actions with any arbitrary probability $p \in [0, 1]$.

The best response for player 1 is therefore:

$$B_1(q) = \begin{cases} 1 & \text{if } q < \frac{1}{2} \\ [0,1] & \text{if } q = \frac{1}{2} \\ 0 & \text{if } q > \frac{1}{2} \end{cases}$$

We can do something similar for the expected payoff for player 2 $u_2(\sigma_1, \sigma_2)$ to get:

$$B_2(p) = \begin{cases} 0 & \text{if } p < \frac{1}{2} \\ [0,1] & \text{if } p = \frac{1}{2} \\ 1 & \text{if } p > \frac{1}{2} \end{cases}$$

We can plot these best responses and find their intersection in Fig. 78. The two best-response correspondences intersect at $p = \frac{1}{2}$ and $q = \frac{1}{2}$. Therefore, the mixed Nash equilibrium is $(p = \frac{1}{2}, q = \frac{1}{2})$.

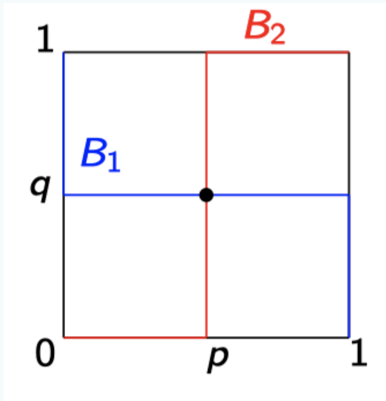

Figure 78: Best response function in matching pennies game.

When $q = \frac{1}{2}$, player 1 is indifferent between playing H and T. Likewise, when $p = \frac{1}{2}$, player 2 is indifferent between H and T. This is an important result that we explain next.

The significance of the previous example is that when a player is choosing a mixed strategy as their best response, they are in fact indifferent between the actions that they are mixing over. If a player chooses a pure action for their strategy, it is because that action leads them to a higher payoff. This leads us to the observation that if a player has a best response that involves randomising between two or more actions, they **must** be indifferent between those two actions (which means that the two actions give them the same payoff). If this was not the case, that is, another action gave a higher payoff, then the player should've just picked that action in the first place!

From this observation, the following result is helpful to find Nash Equilibria in mixed strategies.

Theorem 410 (Indifference Criterion for Mixed Strategy Nash Equilibria)
A mixed strategy $\sigma = (\sigma_1, ..., \sigma_n)$ is a NE if and only if:

(i) The player is indifferent between the pure strategies in the support of the mixed strategy:

$$u_i(s_i, \sigma_{-i}) = u_i(s_i', \sigma_{-i}) \tag{171}$$

for all strategies s_i, s_i' such that they are chosen with positive probability in the mixed strategy, i.e $\sigma_i(s_i) > 0$, $\sigma_i(s_i') > 0$.

(ii) The payoff from playing pure strategies from the strategies in the support of the mixed strategy is higher than payoffs from strategies not in the support of the mixed strategy:

$$u_i(s_i, \sigma_{-i}) \geq u_i(s_i', \sigma_{-i}) \tag{172}$$

for all pure strategies in the support s_i such that $\sigma_i(s_i) > 0$ and for all pure strategies not in the support s_i', such that $\sigma_i(s_i') = 0$.

Proof. $\Rightarrow$) Suppose that $\sigma = (\sigma_1, ..., \sigma_n)$ is a NE in mixed strategies, then $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i})$ for all $\sigma_i' \in \Delta(S_i)$. We proceed in two steps,

1. We first show that $u_i(\sigma_i, \sigma_{-i}) = u_i(s_i, \sigma_{-i})$ for all $s_i \in S_i$ such that $\sigma_i(s_i) > 0$:
   Suppose that this is not the case, i.e., there exists $s_i \in S_i$ with $\sigma_i(s_i) > 0$ such that $u_i(\sigma_i, \sigma_{-i}) > u_i(s_i, \sigma_{-i})$. Then there must exist $s_i' \in S_i$ with $\sigma_i(s_i') > 0$ such that $u_i(s_i', \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i}) > u_i(s_i, \sigma_{-i})$, since $u_i(\sigma_i, \sigma_{-i})$ is a convex combination of the payoffs $u_i(s_i, \sigma_{-i})$ for strategies with $\sigma_i(s_i) > 0$. But then s_i' is a better response to σ_{-i} than σ_i, contradicting the assumption that σ is a NE. Therefore $u_i(s_i, \sigma_{-i}) = u_i(\sigma_i, \sigma_{-i}) = u_i(s_i', \sigma_{-i})$ for all s_i, s_i' with $\sigma_i(s_i) > 0, \sigma_i(s_i') > 0$, which proves (i).

2. Given that σ_i is a best response to σ_{-i}, $u_i(s_i, \sigma_{-i}) = u_i(\sigma_i, \sigma_{-i}) \geq u_i(s_i', \sigma_{-i})$ for all $\sigma_i(s_i) > 0, \sigma_i(s_i') = 0$ which proves (ii).

$\Leftarrow$) Consider a mixed strategy profile σ satisfying (i) and (ii). Then, for all i, and s_i such that $\sigma_i(s_i) > 0$ we have that for all $\sigma_i' \in \Delta(S_i)$:

$$u_i(\sigma_i, \sigma_{-i}) = u_i(s_i, \sigma_{-i}) \geq \sum_{s_i'} \sigma_i'(s_i') u_i(s_i', \sigma_{-i}) = u_i(\sigma_i', \sigma_{-i})$$

and hence, σ_i is a best response to σ_{-i} and σ is a mixed Nash Equilibrium. ☐

Remark 411. *This says that the utility of strategies in our mixture should be greater than the strategies not in our mixture.*

We can apply this indifference criterion to find Nash equilibria.

Example 412
(Battle of the sexes). The battle of the sexes is seen in Fig. 79.

		(q) C	$(1-q)$ P
(p)	C	2, 1	0, 0
$(1-p)$	P	0, 0	1, 2

Figure 79: Battle of sexes.

By the indifference principle, if player 1 were to randomise between playing C or playing P, they need to be indifferent between playing C or playing P. Therefore, fixing the opponent's probability q, player 1 is indifferent between C and P when:

$$\underbrace{(q)u_1(C,C) + (1-q)u_1(C,P)}_{\text{Player 1 plays C}} = \underbrace{(q)u_1(P,C) + (1-q)u_1(P,P)}_{\text{Player 1 plays P}}$$
$$(q)(2) + (1-q)(0) = (q)(0) + (1-q)(1)$$
$$2q = (1-q)$$

which gives us:

$$q = \frac{1}{3}$$

as being the indifference condition for player 1 to be indifferent between playing C or P.

We can repeat a similar process for when player 2 is indifferent between playing C and P when:

$$\underbrace{(p)u_2(C,C) + (1-p)u_2(P,C)}_{\text{Player 2 plays C}} = \underbrace{(p)u_2(C,P) + (1-p)u_2(P,P)}_{\text{Player 2 plays P}}$$
$$(p)(1) + (1-p)(0) = (p)(0) + (1-p)(2)$$
$$p = 2(1-p)$$

which gives us:

$$p = \frac{2}{3}$$

as being the indifference condition for player 2 to be indifferent between playing C or P.

Therefore, we can conclude that the mixed Nash equilibrium is given by:

$$\sigma = \left(\sigma_1 = (\frac{2}{3}, \frac{1}{3}); \sigma_2 = (\frac{1}{3}, \frac{2}{3})\right)$$

as mixing strategies for player 1 and player 2 respectively.

9.3 Existence of Nash Equilibria

To show the existence of a Nash equilibrium, we use Kakutani's fixed-point theorem, the correspondence-valued analogue of Brouwer's fixed-point theorem.

Theorem 413 (Kakutani's Fixed Point Theorem)
Let X be a compact and convex subset of $\mathbb{R}^n$, and let $f : X \rightrightarrows X$ be a correspondence such that

(i) for all $x \in X$, the set $f(x)$ is nonempty and convex, and

(ii) the graph of f is closed: if $y_n \in f(x_n)$, $x_n \longrightarrow x$, and $y_n \longrightarrow y$, then $y \in f(x)$.

Then there exists $x^* \in X$ such that $x^* \in f(x^*)$.

Consider the game $\Gamma_N = [N, \{S_i\}, \{u_i(\cdot)\}]$ with only pure strategies.

We define player i's **best-response correspondence** as $B_i : S_{-i} \rightrightarrows S_i$, which assigns to each $s_{-i} \in S_{-i}$ the set

$$B_i(s_{-i}) = \{s_i \in S_i : u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \text{ for all } s_i' \in S_i\}$$

.

We define the Best Response Correspondence $B : S \longrightarrow S$ as $B(s) = (B_i(s_{-i}))_{i \in N}$, where $B_i(s_{-i})$ is the best response correspondence of player i.
Note that s is a Nash Equilibrium if and only if $s \in B(s)$. In other words s is a fixed point of B. We can now apply Kakutani's fixed point theorem.

Theorem 414 (Existence of Pure Strategy Nash Equilibrium in Infinite Games)
The game $\Gamma_N = [N, \{S_i\}, \{u_i(\cdot)\}]$ has a Nash Equilibrium if for all $i \in \{1, ..., N\}$,

(i) S_i is nonempty, convex and compact,

(ii) $u_i(s)$ is continuous in s and quasiconcave in s_i

Proof. Consider the best-response correspondence $B : S \rightrightarrows S$. For every $i \in \{1, \ldots, N\}$, $B_i(s_{-i})$ is nonempty because u_i is continuous and S_i is compact (Weierstrass theorem). It is convex because u_i is quasiconcave in s_i.
Moreover, B has a closed graph since u_i is continuous for all i. Therefore by Kakutani's fixed point theorem, B has a fixed point. □

Remark 415. *This result we have shown applies to games with an infinite number of strategies. This is because we assumed that the strategy set S_i to be convex, which then implies that there is a infinite number of strategies to choose from.*

We can use the previous result to show that any **finite** game has at least a mixed Nash Equilibrium.

Theorem 416 (Existence of Mixed-Strategy Nash Equilibrium in Finite Games)
Every finite game has a **mixed strategy** NE.

Proof. Consider a game $\Gamma_N = [N, \{S_i\}, \{u_i(\cdot)\}]$ with S_i finite for all i. Consider the mixed extension: $\tilde{\Gamma}_N = [N, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$. Note that if the numbers of strategies in S_i is m_i, then $\Delta(S_i) \equiv \{\sigma_i = (p_1, ..., p_{m_i})$ such that $p_k \geq 0$, and $\sum_{k=1}^{m_i} p_k = 1\}$, hence this set is nonempty, convex and compact. Moreover, $u_i(\sigma)$ is linear in σ, and hence continuous in σ and quasi-concave in σ_i. Therefore, by the previous theorem, the mixed extension of $\tilde{\Gamma}_N$ has a Nash Equilibrium. □

The two existence theorems give sufficient but not necessary conditions for Nash equilibria. The following example does not satisfy the conditions of either theorem but still has a Nash equilibrium.

Example 417

(Nash Demand Game). Suppose that two individuals are arguing over the division of a pot of money. They both simultaneously make demands for what fraction of the pot they want. They receive their demands if and only if both demand can jointly met from the money pot.

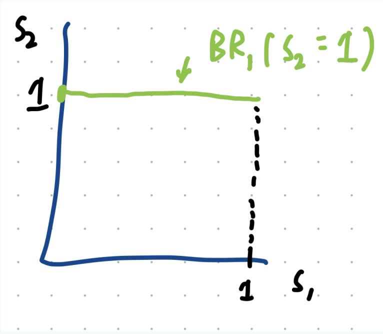

Figure 80: Player i is indifferent among all strategies $s_i \in [0, 1]$ if the opponent chooses $s_j = 1$.

The strategy sets are $S_i = [0, 1]$ as they can demand whatever fraction ranging from nothing (0) to everything (1). The strategy set S_i is nonempty, convex, and compact.

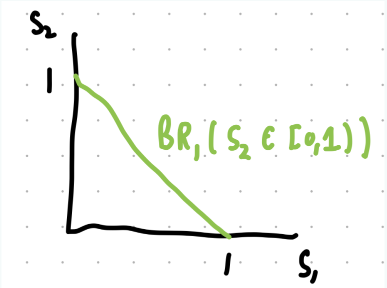

Figure 81: Player i best response is $1 - s_j$.

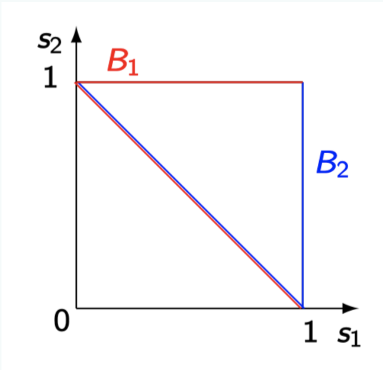

Figure 82: Player i and j best responses.

We assume that if player i receives a share s_i, they get a payoff $v_i(s_i)$ which is continuous and increasing in s_i.

Therefore, the payoff for player i is:

$$u_i(s_i, s_j) = \begin{cases} v_i(s_i) \text{ if } s_i + s_j \leq 1 \\ v_i(0) \text{ if } s_i + s_j > 1 \end{cases}$$

An issue is that even though $v_i(\cdot)$ is continuous, the utility payoff u_i is **not continuous** and hence we cannot apply our previous theorem. However, we can still find a Nash equilibrium.

If the opponent $s_j = 1$ demands everything, then player i gets 0 regardless and hence is indifferent to any strategy $s_i \in [0, 1]$ as illustrated in Fig. 80.

If player j demands anything $s_j < 1$ then the best response for player i is to demand $1 - s_j$ as seen in Fig. 81.

By symmetry for player 2, they will have the same results for their best response function. Therefore, we can plot both players' best responses in Fig. 82. We can see that for any $s \in [0, 1]$, the best response curves of the players intersect. Therefore, for any $s \in [0, 1]$, the strategy profile $(s, 1 - s)$ is a Nash equilibrium.

To summarise, we have seen several techniques for finding Nash equilibria and have shown the existence of mixed-strategy Nash equilibrium in finite games. Still, it can be difficult to justify Nash equilibrium as the appropriate solution concept for a particular game.

9.4 Sequential Rationality and Zermelo's Theorem

We now analyse games with a dynamic component, i.e., games in which there is **sequentiality**. That is, those players who are called to move later on in the game know something about what has happened before. This results in different information sets as compared to simultaneous games as players can now observe and gain information.

We can represent these sequential games in extensive form. However, once a strategy is defined as a complete contingent plan of action, an extensive-form game also has a corresponding **normal-form representation**. We can therefore apply concepts such as Nash equilibrium and dominance to its normal form. The Nash equilibria of the two representations coincide because they describe the same strategies and payoffs.

However, converting an extensive-form game to normal form suppresses information about the game's sequential structure. As a result, Nash equilibria may be supported by **non-credible threats**. We therefore need to refine Nash equilibrium for sequential games. We work with extensive-form representations and, when convenient, behavioural strategies.

Consider first games of perfect information, i.e., all information sets are singletons. In those games, players know perfectly what has happened before when they are called to make a decision.

Example 418

(Predation Game). Firm E (Entrant) is considering whether to enter a market with an incumbent firm, Firm I. If it enters, the incumbent must decide whether to Accommodate or Fight. We represent the game in extensive and normal form in Fig. 83. Firm I has only one information set, so its strategies are simply Fight and Accommodate.

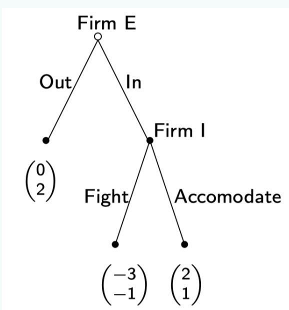

(a) Extensive Form Representation

	Fight	*Accom.*
Out	$\underline{0}$, $\underline{2}$	0, $\underline{2}$
In	-3, -1	$\underline{2}$, $\underline{1}$

(b) Normal Form Representation

Figure 83: The predation game in extensive and normal form.

Looking at the Normal Form representation, we can derive the two (pure) Nash equilibria as being:

$$\{(\text{Out, Fight}); (\text{In, Accommodate})\}.$$

We now assess these Nash equilibria in the extensive form. The equilibrium (Out, Fight) is problematic: firm E stays out because firm I threatens to Fight following entry, which would give the entrant a payoff of -3. Yet if entry occurs and firm I is called upon to act, firm I prefers Accommodate to Fight, receiving 1 rather than -1. Thus, (Out, Fight) relies on a non-credible threat and is not **sequentially rational**.

Example 419

(Sequential Game with 2 Information Sets). Fig. 84 gives the extensive- and normal-form representations of a sequential game in which player 2 (P2) has **two** information sets. Each of P2's pure strategies therefore specifies two actions, one for information set a and one for information set b.

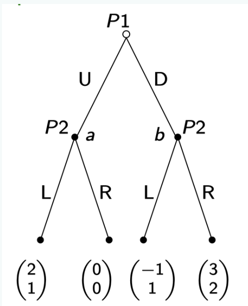

(a) Extensive Form Representation

	LL	*LR*	*RL*	*RR*
U	$\underline{2}$, $\underline{1}$	2, $\underline{1}$	$\underline{0}$, 0	0, 0
D	-1, 1	$\underline{3}$, $\underline{2}$	-1, 1	$\underline{3}$, $\underline{2}$

(b) Normal Form Representation

Figure 84: A sequential game with two information sets, in extensive and normal form.

Looking at the normal-form representation, the Nash equilibria are:

$$\{(\text{U, LL}); (\text{D, LR}), (\text{D, RR})\}.$$

To determine whether these equilibria are sensible, consider the game's **extensive-form** representation. At information set a, P2 prefers L to R (a payoff of 1 rather than 0), so P2 will never choose R there. We can therefore rule out (D, RR), because the first R in P2's strategy prescribes R at information set a.

Likewise, at information set b, we can see that P2 will never choose L as they can earn a higher payoff from playing R (2 versus 1). Therefore, the Nash equilibrium (U, LL) can be ruled out as the strategy for P2, LL, says that P2 will play L if they are at information set b, which we have just shown not to be the case.

Therefore, the Nash equilibria (U, LL) and (D, RR) are **not sequentially rational**.

These examples show that Nash equilibrium can be too weak for sequential games because it need not prescribe optimal play at information sets off the equilibrium path. In particular, an equilibrium may be supported by **non-credible threats**.

These non-credible threats are unappealing because if player i did deviate from the equilibrium path, it will not be optimal for their opponents to follow through on their threats. In order to avoid those non-credible threats, we will require the strategies to be **sequentially rational**, that is, at any information set on or off the equilibrium path, the strategies should specify the optimal continuation play for every player, considering the strategy profile of the other players.

We now seek strategies that satisfy sequential rationality. Zermelo's backward-induction algorithm delivers such strategies at every node of a finite perfect-information game.

Theorem 420 (Backward Induction Algorithm to Find Sequentially Rational Strategies)
Suppose the game is **finite** and has **perfect information**. A sequentially rational strategy can be found as follows:

1 Start at the final decision nodes (those nodes that are the predecessors of terminal nodes) and choose the optimal action for the player associated to those decision nodes.

2 Reduce the game by replacing the final decision nodes by the terminal nodes reached with the optimal actions in the previous stage.

3 Re-apply step 1 to the new reduced game.

Note that this algorithm is well-defined: given that the game is finite and information sets are singletons, we can always find the final decision nodes. Moreover, since at each decision node there is a finite number of actions available, there will always be an optimal action. Finally, given that the game is finite, the algorithm will end in finite time.

We now apply the backward induction to our previous examples.

Example 421

(Predation Game Continued). We apply backward induction to the extensive-form representation from Example 418, as shown in Fig. 85. At firm I's final decision node, Accommodate is preferred to Fight, so we replace that node with the payoff from Accommodate. Moving back to firm E's decision, firm E then prefers In to Out. Thus, (In, Accommodate) is both a Nash equilibrium of the normal form and the sequentially rational outcome. Backward induction rules out (Out, Fight).

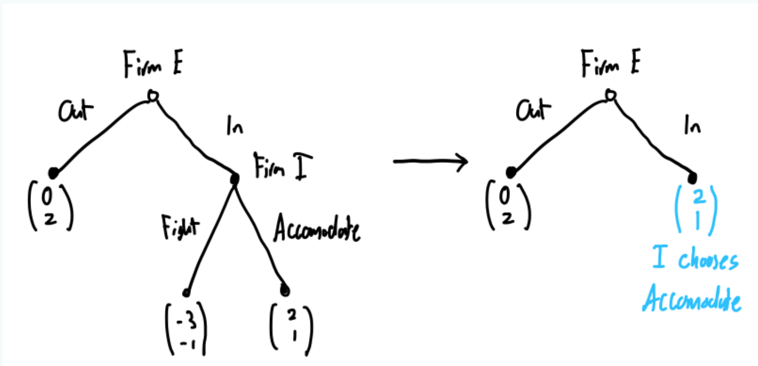

Figure 85: Backward induction applied to predation game.

Example 422

(Sequential Game with 2 Information Sets Continued). We apply backward induction to the game from Example 419, as shown in Fig. 86. At information set a, P2 prefers L to R; at information set b, P2 prefers R to L. Moving back to P1's decision, P1 prefers D to U. Hence, (D, LR) is a sequentially rational Nash equilibrium, while (D, RR) and (U, LL) are not sequentially rational.

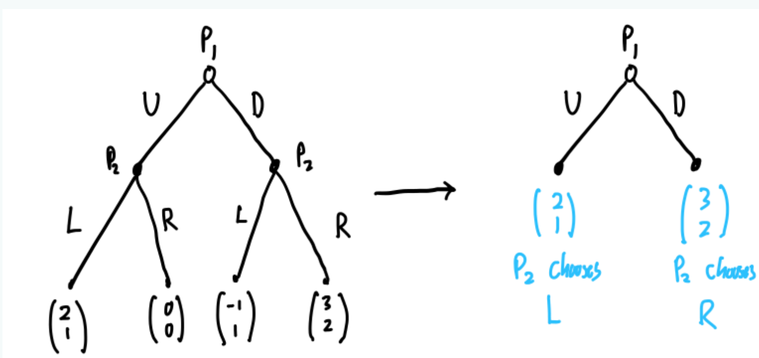

Figure 86: Backward induction applied to sequential game.

Are we guaranteed the existence of a sequentially rational Nash equilibrium? The following theorem gives sufficient conditions and also identifies when the backward-induction outcome is unique.

Theorem 423 (Zermelo's Theorem)
Every **finite** game of **perfect** information has a pure strategy equilibrium that can be derived through backward induction.

Moreover, if each player has distinct payoffs across terminal nodes, backward induction leads to a **unique subgame-perfect Nash equilibrium**.

We can contrast Zermelo's theorem with our previous existence theorem. Zermelo's theorem guarantees a **pure-strategy** equilibrium in finite perfect-information games, whereas Nash's theorem guarantees a **mixed-strategy** equilibrium in every finite normal-form game.

Therefore, every finite perfect-information game has a **pure** Nash equilibrium that satisfies **sequential rationality**. If there are no payoff ties at any decision node, the backward-induction outcome is unique. In practice, however, the game tree may be too large to solve directly, as in chess.

9.5 Subgame Perfect Nash Equilibrium

In this next section we generalise backward induction to dynamic games with imperfect information, in which some information sets may contain more than one node. Zermelo's theorem no longer applies directly because it requires perfect information.

Example 424
(Predation Game with Imperfect Information). We now analyse the predation game with imperfect information, as shown in Fig. 87, in both extensive and normal form. Firm E has two information sets, so each of its pure strategies specifies two actions, e.g. IA. Firm I has one information set, so its strategies specify one action. The normal form has three Nash equilibria: (OA, F), (OF, F), and (IA, A). Ordinary backward induction cannot be applied directly because the game lacks perfect information; we therefore seek a refinement that removes equilibria that are not sequentially rational.

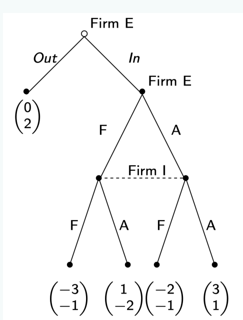

(a) Extensive Form Representation

	A	F
OA	0, $\underline{2}$	$\underline{0}$, $\underline{2}$
OF	0, $\underline{2}$	$\underline{0}$, $\underline{2}$
IA	$\underline{3}$, $\underline{1}$	-2, -1
IF	1, -2	-3, $\underline{-1}$

(b) Normal Form Representation

Figure 87: The predation game with imperfect information, in extensive and normal form.

We want to apply sequential rationality to games of imperfect information. However, we need to do a few things first. We shall see that **subgame perfection** will be the equivalent to sequentially rationality in games of imperfect information. In order to introduce this new concept, we need to define what is a subgame of an extensive form game.

Definition 425 (Subgame)
A **subgame** of an extensive form game Γ_E, is a subset of the game such that:

1. It begins at an information that contains a single decision node x, and contains all the successors of this node, and only those.
2. If the decision node x is in the subgame, every node in the information $I(x)$ is also in the subgame.

In other words, a subgame should start from an information set that is a singleton (this will be the initial node of the subgame), denoted x in the definition. It should contain all x's successors and only those, and it should not cut any information set of the original game. In fact, a subgame is a smaller, well-defined game, that corresponds to a situation that could arise within the original game. The payoffs in the subgame are well-defined and correspond to the payoffs in the original game conditional on reaching that subgame.

Example 426
(Subgame). Look at Fig. 88. The figure on the left shows a game where the only subgame is the game itself. The middle figure shows that there is no subgame starting at P2 as the subgame does not contain all the nodes in the information set for P3 and cuts across information sets. The figure on the right shows that there exists a subgame at P2 now as all nodes in the information set are contained in the subgame.

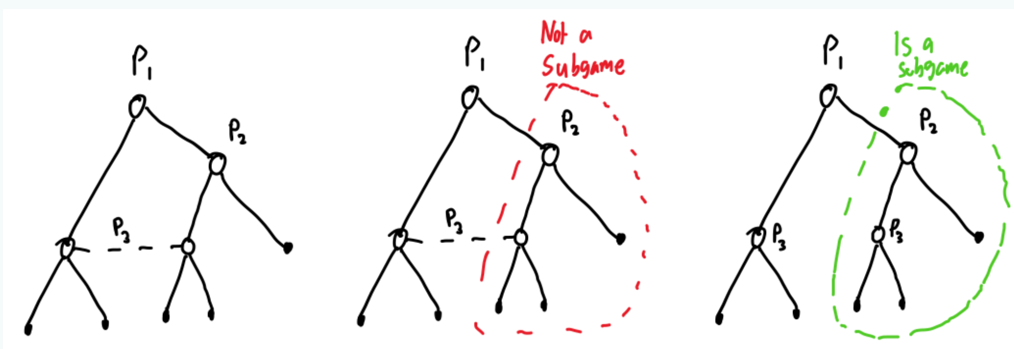

Figure 88: Only one subgame (the game itself) in the game on the left and center. The game on the right has three subgames.

We can now define what we mean by a subgame perfect Nash equilibrium.

Definition 427 (Subgame Perfect Nash Equilibrium)
A behavioral strategy profile σ of an extensive-form game Γ_E is a **subgame perfect Nash equilibrium** (SPNE) if the restriction of σ to any subgame Γ_G is a Nash equilibrium of Γ_G for every subgame Γ_G of Γ_E.

Note that the original game is a subgame of itself.

Proposition 428 (Every SPNE is a NE)
Every subgame perfect Nash equilibrium is a Nash equilibrium. However, the converse does not hold.

Moreover, if the only subgame of the original game is itself, then the set of SPNE coincides with the set of NE. However, if there are more subgames, some NEs might fail to be SPNE.

Proposition 429 (SPNE Induces NE)
Any subgame perfect Nash equilibrium induces a subgame perfect Nash equilibrium at all subgames.

We can use the results we have found so far to characterise another approach to finding SPNEs.

Proposition 430 (Alternative Procedure to Find SPNEs)
If you have a sequential game, find the set of all Nash equilibria in the extensive form of the entire game. Then, find the Nash equilibria in every subgame. Then, the Nash equilibria of the entire game which induces the Nash equilibria found in the subgame is a subgame perfect Nash equilibria.

Remark 431. *Suppose that there is a set of Nash equilibria for an extensive form game. The Nash equilibria which induces a Nash equilibria in every subgame is a subgame perfect Nash equilibria.*

Previously, we showed the existence of mixed strategy NE in a finite game. We can upgrade this result for dynamic games.

Theorem 432 (Existence of Subgame-Perfect Nash Equilibrium)
Every **finite** extensive-form game Γ_E has a subgame-perfect Nash equilibrium.

Proof. (Sketch). Every finite normal-form game has a mixed Nash equilibrium. In a finite extensive-form game, each subgame is also finite. Working backwards through the subgames and selecting a Nash equilibrium in each one yields a strategy profile that is a Nash equilibrium in every subgame, and hence a subgame-perfect Nash equilibrium. □

Remark 433. *Notice that this theorem does* **not** *require perfect information.*

It is easy to see as well that if the game is of perfect information, subgame perfection coincides with the backward induction solution.

Proposition 434 (Subgame Perfect Nash Equilibrium in Perfect Information)
In a game of **perfect information**, a subgame perfect Nash equilibrium coincides with the backward induction solution.

We can generalise the backward induction procedure for games with imperfect information.

Theorem 435 (Generalised Backward Induction Algorithm to Find Sequentially Rational Strategies in Games with Imperfect Information)
Suppose the game we are analysing is a **finite** game with **imperfect information**. The steps for finding sequentially rational strategies are:

1 Start at the final subgames and identify the Nash equilibria of those games.

2 Reduce the game by replacing those subgames with the outcomes prescribed by the selected Nash equilibria.

3 Re-apply step 1 to the new reduced game.

We apply this backward induction algorithm to an example.

Example 436

(Predation Game with Incomplete Information Continued). We apply the generalised backward induction algorithm to Example 424. Look at Fig. 89. We take the extensive form representation of the game and then analyse the subgame of the game starting at Firm E. We then write out the normal form representation of this subgame and compute the Nash equilibrium in the normal form representation. We can see that the Nash equilibrium of this subgame is the strategy profile (A, A).

Now, in Fig. 90, we replace the subgame starting at Firm E with its Nash-equilibrium action profile (A,A) and the associated payoffs. Firm E then prefers In to Out, earning 3 rather than 0.

Therefore, the strategy profile (IA, A) is the **subgame-perfect Nash equilibrium**. It was one of the Nash equilibria previously derived in Example 424; the other two, (OA, F) and (OF, F), are ruled out because they are not sequentially rational.

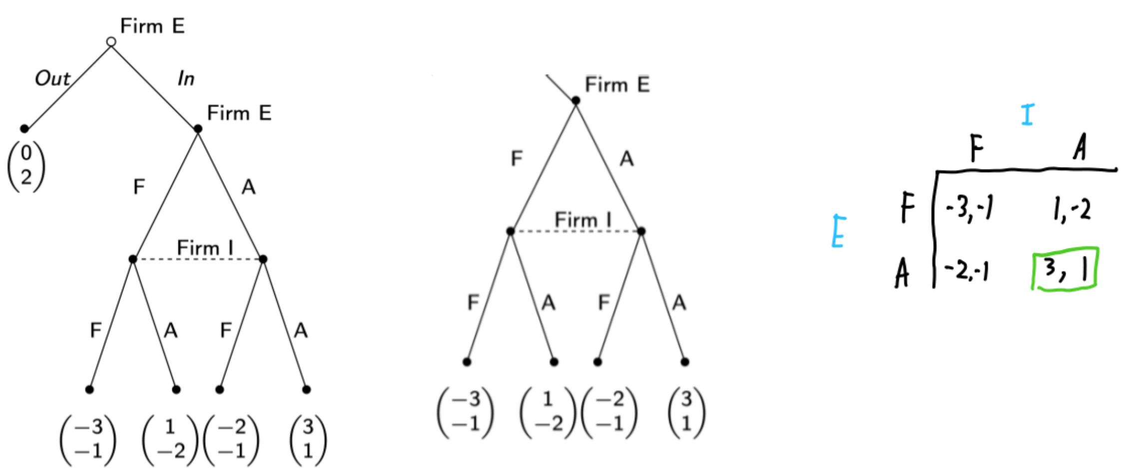

Figure 89: The left figure is the original extensive form game. The middle figure is the subgame starting at Firm E. The right figure is the normal form representation of the subgame starting at firm E. We can see that the Nash equilibrium is given by the strategy profile (A,A).

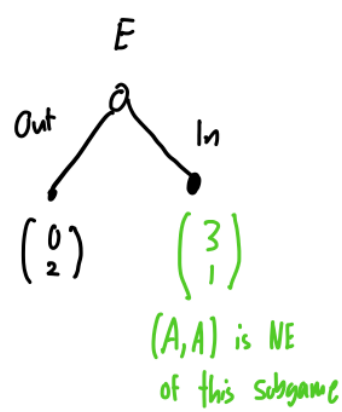

Figure 90: We replace the subgame of the original extensive form game with the Nash equilibrium (A,A) and the associated payoffs (3, 1) that we found in the subgame starting at Firm E.

Whilst subgame-perfect Nash equilibrium may seem like a great solution, there are still many games where the subgame-perfect Nash equilibrium gives us really strange and unintuitive results (see Centipede game).

9.6 Strategic Interactions and Pre-Commitments

We now apply subgame-perfect Nash equilibrium to further industrial-organisation models.

So far, we have seen that there are two static models of competition: the Bertrand model which competed on price and the Cournot model which competed on quantity. In both models, the firms had identical products and made their strategic decision simultaneously. However, the outcomes of the game were quite different:

1. Bertrand: If $n \geq 2$, the perfectly competitive outcome arises: price equals marginal cost and total output equals the perfectly competitive quantity
$$q^{\circ} = \frac{a-c}{b}.$$
2. Cournot: The outcome of the amount produced under the Cournot Q^c is between the monopolistic outcome and the competitive outcome. As the number of firms go to infinity, then we end up with the perfect competition outcome.
$$q^m < Q^c < q^{\circ}$$
where $q^m = \frac{a-c}{2b}$ is the monopoly outcome.

