2022

Microeconomic Theory

Giri Parameswaran
Haverford College, gparames@haverford.edu

Microeconomic Theory

Giri Parameswaran
Haverford College

June 20, 2022

Contents

Part I

Analytical Methods of Economics

Chapter 1

Mathematical Tools

1.1 Review of Single Variable Calculus

A fundamental concern in economics is about how certain quantities change in response to changes in other quantities. For example, suppose the government is considering levying a tax on European cars, which it knows will raise their price. It might want to know how this will affect the quantity of European cars purchased.

Let $y = f(x)$ be a function of a single variable $x \in \mathbb{R}$. The **derivative** of f at x, which we denote $\frac{dy}{dx}$, $\frac{df(x)}{dx}$, or $f'(x)$, measures the rate of change of the output y as the input x changes a small amount. We say that $f'(x)$ is the slope or 'gradient' of the function at the point x.

The gradient of a function can be approximated by:

$$slope \approx \frac{\Delta f(x)}{\Delta x} = \frac{f(x+h) - f(x)}{(x+h) - x}$$

where the approximation becomes more accurate as h is made increasingly small. Thus, we formally define the derivative of f and x by:

$$f'(x) = \frac{df(x)}{dx} = \lim_{h \to 0} \frac{f(x+h) - f(x)}{h}$$

Example 1. Let $f(x) = x^2 + 3x + 5$. Then:

$$\begin{aligned} f'(x) &= \lim_{h \to 0} \frac{[(x+h)^2 + 3(x+h) + 5] - (x^2 + 3x + 5)}{h} \\ &= \lim_{h \to 0} \frac{(2x+3)h + h^2}{h} \\ &= \lim_{h \to 0} (2x + 3 + h) \\ &= 2x + 3 \end{aligned}$$

Of course, it would be cumbersome to have to evaluate these limits whenever we seek to differentiate a function. Thankfully, we have some handy rules for differentiating common functions and combinations of functions.

Derivatives of Common Functions

- Polynomial functions: If $f(x) = ax^n$, then $f'(x) = anx^{n-1}$.
- Exponential functions: If $f(x) = ae^x$, then $f'(x) = ae^x$.
- Logarithmic functions: If $f(x) = a\ln(x)$, then $f'(x) = \frac{a}{x}$.

Derivatives of Composite Functions Let $f(x)$ and $g(x)$ be two functions.

1. Sums Rule: If $h(x) = f(x) + g(x)$, then $h'(x) = f'(x) + g'(x)$.
2. Product Rule: If $h(x) = f(x)g(x)$, then $h'(x) = g(x)f'(x) + f(x)g'(x)$.
3. Quotient Rule: If $h(x) = \frac{f(x)}{g(x)}$, then $h'(x) = \frac{g(x)f'(x) - f(x)g'(x)}{[g(x)]^2}$.
4. Chain Rule: If $h(x) = f(g(x))$, then $h'(x) = f'(g(x)) \cdot g'(x)$

Exercise 1. Do the following:

1. Find the derivatives of each of the following functions:
   - $f(x) = (5 - 2x)^3 + \ln(x)$
   - $g(x) = \frac{-6\ln x}{e^{5-x^2} - 3x}$
   - $h(x) = e^{\ln(x)}$. (Simplify fully!)
2. Use the product rule to derive the quotient rule.
3. Let $f(x) = a^x$ for any constant $a > 0$. Show that $f'(x) = a^x \ln a$.

Example 2 (Percentage Changes)**.** Let $f(t)$ be a function of time that is changing continuously. Let $\Delta\% f(t)$ denote the rate of change in f in percentage terms. Hence $\Delta\% f(t) = \frac{f'(t)}{f(t)}$. By the chain rule, we have:

$$\Delta\% f(t) = \frac{d}{dt} \ln f(t)$$

[Convince yourself that this is true.] Show the following:

- If $h(t) = f(t) \cdot g(t)$, then $\Delta\% h(t) = \Delta\% f(t) + \Delta\% g(t)$.

- If $h(t) = \frac{f(t)}{g(t)}$, then $\Delta\% h(t) = \Delta\% f(t) - \Delta\% g(t)$.
- If $h(t) = [f(t)]^n$, then $\Delta\% h(t) = n\Delta\% f(t)$.
- If $Y(t)$ and $N(t)$ denote GDP and population size at time t, and $y(t) = \frac{Y(t)}{N(t)}$ denotes per-capita GDP, show that per-capita economic growth is economic growth minus population growth.

1.2 Optimization

Suppose we wish to find the optimizer x^* that maximizes or minimizes $f(x)$ on its domain. We can often use calculus to find optimizers. Loosely speaking, to achieve a maximum, a function must increase (have a positive slope) and then turn around and decrease (have a negative slope). Hence, at the maximal point, the function will be flat —its gradient will be zero. Hence, to find maxima, we should search for values of the function where the derivative is zero.

Definition 1. A *critical point* of a function $f(x)$ is a point x at which $f'(x) = 0$ (i.e. the function is flat).

We say the function satisfies the *first order conditions* at its critical points. The critical points are candidate optimizers —although we must be careful; not every critical point is an optimizer, and not every optimizer is at a critical point.

Example 3. Let $f(x) = 20 - 24x + 9x^2 - x^3$ be a function on the domain $\{x \mid x \geq 0\}$. The critical points are given by:

$$\begin{aligned} f'(x) = -24 + 18x - 3x^2 &= 0 \\ x^2 - 6x + 8 &= 0 \\ (x-2)(x-4) &= 0 \\ x &= 2, 4 \end{aligned}$$

Hence, there are two critical points. Evaluating the function at the critical points gives: $f(2) = 0$ and $f(4) = 4$. But note that $f(0) = 20 > f(4)$ and $f(6) = -16 < f(2)$. Hence, the critical points select neither the maximum nor the minimum of the function. Indeed, the maximum is achieved at the corner solution $x^* = 0$. The function has no minimum. However, $x = 2$ is a *local minimum* and $x = 4$ is a local maximum. We can easily see this by plotting the function.

Lemma 1 (Second Order Condition)**.** *Suppose x^* is a critical point of the function f. Then:*

- *x^* is local minimum if $f''(x^*) > 0$.*
- *x^* is a local maximum if $f''(x^*) < 0$.*

Example 4. A firm faces a demand curve: $P(Q) = 100 - Q$. Hence, if it produces quantity Q of output, it must sell them at price $P(Q)$. The firm has production costs given by $C(Q) = \frac{1}{3}Q^3 - 10Q^2 + 156Q + 16/3$.

1. What quantity maximizes its revenue?
2. What quantity maximizes its profit?

To answer (1), note that revenue is $R(Q) = P(Q) \times Q = 100Q - Q^2$. The first order condition implies $R'(Q) = 100 - 2Q = 0$, and so we have $Q = 50$. To answer (2), note that profit is

$$\Pi(Q) = R(Q) - C(Q) = -\frac{1}{3}Q^3 + 9Q^2 - 56Q - 16/3$$

The critical values are the solutions to $\Pi'(Q) = -Q^2 + 18Q - 56 = 0$, which implies $Q = 4$ or $Q = 14$. Now, $\Pi(4) = -\$106.67$, whilst $\Pi(14) = 60$. Clearly $Q = 4$ cannot be profit maximizing. In fact, $Q = 4$ is a profit minimizer.

In fact, we can verify this by checking the second order conditions. $\Pi''(Q) = 18 - 2Q$. Evaluating the second derivative at the critical points gives: $\Pi''(4) = 10 > 0$ and $\Pi''(14) = -10 < 0$. Thus, $Q = 4$ is a local minimum and $Q = 14$ is a local maximum. □

Lemma 2. *Consider a function $f(x)$.*

- *Suppose $f''(x) < 0$ for all x. If x^* is a critical point of f, then x^* is the global maximum.*
- *Suppose $f''(x) > 0$ for all x. If x^* is a critical point of f, then x^* is the global minimum.*

1.3 Multivariate Calculus

In the previous sections, we considered functions of a single variable $f(x)$. We now extend the analysis to multivariate functions —those that take multiple inputs. For example, consider the bi-variate function: $f(x_1, x_2) = x_1^2 - x_1 \ln x_2$.

In single variable calculus, we constructed the derivative by taking small changes in the input (x) and asking how quickly the output $f(x)$ changed in response. Derivatives in the multivariate setting are defined analogously: we take a small change in one of the inputs (either x_1 or x_2) and ask how quickly the output changes, assuming the other input is held fixed at its original level. We call this a **partial derivative**, which we denote $\frac{\partial f}{\partial x_i}$ or $f_i(x_1, x_2)$ where x_i is the variable that is allowed to change.

Naturally, rather than having a single derivative, we now have two partial derivatives —one associated with a change in x_1 and the other associated with a change in x_2. Formally, we have:

$$\frac{\partial f}{\partial x_1} = \lim_{h \to 0} \frac{f(x_1 + h, x_2) - f(x_1, x_2)}{h}$$
$$\frac{\partial f}{\partial x_2} = \lim_{h \to 0} \frac{f(x_1, x_2 + h) - f(x_1, x_2)}{h}$$

The first is the partial derivative of f with respect to x_1, and the second is the partial derivative of f with respect to x_2.

We can compute partial derivatives in much the same way we computed derivatives in the single variable case, using the rules we noted above. When doing so, we treat all variables other than the one that we are taking the derivative with respect to as if they are constants.

Example 5. Consider the bivariate function $f(x_1, x_2) = x_1^2 + x_1 \ln x_2$. The partial derivatives are:

$$f_1(x_1, x_2) = \frac{\partial f}{\partial x_1} = 2x_1 + \ln x_2$$
$$f_2(x_1, x_2) = \frac{\partial f}{\partial x_2} = \frac{x_1}{x_2}$$

When computing the partial derivative w.r.t. x_1, we treat x_2 as if it were constant. So the function is effectively of the form: $x_1^2 + ax_1$, where a stands in for $\ln x_2$ (which we treat as a constant). Clearly the derivative will be of the form $2x_1 + a$. Similarly, when computing the partial derivative w.r.t. x_2, we treat x_1 as if it were constant. So the function is effectively of the form: $b + c \ln x_2$, where b and c stand in for x_1^2 and x_1 (which we treat as constants). Clearly the derivative will be of the form $\frac{c}{x_2}$.

We describe the derivatives as *partial derivatives* because they only show the (partial) effect of a change in the output —that coming from a change in a single input alone. In economics, often-times, several inputs will change simultaneously. The *total derivative* accounts for all of these changes, whereas the partial derivative takes each in isolation.

Just as we can find higher-order (second, third, ...) derivatives of single-variable functions, so can we of multi-variate functions. In single variable calculus, the second derivative is the derivative of the derivative. When the function is bi-variate, we know that that there are two derivatives to consider —and each of these can be differentiated with respect to each of two variables (i.e. we can take the (partial) derivative of $\frac{\partial f}{\partial x_1}$ w.r.t. x_1 or w.r.t. x_2, and we can take the (partial) derivative of $\frac{\partial f}{\partial x_2}$ w.r.t. x_1 or w.r.t. x_2). Thus, there are four second order partial derivatives:

$$
\begin{aligned}
f_{11}(x_1, x_2) &= \frac{\partial^2 f}{\partial x_1^2} = \frac{\partial}{\partial x_1}\left(\frac{\partial f}{\partial x_1}\right) \\
f_{22}(x_1, x_2) &= \frac{\partial^2 f}{\partial x_2^2} = \frac{\partial}{\partial x_2}\left(\frac{\partial f}{\partial x_2}\right) \\
f_{12}(x_1, x_2) &= \frac{\partial^2 f}{\partial x_2 \partial x_1} = \frac{\partial}{\partial x_2}\left(\frac{\partial f}{\partial x_1}\right) \\
f_{21}(x_1, x_2) &= \frac{\partial^2 f}{\partial x_1 \partial x_2} = \frac{\partial}{\partial x_1}\left(\frac{\partial f}{\partial x_2}\right)
\end{aligned}
$$

We refer to the first two as ‘own’ second-order partial derivatives, and the second two as ‘cross’ second-order partial derivatives. The reason for this nomenclature should be obvious.

Example 6. Continuing the previous example, we have: $f_{11} = 2$, $f_{12} = \frac{1}{x_2} = f_{21}$, and $f_{22} = -\frac{x_1}{x_2^2}$.

Exercise 2. For each of the following functions, find all of the first-order and second-order partial derivatives:

1. $f(x_1, x_2) = \frac{2x_1}{\ln x_2} + x_1 x_2 + 5$
2. $g(x, y, z) = x^2 + 2xy - y^2 + 2\frac{y}{z} + z$

1.4 Unconstrained Optimization

We can find local optima of multivariate functions in much the same way we do with univariate functions. The intuition is the same. For a single variable function, if the derivative is non-zero, then there is some small change that will cause the function to increase. Thus, to be at the maximum, it must be that the derivative is zero —there are no further opportunities to increase the function. In the multivariate case, the same logic applies, except we now need *all* of the partial derivatives to be zero. This guarantees that the function cannot be improved by making a small change to *any* of the inputs.

Definition 2. A *critical point* of a function $f(x, y)$ is a point (x, y) at which both $\frac{\partial f}{\partial x} = 0$ and $\frac{\partial f}{\partial y} = 0$.

We say the function satisfies the first order conditions at its critical points. As before, the critical points are candidate optimizers, but not every critical point is an optimizer, and not every optimizer s a critical point.

Lemma 3. *Let $f(x, y)$ be a bivariate function. Suppose (x^*, y^*) is a critical point of f. Then:*

- (x^*, y^*) *is a local minimum if: (i)* $f_{xx} > 0$ *and* $f_{yy} > 0$*, and (ii)* $f_{xx}f_{yy} - f_{xy}f_{yx} > 0$*, where the second-order partial derivatives are all evaluated at* (x^*, y^*)*.*

- (x^*, y^*) *is a local maximum if: (i)* $f_{xx} < 0$ *and* $f_{yy} < 0$*, and (ii)* $f_{xx}f_{yy} - f_{xy}f_{yx} > 0$*, where the second-order partial derivatives are all evaluated at* (x^*, y^*)*.*

Example 7. A firm can produce a quantity of output q by combining k units of capital and l units of labour. The firm's production function is $q = k^{0.25}l^{0.5}$. The firm can hire as much capital and labour as it desires. The per unit cost of capital is 10 and the per worker wage is 5. Firms can sell any quantity of output produced on the market at price 20. How much of each input should it hire? What quantity of output should it produce? What is the firm's optimal profit?

The firm's profit is $\pi = 20q - 10k - 5l$. Since the firm's output is determined by its production function, we have $\pi = 20k^{0.25}l^{0.5} - 10k - 5l$. Thus, the firm's profit is a function of two variables: k and l. The firm's problem is:

$$\max_{k,l} 20k^{0.25}l^{0.5} - 10k - 5l$$

The first order conditions are:

$$\begin{aligned}
\frac{\partial \pi}{\partial k} &= 5k^{-0.75}l^{0.5} - 10 = 0 \\
\frac{\partial \pi}{\partial l} &= 10k^{0.25}l^{-0.5} - 5 = 0
\end{aligned}$$

Combining these gives: $\frac{1}{2}\frac{l}{k} = 2$, and so $l = 4k$. Substituting this into the first condition gives:

$$\begin{aligned}
5k^{-}0.75(4k)^{0.5} &= 10 \\
2k^{-0.25} &= 2 \\
k^* &= 1
\end{aligned}$$

Then, since $l^* = 4k^*$, we have $l^* = 4$.

Next, we must check that the second-order conditions hold so that we can confirm we have indeed found a maximum. Below are each of the second derivatives evaluated at the critical point:

$$\begin{aligned}
\pi_{kk} &= -3.75k^{-1.75}l^{0.5} = -3.75(1)^{-1.75}(4)^{0.5} = -7.5 \\
\pi_{ll} &= -5k^{0.25}l^{-1.5} = -5(1)^{0.25}(4)^{-1.5} = -0.625 \\
\pi_{kl} &= 2.5k^{-0.75}l^{-0.5} = 2.5(1)^{-0.75}(4)^{-0.5} = 1.25 \\
\pi_{lk} &= 2.5k^{-0.75}l^{-0.5} = 2.5(1)^{-0.75}(4)^{-0.5} = 1.25
\end{aligned}$$

Hence, we verify that $\pi_{kk} < 0$, $\pi_{ll} < 0$ and $\pi_{kk}\pi_{ll} - \pi_{kl}\pi_{lk} = -7.5 \times -0.625 - 1.25 \times 1.25 = 3.125 > 0$. The second order conditions are satisfied. □

1.5 Constrained Optimization

Fundamentally, economics is the study of choice. Agents only have genuine choices to make when they face constraints that prevent them from getting everything they want. (Such constraints might include budget constraints, time constraints, technological constraints, institutional constraints, etc.)

Example 8. Two canonical examples of constrained optimization problems:

1. *Utility Maximization subject to a budget constraint.* There is an agent who must choose quantities of goods x and y to purchase to maximize her utility subject to those goods being affordable. We have:

$$\max_{x,y} u(x,y) \,\text{s.t.}\, p_x x + p_y y = I$$

2. *Profit Maximization subject to a technological constraint.* There is a firm that must choose quantities of inputs capital k and labor l to hire, and the quantity of output to produce q, subject to the output being technologically feasible given the inputs. We have:

$$\max_{q,k,l} pq - rk - wl \,\text{s.t.}\, q = f(k,l)$$

Consider a generic problem:

$$\max_{x,y} f(x,y) \,\text{s.t.}\, g(x,y) = 0$$

$f(x,y)$ is the objective function, whilst $g(x,y)$ is the constraint function. (It should be clear that we can express each of the constraints in the above examples in the form $g(\cdot) = 0$.)

Absent the constraint, we know what to do —find the critical points of f by setting all the first-order partial derivatives to zero. But these critical points may not satisfy the constraints (i.e. they may not be feasible). Our solution is to incorporate the constraint into the objective function so that it is already taken into account when we take first order conditions.

Consider a new function:

$$\mathcal{L}(x,y,\lambda) = f(x,y) - \lambda g(x,y)$$

where λ is a new variable called the **Lagrange Multiplier**. The new function $\mathcal{L}$ is called the **Lagrangian**. It turns out that to find the constrained maximizer of the original objective function, it suffices to find the *unconstrained* maximizer of the Lagrangian. To see why, notice that the Lagrangian has the constraint incorporated into it. In fact, it is simply the objective function, augmented with a penalty whenever the constraint is not satisfied. The size of the penalty is given by the Lagrange multiplier.

Notice that the penalty term will be zero whenever the constraint is satisfied, since $g(x,y) = 0$. This ensures that, when the constraint is satisfied, the Lagrangian function simply coincides with the true objective function, so that maximizing the Lagrangian and constrained maximizing the objective function amount to the same thing. If the constraint is not satisfied, then the correction term $(\lambda g(x, y))$ will be positive, which causes the Lagrangian to take a lower value than the objective. If we seek to maximize the Lagrangian then, we will be deterred from settling on values that violate the constraint.

But how do we choose the size of the penalty, λ? We let the calculus do it for us! Notice that $\frac{\partial \mathcal{L}}{\partial \lambda} = -g(x, y)$. Hence, if we take the first order condition with respect to λ, we effectively get the condition $g(x, y) = 0$, which guarantees that the constraint will be satisfied.

In summary, to do constrained optimization, we:

1. Form the Lagrangian function, which incorporates the constraint, in a way that penalizes the maximization whenever the constraint is violated.

2. Find the unconstrained optimum of the Lagrangian. To do so, we take first order conditions with respect to the original variables x and y, as well as the penalty variable λ (i.e. the Lagrange multiplier).

Example 9. Consider the following problem:

$$\max_{x,y} -x^2 - y^2 \text{ s.t. } 2x + y = 5$$

Clearly, the unconstrained maximum is achieved when $\hat{x} = 0 = \hat{y}$ —you can easily check this by taking first order conditions —but this obviously does not satisfy the constraint.

The Lagrangian is:

$$\mathcal{L} = -x^2 - y^2 - \lambda(2x + y - 5)$$

Taking first order conditions, we have:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial x} &= -2x - 2\lambda = 0 \\ \frac{\partial \mathcal{L}}{\partial y} &= -2y - \lambda = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= 5 - 2x - y = 0 \end{aligned}$$

We have a system of 3 equations in 3 unknowns. From the first two equations, after solving for λ, we get:

$$\lambda = -x = -2y$$

which implies that $x = 2y$. Substituting this into the third equation gives $5 - 2(2y) - y = 0$, which implies $y = 1$. Then, since $x = 2y$, it must be that $x = 2$. Finally, since $\lambda = -x = -2y$, we have $\lambda = -2$. The constrained optimum is $(x^*, y^*) = (2, 1)$.

Notice that the penalty $\lambda^* = -2$ is negative. Does this make sense? To get from the constrained optimum to the unconstrained optimum, both x and y would need to decrease, which would cause the constraint $(2x + y - 5)$ to become negative. If λ were positive, the penalty would be negative, and since the penalty is subtracted, this would cause the Lagrangian to be *larger* than the objective. Clearly this cannot be the case. This explains why $\lambda < 0$.

Example 10 (Firm's Choice cont.). Return to Example 8 about a firm choosing which inputs to hire. Recall, the firm faces a production technology $q = k^{0.25}l^{0.5}$, output price $p = 20$, and input prices $r = 10$ and $w = 5$. The firm must choose the quantity of output to produce q and the quantities of each input k and l.

We can reformulate the firm's problem as a constrained optimization problem:

$$\max_{q,k,l} 20q - 10k - 5l \text{ s.t. } k^{0.25}l^{0.5} = q$$

The Lagrangian is:

$$\mathcal{L} = 20q - 10k - 5l - \lambda\left(q - k^{0.25}l^{0.5}\right)$$

The first order conditions are:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial q} &= 20 - \lambda = 0 \\
\frac{\partial \mathcal{L}}{\partial k} &= -10 + \lambda\left(0.25k^{-0.75}l^{0.5}\right) = 0 \\
\frac{\partial \mathcal{L}}{\partial l} &= -5 + \lambda\left(0.5k^{0.25}l^{-0.5}\right) = 0 \\
\frac{\partial \mathcal{L}}{\partial \lambda} &= k^{0.25}l^{0.5} - q = 0
\end{aligned}$$

Using the first condition, we have $\lambda = 20$. Substituting this into the second and third equations gives:

$$\begin{aligned}
\left(0.25k^{-0.75}l^{0.5}\right) &= \frac{1}{2} \\
\left(0.5k^{0.25}l^{-0.5}\right) &= \frac{1}{4}
\end{aligned}$$

This is a system of two equations in two variables identical to what we had in the previous example. Solving, we get $k = 1$ and $l = 4$. Finally, substituting these into the fourth equation gives $q = (1)^{0.25}(4)^{0.5} = 2$.

Hence, the solution is $q^* = 2$, $k^* = 1$, and $l^* = 4$.

Chapter 2

Consumer Theory

Lionel Robbins' famous definition of economics states: "Economics is the science which studies human behaviour as a relationship between ends and scarce means which have alternative uses." At its most basic level, economics is the study of how individuals and societies make choices. In this chapter, we explore the canonical theory of rational choice.

Let X denote the set of possible alternatives. In much of this chapter, X will represent the set of consumption bundles that a consumer may purchase. But the setup is more general. X may also denote:

1. candidates in the 2020 Democratic primary. $X = \{Biden, Buttigieg, Harris, Warren, ...\}$.

2. quadruples of courses for which students can register on Bionic.

3. which of my children I visit this weekend.

4. to which charities to donate, and how much.

In different contexts, only a subset of choices may be available to the agent. (For example, amongst the set of feasible course-quadruples, a student may not choose those for which a pre-requisite is missing, or which includes courses already taken.) Our goal is to develop a theory that explains the agent's choice from different feasible choice sets.

We will do this in two parts. First, we consider **preferences** —which specify an agent's ranking of all possible outcomes (regardless of whether they are 'feasible' or not). Next, we explore how various constraints limit the set of feasible choices. We then combine the two to find the agent's optimal choice.

2.1 Preferences

2.1.1 Baseline Assumptions

Preferences describe an agent's attitude toward the various available choices. Differences in the way agents perceive their options are captured by differences in preferences. In economics, preferences are taken as primitive —two agents are different from one another insofar as they have distinct preferences.

We represent an agent's preference between any two elements $x, y \in X$ using a preference relation $\succeq$. We write $x \succeq y$ to mean that 'x is *at least as preferred* as y', or 'x is weakly preferred to y.'

The weak preference relation is sufficient to capture strict preferences ('x is more preferred than y') and indifference ('x and y are equally preferred'). We denote the strict preference and indifference by $\succ$ and $\sim$, respectively.

The following are true:

- $x \succ y$ iff (if and only if) $x \succeq y$ and $y \not\succeq x$.
- $x \sim y$ iff $x \succeq y$ and $y \succeq x$.

Definition 3 (Rationality)**.** The preference relation $\succeq$ is **rational** if it satisfies:

- **completeness**: for all $x, y \in X$, either $x \succeq y$ or $y \succeq x$ (or both).
- **transitivity**: for all $x, y, z \in X$, if $x \succeq y$ and $y \succeq z$, then $x \succeq z$.

Completeness requires that, when confronted between two alternatives, the agent is always able to assert a preference between them (even if to say that she is indifferent). *Transitivity* requires that the agent's ranking of outcomes be consistent —she cannot prefer x to y, and y to z, and z to x. We say an agent is *rational* if she can always rank outcomes and do so consistently.

Note of caution —our use of the term 'rational' is different from the common usage. Rationality in economics simply means consistency; it does not convey in any sense that the agent makes the 'right' choice. For example, suppose I assert that it is preferable to smoke habitually than to smoke socially; and that it is preferable to smoke socially than to not smoke at all. Then, as long as I also assert that it is preferable to smoke habitually to not smoking at all, I am rational. My ranking of outcomes might not be advisable, but that does not make it 'irrational.'