The Bertrand model is more realistic as firms typically choose prices but the Cournot model's results are in line with what we observe.

9.6.1 Stackelberg Model

Two firms compete in quantities as in the Cournot model but firms will make their decisions **sequentially** rather than simultaneously.

A firm (the leader) chooses a quantity $q_l \in \mathbb{R}$. The second firm (the follower) observes the leader's action and then sets its quantity $q_f \in \mathbb{R}$. We can see here that we have perfect information for the follower.

We now suppose that both firms have the same constant marginal cost c and that price is determined by total quantity following a linear decreasing demand function: $p = a - bQ = a - b(q_l + q_f)$. The profits for the leader and follower is the difference between revenue less cost $q[p-c]$ which becomes:

$$\pi_l = q_l[a - b(q_l + q_f) - c] \quad \text{and} \quad \pi_f = q_f[a - b(q_l + q_f) - c]. \tag{173}$$

Therefore, the strategies are as follows:

1. The leader's strategy is to choose any quantity $q_l \in \mathbb{R}_+$. The leader has only one information set.
2. The follower's strategy is a function $f(q_l)$ that, for each quantity q_l chosen by the leader, specifies a quantity $q_f \in \mathbb{R}_+$. The follower has an uncountable number of information sets.

Theorem 437 (Unique Subgame-Perfect Nash Equilibrium of Stackelberg Model)
The **unique** subgame-perfect Nash equilibrium of the Stackelberg model is:

$$(q_l^s, q_f^s) = (\frac{a-c}{2b}, \frac{a-c}{4b}). \tag{174}$$

Proof. We solve this game by starting from the last stage of the game, which is the final subgame. The follower's program is to solve:

$$\max_{q_f} q_f[a - b(q_l + q_f) - c]$$

We can then take first order conditions:

$$\frac{d\pi_f}{dq_f} : a - bq_l - 2bq_f - c = 0$$

We then re-arrange and solve for q_l to get the best response q_f function

$$q_f = \frac{a - c - bq_l}{2b}.$$

However, we should notice that if the leader produces *too much* q_l, then we don't want the follower to produce a negative amount of output as this is impossible. Therefore, the follower's best reply to the leader's quantity is given by:

$$B_f(q_l) = \max\left\{0, \frac{a - c - bq_l}{2b}\right\}. \tag{175}$$

We now move onto the **leader**. The leader should best reply to the best response of the follower $B_f(\cdot)$. Therefore, we have that:

$$\max_{q_l} q_l[a - b(q_l + B_f(q_l)) - c]$$

We then plug in our expression for the best response for the follower:

$$= \max_{q_l} q_l[a - b(q_l + \frac{a - c - bq_l}{2b}) - c]$$

We then take first order conditions with respect to q_l to get:

$$\frac{d\pi_l}{dq_l} : a - bq_l^s - \frac{2b^2}{2b} q_l^s - c = 0$$

where we can then solve for the Stackelberg optimal output for the leader:

$$q_l^s = \frac{a - c}{2b}$$

Something that is interesting to note is that the Stackelberg optimal output for the leader is to produce the **monopoly output** q^m!

Now that we have computed the optimal output for the leader, we can compute the optimal output for the follower using their best response function in Eq. (175):

$$\begin{aligned}
q_f &= \frac{a - c - bq_l}{2b} = \frac{a - c - b\left[\frac{a-c}{2b}\right]}{2b} \\
&= \frac{a - c - \left[\frac{a-c}{2}\right]}{2b} \\
&= \frac{2a - 2c - a + c}{4b} = \frac{a - c}{4b}
\end{aligned}$$

Therefore, the Stackelberg optimal output for the follower is:

$$q_f^s = \frac{a - c}{4b}.$$

We therefore have that the **unique subgame-perfect equilibrium** is given by:

$$(q_l^s, q_f^s) = (\frac{a-c}{2b}, \frac{a-c}{4b}).$$

□

We note a few important things of our results so far. Recall that under the Cournot model, both firms produce an output that is in between the perfectly competitive output q° and the monopoly output q^m. However, we have already shown that under Stackelberg, the leader is producing at the monopoly output level. Therefore, the leader's profits has increased relative to their profits in the Cournot model. This is because the leader could always choose the Cournot quantity if it was optimal. Additionally, in this linear model in which the best response functions are decreasing, the follower's profits fall. This phenomena is known as **first mover advantage**. We also note that the total quantity produced under the Stackelberg model is:

$$Q = q_f + q_l = \frac{a-c}{4b} + \frac{a-c}{2b} = \frac{3}{4}(\frac{a-c}{b})$$

hence the total quantity produced is:

$$Q = \frac{3}{4}(\frac{a-c}{b}).$$

We see that the total quantity produced under Stackelberg is higher than the amount produced in the monopoly outcome q^m but less than the amount produced under the perfectly competitive outcome q°.

We now give the graphical argument of what we have just shown. In Fig. 91, on the figure on the left, we plot out the best response functions of the leader and follower and denoted that the intersection of the two best response functions was the Cournot outcome (as we have previously shown.) The figure on the right shows the isoprofit curve associated to the Cournot outcome for the leader firm.

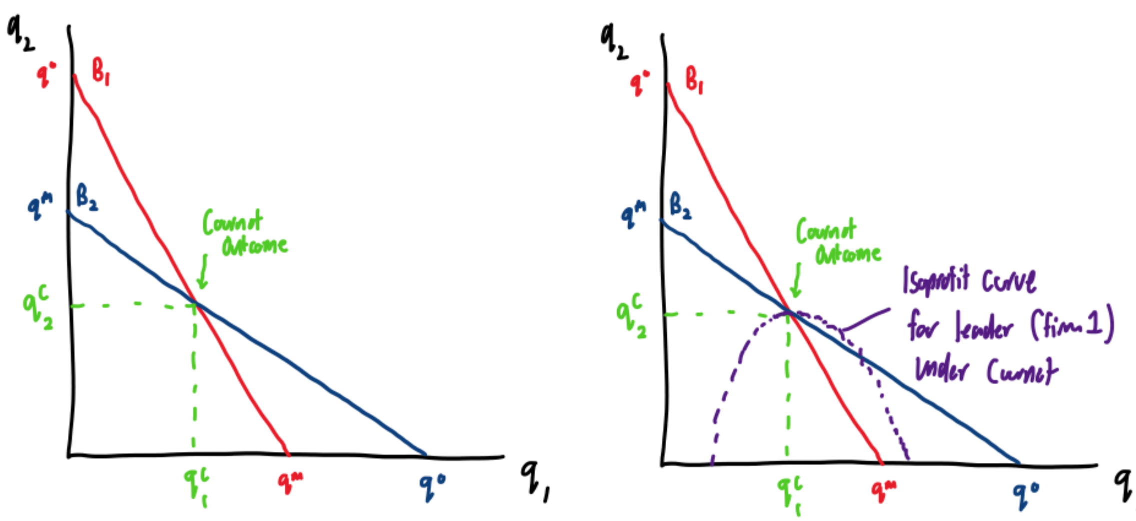

Figure 91: Figure on left denotes the Cournot outcome as the intersection of the two best response function as we have previously shown. Figure on the plot shows the isoprofit curve that is associated to the leader under the Cournot outcome.

Now, under the Stackelberg model, the leader will always chooses first. The leader knows that the follower will **always** respond and produce to whatever the leader produces according to the follower's best response function, which is given by curve $B_2(\cdot)$ (the blue line) since this is the best response and will always be the best response. Therefore, the leader will try to pick a higher an isoprofit line that is associated to higher profits for the leader than the one associated to the Cournot equilibrium (graphically this means a lower isoprofit curve). It turns out, the highest iso-profit curve for the leader is the iso-profit curve that is **tangential** to the best response function of the follower $B_2(\cdot)$. Therefore, the leader will produce an amount q_1^s whereby the isoprofit associated to q_1^s has the same slope as the best response function of the follower $B_2(\cdot)$. We show this in on the left figure in Fig. 92. For the follower, the iso-profit curve associated to the Stackelberg outcome will have lower profits compared to the iso-profit curve associated to the Cournot outcome.

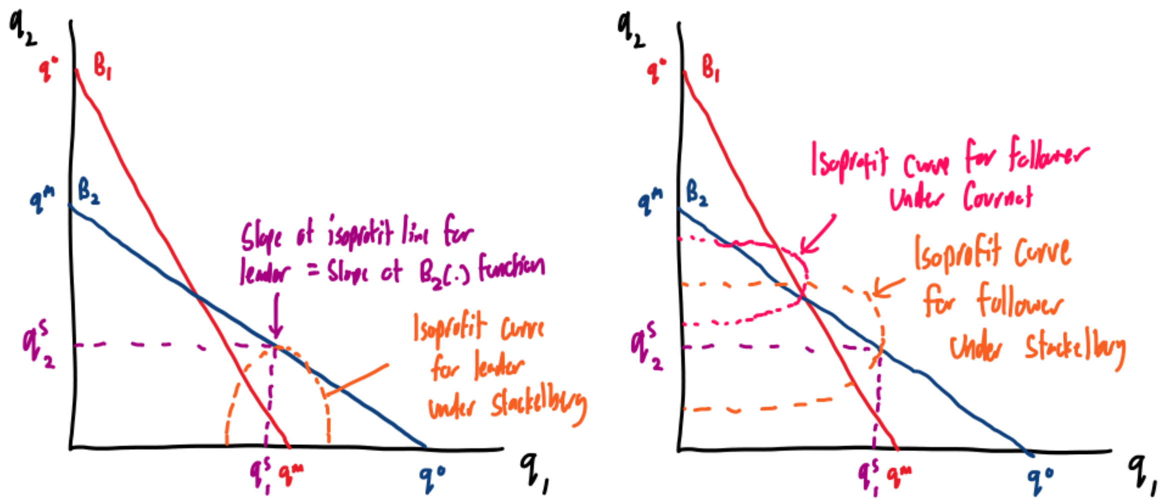

Figure 92: Figure on the left shows the isoprofit curve for the leader firm under the Stackelberg quantities. We see that the slope of the iso-profit curve has the same slope as the best response function of the follower $B_2(\cdot)$. Furthermore, the iso-profit curve for the leader is associated to a higher profit for the leader compared to Cournot. Figure on the right shows the iso-profit curve for the follower under the Stackelberg outcome. This is a lower profit iso-profit curve for the follower compared to the Cournot outcome.

We summarise the Stackelberg outcome and compare it with the Cournot outcome in Fig. 93. As the best response functions are decreasing functions, we can see that in the figure, the quantity produced by the leader (player 1) is higher under the Stackelberg compared to the Cournot $q_1^s > q_1^c$ and hence has higher profits $\pi_1^s > \pi_1^c$. Likewise, we see that the quantity produced by the follower is less under the Stackelberg compared to the Cournot $q_2^s < q_2^c$ and hence has lower profits under Stackelberg $\pi_2^s < \pi_2^c$.

We have shown that there is a unique subgame-perfect Nash equilibrium. However, there are many additional Nash equilibria supported by non-credible threats that the follower would not carry out. We now characterise these equilibria.

Theorem 438 (Nash Equilibria in the Stackelberg Model)
Any leader quantity weakly below the competitive output,

$$q_l = x \leq q^\circ = \frac{a-c}{b},$$

can form part of a Nash equilibrium of the strategic-form representation.

Proof. First, we consider the strategies that are available to the followers based on their best response function. We propose the following strategies that the follower may respond to the leader if the leader sets their output to be less than the perfectly competitive output $x \leq q^\circ$:

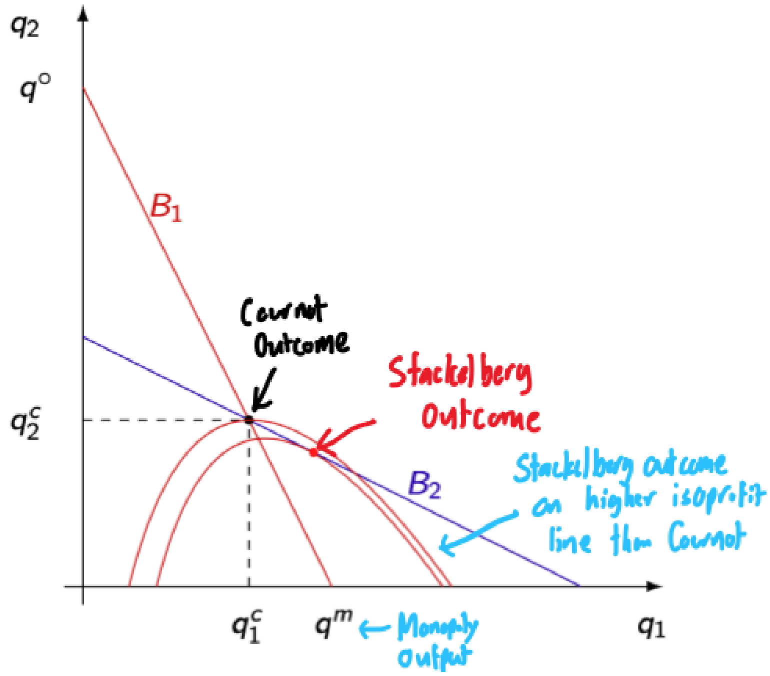

Figure 93: Stackelberg and Cournot outcomes with the best response of the leader B_1 and of the follower B_2. We see that the Stackelberg outcome is on a higher isoprofit curve (for the leader firm) compared to the Cournot outcome.

$$f_x(q_l) = \begin{cases} \frac{a-c-bq_l}{2b} & \text{if } q_l = x \text{ (leader chooses competitive output x)} \\ \\ \max\left\{0, \frac{a-c-bq_l}{b}\right\} & \text{if } q_l \neq x \text{ (leader does not choose output x)} \end{cases} \tag{176}$$

Here, if the leader chooses to produce the competitive quantity $q_l = x$, then the follower can just produce using the best response function we have derived earlier. However, if the leader does not produce the competitive outcome $q_l \neq x$, the follower can now produce a quantity such that the **total quantity** Q is equal to the competitive quantity $Q = q^o$.

We verify that this pair of strategies $(x, f_x(q_l))$ is indeed a Nash equilibrium by seeing that x is a best response to $f_x(q_l)$ and $f_x(q_l)$ is a best response to x.

If the leader chooses $q_l = x$, then the follower's best response to x is to use their best response function. Hence, the follower is doing the best they can already.

Now, we want to see if the leader is best responding to the follower. If the leader produces something different to x, there is a **threat** from the follower that if the leader chooses to produce something other than x, the follower will produce up to the competitive **aggregate quantity** so that both the leader and follower will have zero profits (recall in perfect competition under Cournot, profits are equal to 0!). Therefore, the leader is threatened to produce x. Therefore, the leader has no incentive to deviate away from producing x or else they will be punished by the follower.

We can therefore conclude that $(x, f_x(\cdot))$ is a Nash equilibrium of the Stackelberg model. $\square$

These Nash equilibria are not **sequentially rational**: following a deviation, it is not optimal for the follower to expand output merely to force aggregate output to the competitive level and drive both firms' profits to zero. Thus, the only sequentially rational equilibrium is the Stackelberg outcome derived earlier. This illustrates the value of subgame perfection.

9.6.2 Kreps and Scheinkman, 1983: Bertrand Competition with Capacity Constraints

Suppose that there are two firms $i = 1, 2$ that sell an identical product. There are two stages in the model. They first choose their capacity (investing in machines etc) and then choose what price to charge. That is:

1. Stage 1: Firms choose their **capacity** $q_1, q_2 \in [0, \tilde{q}]$ respectively. Any unit of capacity has a cost k.
2. Stage 2: Firms observe the two capacity constraints (perfect information) and they simultaneously choose which **price** to charge.

Both firms have constant returns to scale with cost per unit c. They face a demand $x(\cdot)$ that is decreasing and such that there exists a price $\tilde{p} > c$ such that $x(\tilde{p}) = 2\tilde{q}$ (there is enough demand to buy all of the firm's production as price is higher than marginal cost). We denote the inverse demand function by $p(\cdot)$.

Given a pair of prices, sales are determined as follows:

1. Customers try to buy from the low-priced firm.
2. If demand exceeds this firm's capacity, consumers are served in descending order of valuation.
3. If the prices are the same, demand is evenly split unless one firm's demand exceeds its capacity. In that case, the extra demand goes to the other firm.

We solve this using backward induction so we start off with the second stage.

Second Stage: Denote prices as (p_1, p_2). Let us look at the sales for the firm i, which we denote by $x_i(p_i, p_j)$. From the first stage, firm i has already chosen a capacity constraint q_i and this is the most that the firm can produce/sell in this second stage. We now have three scenarios to consider regarding firm i's sales.

First, if the price of firm i is less than the other firm $p_i < p_j$, then the firm will either sell to the entire market $x(p_i)$ if it has enough capacity constraint q_i or it may reach its capacity constraint q_i. We can denote this by $\min\{q_i, x_i(p_i)\}$.

Second, suppose that firm i has a higher price than the other firm $p_i > p_j$. If the other firm j has enough capacity, then firm j can sell to the entire market and firm i will sell nothing 0. However, if the other firm j is unable to sell to the entire market as it reaches its capacity constraint q_j, then the excess demand will spillover to firm i $x(p_i) - q_j$. We denote this by $\max\{x(p_i) - q_j, 0\}$. However, with this excess demand, firm i may either be able to meet all this exceed demand or be constrained by its own capacity constraint q_i. Therefore, we have that firm i's sales is: $\min\{q_i, \max\{x(p_i) - q_j, 0\}\}$.

Finally, suppose that both firms charge the same price, $p_i = p_j$. They split demand evenly, $x(p_i)/2$, if both have sufficient capacity. If firm j cannot serve half of demand, firm i also receives the residual demand, $x(p_i) - q_j$. Combining these cases and imposing firm i's own capacity constraint gives $\min\{q_i, \max\{x(p_i)/2, x(p_i) - q_j\}\}$.

Definition 439 (Sales for a firm in second stage)
We can therefore summarise the sales for firm i given (fixed) capacity constraints q_i, q_j and prices as:

$$x_i(p_i, p_j) = \begin{cases} \min\{q_i, x(p_i)\} & \text{if } p_i < p_j \\ \min\{q_i, \max\{x(p_i) - q_j, 0\}\} & \text{if } p_i > p_j \\ \min\{q_i, \max\{\frac{x(p_i)}{2}, x(p_i) - q_j\}\} & \text{if } p_i = p_j \end{cases} \tag{177}$$

If prices are equal to marginal costs, then firms do not make a profit. However, we note that since demand is such that if prices equal marginal cost, there is excess demand (by assumption) $x(c) > q_i$ for $i = 1, 2$ then the prices $p_1 = p_2 = c$ is not an equilibrium as a firm can raise their prices above marginal costs, then the other firm's capacity constraint will bind and the firm with the raised prices can absorb the excess demands and now make positive prices. We can summarise this in the following proposition we have just proven.

Proposition 440 (Bertrand Competition with Capacity Constraint Equilibrium Differs)
The Bertrand Competition with Capacity constraints will have a **different equilibrium** compared to the Bertrand Competition model where prices are equal to marginal cost

$$(p_1, p_2) = (c, c)$$

In fact, each firm is guaranteed to make positive profits by setting a price p_i such that $c < p_i < p(q_j)$ where $p(q_j)$ is the price that will make the capacity constraint of firm j.

We now can derive the Nash equilibrium in the second stage of this game.

Theorem 441 (Nash Equilibrium in Second Stage of Bertrand Model with Capacity Constraints)
Denote $B_i(\cdot)$ as the **Cournot** best response function. If the capacity constraints q_i are such that $q_i \leq B_i(q_j)$ for $i = 1, 2$, then the **unique** pure Nash equilibrium of this price competition duopoly model is:

$$p_1^* = p_2^* = p(q_1 + q_2)$$

where firms set identical prices to the **Cournot equilibrium outcome**.

Remark 442. *The state $q_i \leq B_i(q_j)$ for $i = 1, 2$ says that firm i chooses a capacity constraint q_i that is* **below** *the best response capacity constraint they should have picked if firm j picked capacity constraint q_j. This means that we* **do not have** *an optimal level of capacity constraints.*

Proof. We show that $p_1^* = p_2^* = p(q_1 + q_2)$ is in fact a Nash equilibrium.

First, fix the price set by firm 2 at $p_2^* = p(q_1 + q_2)$ and see what is the best response for firm 1.

Suppose that firm 1 chooses $p_1 \leq p_2^* = p(q_1 + q_2)$ and sells at full capacity. Its profit is then $\pi_1 = (p_1 - c)q_1$. It cannot be optimal to set $p_1 < p(q_1 + q_2)$, because firm 1 could raise its price slightly, continue to sell at full capacity, and earn more. This holds for every $\epsilon > 0$ such that $p_1 + \epsilon < p(q_1 + q_2)$.

We now suppose that firm 1 chooses a price such that $p_1 \geq p_2^* = p(q_1 + q_2)$. Now, firm 2 will sell at full capacity q_2 and firm 1 will absorb this excess demand $x(p_1) - q_2$. Then, we can see what is the optimal price for firm 1 to set by looking at its profit maximisation problem:

$$\max_{x_1 \leq q_1} \left[p(x_1 + q_2) - c\right] x_1$$

whereby firm 1 chooses a quantity to produce x_1 as long as it is less than its capacity constraint q_1 whereby $x_1 + q_2$ is the total production in the market. However, this is exactly the same problem in a Cournot duopoly when firm 2 produces q_2 which we have already solved! By assumption, if $B_1(q_2) \geq q_1$, then the best firm 1 could do is to produce as much as we can, which corresponds to setting $x_1 = q_1$ which is our capacity constraint, and this will correspond to having a price $p(x_1 + q_2) = p(q_1 + q_2)$. Hence, the price that firm 1 should set is $p_1 = p(q_1 + q_2)$.

By a symmetric argument, we can follow the same steps for firm 2. Therefore, $p_1^* = p_2^* = p(q_1 + q_2)$ is a Nash equilibrium of this game.

We now show that this Nash equilibrium is in fact **unique**. First, recall that we are always guaranteed to make a positive profit independently of the other firm's strategy by setting a sufficiently low price that is between the price of the other firm and the marginal cost. Therefore, in the Nash equilibrium, both firms **must** make positive profits or else they have an incentive to deviate to following the strategy we just outlined.

Now, suppose that in equilibrium, firms charge different prices $p_i < p_j$. Since firm j makes positive profits, then firm i **must** be selling at full capacity (or else no excess demand for firm j to absorb and make profits from!). This means that firm i can slightly raise its prices while still selling at full capacity and hence increase its profits. Therefore in equilibrium, firms must charge the same price. We now show that there is only one unique price that this could be.

Suppose instead that $p_1 = p_2 > p(q_1 + q_2)$. At least one firm then sells below capacity and can profitably undercut the common price, either filling its capacity or attracting all consumers from its rival. Hence, this cannot be an equilibrium.

Now suppose we are in the equilibrium where $p_1 = p_2 < p(q_1 + q_2)$ so that both firms are selling at full capacity. Then, each firm could make profits by raising its price to $p(q_1 + q_2)$ and still have binding capacity constraints.

Therefore, we can conclude that $p_1^* = p_2^* = p(q_1 + q_2)$ is the **unique** pure Nash equilibrium of this duopoly.

□

We have solved the second stage of the game whenever we had the non-optimal level of capacity constraints $q_i \leq B_i(q_j)$. For the other scenario where the capacity constraint is too high $q_i > B_i(q_j)$ for some i, the equilibrium in the second stage is more convoluted but we will just take it as given that firm i will never set a price above $p(q_j + B_i(q_j))$ and hence never produce above $B_i(q_j)$. Therefore, they are making profits as if they are producing at the best response level.

First Stage: We now solve the first stage. Firms choose capacities q_i, each unit of which costs k, knowing that in the second period they will play the Nash equilibrium $p_1^* = p_2^* = p(q_1 + q_2)$.

Since each unit of capacity costs the fixed amount k, it will never be optimal for firm i to set a capacity constraint above their best response $q_i > B_i(q_j)$ as they will not use this extra capacity and it will be costly. Therefore, we only to focus on the case where $q_i \leq B_i(q_j)$. However, we have already seen that the Nash equilibrium in the second stage (which we spent a long time solving) where $q_j \leq B_i(q_j)$ is $p_1^* = p_2^* = p(q_1 + q_2)$. Therefore, the problem for firm i in the first period becomes:

$$\max_{q_i} \left[p_i^* - c - k\right] q_i = \max_{q_i} \left[p(q_i + q_j) - c - k\right] q_i \tag{178}$$

Thus, each firm chooses capacity knowing that the stage-2 price is $p(q_1 + q_2)$. Because each unit of capacity also costs k, this is the **Cournot competition model** with marginal cost $c + k$.

Therefore, we can actually conclude that the Bertrand competition model with capacity constraint is actually the Cournot model in disguise!

We have so far been trying to reconcile the results from the Bertrand model (which competed on prices) and the Cournot model (which competed on quantity) whereby the Cournot model had more realistic results but the Bertrand had the more realistic assumption of firms choosing prices. Based off the Bertrand model with capacity constraints, we can rationalise the Cournot competition model as firms actually choosing prices rather than quantities whereby they have a two-stage process of first setting up their capacity constraint, which affects how much they can produce at stage 2, and then setting prices.

We now move onto another extension.

Example 443

(Strategic Pre-Commitment). Suppose that firms can make strategic pre-commitments to make their future threats credible. The setup is as follows:

- In period 1, (only) firm 1 makes a strategic investment $k \in \mathbb{R}$.
- In period 2, firms 1 and 2 engage in competition by simultaneously choosing strategies $s_1 \in \mathbb{R}^+, s_2 \in \mathbb{R}^+$ respectively. This can be prices or quantities. Their profits are given by: $\pi_1(s_1, s_2, k)$ and $\pi_2(s_1, s_2)$. Note

that the only firm one has the strategic investment k affecting their profits.

In the second period, the Nash equilibrium is given by the intersection of the two best responses:

$$(s_1^*, s_2^*) \quad \text{where } s_1^* = B_1(s_2^*, k) \text{ and } s_2^* = B_2(s_1^*)$$

Keeping this Nash equilibrium in mind, then, going to the first period, the **subgame-perfect Nash equilibrium** is firm 1 choosing a strategic investment k to maximise:

$$\max_k \pi_1(s_1^*, s_2^*, k)$$

where (s_1^*, s_2^*) is the Nash equilibrium from the second period. When choosing k, firm 1 considers both the direct cost effect and the strategic effects in stage 2. Applying the chain rule, the effect of investment on equilibrium profit is:

$$\frac{d\pi_1(s_1^*, s_2^*, k)}{dk} = \frac{\partial \pi_1(s_1^*, s_2^*, k)}{\partial k} + \frac{\partial \pi_1(s_1^*, s_2^*, k)}{\partial s_1}\frac{ds_1^*(k)}{dk} + \frac{\partial \pi_1(s_1^*, s_2^*, k)}{\partial s_2}\frac{ds_2^*(k)}{dk}$$

The first term measures how profit changes directly as we vary k; this is the **direct effect**. The second captures the change through firm 1's period-2 strategy s_1, while the third captures the change through firm 2's period-2 strategy s_2.

However, we can invoke the **envelope theorem** since at the optimised variable s_1^*, the indirect effect $\partial \pi_1 / \partial s_1 = 0$ (use first order conditions to verify this). This is because firm 1 chooses s_1 to be at the optimal level s_1^* in the second period. However, we also have that $\partial \pi_1 / \partial s_2 \neq 0$ as s_2 is **not** a variable that firm 1 chooses.

Definition 444 (Change in Equilibrium Profits with Direct and Strategic Effects)

We have that the change in equilibrium profits as a result of a change in investment k is given by:

$$\frac{d\pi_1(s_1^*, s_2^*, k)}{dk} = \underbrace{\frac{\partial \pi_1(s_1^*, s_2^*, k)}{\partial k}}_{\text{Direct Effect}} + \underbrace{\frac{\partial \pi_1(s_1^*, s_2^*, k)}{\partial s_2}\frac{ds_2^*(k)}{dk}}_{\text{Strategic Effect}} \tag{179}$$

We now analyse the strategic effect more carefully. The sign of $\frac{\partial \pi_1(s_1^*, s_2^*, k)}{\partial s_2}\frac{ds_2^*(k)}{dk}$ depends on both how firm 1's profit varies with firm 2's strategy and how firm 2's equilibrium strategy responds to k.

First, consider the sign of $\frac{ds_2^*(k)}{dk}$. At an interior equilibrium, the fixed-point characterisation of mutual best responses gives $s_2^*(k) = B_2(B_1(s_2^*(k), k))$. Implicit differentiation then yields:

$$\frac{ds_2^*(k)}{dk} = \frac{dB_2(B_1(s_2^*(k), k))}{dk} = \frac{dB_2}{ds_1}\left[\frac{\partial B_1}{\partial s_2}\frac{ds_2^*(k)}{dk} + \frac{\partial B_1}{\partial k}\right]$$

where we used the definition that $s_1 = B_1(s_2^*(k), k)$. We can then re-arrange to get the expression:

$$\frac{ds_2^*(k)}{dk} = \frac{\frac{dB_2}{ds_1}\frac{\partial B_1}{\partial k}}{1 - \frac{dB_2}{ds_1}\frac{\partial B_1}{\partial s_2}}$$

In **stable games**, the denominator is positive so:

$$1 - \frac{dB_2}{ds_1}\frac{\partial B_1}{\partial s_2} > 0$$

since the sequence of best responses converge in equilibrium, the product of the best-reply slopes is less than 1.

Therefore, we only need to consider the sign of the numerator:

$$\frac{dB_2}{ds_1} \times \frac{\partial B_1}{\partial k}$$

on whether is firm 2's response positive or negative:

$$\text{Sign}\left[\frac{ds_2^*(k)}{dk}\right] = \text{Sign}\left[\frac{\frac{dB_2}{ds_1}\frac{\partial B_1}{\partial k}}{1 - \frac{dB_2}{ds_1}\frac{\partial B_1}{\partial s_2}}\right] = \text{Sign}\left[\frac{dB_2}{ds_1} \times \frac{\partial B_1}{\partial k}\right]$$

We have therefore found the conditions to characterise the sign of the change in firm 2's strategy $\frac{ds_2^*(k)}{dk}$. We now go back to the **Strategic effect** term in Eq. (179):

$$\frac{\partial \pi_1(s_1^*, s_2^*, k)}{\partial s_2}\frac{ds_2^*(k)}{dk}$$

We have just shown that the sign of $\frac{ds_2^*(k)}{dk}$ depends only on the sign of $\frac{dB_2}{ds_1} \times \frac{\partial B_1}{\partial k}$. We can therefore plug this in to get the sign of the strategic effects.

Theorem 445 (Sign Of Strategic Effects)
The sign of the strategic effects is given by the sign of:

$$\text{Sign}\big[\text{Strategic Effect}\big] = \text{Sign}\left[\frac{\partial \pi_1}{\partial s_2} \times \frac{dB_2}{ds_1} \times \frac{\partial B_1}{\partial k}\right] \tag{180}$$

This says that the sign of the strategic effect depends on three things:

1. Does firm 1's profits increase or decrease with firm 2's strategic choice?
2. Is firm 2's best response function increasing or decreasing in s_1?
3. How does the strategic investment affect firm 1's best-response?

We now apply our example of strategic cost reduction to the Cournot model as we can work out the effect of strategic effects.

Example 446
(Strategic Cost Reduction in Cournot). Suppose that we are under a Cournot model where firms compete under quantities. Then, suppose that there is a situation where k is an investment in marginal cost reduction. Suppose that only firm 1 can make such an investment to lower their marginal cost. Therefore, their profit function is:

$$\pi_1(q_1, q_2, k) = q_1[P(Q) - c(k)] - k$$

where $P(Q)$ is the inverse demand function, k is the investment cost, and $c(k)$ is firm 1's marginal cost which depends on investment cost k and is $c'(k) < 0$. For firm 2, the profit function is:

$$\pi_2(q_1, q_2) = q_2[P(Q) - c]$$

where we don't have the investment cost k and the marginal cost c does **not** depend on k.

Now, assume that the linear inverse demand function is $P(Q) = 1 - Q = 1 - (q_1 + q_2)$. We have already shown that under such a situation, the best response function for firm 1 is given by:

$$B_1(q_2, k) = \frac{1 - q_2 - c(k)}{2}$$

and by symmetry, for firm 2 it is:

$$B_2(q_1, k) = \frac{1 - q_1 - c}{2}$$

We now want to compute the **strategic effect** of firm 1 investing in k. We can use the formula we have derived:

$$\text{Sign}\big[\text{Strategic Effect}\big] = \text{Sign}\left[\frac{\partial \pi_1}{\partial s_2} \times \frac{dB_2}{ds_1} \times \frac{\partial B_1}{\partial k}\right]$$

First, we can take the derivative of firm 1's best response and see that:

$$\frac{\partial B_1}{\partial k} = \frac{\partial}{\partial k}\left[\frac{1 - q_2 - c(k)}{2}\right] = \frac{-c'(k)}{2} > 0$$

as $c'(k) < 0$ by assumption. Therefore, this shifts the best response function of firm 1 out as we will soon see. Looking at the second term, we have that:

$$\frac{\partial B_2}{\partial q_1} = \frac{\partial}{\partial q_1}\left[\frac{1 - q_1 - c}{2}\right] = -\frac{1}{2} < 0.$$

This says that we have **strategic substitutes**. That is, as firm 1 produces more, then firm 2 produces less. This means that the best response for firm 2 decreases as firm 1 produces more.