Example 11. Suppose an individual encounters 2 slices of cake in the kitchen: one mud cake and one cheese cake. They must decide whether to eat both slices (B), one slice (and if so, which (M or C)) or neither ($\emptyset$). Let $X = \{B, M, C, \emptyset\}$ denote the set of possible alternatives for me. Let $\succeq_1$, $\succeq_2$ and $\succeq_3$ be different preference relations, with:

$$B \succ_1 C \sim_1 M \succ_1 \emptyset$$
$$C \succ_2 B \succ_2 \emptyset \succ_2 M$$
$$\emptyset \succ_3 C \sim_3 M \succ_3 B$$

The first preference can be described in words as follows: 'The agent prefers more cake to less. They are indifferent between flavors.' By contrast, the second preference can be described: 'The agent prefers cheese cake and dislikes mud cake. She would rather have both than neither.' We might think of the third preference as capturing a person dedicated to a healthy diet. □

A rational preference relation can be expressed as a list —as we have shown in the example. But listing outcomes can be cumbersome, especially as the list becomes long. It often helps to represent the information in the list using numbers.

Definition 4. A utility function $u : X \to \mathbb{R}$ represents the preference relation $\succeq$ if for all $x, y \in X$, $x \succeq y$ iff $u(x) \geq u(y)$.

A utility function is an assignment of numbers to elements in the choice set X, such that higher numbers are assigned to outcomes that are more preferred, and the same number is assigned to outcomes that are indifferent.

Proposition 1. *A preference relation $\succeq$ can be represented by a utility function u iff it is complete and transitive.*[1]

Exercise 3. We can represent preference relation $\succeq_1$ from Example 11 using the utility function: u_1, where: $u_1(B) = 100$, $u_1(C) = 1$, $u_1(M) = 0$ and $u_1(\emptyset) = -20$. Can you:

- Find functions u_2 and u_3 that represent $\succeq_2$ and $\succeq_3$?
- Find different functions v_1, v_2 and v_3, that also represent $\succeq_1$, $\succeq_2$ and $\succeq_3$, respectively?

As the above exercise makes clear, there is nothing special about the numbers we assign, as long as they respect the appropriate ordering. The binary relation $\succeq$ is *ordinal.* It merely ranks outcomes. It says nothing about the strength of an agent's preference. Similarly, a utility function simply communicates this ranking. The distance between the utility numbers is inconsequential. Suppose $x \succ y \succ z$. We can represent this by: $u(x) = 30$, $u(y) = 20$

[1]For students who have a background in Analysis: Technically, if X is uncountable, we need $\succeq$ to additionally satisfy the assumption of Continuity. Continuity requires that for a sequence of outcomes $x_n \to x$ and $y_n \to y$, if $x_n \succeq y_n$ for each n, then $x \succeq y$.

and $u(z) = 10$. This does NOT mean that y is twice as preferred as z and x is three times as preferred. Simply that y is preferred to z (and so is x).

There is no unique scale to represent utility. It cannot be measured with 'hedonometer'. The same preference relation can be represented by many different utility functions.

Lemma 4. *Suppose u represents $\succeq$ and let $g(\cdot)$ be a strictly increasing function. Then $v(x) = g(u(x))$ also represents $\succeq$.*

Dispelling some myths about utility: The notion of utility in no way commits us to the idea that preferences are over money, wealth, commodities etc. The agent's preferences can be over anything. They may care about things that are not commodifiable. For example, the agent may have preferences over how much time they spend with a loved one (and what they do during that time). Similarly, their preferences may relate to outcomes that affect other agents. For example, the agent might care about the outcomes of their child or spouse or friends. (We say such an agent has 'other-regarding' preferences.) Our theory makes no commitments to ideas of ego-centrism or individualism. It is perfectly consistent with the social aspects of human feelings —love, envy, guilt etc. Moreover, we make no commitments about how the agent ranks outcomes — one agent may prefer more chocolate to less, another may have an intense dislike of chocolate. Individuals may prefer less material goods and more experiential goods, including the experience of quiet, peace, simplicity etc.

Example 12. Return to the cake example. Suppose the individual has a friend, whom they must decide whether to share the cake(s) with. An alternative now is a pair (P, Q) where P denotes what the individual receives, and Q denotes what the friend receives. The set of alternatives is: $X = \{(B, \emptyset), (M, C), (M, \emptyset), (C, M), (C, \emptyset), (\emptyset, B), (\emptyset, M), (\emptyset, C), (, \emptyset)\}$. Clearly there are many possible ways to rank these outcomes. Here are a couple: (For the sake of exposition, suppose both agents have individual preferences given by $B \succ M \succ C \succ \emptyset$).

$$(B,\emptyset) \succ_4 (M,\emptyset) \succ_4 (M,C) \succ_4 (C,\emptyset) \succ_4 (C,M) \succ_4 (\emptyset,\emptyset) \succ_4 (\emptyset,C) \succ_4 (\emptyset,M) \succ_4 (\emptyset,B)$$
$$(B,\emptyset) \succ_5 (M,\emptyset) \succ_5 (C,\emptyset) \succ_5 (\emptyset,\emptyset) \succ_5 (M,C) \succ_5 (\emptyset,C) \succ_5 (C,M) \succ_5 (\emptyset,M) \succ_5 (\emptyset,B)$$
$$(\emptyset,B) \succ_6 (C,M) \succ_6 (\emptyset,M) \succ_6 (M,C) \succ_6 (\emptyset,C) \succ_6 (B,\emptyset) \succ_6 (M,\emptyset) \succ_6 (C,\emptyset) \succ_6 (\emptyset,\emptyset)$$

Agent 4 prefers to get what he individually wants, and conditional upon that, to deny her friend what she wants. Agent 5 prefers to deny his friend what she wants, and conditional upon that, to get what he individually wants. Agent 4 is self-centered and mean-spirited, whilst agent 5 is malicious or vindictive. Agent 6 prefers to first give his friend what she wants before tending to himself. Agent 6 is generous or sacrificing.

2.1.2 Additional Assumptions

Completeness and Transitivity alone are insufficient to build a predictive theory. Why? They are too permissive —any (consistent) ranking of outcomes is possible. But if so, how could we ever predict what choices an agent will make, if in principle, their ranking of outcomes could be anything? To say something more concrete, we need more structure. Two things are worth noting:

- There is no 'correct' way to add additional assumptions. In this course, we will focus on the canonical approach, although we will consider alternative ('behavioral') assumptions from time to time. Each approach has strengths and weaknesses.
- The more structure we add, the stronger predictions our theory can make, but also the more likely it will be that we've introduced false assumptions. Hence, we face a trade-off between making strong predictions and making correct predictions.

The canonical approach adds 2 additional assumptions about the nature of preferences. To reiterate: we do not think these assumptions always hold —but we do think they are reasonable approximations over a large range of contexts. Naturally they will be less reasonable in other contexts.

For concreteness, we will focus most of our attention on a particular context of interest —choosing over bundles of goods. (For example, choosing how much of various goods to buy each month.) We plot the quantity of good X on the horizontal axis and the quantity of good Y on the vertical axis. A consumption bundle is a pair (x, y) which indicates the amounts of good X and good Y being consumed. The **Consumption Space** is the set of all possible consumption bundles - in this case, the non-negative orthant of the Cartesian plane. We assume goods are perfectly divisible - so I can consume $\frac{3}{8}$ of a good, or π goods.

We write '$(x_1, y_1) \geq (x_2, y_2)$' to denote the relationship between bundles (x_1, y_1) and (x_2, y_2) such that $x_1 \geq x_2$ and $y_1 \geq y_2$ —i.e. bundle 1 contains at least as much of each good as bundle 2.

Definition 5. The preference relation $\succeq$ is:

- **monotone** if $(x_1, y_1) \geq (x_2, y_2)$ implies $(x_1, y_1) \succeq (x_2, y_2)$.
- **convex** if $(x_1, y_1) \succeq (x_2, y_2)$ implies $(\lambda x_1 + (1-\lambda)x_2, \lambda y_1 + (1-\lambda)y_2) \succeq (x_2, y_2)$, for all $\lambda \in [0, 1]$.

Monotonicity captures the idea that 'more is better'.[2] *Convexity* captures the idea that agents prefer moderate quantities of all goods to large quantities of some and small quantities of others. (More formally, the (weighted) average of two bundles is at least as good as the worst of the two bundles.)

Let $u(x)$ represent $\succeq$.

- If $\succeq$ is monotone, then u is increasing in each dimension, i.e., $\frac{\partial u}{\partial x} > 0$ and $\frac{\partial u}{\partial y} > 0$.

[2] In fact, monotonicity is stronger than we need, although it will be convenient for our purposes. It suffices to assume that preferences satisfy *local non-satiation*, which states that for every (x_1, y_1), and every $\epsilon > 0$ there exists (x_2, y_2) in an ϵ-neighbourhood of (x_1, y_1) s.t. $(x_2, y_2) \succ (x_1, y_1)$. Local non-satiation says that there is always something (nearby) that is better, whereas monotonicity says that something has 'more'.

- If $\succeq$ is convex, then u is quasi-concave. This implies that the second order conditions are satisfied at any critical point.

We can represent consumer preferences diagrammatically using **indifference curves**. An indifference curve is a line joining all consumption bundles over which the consumer is indifferent, i.e., if two bundles a and b are on the same indifference curve, then $u(a) = u(b)$.

Properties of Indifference Curves

1. Indifference curves are downward sloping (This follows from the monotonocity assumption. By monotonicity, every bundle in quadrant III must be preferred to a. Similarly, every bundle in quadrant I is inferior to a. Hence, bundles which are indifferent to a must lie in either quadrant II or IV - but this implies that the indifference curve is downward sloping.)

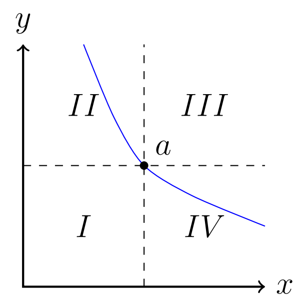

2. Indifference curves do not cross. (We show this by contradiction. Suppose there are two indifference curves which intersect. Let a be the bundle corresponding to the point where they intersect. Since they both lie on IC_1, then a and b must be indifferent. Similarly, since a and c both lie on IC_2, then a and c must be indifferent. Then by transitivity, b and c must be indifferent. But b clearly contains more of both goods than c. Hence by monotonicity, b must be preferred to c. This is a contradiction! Hence, there cannot be two indifference curves which intersect.)

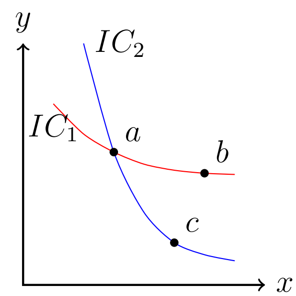

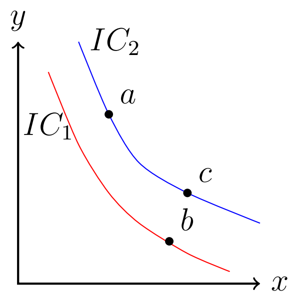

3. Higher indifference curves are preferred to lower ones. In the figure below, notice that a is on a higher indifference curve than b, but our assumptions do not allow us to directly rank them. However, there is another bundle c, lying on the same indifference curve as a (and thus for which $a \sim c$), which we can compare to b using monotonicity ($c \succ b$). Then, by transitivity, it must be that $a \succ b$. Using this logic, we can show that every bundle on a higher indifference curve must be preferred to every bundle on a lower indifference curve.

4. Indifference curves are convex. (This follows from the convexity assumption. Let a and b be indifferent and suppose they lie on IC_1. Draw a straight line connecting a and b. This line represents bundles which are an average of a and b. By the convexity assumption, we know that each of these bundles is strictly preferred to both a and b. Hence, they must all lie above IC_1. Put differently, IC_1 must lie below the line joining a and b. This implies that the indifference curves are convex.)

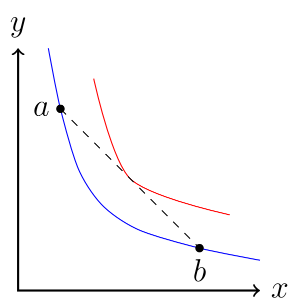

The slope of an indifference curve is known as the **marginal rate of substitution** (MRS). The MRS indicates how much of good Y the consumer is willing to forgo in order to get one extra unit of good X. The MRS is changing along the indifference curve. In fact, as one moves along the indifference curve from left to right, the MRS decreases. (To see why, suppose we are at a point towards the left end of the IC. Such a bundle contains a lot of good Y and relatively little X. By convexity, we know that the consumer prefers moderation to extremes. Hence, the consumer will be willing to give up a lot of Y in order to get a little extra X. Hence the MRS is high. Now, consider a point towards the right end of the IC. This bundle contains a lot of good X and very little Y. Again, since the consumer prefers moderation, they will be very unwilling to give up too much more Y to get an extra unit of

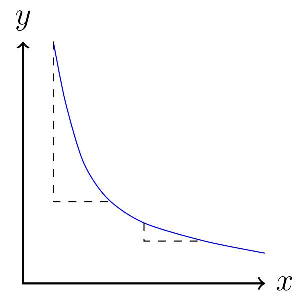

Figure 2.1: The marginal rate of substitution (slope of indifference curve) decreases as we move along the indifference curve.

X - since they already have a lot of X and very little Y. Hence, the MRS is small at this point.)

In fact, we can derive the marginal rate of substitution using calculus. Consider an arbitrary indifference curve, along which the utility level is u_0. Consider a bundle along this indifference curve. If we know the quantity of good x, then we can immediately determine the necessary quantity of good y to ensure the bundle lies on the indifference curve, i.e., the indifference curve defines a function $y(x)$ s.t.

$$u(x, y(x)) = u_0$$

The marginal rate of substitution is simply $\frac{dy}{dx}$. To find this, totally differentiate this expression w.r.t x., We have:

$$\begin{aligned}
\frac{\partial u}{\partial x} + \frac{\partial u}{\partial y}\frac{dy}{dx} &= 0 \\
\frac{dy}{dx} &= -\frac{\partial u/\partial x}{\partial u/\partial y} \\
MRS &= -\frac{MU_x}{MU_y}
\end{aligned}$$

Hence the marginal rate of substitution is the ratio of marginal utilities.

2.1.3 Special Cases

Two special types of indifference curves are worth noting.

- If goods are perfect substitutes (e.g. nickels and dimes), then the indifference curve is linear. (Clearly, I will always be willing to exchange two nickels for a dime, and so the MRS is constant all along the indifference curve.)

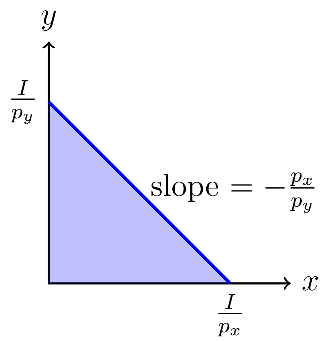

Figure 2.2: The Budget Set

- If goods are perfect complements (e.g. left shoes and right shoes), then the indifference curves are L-shaped. (It should be clear that getting more left shoes is only utility enhancing if I also get more right shoes. Hence, a consumer will be indifferent between the bundles $(2, 2)$ and $(2, 4)$ - since the additional 2 right shoes give them no extra utility. Similarly, a consumer will be indifferent between the bundles $(2, 2)$ and $(5, 2)$, since the additional 3 left shoes provide no additional utility.)

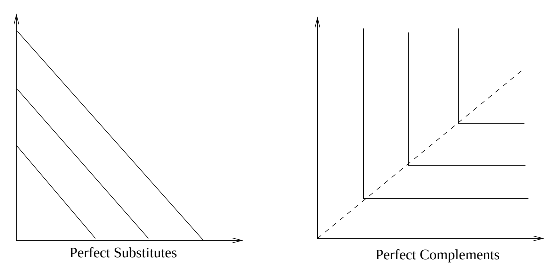

2.2 Budget Set

Consider a consumer with income I and suppose the prices of goods are p_x and p_y, respectively. The **budget set** is the set of all consumption bundles which the consumer can afford given their income and given the prices of goods. The budget set is given by:

$$B(p_x, p_y, I) = \left\{(x, y) \in \mathbb{R}^2_+ \mid p_x x + p_y y \leq I\right\}$$

The frontier of the budget set is known as the **budget constraint** or **budget line**. All consumption bundles along the budget line are just affordable, i.e., if the consumer purchases a bundle along this line, then they will exhaust their income. A point on the interior of the budget set represents a bundle where the consumer is not spending all of their income. Any point along the budget constraint must satisfy:

$$p_x x + p_y y = I$$

It is easy to verify that the intercepts of the budget line are $\left(\frac{I}{p_x}, 0\right)$ and $\left(0, \frac{I}{p_y}\right)$ and that the slope of the budget line is $slope = -\frac{p_x}{p_y}$. The slope of the budget constraint represents the opportunity cost of consuming one extra unit of good X (in terms of the amount of good Y that must be forgone). Note: the slope of the budget constraint is based on the relative price of goods and is independent of the level of income.

Example 13. A consumer has income $I = 24$ and faces prices $p_x = 3$ and $p_y = 4$. If they spend all their income on good X, then they can purchase 8 units of X. (This corresponds to the X- intercept.) Instead, if they spends all of their income on good Y, then they can purchase 6 units of Y. (This corresponds to the Y- intercept.) For each additional unit of good X that they consume, they must forgo $\frac{3}{4}$ units of good Y. This is the opportunity cost of good X. You should verify that the bundles $(4, 3)$ and $(6, 1.5)$ are just affordable, that the bundle $(2.5, 4)$ lies in the interior of the budget set, i.e., the consumer will have surplus income if they purchase this bundle, and that the bundle $(5, 4)$ is not affordable.

2.3 Optimisation

So far, we have considered questions of what consumers can afford, and what they would like to purchase, in isolation. We now bring these two concepts together. Consumers are assumed to be rational utility maximisers. Hence, they will choose the bundle which is most preferred from amongst those which are affordable. Such a bundle must lie on the highest possible indifference curve which is contained within the budget set.

We can reformulate the consumer's problem in the following way:

$$\max_{x,y} U(x, y)$$
$$\text{s.t.} \quad p_x x + p_y y = I$$

We can solve this problem using constrained optimisation techniques. The solution is the consumer's demand.

2.3.1 Two Examples

Example 14. Suppose the consumer's utility function is given by $u(x, y) = \log x + \log y$ and that $p_x = 2$, $p_y = 3$ and $I = 12$. Then we have:

$$\begin{aligned} \max_{x,y} U(x, y) &= \log x + \log y \\ \text{s.t. } 2x + 3y &= 12 \end{aligned}$$

The Lagrangian is:

$$\mathcal{L}(x, y, \lambda) = \log x + \log y - \lambda(2x + 3y - 12)$$

The first order conditions are:

$$\frac{\partial \mathcal{L}}{\partial x} = \frac{1}{x} - 2\lambda = 0$$
$$\frac{\partial \mathcal{L}}{\partial y} = \frac{1}{y} - 3\lambda = 0$$
$$\frac{\partial \mathcal{L}}{\partial \lambda} = 2x + 3y - 12 = 0$$

Using the first two conditions, we have:

$$\lambda = \frac{1}{2x} = \frac{1}{3y}$$
$$x = \frac{3}{2}y$$

Substituting this into the third condition gives:

$$12 - 2\left(\frac{3}{2}y\right) - 3y = 0$$
$$12 - 6y = 0y^* \qquad = 2$$

Then, using the fact that $x = \frac{3}{2}y$, $x^* = 3$.

In the above example, we used specific values for prices and incomes. This generated a specific equilibrium bundle. If we keep the analysis more general, and solve the problem for generic p_x, p_y and I, then we can find the consumer's demand functions for any price and income level.

Example 15. Suppose the consumer's utility function is given by $U(x, y) = -\frac{1}{x} - \frac{1}{y}$ (believe it or not, this satisfies the conditions for a utility function!). We have:

$$\begin{aligned} \max_{x,y} u(x, y) &= -\frac{1}{x} - \frac{1}{y} \\ \text{s.t. } p_x x + p_y y &= I \end{aligned}$$

We use the method of Lagrange Multipliers. We have:

$$\mathcal{L} = -\frac{1}{x} - \frac{1}{y} + \lambda(p_x x + p_y y - I)$$

The first order conditions are:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial x} &= \frac{1}{x^2} - \lambda p_x = 0 \\ \frac{\partial \mathcal{L}}{\partial y} &= \frac{1}{y^2} - \lambda p_y = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= I - p_x x - p_y y = 0 \end{aligned}$$

Combining the first two conditions, we have:

$$\begin{aligned}\lambda &= \frac{1}{x^2 p_x} = \frac{1}{y^2 p_y} \\ \rightarrow & \ y = \sqrt{\frac{p_x}{p_y}} x\end{aligned}$$

Substituting this into the third gives:

$$\begin{aligned}p_x x + p_y \sqrt{\frac{p_x}{p_y}} x &= I \\ x^* \left(p_x, p_y, I\right) &= \frac{I}{\sqrt{p_x}\left(\sqrt{p_x} + \sqrt{p_y}\right)}\end{aligned}$$

and then from the optimality condition:

$$y^* \left(p_x, p_y, I\right) = \frac{I}{\sqrt{p_y}\left(\sqrt{p_x} + \sqrt{p_y}\right)}$$

□

General Insights Return to the generic problem. The problem is:

$$\begin{gathered}\max_{x,y} U\left(x, y\right) \\ \text{s.t.} \quad p_x x + p_y y = I\end{gathered}$$

Again formulate the Lagrangian. We have:

$$\mathcal{L} = u(x, y) + \lambda \left(p_x x + p_y y - I\right)$$

The first order conditions are: The first order conditions are:

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial x} &= \frac{\partial u}{\partial x} - \lambda p_x = 0 \\ \frac{\partial \mathcal{L}}{\partial y} &= \frac{\partial u}{\partial y} - \lambda p_y = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= I - p_x x - p y_y = 0\end{aligned}$$

Combine the first two conditions, and solve out for λ. Then, we have a system of two equations in two variables which we can use to solve for the optimal demands, x and y. We have:

$$\begin{gathered}-\frac{MU_x}{MU_y} = -\frac{p_x}{p_y} \\ p_x x + p_y y = I\end{gathered}$$

Two important insights:

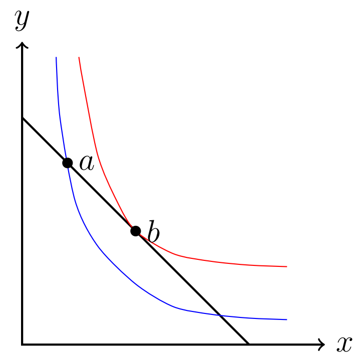

Figure 2.3: Tangency

1. The optimal consumption bundle will always lie on the budget constraint, rather than in the interior of the budget set. (Why? Suppose the consumer chooses a bundle in the interior of the budget set. Then he does not spend all of his income. Such a bundle cannot be optimal, since (by monotonicity) the consumer could do better by spending the surplus income, and purchasing a little more of each good.)

2. The optimal consumption bundle will satisfy the tangency condition:

$$\begin{aligned} MRS &= -\frac{p_x}{p_y} \\ slope\ of\ IC &= slope\ of\ BC \end{aligned}$$

(We see this immediately from the above condition. Alternatively, I offer a proof by contradiction. Suppose the optimal bundle is a, and at this point indifference curve and budget constraint are not tangential. Then, at a, the curves must cross. But then, there is a point along the indifference curve (b, say) which is in the interior of the budget set. Moreover, a and b are indifferent - and so if a is an optimal choice, then so is b. But we know that b cannot be optimal, since it is in the interior. There is a point on the budget line (c, say) which is strictly preferred to b. Furthermore, by monotonicity, c is preferred to a. But then a cannot be an optimal choice, since c is affordable and c is preferred to a. This is a contradiction.)

The second result says that the rate at which the consumer is willing to forgo good Y for good X, i.e., the MRS, is the same as the rate which he must actually exchange the goods in the market, i.e., the opportunity cost. If this is the case, then it is impossible for the consumer to reallocate goods in such a way as to increase his utility. Why? Let $MRS = slopeBC = 2$. Suppose the consumer tries to increase their utility by consuming 1 extra unit of good X. In order to do so, they must forgo 2 units of good Y (since the slope of the budget constraint - which is the opportunity cost of consuming X - equals 2). But since $MRS = 2$, the consumer is exactly willing to forgo 2 units of good Y in order to get 1 extra unit of X, and moreover, they are indifferent between these alternatives. Hence, such a reallocation of consumption leaves the consumer just as well off - but no better off. By the same argument, the consumer cannot increase their utility by consuming 1 fewer unit of good X.

Suppose instead $MRS = 3$ and $slopeBC = 2$, and the consumer decides they want to reallocate goods in such a way to get 1 extra unit of X. To be able to afford this, they must forgo 2 units of Y. But they are willing to forgo 3 units of Y. If forgoing 3 units of Y leaves them equally well off, then forgoing only 2 units will make them better off. Hence, there is an opportunity to reallocate goods in such a way that improves utility. This cannot be optimal. Thus, if $MRS \neq slopeBC$, then there is always an opportunity for the consumer to reallocate their consumption in such a way as to improve their welfare. It follows that, at the optimum, $MRS = slopeBC$.

A different way of interpreting the tangency condition is in terms of marginal-benefit/marginal-cost analysis. Re-write the tangency condition as:

$$MU_x = -\frac{p_x}{p_y} MU_y$$

The left hand term is the marginal benefit of purchasing one more unit of good X. The right hand side is the product of two terms. $-\frac{p_x}{p_y}$ is the number of units of good Y that must be forgone to afford one additional unit of good X, given the budget constraint. MU_y is the utility lost from consuming one fewer unit of good Y. Since the agent must consume $\frac{p_x}{p_y}$ fewer units, the right hand side gives the opportunity cost of consuming one more unit of good X, in terms of lost utility from forgone consumption of good Y. The tangency condition states that at the optimum, the net benefit of consuming one more unit of good X (and commensurately less of good Y) must be zero. Why? If the net benefit were positive, then the agent would do strictly better to reallocate their consumption bundle by adding more of X. By contrast, if the net benefit were negative, then the opposite is true —the agent would do strictly better to reallocate their consumption bundle by adding more of Y.

2.4 Comparative Statics

It should be clear from the above discussion that the agent's consumption choice depends upon the prices of goods and their level of income. We now consider the effect of a change in these factors on consumption choices.

2.4.1 Changes in Income

First, consider the effect of a change in income. For concreteness, suppose there is an increase in income. Since prices are unchanged, the consumer can now afford to purchase more of each good. The budget constraint shifts outward in a parallel fashion. (Why? Recall the slope of the budget constraint is given by the ratio of prices. Since prices have not changed,

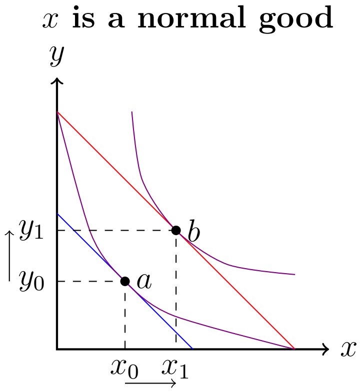

x and y both increase.

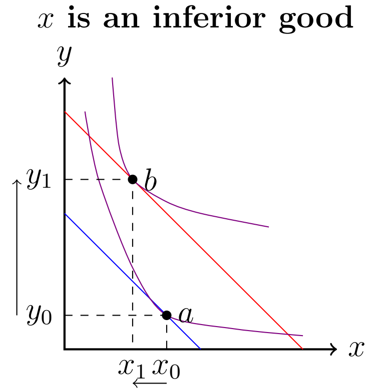

x increases and y decreases.

Figure 2.4: Effect of an increase in income

the slope of the budget constraint remains the same. Hence, there must be a parallel shift in the budget line.)

Although income has increased, it does not immediately follow that agents will increase their consumption of a good. Recall - the effect of income on demand depends on whether the good is a normal good or an inferior good. We illustrate this with Figure 2.4.

Note —in both cases above, good Y behaves as a normal good, i.e., when income increases, the demand for good Y increases. It should be clear that whilst one good (either X or Y) can be inferior, it is not possible for both goods to be inferior. (Why? Suppose both goods were inferior goods. Then, following an increase in income, the demand for both goods would decrease. But this means the consumer would be choosing a bundle on the interior of the budget set. We argued above, that this is never optimal. Hence, when income increases, the demand for at least one good must increase.) We conclude that, in general, not all goods can be inferior - at least one good in the economy must be normal.

2.4.2 Changes in Prices

Next, consider the effect of a change in the price of one good. For concreteness, suppose there is an increase in the price of good X. Since there is a change in the relative price of goods, the slope of the budget constraint will now change. Note further that if the agent spends all their income on good Y, then the amount of good Y they can purchase will not change (since only the price of good X has changed). Hence, the $y-$intercept does not shift. For each quantity of Y chosen, the agent is now able to afford less of good X. Hence, the budget line 'swivels' inwards. We have:

A few comments are worth noting:

- In the above diagrams, the demand for good X obeys the law of demand. Since the price of good X increases, the quantity demanded falls. It is, however, possible for the 'law of demand' to be violated (see below).

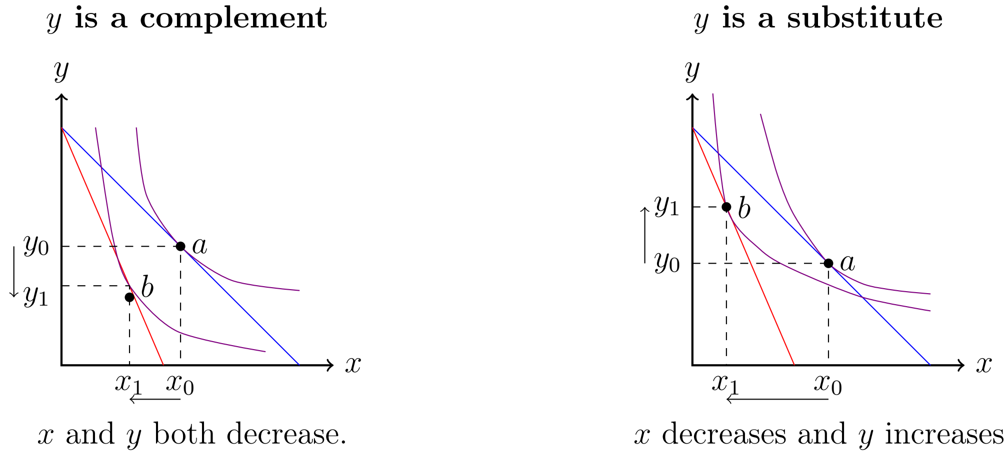

Figure 2.5: Effect of an increase in the price of good x

- Furthermore, in the left hand figure, the quantity demanded of good Y also falls. In this case, it must be that X and Y are *complements*
- In the right hand figure, the quantity demanded of good Y increases. Hence, in this case, X and Y must be *substitutes.*
- In each case, regardless of whether quantities of goods increase or decrease, the agent is forced onto a lower indifference curve. The reason for this should be intuitive. When facing higher prices, there are fewer bundles which are feasible for the consumer to choose. Hence, they must make their choice from a smaller set of alternatives.

Income and Substitution Effects In the above examples, the increase in price resulted in a decrease in quantity demanded for two reasons. First, when the price is higher, consumers have an incentive to substitute their consumption towards cheaper alternatives, since the marginal (opportunity) cost of consuming the good has now increased. Second, when the price is higher, the purchasing power of income decreases, and so demand falls. (To make this latter point most stark, suppose my income is \$100 which I will always spend on DVDs irrespective of price. When the price of DVDs is \$10, I will purchase 10 DVDs. If the price increases to \$20, then the quantity I purchase will fall to 5 - not because I am substituting towards other goods, but simply because the purchasing power of my income is no longer as strong.) We refer to these effects as the substitution and income effects, respectively. In this section, we will make these concepts precise.

Consider the diagram below. The agent originally chooses bundle A when the price is low. After the price increase, the agent actually chooses bundle B, which is on a lower indifference curve. Suppose we (hypothetically) compensate the consumer by increasing their income just enough so that they return to the original indifference curve. (We do this by shifting the new budget constraint BC_1 outwards just until it is tangent with the original indifference curve - i.e. to BC_h.) With this compensation, the agent chooses bundle C, rather than the original bundle A. (We refer to bundle C as the *compensated demand* or *Hicksian demand.*)

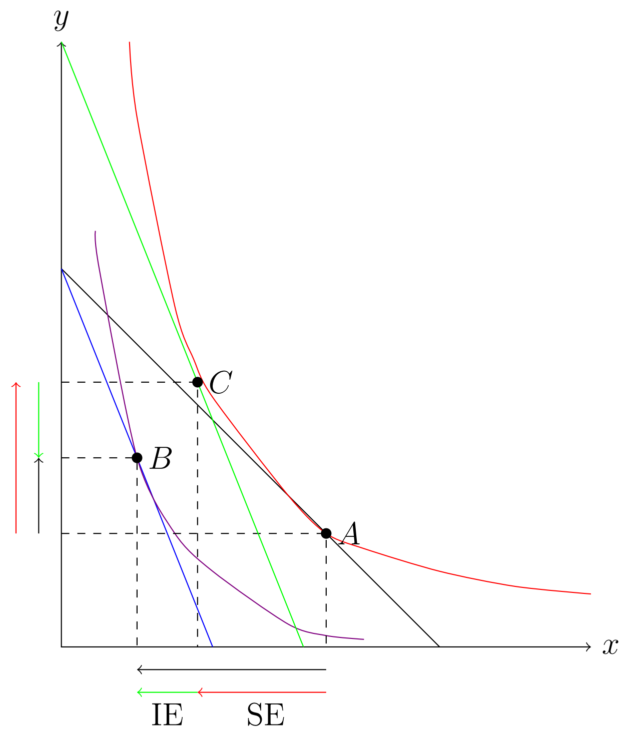

Figure 2.6: Income and Substitution Effects

What does this mean? Since we compensate the agent, 'affordability' is not an issue. The agent simply chooses the bundle of goods which makes them just as happy as they were before the price change. This is the **substitution effect**. It measures the change in demand which arises purely because the agent reallocates their consumption - taking out the effect of affordability considerations. In the above diagram, the substitution effect refers to the fall in demand from x_0 to x_h (where x_h refers to the hypothetical level of demand that arises when the consumer is compensated for the price change).

For any price increase, the substitution effect must be negative, i.e., the hypothetical bundle must contain more of good X and less of good Y. Why? Since the hypothetical bundle is chosen optimally, it must satisfy the tangency condition. We know that:

$$MRS = \frac{p_x}{p_y} = slopeBC$$

Since p_x has increased, the slope of the new budget constraint is higher, and so at the optimum, the MRS must also be higher. But we know that the MRS is decreasing as we move along the demand curve. Hence, the new tangency will occur at a point to the left of bundle A, where MRS is higher. At this new point, the amount of good X will be lower and the amount of good Y will be higher.

The **Law of Demand** is the claim that the *substitution effect* is negative.

Having (hypothetically) compensated the agent and found the substitution effect, we then take this compensation away from the consumer (so that they are returned to the original level of income). This involves a parallel shift in the budget constraint, from the hypothetical budget constraint BC_h, to the new budget constraint BC_1. Note - an increase in price results in a *decrease* in income (since the purchasing power of income is now relatively lower). Since this change involves a pure decrease in income, with no change in prices - we can interpret this as a measure of the change in demand due to considerations of affordability. The income effect is the change in demand from x_h to x_1.

A few comments about the income effect might be appropriate at this stage. Whilst an increase in price results in a decrease in income, this does not immediately imply that quantity demanded will decrease. From the previous section, we know that the effect of a change in income on demand depends on whether a good is normal or inferior. We have:

- If a good is normal, then an increase in price causes a decrease in income, which in turn causes a decrease in quantity demanded, i.e., the income effect is negative. This complements the decrease in demand which arises due to the substitution effect. Hence, for a normal good, an increase in demand unambiguously results in a decrease in quantity demanded.
- If a good is income neutral, then there is no income effect.
- If a good is inferior, then an increase in price causes a decrease in income, but this causes quantity demanded to increase, i.e., the income effect is positive. Then, the income effect counteracts the substitution effect. If the magnitude of the substitution effect is large relative to the income effect, then the overall effect of a price increase will still be for the price to decrease. (See diagram below.)

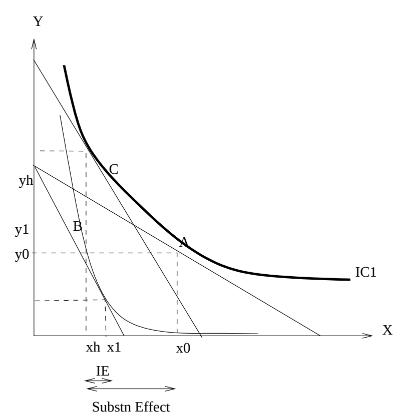

- It is theoretically possible for the income effect to be larger than the substitution effect. If so, then quantity demanded actually rises in response to an increase in price! (See diagram below.) Such a good is known as a **Giffen good**. Whilst their existence is theoretically possible, few examples of Giffen goods exist.

But doesn't this contradict the law of demand? Yes and no. The statement of the law of demand above is somewhat sloppy. The law should more properly state that an increase in

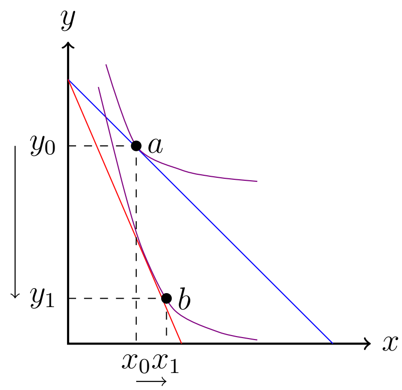

Figure 2.7: x is a Giffen good

prices causes the *compensated* demand to fall, i.e., the substitution effect is negative. The law of demand is silent on the effect of income changes. So in this sense, a positive income effect is perfectly consistent with the law of demand. (It is also worth noting that there are very few examples of Giffen goods. So even though it is not absolutely correct, it turns out that our first version of the law of demand holds in the vast majority of cases.)

2.4.3 Demand Homogeneity

What happens to the optimal consumption bundle if all prices and income double? Then, the relative prices of goods are unchanged (since both prices increase by the same proportion). Similarly, the purchasing power of income is unchanged, since income increases at the same rate as prices. It follows that in this scenario, the budget line is unchanged. (To convince yourself of this, suppose $p_x = 5, p_y = 2$ and $I = 20$ and draw the budget line. Then double all of these, so that $p_x = 10, p_y = 4$ and $I = 40$. Verify that the two budget lines coincide.)

We refer to this property as **demand homogeneity**. If income and all prices increase (or decrease) by the same proportion, then there is no change in the consumer's optimal choice. It follows that an agent's choice does not depend on actual prices and income, but rather on relative prices,i.e., opportunity cost, and real income (purchasing power). Since the actual price levels do not matter, I can scale these up and down as I see fit, and this will not affect the consumer's decision. Often, for convenience, we set the price of some 'base' good equal to one, and then rescale all other prices and income relative to this good. This good is known as the **numeraire**.

2.4.4 Summary

You should have noticed in the above sections that we have a lot of freedom to draw indifference maps in particular ways to generate particular outcomes. Depending on how we 'bend' indifference curves, we can cause goods to be normal or inferior, complements or substitutes, or even Giffen goods. At first glance, it might appear that we're cheating (maybe if you're creative enough in drawing curves, you can show anything?) In response to this criticism, I make the following two points. First, given that different consumers can have vastly different preferences, we should take comfort (rather than concern) at the fact that our theory is so versatile! (Recall - the only assumptions we placed on preferences were that indifference curves were downward sloping and convex, and that they did not intersect. This allows for a wide variety of preferences across goods.)

Second, it is not true that we can show anything simply by bending the indifference curve in the appropriate way. Consumer choices must satisfy a variety of consistency properties. We have already outlined a few:

- The substitution effect must be negative.
- Not all goods can be inferior —there must be at least one normal good.
- Not all goods can be Giffen goods —at least one good must obey the Law of Demand.
- Not all goods can be complements —there must be at least one substitute good.
- Demand homogeneity must be satisfied.

Moreover, there are several other properties that must be satisfied which we will not cover here. (If you are interested, look up the following: Engel Aggregation, Cournot Aggregation and Slutsky Symmetry.)

2.4.5 Deriving the Demand Curve

So far, we have found the quantity demanded for a good, for a given price and income level. If we repeat this procedure for a variety of different prices (holding income and the price of other goods constant), then we can find the relationship between price and quantity demanded for a particular good. This is the demand curve. We illustrate this process diagrammatically. In the top diagram, we plot the budget line and optimal consumption bundles for three different values of p_x (holding p_y and I fixed at some arbitrary level), where $p_0 < p_1 < p_2$. Then, in the bottom diagram, we plot these optimal consumption choices against the price level which induced them.

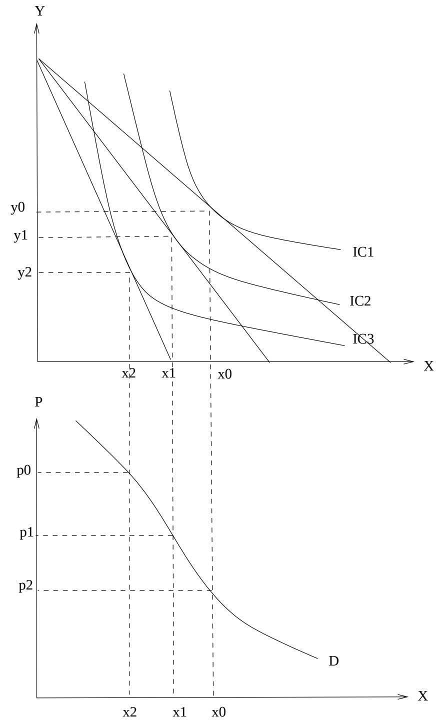

In the first section, we said that a change in income caused a shift in the demand curve. We can easily show this to be the case as well. For concreteness, suppose the good is normal, and there is an increase in income. Then, in the top diagram, each of the 3 budget constraints will shift outwards in a parallel fashion. This will cause the optimal bundles x_0, x_1 and x_2 to also increase. Then, plotting these new demand points on the lower diagram (and noting that prices have not changed), we see that the entire demand curve has shifted to the right. (Verify this for yourself. To simplify the diagram, consider only 2 prices, rather than 3.)

Chapter 3

Applications of Theory

In the previous chapter, we used utility theory to analyse consumer choices between two different goods. In this chapter, we expand the analysis to increase choices over other 'goods' —such as the decision of how much to spend and save, or the decision of how much time to spend working and how much to spend in leisure.

3.1 Optimal Inter-temporal Choice - The Consumption/Saving Decision

Consider an individual who lives for two periods. The consumer earns income Y_1 and Y_2 in periods 1 and 2 (respectively), and must decide how much to consume in each period. In the first period, the consumer can borrow and lend at an interest rate r. The loan is paid back with interest in the second period. (Obviously there can be no borrowing or lending in the second period, as there is no third period in which to settle the loans).

3.1.1 inter-temporal Budget Constraint

First, consider the inter-temporal budget constraint. In period one, suppose the individual consumes C_1 and saves $S = Y_1 - C_1$. (Note —the consumer may consume more than he earns in the first period. If so there is negative saving - i.e. borrowing.) In period two, the individual will receive back his savings with interest (or he will have to pay back his loan with interest). Hence, total consumption in period 2 is $C_2 = Y_2 + (1 + r) S$. (It should be clear that if period one consumption exceeds period one income, then period two consumption must be less than period two income —since the agent must pay back his loan in the second period. The converse is also true.)

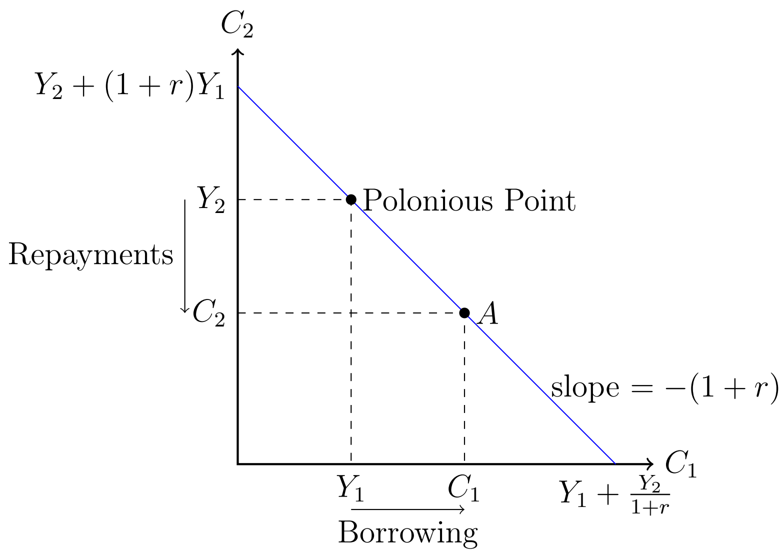

Figure 3.1: Inter-temporal Budget Constraint

Combining these expressions, we have:

$$\begin{aligned} C_2 &= Y_2 + (1+r)(Y_1 - C_1) \\ C_1 + \frac{1}{1+r}C_2 &= Y_1 + \frac{1}{1+r}Y_2 \end{aligned}$$

The right hand side expression is called the *net present value of income* or **net present wealth**. (To see why this should be the case —note that if I have $\frac{1}{1+r}Y_2$ today and I invest it at interest rate r, then (with interest) it will be worth Y_2 in the next period. Hence receiving Y_2 worth of income in the second period is equivalent to receiving $\frac{1}{1+r}Y_2$ worth of income today. This is why we refer to the expression as the *present value*.) It may be helpful to write $W = Y_1 + \frac{1}{1+r}Y_2$. We have:

$$C_1 + \frac{1}{1+r}C_2 = W$$

which is the usual form of our budget constraint. Note that the price of period 1 consumption is 1 (period one consumption is the numeraire), whilst the price of period 2 consumption is $\frac{1}{1+r}$. The slope of the budget constraint is given by $-(1+r)$. (You should confirm this for yourself). We can draw the inter-temporal budget constraint as follows:

The expressions for the intercepts are worth explaining. The vertical intercept refers to a situation in which the agent saves all of his period one income (which generates interest) and the spends this in the second period along with his period two income. Hence, his total period two consumption is $Y_2 + (1+r)Y_1$. On the other hand, the horizontal intercept refers to the situation in which the agent borrows against his entire second period income, in the first period - so that his entire second period income is spent repaying the loan. Clearly this is the maximum amount he can borrow, since if he tried to borrow a higher amount, he would be unable to repay it. The maximum amount he can borrow is $\frac{1}{1+r}Y_2$, since with interest, this amounts to Y_2. In the above diagram, I drew a situation where the agent borrows slightly in the first period (and so consumes more than his first period income), but has to repay this loan in the second period, and so consumes less in the second period than his second period income.

3.1.2 Inter-temporal Preferences

Next, we describe agent's inter-temporal preferences. The usual assumptions over preferences carry over. We assume preferences are monotonic, and so consumers prefer more consumption in either period to less. Furthermore, we assume preferences are convex, and so consumers tend to prefer similar levels of consumption in both periods, rather than large consumption in one period and very little in the next. (This latter behaviour is often referred to as 'consumption smoothing'). Although it is not crucial to the analysis, we assume moreover that consumption is a normal good - so that when wealth increases, consumers will prefer to increase their consumption in both periods.

We can represent inter-temporal preferences using indifference curves. The inter-temporal rate of substitution is the rate at which consumers are willing to forgo future consumption for 1 extra unit of current consumption. (ITRS is the inter-temporal analogue of the MRS). From the above discussion, we can calculate the inter-temporal rate of substitution by:

$$ITRS = \frac{MU_1}{MU_2}$$

where $U\left(C_1, C_2\right)$ is the utility function defined over consumption levels in the two periods, and $MU_i = \frac{\partial U}{\partial C_i}$ for $i = 1, 2$.

It is common to represent inter-temporal preferences by

$$U(C_1, C_2) = u(C_1) + \beta u(C_2)$$

where $\beta \in (0, 1)$ and u is the 'per period' utility function, assumed to be increasing ($u' > 0$) and concave ($u'' < 0$). The former says that more consumption is better, and the second reflects diminishing marginal utility of consumption.

The parameter β is known as the discount factor, and reflects the present-bias of the consumer. To see this, note that now:

$$IRTS = -\frac{u'(C_1)}{\beta u'(C_2)}$$

If $C_1 = C_2$, so that consumption is perfectly smoothed across time, then $IRTS = -\frac{1}{\beta} < -1$. The agent would be willing to forgo more than 1 unit of future consumption in order to get one more unit of current consumption. The agent is biased towards current consumption over future consumption. Of course, if $C_1 > C_2$, so that consumption is already skewed towards the present, then $u'(C_1) < u'(C_2)$ (by the assumption that $u'' < 0$), and so $IRTS$ may be less steep. Since consumption is already skewered, the agent will not be willing to forgo much future consumption to skew it further towards the present.

Sometimes, rather than using the discount factor β, we instead caputre the agent's present bias using the subjective rate of time preference $\rho > 0$. We define: $\beta = \frac{1}{1+\rho}$, so that, for example, if $\beta = 0.95$, then $\rho \approx 0.05$. The agent's lifetime utility then becomes:

$$u(C_1) + \frac{1}{1+\rho} u(C_2)$$

Notice that this has the same flavor of discounting the future as was present in the NPV calculations to find the lifetime budget constraint. The difference is that the agent discounts future utility according to her subjective rate of time preference ρ, whereas she discounts future incomes according to the market interest rate r.

The inter-temporal rate of substitution now becomes:

$$IRTS = -\frac{u'(C_1)}{\frac{1}{1+\rho} u'(C_2)} = -(1+\rho)\frac{u'(C_1)}{u'(C_2)}$$

We can similarly interpret ρ. If $C_1 = C_2$, so that consumption is perfectly smoothed, then the agent requires $1+\rho$ units of consumption in period 2 to compensate for a 1 unit decrease in period 1 consumption.

3.1.3 Equilibrium

The agent's problem is:

$$\max_{C_1, C_2} u(C_1) + \beta u(C_2) \ s.t. \ C_1 + \frac{C_2}{1+r} = Y_1 + \frac{Y_2}{1+r}$$

The Lagrangian is:

$$\mathcal{L} = u(C_1) + \beta u(C_2) - \lambda(C_1 + \frac{C_2}{1+r} - Y_1 - \frac{Y_2}{1+r})$$

The first order conditions are:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial C_1} &= u'(C_1) - \lambda = 0 \\
\frac{\partial \mathcal{L}}{\partial C_2} &= \beta u'(C_2) - \frac{\lambda}{1+r} = 0 \\
\frac{\partial \mathcal{L}}{\partial \lambda} &= Y_1 + \frac{Y_2}{1+r} - C_1 - \frac{C_2}{1+r} = 0
\end{aligned}$$

By the first two equations, we have:

$$\lambda = u'(C_1) = \beta(1+r)u'(C_2)$$

This is *Euler Equation*. If we expressed preferences in terms of ρ rather than β, the Euler Equation becomes:

$$u'(C_1) = \frac{1+r}{1+\rho}u'(C_2)$$

Now, suppose $\beta(1+r) = 1$ (or equivalently $\rho = r$). Then, the rate at which the agent wishes to trade-off consumption across time coincides with the rate at which market demands the trade-off. Then $u'(C_1) = u'(C_2)$, and since $u'' < 0$, this implies that $C_1 = C_2$. The agent perfectly smoothes consumption through time.

By contrast, suppose that $\beta(1+r) > 1$ (or equivalently, $\rho > r$. Then the rate at which the agent can transform period 1 consumption into period 2 consumption in the market rate, is larger than the rate that the agent requires to keep him indifferent. Hence, the agent will want to push her consumption into the future. By the Euler equation, $u'(C_1) > u'(C_2)$, which implies that $C_1 < C_2$ (since $u'' < 0$).

Similarly, if $\beta(1+r) < 1$ (or equivalently, if $\rho < r$), the agent will front-load her consumption.

Example 16. Consider an agent with utility $U = \ln C_1 + \beta \ln C_2$. The agent works during the first period and earns income $Y_1 = 100$. In the second period, the agent retires and but receives a pension of \$50. What is the optimal inter-temporal allocation of consumption? What is the level of saving?