Finally, for the third term, we have:

$$\frac{\partial \pi_1}{\partial q_2} = \frac{\partial}{\partial q_2}\Big[q_1[P(Q) - c(k)] - k\Big] = \frac{\partial}{\partial q_2}\Big[q_1[(1 - (q_1 + q_2)) - c(k)] - k\Big] = -q_1 < 0.$$

That is, as firm 2 produces more, then the profits of firm 1 falls. For the sign of the products of these effects, we have two negatives and one positive, and therefore this implies that the sign of the strategic effect is positive:

$$\text{Sign}\big[\text{Strategic Effect}\big] = \text{Sign}\left[\frac{\partial \pi_1}{\partial s_2} \times \frac{dB_2}{ds_1} \times \frac{\partial B_1}{\partial k}\right] > 0$$

The strategic effect is positive. Investing in cost reduction increases q_1, which induces firm 2 to reduce q_2; this rival response benefits firm 1 relative to a setting in which q_2 did not adjust. We show this in Fig. 94.

Example 447

(Strategic Cost Reduction in Hotelling Model). We look at strategic effects in the Hotelling differentiated-product price model with transport cost $t = 1/2$. Suppose that in period 1, firm 1 can invest k in cost reduction, where its marginal-cost function satisfies $c'(k) < 0$. The profit functions are:

$$\pi_1 = (p_1 - c(k))\big(\frac{1}{2} + p_2 - p_1\big) - k$$

and

$$\pi_2 = (p_2 - c)(\frac{1}{2} + p_1 - p_2)$$

We have also already shown that the respective best response functions are:

$$B_1(p_2, k) = \frac{1}{4} + \frac{p_2 + c(k)}{2} \qquad B_2(p_1) = \frac{1}{4} + \frac{p_1 + c}{2}$$

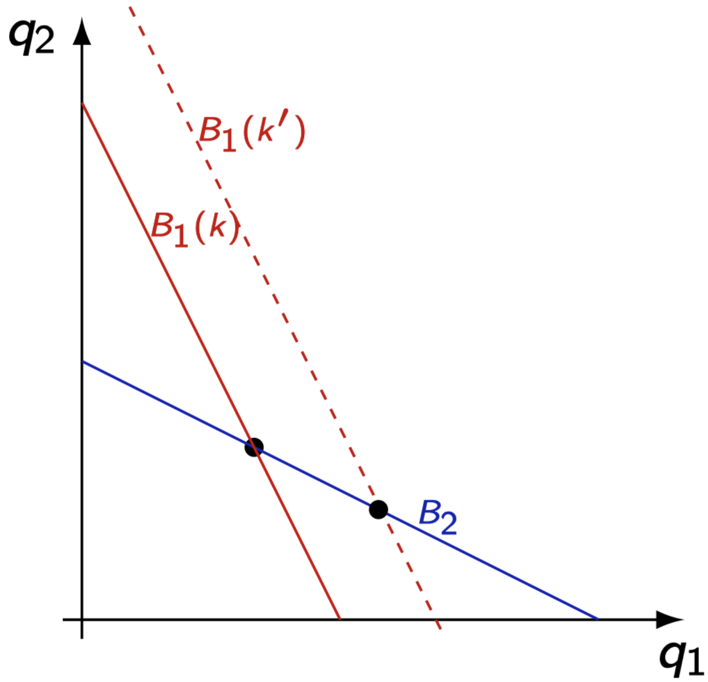

Figure 94: We have shown that $\frac{\partial B_1(k)}{\partial k} > 0$ which implies that the best response function B_1 for firm 1 shifts out as we increase investment k. We then see that at the new intersection of the two best response function, firm 1 is producing more than before whilst firm 2 is now producing less.

We can compute the signs needed for the strategic effects.

$$\frac{\partial B_1}{\partial k} = \frac{\partial}{\partial k}\left[\frac{1}{4} + \frac{p_2 + c(k)}{2}\right] = \frac{c'(k)}{2} < 0$$

$$\frac{\partial B_2}{\partial p_1} = \frac{\partial}{\partial p_1}\left[\frac{1}{4} + \frac{p_1 + c}{2}\right] = \frac{1}{2} > 0$$

$$\frac{\partial \pi_1}{\partial p_2} = \frac{\partial}{\partial p_2}\left[(p_1 - c(k))\left(\frac{1}{2} + p_2 - p_1\right) - k\right] = p_1 - c(k) > 0$$

We therefore have that the sign of the strategic effect is negative:

$$\text{Sign}\left[\text{Strategic Effect}\right] = \text{Sign}\left[\frac{\partial \pi_1}{\partial s_2} \times \frac{dB_2}{ds_1} \times \frac{\partial B_1}{\partial k}\right] < 0$$

The strategic effect is negative. Investing in cost reduction lowers p_1 and also induces firm 2 to lower p_2, which hurts firm 1. The **direct effect** must therefore be sufficiently positive for the investment to be worthwhile.

To summarise this section, we introduced sequential rationality: when making a decision, we should consider subsequent players' best responses. We derived backward-induction procedures for perfect-information games and subgame-perfect Nash equilibria for games with imperfect information. Finally, we used subgame perfection to analyse several strategic effects.

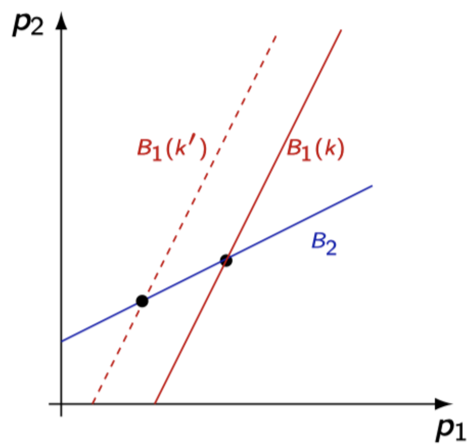

Figure 95: We have shown that $\frac{\partial B_1(k)}{\partial k} < 0$ which implies that the best response function B_1 for firm 1 shifts in as we increase investment k. We then see that at the new intersection of the two best response function, firm 1 sets prices lower than before whilst firm 2 also sets prices to be lower.

We now move onto a new type of sequential game, known as **repeated games**. Here, we will have two or more parties interact repeatedly and also be aware that they will be playing against each other multiple times.

9.7 Finite Multi-Stage Games with Observed Actions

Most of the extensive games we will consider are multi-stage games with observed actions. These games have multiple stages. In each stage, players simultaneously chose some actions. Those actions are then observed by all players before proceeding to the next stage.

The particular structure of these games simplifies the analysis of equilibria that are subgame perfect. Indeed, the subgames of multi-stage games with observed actions correspond to the games that arise after the actions at the end of every stage are revealed. We introduce now some notation that will formalise this.

We will consider games with T stages. The stages will be denoted by $t = 1, ..., T$.

Definition 448 (Action Profiles)
We denote by the vector $a^t = (a_1^t, ..., a_n^t)$ the **action profile** that the n players choose simultaneously at stage t.

Remark 449. *The set of players active at each stage could be different. We can model this by setting that inactive players do nothing at a particular stage.*

At stage t, players do not have perfect information of the stage t as they do not observe what other players are playing at stage t. However, they do get to observe all previous actions made. We call this the history.

Definition 450 (History of Actions)
At the beginning of stage t, all players observe the actions that have been played before. We call h^t the **history of actions before stage** t:

$$h^t = (a^1, ..., a^{t-1}) = \left((a_1^1, ..., a_n^1), ..., (a_1^{t-1}, ..., a_n^{t-1}) \right) \tag{181}$$

whereby $a^j = (a^j_1, ..., a^j_n)$ is the action profile by the n players at stage j.

The set of all **possible histories** before stage t is denoted by H^t.

Remark 451. *Note that h^t is a* **vector of vectors**, *where each entry is an action profile. In stage 1, nothing has happened, so $h^1 = \emptyset$. Finally, H^t is the set of all histories that could have occurred before stage t, whilst h^t is a realised sequence of action profiles.*

One thing that differs to our previous settings is that the actions available to the player now depends on the history.

Definition 452 (Set of Feasible Actions)
The set of actions feasible to player i in stage t after history h^t is $A_i(h^t)$.

Because all past actions are publicly observed, every public history h^t begins a proper subgame. Actions within stage t remain simultaneous, so players do not observe one another's current actions before choosing. A strategy is therefore a contingent plan specifying play **at each stage** after **every possible** public history.

Definition 453 (Strategies)
A **pure strategy** for player i can be seen as a sequence of functions

$$s_i = \{s_i^t\}_{t=1}^T \tag{182}$$

such that $s_i^t(h^t) \in A_i(h^t)$ for all histories $h^t \in H^t$.

A **behavioural** strategy for player i is a sequence of probability measures

$$\{\sigma_i^t\}_{t=1}^T \tag{183}$$

with $\sigma_i^t(h^t) \in \Delta(A_i(h^t))$ where Δ is the set of feasible behavioural strategies.

Remark 454. *At round t, for each history h^t, player i will only be able to choose an action out of all available actions $s_i^t(h^t) \in A_i(h^t)$.*

A strategy profile s induces a **terminal history**. For example, at the first stage $a_s^1 = s^1(\emptyset)$, this tells us what should the first action for player i be given that the player knows nothing since the history is $h^1 = \emptyset$. Then, at the second stage, the player should play an action that is based on the first stage of action $a_s^2 = s^2(a_s^1)$. We repeat this for the third round $a_s^3 = s^3(a_s^1, a_s^2)$ and so on. Therefore, a strategy profile s induces a **terminal history** $h^{T+1} = (a_s^1, a_s^2, ..., a_s^T)$.

Definition 455 (Terminal History)
A **terminal history** h^{T+1} is a history of actions that is induced by a strategy profile s:

$$h^{T+1} = (a_s^1, a_s^2, ..., a_s^T). \tag{184}$$

We can now define payoff functions that are based on the history that the strategy profile s induced.

Definition 456 (Payoff Functions)
The **payoff functions** are defined as functions of terminal histories:

$$u_i : H^{T+1} \to \mathbb{R}. \tag{185}$$

The payoff induced by strategy s for player i:

$$u_i(s) \equiv u_i(a_s) \tag{186}$$

as the utility from the history of actions.

Remark 457. *Note that the payoff functions map from the domain of* **all possible histories**.

Therefore, we can see that a strategy profile will determine the final path of actions and generate a particular payoff. We also want to be able to analyse the subgames of this game. Note that we observe all past actions, every information set is a singleton set. This makes it easy to define subgames.

Definition 458 (Subgame)
$G(h^t)$ denotes the **subgame** induced by history h^t.

Remark 459. *This is the game that starts at stage t once history h^t has been realised.*

From our setup so far, a strategy profile s (or σ) for the whole game induces a strategy profile $s|h^t$ (or $\sigma|h^t$) on the subgame $G(h^t)$.

The definition of a Nash Equilibrium (NE) is the same.

Definition 460 (Nash Equilibrium)
A strategy profile s is a Nash equilibrium if:

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \tag{187}$$

for all s_i' and every player i.

With this notation, the concept of subgame perfection can be written as:

Definition 461 (Subgame Perfect Nash Equilibrium)
A pure strategy profile s of a multi-stage game with observed actions is a **subgame-perfect Nash equilibrium** (SPNE) if the induced strategy profile $s|h^t$ is a Nash equilibrium of the subgame $G(h^t)$ for every history h^t.

Remark 462. *We see that the strategy profile is conditioned on observing history h^t. We can also extend this definition to behavioural strategies whereby we mix among actions that are available to you at the information set.*

Hence, we can see that every dynamic game can be represented as a simultaneous move game where strategies are complete contingent plans.

9.8 Finite Concatenated Games

We now look at a particular example of a multi-stage game with observed actions. **Finite concatenated games** are an example.

Definition 463 (Finite Concatenated and Repeated Games)
A **finite concatenated game** is a multi-stage game with observed actions in which, at each stage, players play a simultaneous-move game $\Gamma_t = [N, \{A_i^t\}, \{u_i^t\}]$ and the payoff of the whole game is the sum of stage payoffs:

$$u_i(s) = \sum_{t=1}^{T} u_i^t(a_s^t). \tag{188}$$

It is a **repeated game** when the same stage game is played in every period. This models repeated interactions between agents.

Example 464
(Finite Repeated Game). Suppose we have a two-period game with two players, where the same game is played in both periods. Before the second interaction, players observe both first-period actions. The stage-game payoffs are shown in Fig. 96. The **unique** Nash equilibrium is (M, C), with payoff $(3, 3)$. A more attractive path would play (T, R) in the first period and (B, R) in the second. Consider the following strategy profile for the twice-repeated game:

1. Player 1 Strategy: Play T in the first stage. In second stage, play B if observed history $h^2 = (T, R)$, otherwise play M.
2. Player 2 Strategy: Play R in the first stage. In second stage, play R if observed history $h^2 = (T, R)$, otherwise play C.

Under this strategy, the payoff $u_i(s) = 7$ for both players compared to 6 if they played (M, C) both stages. However, this strategy is not sequentially rational. Player 2 is incentivised to deviate in the first turn from this proposed strategy and instead play L to earn a payoff of 7 in the first turn and then receive a payoff of 3 in the second turn (M, C).

	L	C	R
T	0, 7	2, 5	7, 0
M	5, 2	3, 3	5, 2
B	7, 0	2, 5	0, 7

Figure 96: Stage game for 2 period game. Unique Nash equilibrium in stage game is (M, C).

In Example 464, we proposed a Pareto-improving path. Unfortunately, the associated strategy profile is not sequentially rational.

Theorem 465 (Unique SPNE is a Sequence of Unique NEs)
Consider an extensive-form game Γ_E in which n players play $T < \infty$ concatenated simultaneous-move games Γ_t, and suppose that each Γ_t has a unique Nash equilibrium.

Then, the extensive form game Γ_E has a **unique subgame-perfect Nash equilibrium**, and the strategies are the concatenated equilibrium strategies of the individual games.

Proof. Denote the unique Nash equilibrium of game Γ_t by $\hat{\sigma}^t$. The proof is by induction on the number of games T.

If $T = 1$, then the unique Nash equilibrium of the simultaneous move game is the unique SPNE.

Suppose that the statement is true for $T = M - 1$ concatenated games: if we have $M - 1$ concatenated games, the unique SPNE of the game is σ^* such that $\sigma^*(h^t) = \hat{\sigma}^t$ for all h^t.

Consider $T = M$ concatenated games, and let σ be an SPNE of the whole game. Then $\sigma|h^2$ must be an SPNE in $G(h^2)$ for every history h^2. But $G(h^2)$ is a concatenated game with $M - 1$ stages, so by the induction hypothesis, $\sigma|h^2(h^t) = \hat{\sigma}^t$ for all h^t and $t \geq 2$.

Let $U_i = \sum_{t=2}^{M} u_i^t(\hat{\sigma}^t)$ denote player i's total equilibrium payoff from the last $M - 1$ games. The reduced game that replaces those games by their equilibrium continuation payoff has payoffs: $u_i^1(a_i^1, a_{-i}^1) + U_i$ where (a_i^1, a_{-i}^1) is the profile of actions played in game Γ_1. The Nash equilibrium of this reduced game therefore coincides with the Nash equilibrium of Γ_1, which is unique by assumption.

Hence $\sigma(h^t) = \hat{\sigma}^t$ for every h^t and $t \geq 1$, so the result holds for $T = M$. □

Remark 466. *If the games Γ_t have an* **unique equilibrium***, then finitely repeated interactions cannot change the behaviour of the players in the individual games. Moreover, the strategies of the unique SPNE cannot be history dependent.*

The unique SPNE plays the unique NE at every stage. If some stage games have multiple Nash equilibria, there may instead be multiple SPNE, including equilibria that do not induce a stage-game equilibrium in an early period.

Example 467

(Finite Repeated Game Continued). We modify the payoffs from Example 464, as shown in Fig. 97. The stage game now has three Nash equilibria:

1. (T, C) where the payoffs are $u(T, C) = (3, 4)$
2. (M, L) where the payoffs are $u(M, L) = (4, 3)$
3. The mixed strategy $(\sigma_1 = (\frac{3}{7}, \frac{4}{7}, 0); \sigma_2 = (\frac{3}{7}, \frac{4}{7}, 0))$ with the payoff $u(\sigma_1, \sigma_2) = (\frac{12}{7}, \frac{12}{7})$.

However, (B, R) Pareto dominates all three equilibria. Can this Pareto-superior outcome be supported in the first period as part of an SPNE? The answer is yes, but we first need another concept.

	L	C	R
T	0, 0	3, 4	6, 0
M	4, 3	0, 0	0, 0
B	0, 6	0, 0	$\underbrace{5, 5}$

This is pareto

Figure 97: Pareto outcome for repeated game.

When analysing subgame-perfect Nash equilibria, there are many possible histories that might be generated in a multi-stage game with observed actions. This is because as we increase the number of stages, the number of histories grows exponentially. As a result, we would need to check deviations that might differ in any of those subgames. The following result will help us reduce the number of deviations that we need to check.

Definition 468 (One-Shot Deviation)
Given a strategy s_i, $\hat{s}_i$ is a **one-shot deviation** from s_i for player i if there is exactly one history h^t at which $\hat{s}_i(h^t) \neq s_i(h^t)$; at every other history, the two strategies agree.

Remark 469. *Thus, a one-shot deviation changes the prescribed action at one history only.*

It turns out that instead of comparing across all possible histories, we can in fact just check one-shot deviations.

Theorem 470 (One-deviation Principle)
In a finite multi-stage game with observed actions, a strategy profile $s = (s_i, s_{-i})$ is an SPNE if and only if no player i can gain from a one-shot deviation from s_i.

Proof. ($\Rightarrow$): If player i has a profitable one-shot deviation, the strategy is not subgame perfect, because player i has a profitable deviation in the subgame $G(h^t)$ beginning at the history where the one-shot deviation differs.
($\Leftarrow$): Suppose by contradiction that s is not an SPNE but there are no profitable one-shot deviations from s.
Since s is not a SPNE, then there exists a player i and a history h^t, such that s does not induce a NE at $G(h^t)$. In other words, conditional on being at h^t, $\hat{s}_i$ is a better response to s_{-i} than s_i. $(u_i(\hat{s}_i|h^t, s_{-i}|h^t) > u_i(s_i|h^t, s_{-i}|h^t))$.

1. Denote by t' the largest integer such that for some history $h^{t'}$, $s_i(h^{t'}) \neq \hat{s}_i(h^{t'})$. In other words, $\hat{s}_i$ agrees with s_i from $t'+1$ onwards. Note that $t' > t$ or otherwise there would be a one-shot profitable deviation starting from h^t.
2. Consider the alternative strategy $\bar{s}_i$ that agrees with $\hat{s}_i$ at all $t < t'$ and agrees with s_i from t' onwards. In particular, conditional on being at t', $\bar{s}_i$ (which coincides with s_i) and $\hat{s}_i$ only disagree in stage t'. Since there are no profitable one-shot deviations from s_i, $\bar{s}_i$ must do at least as well as $\hat{s}_i$ from that history. Because $\bar{s}_i$ and $\hat{s}_i$ coincide at all stages from t to $t'-1$, $\bar{s}_i$ is also a profitable deviation in $G(h^t)$.
3. We can now replace $\hat{s}_i$ by $\bar{s}_i$ and restart point 1. Repeating this argument eventually gives a profitable one-shot deviation, a contradiction.

□

In other words, to check whether s is an SPNE, we need only verify that no one-shot deviation $\hat{s}_i$ is a better response than s_i in any subgame $G(h^t)$. In terms of payoffs:

$$u_i(s_i|h^t, s_{-i}|h^t) \geq u_i(\hat{s}_i|h^t, s_{-i}|h^t).$$

We now apply it to our example Example 467.

Finite Repeated Game Continued

The payoffs are given by Fig. 97. Recall that the stage game has three Nash equilibria:

1. (T, C) where the payoffs are $u(T, C) = (3, 4)$
2. (M, L) where the payoffs are $u(M, L) = (4, 3)$
3. The mixed strategy $(\sigma_1 = (\frac{3}{7}, \frac{4}{7}, 0); \sigma_2 = (\frac{3}{7}, \frac{4}{7}, 0))$ with the payoff $u(\sigma_1, \sigma_2) = (\frac{12}{7}, \frac{12}{7})$.

We now show that the Pareto-optimal outcome (B, R) can be supported. We claim that the following strategy is an SPNE.

	L	C	R
T	0, 0	3, 4	6, 0
M	4, 3	0, 0	0, 0
B	0, 6	0, 0	5, 5

This is pareto

Figure 98: Pareto outcome for repeated game.

Claim 471. *The following strategy profile is a subgame-perfect Nash equilibrium:*

1. *P_1: Play B in the first period, If (B, R) is observed, play M. Otherwise, play σ_1*
2. *P_2: Play R in the first period, If (B, R) is observed, play L. Otherwise, play σ_2.*

Proof. To verify this, we can just check single deviations for each player at each stage.

In the final stage, there is no profitable deviation at the last stage since both potential strategy profiles (M, L) and (σ_1, σ_2) are Nash equilibria.

We now look at the first stage strategy (B, R). We only need to verify that the strategy does better than any other most profitable deviations for the players.

For P_1, they can deviate from playing B to playing T to earn a payoff of 6 in the first round. However, as a result, P_2 will punish in the second round and the (σ_1, σ_2) occurs, leading to a payoff of 12/7. Therefore, the payoffs between the one-shot deviation and strategy proposed is:

$$u_1(T, R) + u_1(\sigma_1, \sigma_2) = 6 + \frac{12}{7} < 5 + 4 = u_1(B, R) + u_1(M, L)$$

Hence no incentive to deviate for player 1.

For P_2, they could deviate from playing (B, R) and play L instead. Then, in the second stage, player 1 will punish by playing the "bad" strategy (σ_1, σ_2). Hence, the payoffs between the one-shot deviation and strategy proposed is:

$$u_2(B, L) + u_2(\sigma_1, \sigma_2) = 6 + \frac{12}{7} < 5 + 3 = u_2(B, R) + u_2(M, L)$$

Therefore no incentive to deviate for player 2.

Hence, the proposed strategy profile is indeed an SPNE in which (B, R) is played in the first period. □

Notice that (B, R) is **not** a stage-game NE, yet it is played in the first period of this SPNE.

In our example, multiple stage-game equilibria allowed outcomes that were not themselves Nash equilibria to be supported in an early period. Players could credibly threaten to switch to the **bad** Nash equilibrium after a deviation from the Pareto-superior path.

To summarise, the one-shot deviation principle simplifies the verification of SPNE. In a finite concatenated game, if every stage game has a unique NE, the whole game has a unique SPNE. With multiple stage-game equilibria, there may be multiple SPNE, some of which induce early-period actions that are not part of a stage-game NE.

We now turn to the final class of complete-information games. So far, we have studied repeated games and the more general case of finitely concatenated games, where stage games need not be identical. With a unique NE in every stage game, finite repetition yields a unique SPNE that plays those equilibria. With multiple stage-game equilibria, credible equilibrium-selection threats may support early-period actions that are not themselves stage-game equilibria.

In infinitely repeated games, even a stage game with a unique NE may support many different SPNE outcomes.

9.9 Repeated Games with Discounting

Sometimes players might interact for infinite number of periods. In order to analyse infinitely repeated interactions we will need to add some discounting to the payoffs.

Definition 472 (Average Discounted Payoff of Repeated Game)
Consider a normal-form game $\Gamma = [N, \{A_i\}, \{u_i\}]$. Players repeat the same game for T periods (potentially forever) and discount the future with a discount factor $\delta \in (0, 1)$. The payoff of the **finite** repeated game is the average discounted payoff:

$$U_i(s) = \frac{1-\delta}{1-\delta^T} \sum_{t=1}^{T} \delta^{t-1} u_i(a_s^t) \tag{189}$$

and the payoff of the **infinitely repeated** game is the average discounted payoff:

$$U_i = (1-\delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(a_s^t) \tag{190}$$

where $a_s^t \in A$ is the action profile induced in period t by the strategy profile s.

Example 473
(Payoff Example). Suppose we had a stage that pays off 7 at every stage. Then, the discounted payoff is:

$$\sum_{t=0}^{\infty} \delta^t 7 = \frac{7}{1-\delta}$$

However, this doesn't allow us to compare payoffs against stage games. Hence, we can normalise this by going:

$$(1-\delta) \sum_{t=0}^{\infty} \delta^t 7 = 7.$$

The one-deviation principle can be extended to infinitely repeated games with discounting. Therefore, to check if a given strategy is a subgame-perfect Nash equilibrium, it is enough to check against one-shot deviations.

Theorem 474 (One-Shot Deviation Principle for Infinitely Repeated Games)
In an infinite extensive-form game with discounting, observed actions, and bounded stage payoffs, a (potentially mixed) strategy profile σ is a subgame-perfect equilibrium if and only if no player has a profitable one-shot deviation after any history.

Proof. ($\Rightarrow$): Trivial.

($\Leftarrow$): We have shown before that if at some history h^t there is a profitable **finite** deviation, then there is a profitable one-shot deviation.

Consider a history h^t and a profitable deviation $\hat{s}_i$ that differs from s_i in an infinite number of stages. To simplify the notation, denote by $(\hat{a}_i^{t+\tilde{t}}, a_{-i}^{t+\tilde{t}})$ the action profile at stage $t + \tilde{t}$ induced by the strategy profile $(\hat{s}_i|h^t, s_{-i}|h^t)$.

Analogously, $a^{t+\tilde{t}}$ is the action profile at stage $t+\tilde{t}$ induced by the strategy profile $s|h^t$. Then

$$\sum_{\tilde{t}=1}^{\infty} \delta^{\tilde{t}-1} u_i(\hat{a}_i^{t+\tilde{t}}, a_{-i}^{t+\tilde{t}}) > \sum_{\tilde{t}=1}^{\infty} \delta^{\tilde{t}-1} u_i(a^{t+\tilde{t}})$$

Given the discounting, there exists a sufficiently large T, such that:

$$\delta^T \sum_{\tilde{t}=1}^{\infty} \delta^{\tilde{t}-1} \left[u_i(\hat{a}_i^{T+t+\tilde{t}}, a_{-i}^{T+t+\tilde{t}}) - u_i(a^{T+t+\tilde{t}}) \right] < \epsilon$$

In other words, the gains far in the future, (from T onwards) will be small, if we go sufficiently far. Therefore,

$$\sum_{\tilde{t}=1}^{T} \delta^{\tilde{t}-1} u_i(\hat{a}_i^{t+\tilde{t}}, a_{-i}^{t+\tilde{t}}) > \sum_{\tilde{t}=1}^{T} \delta^{\tilde{t}-1} u_i(a^{t+\tilde{t}})$$

and we can define $\tilde{s}_i$ such that

$$\tilde{s}_i^{\tilde{t}} = \begin{cases} \hat{s}_i^{\tilde{t}} & if \quad t \leq \tilde{t} < T+t \\ s_i^{\tilde{t}} & \quad otherwise \end{cases}$$

And $\tilde{s}$ is a profitable finite deviation from s_i conditional on being at h^t. □

Therefore, it is enough to check one-shot deviations from equilibrium candidates. However, this can still result in complicated analysis if we have an infinite strategy space.

9.10 Reverse Folk Theorem

In contrast with finitely repeated games, infinitely repeated games may have many payoffs achievable through SPNE, even when the stage game has a unique NE.

Definition 475 (Feasible Payoffs)
Given a stage game $\Gamma = [N, \{A_i\}, \{u_i\}]$, we define the **feasible** set of payoffs of game Γ as the convex hull of the payoffs achieved through pure strategy profiles,

$$co(\{u(a) \mid a \in A\}). \tag{191}$$

In a repeated game, these convex combinations can be generated by time-sharing across action profiles or by suitable public randomisation.

Infinitely Repeated Prisoner's Dilemma

We now consider an infinitely repeated Prisoner's Dilemma. The set of feasible payoffs is shown in Fig. 99.

We now want to ask the question on which payoffs in the convex hull can be supported in a subgame-perfect Nash equilibrium of an infinitely repeated game. In particular, the cooperation equilibrium (c, c) is a Pareto outcome but was not actually achieved when we previously analysed this game. However, we can show that under certain conditions, this equilibrium can be supported via credible threats.

Claim 476. *The co-operation strategy (c, c) can be supported in a subgame-perfect Nash equilibrium under an infinitely repeated game if agents are patient enough $\delta > \frac{1}{2}$.*

	c	d
c	3, 3	0, 5
d	5, 0	1, 1

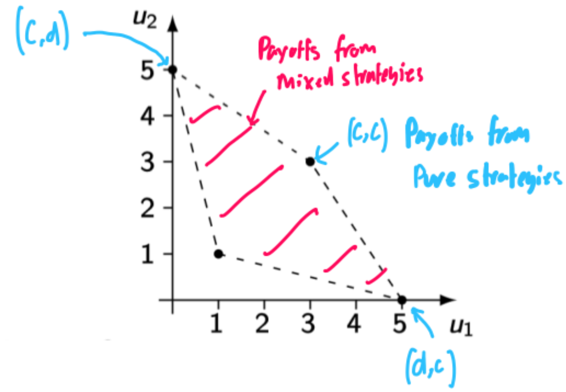

Figure 99: Set of feasible payoffs for the Prisoner's Dilemma.

Proof. We propose that the following **trigger strategy** will support this equilibrium. For player i, play the action c. If (c, c) has been played in **all previous periods**, then continue to play c. Otherwise, play d. So the trigger strategy for player i is:

$$s_i = \begin{cases} c & \text{if (c,c) has been played in all previous periods} \\ d & \text{otherwise} \end{cases}$$

We now check whether this trigger strategy supports (c, c) as a subgame-perfect Nash equilibrium. By the **one-shot deviation principle**, it is sufficient to check one-shot deviations from the strategy profile.

First, we note that even if the strategies are infinite, there are only two distinct ways of playing since there are only two groups of histories:

1. $s_i(h^t) = c$ if $h^t = ((c, c), (c, c), ...)$ or $h^t = h^1$. That is, no one has ever deviated before.

2. $s_i(h^t) = d$ if someone has deviated at some point.

Therefore, we only have two types of subgames where either no one has ever deviated in the history or a subgame where someone has deviated at some point in the past. We now compare the payoffs between these two subgames.

Subgame 1: Let us look at the subgame where that no one has ever deviated in the past. If a player continue to comply c for the rest of the game, then their payoff is the discounted sum of the payoff from (c, c) which is:

$$3 + 3\delta + 3\delta^2 + ... = \sum_{t=0}^{\infty} \delta^t 3 = \frac{3}{1-\delta}$$

We now analyse the one-shot deviation. If a player decides to deviate, the highest payoff deviation is if they deviate in the immediate turn (since no discounting). Then, they will earn a payoff of 5 from deviating. However, they will then be punished by the outcome (d,d) for the rest of the game. Therefore, the payoff from this one-shot deviation strategy is:

$$5 + 1\delta + 1\delta^2 + ... = 5 + \sum_{t=1}^{\infty} \delta^t = 5 + \frac{\delta}{1-\delta}$$

So we now compare the payoffs from the trigger strategy and the most profitable one-shot deviation in this subgame:

$$\underbrace{3 + 3\delta + 3\delta^2 + \dots}_{\text{Comply}} \geq \underbrace{5 + 1\delta + 1\delta^2 + \dots}_{\text{Deviate}}$$
$$\frac{3}{1-\delta} \geq 5 + \frac{\delta}{1-\delta}$$
$$\delta \geq \frac{1}{2}$$

Therefore, we can see that complying is better than deviating if the discount factor is greater than $\frac{1}{2}$. This says that if agents are patient enough and care about being punished, then they have no incentive to deviate.

Subgame 2: Now let us look at the subgame where someone has defected in the past. Then, by our trigger strategy, agents should play d and hence the NE (d,d) arises. The payoff from complying from playing d every round is:

$$1 + 1\delta + 1\delta^2 + \dots = \sum_{t=0}^{\infty} \delta^t = \frac{1}{1-\delta}.$$

Now, let us look at the one-shot deviation from playing d every turn. If an agent deviates and instead plays c, whilst the other player still plays d, then the deviator will earn a payoff of 0. They then go back to playing d every turn afterwards and earn a payoff of 1. Therefore, the payoff from a one-shot deviation is:

$$0 + 1\delta + 1\delta^2 + \dots = 0 + \frac{\delta}{1-\delta}$$

Hence, we can clearly see that complying and playing d earns a higher payoff than deviating:

$$\underbrace{\frac{1}{1-\delta}}_{\text{Comply}} > \underbrace{0 + \frac{\delta}{1-\delta}}_{\text{Deviate}}.$$

Therefore, the player should comply with the strategy.

Therefore, we have shown that for $\delta > \frac{1}{2}$, we can support the outcome (c, c) through our trigger strategy to earn an average expected payoff of $(3, 3)$. $\square$

If individuals are patient enough, we can support the Pareto outcome (c, c) through a trigger strategy!

We have shown that a Pareto-superior outcome that is not a stage-game NE can be supported when players are sufficiently patient. The next theorem generalises this result to feasible payoffs that strictly exceed a stage-game equilibrium payoff for every player.

Theorem 477 (Nash Reversion Folk Theorem - Friedman 1971)
Let s^* be a NE of the stage game Γ, with payoff profile u^*. For any feasible payoff profile u such that $u_i > u_i^*$ for every $i \in N$, there exists $\underline{\delta} \in (0, 1)$ such that, for every $\delta \in (\underline{\delta}, 1)$, an **SPNE** of the infinitely repeated game has payoff u.

Remark 478. *This says that feasible payoffs greater than the NE for all players of the stage game can be supported provided that agents are patient enough.*

Intuitively, players follow a path that generates the target feasible payoff unless someone deviates, after which all players revert forever to the stage-game NE. If players are patient enough, any one-period gain from deviation is outweighed by the future losses from reversion.

Example 479

(Repeated Prisoner's Dilemma). The stage-game NE is (D, D), with payoff $(1, 1)$.

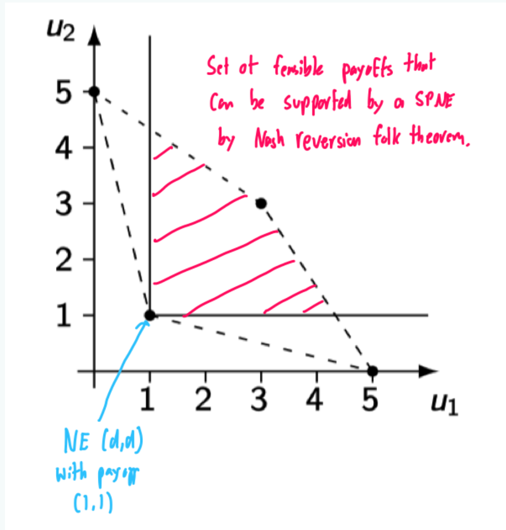

Figure 100: Set of feasible payoffs supportable by an SPNE in the Prisoner's Dilemma.

By the Nash-reversion folk theorem, Fig. 100 shows the set of feasible payoffs that can be supported by an SPNE. The target payoff must exceed the stage-game NE payoff for **every player**.

Note that these strategies are SPNE since both the equilibrium path and the potential punishment if someone deviates are NE in their respective subgames.

9.10.1 Collusion

We now apply our ideas to collusion. For a collusion to be achievable for firms, there should be no incentive to deviate from the collusion. This collusion should also be supported by credible threats.

Example 480

(Collusion in Bertrand Competition). Consider n identical firms with constant marginal c and that these firms compete on prices. Recall that the **unique NE** of the stage is for firms to set prices equal to marginal cost $p_i = c$. Therefore, by our earlier theorem, if firms compete for a finite number of periods, the only competitive outcome that can be supported as a subgame-perfect Nash equilibrium is for firms to set prices equal marginal cost every period $p_i = c$.

However, what happens if firms play for an infinite period? Then by the Nash Reversion folk theorems, more outcomes can be sustained. Firms can employ trigger strategies and then by the Nash reversion folk theorem, outcomes besides the Nash equilibrium of $p_i = c$ can be sustained.

Definition 481 (Trigger Strategies in Bertrand Competition)
The following is a **trigger strategy** for firms in the Bertrand competition model:

$$p_j^t(h^t) = \begin{cases} p^m & \text{if } h^t = \{a^1, ..., a^{t-1}\} \text{ where } a^r = (p^m, ..., p^m) \\ c & \text{otherwise} \end{cases} \tag{192}$$

where p^m is the monopoly price, c is the marginal cost, and p_j^t is the price that firm j should set at period t.

This trigger strategy says that a firm should play the monopoly price if in all previous history, everyone set their prices to the monopoly price p^m. Otherwise, if someone deviated from the monopoly price, the firm should set price equal to marginal cost c.

Theorem 482 (SPNE for the Infinitely Repeated Bertrand Model)
The trigger strategies for the Bertrand model constitute a subgame-perfect Nash equilibrium of the infinitely repeated game if and only if the discount factor is sufficiently high:

$$\delta \geq \delta^b \equiv 1 - \frac{1}{n}.$$

The trigger strategy can support the monopoly price p^m provided that firms are sufficiently patient. The threshold is $\delta^b \equiv 1 - \frac{1}{n}$. As the number of firms rises, this threshold approaches 1, so collusion becomes harder—but not impossible—to sustain.

We now prove where this threshold comes from.

Proof. By the one-shot deviation principle, we need to consider only two classes of histories: those after which no firm has deviated and those after which at least one firm has deviated.

If no one has previously deviated, a firm can undercut the monopoly price by an arbitrarily small amount and obtain a payoff arbitrarily close to the full monopoly profit π^m. Thereafter, the trigger strategies prescribe marginal-cost pricing, which yields zero profit. The limiting deviation payoff is therefore:

$$\pi^m + 0\delta + 0\delta^2 + ... = \pi^m + 0 = \pi^m$$

If they comply and set their price equal to the monopoly price, then they get a fraction of the total profits π^m/n where n is the number of firms setting their prices equal to the monopoly price. Therefore, the payoff from continually complying is:

$$\frac{\pi^m}{n} + \frac{\pi^m}{n}\delta + \frac{\pi^m}{n}\delta^2 + ... = \frac{1}{1-\delta}\frac{\pi^m}{n}.$$

Therefore, deviation is **not profitable** if the payoff from complying is greater than deviation:

$$\underbrace{\frac{1}{1-\delta}\frac{\pi^m}{n}}_{\text{Comply}} \geq \underbrace{\pi^m}_{\text{Deviate}}$$

Re-arranging this, we get the discount threshold:

$$\delta \geq 1 - \frac{1}{n} \equiv \delta^b.$$

In the subgame where someone has previously deviated, no one can gain from moving from $p = c$ as it is already a NE of the stage game. So the player should continue to play the trigger strategy.

Therefore, we have shown that the Bertrand trigger strategy is an SPNE. $\square$

Hence, the minimum discount factor δ^b rises with the number of firms. Each firm receives only π^m/n per collusive period, while a successful one-shot undercut captures almost the entire monopoly profit. As n increases, the foregone collusive share becomes smaller relative to the deviation gain, so greater patience is required.

We now look at Cournot model that is infinitely repeated.

Example 483

(Collusion in Cournot Competition). Suppose there are n identical firms with constant marginal cost c, competing in quantities q_i. Inverse demand is $p(Q) = a - bQ$, where Q is aggregate output and $a > c$.

We know that the monopolistic quantity under Cournot is given by $q = \frac{a-c}{2b}$. The Cournot equilibrium was also characterised by the first order conditions:

$$\frac{Q^c}{n} p'(Q^c) + p(Q^c) = c.$$

See Example 401 for a reminder of these results. We then re-arrange this FOC to get:

$$\frac{Q^c}{n} = -\frac{p(Q^c) - c}{p'(Q^c)}$$

We see that the Cournot quantity produced by a single firm is the Cournot aggregate divided by the number of firms:

$$q^c = \frac{Q^c}{n} = -\frac{p(Q^c) - c}{p'(Q^c)} = \frac{1}{b}\left(\frac{a-c}{n+1}\right)$$

where we plugged in the inverse demand function. This then implies that the profit each firm will get is:

$$\pi^c = \frac{1}{b}\left(\frac{a-c}{n+1}\right)^2$$

We now define the trigger strategy in the Cournot competition.

Definition 484 (Trigger Strategies in Cournot Competition)

The following is a **trigger strategy** for firms in the Cournot competition model:

$$q_j^t(h^t) = \begin{cases} \frac{q^m}{n} & \text{if } h^t = \{a^1, \ldots, a^{t-1}\} \text{ where } a^r = (\frac{q^m}{n}, \ldots, \frac{q^m}{n}) \\ q^c & \text{otherwise} \end{cases} \tag{193}$$

where q^m is the monopoly output, q^c is the Cournot equilibrium output, and q_j^t is the quantity that firm j should set at period t.

We can now state when these strategies form an SPNE of the infinitely repeated Cournot model.

Theorem 485 (SPNE for the Infinitely Repeated Cournot Model)

The trigger strategies for the Cournot model constitute a subgame-perfect Nash equilibrium of the infinitely repeated game if and only if the discount factor is sufficiently high:

$$\delta \geq \delta^c \equiv 1 - \frac{4n}{(n+1)^2 + 4n}.$$

Proof. If no one deviated before, the optimal deviation for a firm would be to respond to the total quantity produced by **other** firms $(n-1)\frac{q^m}{n}$ by producing the deviating amount:

$$q_i^d = B_i((n-1)\frac{q^m}{n}) = \frac{(a-c)(n+1)}{4bn}$$

Then, the total quantity in the market will be:

$$Q^d = (n-1)\frac{q^m}{n} + \frac{(a-c)(n+1)}{4bn} = \frac{(a-c)}{4bn}(3n-1).$$

Then, by the linear demand function, this leads to a price:

$$p^d = a - \frac{(a-c)(3n-1)}{4n}.$$

Hence, the deviating firm achieves a payoff that is given by:

$$\pi_i^d = (p^d - c)q_i^d = \frac{1}{b}\left(\frac{(a-c)(n+1)}{4n}\right)^2$$

Deviation is not profitable if the net present value of monopolistic profits is greater than the profits from deviation plus competitive outcome profits:

$$\frac{1}{1-\delta}\frac{\pi^m}{n} \geq \pi_i^d + \frac{\delta}{1-\delta}\pi_i^c.$$

By algebra, it can be shown that this is equivalent to:

$$\delta \geq \delta^c \equiv \frac{(n+1)^2}{(n+1)^2+4n} = 1 - \frac{4n}{(n+1)^2+4n}.$$

In the subgame where someone has deviated and everyone produces the Cournot output q^c, then there is no incentive to deviate as this is a Nash equilibrium. □

We see that as $n \to \infty$, there is also less incentive to collude in the Cournot model.

We now look at the discount thresholds in the Bertrand and Cournot model:

$$\delta^b \equiv 1 - \frac{1}{n}.$$

$$\delta^c \equiv 1 - \frac{4n}{(n+1)^2+4n}.$$

Thus, the strategic variable—price or quantity—affects the viability of collusion. For $n \geq 3$, $\delta^b > \delta^c$, so this trigger-strategy construction requires greater patience under Bertrand competition. In Bertrand, the collusive share falls with n while the one-shot gain from undercutting remains close to the full monopoly profit. Under Cournot competition, the collusive payoff, deviation payoff, and punishment payoff all vary with n.

Under Nash-reversion strategies, the sets of supportable payoffs differ between the two models, as illustrated in Fig. 101. The figure suggests that this construction supports a smaller payoff region in the Cournot model. We next ask whether payoffs **below** the Cournot equilibrium in the figure can also be supported; answering this requires the more general folk theorem.

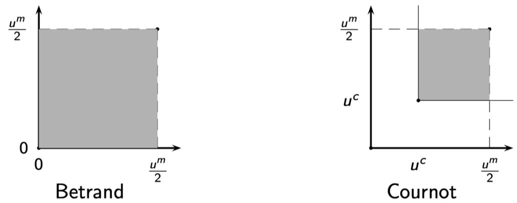

Figure 101: Shaded area represents feasible payoffs achievable through SPNE via collusion.

9.11 Folk Theorem

In fact, we can achieve even more payoffs through SPNE strategies. First, we define what we mean by a minmax payoff.

Definition 486 (Minmax Payoffs)
Given the stage game Γ, the **minmax payoff** for player i is the lowest best-response payoff that the other players can impose on player i.

$$\underline{u}_i = \min_{\sigma_{-i}} \max_{a_i} u_i(a_i, \sigma_{-i}). \tag{194}$$

This is the worst punishment for player i imposed by all other players.

Example 487
(Battle of the Sexes). We compute the minmax payoffs in the battle of the sexes shown in Fig. 102. If player 2 plays L with probability q, player 1 obtains $2q$ from U and $1-q$ from D. Player 2 minimises player 1's best-response payoff by setting $2q = 1 - q$, so $q = 1/3$ and

$$\underline{u}_1 = \frac{2}{3}.$$

Symmetrically, player 1 mixes so as to make player 2 indifferent, giving

$$\underline{u}_2 = \frac{2}{3}.$$

Restricting punishments to pure strategies would instead give the higher value 1, but the minmax definition permits mixed strategies.

	L	R	
U	2, 1	0, 0	1
D	0, 0	1, 2	2
	2	1	

Figure 102: Minmax payoffs for battle of the sexes.

Example 488

(Minmax payoffs with mixed strategies). We now look at the minmax payoffs for Fig. 103. The minmax payoff for player 1 is $\underline{u}_1 = 1$ under the strategy profiles (M, L) and (U, R). The minmax payoff for player 2 is $\underline{u}_2 = 1$ under the strategy profiles $(D, L), (D, R)$.

However, we can extend minmax payoffs to mixed strategy. If player 2 randomises with probability q on L, then the payoffs are:

$$u_1(U, q) = -2q + 1(1-q) = 1 - 3q$$

$$u_1(M, q) = 1q + (-1)(1-q) = 2q - 1$$

$$u_1(D, q) = 0q + 0(1-q) = 0$$

Therefore, if player 2 mixes with probability $q \in [\frac{1}{3}, \frac{1}{2}]$, then they can impose a payoff of $\underline{u}_1 = 0$ on player 1.

	L	R	
U	-2, 2	1, -2	2
M	1, -2	-1, 2	2
D	0, 1	0, 1	1
	1	1	

Figure 103: Minmax payoffs with mixed strategies.

Given a minmax payoff, we can define payoffs that are greater than the minmax payoff.

Definition 489 (Individually Rational)

The stage game payoff u_i for player i is **individually rational** if $u_i > \underline{u}_i$ where $\underline{u}_i$ is the minmax payoff for player i.

We now state an important result that allows us to expand the set of supported payoffs through SPNE.

Theorem 490 (Folk Theorem—Fudenberg and Maskin (1986))

Consider a finite n-player stage game with perfect monitoring. Suppose that the feasible payoff set has full dimension. For any feasible payoff vector v satisfying $v_i > \underline{u}_i$ for every player i, there exists $\underline{\delta} \in (0, 1)$ such that, for every $\delta \in (\underline{\delta}, 1)$, the infinitely repeated game has a subgame-perfect equilibrium with payoff vector v.

Proof. See FT- CH5 (Theorem 5.4). □

Thus, any feasible and strictly individually rational payoff can be supported by an SPNE when players are sufficiently patient. Unlike simple Nash reversion, imposing a deviator's minmax payoff forever need not be credible because punishment may also be costly for the punishers. The construction therefore uses continuation payoffs that reward players for carrying out prescribed punishments. A player who refuses to punish is treated as the next deviator and loses that future reward. This combination of punishment and continuation rewards makes the entire strategy sequentially rational.

We now go back and apply this to the Cournot model.

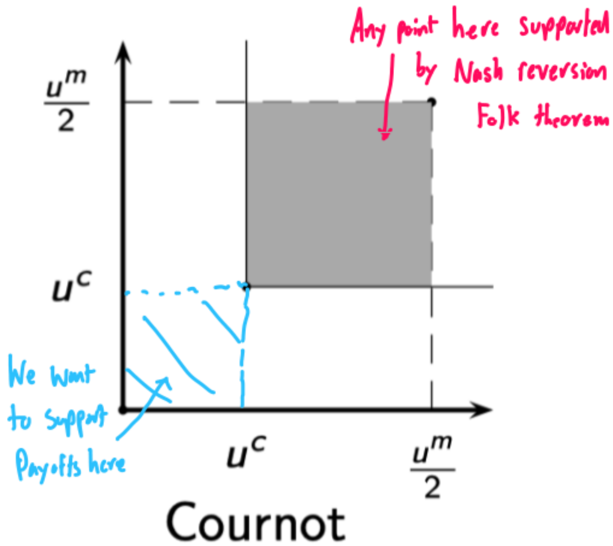

Figure 104: Payoffs in Cournot model.

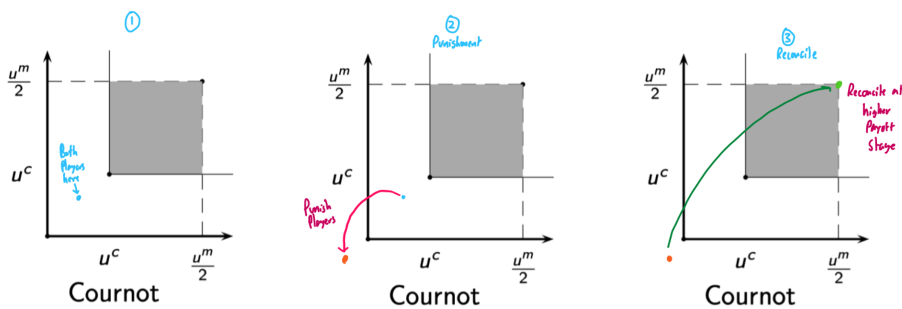

Figure 105: Cournot model where payoffs below Cournot equilibrium are supported by SPNE.

Example 491

(Folk Theorem Applied to Cournot Model).

Recall the Cournot payoff set in Fig. 104. The strategy illustrated in Fig. 105 begins with both firms producing below the Cournot quantity. Following a deviation, play moves to a punishment phase with lower continuation payoffs for the deviator. The firms then enter a reconciliation phase that rewards those who carried out the punishment. When the discount factor is sufficiently high, these continuation incentives can support the target feasible and individually rational payoff.

We have used repeated games to explain how an oligopoly can sustain collusive outcomes and how more nuanced strategies can support a large set of payoffs.

10 Games of Incomplete Information

10.1 Weak Perfect Bayesian Equilibrium

So far we have been dealing with games of perfect information or with multi-stage games with observed actions. In those games, the concept of subgame perfection was successful in eliminating the Nash equilibria that were not sequentially rational, or in other words, that relied on non-credible threats.

However, sometimes the amount of imperfect information is substantial, for example in some cases the unique subgame of a game is itself. In those cases, subgame perfection does not have any bite. Still, we would like to rule out strategies that induce non-optimal play in some information sets. But how could we determine what is an optimal action in an information set?

Example 492

(Incumbent Firm Battle). Consider the game in Fig. 106. The two decision nodes of firm I belong to the same information set, so neither can begin a proper subgame. Hence, the only subgame is the whole game, and every Nash equilibrium is also subgame perfect. The figure also gives the normal-form representation, from which we obtain:

$$(\text{Out}, \text{F}) \quad (\text{In}_1, \text{A})$$

are pure strategy NE whilst

$$\left(\text{Out}, \sigma_2(F) = p\right) \quad \text{with } p \geq \frac{3}{4}$$

describes the remaining Nash equilibria of the game.

However, by sequential rationality, if firm I is ever called to play, it should always play A, since A gives it a higher payoff than F at either node. Applying backward induction, since firm E knows that firm I will always play A, firm E should play In_1, as $(\text{In}_1, A) \succ (\text{In}_2, A)$. Therefore, the Nash equilibrium (Out, F) is supported by behaviour that is not sequentially rational. We therefore need to strengthen our equilibrium concept.

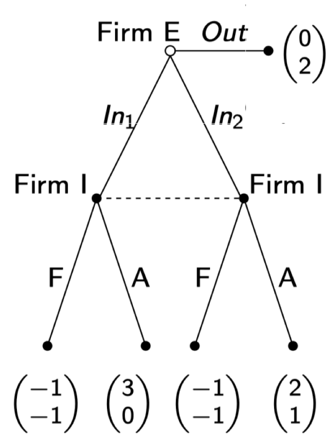

	F	A
Out	$\underline{0}, \underline{2}$	$0, \underline{2}$
In_1	-1, -1	$\underline{3}, \underline{0}$
In_2	-1, -1	$2, \underline{1}$

Figure 106: Entrant vs incumbent game. Extensive form representation on the left and normal form representation on the right.

In order to generalise the concept of sequential rationality to games with imperfect information, we introduce beliefs in information sets. With these beliefs, we require players to behave optimally at each information set given those beliefs.

Definition 493 (System of Beliefs)
A **system of beliefs** in an extensive form game Γ_E specifies a probability distribution μ for each information set I, where $\mu(x) \geq 0$ and $\sum_{x \in I} \mu(x) = 1$.

The term $\mu(x)$ represents the probability that the player who moves at information set $I(x)$ allocates to node x.

Definition 494 (Assessment)
An **assessment** is a strategy profile **and** a system of beliefs (σ, μ).

We will now work with assessments rather than strategy profiles. We require two conditions on these assessments in order to capture sequential rationality. Intuitively, the first says that the strategy σ must be optimal given the beliefs μ. The second says that the beliefs must be consistent with the play induced by σ.

Definition 495 (Expected Utility at an Information Set with Beliefs)
Denote by $E[u_i \mid I_i, \mu, \sigma_i, \sigma_{-i}]$ the **expected utility** for player i starting at their information set I_i, if her beliefs regarding the conditional probabilities of each node in the information set are given by μ, and the strategy profile is (σ_i, σ_{-i}).

Remark 496. *This is the expected payoff for player i given their beliefs and strategy profile.*

We now define the conditions needed on beliefs and strategies.

Definition 497 (Sequentially Rational Strategies)
A behavioural strategy profile σ is **sequentially rational given a system of beliefs** μ if for every player i and any information set I_i of player i, we have that

$$E[u_i \mid I_i, \mu, \sigma_i, \sigma_{-i}] \geq E[u_i \mid I_i, \mu, \sigma_i', \sigma_{-i}] \text{ for all } \sigma_i' \in \Delta(S_i). \tag{195}$$

The above definition of sequential rationality captures the idea that the player should choose her actions optimally given the beliefs she holds at every information set. We will also require the system of beliefs to be consistent with the strategy profile.

Definition 498 (Consistent Beliefs)
The system of beliefs μ is **consistent with the strategy profile** σ if for any information set I reached with positive probability $(Pr(I \mid \sigma) > 0)$ we have

$$\mu(x) = \frac{Pr(x \mid \sigma)}{Pr(I \mid \sigma)} \text{ for all } x \in I. \tag{196}$$

This says that the beliefs $\mu(x)$ should be equal to the probability of reaching node x through strategy σ, $Pr(x|\sigma)$, divided by the probability of reaching the information set I that contains node x through the strategy σ, denoted by $Pr(I|\sigma)$. This means that if an information set I is reached through the strategies, then the belief $\mu(x)$ for a node $x \in I$ should agree with the probability that the strategies induce on that node conditional on reaching the information set. In other words, we can use Bayes' rule at those information sets to determine the beliefs.

From now on, instead of working with strategy profiles, we will be working with *assessments* (σ, μ), i.e., pairs of strategy profiles and system of beliefs. Intuitively we will require the strategies and the beliefs in an assessment to be intrinsically connected in the sense that the strategies should be sequentially rational given the beliefs, and the beliefs should be consistent with the strategies. This is what we call a weak Perfect Bayesian Equilibrium.

Definition 499 (Weak Perfect Bayesian Equilibrium)
An assessment (σ, μ) is a **weak perfect Bayesian equilibrium (weak PBE)** if:

1. The strategy profile σ is **sequentially rational** given the system of beliefs μ.
2. The system of beliefs μ is **consistent with the strategy profile** σ.

Note that the consistency requirement is imposed only on information sets that are reached through the strategy profile. However, even in those information sets that are not reached, there should be a belief, and the strategies should best respond to it.

Proposition 500 (Weak PBE and NE)
If an assessment (σ, μ) is a weak PBE, then σ is a Nash equilibrium.

Proof. Along the equilibrium path, beliefs are completely pinned down by Bayes' rule, and the strategies must be best responding to what the others do in equilibrium. □

Remark 501. *However, the reverse is not true. A NE can be supported by a threat of doing something that is never optimal (there is no belief that could justify it), and this would not satisfy the sequential rationality condition in the definition.*

Example 502
(Incumbent Firm Battle Continued). We return to the game in Example 492. Working backwards, for any belief $\mu \in [0, 1]$, firm I should always play A because A is a dominant strategy. Firm E's decision then looks like Fig. 108, where it should play In_1 because this gives it a higher payoff. Given the strategy profile (In_1, A), the only **consistent belief** for firm I is $\mu = 1$: firm I assigns probability one to the leftmost node. Therefore, the assessment $\{(\text{In}_1, A), \mu = 1\}$ is the only assessment in which (In_1, A) is sequentially rational and the beliefs are consistent.

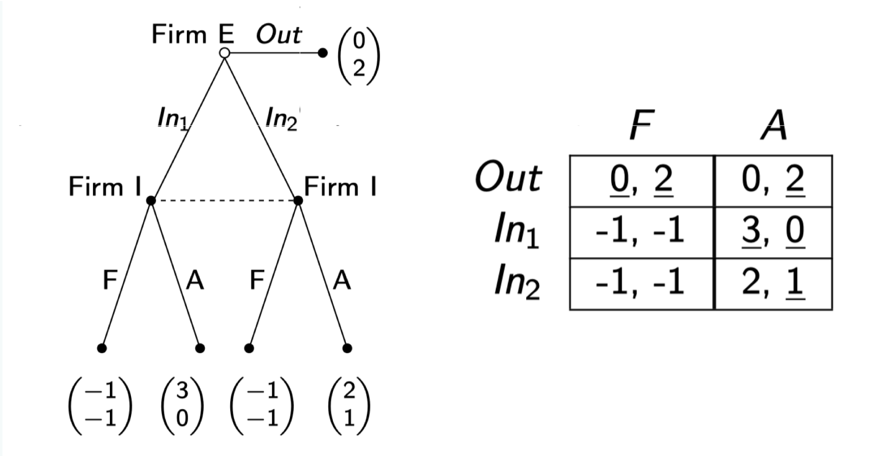

	F	A
Out	0, 2	0, 2
In_1	-1, -1	3, 0
In_2	-1, -1	2, 1

Figure 107: Entrant vs incumbent game. Extensive form representation on the left and normal form representation on the right.

The assessment $\{(\text{In}_1, A), \mu = 1\}$ is a weak perfect Bayesian equilibrium. Notice that we have ruled out the Nash equilibrium (Out, F), as desired. Meanwhile, we have shown that the Nash equilibrium (In_1, A) satisfies the stronger weak-PBE requirements.

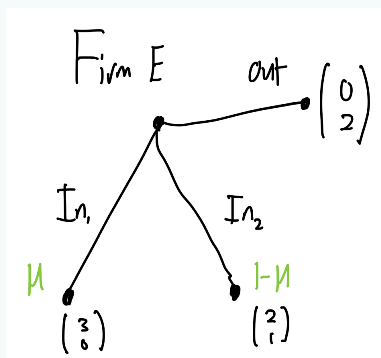

Figure 108: Entrant vs incumbent game. Choice for firm E.

This example was quite simple because firm I had a dominant strategy to accommodate. We can apply weak perfect Bayesian equilibrium even without a dominant strategy.

Modified Incumbent Firm Battle

Suppose that we have a modified incumbent firm battle as in Fig. 109.

We start off at the information set of Firm I. First, we see that for firm I, playing F as a Best Response function only holds if:

$$\mu(-1) + (1-\mu)(-1) \geq \mu(-2) + (1-\mu)(1)$$

which says that F gives at least as much expected payoff as A under belief μ. Solving the inequality, F is a best response if and only if:

$$\mu \geq \frac{2}{3}.$$

From this, we have 3 different cases to consider depending on what value μ is.

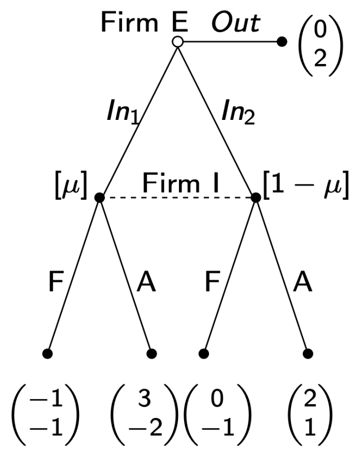

Figure 109: Modified entrant vs incumbent game's extensive form representation.

$\mu > \frac{2}{3}$

When $\mu > 2/3$, firm I plays F with probability 1. We then move up the tree to consider firm E's choice, shown in Fig. 110. Firm E is indifferent between Out and In_2, since it receives a payoff of 0 from either action. This leads to three possible cases.

Suppose that firm E plays In_2 with positive probability and places zero probability on In_1. This implies that the first node in firm I's information set is never reached, so $\mu = 0$ is the only consistent belief. However, this contradicts the assumption that $\mu > 2/3$, so this cannot be the case. Hence, firm E will not put any probability on playing In_2.

Alternatively, if firm E chooses Out with probability 1, there are no restrictions on beliefs at firm I's information set, so any pair $(\mu, 1-\mu)$ is allowed. Setting $\mu \geq 2/3$ makes F sequentially rational for firm I. Hence, the assessment

$$\{(Out, F); (\mu, 1-\mu), \text{for any } \mu \geq \frac{2}{3}\}$$

is a **weak Perfect Bayesian Equilibrium**.

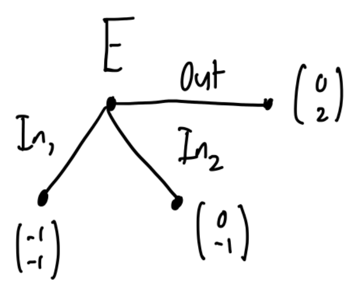

Figure 110: Modified entrant vs incumbent game's extensive form representation for $\mu > 2/3$.

$\mu < \frac{2}{3}$

When $\mu < 2/3$, firm I will play A with probability 1. We show this in Fig. 111. Clearly, firm E will play In_1, as it gives it the highest payoff and dominates the other two strategies, In_2 and Out. This then implies that $\mu = 1$. However, this contradicts the assumption that $\mu < 2/3$. Hence, this is not possible.

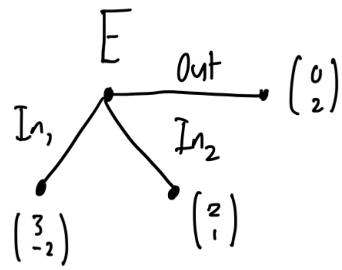

Figure 111: Modified entrant vs incumbent game's extensive form representation for $\mu < 2/3$.

$\mu = \frac{2}{3}$

When $\mu = 2/3$, firm I is indifferent between playing A and F and may mix between the two strategies. Let firm I play F with probability q. We draw this in Fig. 112. Looking at firm E's decision, firm E will play In_1 rather than In_2 if and only if:

$$(-1)q + (3)(1-q) > (0)q + (2)(1-q)$$

depending on firm I's mixed strategy. Solving for this, firm E will play In_1 if and only if firm I plays F with probability

$$q < \frac{1}{2}.$$

This implies that firm E will play In_2 rather than In_1 if $q > 1/2$, and will be indifferent between them if $q = 1/2$. The first two scenarios cannot occur in equilibrium: playing In_1 for sure implies $\mu = 1$, while playing In_2 for sure implies $\mu = 0$, both contradicting $\mu = 2/3$. Therefore, firm E must mix between In_1 and In_2. It assigns probabilities 2/3 and 1/3, respectively, while firm I assigns probability 1/2 to each of F and A. Finally, we can verify that this mixed strategy gives firm E more than playing Out:

$$2(1-q) = 2(0.5) = 1 > 0.$$

Therefore, we have shown that the following assessment of mixed strategies:

$$\left\{ \left(\sigma_1(\text{Out}) = 0,\ \sigma_1(\text{In}_1) = \frac{2}{3},\ \sigma_1(\text{In}_2) = \frac{1}{3};\ \sigma_2(F) = \frac{1}{2},\ \sigma_2(A) = \frac{1}{2} \right);\ \mu = \frac{2}{3} \right\}$$

is a **weak Perfect Bayesian Equilibrium**.

Selten's Horse

We look at another game called Selten's Horse. We represent it in Fig. 113. Looking at the normal form, we can deduce that the two pure Nash equilibria are:

$$(D, a, L) \quad (A, a, R)$$

Furthermore, since there are no subgames, these Nash equilibria are also subgame-perfect Nash equilibria. However, if we look at the equilibrium (D, a, L), we see that there is an issue with this being an equilibrium. It is not sequentially rational for player 2 to play a and should instead play d. This is because in this equilibrium, player 2 knows that player

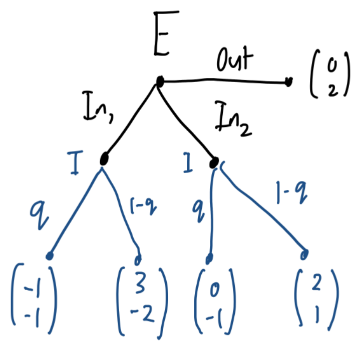

Figure 112: Modified entrant vs incumbent game's extensive form representation for $\mu = 2/3$.

3 will play L. Therefore, the threat of player 3 playing R is not a credible threat. Hence, player 2 should play d if they know that player 3 will play L.

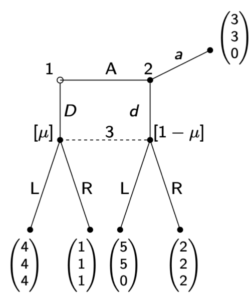

L

	a	*d*
A	3, 3, $\underline{0}$	$\underline{5}$, $\underline{5}$, 0
D	$\underline{4}$, $\underline{4}$, $\underline{4}$	4, $\underline{4}$, $\underline{4}$

R

	a	*d*
A	$\underline{3}$, $\underline{3}$, $\underline{0}$	$\underline{2}$, 2, $\underline{2}$
D	1, $\underline{1}$, 1	1, $\underline{1}$, 1

Figure 113: Selten's Horse game with the extensive form on the left and the normal form on the right.

First, let us start player 3's information set. We see that playing L is a best response for player 3 compared to playing R if and only if:

$$(4)\mu + (0)(1-\mu) \geq (1)\mu + (2)(1-\mu)$$

which depends on which node is player 3 at in the information set, which we denoted by μ. Solving for this, we get that playing L is a best response for player 3 if and only if:

$$\mu \geq \frac{2}{5}.$$

Now, we move up the tree and look at player 2's options. We have three possible scenarios depending on the value of μ which we need to analyse.

$\mu > \frac{2}{5}$

Suppose that $\mu > 2/5$. Then, this means that player 3 will play L for sure. Then, for player 2, they should play d as this gives them a payoff of 5 compared to a payoff of 3 if they played a. Then, for player 1, they should play A as this will earn them a payoff of 5 compared to a payoff of 4 if they played D (since player 3 will play L). Given these strategies, the only consistent belief is when $\mu = 0$ as the first node in player 3's information set is never reached. However, this contradicts our assumption that $\mu > 2/5$. Therefore, we rule out this scenario.

$\mu < \frac{2}{5}$

Suppose that $\mu < 2/5$. Player 3 then plays R for sure, so player 2 plays a and earns 3 rather than the payoff 2 from d. Player 1 plays A, earning 3 rather than the payoff 1 from D. Given these strategies, neither node in player 3's information set is reached, so beliefs are unrestricted. We may therefore choose any $\mu < 2/5$, giving the assessment:

$$\{(A, a, R); \mu < \frac{2}{5}\}$$

is a **weak Perfect Bayesian Equilibrium**.

$\mu = \frac{2}{5}$

Suppose that $\mu = 2/5$. Then, player 3 is indifferent to playing L or R. Suppose that player 3 plays L with probability q.