The Euler condition implies:

$$\begin{aligned}
\frac{1}{C_1} &= \beta(1+r)\frac{1}{C_2} \\
C_2 &= \beta(1+r)C_1
\end{aligned}$$

Substituting this into the budget constraint, we have:

$$\begin{aligned}
C_1 + \frac{\beta(1+r)C_1}{1+r} &= Y_1 + \frac{Y_2}{1+r} \\
(1+\beta)C_1 &= Y_1 + \frac{Y_2}{1+r} \\
C_1^* &= \frac{1}{1+\beta}Y_1 + \frac{1}{(1+\beta)}\frac{Y_2}{1+r}
\end{aligned}$$

Then: $c_2^* = \frac{\beta}{1+\beta}(1+r)Y_1 + \frac{\beta}{1+\beta}Y_2$ and:

$$\begin{aligned}
S^* &= Y_1 - C_1^* \\
&= (1 - \frac{1}{1+\beta})Y_1 - \frac{1}{1+\beta}\frac{Y_2}{1+r} \\
&= \frac{\beta}{1+\beta}Y_1 - \frac{1}{1+\beta}\frac{Y_2}{1+r}
\end{aligned}$$

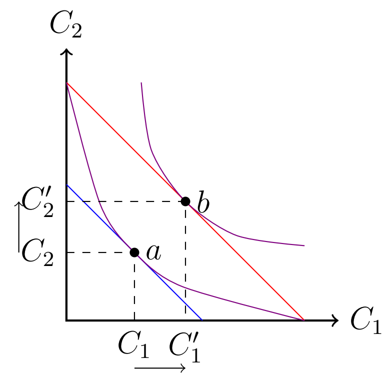

Figure 3.2: Effect of an Increase in Wealth (due to an increase in either either Y_1 or Y_2.

When does the agent save? $S^* > 0$ is:

$$\frac{\beta}{1+\beta}Y_1 > \frac{1}{1+\beta}\frac{Y_2}{1+r}$$
$$Y_2 < \beta(1+r)Y_1$$

Hence, if future income is small relative to current income, the agent will save, and vice versa. □

3.1.4 Comparative Statics

What happens when income increases? Suppose there is an increase in Y_1. This causes net present wealth to increase, and so the inter-temporal budget constraint shifts outward. Since, consumption is normal good, this will cause consumption to increase in both periods.

Note - although the agent's income only increases in period 1, his consumption increases in both periods, because of the incentive to smooth consumption. In the above diagram, when first period income increases, the agent further increases his level of saving, in order to sustain a high level of period 2 consumption. It should be clear that if Y_2 had increased rather than Y_1, the consumer would still increase consumption in both periods - but in this case, he would borrow in the first period against the higher second period income.

What happens when the interest rate changes? (Recall - the interest rate determines the relative price of second period consumption in terms of first period consumption.) Suppose there is an increase in the interest rate. The budget constraint swivels about the point (Y_1, Y_2). (To see why, note that if the consumer did not borrow or lend, and simply consumed his income in each period, then his choice is unaffected by the interest rate. Hence, the point (Y_1, Y_2) must lie on every budget constraint, regardless of the level of interest rates.) We illustrate this graphically:

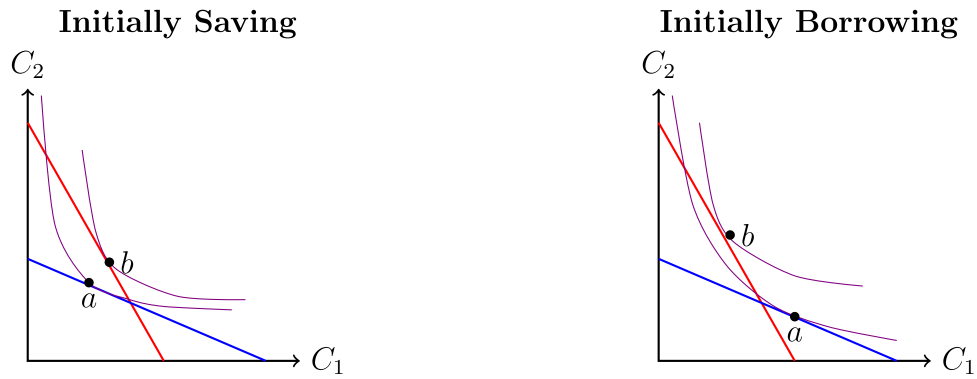

Figure 3.3: Effect of an Increase in Interest Rates

3.2 Allocation of Time - The Consumption/Leisure decision

3.2.1 Time and Budget Constraints

Consider an individual with T hours of discretionary time each week, which he may allocate between work (Z) and leisure (L). Leisure does not involve any monetary cost - it is simply time not spent working. The agent has unearned income of M, and earns an hourly wage rate w for each hour worked. The time constraint is given by:

$$z + L = T$$

The agent works in order to earn income to purchase goods and services. We set the price of consumption equal to one, and so the budget constraint is:

$$\begin{aligned} C &= M + wZ \\ C &= M + w(T - L) \\ C + wL &= M + wT \end{aligned}$$

This last expression is particularly insightful. The expression on the right hand side is known as the level of **full income**. It is the total possible income the consumer could have, if he spent all of his discretionary time working. For each hour not worked (i.e. spent in leisure), the consumer forgoes w worth of consumption. Hence, the wage rate w is the *shadow price* (opportunity cost) of leisure. We can think of the consumer 'purchasing' leisure at a cost of w, for each hour he chooses not to work.

Note —whilst it appears that there are 3 choice variables for the agent, C, L and Z - in reality the consumer only chooses one of these. The others follow immediately. For example - if the agent decides that he will only work for 20 hours per week, then this both determines his income (and hence consumption), as well as the amount of time he has left over for leisure.

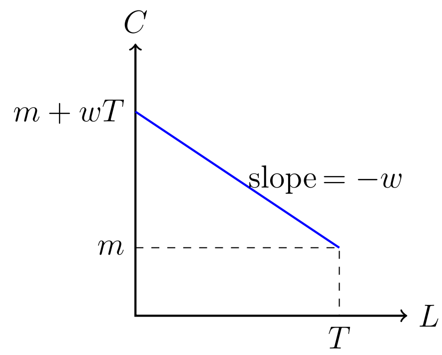

Figure 3.4: Allocation of Time and the Budget Constraint

3.2.2 Preferences

We assume that agents derive utility out of both consumption and leisure. Agents actually derive disutility from working, but they may still choose to work in order to be able to consume more goods. The usual indifference map follows. Consumption and leisure are both 'goods', in the sense that the agent prefers to bundles which have more of either, to ones which have less. Moreoever, we maintain our usual convexity assumption. Agents prefer to have moderate levels of both leisure and consumption - and this is preferable to being rich but having no leisure, or to having a lot of leisure time, but not being able to clothe or feed oneself. Finally, we assume that consumption and leisure are both normal goods —when unearned income increases, agents will increase both their material consumption and the amount of time spent in leisure.

3.2.3 Equilibrium

The equilibrium consumption/leisure choice must satisfy the tangency condition (as usual). Diagrammatically, we have:

Comparative Statics Consider the effect of an increase in unearned income m. For each level of work effort, the consumer can now purchase more goods and services. Hence the budget constraint shifts vertically upwards. There is a pure income effect. Since both consumption and leisure are assumed to be normal goods, then this will cause both to increase. It follows that work effort falls. Why? The agent can now satisfy more of his consumption using his unearned income - hence he has less need to work - and is able to devote more time to leisure. The agent achieves a higher level of utility, since he has a greater level of income at his disposal (and hence can choose amongst a greater range of bundles).

Next, consider the effect of an increase in the wage rate w. This causes the relative price of consumption/leisure to change, and so the slope of the budget constraint changes. More precisely, leisure now becomes more expensive, since for every hour spent in leisure, the agent

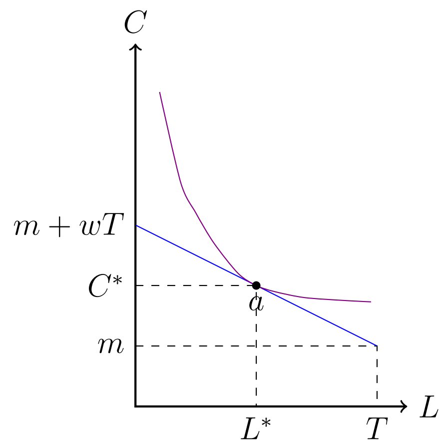

Figure 3.5: Caption

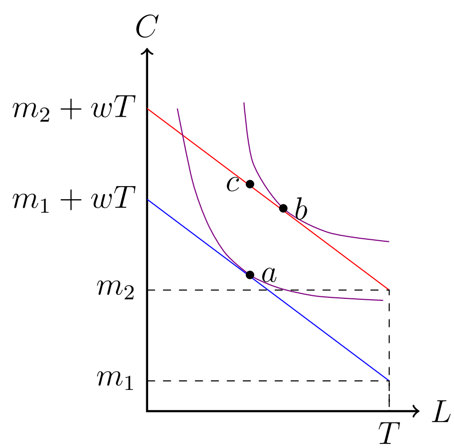

Figure 3.6: Effect of an Increase in Unearned Income

forgoes a greater amount of income/consumption. Hence, the budget constraint becomes steeper. There are two effects at play:

1. *Substitution Effect* - since leisure is now more costly, agents will tend to substitute from leisure to work. This increases the level of work and consumption and decreases the level of leisure.

2. *'Income' Effect* - for each hour worked, the agent now receives more income. Hence, can satisfy his consumption needs without needing to work as many hours. This increases the level of consumption and leisure, but decreases the level of work effort.

There is an unambiguous increase in the level of consumption. The effect of a wage increase in work effort is ambiguous. On the one hand, the substitution effect causes work effort to increase, whilst on the other, the income effect causes the level of work effort to decreases. The net effect will depend on the relative magnitudes of these two effects. In the above

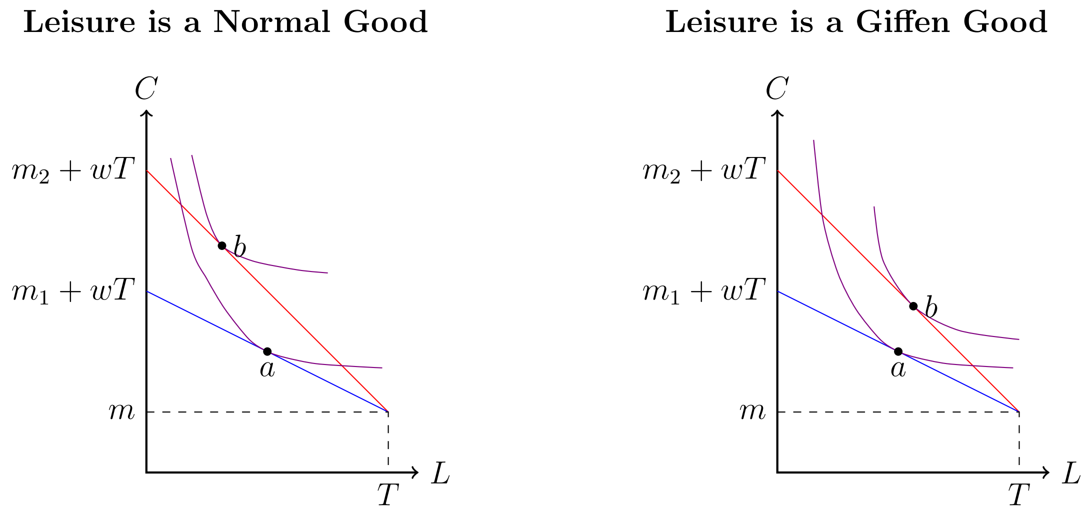

Figure 3.7: Effect of an increase in the wage.

diagram, the substitution effect is larger than the income effect, and so the labour supply increases (and leisure decreases).

Note —in the above example, we generated the usual story regarding the effect of wages on labour supply. If agents are paid more, then they will work harder. However, as we argued above, the effect of an increase in wages on work effort is ambiguous. In the diagram below, the the income effect outweighs the substitution, and so work effort actually decreases in response to a wage increase! Moreover, this implies that as the price of leisure increases, the amount of leisure consumed increases. I.e. the demand curve for leisure is upward sloping - leisure can be a Giffen good!

3.2.4 Deriving the Labour Supply Curve

In the above section, we found the optimal work level for two different wage levels. If we find the optimal labour choice for various different wage levels and plot these wage-effort combinations, we have the labour supply curve. (The process is identical to the process we used to derive the demand curve for an individual good, above.) I illustrate this in the diagram below:

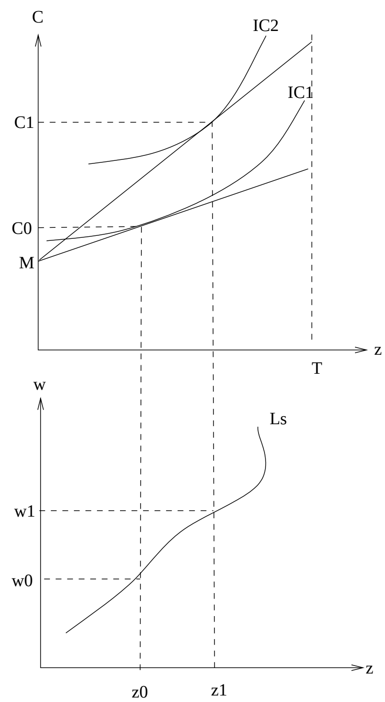

There is some empirical evidence to suggest that when the wage is low, the substitution effect tends to dominate the wealth effect, and so an increase in wages causes an increase in labour supply (i.e. L^s is upward sloping). But for high wages, the wealth effect dominates, and so a further increase in wages will actually cause labour supply to decrease. We say the labour supply curve is backward bending.

Example 17 (Mathematical Derivation of Labour Supply)**.** Consider a consumer with utility function $U(C,L) = \log C + \frac{1}{2}\log L$ and facing the budget constraint $pC + wL = M + wT$. Derive the consumption demand function and labour supply function.The consumer's problem is:

$$\begin{aligned} \max_{C,z} U(C,z) &= \log C + \frac{1}{2}\log L \\ \text{s.t. } pC + wL &= M + wT \end{aligned}$$

The Lagrangian is:

$$\mathcal{L} = \log C + \frac{1}{2}\log L - \lambda(pC + wL - M - wT)$$

The first order conditions are:

$$\begin{aligned} \frac{d\mathcal{L}}{dC} &= \frac{1}{C} - \lambda p = 0 \\ \frac{d\mathcal{L}}{dL} &= \frac{1}{2L} - \lambda w = 0 \\ \frac{d\mathcal{L}}{d\lambda} &= M + wT - pC - wL = 0 \end{aligned}$$

Combining the first two conditions gives:

$$\frac{1}{\lambda} = pC = 2wL$$

Then from the third condition (budget constraint), we have:

$$(2wL) + wL = M + wTL^* \qquad = \frac{M + wT}{2w} = \frac{T}{2} + \frac{M}{2w}$$

which provides an expression for the demand for leisure. Substituting back into the tangency condition gives:

$$C^* = \frac{M + wT}{p}$$

Finally, substituting into the time constraint gives the labour supply function:

$$z^* = \frac{T}{2} - \frac{M}{2w}$$

We notice several features:

- Consumption is increasing in unearned income M, wages w, and decreasing in consumption prices p.
- Labor supply is decreasing in unearned income M and increasing in the wage w (provided that $M > 0$).
- If there is no unearned income ($M = 0$), the agent works for half the discretionary time, regardless of the wage. (In this case, the 'income' and substitution effects exactly balance.

3.3 Vouchers and Foodstamps

Consider the following scenario - there is an economy with two goods: fruit and beer. The price of fruit is \$2 per kg, and the price of beer is \$5 per pint. Unemployed persons currently receive welfare from the government in the form of a \$50 weekly allowance. Suppose there are two types of welfare recipients - those who prefer fruit to beer, and those who prefer beer to fruit. We illustrate the consumption decision of each type of agent in the diagram below:

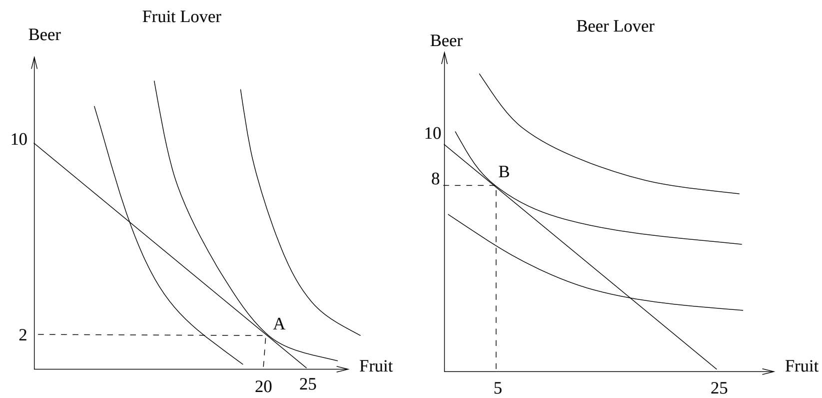

Since fruit lovers have a strong preference for fruit, they were willing to give up a lot of beer to receive an additional unit of fruit. Hence, fruit lovers have relatively steep indifference curves. This results in a tangency at point A, where the amount of fruit consumed is high and the amount of beer is low. By contrast, beer lovers are willing to give up very little beer to receive an additional unit of fruit. Hence, their indifference curves are relatively flat. This results in a tangency at point B, where the amount of fruit consumed is low and the amount of beer is high.

The government is concerned that the consumption decisions of the beer-lovers is a wasteful use of public funds. The government proposes a different welfare scheme whereby the unemployed receive a \$30 weekly allowance and 10 units worth of fruit vouchers. What is the effect of such a policy? There is a change in the budget constraint. (The decrease in cash income causes an inwards parallel shift in the budget line. In addition, consumers can now purchase 10 additional units of fruit, and so the budget constraint shifts 10 units to the right - but not upwards). We have:

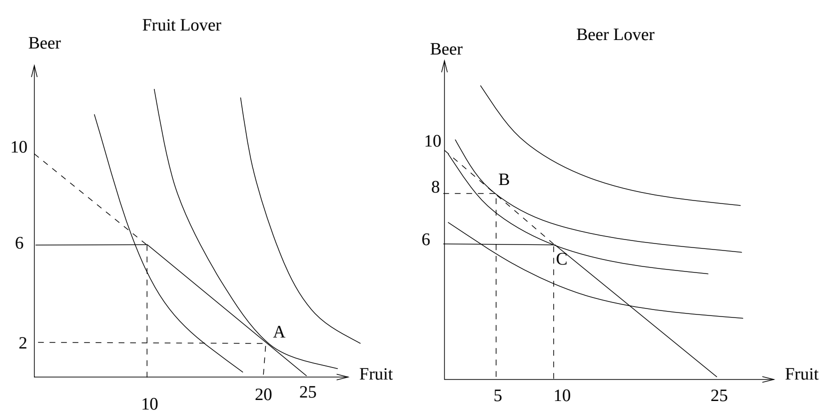

The new budget line has the same slope as the original budget line, but it has a kink at the point $(10, 6)$. The reason for this is as follows. Since the relative price of beer/fruit has not changed, the slope of the budget constraint is unchanged. Moreover, whilst his income has fallen by \$20, he also receives \$20 in fruit vouchers. Hence, the consumer can still purchase up to \$50 worth of fruit if he so chooses. The only difference is that with vouchers, the consumer can now purchase a maximum of only \$30 of beer. Once he has purchased \$30 (or 6 units) of beer, then he has exhausted all his cash income. All he can then do is to cash in some/all of his fruit vouchers.

From the diagram above, it should be clear that for fruit lovers, the introduction of vouchers is inconsequential. (Although it causes the budget constraint to change, these consumers were never going to purchase goods along the part of the budget constraint which is affected.) Fruit lovers will continue to purchase bundle A. By contrast, the introduction of vouchers does affect the behaviour of beer lovers. These consumers can no longer afford their preferred bundle (point B). Instead they maximise their utility by consuming at point C. This bundle contains the maximum amount of beer which they can afford. The policy is effective in the sense that it forces beer lovers to consume less beer and more fruit. Clearly bundle C is on a lower indifference curve than bundle B, and so beer lovers experience lower utility as a result of the government policy. (Obviously, one can argue that this might be a good policy, notwithstanding the disutility experienced by beer lovers.)

Important note —From the above example, it should be clear that restrictions/constraints on consumer choices can *never* increase utility. If the constraint is non-binding (i.e. if it is inconsequential to the agent), then obviously it has no effect on utility. However, if the constraint is binding, then it restricts the consumers choice set - and this has the effect of reducing utility.

Chapter 4

Welfare

In the previous chapter, we analyzed the effect of a change in prices or income on utility, i.e., welfare. However, we were only able to talk about these changes in welfare in a qualitative sense - a utility increase or decrease. In this chapter, we seek to quantify the effect of policy changes on welfare. By how much does welfare change?

How should we measure the magnitude of welfare changes? At first glance one might expect that we should simply calculate the amount by which utility changes. However, this approach has two problems:

1. The utility function is an ordinal function. This means it simply ranks bundles according to preferences —but says nothing about the strength of preferences. (To make this point stark, suppose $u(a) = 10$, $u(b) = 20$ and $u(c) = 30$. Then all we can say is that c is preferred to b which is in turn preferred to a. It does not follow that b is twice as preferred as a. Similarly, it does not follow that the increase in welfare in going from a to b, is the same as the increase in welfare in going from b to c.) As should be clear, the size of the change in utility conveys very little information.

2. There is no standard unit of measure for utility. For example. Consider two individuals with the same preferences. Suppose individual A has utility given by the function $u(x, y)$. Must individual B have the same utility function? No. e.g. suppose B's utility function is $v(x, y)$ where $v(x, y) = [u(x, y)]^3$. These utility functions both represent the same preferences, since whenever $u(a) > u(b)$, $v(a) > v(b)$. But then, if we measure the change in utility associated with a given policy, we will get different answers for the two individuals - even though they both have the same preferences!

We need another way of measuring the welfare effect of policy changes. A measure which avoids this problem is to use the consumer's own monetary valuation of the policy change. (i.e. how much would they pay to avoid the policy change? or how much money do they require as compensation for the policy change.) Since this measure is purely in terms of money, it does not suffer from the problems of comparison which we noted above.

4.1 Compensating and Equivalence Variations

We focus on the effect of changes in the price of a single good X (which we plot on the horizontal axis). For simplicity, we assume that the good on the vertical axis is a composite bundle of 'all other goods', whose price is normalised to 1. Hence, the consumer's choice involves how much of good X to purchase, and how much to spend (in dollar terms) on all other goods. This has the convenient property that the unit of measurement on the vertical axis is simply dollars.

4.1.1 Compensating Variation

Suppose the price of good X increases. The consumer originally chose bundle A on IC_0. After the price change, he chooses bundle B on IC_1. Clearly the consumer is worse off after the price increase. The **Compensating Variation** asks how much money we need give the consumer to compensate him for the loss in utility. What is this amount? It is just enough money, so that the consumer can afford a bundle which is on the original indifference curve IC_0, although not necessarily the one he previously had purchased. (This leaves him just as well off after the price change, as before it.)

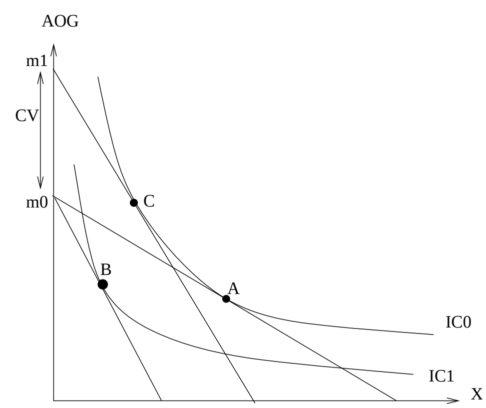

If we increase income from m_0 to m_1, then at the new price, the consumer can afford bundle C which lies on the original indifference curve IC_0. Notice that at point C, IC_0 is tangential to the new budget constraint. (This should be obvious, since the consumer will choose the new consumption bundle optimally, given his new level of income and the higher prices.) You should have noticed that the move from bundle A to bundle C, corresponds to the substitution effect. Since the price of good X is now higher, the consumer can achieve the same utility by consuming less of good X, and more of the composite bundle. The additional income we had to give the consumer $(m_1 - m_0)$ is referred to as the **compensating variation.**

The compensating variation is the minimum amount of income we must give (or the maximum amount we must take away from) the consumer to compensate them for an actual change in the price. Clearly, if the price of good X fell, then the consumer would be better off, so we would have to take income away from them, to bring them back to the original level of utility. In this case, the compensating variation is negative.

4.1.2 Equivalence Variation

Rather than imposing a price change and then compensating the consumer, we could ask the following question. What is the maximum amount the consumer be willing to pay to avoid the price increase altogether? If the price increase is implemented, the consumer's utility falls to IC_1. For each dollar the consumer pays to avoid the price change, his utility falls (since he has less income to spend on consumption). The consumer will be willing to pay to avoid the price change, so long as the utility he achieves with the income he has remaining is no lower than IC_1. Hence, the maximum amount he is willing to pay, is precisely the amount that would cause him to choose a bundle along IC_1. We refer to this amount as the **equivalence variation**.

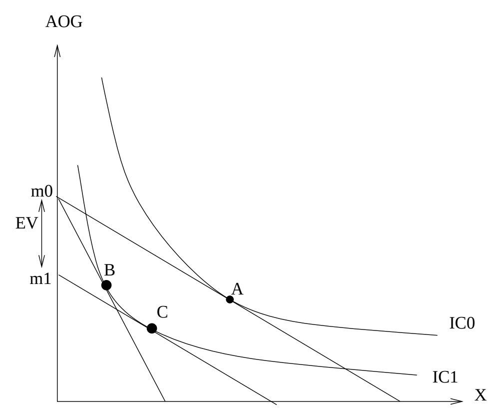

The equivalence variation is the maximum amount of income the consumer is willing to pay (or the minimum amount which he will demand), to avoid the price change from occurring. Clearly, if the price of good X fell, then the consumer would be better off, so he would demand payment in lieu of the utility increase he would enjoy if the price decrease actually went ahead.