For player 2, playing a gives a payoff of 3, while playing d gives $(5)q + 2(1 - q)$. Therefore, player 2 will choose d if and only if:

$$(5)q + 2(1 - q) \geq 3$$

Solving for q, d is a best response function for player 2 if and only if:

$$q \geq \frac{1}{3}.$$

Now looking at player 1, suppose that $q < 1/3$, so then player 2 will play a. Then, if player 1 chooses to play D, they will earn a payoff of $(4)q + (1)(1 - q)$ since player 3 mixes between L and R with probability q. If player chooses to play A, then they will earn a payoff of 3 as player 2 plays a if $q < 1/3$. Therefore, player 1 will play D if and only if:

$$(4)q + (1)(1 - q) \geq 3$$

which gives us:

$$q \geq \frac{2}{3}.$$

This condition contradicts $q < 1/3$, so D cannot be optimal in this case.

Therefore, it is optimal for player 1 to play A and for player 2 to play a. Player 3's information set is then not reached, so we may set $\mu = \frac{2}{5}$. Player 3 can mix with $q \leq \frac{1}{3}$ while a remains a best response for player 2. Thus, another assessment is:

$$\{(A, a, \sigma_3(L) = q);\ q \leq \tfrac{1}{3};\ \mu = \tfrac{2}{5}\}.$$

10.2 Sequential Equilibrium

The concept of weak PBE is only an intermediate step towards the concept of Sequential Equilibrium. It provides the key intuition of the relationship between beliefs and strategies. However it has some flaws. The main issue is that there is not enough structure on the beliefs as any information set that is not reached in equilibrium can have any beliefs over it. We will give two examples of why this is bad. The first one we shall show is that beliefs in some information may not be consistent with the structure of the game, for example with the probabilities induced by Nature.

Example 503

(Structural Consistency). Consider the game in Fig. 114, where Nature (N) chooses each of player 1's initial nodes with probability 0.5. Player 1 does not know Nature's choice. Player 2 observes whether player 1 chose U or D, but not Nature's choice.

Now, let us consider the following assessment:

$$\{(U, L), (\mu_1 = 0.5, \mu_2 = 0.9)\}$$

Since player 1 chooses U, player 2's information set is never reached and any belief μ_2 is permitted by weak PBE. Suppose $\mu_2 = 0.9$. Player 2 prefers L to R if and only if:

$$(5)\mu_2 + (5)(1 - \mu_2) \geq (2)\mu_2 + (10)(1 - \mu_2)$$

Plugging in $\mu_2 = 0.9$, we see that:

$$5(0.9) + 5(0.1) > (2)(0.9) + (10)(0.1)$$

so then P2 will play L if $\mu_2 = 0.9$.

Now going up the tree, we can replace the subgame with the outcome that player 2 will always choose L as seen in Fig. 115. Clearly, we can see that P1 will prefer to play U over D as they earn a payoff of 2 versus 0. Therefore, playing U is sequentially rational for P1.

Therefore, $\{(U, L), (\mu_1 = 0.5, \mu_2 = 0.9)\}$ is a valid weak perfect Bayesian equilibrium. However, the belief $\mu_2 = 0.9$ is difficult to justify: it assigns greater probability to the left subtree even though Nature selects either subtree with probability 0.5.

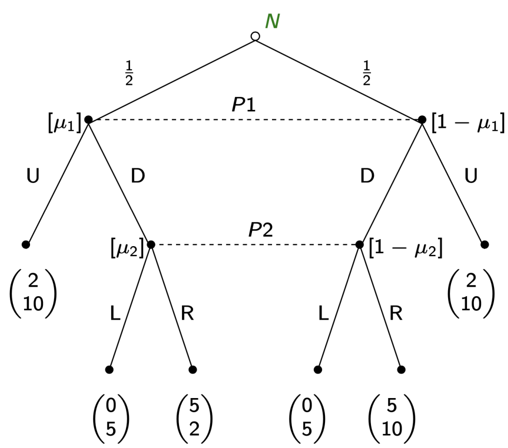

Figure 114: Game with nature choosing.

Incumbent Firm vs Entrant Firm

We look at the game in Fig. 116. We consider the following assessment:

$$\{((O, A), F); \mu = 1\}$$

As firm E plays Out, none of the nodes in firm I's information set is reached. Therefore, there are no restrictions on the belief μ, and $\mu = 1$ is admissible. Given $\mu = 1$, firm I prefers F to A. At its later decision node, firm E then plays A if it anticipates F. At the initial node, entering would lead to (A, F) and a payoff of -2 for firm E, so it instead plays Out and earns 0. Therefore, the assessment $\{((O, A), F); \mu = 1\}$ is sequentially rational and hence a weak perfect Bayesian equilibrium.

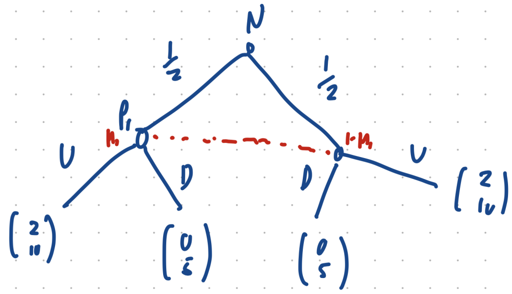

Figure 115: Game with nature choosing.

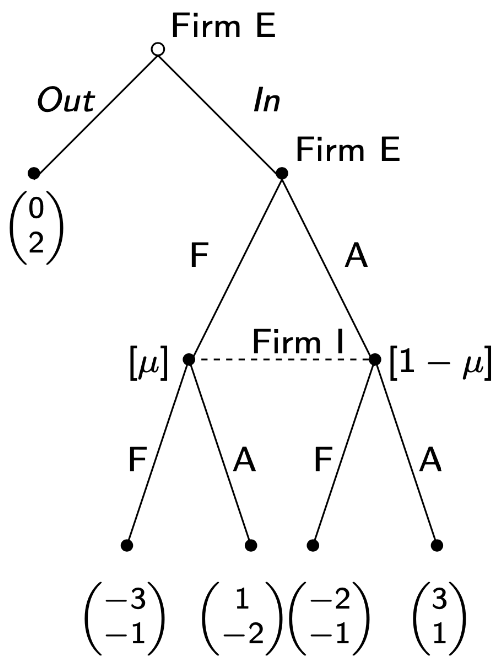

Figure 116: Entrant vs incumbent firm extensive form.

The issue is that the assessment's strategy profile is **not** a subgame-perfect Nash equilibrium. Consider the subgame starting at firm E in Fig. 117. Its unique Nash equilibrium is (A, A): A dominates F for firm I, and firm E's best response to A is also A. Any subgame-perfect Nash equilibrium must therefore induce (A, A) in this subgame. The assessment $\{((O, A), F); \mu = 1\}$ does not do so and is not subgame perfect.

Hence, the two examples show that a weak PBE need not be subgame perfect and that off-path beliefs need not respect the structure of the game, such as probabilities induced by Nature. These flaws arise because little is imposed on beliefs at information sets that are not reached on the equilibrium path. **To eliminate these flaws, we need to impose more structure on beliefs.**

We shall introduce **sequential equilibrium**, which strengthens the consistency requirement by introducing small *mistakes*. Under each perturbed, completely mixed strategy profile, every information set is reached and beliefs are pinned down by Bayes' rule. We then study the limit of these beliefs as the mistakes go to zero.

Consider a sequence of completely mixed behavioural strategies that converge to the strategy profile in the assessment. For each completely mixed behavioural strategy in the sequence, we can construct the system of beliefs by

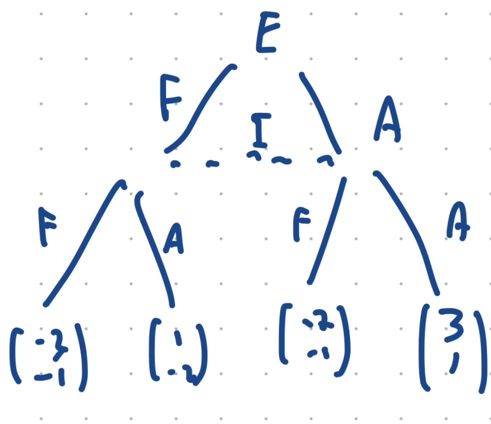

Figure 117: Entrant vs incumbent firm subgame starting at firm E.

using Bayes' rule. Note that under this construction, all information sets are reached with positive probability if the strategies are completely mixed. Then, we consider the limit of those sequences of systems of beliefs.

Definition 504 (Sequential Equilibrium)
An assessment (σ, μ) is a **sequential equilibrium** if:

1. The strategy profile σ is sequentially rational given the belief system μ.
2. There exists a sequence of completely mixed strategies $\{\sigma^k\}_{k=1}^{\infty}$ with $\sigma^k \to \sigma$. If μ^k denotes the beliefs generated from σ^k by Bayes' rule, then $\mu^k \to \mu$.

Remark 505. *The completely mixed perturbations used to construct a sequential equilibrium reach every information set, including those that are off path under the limiting strategy profile.*

One can interpret this sequence of completely mixed strategies as players making small mistakes. Every information set is reached under each perturbation, but those on the equilibrium path are much more likely. We require the assessment's beliefs to be the limit of the beliefs generated as these mistakes vanish.

Proposition 506 (First Procedure to Check for Sequential Equilibrium)
The first two-step procedure for checking sequential equilibrium is:

1. First find all assessments that are weak perfect Bayesian equilibria.
2. Find the information sets that are not reached and introduce *trembles* to see whether beliefs associated with these weak perfect Bayesian equilibria can arise as limits of Bayesian updating. If so, they are sequential equilibria.

Example 507
(Structural Consistency Revisited). We revisit Example 503, where off-path beliefs could conflict with Nature's probabilities. Because player 1 chose U, neither node in player 2's information set was reached, allowing arbitrary values of μ_2. The game is shown in Fig. 118. Consider the strategy profile (U, L) again and ask whether some beliefs can make it part of a sequential equilibrium.

First, instead of player 1 playing U, we now let player 1 play the mixed behavioural strategy which is given by:

$$(\sigma_1(U) = 1 - \epsilon_1, \sigma_1(D) = \epsilon_1)$$

This mixed behavioural strategy converges to the pure strategy U as ϵ_1 goes to 0:

$$(\sigma_1(U) = 1 - \epsilon_1, \sigma_1(D) = \epsilon_1) \to U \quad \text{as } \epsilon_1 \to 0.$$

Under this behavioural strategy, the nodes in player 2's information set are reached. We now calculate the belief μ_2 consistent with player 1's strategy. By Bayes' rule, it is the probability of reaching the relevant node divided by the probability of reaching information set I_2. Nature selects each branch with probability 1/2, and player 1 chooses D with probability ϵ_1 on either branch. Therefore,

$$\mu_2^\epsilon = \frac{\frac{1}{2}\epsilon_1}{\frac{1}{2}\epsilon_1 + \frac{1}{2}\epsilon_1} = \frac{1}{2}.$$

Given such a belief, L is not sequentially rational for player 2, since player 2 prefers L to R only if:

$$(5)\mu_2 + (5)(1 - \mu_2) > (2)\mu_2 + (10)(1 - \mu_2)$$

In this, case for $\mu_2^\epsilon = 1/2$, we have that:

$$(5)\frac{1}{2} + (5)(\frac{1}{2}) < (2)\frac{1}{2} + (10)(\frac{1}{2})$$

Therefore, player 2 should play R, not L.

Therefore, we can conclude that (U, L) is **not** part of a sequential equilibrium. To conclude, we have shown that (U,L) is part of a weak Perfect Bayesian Equilibrium but not a sequential equilibrium.

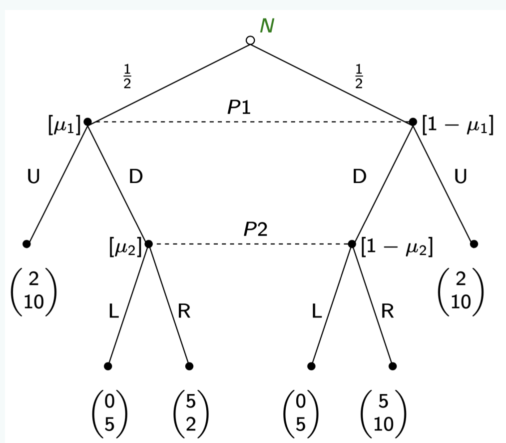

Figure 118: Game with nature choosing. We introduce trembles in player 1's strategy.

Example 508

(Subgame Perfection Revisited). We now revisit the game in Fig. 116. We showed that the weak perfect Bayesian equilibrium with strategy profile $((O, A), F)$ was **not** subgame perfect. Because firm E played Out, firm I's information set was not reached and arbitrary beliefs could be imposed there. We now introduce completely mixed behavioural strategies for firm E, as shown in Fig. 119.

The mixed behavioural strategy for firm E at node a is now given by:

$$\sigma_1^a(In) = \epsilon_1 \quad \sigma_1^a(Out) = 1 - \epsilon_1$$

and the mixed behavioural strategy for firm E (again) at node b is given by:

$$\sigma_1^b(F) = \epsilon_2 \quad \sigma_1^b(A) = 1 - \epsilon_2$$

These mixed behavioural strategies converge to player 1's pure strategy that we have shown earlier:

$$\big((\sigma_1^a(In), \sigma_1^a(Out)), (\sigma_1^b(F), \sigma_1^b(A))\big) \to (O, A) \quad \text{as } \epsilon_1 \to 0,\ \epsilon_2 \to 0.$$

With these trembles, firm I's information set is reached and its beliefs are pinned down by Bayes' rule. To reach the node associated with μ, firm E must play In and then F, which occurs with probability $\epsilon_1\epsilon_2$. The information set is reached if firm E plays either (In, F) or (In, A), with probabilities $\epsilon_1\epsilon_2$ and $\epsilon_1(1-\epsilon_2)$, respectively. Therefore,

$$\mu^\epsilon = \frac{\epsilon_1\epsilon_2}{(\epsilon_1\epsilon_2) + (\epsilon_1(1-\epsilon_2))} = \frac{\epsilon_1\epsilon_2}{\epsilon_1} = \epsilon_2$$

However, by definition of sequential equilibrium, we let $\epsilon_2 \to 0$ which means that:

$$\mu^\epsilon \to 0 \quad \text{as } \epsilon_2 \to 0.$$

However, if $\mu = 0$, it is not sequentially rational for Firm I to play F as they can earn a payoff of 1 playing A compared to a payoff -1 if they played F when $\mu = 0$.

Therefore, the strategy profile $((O, A), F)$ cannot be part of a sequential equilibrium, although it can be part of a weak perfect Bayesian equilibrium.

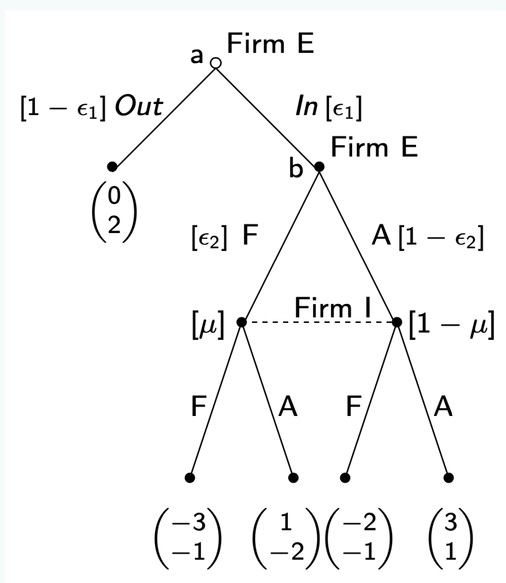

Figure 119: Entrant vs incumbent firm subgame with trembles introduced for firm E.

This process ensures that every sequential equilibrium is an SPNE (the reverse is not true) and that beliefs respect probabilities induced by Nature. Moreover, the assessment satisfies sequential rationality and consistency, so every sequential equilibrium is a weak PBE, while the reverse is again false.

Proposition 509 (Sequential Equilibria are SPNEs and Weak PBEs)
Every sequential equilibrium is a subgame-perfect Nash equilibrium and a weak perfect Bayesian equilibrium.

Since we now know that sequential equilibria are also subgame perfect Nash equilibria, this gives us another way to find sequential equilibria.

Proposition 510 (Second Procedure to Check for Sequential Equilibrium)
The second two-step procedure for checking sequential equilibrium is:

1. First find all assessments whose strategies form a **subgame-perfect Nash equilibrium**.
2. Find the information sets that are not reached and introduce *trembles* to see whether the associated beliefs can arise as limits of Bayesian updating. If so, the assessments are sequential equilibria.

Remark 511. *This now involves SPNEs rather than weak PBEs that we saw before.*

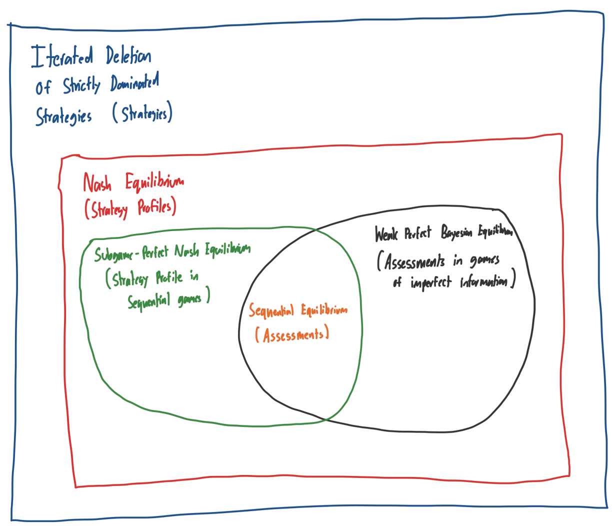

Figure 120: Relationships between different equilibria concepts.

Finally, the following theorem shows existence of sequential equilibrium in finite games.

Theorem 512 (Kreps and Wilson, 1982)
Every finite game has a sequential equilibrium.

From this result, we get another really useful result that will help us later.

Proposition 513 (Existence of Weak Perfect Bayesian Equilibrium in Finite Games)
Every finite game has a weak Perfect Bayesian equilibrium.

Proof. A sequential equilibrium is a weak Perfect Bayesian equilibrium and we know that every finite game has a sequential equilibrium. □

We have already seen an existence theorem for subgame perfect Nash equilibria, but we repeat it again.

Proposition 514 (Existence of Subgame Perfect Nash Equilibrium in Finite Games)
Every finite game has a subgame perfect Nash equilibrium.

Proof. Every sequential equilibrium is an SPNE, and every finite game has a sequential equilibrium. Therefore, every finite game has a subgame-perfect Nash equilibrium. $\square$

Finally, we note one more thing.

Proposition 515 (Completely Mixed NE)
Any completely mixed Nash equilibrium is a sequential equilibrium.

Remark 516. *All information sets are reached, hence the strategies will have to pin down the beliefs and the strategies are best responses to the beliefs.*

We can summarise and collate all the equilibria we have covered so far in Fig. 120.

10.3 Bayesian Nash Equilibrium

Up to now, we have assumed that each player knew the structure of the game—the players, timing, available actions, and how outcomes determine every player's payoff. These were games of **complete information**. This concept is distinct from perfect information.

- **Perfect Information**: Players perfectly observe the actions that have been chosen before. Here, all information sets are singleton sets.
- **Complete Information**: The structure of the game and the relationship between outcomes and payoffs are known to all. Information sets need not be singletons.

In reality, most games have incomplete information; one player might not know the choices available to others, or might not know the payoff of others after some interaction. These games are called games of **incomplete information**. Sometimes they are also called games of **asymmetric information** because typically a player will know more about her own payoff structure or her own available actions than her opponents. In fact a complete information game can be thought as a **symmetric information** game since there, all players hold the same information about the structure of the game.

Analysing a game in which players might not know which actions might be played, or how a profile of actions will translate into the payoffs for every player, is quite challenging. How would you predict how others would behave if you don't know their payoffs?

Example 517

(Bayesian Battle of the Sexes). Suppose that two students are deciding whether to go to the Pub (P) or Cinema (C). Player 1 likes player 2, and this is common knowledge. However, player 1 does not know whether player 2 likes or dislikes player 1. Hence, player 2 has two potential types. We represent this scenario with two payoff matrices, one for each type of player 2, in Fig. 121.

Like

	C	P
C	4, 3	1, 1
P	0, 0	3, 4

Dislike

	C	P
C	4, 0	1, 4
P	0, 3	3, 1

Figure 121: Bayesian battle of the sexes. The left payoff matrix is for when player 2 likes player 1. The right payoff matrix is for when player 2 dislikes player 1.

Harsanyi (1967) proposed modelling these games by introducing **types**, with one type for each relevant specification of the game that is unknown to some players. For example, if player 2 does not know whether player 1 has high or low marginal cost, we give player 1 a high-cost type and a low-cost type. Nature moves first and selects the realised type. Player 1 observes this choice, but player 2 does not. We have thereby transformed a game of **incomplete information** into a game of **complete but imperfect information**, which we know how to analyse.

Example 518

(Bayesian Battle of the Sexes Continued). We continue on with Example 517. Here, we apply Harsanyi's transformation to turn the game from a game of incomplete information into one with complete but imperfect information by having nature choose the type of player 2 with probability 0.5. We represent this in Fig. 122.

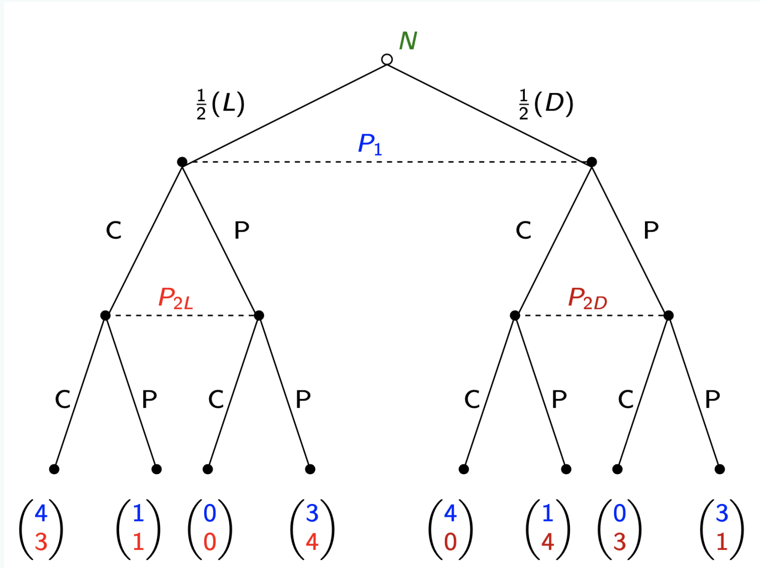

Figure 122: Bayesian battle of the sexes game being represented in extensive form. Nature chooses the type of player 2 with probability 0.5.

We will soon see that the Bayesian Nash equilibria will be the Nash Equilibria applied to the transformed form of the incomplete information game.

There is a hidden assumption that we are making when we do this transformation.

Assumption (Common Prior Assumption)
The **common prior assumption** is that the probability distribution that Nature uses to choose types is common knowledge to all players.

We are now ready to formally define Bayesian games, which are incomplete information games with the common prior assumption.

Definition 519 (Bayesian Game)
A **Bayesian Game** $\Gamma = [N, \{\Theta_i\}, \{S_i\}, p, \{u_i\}]$, consists of:

1. A finite set of players N.
2. For each player $i \in N$, a set of types Θ_i that the player could be.
3. For each player $i \in N$, a set of available actions S_i.
4. A joint probability distribution over the agents' types Θ, which we denote by $p(\theta_1, \theta_2, ..., \theta_N)$.
5. For each $i \in N$ a vNM utility function $u_i : S \times \Theta \longrightarrow \mathbb{R}$.

We make a few comments about this Bayesian game.

Player i's type θ_i is observed only by player i; it is private information. The distribution p of players' types is common knowledge (the common-prior assumption), and hence $p(\theta_{-i} \mid \theta_i)$ is player i's conditional distribution over the other players' types. We assume that the marginal $p_i(\theta_i)$ is strictly positive for each $\theta_i \in \Theta_i$, so we consider only types that arise with positive probability.

Definition 520 (Independent Types)
If the type of a player doesn't provide any further information about the types of the other players,

$$p(\theta_{-i}|\theta_i) = p(\theta_{-i})$$

then we say that types are **independent**.

The payoff function $u_i(s_1, ..., s_n; \theta_1, ..., \theta_n)$ depends not only on the strategy profile but also on the realisation of the types (both the player's type and other players' types).

Definition 521 (Private Types)
If a player's preferences depend only on their own type and the actions played, but not on the other players' types, then we say that the types are **private**:

$$u_i(s; \theta) = u_i(s; \theta_i).$$

Remark 522. *In some auction and signalling models, other agents' types may also enter a player's payoff.*

We can now define what strategies are in Bayesian games.

Definition 523 (Strategies in Bayesian Games)
A **pure strategy** for player i in a Bayesian game is a function $s_i : \Theta_i \longrightarrow S_i$, which for each type realisation θ_i selects a strategy $s_i(\theta_i) \in S_i$.

A **mixed strategy** for player i in a Bayesian game defines for every possible type, a probability distribution over pure strategies: $\sigma_i : \Theta_i \longrightarrow \Delta(S_i)$.

Remark 524. *It is important to realise that a pure strategy for Bayesian games differ as it is now a function mapping from the space of potential types for player i: Θ_i. Before, it was a mapping from a player's information set! Hence, we can think of this as we are assigning each information set to a player's type.*

Note that although each player is aware of her type, a strategy for player i needs to specify the action chosen for every possible realisation of their type. This is similar to how previously strategies still needed to be specified for information sets that were never reached.

We can now define a Nash equilibrium in a Bayesian game.

Definition 525 (Bayesian Nash Equilibrium)
A **Bayesian Nash equilibrium** (BNE) is a profile of strategies $(\sigma_1(\cdot), \dots, \sigma_n(\cdot))$ such that no player, **conditional on their type**, has a profitable deviation; that is, for all i, all $\theta_i \in \Theta_i$, and all $\sigma_i' \in \Delta(S_i)$:

$$\sum_{\theta_{-i}} p(\theta_{-i}|\theta_i) u_i(\sigma_i(\theta_i), \sigma_{-i}(\theta_{-i}); (\theta_i, \theta_{-i})) \geq \sum_{\theta_{-i}} p(\theta_{-i}|\theta_i) u_i(\sigma_i'(\theta_i), \sigma_{-i}(\theta_{-i}); (\theta_i, \theta_{-i}))$$

Thus, conditional on type θ_i, player i obtains at least as much expected utility from $\sigma_i(\theta_i)$ as from any alternative $\sigma_i'(\theta_i)$.

Remark 526. *Notice that the two terms differ by the expression* $\sigma_i(\theta_i)$ *vs* $\sigma_i'(\theta_i)$.

Note that if you transformed this Bayesian game into an imperfect but complete information game as we said above, the definition of the Bayesian Nash Equilibrium coincides with the definition of Nash Equilibrium in the transformed game. In particular, since we know that in any finite game there exists a (mixed) Nash equilibrium, we can conclude that in any finite Bayesian game there exists a BNE.

Proposition 527 (Existence of Bayesian Nash Equilibrium)
Every finite Bayesian game has a Bayesian Nash Equilibrium.

Bayesian Battle of the Sexes Continued

We now solve the Bayesian battle of the sexes to find the Bayesian Nash equilibrium. First, recall that the game had the structure in Fig. 121. First, note that we have two players $N = \{1, 2\}$. Player 1 has only one type $\Theta_1 = \{1\}$ which is common knowledge. Player 2 has two types $\Theta_2 = \{L, D\}$. Each player has two possible strategies of either going to the cinema or the pub:

$$S_1 = S_2 = \{C, P\}.$$

The probability distribution over the types of the player are:

$$P(1, L) = P(1, D) = \frac{1}{2}$$

Thus, player 2 has either type L or type D with probability one half, while player 1 has only one possible type. The types are independent, and the payoffs are described in the two payoff matrices. Types are also private: each player's payoff depends on their own type and the action profile, but not directly on the other player's type.

Like

	C	P
C	4, 3	1, 1
P	0, 0	3, 4

Dislike

	C	P
C	4, 0	1, 4
P	0, 3	3, 1

Figure 123: Bayesian battle of the sexes. The left payoff matrix is for when player 2 likes player 1. The right payoff matrix is for when player 2 dislikes player 1.

Now, a pure Bayesian Nash Equilibrium is given by the strategy profile:

$$(s_1; (s_{2L}, s_{2D}))$$

Notice that for player 1, they have only one type so they require only one strategy s_1. However, player 2 has two types (L or D) and hence need to specify two strategies s_{2L} and s_{2D}. For a strategy profile to be a Nash equilibrium, the strategies must be a **best response** to the other players' strategies:

$$s_1 = BR_1(s_{2L}, s_{2D}) \quad \text{and} \quad (s_{2L}, s_{2D}) = BR_2(s_1)$$

Let us figure out what the best response functions are.

P_1 plays C

If P_1 plays C, then for P_{2L}, their best response is to play C as they earn a payoff of 3 compared to a payoff of 1 (look at the matrix on the left).

If P_1 plays C, then for P_{2D}, their best response is to play P as they earn a payoff of 4 compared to a payoff of 0 (look at the matrix on the right).

P_{2L} plays C and P_{2D} plays P

Analysing P_1's best response is slightly more involved because P_2 could be either type L or type D. We must therefore compare P_1's expected payoffs over player 2's type.

Recall that we previously found (C, P) as the respective best responses for P_{2L} and P_{2D}.

If P_1 plays C, then their payoff is the weighted payoffs of (C,C) if P_2 is type L and (C,P) if P_2 is type D (where there is a probability of 0.5 of P_2's type):

$$0.5(u_1(C, C)) + 0.5(u_1(C, P)) = 0.5(4) + 0.5(1) = \frac{5}{2}.$$

If P_1 plays P, then their payoff is:

$$0.5(u_1(P, C)) + 0.5(u_1(P, P)) = 0.5(0) + 0.5(3) = \frac{3}{2}.$$

Therefore, P_1 should play C.

Hence, we have found a pure Bayesian Nash equilibrium in which the strategies are mutual best responses:

$$(s_1; (s_{2L}, s_{2D})) = (C, (C, P)).$$

This is the only pure Bayesian Nash equilibrium.

Mixed Bayesian Nash Equilibrium

We now look for a mixed Bayesian Nash equilibrium. We suppose that the three players P_1, P_{2L}, P_{2D} mix between playing C and P with the following probabilities:

1. $p = \sigma_1(C)$
2. $q = \sigma_{2L}(C)$
3. $r = \sigma_{2D}(C)$

We again compute the best responses for the three players.

Assume that P_1 is playing the mixed strategy $p = \sigma_1(C)$. Then, for P_{2L}, their respective payoffs from playing C or P (by looking at the Like Matrix) is:

$$u_{2L}(\sigma_1, C) = (3)p + (0)(1 - p) \qquad u_{2L}(\sigma_1, P) = (1)p + (4)(1 - p)$$

From this, we see that P_{2L} should play C if and only if:

$$u_{2L}(\sigma_1, C) \geq u_{2L}(\sigma_1, P)$$

$$\leftrightarrow 3p \geq p + 4(1 - p)$$

$$\Rightarrow p \geq \frac{2}{3}$$

For P_{2D}, assuming that P_1 plays mixed strategy p, then their payoffs are:

$$u_{2D}(\sigma_1, C) = (0)p + (3)(1-p) \qquad u_{2D}(\sigma_1, P) = (4)p + (1)(1-p)$$

Therefore, P_{2D} should play C if and only if:

$$u_{2D}(\sigma_1, C) \geq u_{2D}(\sigma_1, P)$$

$$p \leq \frac{1}{3}.$$

Now, for P_1, we assume that P_{2L} plays mixed strategy q and P_{2D} plays mixed strategy r. Then, P_1 payoff from playing C is:

$$u_1(C, \sigma_2) = \frac{1}{2}\big((4)q + (1)(1-q)\big) + \frac{1}{2}\big((4)r + (1)(1-r)\big)$$
$$u_1(P, \sigma_2) = \frac{1}{2}\big((0)q + (3)(1-q)\big) + \frac{1}{2}\big((0)r + (3)(1-r)\big)$$

Therefore, P_1 should play C if and only if:

$$u_1(C, \sigma_2) \geq u_1(P, \sigma_2)$$

which is

$$\frac{1}{2}\big((4)q + (1)(1-q)\big) + \frac{1}{2}\big((4)r + (1)(1-r)\big) \geq \frac{1}{2}\big((0)q + (3)(1-q)\big) + \frac{1}{2}\big((0)r + (3)(1-r)\big)$$

Solving this gives us that P_1 should play C if and only if:

$$q + r \geq \frac{2}{3}.$$

We can now summarise the best responses for the players' mixed strategies and label the conditions for the best responses as such:

(*) P_1: Play C if and only if $q + r \geq \frac{2}{3}$

(**) P_{2L}: Play C if and only if $p \geq \frac{2}{3}$

(***) P_{2D}: Play C if and only if $p \leq \frac{1}{3}$

Notice that P_1's decision depends on P_{2L} and P_{2D}'s mixed strategy q and r. Meanwhile, the other 2 player's decision depends on P_1's mixed strategy p.

We now bring all this together and check the five possible ranges for p.

$p < \frac{1}{3}$

Under this scenario, P_{2L} will not play C as (**) is not satisfied. Therefore,

$$q = \sigma_{2L}(C) = 0.$$

However, player P_{2D} will play C as (***) is satisfied. Therefore,

$$r = \sigma_{2D}(C) = 1.$$

Therefore, combining these two mixed strategies:

$$q + r = 0 + 1 = 1$$

But, we then have that $q + r = 1 > \frac{2}{3}$ so (*) is satisfied. Therefore, P_1 will play C for sure:

$$p = \sigma_1(C) = 1.$$

However, this contradicts our initial assumption that $p < \frac{1}{3}$ so this scenario is not possible!

$p = \frac{1}{3}$

For P_{2L}, (**) is not satisfied, so P_{2L} will not play C:

$$q = \sigma_{2L}(C) = 0.$$

However, P_{2D} is indifferent between C and P because (***) holds with equality. Therefore, any $r \in [0, 1]$ is a best response. We choose r so that P_1 is also indifferent between C and P. Thus condition (*) must hold with equality:

$$(*) = q + r = \frac{2}{3}$$

and as $q = 0$, we therefore have that:

$$r = \frac{2}{3}$$

Therefore, if $r = \frac{2}{3}$, P_1 will be indifferent to playing C. Then, this means that p can take any value $p \in [0, 1]$. However, to be consistent with our assumption, we set:

$$p = \frac{1}{3}$$

Therefore, we have found the mixed Bayesian Nash equilibrium:

$$\big(p = \frac{1}{3}, (q = 0, r = \frac{2}{3})\big).$$

$\frac{1}{3} < p < \frac{2}{3}$

For P_{2L}, we see that (**) is still not satisfied. Therefore, P_{2L} will play P so that we have :

$$q = \sigma_{2L}(C) = 0.$$

For P_{2D}, (***) is no longer satisfied. Therefore, P_{2D} will play P, and hence:

$$r = \sigma_{2D}(C) = 0.$$

Now for P_1, looking at (*), we see that it is not satisfied as:

$$q + r = 0 + 0 < \frac{2}{3}$$

Therefore, P_1 will not play C so therefore:

$$p = \sigma_1(C) = 0.$$

However, this contradicts our assumption that $\frac{1}{3} < p < \frac{2}{3}$. Therefore, such a scenario is not possible.

$p = \frac{2}{3}$

For P_{2L}, looking at (**), we see that the inequality condition holds with equality. Therefore, P_{2L} is indifferent between playing C and P. Therefore, we have that:

$$q = \sigma_{2L}(C) \in [0, 1].$$

For P_{2D}, looking at (***), we see that this condition does not hold. Therefore P_{2D} will not play C and hence:

$$r = \sigma_{2D}(C) = 0.$$

For the assumption $p = \frac{2}{3}$ to hold, P_1 must be indifferent between C and P. Thus condition (*) must hold with equality. Since $r = 0$, we require:

$$q + r = \frac{2}{3}$$

which holds only if $q = \frac{2}{3}$. At this value, P_1 is indifferent between C and P. Therefore,

$$p = \sigma_1(C) \in [0, 1]$$

and hence to be consistent with our assumption, we can set:

$$p = \frac{2}{3}.$$

Therefore, we have found another mixed Bayesian Nash equilibrium:

$$(p = \frac{2}{3}, (q = \frac{2}{3}, r = 0)).$$

$p > \frac{2}{3}$

For P_{2L}, looking at (**), we see that they will play C so then:

$$q = \sigma_{2L}(C) = 1.$$

For P_{2D}, looking at (***), we see that this condition does not hold. Therefore P_{2D} will not play C and hence:

$$r = \sigma_{2D}(C) = 0.$$

For P_1, looking at (*), we see that:

$$q + r = 1 + 0 > \frac{2}{3}$$

and therefore P_1 will play C. Hence,

$$p = \sigma_1(C) = 1$$

Having $p = 1$ is consistent with our assumption that $p > 2/3$. Therefore, we have found another equilibrium:

$$\big(p = 1, (q = 1, r = 0)\big)$$

However, notice that these are actually **pure strategies**:

$$\big(C, (C, P)\big).$$

This was the pure strategy Bayesian Nash equilibrium we found earlier!

Hence, to summarise, we have that in the Bayesian Battle of the Sexes, the pure strategy Bayesian Nash Equilibrium is:

$$(C, (C, P))$$

and the two mixed strategy Bayesian Nash equilibria are:

$$(p = \frac{1}{3}, (q = 0, r = \frac{2}{3})).$$

$$(p = \frac{2}{3}, (q = \frac{2}{3}, r = 0)).$$

10.3.1 Bayesian Nash Equilibrium with Infinite Types

We have shown that finite Bayesian games have a Nash equilibrium. We now consider a Bayesian game with infinitely many possible types.

Example 528

(Public Good Provision). Consider two players, each with two possible actions but a continuum of types. Player i decides whether to contribute ($s_i = 1$) or not ($s_i = 0$) to a public good. The contribution cost $c_i \in [0, \overline{c}]$ is player i's private information. For both players, c_i has cumulative distribution function $F(\cdot)$, and c_1 and c_2 are independent.

We now write out the payoff for player i.

$$u_i(s_i, s_j; c_i) = \max\{s_1, s_2\} - c_i s_i = \begin{cases} 1 - c_i & \text{if } s_i = 1, s_j \in \{0, 1\} \\ 1 & \text{if } s_i = 0, s_j = 1 \\ 0 & \text{if } s_i = 0 = s_j \end{cases} \tag{197}$$

where $\max\{s_1, s_2\}$ indicates whether the public good is provided. Player i's payoff depends on both actions and on their private cost. Types are private because the payoff does not depend directly on the other player's cost.

We now solve for the Bayesian Nash equilibrium. Look at player i. Take the strategy of the opponent $s_j(\cdot)$ as given. Now, denote by z_j:

$$z_j = \mathbb{P}\big(s_j(c_j) = 1\big),$$

which is the probability that player j contributes. This probability depends on player j's equilibrium strategy $s_j(\cdot)$, which remains to be determined.

If player i does not contribute, $s_i = 0$, their expected payoff equals the probability that the other player contributes:

$$\mathbb{E}_{c_j}\Big[u_i(0, s_j(c_j); c_i)\Big] = z_j$$

If player i did contribute, then their expected payoff is:

$$\mathbb{E}_{c_j}\Big[u_i(1, s_j(c_j); c_i)\Big] = 1 - c_i$$

which is the payoff player i gets regardless of the player j's contribution as the public good is provided.

Therefore, the best response function for player i is:

$$s_i(c_i) = \begin{cases} 1 & \text{if } c_i < 1 - z_j \\ 0 & \text{if } c_i \geq 1 - z_j \end{cases}$$

This says that the probability that player i contributes to the public good depends on their cost c_i and the probability that player j does **not** contribute $(1 - z_j)$. If player i's cost is lower than the probability that player j does not contribute, then player i should just go ahead and contribute!

Therefore, the probability that player i contributes to the public good is given by:

$$\mathbb{P}(s_i(c_i) = 1) = \mathbb{P}(c_i < 1 - z_j)$$

Equivalently, as $c_i \sim F(\cdot)$, this can be written as:

$$z_i = F(1 - z_j).$$

Thus, player i's probability of contributing is the CDF evaluated at the cutoff $1 - z_j$.

By symmetry, player 2's probability of contributing is:

$$z_2 = F(1 - z_1)$$

Now, suppose that (z_1^*, z_2^*) solves $z_1 = F(1 - z_2)$ and $z_2 = F(1 - z_1)$. Then, the Bayesian Nash equilibrium is given by:

$$s_i(c_i) = \begin{cases} 1 & \text{if } c_i < 1 - z_j^* \\ 0 & \text{if } c_i \geq 1 - z_j^* \end{cases}$$

We can then look for a symmetric equilibrium and in this case, z^* would solve $z = F(1 - z)$.

10.3.2 Bayesian Nash Equilibrium with Infinite Actions and Types

We now consider a setting in which agents have infinitely many types and actions.

Example 529

(War of Attrition).

Let us suppose that two animals are fighting over a prey which they both value. Let us denote $\theta_i \in [0, \infty)$ the value of the prey for player i. This value θ_i is private information so it is only known by player i. Suppose that this type $\theta_i \sim F(\cdot)$ where $F(\cdot)$ is a cumulative distribution function that is the same for both players. As a distribution function, we assume that it is differentiable and has a probability density function $f(\cdot)$.

The animal that first stops fighting loses. The strategy for the player is the length of time the animal decides to fight for, $s_i \in [0, \infty)$. However, fighting is costly for the animal and hence their cost is the amount of time they fight for. This is a static game as there is no arrival of information during the game.

If animal i is the first to give up, $s_i \leq s_j$, its payoff is the negative cost of the time spent fighting, $-s_i$. If animal j gives up first, $s_i > s_j$, animal i obtains the prey and pays the cost incurred until time s_j, giving payoff $\theta_i - s_j$. Thus,

$$u_i(s_i, s_j; \theta_i) = \begin{cases} -s_i & \text{if } s_j \geq s_i \\ \theta_i - s_j & \text{if } s_j < s_i; \end{cases} \tag{198}$$

We now seek a pure-strategy Bayesian Nash equilibrium $(s_1(\cdot), s_2(\cdot))$.

The expected payoff for animal i given their strategy s_i, their type θ_i, and animal j's strategy s_j, is the probability that player j fights longer than player i times the negative cost for player i plus the probability that player i fights longer than player j times the expected payoff from the value of the prey less the cost. This can be written as:

$$\mathbb{E}_{\theta_j}\left[u_i(s_i, s_j(\theta_j); \theta_i)\right] = \underbrace{\mathbb{P}(s_j(\theta_j) \geq s_i)(-s_i)}_{(A)} + \underbrace{\mathbb{P}(s_j(\theta_j) < s_i)\mathbb{E}_{\theta_j}\left[\theta_i - s_j(\theta_j) \mid s_j(\theta_j) < s_i\right]}_{(B)} \tag{199}$$

Assume that animal j's fight time $s_j(\cdot)$ is a strictly increasing and continuous function of its valuation θ_j. Intuitively, an animal that values the prey more is willing to fight longer. The strategy therefore has an inverse on its range, which we denote by:

$$\phi_j(\cdot) = s_j(\cdot)^{-1}$$

whereby by the definition of an inverse:

$$\phi_j(s_j(\theta_j)) = (\phi_j \circ s_j)(\theta_j) = (s_j^{-1} \circ s_j)(\theta_j) = \theta_j$$

Term (A)

We now analyse term A in Eq. (199):

$$\mathbb{P}(s_j(\theta_j) \geq s_i)(-s_i)$$

We make use the inverse of s_j to rewrite this as:

$$\begin{aligned}
\mathbb{P}(s_j(\theta_j) \geq s_i)(-s_i) &= \mathbb{P}(\theta_j \geq s_j^{-1}(s_i))(-s_i) \\
&= \mathbb{P}(\theta_j \geq \phi_j(s_i))(-s_i) \\
&=_{(1)} [1 - F(\phi_j(s_i))](-s_i)
\end{aligned}$$

where $=_{(1)}$ arises as $\theta_j \sim F(\cdot)$.

Term (B)

We look at term B in Eq. (199):

$$\mathbb{P}(s_j(\theta_j) < s_i)\mathbb{E}_{\theta_j}\left[\theta_i - s_j(\theta_j) | s_j(\theta_j) < s_i\right]$$

We rewrite this using the inverse:

$$\begin{aligned}
\mathbb{P}(s_j(\theta_j) < s_i)\mathbb{E}_{\theta_j}\left[\theta_i - s_j(\theta_j) | s_j(\theta_j) < s_i\right] &= \mathbb{P}(\theta_j < s_j^{-1}(s_i))\mathbb{E}_{\theta_j}[\theta_i - s_j(\theta_j) | \theta_j < s_j^{-1}(s_i)] \\
&=_{(1)} \mathbb{P}(\theta_j < \phi_j(s_i))\mathbb{E}_{\theta_j}[\theta_i - s_j(\theta_j) | \theta_j < \phi_j(s_i)] \\
&= \int_0^{\phi_j(s_i)} (\theta_i - s_j(\theta_j)) f(\theta_j) d\theta_j
\end{aligned}$$

where $=_{(1)}$ uses the conditional-expectation identity

$$\mathbb{P}(X \leq c)\mathbb{E}[h(X) \mid X \leq c] = \int_{-\infty}^{c} h(x) f(x)\, dx.$$

We combine our results for (A) and (B) into Eq. (199) to get:

$$\mathbb{E}_{\theta_j}\left[u_i(s_i, s_j(\theta_j); \theta_i)\right] = -s_i[1 - F(\phi_j(s_i))] + \underbrace{\int_0^{\phi_j(s_i)} (\theta_i - s_j(\theta_j))f(\theta_j)d\theta_j}_{(C)} \tag{200}$$

We will re-express term (C) to make it easier to work with by using a change of variable. Let $x_j = s_j(\theta_j)$. This implies that:

$$\theta_j = s_j^{-1}(x_j) \leftrightarrow \theta_j = \phi_j(x_j).$$

Furthermore, we notice that the limits of integration in (C) changes as:

$$x_j = s_j(\phi_j(s_i)) = (s_j \circ \phi_j)(s_i) = s_i.$$

Now, we recall an important result from Statistics.

Theorem 530 (Density of Transformed Random Variable)
If the random variable X has a density function f_X and $g(\cdot)$ is a one-to-one function, then the random variable $Y = g(X)$ has the probability density function:

$$f_Y(y) = f_X\left(g^{-1}(y)\right)\left|(g^{-1})'(y)\right|. \tag{201}$$

We make use of this result by applying this to (C) with our change of variable, we end up with:

$$(C) = \int_0^{s_i} (\theta_i - x_j)f(\phi_j(x_j))\phi_j'(x_j)dx_j$$

Plugging this back into term (C) in Eq. (200), we end up with:

$$\mathbb{E}_{\theta_j}\left[u_i(s_i, s_j(\theta_j); \theta_i)\right] = -s_i[1 - F(\phi_j(s_i))] + \int_0^{s_i} (\theta_i - x_j)f(\phi_j(x_j))\phi_j'(x_j)dx_j \tag{202}$$

which is the expected payoff for player i of type θ_i when player i chooses s_i and player j uses $s_j(\theta_j)$.

Player i finds a best response by maximising the expected payoff in Eq. (202) with respect to the fight time s_i:

$$\begin{aligned}
\frac{\partial}{\partial s_i}\left[\mathbb{E}_{\theta_j}\left[u_i(s_i, s_j(\theta_j); \theta_i)\right]\right] &= \frac{\partial}{\partial s_i}\left[-s_i[1 - F(\phi_j(s_i))] + \int_0^{s_i} (\theta_i - x_j)f(\phi_j(x_j))\phi_j'(x_j)dx_j\right] \\
&=_{(1)} -(1 - F(\phi_j(s_i))) + s_i f(\phi_j(s_i))\phi_j'(s_i) + (\theta_i - s_i)f(\phi_j(s_i))\phi_j'(s_i) \\
&= -(1 - F(\phi_j(s_i))) + \theta_i f(\phi_j(s_i))\phi_j'(s_i).
\end{aligned}$$

The second line uses the product rule for the first term and the Leibniz rule for the integral. At an interior optimum, we set this derivative equal to zero.

Therefore, we can rewrite this to get:

$$\underbrace{1 - F(\phi_j(s_i))}_{\text{Marginal cost of continuing in the fight}} = \underbrace{\theta_i f(\phi_j(s_i))\phi_j'(s_i)}_{\text{Marginal benefit of continuing in the fight}} \tag{203}$$

Therefore, the strategy $s_i(\theta_i)$ must solve the first-order condition in Eq. (203).

To make life easier, we can impose that we have a **symmetric equilibrium** so that the optimal strategies are the same for both animals:

$$s_i(\theta) = s(\theta) \quad \text{and} \quad \phi_j(x) = \phi(x)$$

The first result we get from this is that we can see that the inverse function of the opponent types becomes:

$$\phi_j(s_i(\theta_i)) = \phi(s(\theta_i)) = (\phi \circ s)(\theta_i) = \theta_i$$

The second result we get first uses the fact that:

$$\phi(s(\theta_i)) = \theta_i$$

and differentiating on both sides with respect to θ_i (and the chain rule):

$$\phi'(s(\theta_i)) \cdot s'(\theta_i) = 1$$

so then we have that the second result we get from a symmetric equilibrium is:

$$\phi'(s(\theta_i)) = \frac{1}{s'(\theta_i)}.$$

So now if we go back to our first order condition in Eq. (203),

$$1 - F(\phi_j(s_i)) = \theta_i f(\phi_j(s_i))\phi_j'(s_i)$$

this can now be rewritten as:

$$1 - F(\phi(s(\theta_i))) = \theta_i f(\phi(s(\theta_i)))\phi'(s(\theta_i))$$

and using our two results, this becomes:

$$1 - F(\theta_i) = \theta_i f(\theta_i)\frac{1}{s'(\theta_i)}$$

Rearranging for $s'(\theta_i)$ gives

$$s'(\theta_i) = \theta_i \frac{f(\theta_i)}{1 - F(\theta_i)}$$

Integrating this expression and using the initial condition $s(0) = 0$ gives the desired result.

Theorem 531 (Optimal Strategy in Attrition Game)

The optimal strategy for player i under a symmetric equilibrium in the attrition game is:

$$s(\theta_i) = \int_0^{\theta_i} t\frac{f(t)}{1 - F(t)}dt.$$

The strategy profile $(s(\theta_i), s(\theta_j))$ is the **Bayesian Nash Equilibrium** of the attrition game.

Now, we can specify a cumulative distribution function $F(\cdot)$ for the type and immediately get a result for the strategy. Suppose that we assumed that the type of a player is given by the exponential distribution:

$$F(\theta_i) = 1 - e^{-\theta_i}.$$

Then, plugging this into our optimal strategy formula, this gives us:

$$s(\theta_i) = \int_0^{\theta_i} t\frac{e^{-t}}{e^{-t}}dt$$

and this can be solved to get the Bayesian Nash equilibrium:

$$s(\theta_i) = \frac{\theta_i^2}{2} \tag{204}$$

10.4 Purification

When we discussed mixed strategies so far, we acknowledged that the concept required players to choose a very specific mixing among pure strategies, even though they were completely indifferent about all mixing with the same (or smaller) support. The specific mixing corresponded to the one that made the **other** players indifferent among their strategies. It is not clear how we can justify such a choice.

We are now going to introduce a result that explains that a mixed equilibrium can be seen as the limit of a pure equilibrium in a sequence of slightly perturbed Bayesian games.

Consider a finite complete-information game $\Gamma = [N, \{\Delta(S_i)\}, \{u_i\}]$. We construct perturbed Bayesian games as follows. For each player and pure action, let θ_i^s be a payoff shock with support $[-1, 1]$ (or any closed interval), and let $\epsilon > 0$. The ϵ-perturbed game is $\Gamma_\epsilon = [N, \{\Theta_i\}, \{\Delta(S_i)\}, p, \{\tilde{u}_i\}]$, where

$$\tilde{u}_i(s, \theta) = u_i(s) + \epsilon\theta_i^s,$$

and the shocks are independently and continuously distributed.

Purification in Battle of the Sexes

We look at the perturbed battle of the sexes game in Fig. 124. Here, $\epsilon\theta_1$ and $\epsilon\theta_2$ represents the perturbations in the players' preferences. We have that θ_i is the private information to player i whereby $\theta_1, \theta_2 \sim U[0, 1]$ are independent random variables. We perturb the game by a small ϵ.

	C	P
C	$2 + \epsilon\theta_1, 1$	$\epsilon\theta_1, \epsilon\theta_2$
P	$0, 0$	$1, 2 + \epsilon\theta_2$

Figure 124: Purification applied to battle of the sexes.

Recall that in a Bayesian game, a strategy is a function of the players' type:

$$s_i : \theta_i \to \{C, P\}$$

We now fix player 2's strategy $s_2(\cdot)$. Then player 1 should play C if and only if:

$$u_1(C, C)\mathbb{P}[s_2(\theta_2) = C] + u_1(C, P)\mathbb{P}[s_2(\theta_2) = P] \geq u_1(P, C)\mathbb{P}[s_2(\theta_2) = C] + u_1(P, P)\mathbb{P}[s_2(\theta_2) = P]$$

which becomes:

$$(2 + \epsilon\theta_1)\mathbb{P}[s_2(\theta_2) = C] + (\epsilon\theta_1)\mathbb{P}[s_2(\theta_2) = P] \geq (0)\mathbb{P}[s_2(\theta_2) = C] + (1)\mathbb{P}[s_2(\theta_2) = P]$$

Note that by the complementary property of probability, we can write: $\mathbb{P}[s_2(\theta_2) = P] \leftrightarrow 1 - \mathbb{P}[s_2(\theta_2) = C]$ and therefore we have:

$$(2 + \epsilon\theta_1)\mathbb{P}[s_2(\theta_2) = C] + (\epsilon\theta_1)\mathbb{P}[s_2(\theta_2) = P] \geq 1 - \mathbb{P}[s_2(\theta_2) = C]$$

$$2\mathbb{P}[s_2(\theta_2) = C] + \epsilon\theta_1\mathbb{P}[s_2(\theta_2) = C] + \epsilon\theta_1 - \epsilon\theta_1\mathbb{P}[s_2(\theta_2) = C] \geq 1 - \mathbb{P}[s_2(\theta_2) = C]$$

Hence, player 1 should play C if and only if:

$$\epsilon\theta_1 \geq 1 - 3\mathbb{P}[s_2(\theta_2) = C] \equiv \epsilon\bar{\theta}_1.$$

Now, for player 2, they should play C if and only if:

$$u_2(C, C)\mathbb{P}[s_1(\theta_1) = C] + u_2(P, C)\mathbb{P}[s_1(\theta_1) = P] \geq u_2(C, P)\mathbb{P}[s_1(\theta_1) = C] + u_2(P, P)\mathbb{P}[s_1(\theta_1) = P]$$

which becomes:

$$(1)\mathbb{P}[s_1(\theta_1) = C] + (0)\mathbb{P}[s_1(\theta_1) = P] \geq (\epsilon\theta_2)\mathbb{P}[s_1(\theta_1) = C] + (2 + \epsilon\theta_2)\mathbb{P}[s_1(\theta_1) = P]$$

Solving and re-arranging, we get:

$$\epsilon\theta_2 \leq 3\mathbb{P}[s_1(\theta_1) = C] - 2 \equiv \epsilon\bar{\theta}_2$$

So we have the 2 conditions whereby player 1 should play C if and only if:

$$\epsilon\theta_1 \geq 1 - 3\mathbb{P}[s_2(\theta_2) = C] \equiv \epsilon\bar{\theta}_1 \tag{205}$$

and player 2 should play C if and only if:

$$\epsilon\theta_2 \leq 3\mathbb{P}[s_1(\theta_1) = C] - 2 \equiv \epsilon\bar{\theta}_2 \tag{206}$$

These are threshold strategies! That is, strategies that are based on the thresholds. We want to now retrieve what these thresholds are.

For $\mathbb{P}[s_2(\theta_2) = C]$, recall that player 2 plays C if and only if:

$$\epsilon\theta_2 \leq \epsilon\bar{\theta}_2$$

We assumed that $\theta_2 \sim U[0, 1]$ and hence the CDF of $\mathbb{P}[\theta_2 \leq \bar{\theta}_2] = \bar{\theta}_2$. Going back to player 1's threshold, we therefore have that:

$$\epsilon\bar{\theta}_1 = 1 - 3\mathbb{P}[s_2(\theta_2) = C] = 1 - 3\bar{\theta}_2$$

We now want to retrieve $\epsilon\bar{\theta}_2$. For this, we need to know what $\mathbb{P}[s_1(\theta_1) = C]$ is. Since player 1 will play C if and only if $\epsilon\theta_1 \geq \epsilon\bar{\theta}_1$ and $\theta_1 \sim U[0, 1]$, this implies that:

$$\mathbb{P}[\epsilon\theta_1 \geq \epsilon\bar{\theta}_1] = \mathbb{P}[\theta_1 \geq \bar{\theta}_1] = 1 - \mathbb{P}[\theta_1 \leq \bar{\theta}_1] = 1 - \bar{\theta}_1$$

So then,

$$\begin{aligned}
\epsilon\bar{\theta}_2 &= [3\mathbb{P}[s_1(\theta_1) = C]] - 2 \\
&= [(1 - \bar{\theta}_1)3] - 2 \\
&= 3 - 3\bar{\theta}_1 - 2 \\
&= 1 - 3\bar{\theta}_1
\end{aligned}$$

Therefore, combining the two thresholds:

$$\epsilon\bar{\theta}_1 = 1 - 3\bar{\theta}_2 \quad \text{and} \quad \epsilon\bar{\theta}_2 = 1 - 3\bar{\theta}_1$$

solving for the thresholds gives us:

$$\bar{\theta}_1 = \bar{\theta}_2 = \frac{1}{\epsilon + 3}$$

So the pure Bayesian Nash equilibrium of the perturbed game is $(s_1(\cdot), s_2(\cdot))$ whereby we have the threshold strategies:

$$s_1(\theta_1) = \begin{cases} C & \text{if } \theta_1 \geq \frac{1}{\epsilon+3} \\ P & \text{if } \theta_1 < \frac{1}{\epsilon+3} \end{cases}$$

$$s_2(\theta_2) = \begin{cases} C & \text{if } \theta_2 \leq \frac{1}{\epsilon+3} \\ P & \text{if } \theta_2 > \frac{1}{\epsilon+3} \end{cases}$$

The probabilities that the players will play C can be read off from these strategies:

$$\mathbb{P}[\theta_1 \geq \bar{\theta}_1] = 1 - \frac{1}{\epsilon + 3} \qquad \mathbb{P}[\theta_2 \leq \bar{\theta}_2] = \frac{1}{\epsilon + 3}$$

which converges to

$$(\frac{2}{3}, \frac{1}{3})$$

as $\epsilon \to 0$, which corresponds to the exact probabilities in the mixed equilibria of the unperturbed game. Hence, we have found another way to describe mixed strategies from the perspective of playing pure strategies in perturbed Bayesian games.

As an outsider, it looks like players are randomising, but instead, they are in fact picking pure strategies since they know their own private valuations.

10.4.1 Purification Theorem

We can in fact generalise our result.

> **Theorem 532** (Purification (Harsanyi 1973))
> For almost all finite complete information games Γ, there exists a sequence of perturbed games Γ_ϵ such that any **Nash equilibrium** of Γ is the limit as $\epsilon \to 0$ of a sequence of **pure Bayesian equilibria** of the perturbed games Γ_ϵ.

Remark 533. *The qualification* almost all *is needed because equilibria that rely on weakly dominated strategies cannot generally be recovered this way. Such equilibria require payoff ties and arise only in a nongeneric set of games.*

If we think that complete information is a simplification and that in general, players have at least a small uncertainty about the other's payoffs, then the theorem states that the distinction between pure and mixed strategies is artificial. All of them are the limits of pure strategies of the Bayesian games players might be playing.

11 Dynamic Games of Incomplete Information

We now consider dynamic games of incomplete information. In such games, players may not know features such as their opponents' payoffs. We focus on cases in which players have privately observed types that determine payoffs. As in static games, we can transform a dynamic **incomplete-information game** into a game of **complete but imperfect information**: Nature moves first, selects the players' types according to a common prior, and privately reveals each type to the relevant player.

Because types are private information, the transformed game may contain substantial imperfect information, leaving subgame perfection with little bite. We therefore use weak perfect Bayesian equilibrium or sequential equilibrium.

We could stop here. We have already seen what a sequential equilibrium is, and hence we can analyse dynamic games of incomplete information just by transforming the game into a complete information one and look for sequential equilibria. However, as we did for static games, rather than working with the transformed game, we will introduce a new concept which will be easier to use for these incomplete information games. This is the concept of **Perfect Bayesian Equilibrium**.

Intuitively, what we would like to do is to consider an **assessment** and impose that the strategies in the assessment constitute a Bayesian Nash equilibrium in the **continuation game**, given the system of belief in the assessment. Notice that we did not say subgame, but continuation game.

We will look at a particular type of dynamic game with incomplete information called **signalling games** as seen in Fig. 125.

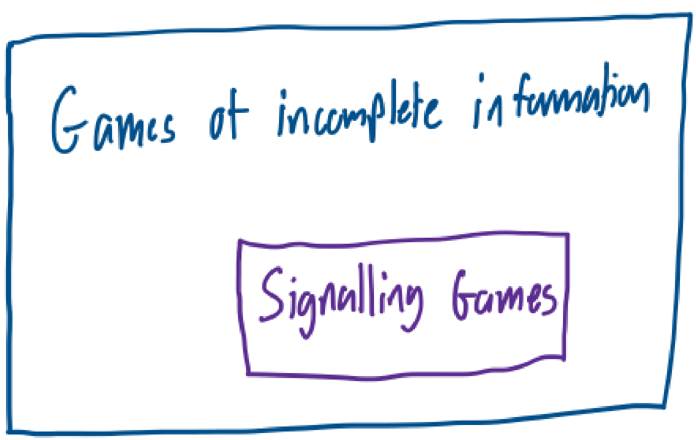

Figure 125: A signalling game is a dynamic game with incomplete information.

11.1 Signalling Games

In a signalling game there are two players. Player 1 has **private information** and is the **first** one to move. Player 2 doesn't have private information and moves second. We call these games signalling games or sender-receiver games because the action of Player 1 acts as a signal about her type.

Definition 534 (Signalling Game)
Formally, a **signalling game** has the following elements and timing:

- Player 1 has private information, modelled by different **types**, $\theta \in \Theta$. There is a common prior distribution on Θ by $p(\cdot)$.
- Player 1 chooses an action $a_1 \in A_1$.
- Player 2 doesn't have private information. She observes action a_1 and chooses $a_2 \in A_2$.
- Given a type θ and some actions a_1, a_2, the payoffs to the players are given by $u_i(a_1, a_2, \theta)$.

There are a few things to dwell on in the set up. First, only player 1 has types (player 2 only has 1 type). So the common prior distribution $p(\cdot)$ reflects the probabilities of what type player 1 can take on. Furthermore, player 2 can only observe player 1's action a_1 but not their type θ. This is important because player 2's payoff will depend on the type of player 1 in addition to the actions taken $u_2(a_1, a_2, \theta)$ and hence intuitively, player 2 will try to figure out player 1's type from their action a_1.

Remark 535. *Player 1 can also be called the* **sender** *because their action choice sends information about their type. Player 2 can then also be called the* **receiver**.

As with sequential equilibria we will need to introduce a system of belief for Player 2 alongside the strategy profile. However, we will see that there will be some differences in how we define an assessment compared to how we previously defined assessments.

Definition 536 (Assessment in Signalling Game)
An **assessment** (σ, μ) of a signalling game consists of:

1. A mixed strategy for player 1, for each possible realisation of their type θ: $\sigma_1(\cdot|\theta)$, $\forall\theta \in \Theta$.
2. A belief for player 2 after each possible action a_1: $\mu(\cdot|a_1)$, $\forall a_1 \in A_1$.
3. A mixed strategy for player 2, for each possible action a_1: $\sigma_2(\cdot|a_1)$, $\forall a_1 \in A_1$.

There are a few things to note here. First, a mixed strategy for player 1 involves choosing a strategy for every possible **type** that player 1 can take on $\theta \in \Theta$ rather than for each of player 1's information set. However, in this setting, every information set for player 1 corresponds to a particular type that player 1 can take and therefore we can switch back and forth between information sets for player 1 and player 1's types. Also recall that these types for player 1 are assigned to player 1 by nature.

The second component is player 2's belief at each information set. If player 2 has information sets I_X and I_Y, the corresponding beliefs μ_X and μ_Y must each sum to one: $\sum_{x\in I_X} \mu_X(x) = 1$ and $\sum_{y\in I_Y} \mu_Y(y) = 1$. Because player 2 observes player 1's **action** but not their type, each information set of player 2 corresponds to an action a_1. We therefore index beliefs by the observed action, writing $\mu(\cdot \mid a_1)$ for every $a_1 \in A_1$.

Finally, player 2's mixed strategy must be defined at every information set. Since each of player 2's information sets corresponds uniquely to an action $a_1 \in A_1$ by player 1, the strategy specifies what player 2 does after every possible action by player 1.

We illustrate all these definitions to Fig. 126.

We are now ready to define a **Perfect Bayesian Equilibrium**.

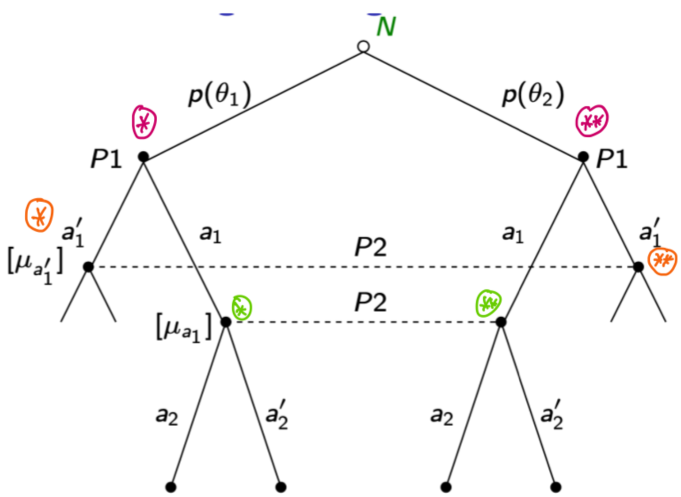

(*) and (**) correspond to 2 different information sets for P1. However, notice that information set (*) corresponds to P1 being type θ_1 and (**) correspond to P1 being type θ_2.

Nodes (*) and (**) correspond to being in the same information set for player 2. However, notice that they both are reached by the same action of P1 (playing a_1') but (*) is for type θ_1 and (**) for type θ_2. Hence, $\mu_{a_1'}$ correspond to player 2's belief that P1 is type θ_1 and $1-\mu_{a_1'}$ to type θ_2

(*) and (**) correspond to action a_1. Furthermore, μ_{a_1} is belief that P1 is type θ_1 and $1-\mu_{a_1}$ is type θ_2.

Figure 126: All concepts applied to a signalling game.

Definition 537 (Perfect Bayesian Equilibrium)

A **Perfect Bayesian equilibrium** (PBE) of a signalling game is a strategy profile σ and a system of beliefs μ such that:

1. For all θ, and all a_1^* such that $\sigma_1(a_1^*|\theta) > 0$,

$$a_1^* \in \arg\max_{a_1} \sum_{a_2} \sigma_2(a_2|a_1) u_1(a_1, a_2, \theta) \tag{207}$$

2. For all a_1, and all a_2^* such that $\sigma_2(a_2^*|a_1) > 0$,

$$a_2^* \in \arg\max_{a_2} \sum_{\theta} \mu(\theta|a_1) u_2(a_1, a_2, \theta) \tag{208}$$

3. If $\sum_{\theta' \in \Theta} p(\theta')\sigma_1(a_1|\theta') > 0$,

$$\mu(\theta|a_1) = \frac{p(\theta)\sigma_1(a_1|\theta)}{\sum_{\theta' \in \Theta} p(\theta')\sigma_1(a_1|\theta')} \tag{209}$$

Let's go through each part of a Perfect Bayesian Equilibrium.