Important note —the compensating variation is calculated assuming that the price change actually happens. So when we draw the budget lines, we draw them for the new prices. The equivalence variation, on the other hand, is calculated assuming that the price change does not happen. Hence, the budget lines are drawn using the old prices.

4.2 Applications

We have two different measures of welfare - CV and EV. Which should we use? The answer depends in part on the context we wish to analyse. Consider the following example:

The government is concerned that the price of fresh fruit and vegetables is too high and that this is affecting the eating habits of low/middle income earners. The government proposes to subsidise the cost of fresh fruits and vegetables, and to pay for the subsidy by levying a lump-sum income tax. Then, the appropriate measure of welfare would be the compensating variation. The compensating variation measures how much income the consumer is willing to give up in return for receiving the subsidy. If the lump-sum income tax is less than CV, the consumers are unambiguously better off. If the income tax is larger than CV, then consumers are worse off.

I illustrate these scenarios in the diagram below:

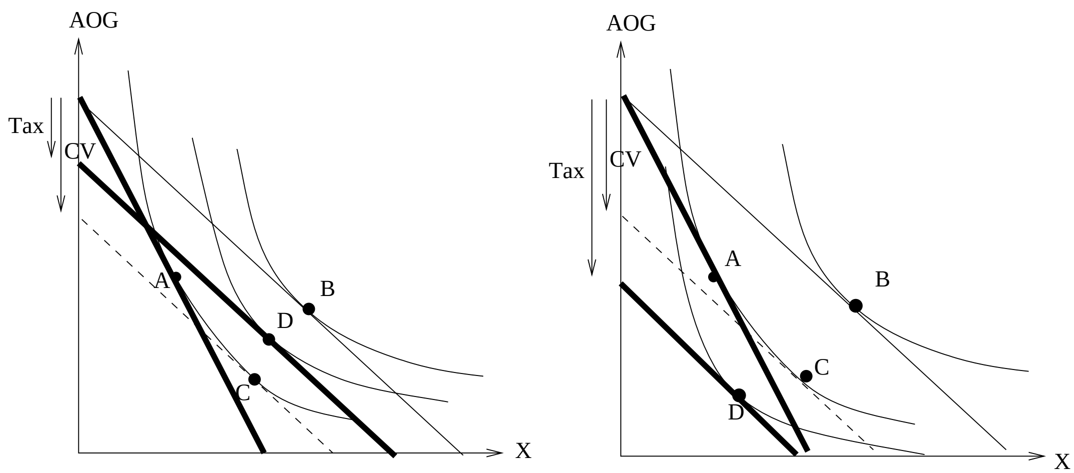

In the above diagram, the thick lines represent the actual budget constraints. The steep line is the budget constraint before the price decrease. The flatter line is the budget constraint after the change in price and income. (In the left hand panel, the magnitude of the income tax is small, relative to the right hand panel.) The thin line is the budget constraint if there was a price decrease (only) and no decrease in income. The dotted line is the budget constraint after the price change, and if just enough income is taken away from the consumer, so that they return to the original level of utility. In both diagrams, point A is the original consumption bundle; point B is the bundle that would've been chosen if there was only a price decrease, but not corresponding decrease in income; point C is the bundle that would be chosen at the new prices, if just enough income was taken away from the consumer to leave them as well off as they were prior to the price change; and point D is the actual new consumption bundle. Note that points A, B and C are identical in both panels - only point D differs. Clearly, if $tax > CV$ then the consumer is worse off, and consumes a bundle on a lower indifference curve. Conversely, if $tax < CV$, then the consumer is better off, and consumes a bundle on a higher indifference curve.

Suppose, on the other hand, the government decides to not decreases prices, but instead determines to supplement consumers' incomes, so that they can better afford fresh foods. Then, the appropriate measure of welfare would be the equivalence variation. The equivalence variation measures how much income the consumer would need to receive to make him as well off as he would've been had the subsidy been implemented. If the magnitude of the income supplement is larger than the EV, then clearly consumers are better off. Conversely, if the income supplement is smaller than EV, the consumer would've been better off with the subsidy.

4.2.1 Taxation/Subsidies and Excess Burden

We know that the government can affect consumer welfare by intervening in the market by providing price subsidies funded by lump-sum income taxes (or conversely by taxing goods and compensating consumers with income supplements). A natural question is this: can the government in such a way that consumers are better off AND there is no net cost to the government? If so, then there is a strong argument for government intervention in the market.

We define the **excess burden** as the net cost to the government of implementing some policy, such that consumer welfare is unchanged. (If the excess burden is negative, then the government can - by implementing some policy - raises a net revenue, whilst leaving consumer welfare unchanged.

Example: School Vouchers Consider the following example. The government wishes to increase participation in education. It can do this by subsidising the cost of education (and hence reduce the price from p_0 to p_1), with the subsidy being funded through a lump-sum income tax. Suppose the government levies the largest possible tax that does not make consumers worse off (relative to their utility prior to the implementation of the price subsidy and income tax) - this is the compensating variation. (Recall that since the subsidy makes consumers better off, the compensating variation is negative.) Consumers originally choose bundle A, and after the price and income changes, choose bundle B. The amount of revenue that the government raises through the lump-sum income tax is given by $CV = I_0 - I_1$. Note that the consumer's utility is unchanged.

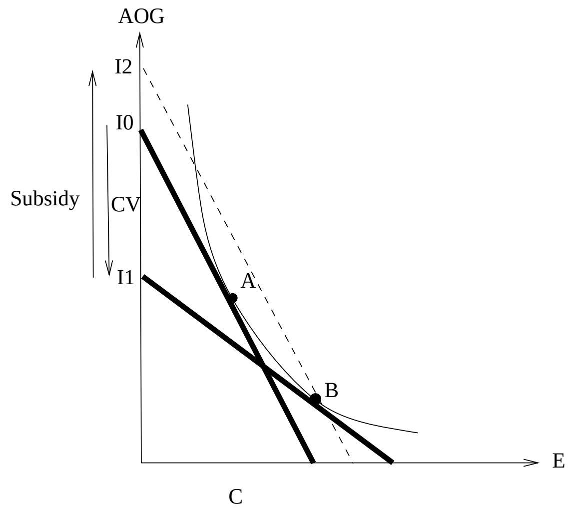

In the above diagram, the dashed line is parallel to the old budget line, but contains the new consumption bundle B. The cost of the subsidy to the government is given by $I_2 - I_1$. (This is the difference in intercepts of two budget lines that contain point B - one reflecting the old prices and one reflecting the new prices.) To see this formally, note that the formula for the new budget constraint is:

$$p_1 E_1 + y_1 = I_1$$

and the formula for the dotted line is:

$$p_0 E_1 + y_1 = I_2$$

where (E_1, y_1) represents the amount of education and all other goods after the policy change, and p_0 and p_1 represent the original and final prices of education. (Recall the price of all-other-goods is normalised to 1). Subtracting these two equations gives:

$$(p_0 - p_1) \times E_1 = I_2 - I_1$$

The left hand side of the expression is exactly the cost of the government subsidy. (The government provides the supplier with $(p_0 - p_1)$ for each unit produced, and the suppliers sell E_B units.) Hence, the cost of the subsidy can be represented by the distance between I_2 and I_1.

Note that the cost of the subsidy is more than the amount of revenue that the government can generate using income taxes. The excess burden of costs over revenues is given by $I_2 - I_0$. This illustrates a more general principle, that government policies that distort relative prices are less efficient than lump-sum transfers.

We could think of this problem slightly differently. Suppose the government, after implementing the subsidy, decided to raise taxes sufficiently to pay for the cost of the subsidy. Then, it must raise taxes by *more* than the compensating variation (since we have just shown that, if it raises taxes just by the level of CV, this will be insufficient to cover the entire cost of the subsidy). But doing so makes consumers worse-off than they were prior to the subsidy. Hence, it is not possible for the government to affect a policy through subsidies, which both maintains the welfare of consumers and keeps the government's budget in balance.

Chapter 5

Producer Theory

5.1 Technology - the Production function

In this chapter, we study decision making by firms. We identify firms by their production technology —which describes the relationship between the inputs consumed in the production process and the output generated. Consider a firm which produces a single output y using two inputs, (x_1, x_2). The **Production Function** describes the maximum amount of the output which can be produced, given a set of inputs, if they are used in their most efficient way. We have:

$$y = f(x_1, x_2)$$

The firm's production technology and production function are analogous to the consumer's preference relation and utility function. Both are taken as primitives —they describe the fundamental characteristics of firms and consumers, respectively. As in the Consumer Theory section, we seek to build a theory of firm decision making, taking its technology as given.

As with the chapter on consumer theory, we assume that quantities of inputs and output are perfectly divisible, so that firms can employ fractional units of inputs and produce fractional outputs. The **marginal product of input 1** is the amount of extra output that is produced when one additional unit of input 1 is added to the production process, holding constant all other inputs. We have

$$MP_1 = \frac{\partial f}{\partial x_1}$$

The marginal product of input 2 is similarly defined.

In general, we assume $MP_1 \geq 0$. (i.e. adding inputs to the production process usually does not decrease output. I note that in some cases, adding extra inputs may cause the process to become so congested and actually cause output to decrease - as the old adage goes: 'too many cooks spoil the broth'.) Furthermore, we assume that marginal product decreases as we add more units of an input to the production process, holding constant the level of all

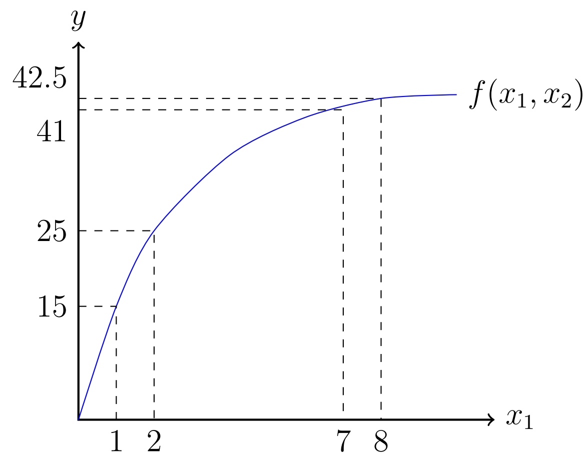

Figure 5.1: Law of Diminishing Marginal Returns

other inputs. (i.e. the additional contribution to output of each successive unit of an input is less than the previous unit.) This implies:

$$\frac{\partial MP_1}{\partial x_1} = \frac{\partial^2 f}{\partial x_1^2} < 0$$

We refer to this property as the **Law of Diminishing Marginal Returns**. Diminishing marginal returns implies that we cannot continue to increase output at a rapid rate by simply increasing the level of one input, ignoring the others. (For example, we cannot produce enough food to feed the world by simply adding more and more water to a single flower pot.) We illustrate the notion of diminishing returns in Figure 9.1. Notice that when the quantity of input x_1 used in the production process increases from 1 to 2, the increase in output is significantly larger than the corresponding increase in output when the quantity of input x_1 increases from 7 to 8. (The quantity of x_2 is held constant throughout.)

We also typically assume that the production function exhibits **decreasing returns to scale.** This means that if we increase the level of all inputs by some proportion then output will increase, but by a proportion less than the increase in inputs. I.e. if we double (or treble) the level of each input, output will increase, but will not double (or treble). Mathematically, we have:

$$f(\lambda x_1, \lambda x_2) < \lambda f(x_1, x_2) \qquad \text{for any } \lambda > 0$$

We refer top this property as **diseconomies of scale.** (Traditionally we assume that up to a particular point, when a firm increases its scale, it can produce goods more efficiently by employing more specialised inputs and exploiting efficiencies that arise when producing goods in bulk. If so, the technology exhibits *increasing returns to scale* or *economies of scale*. However, economies of scale cannot continue to persist forever. At some point, the size of a firm becomes so large, that production becomes efficient - perhaps due to the growth of sinusoids. Hence - eventually - we assume that the technology will exhibit diseconomies of scale.

Comments:

1. When considering the firm's production decisions, we distinguish between the short-run and the long-run. The **short-run** is defined as the period short enough that the firm cannot vary all of the inputs used in the production process. (e.g. the size of a factory assembly line cannot be changed over-night. In the short run, the size of this capital input is fixed.) The **long-run** is defined as the period long enough that the firm can vary all of its inputs. (e.g. continuing with the previous example, given sufficient time, the firm can increase the size of its assembly line to meet growing demand, or downsize a plant if needed.)

2. It is important to distinguish the concepts of *diminishing marginal product* and *diseconomies of scale*. Diminishing marginal product is a short run phenomenon. It says that if increasing quantities of a variable input are added to a fixed input, the marginal contribution of each additional variable input falls. Diseconomies of scale on the other hand, is a long run phenomenon. It says that even if we scale up all inputs in proportion, output wont increase by as much.

5.2 Cost Minimization

Our theory of the firm has two components. We must determine: (i) the quantity of output that the firm should produces, and (ii) the quantities of inputs that the firm should utilize in the production process. Ultimately we assume that firms are profit maximizers, and will produce the quantity that maximizes profit.

But there are multiple input combinations that produce this optimal output. Which combination should the firm choose? When contemplating producing *any* quantity of output y, we assume that firms choose the input combination that minimizes the cost of producing that output. To be clear —we are not saying that the firm's ultimate goal is to minimize costs; their goal is to maximize profits. Cost minimization is simply the criterion to choose amongst input combinations that achieve the same output. The cost minimizing input combination is the one that produces the desired output in the most efficient way. Suppose a firm wishes to produce y units of output. Let w_1 and w_2 be the input prices of labor and capital, respectively. An **isoquant** represents the set of all inputs which can be combined to produce the same level of output. (We can think of an isoquant as the producer theory analogue of an indifference curve.)

Properties of Isoquants

1. An isoquant must be downward sloping. (This follows from the assumption that $MP_i > 0$ for each input. Consider the diagram, below. Every input bundle in quadrant III must produce more output than a, since it consists of at least as much of each input

as a. Similarly, every input bundle in quadrant I must produce less output than a. Hence, bundles which produce the same level of output as a must lie in quadrants II and IV. But this implies that the isoquant is downward sloping.)

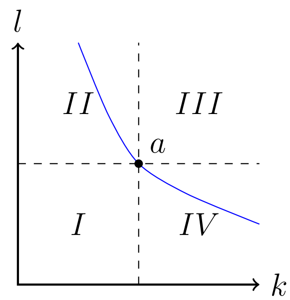

2. Higher isoquants represent input bundles which produce a higher level of output. (This again follows from the fact that $MP_i > 0$.)

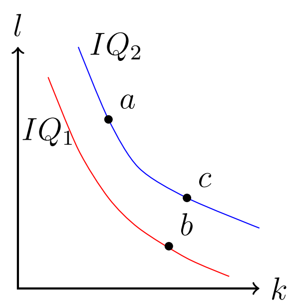

3. Isoquants cannot cross. (To see why, suppose two isoquants did indeed cross, and let a be the point at which they cross. Then a is a bundle which produces two distinct levels of output - and this obviously cannot be the case.)

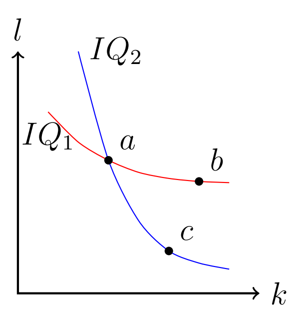

4. Isoquants are convex. (This follows from the assumption of diminishing marginal returns.)

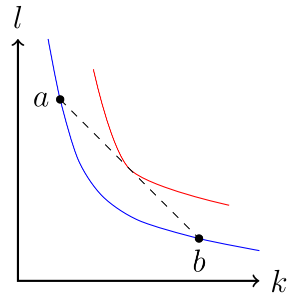

The slope of an isoquant is called the **marginal rate of technical substitution** (MRTS). (Again, the $MRTS$ is the producer theory analogue of the marginal rate of substitution.) $MRTS$ is the amount by which we must reduce input 2 if we increase input 1 by one unit, and wish to keep overall output unchanged.

Let $x_2(x_1)$ define an isoquant. For each x_1, it tells us how much x_2 is needed to achieve the desired output level y. Naturally, along this indifference curve, $f(x_1, x_2(x_1)) = y$.

Totally differentiate w.r.t. x_1:

$$\begin{aligned} \frac{\partial f}{\partial x_1} + \frac{\partial f}{\partial x_2}\frac{dx_2(x_1)}{dx_1} &= 0 \\ \frac{dx_2}{dx_1} &= -\frac{\partial f/\partial x_1}{\partial f/\partial x_2} \\ MRTS &= -\frac{MP_1}{MP_2} \end{aligned}$$

The marginal rate of technical substitution decreases as one moves long an isoquant. This follows from the law of diminishing marginal product. (To see why, suppose we are at the top of an isoquant, where there is a large amount of input 2 and very little of input 1. Then the additional contribution to output of each unit of x_2 must be small relative to the additional contribution of each unit of x_1. Hence, if we increase the amount of L by one unit, we can afford to take many units of x_2 out of production and keep output at the same level. Hence the isoquant is relatively steep. Now consider a point towards the bottom of the isoquant, where there is a large input of x_1 and very little input of x_2. Then, the additional contribution to output of a unit of x_1 will be very small, whereas the additional contribution to output of each unit of x_2 is relatively large by comparison. Hence, if we add one additional unit of x_1 to the production process, we can only afford to release a small amount of x_2 in order to keep the production level unchanged.)

An **isocost** line represents the set of all input bundles (x_1, x_2) which have the same cost. The isocost line is given by:

$$IC = \{(x_1, x_2) \,|\, w_1 x_1 + w_2 x_2 = c\}$$

The slope of the isocost function is given by *slope* $IC = -\frac{w_1}{w_2}$. The slope of the isocost line is the amount of input 2 that the firm must take out of the production process when it increases the amount of input 1 by 1 unit —if it wishes to keep total costs unchanged. Higher isocost lines represent bundles with higher costs.

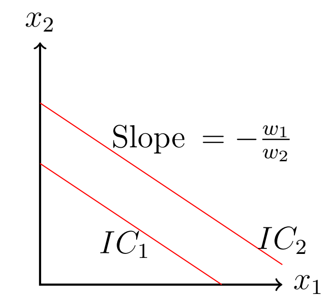

5.2.1 Optimisation

We want to find the input bundle which minimises the cost of producing quantity y of output. (This amounts to find the the bundle on the isoquant corresponding to output quantity y, which lies on the lowest achievable isoquant.) Formally, we have: $\min_{x_1,x_2} w_1x_1 + w_2x_2$ s.t. $y = f(x_1, x_2)$.

The Lagrangian is:

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

The first order conditions are:

$$\frac{\partial \mathcal{L}}{\partial x_1} = w_1 - \lambda \frac{}{x_1} = 0$$
$$\frac{\partial \mathcal{L}}{\partial x_2} = w_2 - \lambda \frac{}{x_2} = 0$$
$$\frac{\partial \mathcal{L}}{\partial \lambda} = y - f(x_1, x_2) = 0$$

Combining the first two conditions gives:

$$\lambda = \frac{w_1}{MP_1} = \frac{w_2}{MP_2} \tag{5.1}$$

Two important insights:

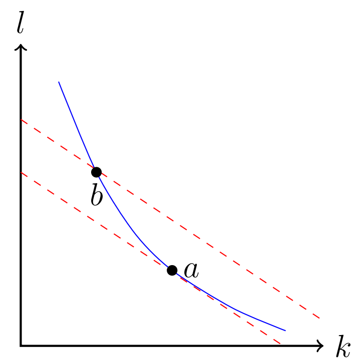

Figure 5.2: Optimal Input Choice

1. The optimal input bundle must lie on the isoquant which corresponds to the quantity y which we seek to produce.(Why? If the bundle lies on a different isoquant, then we are either producing more output than we need to, or not producing enough to satisfy our needs.)

2. The optimal consumption bundle will satisfy the tangency condition:

$$\begin{aligned} \textit{slope of IQ} &= \textit{slope of IC} \\ MRTS &= -\frac{MP_1}{MP_2} = -\frac{w_1}{w_2} \end{aligned}$$

   which follows from (9.1). (We can also prove this by contradiction. Suppose the optimal bundle is a, and at this point, the isoquant and isocost curves are not tangential. Then, at a, the curves must cross. But then, there is a point along the isoquant curve (b, say) which achieves a lower isocost line than a does. Since input bundles a and b both produce y units of output, and b is cheaper than a, then it cannot be optimal for a firm to choose input combination a. We have a contradiction.)

The tangency condition says that the rate at which producers are able to substitute between inputs and keep total output unchanged (i.e. $MRTS$) must be the same as the rate at which they are able to a substitute between inputs and keep total cost unchanged (i.e. slope of IC). If this is the case, then it is impossible for the firm to reallocate inputs in the production process in such a way as to reduce costs. Why? Let $MRTS = slopeIC = -2$. Suppose the firm tries to reallocate production by using one additional unit of x_1. Then to keep output unchanged, the firm must reduce the amount of x_2 utilised by 2 units. But since the slope of IC is -2, then if the firm increases x_1 by 1 unit and reduces x_2 by 2 units, then its total cost is unchanged. Hence, the firm cannot reduce costs by increasing the amount of input x_1 in the production process. By a similar argument, there is no gain to reducing the amount of x_1 in the production process.

By contrast, suppose in $MRTS = -3$ and $slopeIC = -2$, and the firm decides to reallocate inputs by using 1 additional unit of 1. This requires that he uses 3 fewer units of input 2. But for a unit increase in x_1, reducing the amount of x_2 by 2 units leaves costs unchanged -

so reducing the amount of x_2 by 3 units will cause costs to fall. Hence, it is possible for the firm to reallocate inputs in the production process in such a way as to reduce costs. But if so, the original input choice could not have been optimal. Hence, if $MRTS \neq slopeIC$, then there is always an opportunity for the firm to reallocate inputs in such a way as to reduce costs. It follows that, at the optimum, $MRTS = slopeIC$.

We note that the tangency condition can be restated as:

$$\frac{MP_1}{w_1} = \frac{MP_2}{w_2}$$

How do we interpret this? MP_1 is the additional output produced when 1 additional unit of x_1 is added to the production process. If the firm employs \$1 of x_1, then additional quantity of x_1 which it hires is $\frac{1}{w_1}$. Hence, the last dollar that the firm spent to employ input 1 generated $\frac{MP_1}{w_1}$ additional units of output. We interpret $\frac{MP_2}{w_2}$ in the same way. We require that for the last dollar spent on each input, the additional contribution to output be the same. Why? Suppose this was not the case. For concreteness, suppose $\frac{MP_1}{w_1} > \frac{MP_2}{w_2}$. Then the last dollar spent on labor contributed more to output than the last dollar spent on capital. If we take the last dollar spent on capital and instead spend it on labor, we lose $\frac{MP_1}{w_2}$ of output and gain $\frac{MP_1}{w_1}$ of output. Since $\frac{MP_1}{w_1} > \frac{MP_2}{w_2}$, total output has increased, whilst our costs have remained the same (since we are merely reallocating a dollar of expenses). Clearly, then, the original input choice could not have been optimal.

We can also restate the tangency condition as:

$$\frac{w_1}{MP_1} = \frac{w_2}{MP_2}$$

How do we interpret this? $\frac{1}{MP_1}$ is the amount of additional units of input 1 that needs to be hired in order to produce 1 additional unit of output. (If you unsure of this, consider the following numerical examples. Suppose $MP_1 = 3$. Then 1 additional unit of input 1 generates 3 additional units of output. Hence $\frac{1}{3}$ additional unit of input 1 will suffice to increase output by one unit. Similarly, suppose $MP_K = \frac{1}{2}$. Then 1 additional unit of x_2 generates $\frac{1}{2}$ a unit of additional output. Hence, 2 extra units of x_2 are needed to increase output by one unit.) Then $\frac{w}{MP_1}$ is the cost of producing one additional unit of output by increasing the amount of input 1 alone. The tangency condition written in this way says that the cost of producing one extra unit of output should be the same, regardless if we do this by using more of input 1, or more of input 2, or any combination of these inputs.

5.2.2 Comparative Statics

From the above discussion, it should be clear that the firm's optimal input bundle choice depends upon the input prices (w_1 and w_2) and the quantity of output being produced (y). How does the optimal input bundle change when there are changes in these factors?

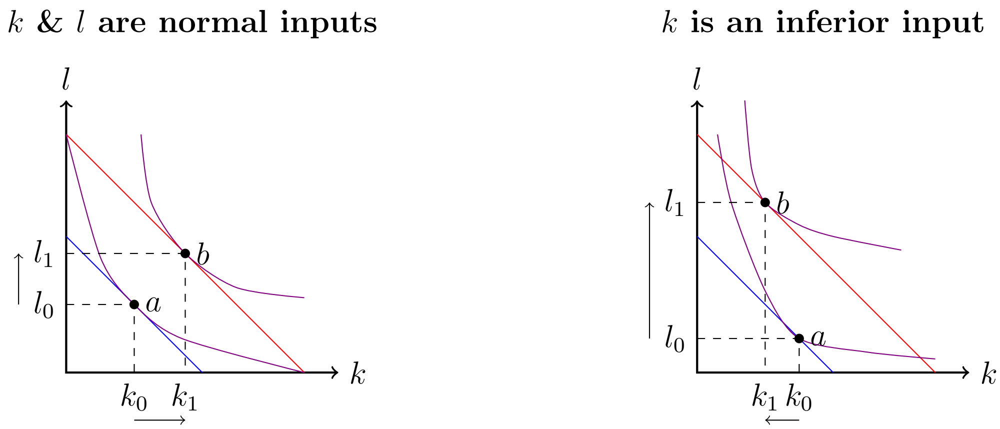

Figure 5.3: Effect of an increase in target output on conditional factor demands.

Changes in Quantity If the quantity being produced increases, then production must occur along a higher isoquant. In the diagram, below, we illustrate the effect of changes in y on the optimal input-bundle choice. The **expansion path** is the locus of optimal input bundles, traced out as the quantity of output required increases, but holding input prices constant. The expansion path describes how the optimal input bundle changes as the firm increases or decreases the scale of its production.