The first part is a statement on player 1's best reply function being **sequentially rational**. It says that for all of player 1's potential types θ, and all of player 1's actions a_1^* such that they are played with positive probability by player 1, $\sigma_1(a_1^*|\theta) > 0$, we require that a_1^* maximises player 1's expected utility $u_1(a_1, a_2, \theta)$ where the expectation is taken over all possible strategies that **player 2** may take $\sigma_2(a_2|a_1)$. Hence, this is written compactly as:

$$a_1^* \in \arg\max_{a_1} \sum_{a_2} \sigma_2(a_2|a_1) u_1(a_1, a_2, \theta)$$

The second part says that player 2's strategy is also sequentially rational. After every observed action a_1, each action a_2^* played with positive probability must maximise player 2's expected utility under the posterior belief $\mu(\theta \mid a_1)$:

$$a_2^* \in \arg\max_{a_2} \sum_{\theta} \mu(\theta|a_1) u_2(a_1, a_2, \theta)$$

The final part requires consistent beliefs. If action a_1 is played with positive probability, so $\sum_{\theta' \in \Theta} p(\theta')\sigma_1(a_1 \mid \theta') > 0$, player 2 updates by Bayes' rule. The numerator is the prior probability that player 1 has type θ times the probability that this type chooses a_1; the denominator sums this joint probability over all possible types.

In signalling games, weak perfect Bayesian equilibrium and sequential equilibrium coincide, and a perfect Bayesian equilibrium can be viewed as a sequential equilibrium of the transformed game. Since every finite game has a sequential equilibrium, every finite signalling game has a PBE.

Example 538

(Reputation Game—Kreps and Wilson, 1982). Suppose that there is an incumbent monopolist and a potential challenger. There are two types of incumbent: a **crazy** type that always likes to fight on price and a **sane** type that would rather accommodate entry than fight. The entrant does not know the incumbent's type.

The structure of the game is:

1. In the first period, the incumbent can engage in a price battle (prey) or accommodate the entrant.

2. In the second period, the challenger decides whether to stay or exit the market.

If the challenger stays, the crazy incumbent keeps fighting, while a sane incumbent accommodates the challenger.

Mathematically, we represent the Incumbent (I) and Challenger (C) as playing a game over two periods. The incumbent has one of two types:

$$\theta \in \{\text{sane, crazy}\}$$

with a prior distribution of the type being given by:

$$\mathbb{P}[\text{sane}] = p.$$

In the game, the incumbent chooses $a_1 \in \{\text{prey}, \text{accommodate}\}$ and then the challenger chooses $a_2 \in \{\text{stay}, \text{exit}\}$.

We represent the game structure and payoffs in Fig. 127.

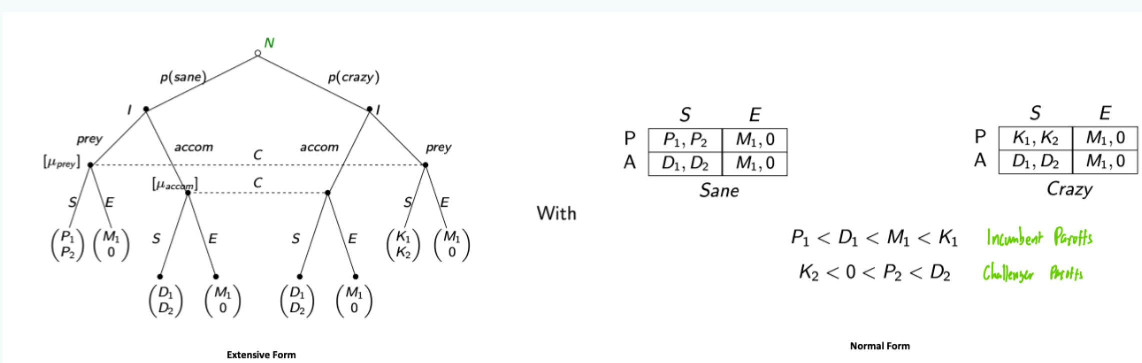

Figure 127: Extensive and normal form representation of reputation game. μ_{prey} are C's beliefs of I's type after observing I has played prey. Likewise, μ_{accom} are C's belief of I's type after observing I has played accom.

We can impose the assumption on the proportion of crazy types which we shall justify later.

Assumption (Proportion of Crazy Types)
The proportion of crazy types $(1 - p)$ is not too high:

$$pP_2 \geq -(1 - p)K_2 \tag{210}$$

If we look at the normal form representation of the payoff matrix for the crazy type in Fig. 127, we see that the crazy type has a weakly dominant strategy to prey (P) as $K_1 > D_1$ and $M_1 = M_1$. Therefore, the crazy type will always choose to prey and we can write:

$$\sigma_1(\text{prey}|\text{crazy}) = 1.$$

We now consider the sane type's choice. It would like to accommodate, but preying may convince the challenger that it is crazy and induce exit. To analyse this reputation motive, define the probability that a sane incumbent preys in the first period as:

$$\alpha \equiv \sigma_1(\text{prey}|\text{sane}). \tag{211}$$

There are three possible equilibria that can arise which we shall examine.

Separating Equilibrium

Under a separating equilibrium,

$$\alpha = \sigma_1(\text{prey} \mid \text{sane}) = 0,$$

so the sane type always accommodates. The crazy type always preys, $\sigma_1(\text{prey} \mid \text{crazy}) = 1$. The two types therefore take different actions, and the challenger learns the incumbent's type from its action.

Pooling Equilibrium

Under a pooling equilibrium,

$$\alpha = \sigma_1(\text{prey} \mid \text{sane}) = 1,$$

so the sane type always preys, just like the crazy type. Since both types choose the same action, the action transmits no information about the incumbent's type.

Semi-Separating Equilibrium

Under a semi-separating equilibrium,

$$\alpha = \sigma_1(\text{prey} \mid \text{sane}) \in (0, 1),$$

so the sane type randomises while the crazy type always preys. Observing accommodation reveals that the incumbent is sane, whereas observing prey does not reveal its type with certainty. Hence, information is transmitted only partially.

Definition 539 (Separating Equilibrium)
In a **separating equilibrium**, all types for player 1 will do different actions.

Definition 540 (Pooling Equilibrium)
In a **pooling equilibrium**, all types for player 1 will do the same action.

Definition 541 (Semi-Separating Equilibrium)
In a **semi-separating equilibrium**, **some** types of player 1 will do the same action whilst other types will do different actions.

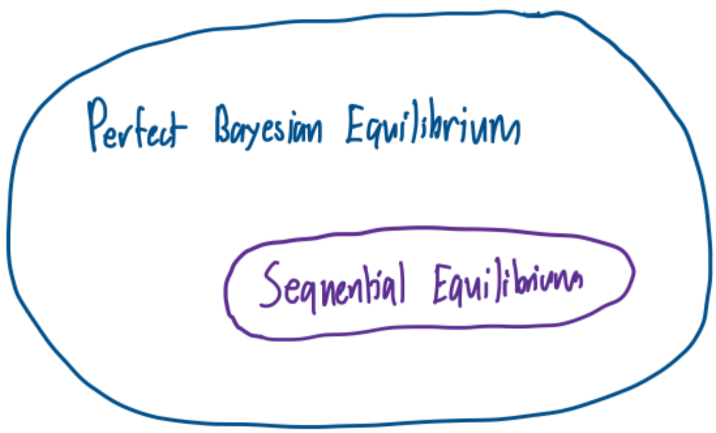

Figure 128: In general, a sequential equilibrium is a stronger concept than a perfect Bayesian equilibrium.

11.2 Multi-stage games with observed actions and incomplete information

There is a more complicated definition of PBE for multi-stage games with observed actions and incomplete information (this would be useful for example if we wanted to study repeated games with incomplete information). We will not cover this definition in this course, but if interested you can find it in FT Ch 8.2.3. The key difference is that some more consistency requirements of beliefs are added. For example, when types are *independent*, we will require all the types of a player to hold the same beliefs about others, given the same actions. We might also impose that player i and j have the same beliefs about player k, or that a player should not be able to signal something that she doesn't know.

For those cases the concept of PBE is stronger than the corresponding weak PBE of the transformed game, but weaker than the Sequential equilibrium:

$$\{\text{Sequential equilibrium}\} \subset \{\text{PBE}\} \subset \{\text{weak PBE}\}$$

Theorem 542 (Fudenberg and Tirole, 1991)
If each player has at most two types $|\Theta_i| = 2$, or there are at most two periods, then the set of Perfect Bayesian equilibrium coincides with the set of sequential equilibria.

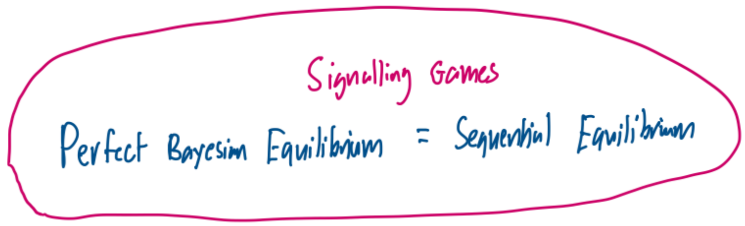

Figure 129: In signalling games, a sequential equilibrium is equivalent to a perfect Bayesian equilibrium.

11.3 Refinements

The concept of PBE (and indeed the concept of SE) does not impose any (or nearly any) restrictions on out-of-equilibrium beliefs. Sometimes though the set of reasonable beliefs on information sets that are not reached through equilibrium can be narrowed further.

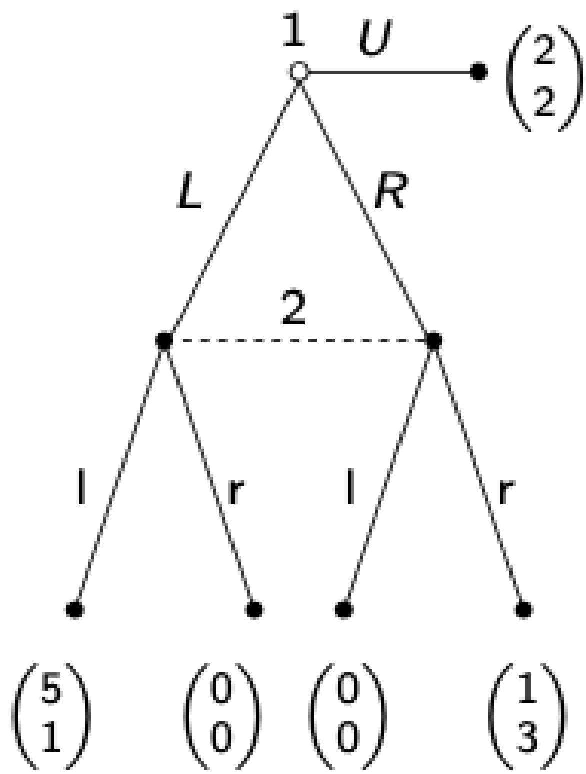

Figure 130: Extensive form game.

Look at Fig. 130. We can show that $\{(U, r); \mu = \frac{1}{2}\}$ is a sequential equilibrium. However, player 2's information set is not reached, so the belief $\mu = 0.5$ is not disciplined by on-path play. Instead of treating every deviation as a mistake, we can ask which deviations might be intentional. If player 2 is called to play, it is reasonable to infer that player 1 chose a deviation that could improve their payoff. In this game, that points to L rather than R and supports the action profile (L, l), ruling out the original assessment in which player 2 chooses r. Such belief-based refinements use *forward induction*: they evaluate what a deviator gives up, and could gain, by deviating.

If there is a subset of types that could potentially benefit from a deviation (if the others hold some particular beliefs), whereas another subset of types would always be worse off by deviating that what they get in equilibrium (independently of the beliefs held by others), then the beliefs at that information set should allocate zero weight to the latter group. If given those beliefs, the assessment is still an equilibrium then such assessment has survived this *intuitive criterion*. We now look at the intuitive criterion.

Definition 543 (Intuitive Criterion)
Formally, suppose that we have a signalling game with two players and equilibrium strategies (σ_1^*, σ_2^*). For each $\theta \in \Theta$, let

$$u_1^*(\theta) = \max_{a_1} \sum_{a_2} u_1(a_1, a_2, \theta)\sigma_2^*(a_2 \mid a_1).$$

In other words, $u_1^*(\theta)$ is the equilibrium payoff of a type θ of player 1. We say that a_1' is *equilibrium dominated at* θ, if $u_1^*(\theta) > \max_{a_2} u(a_1', a_2, \theta)$. That is, the maximum that a type θ can get by playing a_1' is smaller than what that type achieves in equilibrium.

After an off-path action a_1', the intuitive criterion requires beliefs to assign zero probability to every type for which a_1' is equilibrium dominated (provided some other type could benefit from the deviation). Thus an assessment fails this belief restriction if $\mu(\theta \mid a_1') > 0$ for such a dominated type.

Remark 544. *The intuitive criterion states that upon observing a deviation from the equilibrium, put zero weights on types whose equilibrium payoff exceeds all possible payoffs from deviating.*

We give an example to apply the intuitive criterion.

Example 545
(Beer-Quiche Game; Cho and Kreps, 1987).

Suppose we have a pub-goer (P1) who orders either Beer or Quiche. A bully (P2) sees this and decides whether to fight him or not. P2 wants to fight if P1 is weak but not when P1 is strong. Only P1 knows his own type. Let μ_B be player 2's posterior probability that player 1 is strong after observing Beer, and let μ_Q be the corresponding posterior after observing Quiche. The prior probabilities are $p(w) = 0.2$ and $p(s) = 0.8$. The weak type prefers Quiche, the strong type prefers Beer, and neither type wants to fight. We represent the game in Fig. 131. Thus, player 1's choice may signal their type to player 2.

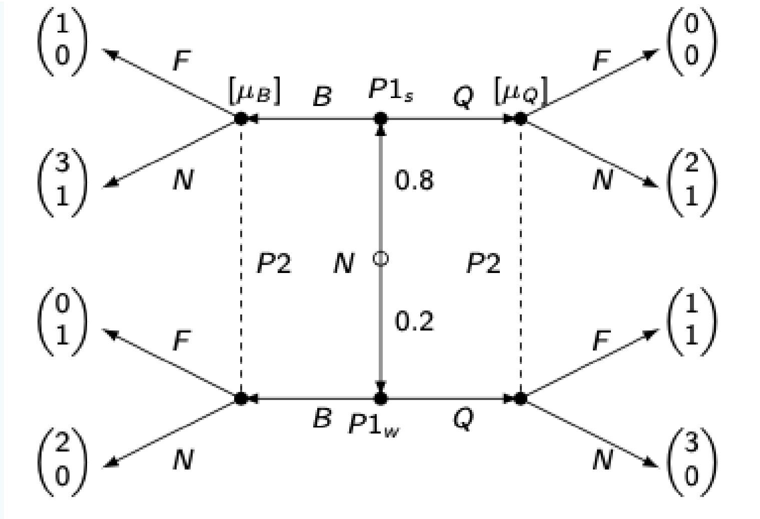

Figure 131: Beer-quiche extensive form game.

The strategy for P1 is to choose an action based on their type:

$$a_1(\theta) \in \{B, Q\}$$

for either type that they may be $\theta \in \{s, w\}$. The strategy for P2 is to choose whether to fight or not after having observed P1's action:

$$a_2(a_1) \in \{F, N\}$$

for either action $a_1 \in \{B, Q\}$. We now find the perfect Bayesian equilibria of this game.

Separating Equilibrium

In a separating equilibrium, we first look at what happens if the strong type always orders Beer and the weak type always orders Quiche. That is:

$$a_1(s) = B, a_1(w) = Q.$$

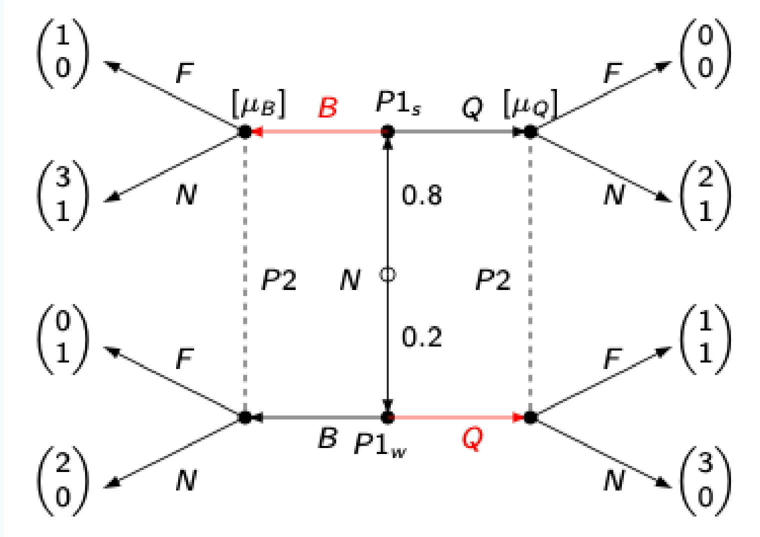

Figure 132: Beer-quiche extensive form game under Separating Equilibrium.

Now, P2 will learn in equilibrium that when a Beer is ordered, P1 is a strong type. Therefore, the beliefs that are consistent with this are:

$$\mu_B = 1, \mu_Q = 0.$$

For P2 to be sequentially rational given these beliefs, they must have the strategies:

$$a_2(B) = N, a_2(Q) = F$$

as P2 does not want to fight a strong type, which they suspect P1 is if they order a Beer.

However, from this, the weak type P1 will clearly want to deviate and order a beer too:

$$a_1(w) = B$$

as the payoff from ordering a Beer and not fighting for the weak type is 2 compared to ordering a Quiche and fighting which gives a payoff of 1. Hence, there is a deviation and therefore there is no separating equilibrium. We illustrate this in Fig. 132.

The other possible separating equilibrium is when the weak type order a Beer and the strong type orders a Quiche. However, the strong type has incentives to deviate and therefore this is not a separating equilibrium either.

Pooling Equilibrium with Both Beer

We now look at the pooling equilibrium where both types order Beer.

$$a_1(s) = a_1(w) = B.$$

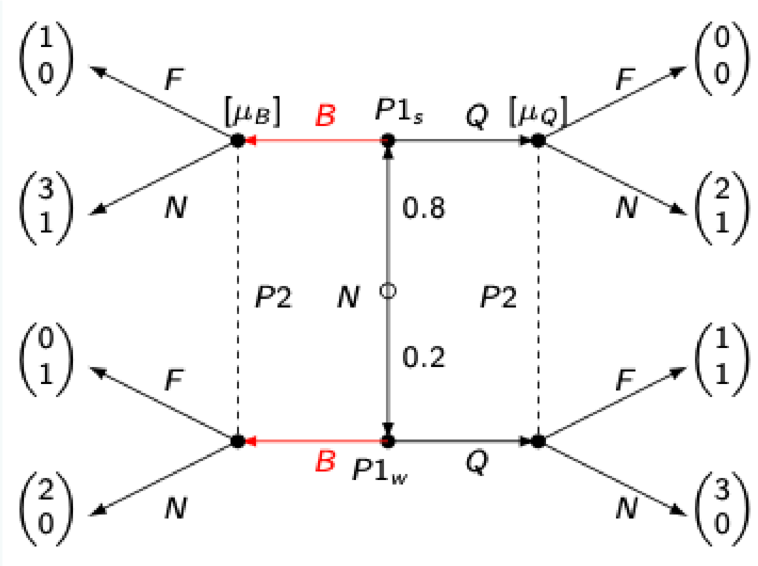

Figure 133: Beer-quiche extensive form game under first Pooling Equilibrium.

Since P2 gains no information from observing P1's actions, then they will set their beliefs identical to the prior:

$$\mu_B = 0.8.$$

However, we never reach the information set associated to the beliefs μ_Q and therefore this belief is unrestricted:

$$\mu_Q \in [0, 1].$$

We show this in Fig. 133. Now, we figure out the actions that are sequentially rational for P2 given these beliefs. If player 2 wants to fight, they will get a payoff of:

$$0.8(0) + 0.2(1) = 0.2$$

whilst if they don't want to fight, they get a payoff:

$$0.8(1) + 0.2(0) = 0.8.$$

Therefore, if P2 observes P1 ordering a Beer, they should not fight:

$$a_2(B) = N.$$

Now, if P2 observes Quiche being ordered, if they fight, they will get a payoff:

$$\mu_Q(0) + (1 - \mu_Q)(1) = 1 - \mu_Q$$

and if they don't fight, they get a payoff:

$$\mu_Q(1) + (1 - \mu_Q)(0) = \mu_Q.$$

Therefore, they will fight if and only if:

$$1 - \mu_Q \geq \mu_Q \quad \Rightarrow \mu_Q \leq \frac{1}{2}.$$

Now, if $a_2(Q) = F$, then for a weak P1 type, they would rather order a Beer:

$$a_1(W) = B$$

and hence have no incentive to deviate. Therefore, we found a Pooling Bayesian Equilibrium where both types order Beer:

$$\{a_1(s) = a_1(w) = B;\ a_2(B) = N,\ a_2(Q) = F;\ \mu_B = 0.8,\ \mu_Q \leq \tfrac{1}{2}\}.$$

Pooling Equilibrium with Both Quiche

We now look at the pooling equilibrium where both types order Quiche.

$$a_1(s) = a_1(w) = Q$$

which we show in Fig. 134.

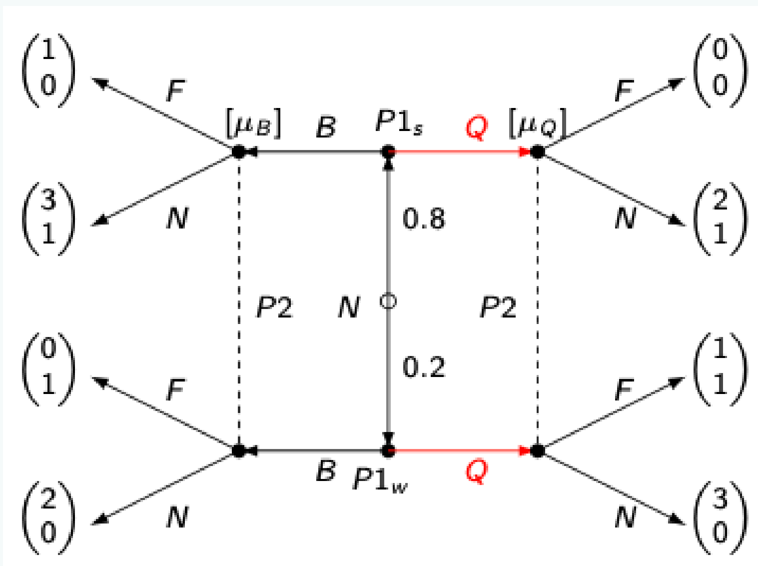

Figure 134: Beer-quiche extensive form game under second Pooling Equilibrium.

Since P2 gains no information from observing P1's actions, then they will set their beliefs identical to the prior:

$$\mu_Q = 0.8.$$

However, we never reach the information set associated to the beliefs μ_B and therefore this belief is unrestricted:

$$\mu_B \in [0, 1].$$

Now, suppose that P2 observes P1 ordering a Quiche. Then, the payoff if P2 chooses to Fight is:

$$0.8(0) + 0.2(1) = 0.2$$

and the payoff if P2 chooses to Not Fight is:

$$0.8(1) + 0.2(0) = 0.8.$$

Therefore, if P2 observes Quiche, they will Not Fight:

$$a_2(Q) = N.$$

Now, suppose P2 observes Beer. Then, the payoff they get if P2 fights is:

$$\mu_B(0) + (1 - \mu_B)(1) = 1 - \mu_B.$$

The payoff that P2 gets if P2 doesn't fight is:

$$\mu_B(1) + (1 - \mu_B)(0) = \mu_B.$$

Therefore, P2 will only fight if:

$$1 - \mu_B \geq \mu_B \quad \Rightarrow \mu_B \leq \frac{1}{2}.$$

Therefore, we have P2 will fight if they observe Beer:

$$a_2(B) = F.$$

Now, we look at the P1 Strong Type who is currently ordering Quiche. We can see that the Strong P1 has no incentive from deviating from ordering Quiche and not fighting (payoff 2) to ordering Beer and fighting (payoff of 1). Therefore, the strong type has no incentive to deviate due to fear of fighting.

Therefore, we found another Pooling Bayesian Equilibrium:

$$\{a_1(s) = a_1(w) = Q;\ a_2(B) = F,\ a_2(Q) = N;\ \mu_B \leq \tfrac{1}{2},\ \mu_Q = 0.8\}.$$

However, this is a strange PBE. This is because since $\mu_B \leq \frac{1}{2}$, that means then P2's belief is that if they observe P1 ordering a B, it is more likely that P1 is a weak type and therefore P2 should fight. However, why would the Weak type deviate from ordering Quiche (which is their preference) and order Beer instead? Therefore, the **intuitive criterion** says that we can impose that:

$$\mu_B = 1$$

where if they ever observe P1 ordering Beer, it **must** be the strong type as the weak type has no incentive to deviate from an action that is best for them.

This refinement reduces the set of equilibria in many signalling games. However, sometimes even this refinement is not powerful enough.

12 Auction Theory

12.1 Math Prerequisites

Here we recall some math results.

Definition 546 (Expectation of a Random Variable)
Let X be a continuous random variable with density f on $\mathbb{R}$ and a finite expectation. Then:

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} xf(x)dx. \tag{212}$$

We can extend this definition to functions of random variables.

Proposition 547 (Expectation of a Continuous Function of a Random Variable)
Let X be a random variable and let $g(\cdot)$ be a continuous function. Then the expectation of the function $g(X)$ is:

$$\mathbb{E}[g(X)] = \int_{-\infty}^{\infty} g(x)f(x)dx. \tag{213}$$

We now derive conditional expectations.

Definition 548 (Conditional Expectation)
Let X be a continuous random variable. For $F(x) > 0$, its lower-tail conditional expectation is:

$$\mathbb{E}[X|X < x] = \frac{1}{F(x)} \int_{-\infty}^{x} tf(t)dt \tag{214}$$

where $F(\cdot)$ is the CDF of X.

Proof. First recall that the definition of conditional density is:

$$f_Y(y|X = x) = \frac{f(x,y)}{f_X(x)}$$

where $f(x,y)$ is the joint density and $f_X(x)$ is the marginal density. We have something similar for conditional expectation.

We are interested in the expected value of X given that $X \in (-\infty, x)$. Denote this event by $A = (-\infty, x)$. We then define the indicator function:

$$1_A(t) = \begin{cases} 1 & \text{if } t \in (-\infty, x) \\ 0 & \text{otherwise.} \end{cases}$$

Now, we can write the conditional expectation as:

$$
\begin{aligned}
\mathbb{E}[X|A] &= \frac{\mathbb{E}[X \cdot 1_A]}{\mathbb{P}(A)} \\
&= \frac{\int_{-\infty}^{\infty} \Big(t \cdot 1_A(t)\Big) f(t)dt}{\mathbb{P}(A)} \\
&= \frac{\int_{-\infty}^{x} tf(t)dt}{\mathbb{P}(A)} \\
&= \frac{\int_{-\infty}^{x} tf(t)dt}{\int_{-\infty}^{x} f(t)dt} \\
&= \frac{\int_{-\infty}^{x} tf(t)dt}{F(x)}.
\end{aligned}
$$

□

Theorem 549 (Corollary of Conditional Expectation)

We can express conditional expectation as:

$$F(x)\mathbb{E}[X \mid X < x] = xF(x) - \int_{-\infty}^{x} F(t)dt, \tag{215}$$

Proof. First recall that we showed:

$$\mathbb{E}[X|X < x] = \frac{1}{F(x)} \int_{-\infty}^{x} tf(t)dt.$$

Multiply by $F(x)$ on both sides:

$$F(x)\mathbb{E}[X|X < x] = \int_{-\infty}^{x} tf(t)dt.$$

We can use integration by parts:

$$F(x)\mathbb{E}[X|X < x] = xF(x) - \int_{-\infty}^{x} F(t)dt.$$

□

We now look at order statistics. Suppose that $X_1, \dots, X_n$ are i.i.d. random variables. Write $X_{(1)}$ for the smallest observation and $X_{(n)}$ for the largest; in general, $X_{(k)}$ denotes the k-th smallest observation. When descending ranks are useful, we write $X_{[k]}$ for the k-th highest observation, so $X_{[k]} = X_{(n+1-k)}$.

Definition 550 (Order Statistic)

Let $X_1, \dots, X_n$ be continuous i.i.d. random variables with density f and cumulative distribution function F. Denote the k-th smallest order statistic by $X_{(k)}$. Its cumulative distribution function is:

$$\mathbb{P}(X_{(k)} \leq x) = \sum_{j=k}^{n} \binom{n}{j} F(x)^j (1 - F(x))^{n-j}. \tag{216}$$

The probability density function of $X_{(k)}$ is given by:

$$f_{X_{(k)}}(x) = n\binom{n-1}{k-1} F(x)^{k-1}(1 - F(x))^{n-k} f(x). \tag{217}$$

We can think about the CDF intuitively:

$$\mathbb{P}(X_{(k)} \leq x) = \sum_{j=k}^{n} \binom{n}{j} F(x)^j (1 - F(x))^{n-j}.$$

First, we want to know that the probability that:

$$\mathbb{P}[X_{(k)} \leq x] = \mathbb{P}\Big[\text{at least k of } X_i \leq x\Big].$$

This is a binomial probability: a "success" is the event $X_i \leq x$, which has probability $F(x)$. For each $j \geq k$, there are $\binom{n}{j}$ ways to choose the j observations below x, and we sum from j=k to n.

Now, let us look at the probability density function:

$$f_{X_{(k)}}(x) = n\binom{n-1}{k-1} F(x)^{k-1}(1 - F(x))^{n-k} f(x).$$

Fix $x \in \mathbb{R}$ and a small interval around x. Any one of the n observations can lie in this interval. Of the remaining $n-1$ observations, choose $k-1$ to lie below x, giving the factor $\binom{n-1}{k-1}F(x)^{k-1}$; the remaining $n-k$ lie above x, giving $(1 - F(x))^{n-k}$.

Theorem 551 (Maximum Order Statistic)
Let $X_{(n)}$ be the maximum order statistic. Its cumulative distribution function is:

$$F_{X_n}(x) = F(x)^n. \tag{218}$$

The associated probability density function is:

$$f_{X_{(n)}}(x) = nF(x)^{n-1} f(x). \tag{219}$$

Proof. Apply the cumulative distribution formula and the probability density function formula. □

Theorem 552 (Second-Highest Order Statistic)
Let $X_{(n-1)}$ be the second-highest order statistic. Its cumulative distribution function is:

$$F_{X_{(n-1)}}(x) = F(x)^n + nF(x)^{n-1}(1 - F(x)). \tag{220}$$

The associated probability density function is:

$$f_{X_{(n-1)}}(x) = n(n-1)F(x)^{n-2}(1 - F(x))f(x). \tag{221}$$

Proof. For the cumulative distribution, we use the formula to get for k = n:

$$\binom{n}{n} F(x)^n = F(x)^n.$$

For $k = n - 1$, we have:

$$\binom{n}{n-1} F(x)^{n-1}(1 - F(x)) = nF(x)^{n-1}(1 - F(x))$$

where we use the binomial-coefficient formula $\binom{n}{k} = \frac{n!}{k!(n-k)!}$, so that:

$$\binom{n}{n-1} = \frac{n!}{(n-1)!(n-n+1)!} = \frac{n!}{(n-1)!} = n.$$

For the probability density function:

$$\begin{aligned} n\binom{n-1}{n-2}F(x)^{n-2}(1-F(x))f(x) &= n\frac{(n-1)!}{(n-2)!}F(x)^{n-2}(1-F(x))f(x) \\ &= n(n-1)F(x)^{n-2}(1-F(x))f(x). \end{aligned}$$

□

We now want to apply order statistics to uniform random variables.

Definition 553 (Standard Uniform Random Variable)
The random variable $X \sim U[0,1]$ if it has the probability density function:

$$f(x) = 1 \quad \text{for } x \in [0,1]. \tag{222}$$

Theorem 554 (Uniform Order Statistics and Beta Random Variables)
Let $U_1, \ldots, U_n$ be i.i.d. $U[0,1]$ random variables. Let $U_{(k)}$ be the k-th smallest order statistic. Then $U_{(k)} \sim \text{Beta}(k, n+1-k)$.

Proof. Using the density formula for the k-th smallest order statistic:

$$f_{X_{(k)}}(x) = n\binom{n-1}{k-1}F(x)^{k-1}(1-F(x))^{n-k}f(x)$$

where $X \sim U[0,1]$ so we end up with:

$$f_{X_{(k)}}(x) = n\binom{n-1}{k-1}x^{k-1}(1-x)^{n-k}$$

where $x \in (0,1)$. This is the probability density function of a Beta random variable. □

Proposition 555 (Expectation of Beta Random Variable)
Let $X \sim Beta(k, n-k+1)$. The expectation of the random variable X is:

$$\mathbb{E}[X] = \frac{k}{n+1}. \tag{223}$$

Proof. The expectation of a Beta distribution is:

$$\mathbb{E}[X] = \frac{k}{k+(n-k+1)} = \frac{k}{n+1}.$$

□

Definition 556 (Uniform Random Variable)
The random variable V has a uniform distribution $U[\underline{V}, \overline{V}]$ if it has the cumulative distribution function:

$$F(v) = \frac{v - \underline{V}}{\overline{V} - \underline{V}}. \tag{224}$$

The associated probability density function is:

$$f(v) = \frac{1}{\overline{V} - \underline{V}} \tag{225}$$

for $v \in (\underline{V}, \overline{V})$ and 0 otherwise.

We now come to a result that we will need often.