We note that when the quantity required increases, the amounts of inputs utilised do not necessarily increase in proportion. For example, when quantity increases from y_1 to y_2, the quantity of input x_2 utilised increases by more than the amount of input x_1. On the other hand, when the quantity increases from y_2 to y_3, the amount of input x_1 used increases, but the amount of input x_2 used decreases. We classify inputs in the following way:

- An input is *normal* if the quantity employed increases when output increases.
- An input is *inferior* if the quantity employed decreases when output increases.

Changes in input prices Next, consider the effect of an increase in the price of an input. For concreteness, suppose w_1 (the price of input x_1) increases. This causes the slope of the isocost cost curve to become steeper. In fact, the isocost line will become steeper. We illustrate the effect of an increase in input prices in the diagram below:

In the above diagram, the increase in the price of input x_1 causes the isocost lines to become steeper. As is apparent, an increase in the price of an input causes the amount of that input used in the production process to decrease. We refer to this as the **input substitution effect.** The input substitution effect must be negative. (To see why, note that an increase in the price of an input causes the isocost line to become steeper. But since the optimal input bundle must satisfy the tangency condition, then the optimum will occur at a steeper part of the isoquant. But since the isoquant is convex (i.e. since the slope of the isoquant decreases as we move along it), this implies that the new tangency point is at a point higher up the isoquant. Hence, the optimal bundle will consist of less of the input whose price increased, and more of the other inputs.)

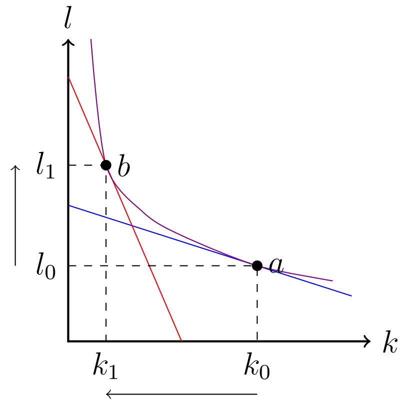

Figure 5.4: Effect of an increase in input price (r) on conditional factor demands.

5.2.3 Deriving the Conditional Factor Demand Functions

In the previous section, we illustrated the effect of an in the price of input x_1 on the optimal input bundle. If we repeat this procedure for various different input prices (holding constant the prices of other inputs and the quantity of output produced), we can find the relationship between input price and the quantity of that input We refer to this as the **conditional factor demand function**. (The factor demand is 'conditional', because it is defined for a particular quantity y. As we will show in a subsequent section, the optimal quantity is itself a function of factor prices.)

Suppose the price of input x_1 increases from w to w' to w'' (where $w < w' < w''$). This causes the level of input x_1 consumed in the production process to decrease from x_1 to x_1' to x_1''. We represent this situation diagrammatically:

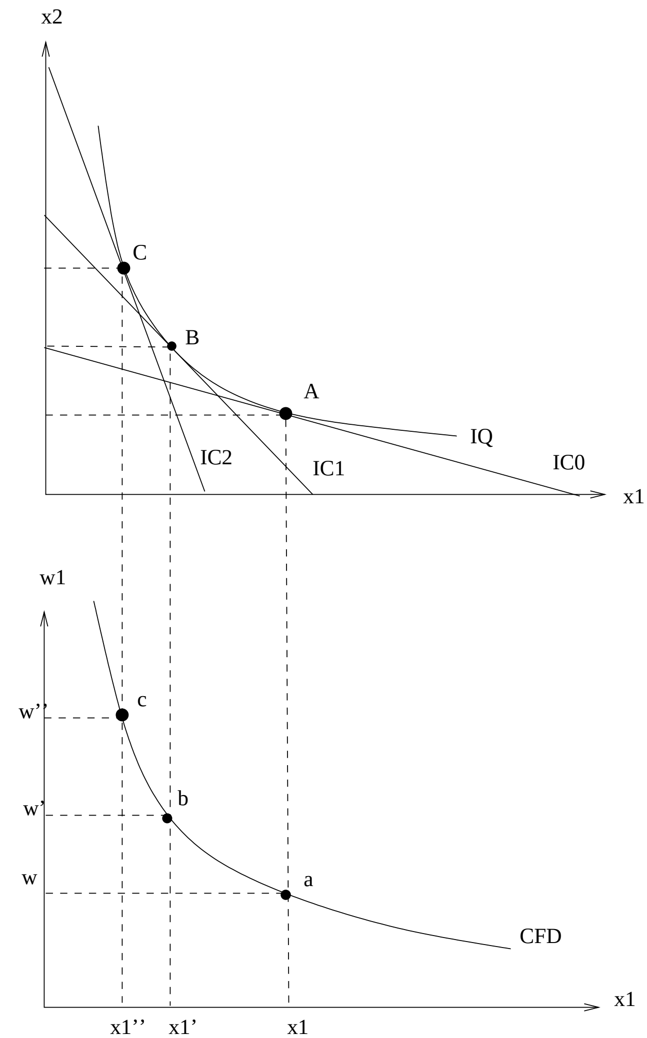

where IC_0, IC_1 and IC_2 are drawn for factor prices w_1, w_1' and w_1'' respectively. (Since $w_1' > w_1$, the slope of IC_1 must be steeper than the slope of IC_0.) The conditional factor demand function is drawn by plotting all the price-quantity combinations for factor x_1, holding constant the prices of other inputs and the quantity of output. Since the input substitution effect is negative, the conditional factor demand function must be downward sloping.

As with usual demand functions, changes in the price of the input cause movements along the conditional factor demand curve, whilst changes in other factors cause the entire factor demand curve to shift. An increase in the price of another input will cause the factor demand curve to shift to the right if the two inputs are substitutes in production, and it will cause the factor demand curve to shift to the left, if the two inputs are complements in production. Similarly, an increase in the quantity of output to be produced will cause the conditional factor demand to shift to the right, if the input is normal, and it will cause demand to shift to the left if the input is inferior.

5.2.4 Cost Function

The conditional factor demand functions give the input bundles which minimise the cost of producing a particular quantity of output y, given input prices w. The **cost function** specifies this minimised cost, for each potential output quantity.

Given the conditional factor demand functions ($\hat{x}_1$ and $\hat{x}_2$) it is easy to calculate the minimised cost. We have:

$$C(w_1, w_2, y) = w_1\hat{x}_1 + w_2\hat{x}_2$$

The **marginal cost**, is the cost of producing one extra unit of output. Formally, we have:

$$MC = \frac{\partial C(w, y)}{\partial y}$$

We noted above, that the cost of producing one extra unit of output is the same whether we do so by increasing the amount of input 1 only, or the amount of input 2 only, or some combination of the two. Hence, we have:

$$MC = \frac{w_1}{MP_1} = \frac{w_2}{MP_2}$$

Properties of the Cost and Marginal Cost Functions

1. $C(w, y) \geq 0$. (This follows straightforwardly from the fact that factor prices and conditional factor demands cannot be negative)

2. $MC = \frac{\partial C}{\partial y} > 0$ which implies the cost function is upward sloping in y. (This follows from the fact that MC is the ratio of input prices and marginal products, both of which are non-negative. This result should be relatively intuitive. The cost function gives the minimum cost of producing a particular level of output when inputs are combined in the least cost way. In order to increase the level of output produced, we must use more inputs in the production process - and this will clearly increase the cost of production.)

3. $\frac{\partial MC}{\partial y} = \frac{\partial^2 C}{\partial y^2} > 0$. (i.e the MC curve is upward sloping) This implies that the cost of producing additional output increases at an increasing rate. (Why? We know that MC is inversely proportional to marginal product, and that MP decreases by virtue of the law of diminishing marginal returns. Since adding additional units of an input to the production process has a increasing lower effect on output, we must add increasingly more inputs to the production process to increase output, as the level of existing output becomes large. But then the cost of increasing output becomes increasingly expensive.) Note —to the extent that the production function may exhibit increasing returns when the level of input is low, the marginal cost may be downward sloping for low levels of output. But this effect will disappear, as soon as diminishing returns begins to take effect).

We illustrate the shape of the marginal cost curve in the diagram below. In addition to marginal cost, we define the **average cost** of production:

$$AC = \frac{C(w, y)}{y}$$

With a little reflection, it should be apparent that the average cost and marginal cost functions are intimately related. In particular, suppose $MC < AC$. This implies that the cost of producing one additional unit of output (MC) is less than the average cost of producing all the previous units of output. Then, if this additional unit is produced, it will cause the average cost of production to fall (since the final unit cost less than the previous ones, on average). Hence, if $MC < AC$, then the average cost will tend to fall. Similarly, suppose $MC > AC$. Then the cost of producing one additional unit of output is more than the average of the costs of producing the previous units. Then, producing this additional unit causes the average cost to rise. Finally, if $MC = AC$, then cost of producing an additional unit of output is the same as the average cost of producing all the previous units. Producing this additional unit does not cause the average to change.

We verify this relationship mathematically, by noting that:

$$\begin{aligned}
\frac{\partial AC}{\partial y} &= \frac{y\frac{\partial C}{\partial y} - C}{y^2} \\
&= \frac{\frac{\partial C}{\partial y} - \frac{C}{y}}{y} \\
&= \frac{MC - AC}{y}
\end{aligned}$$

(Clearly if $MC > AC$, then $\frac{\partial AC}{\partial y} > 0$ and so AC is increasing and so on.)

It follows that:

- If MC is above AC, then AC must be increasing (upward sloping)
- If MC is below AC, then AC must be decreasing (downward sloping).
- If MC and AC intersect, then AC must be at a critical point. We can verify that this will be a minimum point of the average cost function.

Using these facts, we plot the average and marginal cost functions:

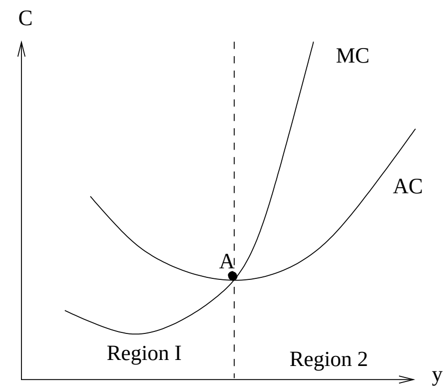

Consider the above diagram. In region 1, $MC > AC$ and so AC is decreasing. By contrast, in region 2, since $MC < AC$, it follows that AC must be increasing. AC and MC intersect at point A. At this point AC is neither increasing nor decreasing. This is represents the turning point of the AC curve - it is the level of output where average costs are minimised.

5.3 Profit Maximization

Having derived the cost function, we can now analyse the firm's profit maximisation problem. Suppose the price of output is p. Then, the firm chooses the level of output y that maximises its profit:

$$\pi = py - C(w_1, w_2, y)$$

The optimal level of output y^* must satisfy the first order condition:

$$p = MC(w_1, w_2, y^*)$$

(Why must this be the case? Suppose y is profit maximising, but at y, $p \neq MC(y)$. For concreteness, suppose $p > MC$. Then, if the firm produces one additional unit of output, it will receive an additional revenue of p and incur an additional cost of MC. Since $p > MC$, producing this additional unit has the net effect of increasing the firm's profit. But if so, then the original output y could not have been profit maximising. Similarly, suppose $p < MC$. Then, the last unit of output produced generated additional revenue of p, but incurred an additional cost of MC. Since $p < MC$, the last unit produced generated a net loss for the firm. The firm could increase it's profits by decreasing it's output. But if so, then the original output y could not have been profit maximising. Hence, at the optimum, we must have $p = MC$.)

We illustrate this situation diagrammatically:

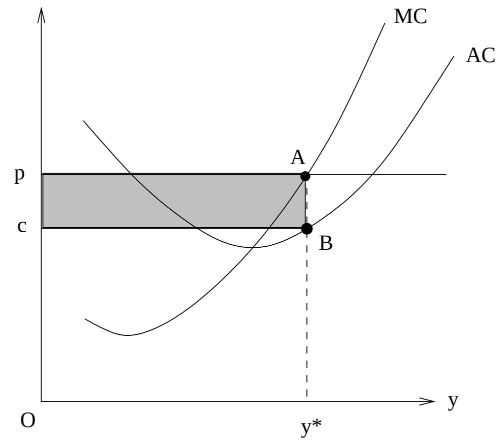

The MC curve intersects the price line at point A. This implies that the profit maximising level of output is y^*. At this level of output, the firm's revenue is py^* which is indicated in the diagram by the area $OPAy^*$. Furthermore, for this level of output, the average cost of production is $AC(y^*) = c$. Then, the total cost of production is given by cy^* which is indicated by the area $OcBy^*$. The firm's profit is its revenue less its cost - this is indicated by the shaded area $cPAB$.

5.4 Firm Supply

The firm's **supply function** $y^*(p, w_1, w_2)$ gives the profit maximising level of output as a function of the output and input prices. The supply function has the following properties:

- The supply function must be upward sloping $\left(\frac{\partial y^*}{\partial p} > 0\right)$. (Why? We know that the marginal cost function is increasing in output y. If the output price p increases, then the profit maximisation condition requires MC to be higher. But this requires an increase in y. More formally, by totalling differentiating the profit maximisation condition with respect to p, we note that;

$$\begin{aligned} p &= MC(w_1, w_2, y^*) \\ 1 &= \frac{\partial MC}{\partial y^*} \cdot \frac{\partial y^*}{\partial p} \\ \frac{\partial y^*}{\partial p} &= \frac{1}{\frac{\partial MC}{\partial y^*}} \end{aligned}$$

  Then, since $\frac{\partial MC}{\partial y^*} > 0$, it follows that $\frac{\partial y^*}{\partial p} > 0$.

- If an input is normal (in the sense that when output increases, more of the input is consumed in the production process), then an increase in the factor price will cause supply to decrease. (i.e. $\frac{\partial y^*}{\partial w_i} < 0$ if w_i is normal) An increase in the price of a normal input causes the supply curve to shift to the left.

- If an input is inferior (in the sense that when output increases, less of the input is consumed in the production process), then an increase in the factor price will cause supply to increase (i.e. $\frac{\partial y^*}{\partial w_i} > 0$ if w_i is inferior). An increase in the price of an inferior input causes the supply curve to shift to the right.

Example 18. Suppose the production function is given by $y = 2x_1^{0.5}x_2^{0.25}$. Given generic output and factor prices (p, w_1, w_2), find the firm's supply function. Verify that it is increasing in the output price p. What is the effect of an increase in factor prices w?

From the previous example, we know that the cost function is given by $C(w_1, w_2, y) = \frac{3}{4}w_1^{\frac{2}{3}}w_2^{\frac{1}{3}}y^{\frac{4}{3}}$. Then, the marginal cost function is:

$$\begin{aligned} MC(w_1, w_2, y) &= \frac{\partial C}{\partial y} \\ &= w_1^{\frac{2}{3}}w_2^{\frac{1}{3}}y^{\frac{1}{3}} \end{aligned}$$

Profit maximisation requires:

$$\begin{aligned} p &= MC(w_1, w_2, y) \\ p &= w_1^{\frac{2}{3}}w_2^{\frac{1}{3}}y^{\frac{1}{3}} \\ y^{\frac{1}{3}} &= pw_1^{-\frac{2}{3}}w_2^{-\frac{1}{3}} \\ y^* &= \frac{p^3}{w_1^2 w_2} \end{aligned}$$

Clearly, the supply function is increasing in p and decreasing in w_1 and w_2. It follows that x_1 and x_2 must both be normal inputs. (We confirmed that this was the case in the previous example.)

5.5 Factor Demand

The profit maximisation condition also informs the firm's decision about what combination of inputs to use in the production process. Recall in the section on cost minimisation, we showed that the marginal cost could be expressed as:

$$MC = \frac{w_i}{MP_i}$$

where w_i is the price of input x_i and $MP_i = \frac{\partial f}{\partial x_i}$ is the marginal product of that factor. Substituting this expression into the profit maximisation condition, we get:

$$p = \frac{w_i}{MP_i}$$

which implies both of the following:

$$p \cdot MP_i = w_i \tag{5.2}$$

$$MP_i = \frac{w_i}{p} \tag{5.3}$$

What do these conditions mean? Suppose the firm hires one additional unit of input i. Then, its costs increase by w_i. Furthermore, by hiring the additional input, the firm produces MP_i additional units of output, which generates total additional revenue of $p \cdot MP_i$. (We refer to this latter quantity as the *marginal revenue product.*) Condition (5.2) states that the additional revenue gained from hiring another unit of any given input must equal the cost of hiring that input. (Why? Suppose not. For concreteness, suppose $p \cdot MP_i > w_i$. Then the additional revenue gained by hiring one extra unit of input i exceeds the cost of hiring that input. The firm could increase its profit by hiring more of input i. But if so, then the original input decision could not have been optimal. The opposite is true if $p \cdot MP_i < w_i$. (The firm could increase its profit by hiring less of input i.) Hence, at the optimum, we must have $p \cdot MP_i = w_i$, for each input i)

Condition (5.3) restates this condition in real (rather than monetary) terms. The right hand side of (5.3) $\frac{w_i}{p}$ is the *real wage* that the firm must pay input i. The real wage is the number of units of output which the owner of the input can afford to purchase with the wage he receives for contributing one unit of the input. (e.g. if my hourly wage is $\$5/hr$ and the output cost is \$2, then my real wage is the equivalent of 2.5 units of output.) Profit maximisation requires that at the optimum, the real wage paid to each input is equal to the marginal product of that input. (Why? Suppose this was not the case. For concreteness that $MP_i > \frac{w_i}{p}$. Then, hiring one additional unit of input i generates MP_i units of additional output for the firm. The firm need only pay the input a wage equivalent to $\frac{w_i}{p}$ units of output. Hence, by hiring one additional unit of output, the firm can increase it's production at no net cost. But if so, the original input bundle could not have been optimal. The opposite argument holds if $MP_i < \frac{w_i}{p}$.)

We can use conditions (5.2) or (5.3) to derive the firm's **factor demand** functions $x^* (p, w_1, w_2)$. Consider the following diagram:

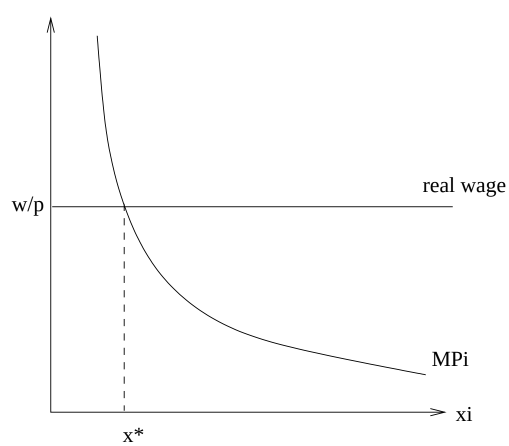

The factor demand functions have the following properties:

- The factor demand function for input i is decreasing in the factor price w_i. (i.e. $\frac{\partial x_i^*}{\partial w_i} < 0$). (To see why, keeping the output price constant, an increase in the input price w_i causes the real wage $\frac{w_i}{p}$ to increase. Since the marginal productivity must equal the real wage at the optimum, this requires MP_i to increase. But this implies the quantity of input employed must decrease, since MP_i is decreasing in the level of input employed. More formally, taking the total derivative of (5.3) with respect to w_i, we have:

$$\begin{aligned} \frac{\partial MP_i}{\partial x_i^*}\frac{\partial x_i^*}{\partial w_i} &= \frac{1}{p} \\ \frac{\partial x_i^*}{\partial w_i} &= \frac{1}{p\frac{\partial MP_i}{\partial x_i^*}} < 0 \end{aligned}$$

  since $\frac{\partial MP_i}{\partial x_i^*} < 0$ by diminishing marginal returns.)

- If an input is normal, then an increase in the output price p will cause the demand for that input to increase. (i.e. $\frac{\partial x_i^*}{\partial p} > 0$ if input i is normal. To see why, note that an increase in p causes optimal output y^* to increase. But since input i is normal, this requires more of x_i to be contributed to the production process). An increase in p causes the factor demand curve to shift to the right.

- If an input is inferior, then an increase in the output price p will cause the demand for that input to decrease. (i.e. $\frac{\partial x_i^*}{\partial p} < 0$ if input i is inferior. To see why, note that an increase in p causes optimal output y^* to increase. But since input i is inferior, this requires less of x_i to be contributed to the production process). An increase in p causes the factor demand curve to shift to the left.

Example 19. We calculate the factor demand function for the firm with production function given in the above example: $y = 2x_1^{0.5}x_2^{0.25}$. Recall, the marginal product functions are given by:

$$\begin{aligned} MP_1 &= x_1^{-0.5}x_2^{0.25} \\ MP_2 &= \frac{1}{2}x_1^{0.5}x_2^{-0.75} \end{aligned}$$

Applying condition (5.3), we have the following two equations:

$$x_1^{-0.5}x_2^{0.25} = \frac{w_1}{p} \tag{5.4}$$

$$\frac{1}{2}x_1^{0.5}x_2^{-0.75} = \frac{w_2}{p} \tag{5.5}$$

Rearranging (9.2) gives:

$$x_2 = \left(\frac{w_1}{p}\right)^4 x_1^2 \tag{5.6}$$

Substituting this into (9.3), we have:

$$\begin{aligned}\frac{1}{2}x_1^{0.5}\left[\left(\frac{w_1}{p}\right)^4 x_1^2\right]^{-0.75} &= \frac{w_2}{p}\\ \frac{1}{2}x_1^{-1}\left(\frac{w_1}{p}\right)^{-3} &= \frac{w_2}{p}\\ x_1^* &= \frac{p^4}{2w_1^3w_2}\end{aligned}$$

Then, by (9.4), we have:

$$\begin{aligned}x_2^* &= \left(\frac{w_1}{p}\right)^4 x_1^2\\ &= \left(\frac{w_1}{p}\right)^4\left[\frac{p^4}{2w_1^3w_2}\right]^2\\ &= \frac{p^4}{4w_1^2w_2^2}\end{aligned}$$

We confirm that the factor demand functions are decreasing in their own factor prices, and increasing in the output price p (since - as we showed previously - these inputs are normal).

In the section on cost minimisation, we derived the conditional factor demand functions $\hat{x}_i(w_1, w_2, y)$. How are these distinct from the factor demand functions $x_i^*(w_1, w_2, p)$? The conditional factor demand functions gave the input bundles that a firm would employ if it's objective was to produce some arbitrary level of output y, in the least cost way. This output level need not be the profit-maximising level - indeed, we do not take output prices into account when constructing the conditional factor demand functions. The conditional factor demand functions are the producer theory analogue of the compensated (Hicksian) demand curves. Recall the the compensated demand functions gave the optimal consumption bundles that a consumer would choose, if his goal was to achieve some arbitrary level of utility, and if affordability wasn't a consideration. The conditional factor demand functions give the optimal input bundle that a firm would choose, if its goal was to achieve some arbitrary level of output, and if profitability wasn't a consideration. To this extent, conditional factor demand functions - like the compensated demand functions - are a theoretical construct. What we actually observe are the factor demand functions.

Note - we can calculate the factor demand functions using the conditional factor demand functions and the supply function. Intuitively, we should have:

$$x_i^*(w_1, w_2, p) = \hat{x}_i(w_1, w_2, y^*(w_1, w_2, p))$$

I.e. The factor demand functions are simply the conditional factor demand functions, evaluated at the optimal output level (which is given by the supply function). We illustrate this numerically:

Again consider the production technology in the example above. We showed in a previous example that the supply function was given by:

$$y^*(w_1, w_2, p) = \frac{p^3}{w_1^2 w_2}$$

and the conditional factor demand functions were given by:

$$\begin{aligned}
\hat{x}_1(w_1, w_2, y) &= \frac{1}{2}\left(\frac{w2}{w_1}\right)^{\frac{1}{3}} y^{\frac{4}{3}} \\
\hat{x}_2(w_1, w_2, y) &= \left(\frac{1}{2}\right)^2 \left(\frac{w_1}{w_2}\right)^{\frac{2}{3}} y^{\frac{4}{3}}
\end{aligned}$$

It follows that the factor demand functions are given by:

$$\begin{aligned}
x_1^*(w_1, w_2, p) &= \frac{1}{2}\left(\frac{w2}{w_1}\right)^{\frac{1}{3}} \left[\frac{p^3}{w_1^2 w_2}\right]^{\frac{4}{3}} \\
&= \frac{p^4}{2w_1^3 w_2} \\
x_2^*(w_1, w_2, p) &= \left(\frac{1}{2}\right)^2 \left(\frac{w_1}{w_2}\right)^{\frac{2}{3}} \left[\frac{p^3}{w_1^2 w_2}\right]^{\frac{4}{3}} \\
&= \frac{p^4}{4w_1^2 w_2^2}
\end{aligned}$$

which correspond to the factor demand functions which we calculated in the previous example.

Chapter 6

Choice Under Uncertainty

6.1 Expected Utility

6.1.1 Motivation

In the previous models, the agent was required to make choices over known outcomes. However, in reality, we often have to make choices whose outcomes are uncertain. For example —suppose I am running late for an interview. Should I speed or not? If I speed, then, there is some chance that I will get caught, and have to pay a fine (and will be even later). However there is also a chance that I will not get caught, and I will arrive at the interview on time. On the other hand, if I do not speed, I will definitely arrive late. But this could either be a problem (if my lateness is noticed), or not a problem (if the interview panel is itself running over time). What should I do?

In this chapter, we develop a theory of how agents make choices when the outcomes are uncertain. In a model with uncertainty, agents are no longer making choices over the set of available consumption bundles. Instead, they choose amongst the set of available **lotteries.**

A *lottery* P is an uncertain prospect, which specifies outcomes and their associated probabilities. More formally, let $x = (x_1, ..., x_n)$ be a set of mutually exclusive and exhaustive outcomes. (I.e. we know that when the uncertainty is resolved, the agent will receive exactly one of the prizes from amongst the set $\{x_1, ..., x_n\}$.) A lottery P, assigns to each outcome x_i, a probability p_i of that outcome occurring. The assignment of probabilities must satisfy the usual assumptions on probability measures. I.e.