Theorem 557 (Expected Value of K-th Highest Uniform Order Statistic)
Suppose that $X_1, \ldots, X_n$ are i.i.d. $U[\underline{V}, \overline{V}]$. Let $X_{[k]}$ be the k-th highest order statistic. Then:

$$\mathbb{E}[X_{[k]}] = \underline{V} + \left(\frac{n+1-k}{n+1}\right)(\overline{V} - \underline{V}). \tag{226}$$

12.2 Independent Private Values Auctions

We are interested in auctions where each bidder has an independent valuation for a single object on offer. Furthermore, these valuations are private information to each bidder. Mathematically, we have n bidders competing for a **single unit**. Bidder i values the unit at v_i, which is **private information** to them. We make some assumptions surrounding the bidders before we proceed.

Assumption (Distribution of Bidders' Values)
The values $v_1, \ldots, v_n$ are drawn **independently** from the same continuous distribution $F(v)$ on $[\underline{V}, \overline{V}]$, with $F(\underline{V}) = 0$.

Remark 558. *This distribution that all values are drawn from is common knowledge to all players.*

Assumption (Risk-Neutrality of Bidders)
All bidders are risk-neutral.

Each bidder therefore maximises expected surplus. We ask how bidders behave under different auction formats and what expected revenue each format generates for the seller.

We have 4 different types of auctions we are interested in. Note that for each auction, it is important we specify the rules of the auction, the allocation rule (who is awarded the object) and the payment rule (what do bidders pay).

Definition 559 (First-Price Sealed Bid Auction)
Each bidder submits a single bid without observing competitors' bids. The highest bidder wins and pays their own bid.

Definition 560 (Dutch Open Descending Auction)
The price starts high and gradually falls until one bidder agrees to buy the object at the price.

We can show that these two auctions are actually equivalent to one another.

Definition 561 (Strategically Equivalent Games)
Two games are strategically equivalent if they have the same *normal form representation*.

Remark 562. *Heuristically, this means that for every strategy in one game, a player has a strategy in the other game which results in the same outcome.*

Proposition 563 (Equivalence of Open Descending and First-Price Sealed Bid Auction)
The Dutch open descending price auction is strategically equivalent to the first-price sealed bid auction.

Proof. In a first-price sealed bid auction, a bidder's strategy maps his private information into a bid.

In a Dutch auction, one gains no information from playing the game on other player's private values. This is because the only information that is available is that a bidder has agreed to buy the good at the current price. However, once that has happened, then the auction ends.

We can therefore conclude for every strategy in a first-price auction, there is an equivalent strategy in the Dutch auction and vice-versa. □

We now look at two more auctions.

Definition 564 (Japanese Ascending Auction)
The price starts low and gradually rises until only one bidder remains. The remaining bidder pays the price at which the penultimate bidder dropped out.

Remark 565. *We assume that once a bidder drops out of the bidding process, they cannot enter back into the bidding process.*

Remark 566. *The Japanese auction is a special case of an English auction in which the price rises continuously until only one bidder remains. A general English auction need not specify this precise price-adjustment rule.*

Definition 567 (Second-Price Sealed Bid Auction)
Bidders simultaneously write down their bids. Highest bidder wins the object and pays the second highest bid made.

We can show that these two auctions are also equivalent.

Proposition 568 (Japanese Ascending Price and Second-Price Sealed Bid Auction)
Under private values, the Japanese ascending price auction is strategically equivalent to the second-price sealed bid auction.

Proof. In the Japanese ascending price auction, we gain information as other bidders drop out. Through observing this, we can infer something about the other bidder's privately known information. It is clear that one should continue bidding up until they reach their value for the good and then drop out, as staying in will cause only negative utility if they do win. Similarly, in the second-price sealed bid auction, it is a dominant strategy to bid exactly their private value. Hence, these two auctions have optimal strategies only if the values are private. □

These are two measures for how we evaluate auctions.

Definition 569 (Criteria for Evaluating Auctions)
The first criterion to evaluate an auction is the **revenue** which is the expected selling price that the auction will fetch.

The second criterion is the **efficiency** of the auction, which measures who ends up with the object.

Before proceeding, recall that every bidder's value v_i is drawn independently from the **same distribution** $F(\cdot)$, and this is common knowledge. It is therefore natural to focus on symmetric equilibria.

Definition 570 (Symmetric Equilibrium)
A symmetric equilibrium is an equilibrium where all players use the same (mixed) strategy.

12.3 First Price Sealed Bid Auction

We first look at the first price sealed bid auction. Remember that this is strategically equivalent to the Dutch open descending auction so the results in this section will hold for that auction too.

We now determine what a bidder of value v bids in a symmetric Nash equilibrium. We first need an assumption.

Assumption (Continuous and Increasing Bidding Function)
The bidding function $b(v)$ is a continuous strictly increasing function of their value v.

We can state the payoffs for the bidder under this auction setting.

Definition 571 (Payoff For the Bidder under First Price Sealed Bid Auction)
Suppose we have n bidders. For bidder i with value v_i and bid $b(v_i)$, their payoffs in a first price sealed bid auction is:

$$\Pi_i = \begin{cases} v_i - b(v_i) & \text{when } b(v_i) > \max_{j \neq i} b(v_j) \\ 0 & \text{otherwise.} \end{cases} \tag{227}$$

We will often need to analyse the highest order statistic of the other bidders' values.

Definition 572 (Highest Remaining Order Statistic)
Suppose we have n bidders with i.i.d values $\{v_1, ..., v_n\}$ drawn from the distribution function $F(\cdot)$. Fix a bidder i and remove their value v_i from the n values $\{v_1, v_2, ..., v_{i-1}, v_{i+1}, ..., v_n\}$. Then, the highest order statistic among the remaining values is:

$$v_{[1]}^{-i} = \max_{j \neq i}\{v_j\}. \tag{228}$$

The cumulative distribution function of $v_{[1]}^{-i}$ is:

$$F_{v_{[1]}^{-i}}(x) = \left[F(x)\right]^{n-1} \tag{229}$$

and the associated probability density function is:

$$f_{v_{[1]}^{-i}}(x) = (n-1)\left[F(x)\right]^{n-2} f(x). \tag{230}$$

Proof. The cumulative distribution function of $v_{[1]}^{-i}$,

$$F_{v_{[1]}^{-i}}(x) = F(x)^{n-1},$$

comes from the cumulative distribution function of the maximum order statistic:

$$F_{X_{(n)}}(x) = F(x)^n$$

where we now have $n-1$ observations. □

We can now state the probability of a bidder winning in a first price sealed bid auction.

Proposition 573 (Probability of Bidder Winning First Price Sealed Bid Auction)
Suppose we have n bidders participating in a first price sealed bid auction. Let i be a bidder with valuation v_i and bidding function $b(v_i)$. Then the probability that bidder i will win the auction in equilibrium $P_i(v_i)$ is given by:

$$P_i(v_i) = \left[F(v_i)\right]^{n-1}. \tag{231}$$

Proof. Bidder i will win the auction if they have the highest bid $b(v_i)$:

$$b(v_i) \geq \max_{j \neq i} b(v_j).$$

By assumption, the bidding function is a continuous increasing function so then:

$$\max_{j \neq i} b(v_j) = b(\max_{j \neq i} v_j) = b(v_{[1]}^{-i})$$

based on our earlier definition of $v_{[1]}^{-i}$. We can therefore rewrite the condition for bidder i to win:

$$\begin{aligned} & b(v_i) \geq \max_{j \neq i} b(v_j) \\ & \leftrightarrow b(v_i) \geq b(v_{[1]}^{-i}) \\ & \leftrightarrow v_i \geq v_{[1]}^{-i}. \end{aligned}$$

The last equivalence follows because $b(\cdot)$ is strictly increasing. The probability of this event is the CDF of $v_{[1]}^{-i}$:

$$\mathbb{P}\left[v_{[1]}^{-i} \leq v_i\right] = F_{v_{[1]}^{-i}}(v_i) = \left[F(v_i)\right]^{n-1}.$$

Therefore, the probability that bidder i with value v_i wins is:

$$P_i(v_i) = \left[F(v_i)\right]^{n-1}.$$

□

We know the payoffs for the bidders and their probability of winning. We can therefore retrieve their expected surplus (payoff).

Proposition 574 (Expected Surplus in First-Price Sealed Bid Auction)
The expected surplus for bidder i with value v_i is:

$$S(v_i) = \big(v_i - b(v_i)\big) \times \big[F(v_i)\big]^{n-1}. \tag{232}$$

Proof. Expected surplus is the payoff from winning Π_i times the probability that you win $P_i(v_i)$. We therefore plug in our earlier results:

$$\begin{aligned} \Pi_i &= v_i - b(v_i) \quad \text{i has highest bid} \\ P_i(v_i) &= \big[F(v_i)\big]^{n-1}. \end{aligned}$$

□

The tricky part to analysing this is that if bidder i bids their true value v_i, they will get a payoff of 0 if they win. Furthermore, they will bid not above their true value v_i as they will get a negative payoff if they win. Therefore, they will be bidding below their true value v_i. However, there is a tradeoff whereby the higher their bid, the higher their probability of winning the auction is but they will have reduced their gains from winning their auction. Therefore, the optimal bid for a bidder in a first price sealed bid auction is trading off the probability of winning against the gains from winning.

We can in fact state the optimal bid in such a situation.

Theorem 575 (Optimal Bid in First Price Sealed Bid Auction)
Let i be a bidder with value v_i. Then, the optimal bid for i in a first price sealed bid auction is:

$$b(v_i) = \mathbb{E}[v_{[1]}^{-i} \mid v_{[1]}^{-i} \leq v_i] \tag{233}$$

where $v_{[1]}^{-i} = \max_{j \neq i}\{v_j\}$.

Proof. See Krishna textbook chapter 2. □

The next result characterises how far bidder i shades their bid below the true value v_i.

Theorem 576 (Bid Shading in First Price Sealed Auction)
Let i be a bidder with value v_i. Then, the amount they will bid in equilibrium in a first price sealed auction is:

$$b(v_i) = v_i - \int_{\underline{V}}^{v_i} \left[\frac{F(y)}{F(v_i)}\right]^{n-1} dy. \tag{234}$$

Proof. Let $G(x) = F_{v_{[1]}^{-i}}(x) = [F(x)]^{n-1}$. Integration by parts gives the lower-tail conditional-mean identity

$$\mathbb{E}[X \mid X \leq x] = x - \int_{\underline{V}}^{x} \frac{G(y)}{G(x)}\, dy$$

for a random variable X with lower support point $\underline{V}$ and CDF G. Applying this to $X = v_{[1]}^{-i}$ and $x = v_i$ yields

$$\begin{aligned} b(v_i) &= \mathbb{E}[v_{[1]}^{-i} \mid v_{[1]}^{-i} \leq v_i] \\ &= v_i - \int_{\underline{V}}^{v_i} \frac{[F(y)]^{n-1}}{[F(v_i)]^{n-1}} \, dy \\ &= v_i - \int_{\underline{V}}^{v_i} \left[\frac{F(y)}{F(v_i)} \right]^{n-1} dy. \end{aligned}$$

□

Consequently, bidder i's interim expected payment in the first-price auction is

$$\mathbb{E}[m_1(v_i)] = [F(v_i)]^{n-1} b(v_i) = [F(v_i)]^{n-1} v_i - \int_{\underline{V}}^{v_i} [F(y)]^{n-1} \, dy. \tag{235}$$

Bidder i bids below their true value because the second term is positive. This difference is the degree of bid shading. As the number of bidders n increases, the bid moves closer to v_i because competition raises the cost of shading.

12.3.1 Mimicking Strategy

We now give a different derivation. Suppose all other bidders use the equilibrium function $b(v)$ and bidder i submits $\tilde{b}$. Let $\tilde{v}$ be the type whose equilibrium bid equals this deviation:

$$b(\tilde{v}) = \tilde{b}.$$

We ask which value $\tilde{v}$ bidder i should mimic in a first-price sealed-bid auction.

Proposition 577 (Expected Surplus Under Mimicking in a First-Price Sealed-Bid Auction)
The expected surplus for bidder i of value v_i from mimicking type $\tilde{v}$ is:

$$S(v_i, \tilde{v}) = (v_i - b(\tilde{v})) \big[F(\tilde{v}) \big]^{n-1}. \tag{236}$$

Proof. If bidder i bids $\tilde{b} = b(\tilde{v})$, they mimic how a type $\tilde{v}$ bids. The probability that i beats any other bidder k is the probability that k's valuation is below $\tilde{v}$:

$$v_k \leq \tilde{v}$$

because bidding function are increasing in type so this would imply that:

$$b(v_k) \leq b(\tilde{v}) = \tilde{b}.$$

We know the distribution of v_k is $F(\cdot)$ so the probability of this happening is:

$$\mathbb{P}\big[v_k \leq \tilde{v}\big] = F(\tilde{v}).$$

Therefore, if bidder i mimics $\tilde{v}$, they will only win the object by beating all other $(n-1)$ bidders with probability:

$$\big[F(\tilde{v})\big]^{n-1}.$$

The surplus if bidder i wins is their true value less the bid associated with the mimicked type:

$$v_i - \tilde{b} = v_i - b(\tilde{v}).$$

Therefore, expected surplus is the surplus if they win times the probability that they win.

□

The expected surplus again displays the tradeoff between the probability of winning, $[F(\tilde{v})]^{n-1}$, and the surplus conditional on winning, $v_i - b(\tilde{v})$.

Theorem 578 (Optimal Bidding Function in First Price Sealed Bid Auction)
The optimal bidding function $b(\cdot)$ must satisfy the first-order differential equation:

$$b'(v_i) = (v_i - b(v_i))(n-1)\frac{f(v_i)}{F(v_i)} \tag{237}$$

$$b(\underline{V}) = \underline{V}. \tag{238}$$

Proof. Look at the expected surplus:

$$S(v_i, \tilde{v}) = (v_i - b(\tilde{v}))\big[F(\tilde{v})\big]^{n-1}.$$

Differentiate with respect to $\tilde{v}$:

$$\frac{\partial S(v_i, \tilde{v})}{\partial \tilde{v}} = -b'(\tilde{v})\big[F(\tilde{v})\big]^{n-1} + (v_i - b(\tilde{v}))(n-1)\big[F(\tilde{v})\big]^{n-2} f(\tilde{v}).$$

In a symmetric equilibrium, bidder i optimally mimics their own type, so $\tilde{v} = v_i$ and

$$\frac{\partial S(v_i, \tilde{v})}{\partial \tilde{v}} = 0.$$

The differential equation then becomes:

$$b'(v_i) = (v_i - b(v_i))(n-1)\frac{f(v_i)}{F(v_i)}.$$

The boundary condition

$$b(\underline{V}) = \underline{V}$$

is the limiting bid of the lowest type implied by the equilibrium formula (equivalently, the usual boundary condition when $\underline{V}$ is the lowest admissible bid). □

To proceed further, we need to impose assumptions on the distribution of the valuations v_i. We will use a uniform $U[\underline{V}, \overline{V}]$ to do so.

Theorem 579 (Bidding Strategy Under Uniform Values)
Assume that the agent's value $v_i \sim U[\underline{V}, \overline{V}]$. Then, the differential equation becomes:

$$b'(v_i) = (v_i - b(v_i))(n-1)\left(\frac{1}{v_i - \underline{V}}\right) \tag{239}$$

The optimal bidding strategy that satisfies the differential equation is:

$$b(v_i) = \underline{V} + \left(\frac{n-1}{n}\right)(v_i - \underline{V}). \tag{240}$$

Proof. We guess that the bidding function $b(\cdot)$ is a linear function of value v_i:

$$b(v_i) = Av_i + B.$$

We can see that:

$$b'(v_i) = A.$$

Now, we plug this guess into:

$$b'(v_i) = (v_i - b(v_i))(n-1)\left(\frac{1}{v_i - \underline{V}}\right)$$

and solve for A and B.

$$\begin{aligned}
& b'(v_i) = (v_i - b(v_i))(n-1)\left(\frac{1}{v_i - \underline{V}}\right) \\
& \leftrightarrow A = (v_i - \left[Av_i + B\right])(n-1)\left(\frac{1}{v_i - \underline{V}}\right) \\
& \leftrightarrow A = (v_i(1-A) - B)(n-1)\left(\frac{1}{v_i - \underline{V}}\right) \\
& \leftrightarrow A(v_i - \underline{V}) = (v_i(1-A) - B)(n-1) \\
& \leftrightarrow Av_i - A\underline{V} = (1-A)(n-1)v_i - (n-1)B
\end{aligned}$$

Now, we equate coefficients of v_i and the constant on both sides to get:

$$\begin{aligned}
& A = (1-A)(n-1) \\
& -A\underline{V} = -(n-1)B
\end{aligned}$$

We solve for the coefficient A in the first line:

$$\begin{aligned}
& A = (1-A)(n-1) \\
& \leftrightarrow A + A(n-1) = (n-1) \\
& \leftrightarrow A = \frac{(n-1)}{1+n-1} \\
& \Rightarrow A = \frac{(n-1)}{n}
\end{aligned}$$

We now solve for the coefficient of B:

$$\begin{aligned}
& -A\underline{V} = -(n-1)B \\
& \leftrightarrow -\left[\frac{(n-1)}{n}\right]\underline{V} = -(n-1)B \\
& \leftrightarrow \left[\frac{(n-1)}{n}\right]\underline{V} = (n-1)B \\
& \Rightarrow B = \frac{\underline{V}}{n}.
\end{aligned}$$

We plug these values into the bidding function $b(\cdot)$ to get:

$$b(v_i) = \left(\frac{n-1}{n}\right)v_i + \frac{\underline{V}}{n}.$$

Now, we do a final trick to get it into the right form:

$$
\begin{aligned}
b(v_i) &= \left(\frac{n-1}{n}\right) v_i + \frac{\underline{V}}{n} \\
\leftrightarrow b(v_i) &= \left(\frac{n-1}{n}\right) v_i - \frac{(n-1)\underline{V} - n\underline{V}}{n} \\
\leftrightarrow b(v_i) &= \left(\frac{n-1}{n}\right)(v_i - \underline{V}) + \frac{n\underline{V}}{n} \\
\Rightarrow b(v_i) &= \left(\frac{n-1}{n}\right)(v_i - \underline{V}) + \underline{V}
\end{aligned}
$$

□

Now that we have the amount that an agent will bid, we can then derive how much expected revenue the seller expects to receive.

Theorem 580 (Expected Revenue of Seller in First Price Sealed Bid Auction under Uniform Values)
Suppose we have n bidders with i.i.d valuations $v_i \sim U[\underline{V}, \overline{V}]$. Then the expected revenue $\mathbb{E}[R]$ of the seller in a first price sealed bid auction is:

$$
\mathbb{E}[R] = \underline{V} + \left(\frac{n-1}{n+1}\right)(\overline{V} - \underline{V}). \tag{241}
$$

Proof. Intuitively, the seller's expected revenue is the expected maximum bid by the bidders. These bids $b(\cdot)$ are functions of random variables of types v_i. Therefore:

$$
\begin{aligned}
\mathbb{E}[R] &= \mathbb{E}\left[\max_{i=1,\dots,n} b(v_i)\right] \\
&= \mathbb{E}\left[b(\max_{i=1,\dots,n} v_i)\right] =_{(1)} \mathbb{E}\left[\underline{V} + \left(\frac{n-1}{n}\right)\left(\max_{i=1,\dots,n} v_i - \underline{V}\right)\right] \\
&= \underline{V} + \left(\frac{n-1}{n}\right)\left(\mathbb{E}\left[\max_{i=1,\dots,n} v_i\right] - \underline{V}\right)
\end{aligned}
$$

where in $=_{(1)}$ we used the equilibrium bidding function. From before, the expected k-th highest order statistic $v_{[k]}$ is:

$$
\mathbb{E}[v_{[k]}] = \underline{V} + \left(\frac{n+1-k}{n+1}\right)(\overline{V} - \underline{V}).
$$

The highest valuation corresponds to $k = 1$:

$$
\mathbb{E}\left[\max_{i=1,\dots,n} v_i\right] = \mathbb{E}[v_{[1]}] = \underline{V} + \left(\frac{n}{n+1}\right)(\overline{V} - \underline{V}).
$$

Plug this highest valuation into seller's expected revenue:

$$\begin{aligned}
\mathbb{E}[R_1] &= \mathbb{E}\left[\max_{i=1,\dots,n} b(v_i)\right] \\
&= \mathbb{E}\left[\underline{V} + \left(\frac{n-1}{n}\right)\left(\max_{i=1,\dots,n} v_i - \underline{V}\right)\right] \\
&=_{(1)} \underline{V} + \left(\frac{n-1}{n}\right)\left(\mathbb{E}\left[\max_{i=1,\dots,n} v_i\right] - \underline{V}\right) \\
&=_{(2)} \underline{V} + \left(\frac{n-1}{n}\right)\left(\left(\underline{V} + \left(\frac{n}{n+1}\right)(\overline{V} - \underline{V})\right) - \underline{V}\right) \\
&= \underline{V} + \left(\frac{n-1}{n}\right)\left(\left(\frac{n}{n+1}\right)(\overline{V} - \underline{V})\right) \\
&=_{(3)} \underline{V} + \left(\frac{n-1}{n+1}\right)\left(\overline{V} - \underline{V}\right)
\end{aligned}$$

$=_{(1)}$ all terms are constants except v_i, $=_{(2)}$ substitute in the highest valuation expression, and $=_{(3)}$ where n terms cancel out.

□

We can now derive the expected payment by a bidder.

Proposition 581 (Ex Ante Expected Payment in a First-Price Sealed-Bid Auction)

Suppose there are n bidders with i.i.d. valuations $v_i \sim U[\underline{V}, \overline{V}]$. By symmetry, the ex ante expected payment of each bidder is:

$$\mathbb{E}[m_i] = \frac{\underline{V}}{n} + \left(\frac{n-1}{n(n+1)}\right)\left(\overline{V} - \underline{V}\right). \tag{242}$$

Proof. Expected payment is equal to expected revenue divided by n. We know that expected revenue in a first price sealed bid auction is:

$$\mathbb{E}[R] = \underline{V} + \left(\frac{n-1}{n+1}\right)(\overline{V} - \underline{V}).$$

We then divide this by n to get expected payment:

$$\mathbb{E}[m_i] = \frac{\underline{V}}{n} + \left(\frac{n-1}{n(n+1)}\right)\left(\overline{V} - \underline{V}\right).$$

□

12.4 Second Price Sealed Bid Auction

Suppose there is one seller of an indivisible good and n potential buyers. Bidders submit bids simultaneously. The highest bidder wins and pays the **second-highest bid**; losing bidders pay nothing. If bidder i wins, their payoff is $v_i - \max_{j \neq i} b_j(v_j)$, and otherwise it is zero. We omit ties, which occur with probability zero under continuous values and strictly increasing strategies.

Definition 582 (Payoff For the Bidder under Second Price Sealed Bid Auction)

Suppose we have n bidders. For bidder i with value v_i and bid $b(v_i)$, their payoffs in a second price sealed bid auction is:

$$\Pi_i = \begin{cases} v_i - \max_{j \neq i} b(v_j) & \text{when } b(v_i) > \max_{j \neq i} b(v_j) \\ 0 & \text{otherwise.} \end{cases} \tag{243}$$

We can immediately see the weakly dominant strategy in this auction.

Theorem 583 (Weakly Dominant Strategy in Second Price Sealed Bid Auction)
In a second-price sealed-bid auction, it is a **weakly dominant strategy** for bidder i to bid their true value v_i:

$$b(v_i) = v_i. \tag{244}$$

Proof. To show this claim, we need to show that bidding the true valuation $b_i \equiv b(v_i) = v_i$ for bidder i does weakly better than any other bid bidder i could have made $b_i' \equiv b'(v_i)$ against any other opponents' bids $b_{-i} \equiv b(v_{-i})$. For notational purposes, denote $r_i \equiv \max_{j \neq i} b(v_j)$ to be the highest bid that is **not** bidder i's bid.

Suppose that we are in the case that i bids higher than their valuation $b_i' > v_i$. There are then 3 cases to consider based on what the highest bid that was not i's bid did, which we denoted by r_i.

1. If $r_i > b_i'$, bidder i loses and receives 0. Truthful bidding also loses and gives 0.
2. If $r_i < b_i'$ and $r_i \leq v_i$, bidder i wins and receives $v_i - r_i \geq 0$. Truthful bidding yields the same outcome and payoff.
3. If $v_i < r_i < b_i'$, bidder i wins but receives $v_i - r_i < 0$. Truthful bidding loses and yields 0, which is strictly better.

Thus, bidding $b_i = v_i$ weakly dominates every bid above v_i.

A symmetric argument applies to bids below v_i. If $b_i' < r_i < v_i$, the low bid loses and yields 0, whereas truthful bidding wins and yields $v_i - r_i > 0$; in all other cases, the two bids give the same payoff.

Therefore, $b_i = v_i$ is a **weakly dominant strategy**, and truthful bidding by every bidder is a dominant-strategy equilibrium. ☐

Proposition 584 (Probability of Winning a Second-Price Sealed-Bid Auction)
Suppose n bidders participate in a second-price sealed-bid auction. Conditional on bidder i's valuation v_i, their probability of winning in the truthful equilibrium is:

$$P_i(v_i) = \left[F(v_i)\right]^{n-1}. \tag{245}$$

Proof. It is weakly dominant for each bidder to bid their true value:

$$b(v_i) = v_i.$$

Bidder i therefore wins exactly when every other bidder has value at most v_i. Independence gives

$$\prod_{j \neq i} \mathbb{P}[v_j \leq v_i] = F(v_i)^{n-1}.$$

☐

We now can state the expected payment of a bidder.

Theorem 585 (Expected Payment by a Bidder in a Second-Price Sealed-Bid Auction)
Suppose n bidders participate in a second-price sealed-bid auction. Let $v_{[1]}^{-i}$ denote the highest valuation among bidders other than i. Conditional on bidder i's value v_i, their expected payment is:

$$\mathbb{E}[m_2(v_i)] = \big[F(v_i)\big]^{n-1}\mathbb{E}\big[v_{[1]}^{-i} \mid v_{[1]}^{-i} \leq v_i\big]. \tag{246}$$

Proof. The expected payment is the probability of winning times the expected second-highest bid conditional on winning:

$$\mathbb{E}[m_2(v_i)] = \mathbb{P}[\text{i wins} \mid v_i] \times \mathbb{E}\big[\text{second-highest bid} \mid \text{i wins}, v_i\big].$$

We have already shown the probability that bidder i wins conditional on v_i:

$$\mathbb{P}[\text{i wins} \mid v_i] = P_i(v_i) = \big[F(v_i)\big]^{n-1}.$$

Because truthful bidding is weakly dominant, the event that i has the highest bid is equivalent, up to zero-probability ties, to i having the highest value:

$$\mathbb{E}\big[\text{second-highest bid} \mid \text{i has highest bid}\big] = \mathbb{E}\big[\text{second-highest bid} \mid \text{i has highest value}\big].$$

Under truthful bidding, conditional on i winning, the second-highest bid is $v_{[1]}^{-i}$, the highest value among the remaining bidders. The winning event is $v_{[1]}^{-i} \leq v_i$. Hence,

$$\mathbb{E}\big[\text{second-highest bid} \mid \text{i wins}, v_i\big] = \mathbb{E}\big[v_{[1]}^{-i} \mid v_{[1]}^{-i} \leq v_i\big].$$

Now, we plug in both our expression to get the expected payment by bidder i:

$$\mathbb{E}[m_2(v_i)] = \big[F(v_i)\big]^{n-1}\mathbb{E}\big[v_{[1]}^{-i} \mid v_{[1]}^{-i} \leq v_i\big].$$

□

Theorem 586 (Expected Payments Are Equivalent)
Conditional on valuation v_i, bidder i makes the same expected payment in the symmetric equilibria of the first-price and second-price sealed-bid auctions.

Proof. In the first-price auction, Eqs. (233) and (235) give

$$\mathbb{E}[m_1(v_i)] = [F(v_i)]^{n-1}b(v_i) = [F(v_i)]^{n-1}\mathbb{E}[v_{[1]}^{-i} \mid v_{[1]}^{-i} \leq v_i].$$

In the second-price auction, truthful bidding implies that the winner pays $v_{[1]}^{-i}$, so Eq. (246) gives exactly the same expression:

$$\mathbb{E}[m_2(v_i)] = [F(v_i)]^{n-1}\mathbb{E}[v_{[1]}^{-i} \mid v_{[1]}^{-i} \leq v_i].$$

□

As the expected payments in both auctions are the same, then the expected revenues are the same.

Theorem 587 (Expected Revenue of Seller in Second Price Sealed Bid Auction under Uniform Values)
Suppose we have n bidders with i.i.d valuations $v_i \sim U[\underline{V}, \overline{V}]$. Then the expected revenue $\mathbb{E}[R]$ of the seller in a

second price sealed bid auction is:

$$\mathbb{E}[R] = \underline{V} + \left(\frac{n-1}{n+1}\right)(\overline{V} - \underline{V}). \tag{247}$$

The expected revenue of the second price sealed bid auction is the same as the expected revenue of the first price sealed bid auction.

Proof. Interim expected payments coincide for every value by the preceding theorem. Integrating over values and summing over the n symmetric bidders therefore gives the same expected revenue as in the first-price auction. ☐

12.5 Revenue Equivalence Theorem

Suppose we have n bidders and one good on sale. Suppose we have a mechanism that will allocate the good. We now want to derive a very powerful result that the expected payment and expected revenue will be the same in a certain class of auctions. We state the first important condition we need.

Definition 588 (Incentive Compatibility Constraint)
Let v be bidder i's type, let $P_i(v)$ be their probability of winning, and let $S_i(v)$ be their equilibrium expected surplus. The incentive-compatibility constraint is:

$$S_i(v) \geq S_i(\tilde{v}) + (v - \tilde{v})P_i(\tilde{v}). \tag{248}$$

Remark 589. *This incentive-compatibility condition says that reporting the true type v gives bidder i at least as much expected surplus as mimicking type $\tilde{v}$. A true type v who mimics $\tilde{v}$ obtains the equilibrium surplus of type $\tilde{v}$, adjusted by $(v - \tilde{v})P_i(\tilde{v})$ because their value differs whenever the mimicking bid wins.*

We now can derive the utility equivalence theorem.

Theorem 590 (Utility Equivalence Theorem)
The expected surplus of bidder i is given by:

$$S_i(v) = S_i(\underline{V}) + \int_{x=\underline{V}}^{v} P_i(x)dx \tag{249}$$

where $P_i(x)$ is the probability that bidder i wins the object when pretending to be type x.

Proof. Let $\tilde{v} = v + dv$.

$$S_i(v) \geq S_i(v + dv) + (v - v - dv)P_i(v + dv) = S_i(v + dv) - dvP_i(v + dv)$$

We also have that:

$$S_i(v + dv) \geq S_i(v) + dvP_i(v).$$

Combine these equations to get:

$$P_i(v + dv) \geq \frac{S_i(v + dv) - S_i(v)}{dv} \geq P_i(v).$$

Take limits as $dv \to 0$ and use the definition of a derivative:

$$\frac{dS_i}{dv} = P_i(v).$$

We now want to solve this differential equation:

$$\frac{dS_i(v)}{dv} = P_i(v).$$

First, we integrate both sides to get:

$$S_i(v) = \int_{\underline{V}}^{v} P_i(x)dx + C.$$

Now, let $v = \underline{V}$ and

$$S_i(\underline{V}) = \int_{\underline{V}}^{\underline{V}} P_i(x)dx + C.$$

$$\Rightarrow C = S_i(\underline{V}).$$

We plug this into our earlier expression to conclude:

$$S_i(v) = \int_{\underline{V}}^{v} P_i(x)dx + C$$
$$\Rightarrow S_i(v) = \int_{\underline{V}}^{v} P_i(x)dx + S_i(\underline{V}).$$

□

Theorem 591 (Revenue Equivalence Theorem (IPV Case))
Assume each of n risk-neutral potential buyers has an independent private value drawn from a common distribution $F(v)$ that is strictly increasing and atomless on $[\underline{V}, \overline{V}]$. Consider auction mechanisms satisfying:

1. The object always goes to the buyer with the highest value:
$$P_i(v) = \left[F(v)\right]^{n-1}. \tag{250}$$
2. A bidder with the lowest value $\underline{V}$ expects zero surplus:
$$S_i(\underline{V}) = 0. \tag{251}$$

Then every auction mechanism in this class has the **same expected revenue**, and a buyer of value v makes expected payment

$$\mathbb{E}[m(v)] = (F(v))^{n-1}v - \int_{x=\underline{V}}^{v} (F(x))^{n-1}dx. \tag{252}$$

Proof. Expected surplus is:

$$S_i(v) = v \cdot P_i(v) - \mathbb{E}[m(v)]$$

where $\mathbb{E}[m(v)]$ is expected payment, $vP_i(v)$ is expected gross value, and $S_i(v)$ is expected surplus. Rearranging for expected payment gives:

$$\mathbb{E}[m(v)] = v \cdot P_i(v) - S_i(v).$$

Plug in our expression from the utility equivalence theorem:

$$\begin{aligned}\mathbb{E}[m(v)] &= v \cdot P_i(v) - S_i(v) \\ &= v \cdot P_i(v) - \left[S_i(\underline{V}) + \int_{x=\underline{V}}^{v} P_i(x)dx\right] \\ &= v \cdot P_i(v) - 0 - \int_{x=\underline{V}}^{v} P_i(x)dx\end{aligned}$$

where we used the boundary condition $S_i(\underline{V}) = 0$.

Expected payment is:

$$P_i(v)v - \int_{x=\underline{V}}^{v} P_i(x)dx$$

where $P_i(x)$ is the probability of winning if you have value x. Plug in

$$P_i(v) = \left[F(v)\right]^{n-1}$$

and we are done.

$\square$

We get a useful result from the revenue equivalence theorem.

Theorem 592 (First-Price Expected Payment Expression)
In any auction in which only the winner pays and pays their own bid, the expected payment of bidder i with value v is their equilibrium bid times their probability of winning:

$$\mathbb{E}[m(v)] = b(v) \times P_i(v). \tag{253}$$