1. $0 \leq p_i \leq 1 \ \forall i$
2. $\sum_{i=1}^{n} p_i = 1$

The following are examples of lotteries:

- P_1 : (betting on number 23 in roulette) Win \$175 with probability $\frac{1}{37}$ and lose \$5 with probability $\frac{36}{37}$.
- P_2 : (when discussing politics with my father) Get into a heated argument with probability 0.71 or talk amiably with probability 0.29.
- P_3 : (when buying an employment insurance policy) Receive a guaranteed \$50,000 level of income with probability 1 (i.e. regardless of whether I lose my job or not)

Comments

- The probability of outcomes in a lottery can be either objective (as in the case of P_1) or subjective (as in the case of P_2).
- P_3 is a special type of lottery - in which the outcome is certain. We refer to this as a *degenerate lottery.*
- In this course, we will focus on lotteries which have monetary prizes (outcomes).

The *Expected Value* $E[X]$ of a lottery is given by:

$$E[X] = \sum_{i=1}^{n} p_i \cdot x_i = p_1 x_1 + ... + p_n x_n$$

We can think of the expected value of a lottery as the average payoff which an agent would receive if he played the lottery repeatedly.

Example 20. Consider the following gamble. I roll a fair die. If I roll a 6, then I win \$11. Otherwise, I lose \$1. The outcomes are $\{-1, 11\}$ and the probabilities of these events are $\Pr(-1) = \frac{5}{6}$ and $\Pr(11) = \frac{1}{6}$. The expected value of this lottery is:

$$E[X] = \frac{5}{6}(-1) + \frac{1}{6}(11) = 1$$

Hence, if I play this game repeatedly, then on average I will win \$1 —although in any given round, I will either win \$11, or lose \$1.

We classify lotteries in the following way:

- A lottery is said to be a **fair bet** if its expected value is 0.
- A lottery is said to be a **favourable bet** if its expected value is positive.
- A lottery is said to be an **unfavourable bet** if its expected value is negative.

The lottery in example 20 is a favourable bet since the expected value is positive. By contrast, lottery P_1 above is an unfavourable bet. To see this, note that:

$$E[P_1] = \frac{1}{37}(175) + \frac{36}{37}(-5) = -0.135$$

Agents will accept all favourable bets and reject all unfavourable bets. They will be indifferent accepting and rejecting fair bets.

Let us test this conjecture. Consider the following two gambles:

1. A fair coin is tossed. If the coin shows heads, you win \$100. If the coin shows tails, you lose \$100.

2. A fair coin is tossed. If the coin shows heads, you win \$105. If the coin shows tails, you lose \$100.

Which of these gambles would you be willing to accept? The first gamble is a fair bet - and so according to the conjecture, agents should be indifferent between accepting and rejecting this gamble. If so, we should observe some people choosing to accept the gamble and others choosing to reject. In reality, almost every person would reject such a gamble! The second gamble is a favourable bet. Then according to our conjecture, agents should always accept this gamble. In reality, most people would reject this gamble, although a small fraction do accept.

Our conjecture doesn't seem to quite work. To make the argument even more stark, consider the following famous example:

Example 21 (St. Petersburg Paradox)**.** A fair coin is tossed repeatedly until the first heads comes up. If heads comes up on the first toss, you receive \$2. If heads comes up on the second toss, you receive \$4. If heads comes up on the n^{th} toss, you receive $\$2^n$. What is the expected value of this game? Would you be willing to pay \$5 to participate in this game? What about \$20? What is the maximum amount you would be willing to pay to participate in such a game?

First, let us calculate the probability of each outcome. The probability that heads comes up on the first toss is $\frac{1}{2}$. The probability that heads first comes up on the second toss is $\frac{1}{4}$ (Why? This requires that tails came up on the first toss, and heads came up on the second toss. Each of these events has probability $\frac{1}{2}$ and so the joint probability is $\frac{1}{4}$.) The probability that heads first comes up on the n^{th} toss is $\left(\frac{1}{2}\right)^n$. Then the expected value of the lottery is:

$$\begin{aligned} E[X] &= \frac{1}{2}(\$2) + \frac{1}{4}(\$4) + \frac{1}{8} + ... + \frac{1}{2^n}(\$2^n) + ... \\ &= 1 + 1 + 1 + ... + 1 + ... \\ &= \infty \end{aligned}$$

Hence, if we play this game repeatedly, on average our payoff will be infinite. If so, we should be willing to pay ANY amount to take part in this game (since whatever we pay, the expected payoff will still be greater). And yet some people would not even pay \$5 to participate in the game, and most people would not pay \$20.

6.1.2 von Neumann-Morgenstern Expected Utility

Clearly our conjecture doesn't work. What is happening? Our conjecture assumes that agents are concerned with dollar amounts of the prizes they receive - and that they will choose the lottery which has the highest expected dollar prize. However, it is clear that agents do not behave in this way. Rather, agents are concerned with the utility they derive when they receive various prizes. They will choose the lottery which has the highest expected utility.

Example 22. Suppose has \$100, and with this money he can feed himself for a week. The agent is offered a gamble in which he wins \$10 with probability $\frac{1}{2}$ and loses \$10 with probability $\frac{1}{2}$. If he wins \$10, then his total wealth is \$110, and now he can afford to have a glass of wine with his dinner. If he loses \$10, then his total wealth is \$90, and he can no longer afford dinner every night. (He must go hungry for one night.) Although the monetary value of his winnings/losses are the same (\$10), the utility loss associated with forgoing one dinner, must be larger than the utility gain associated with having wine with dinner. Hence, the agent will reject the gamble.

Let P be a lottery with prizes $\{x_1, ..., x_n\}$ which are received with probabilities $\{p_1, ..., p_n\}$. Let $u(x_i)$ describe the utility which the agent derives upon receiving prize x_i. We refer to $u(x_i)$ as the utility index over prizes. Then, the agent's preference for the lottery is given by the utility function:

$$U(X) = \sum_{i=1}^{n} p_i u(x_i) = p_1 u(x_1) + ... + p_n u(x_n)$$

$U(P)$ is referred to as the **expected utility function**, or the **von Neumann-Morgernstern** utility function. (It is name for John von Neumann and Oskar Morgenstern who pioneered the research into expected utility theory.)

We need to be careful to distinguish two different notions of utility. The utility index $u(\cdot)$ refers to the utility which the agent derives when he receives a particular prize. But agents do not choose amongst prizes. Rather - they choose amongst lotteries, each of which specifies the probability with which they will receive a given prize. The utility associated with the lottery is given by $U(\cdot)$. (Note - at the end of the day - when the uncertainty is resolved, then the agent will actually receive a single prize, and his utility will be given by the function $u(\cdot)$. However, he must make his choice over lotteries *before* knowing which prize he will receive. Hence, he chooses the lottery which maximises his utility, in expectation.)

Example 23. An agent with initial wealth of \$100 is faced with the following lotteries:

- P_1 : receive \$20 with probability $\frac{1}{3}$ and lose \$10 with probability $\frac{2}{3}$.
- P_2 : receive \$25 with probability $\frac{1}{6}$ and lose \$5 with probability $\frac{5}{6}$

The agent's utility index over his wealth level is given by $u(w) = \sqrt{w}$. Which lottery should he choose?

We calculate the expected utility for each lottery. We have:

$$\begin{aligned} U(P_1) &= \frac{1}{3}\sqrt{120} + \frac{2}{3}\sqrt{90} = 9.976 \\ U(P_2) &= \frac{1}{6}\sqrt{125} + \frac{5}{6}\sqrt{95} = 9.986 \end{aligned}$$

Hence, the agent will choose the second lottery, since this gives him a higher expected utility.

Two features of the above example are worth noting.

- The expected value of both lotteries is the same. (They are both fair bets, and so they both have zero expected value.) However, the agent still prefers P_2 over P_1. In this case, the agent seems to be averse to losses. Whilst P_2 has a larger chance of incurring a loss, the size of this loss is lower.

- The agent would be better off, if he could avoid the gambles altogether. To see this, suppose the agent had the choice of not gambling, and keeping his \$100 with certainty. (We refer to this as the degenerate lottery P_3) Then his utility is given by $U(P_3) = 1 \cdot \sqrt{100} = 10$. Clearly, this is strictly larger than the expected utility associated with each of the gambles offered. (As we will show below, the agent may even wish to pay a small amount, to avoid having to accept either gamble.)

6.1.3 Risk Preferences

The utility index $u(\cdot)$ represents an agent's attitude towards risk. We characterise risk preferences into 3 categories: risk averse, risk neutral and risk loving. We explore each of these preferences below:

Risk Averse Preferences An agent is **risk averse** if the loss in utility associated with a decrease in wealth is larger than the gain in utility associated with an equivalent increase in wealth. The utility index for risk averse agents is concave in wealth. As is apparent from the diagram, the utility loss associated with a decrease in wealth is larger than the utility gain from an equivalent increase in wealth.

Recall —the utility index in the diagram above represents the utility that the consumer derives when he receives a prize with W with certainty. How do we represent the expected utility over lotteries? Consider a lottery with 2 prizes, such that the final wealth levels are W_1 and W_2 (with $W_1 < W_2$). Let $u(W_1)$ and $u(W_2)$ be the corresponding utility levels associated with these prizes. Suppose the probability of receiving the good prize is p. Then,

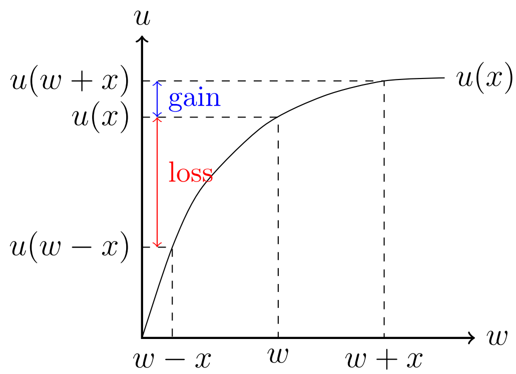

Figure 6.1: Utility Index for a Risk Averse Agent

the expected utility is simply a convex combination of the utilities associated with the final prizes. I.e.

$$U[P] = (1-p)\,u(W_1) + pu(W_2)$$

As before, the expected value of the lottery is given by:

$$E[P] = (1-p)\,W_1 + pW_2$$

Example 24. A consumer has initial wealth of w_0. Let $w_L < w_0 < w_H$. The consumer can either keep his initial wealth for sure (I denote this by the degenerate lottery X_0) or he can accept one of the following lotteries

- X_1 : receive w_H with probability $\frac{1}{4}$ and receive w_L with probability $\frac{3}{4}$
- X_2 : receive w_H with probability $\frac{1}{2}$ and receive w_L with probability $\frac{1}{2}$
- X_3 : receive w_H with probability $\frac{2}{3}$ and receive w_L with probability $\frac{1}{3}$
- X_4 : receive w_H with probability $\frac{9}{10}$ and receive w_L with probability $\frac{1}{10}$

For different lotteries, we represent the expected value along the horizontal axis and the expected utility along the vertical axis. The blue line associates expected values and corresponding expected utilities. The higher is the probability of the good outcome, the further along the blue line is the lottery. A 50-50 lottery appears half way along the line. A lottery with a $\frac{2}{3}$ probability of the good outcome, and a $\frac{1}{3}$ probability of the bad outcome (i.e. lottery X_3) appears $\frac{2}{3}$ of the way along the line. If the agent chooses to keep his original wealth for sure, then he can guarantee himself a utility level $u(w_0)$. It follows that he will accept any of the above lotteries, if the expected utility of the lottery is at least as large as $u(w_0)$.

We note the following properties of the above lotteries. X_1 represents an unfavourable bet, and gives the consumer less utility than if he took w_0 for sure. X_2 represents a fair bet. The consumer receives w_0 in expectation. However, the agent is better off is he rejects the gamble, and keeps his initial wealth for sure. (We see this in the fact that the expected

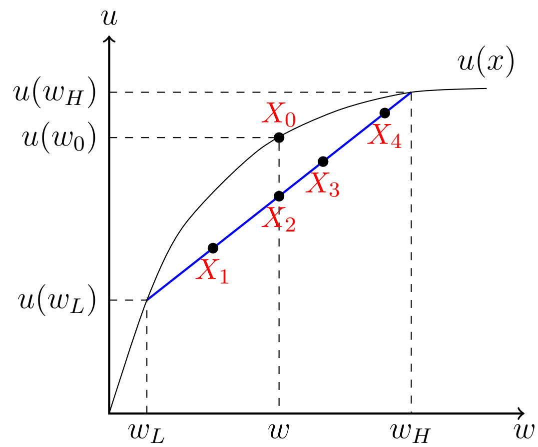

Figure 6.2: Comparing Expected Values and Expected Utilities

utility associated with X_2 is below $u(w_0)$). X_3 represents a favourable bet. However, the agent is still better off accepting w_0 for sure. Finally, X_4 represents a favourable bet, which the consumer actually prefers to the sure thing. (In this case, although there is a still a chance of receiving only w_L, this is so small in comparison to the chance of receiving w_H, that the agent is willing to accept the risk.) □

We note the following results:

- A risk averse agent rejects all unfavourable bets. (Why? The agent is averse to losses, and he is more likely than not to lose in an unfavourable bet.)

- A risk averse agent rejects all fair bets. (Why? The agent values losses more than equivalent gains.) To see this most clearly, note that a fair bet has the same expected payoff as the sure thing. But the utility associated with the sure thing lies on the curved utility function, whilst the utility associated with the gamble lies on the straight line. Since the curve always lies above the line, a risk averse agent always prefers the sure thing to a fair bet.

- A risk averse agent will accept some favourable bets (if they are favourable enough), and reject others.

Note —risk averse agents may differ in their degree of aversion to risk. This is reflected in the curvature of the utility index. An agent who is only slightly risk averse will have a utility index which curves very gently, whereas an agent who is very averse to risk will have a highly 'bowed' utility function. We represent this in the diagram below, where agent A is slightly risk averse and agent B is highly risk averse:

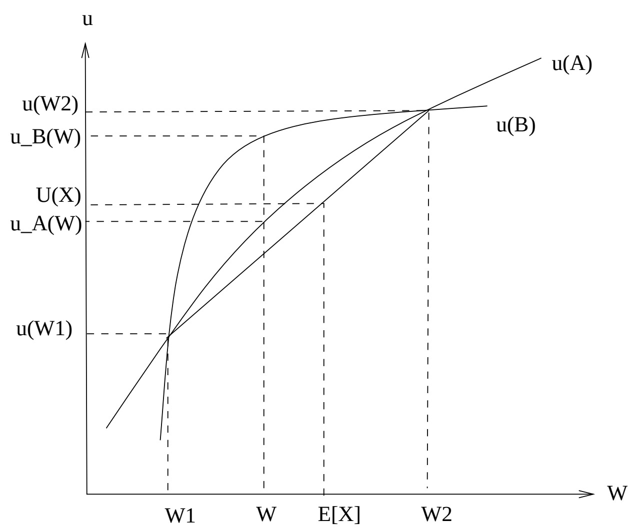

The above diagram considers a situation with two risk averse agents (A and B), both of whom have initial wealth W. Agent B is more risk averse than agent A. Both agents are offered a favourable bet X. As should be clear from the diagram, agent A will accept the lottery, since it improves his utility from $u_A(W)$ to $U(X)$, whilst agent B rejects the same lottery. As agents become more risk averse (i.e. as the utility index becomes more "steeply curved"), agents will be less likely to accept favourable bets.

Risk Neutral Preferences An agent is **risk neutral** if the loss in utility associated with a decrease in wealth exactly matches the gain in utility associated with an equivalent increase in wealth. It follows that the utility index associated with risk neutral preferences is linear in wealth. Suppose $u(w) = \alpha + \beta w$. Then the expected utility of a lottery is given by:

$$U(P) = E[\alpha + \beta w] = \alpha + \beta E[w]$$

It follows that the utility index over prizes and the utility function over lotteries coincide! Diagrammatically, we have:

We note the following results:

- A risk neutral agent rejects all unfavourable bets.
- A risk neutral agent is indifferent between fair bets.
- A risk neutral agent accepts all favourable bets.

Risk Loving Preference An agent is **risk loving** if the loss in utility associated with a decrease in wealth is smaller than the gain in utility associated with an equivalent increase in wealth. The utility index for risk averse agents is convex in wealth. As is apparent from

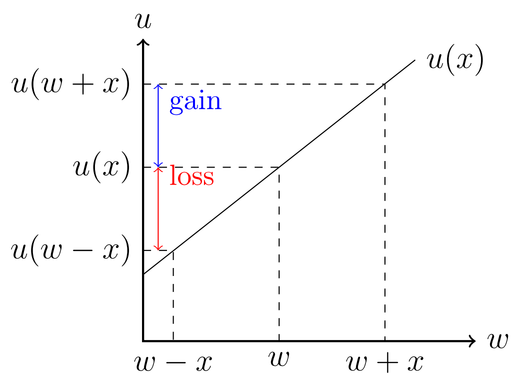

Figure 6.3: Utility Index for a Risk Neutral Agent

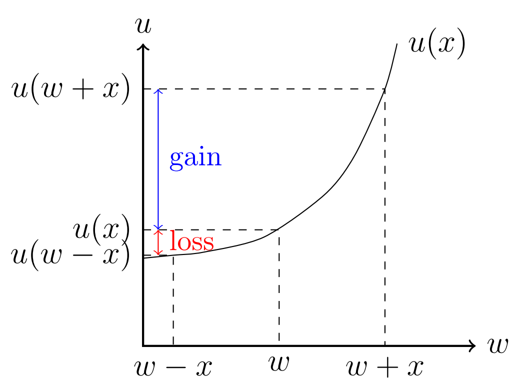

Figure 6.4: Utility Index for a Risk Loving Agent

the diagram, the utility loss associated with a decrease in wealth is smaller than the utility gain from an equivalent increase in wealth.

In the above diagram, X_0 represents the degenerate lottery in which the original wealth level is received with certainty. X_1 represents a fair bet (in which the agent has a 50-50 chance of gaining or losing $\$x$). As we can see, a risk-loving agent prefers the gamble to the sure thing! (Why? the agent savours the prospect of a wealth gain more than the prospect of an equivalent wealth loss.) We have the following results:

- A risk loving agent rejects some unfavourable bets (if they are unfavourable enough) and accepts others.
- A risk loving agent accepts all fair bets.
- A risk loving agent accepts all favourable bets.

6.2 Applications of Expected Utility Theory

6.2.1 Insurance & Risk Premia

Certainty Equivalent In the previous section, we saw that agents are not necessarily indifferent between prizes received with certainty, and lotteries that yield the same expected value. Risk averse agents prefer certain prizes to lotteries with the same expected value. Risk loving agents, on the other hand, prefer lotteries to the certain prize. It follows that if a risk averse agent is indifferent between a lottery and a certain prospect, then the expected value of the lottery must be larger than the prize received with certainty. Similarly, if a risk loving agent is indifferent between a lottery and a prize received with certainty, then the value of the prize must exceed the expected value of the lottery.

Let X be any lottery.

Definition 6. The **certainty equivalent** of X, denoted $CE(X)$ is the monetary value of the prize which if received with certainty, yields the same utility as the lottery itself.

By definition, we must have:

$$u(CE(X)) = U(X)$$

As we argued above:

- for a risk averse agent: $CE(X) < E[X]$

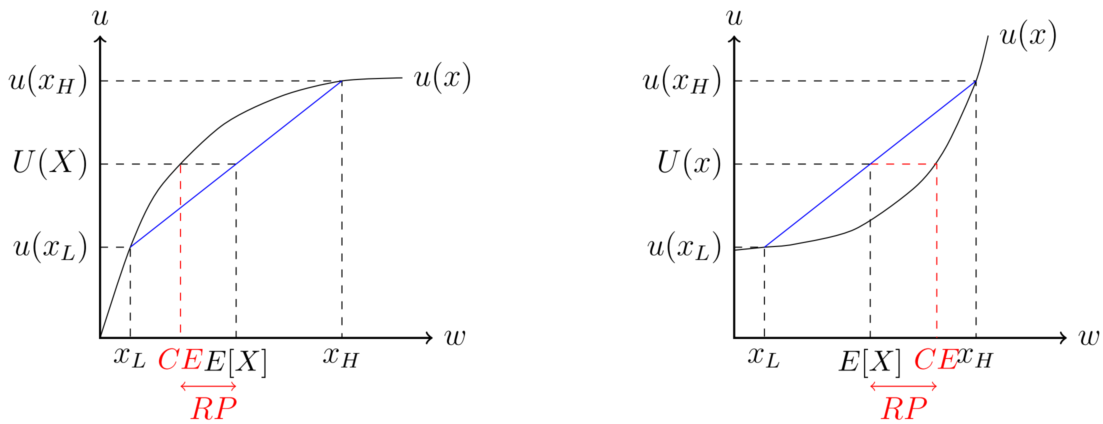

Figure 6.5: Certainty Equivalents and Risk Premia for Risk Averse and Risk Loving Agents.

- for a risk neutral agent: $CE\,(X) = E\,[X]$
- for a risk loving agent: $CE\,(X) > E\,[X]$

We represent the certainty equivalent in the following diagrams:

In each case, the agent faces the same lottery X (which gives prize W_1 with probability $1-p$ and prize W_2 with probability p). The expected value of the lottery is clearly the same for both agents. However, the risk averse agent is indifferent between receiving a prize $W^* < E\,[X]$ for sure, and accepting the outcome of the lottery. On the other hand, the risk loving agent is indifferent between receiving a prize $W^{**} > E\,[X]$ for sure, and accepting the outcome of the lottery. W^* and W^{**} are the certainty equivalents for the risk averse and risk loving agents, respectively.

Example 25 (Finding the Certainty Equivalent)**.** Bob has utility index over monetary prizes given by $u\,(W) = \ln W$. Bob is offered a job which pays in bonuses. He knows that he will earn \$150, 000 with probability 0.6 (if business is good), but only \$50, 000 with probability 0.4 (if business is bad). What is his certainty equivalent?

The agent's income is a lottery, with expected value:

$$E\,[X] = \$150,000 \times 0.6 + \$50,000 \times 0.4 = 110,000$$

His expected utility is:

$$U\,(X) = \ln(150000) \times 0.6 + \ln(50000) \times 0.4 = 11.479$$

The certainty equivalent W^*, is a level of income which is received with certainty and yields the same utility as the lottery. Hence, we have:

$$\begin{aligned} u\,(W^*) &= U\,(X) \\ \ln W^* &= 11.479 \\ W^* &= \$96,659 \end{aligned}$$

Hence, Bob is indifferent between accepting the risky job, and a job which pays a sure income if \$96,659. □

Risk Premium Consider an agent who faces a risky prospect X, with expected value $E[X]$. Suppose the agent is offered the option to replace the risky prospect off the agent's hands with a certain prize of $E[X]$. (This is clearly advantageous for the risk averse agent. In monetary terms, he is just as well-off in expectation, but now faces less risk.) How much will the agent be willing to pay for such an insurance service?

The agent will at most pay an amount that causes his utility from the sure thing (less the risk premium) to equal the utility from continuing to face the risky prospect. I.e. the agent will accept the switch as long as:

$$u(E[X] - RP) \geq U(X)$$

It follows that the maximum price the agent will be willing to pay is RP^*, where RP^* satisfies:

$$u(E[X] - RP^*) = U(X)$$

But we know that $u(CE(X)) = U(X)$. Hence, we have:

$$\begin{aligned} CE(X) &= E[X] - RP^* \\ RP^* &= E[X] - CE(X) \end{aligned}$$

The risk premium associated with a lottery is the amount that causes his net wealth to fall to the certainty equivalent. Why? If he pays any more, then the utility of the certain prospect will be less than the utility associated with the gamble. The agent would rather take the gamble and get the potentially high payoff, than pay to avoid the risk;

Example 26. Consider the example of Bob's income (above). Bob's expected income is $\$110,000$ and his certainty equivalent is $\$96,659$. An insurer offers to guarantee Bob a certain income of $\$110,000$. (I.e. if business is good, the insurer gets $\$150,000$ and pays $\$110,000$ to Bob, whilst if business is bad, the insurer gets $\$50,000$ and still pays $\$110,000$ to Bob). How much will Bob pay for this insurance. Since Bob is indifferent between receiving $\$110,000$ with uncertainty, and $\$96,659$ for sure, he will be willing to pay up to $\$13,341$ to insure against this risk. (If he pays this amount in full, then he will be just as well off as he was facing the risk. If he pays less than the full amount, then he is unambiguously better off).

Since a risk averse agent prefers certain to uncertain outcomes, he will always be willing to pay some amount to insure himself against a risk. (Mathematically, this follows from the fact that $CE(X) < E[X]$).

Consider the following example:

Example 27. Ben has a utility index over wealth given by $u(W) = \ln W$. At his current job, Ben receives a fixed income of \$100,000. His boss offers him a promotion to a managerial position that pays according to performance. Ben knows that he will receive \$50,000 with probability $\frac{1}{3}$ (if the company performs badly) and \$150,000 with probability $\frac{2}{3}$ (if the company performs well). Should Ben accept the promotion? If he does how much will he be willing to pay to insure himself against the risk? Suppose Ben keeps his current job. Then his utility is $u(100,000) = \ln(100,000) = 11.513$. If he accepts the promotion, then his expected utility is:

$$U(X) = \frac{1}{3}\ln(50,000) + \frac{2}{3}\ln(150,000) = 11.552$$

Hence, Ben can increase his utility by accepting the promotion. Having accepted the promotion, Ben now faces a risky prospect. Since he is risk averse, there is an incentive to insure himself against this risk. To find the maximum insurance premium he is willing to pay, we must find the certainty equivalent W^*. We have:

$$\begin{aligned} u(W^*) &= U(X) \\ \ln W^* &= \frac{1}{3}\ln(50,000) + \frac{2}{3}\ln(150,000) \\ W^* &= \$104,004 \end{aligned}$$

Hence, Ben is indifferent between the lottery and receiving a sure income of \$104,004. His expected salary is $\frac{1}{3} \cdot 50000 + \frac{2}{4} \cdot 150000 = \$116,667$. The maximum insurance premium he will be willing to pay is \$12,662. See Figure 27 below.

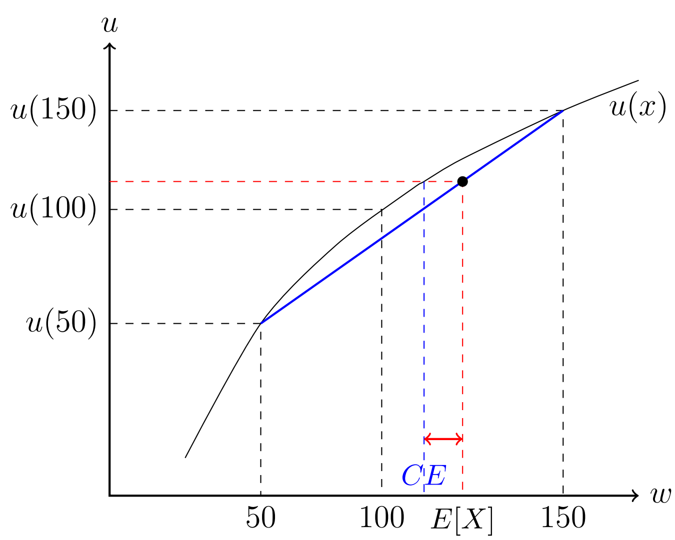

Figure 6.6: Ben's Promotion and Insurance Decision

In the above example, the risk averse agent voluntarily accepts a risky prospect, and then insures himself against it! At first blush it may appear contradictory that an agent would simultaneously accept risk and insure himself against it. The above example point highlights two important facts:

1. Risk averse agents will accept risks if they are favourable enough. (In the above example, the gamble was weighted so far in favour of the good outcome that the agent was willing to voluntarily accept the risk).
2. Whenever a risk averse agent faces a gamble, there is an incentive to insure against it - even if the agent accepted the gamble voluntarily.

6.3 The State-Space Approach

In the above section, we describe qualitatively whether an agent would accept a particular gamble —and showed that the decision to do so depended on both the favourability of the bet and the agent's attitude towards risk. In this section, we seek to quantify this analysis. In particular, we ask the question - if an agent is given the option of gambling some or all of his wealth, how much will he choose to risk?

6.3.1 The Model

Budget Constraint Suppose the agent begins with a certain level of wealth. The agent is offered a gamble in which he loses $\$b$ with probability π_b and wins $\$g$ with probability π_g (where $\pi_b + \pi_g = 1$). There are two states of the world —C_b (the bad state) and C_g (the good state) —which refer to the agent's consumption, given the outcome of the lottery. The agent can select multiple units of the gamble. (i.e. If he chooses 3 'units', then he wins $\$3g$ in the good state and loses $\$3b$ in the bad state. Similarly, if he chooses $\frac{1}{2}$ units, then he wins $\$\frac{1}{2}g$ units in the good state and loses $\$\frac{1}{2}b$ in the bad state.) By choosing how much to gamble, the agent can achieve different levels of consumption in the good and bad states.

The state space approach to choice under uncertainty is analogous to the other models of agent decision making we saw in the sections on Consumer Theory and its Applications . However, rather than choosing between two goods (as in traditional consumer theory) or between consumption in two periods, or between consumption and leisure, in this context, the agent is required to choose between consumption in the good and bad states of the world. The budget constraint shows the various combinations of C_b and C_g that the agent can achieve by choosing to gamble different amounts.

Suppose the agent purchases x units of the gamble. Then, in the good state his consumption is:

$$C_g = W + xg \tag{6.1}$$

, whilst in the bad state his consumption is:

$$C_b = W - xb \tag{6.2}$$

From (6.2), we have $x = \frac{W-C_b}{b}$. Substituting this into (6.1) we have:

$$\begin{aligned} C_g &= W + \left(\frac{W - C_b}{b}\right) g \\ C_g + \frac{g}{b} C_b &= \frac{b+g}{b} W \end{aligned} \tag{6.3}$$

where (6.3) gives an expression for the budget constraint. The slope of the budget constraint is:

$$slope\ BC = -\frac{g}{b}$$

which implies that for each additional unit of consumption the agent stands to lose in the bad state (by gambling a little more), the agent stands to gain $\frac{g}{b}$ units of consumption in the good state. (We can think of $\frac{g}{b}$ as the price of consumption in the good state, in terms of consumption in the bad state which we must forgo).

The **fair-odds line** represents the combinations of consumptions in the good and bad states in which the agent's expected consumption is unchanged. (The fair-odds line can be thought of the budget line that would arise if the agent was offered a fair-bet.) The fair-odds line is defined by:

$$\pi_b C_b + \pi_g C_g = W$$

The slope of the fair-odds line is given by:

$$slope\ FO = -\frac{\pi_b}{\pi_g}$$

Suppose the actual gamble represents a favourable bet. Then we have:

$$\begin{aligned} \pi_b(-b) + \pi_g(g) &> 0 \\ \frac{g}{b} &> \frac{\pi_b}{\pi_g} \end{aligned}$$

which implies that the slope of the budget line is steeper than the slope of the fair-odds line. We can similarly verify that if the actual gamble is unfavourable, then the slope of the budget line will be flatter than the slope of the fair odds line.

The **certainty line** represents the set of consumption profiles in which the agent's consumption is certain (i.e. it does not depend on the state of the world). Clearly, along the certainty line, we must have $C_g = C_b$.

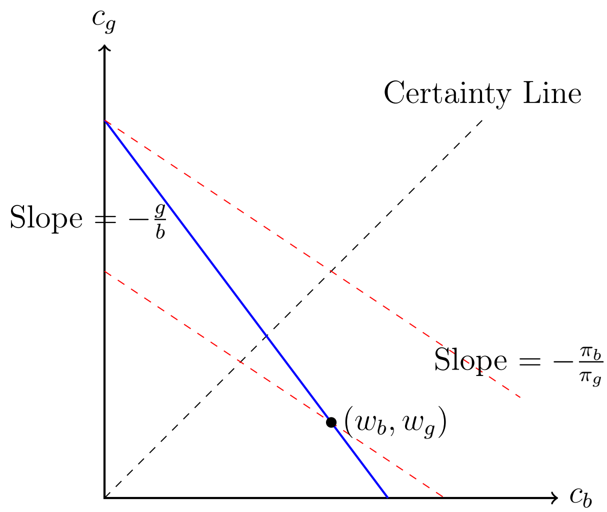

Figure 6.7: Budget Constraint over Contingent Consumption Plans

Preferences We now characterise preferences over uncertain prospects. We can represent preferences using indifference curves over consumption in the good and bad states. The usual assumptions from consumer theory carry over. By monotonicity, we know that consumers strictly prefer more consumption in both states to less - which implies that indifference curves must be downward sloping. Moreover, indifference curves which are further from the origin represent more preferred consumption profiles.

Recall, an agent's expected utility over consumption profiles is given by:

$$U(X) = \pi_b u(C_b) + \pi_g u(C_g)$$

Then, an indifference curve defines the set of consumption profiles (C_b, C_g) which give the same level of expected utility. This implies:

$$\begin{aligned} dU &= \pi_b u'(C_b)\, dC_b + \pi_g u'(C_g)\, dC_g = 0 \\ \Rightarrow \quad \frac{dC_g}{dC_b} &= -\frac{\pi_b}{\pi_g} \cdot \frac{u'(C_b)}{u'(C_g)} \end{aligned}$$

where this latter expression gives the slope of the indifference curve at any consumption profile.

We note that along the certainty line (i.e. when $C_g = C_b$), the slope of any indifference curve is given by:

$$\frac{dC_g}{dC_b} = -\frac{\pi_b}{\pi_g} = slope\ FO$$

Hence, along the certainty line, the indifference curve must be tangent to the fair-odds line.

The curvature of indifference curves depends upon the agent's attitudes towards risk. Consider a risk-averse agent. Recall, a risk averse agent values a loss in consumption by more

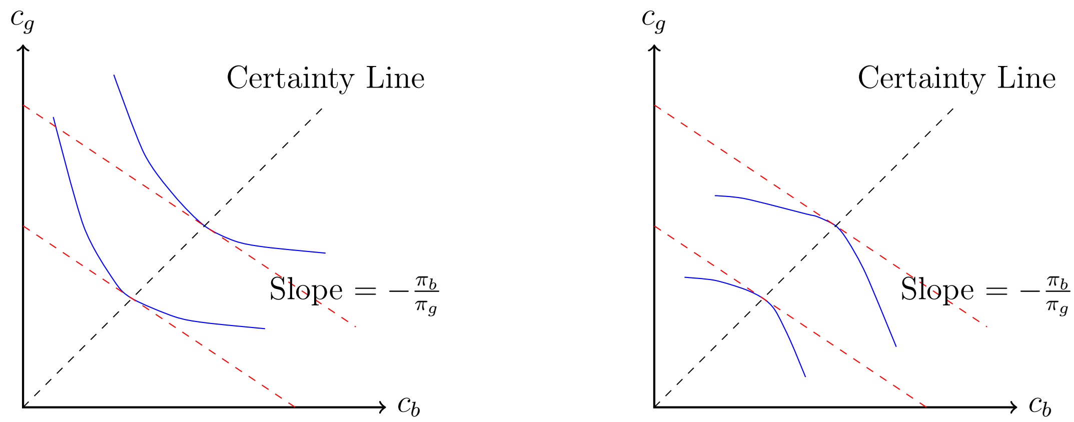

Figure 6.8: Risk preferences and the shape indifference curves. The left panel shows indifference curves for a risk averse agent. The right panel shows indifference curves for a risk loving agent. Indifference curves for a risk-neutral agent are linear and have the same slope as fair odds. Notice that, in all 3 cases, indifference curves have the same slope as fair odds at the certainty line.

than an equivalent gain in consumption. Hence, a risk-averse agent requires a lot of extra consumption in the good state to compensate for the loss of consumption in the bad state. Moreover, since risk averse agents are more wary of large losses than small losses, the size of the consumption necessary in the good start is increasing in the amount lost in the bad state. This implies that risk-averse agents have convex indifference curves (as shown in Figure 6.8).

Risk-neutral agents value gains and losses in consumption equally. Hence, risk-neutral agents have linear indifference curves, whose slope corresponds to the slope of the fair-odds line. Finally, we noted above that risk-loving agents value losses in consumption by less than equivalent gains in consumption. It follows that risk-loving agents have concave indifference curves.

Optimisation The agent maximises his expected utility subject to the budget constraint. He will choose to gamble the amount which causes his consumption profile to achieve the highest possible indifference curve. In the usual way, such a consumption profile must satisfy the tangency condition.

Suppose the agent is risk averse and is offered a favourable bet. Will the agent accept the gamble? If so, how much will he choose to gamble? We have two important results:

Lemma 5. *Suppose a risk averse agent faces a fair-odds gamble. Her optimal contingent consumption plan will be certain (i.e.* $c_b = c_g$*).*

To understand why, note that if a bet is fair, then the budget line corresponds to the fair odds line. But we know that indifference curve is tangent to the fair-odds line at the certainty line. Hence, the risk averse agent maximises his expected utility by choosing not to gamble.

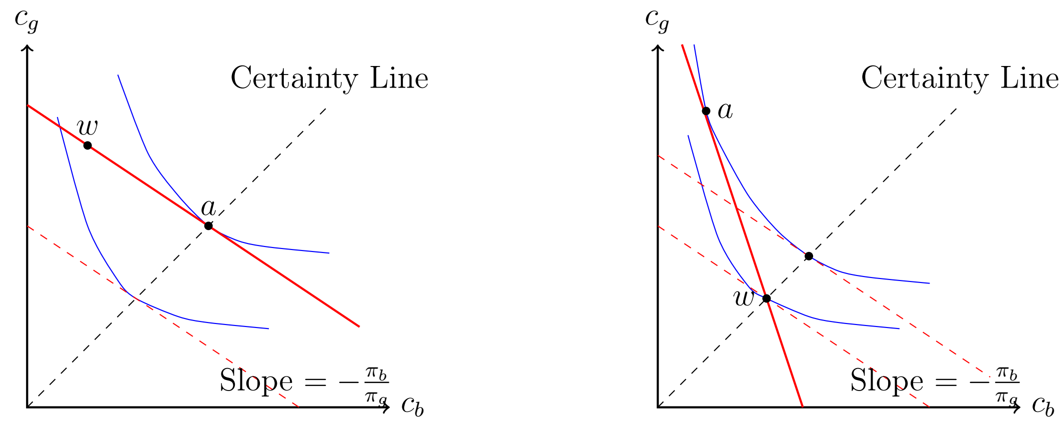

Figure 6.9: Decision making by a risk averse agent. The left panel shows that the risk averse agent facing a fair-odds gamble will trade away all risk so that her consumption plan is non-state contingent. The right panel shows that a risk averse agent facing a favorable gamble will always leave some residual risk in her consumption plan. She will not fully insure herself.

We demonstrate this in Figure 6.9. It is also easy to verify that a risk-averse agent will reject any unfavourable bet. To see why, note that the budget line for an unfavourable bet lies below the fair-odds line. If so, by accepting an unfavourable bet, the agent guarantees himself a strictly lower level of utility than if he remained on the certainty line.

A risk-averse agent will always accept some level of risk when faced with a favourable bet. (Why? If the agent does not gamble, then he will be on an indifference curve which is tangent to the fair-odds line. But since the bet is favourable, the budget line lies above the fair-odds line. Hence, the original indifference curve lies in the interior of the budget set. But we know that this cannot be optimal. The agent can achieve a higher indifference curve by gambling a small amount.) In this case, the consumer will be willing to gamble $\$(W - C_b)$ in the hope of winning $\$(C_g - W)$.

Note —in the previous section, we said that a risk-averse agent would accept a favourable bet only if it was favourable enough, and would often reject favourable bets. But now it seems that a risk averse agent will always accept a favourable bet. Why is there a difference? In the previous section we assumed that the agent was made an offer to purchase a single unit of a gamble and to take the winnings/losses accordingly. In this section, we assume that agents can purchase multiple units or fractional units of gambles. If at the optimum, the agent is willing to purchase a fraction of a gamble (but not the entire gamble) —then in the previous section we would've concluded that the agent would reject the offer. It should be clear that the two approaches are not inconsistent —they just differ slightly in the assumptions about how much is gambled.

Intuitively, it should be apparent that the degree to which risk averse agents are willing to accept risks depends on the degree of risk aversion. The degree of risk aversion is reflected in the curvature of the indifference curves. Agents who are highly risk averse will have "strongly bowed" indifference curves, whilst agents who are only somewhat risk averse will have less bowed indifference curves. (Note - regardless of the extent of risk aversion, it must still be the case that the slope of the indifference curve along the certainty line coincides with the slope of the fair odds line.) Agents who are more risk averse will gamble a smaller portion of their wealth than agents are who are less risk averse.

Example 28 (Numerical)**.** A risk averse agent has utility index $u(W) = \sqrt{W}$. The agent originally has \$100. She is offered the following gamble: A fair die is rolled. If the outcome is high (either 5 or 6), then she wins \$3. Otherwise she loses \$1. How much will gamble?

Since the gamble is a favourable bet, we know that the agent will gamble a positive amount. The budget constraint is given by:

$$\begin{aligned} C_g + \frac{g}{b}C_b &= \frac{b+g}{b}W \\ C_g + \frac{3}{1}C_b &= \frac{3+1}{1}W \\ C_g + 3C_b &= 400 \end{aligned}$$

and so, the slope of the budget constraint is:

$$slope\ BC = -\frac{g}{b} = -3$$

Since $u(W) = \sqrt{W}$, we know that $u'(W) = \frac{1}{2\sqrt{W}}$. The slope of the indifference curve is given by:

$$\begin{aligned} slope\ IC &= -\frac{\pi_b}{\pi_g} \cdot \frac{u'(C_b)}{u'(C_g)} \\ &= -\frac{\frac{2}{3}}{\frac{1}{3}} \cdot \frac{\frac{1}{2\sqrt{C_b}}}{\frac{1}{2\sqrt{C_g}}} \\ &= -2\sqrt{\frac{C_g}{C_b}} \end{aligned}$$

From the tangency condition, we have:

$$\begin{aligned} slope\ IC &= slope\ BC \\ -2\sqrt{\frac{C_g}{C_b}} &= -3 \\ C_g &= \frac{9}{4}C_b \end{aligned}$$

Substituting this into the budget constraint gives:

$$\begin{aligned} \left(\frac{9}{4}C_b\right) + 3C_b &= 400 \\ \frac{21}{4}C_b &= 400 \\ C_b^* &= \frac{1600}{21} = 76.19 \end{aligned}$$

It follows that:

$$C_g^* = \frac{9}{4} C_b^* = \frac{1200}{7} = 171.43$$

The agent is willing to gamble up to 23.8% of his wealth, in the hope of getting the good outcome in the favourable bet. □

What about risk-neutral and risk-loving agents? We have already shown that a risk-neutral agent is indifferent between fair bets (since the agent's indifference curve and the budget line for a fair bet, coincide). When faced with a favourable bet, a risk-neutral agent will gamble his entire wealth. Conversely, the risk averse agent maximises his utility by rejecting unfavourable bets. See Figure 6.10.

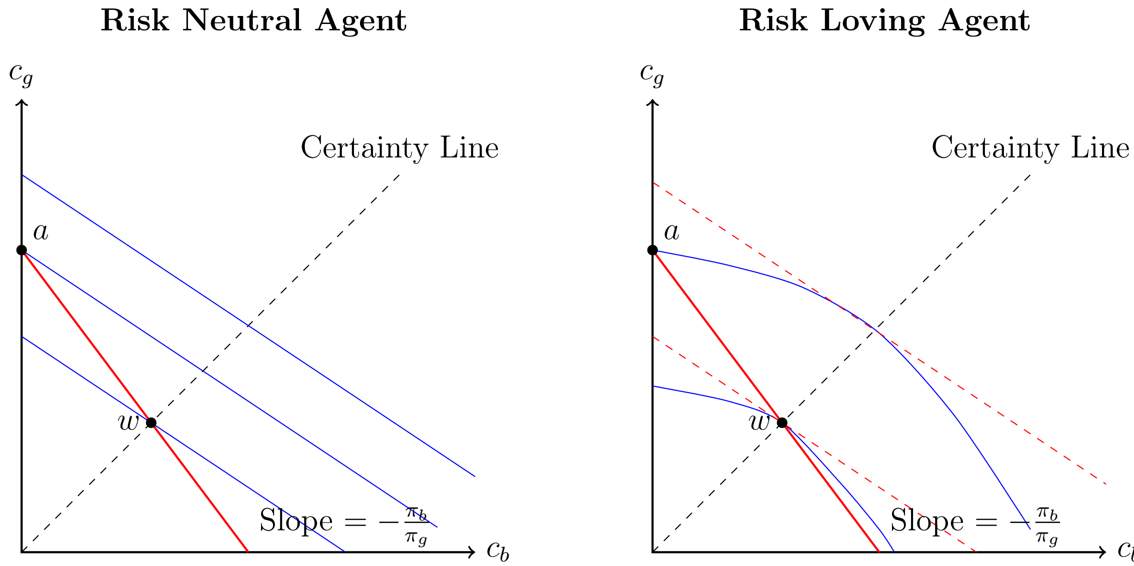

Figure 6.10: Both a risk neutral (left panel) and risk loving (right panel) agent will gamble their entire wealth when confronted with a favorable bet.

A risk-loving agent will gamble his entire wealth when faced with a fair or favourable bet. See Figure 6.10. Furthermore, a risk-loving agent will gamble his entire wealth for some unfavourable bets, but in others he will not gamble at all. (A risk-loving agent will never choose to partially gamble. His participation in a gamble is either all or nothing).

6.3.2 Application: Optimal Participation in Crime

Consider a risk-averse individual who works T hours each week. The agent can choose to allocate his working time between legitimate and illegitimate activities. Suppose for each hour of legitimate work, the agent receives a wage w, and for each hour of illegitimate work, he receives a wage m with $m > w$. The authorities can only police illegal activity imperfectly, and so they catch workers doing illegitimate work with probability p. If caught, the worker forfeits his wage for each hour of illegal work, and in addition must pay a fine f for each hour of illegitimate work. Suppose the agent has logarithmic utility index over consumption. i.e. $u(C) = \ln C$. (It follows that $u'(C) = \frac{1}{C}$)

Let the bad and good states correspond to the states in which he is caught and not-caught, respectively. What is the agent's budget constraint? If the agent performs only legitimate work, then his consumption will be the same in both periods - and he consumes $C_b = C_g = Tw$. Suppose the agent spends x hours in illegal activity. Then, if he is caught, his consumption is:

$$C_b = w(T - x) - fx = wT - (w + f)x \tag{6.4}$$

If he is not caught, his consumption is:

$$C_g = w(T - x) + mx = wT + (m - w)x \tag{6.5}$$

Rearranging (6.4) we have:

$$x = \frac{wT - C_b}{w + f} \tag{6.6}$$

and substituting this into (6.5) gives:

$$\begin{aligned} C_g &= wT + (m - w)\left(\frac{wT - C_b}{w + f}\right) \\ C_g + \frac{m - w}{w + f}C_b &= \left(\frac{m + f}{w + f}\right)wT \end{aligned}$$

which is our expression for the budget constraint. Moreover, the slope of the budget constraint is given by:

$$slope\ BC = -\frac{m - w}{w + f}$$

Will the agent choose allocate any of his time to illegal work? Since the agent is risk averse, we know that he will engage in illegal work only if the payoff to illegitimate activity is a favourable bet. (If the payoff to crime is either a fair or unfavourable bet, then the agent maximises his utility by not accepting the gamble (i.e. by only engaging in legal activity). How do check if the gamble is favourable or not? Recall, a gamble is favourable if the slope of the budget constraint is steeper than the slope of the fair odds line.

What is the fair-odds line? Recall, the agent's expected consumption does not change along the fair odds line. Hence, we have:

$$(1 - p)C_g + pC_b = wT$$

The slope of the fair odds line is given by:

$$slope\ FO = -\frac{p}{1 - p}$$

Hence, illegal activity is favourable if:

$$\begin{aligned}
|slope\ BC| &> |slope\ FO| \\
\frac{m-w}{w+f} &> \frac{p}{1-p} \\
f &< \frac{m-w}{p} - m \qquad (6.7)
\end{aligned}$$

Comments

- Equation (6.7) states that the agent will participate in criminal activity if the fine f is not too large (relative to the payoff to criminal activity and given the probability of being caught).
- How large must the fine be if it is to be effective in deterring criminal activity? We need the inequality to reverse, so that criminal activity is either a fair or unfavourable gamble. The minimum effective fine is given by:

$$f^* = \frac{m-w}{p} - m$$

- We note that f^* is increasing in $(m-w)$,which implies that the fine must be larger when the net payoff to crime $(m-w)$ is higher.
- We note further that f^* is decreasing in p, which implies that a smaller fine can still be effective if the probability of being caught is larger. (i.e. if there is a more extensive/efficient police force).

Now, suppose (6.7) is satisfied, so that even a risk averse agent has an incentive to engage in criminal activity. How many hours of illegitimate work will he choose? The optimal level of criminal activity must satisfy the tangency condition. The slope of the indifference curve is given by:

$$\begin{aligned}
slope\ IC &= -\frac{p}{1-p}\frac{u'(C_b)}{u'(C_g)} \\
&= -\frac{p}{1-p}\frac{\frac{1}{C_b}}{\frac{1}{C_g}} \\
&= -\frac{p}{1-p}\frac{C_g}{C_b}
\end{aligned}$$

Then, tangency requires:

$$\begin{aligned}
slope\ IC &= slope\ BC \\
-\frac{p}{1-p}\frac{C_g}{C_b} &= -\frac{m-w}{w+f} \\
C_g &= \frac{1-p}{p} \cdot \frac{m-w}{w+f} C_b
\end{aligned}$$

Substituting this into the budget constraint, we have:

$$\begin{aligned}
\left(\frac{1-p}{p}\cdot\frac{m-w}{w+f}C_b\right)+\frac{m-w}{w+f}C_b &= \left(\frac{m+f}{w+f}\right)wT \\
\frac{1}{p}\frac{m-w}{w+f}C_b &= \left(\frac{m+f}{w+f}\right)wT \\
C_b &= p\frac{m+f}{m-w}wT
\end{aligned}$$

It follows that:

$$\begin{aligned}
C_g &= \frac{1-p}{p}\cdot\frac{m-w}{w+f}C_b \\
&= \frac{1-p}{p}\cdot\frac{m-w}{w+f}\left[p\frac{m+f}{m-w}wT\right] \\
&= (1-p)\frac{m+f}{w+f}wT
\end{aligned}$$

Finally, by (6.6), the number of hours spent in illegal activity is:

$$\begin{aligned}
x^* &= \frac{wT-C_b}{w+f} \\
&= \left[1-p\frac{m+f}{m-w}\right]\frac{w}{w+f}T
\end{aligned}$$

We verify that the number of hours spent in criminal activity is:

- decreasing in p - as the probability of getting caught increases, agents substitute away from illegal activity, towards legitimate work.
- decreasing in f - as the punishment for illegal activity increases, agents spend less time pursuing illegitimate activities
- increasing in $(m-w)$ which is the net payoff to criminal activity. (It follows that one method of disincentivising criminal activity is to increase the payoff to legitimate work, by increasing w)

Example 29 (Calibration)**.** Suppose $w = 20, m = 50, T = 40$,and $p = 0.27$. Then, the minimum fine needed to disincentivise criminal behaviour is given by:

$$\begin{aligned}
f^* &= \frac{m-w}{p}-m \\
&= \frac{50-20}{0.27}-50 \\
&= \$61.11
\end{aligned}$$

Suppose the fine is only $f = 30$. Then, the fine is sufficiently low that even risk averse agents will choose to engage in some amount of criminal activity. How much will they choose? We have:

$$\begin{aligned} x^* &= \left[1 - p\frac{m+f}{m-w}\right]\frac{w}{w+f}T \\ &= \left[1 - 0.27 \cdot \frac{50+30}{50-20}\right] \cdot \frac{20}{20+30} \cdot 40 \\ &= 4.48 \end{aligned}$$

Hence, there is an incentive for the typical risk averse worker to shirk for $\frac{1}{2}$ a day each week, and engage in illegitimate work activity. □

Exercise 4 (Further Calibration)**.** In the above application we assumed that agents have log utility index over consumption. This represents a particular degree of aversion towards risk. Suppose instead that the agent has utility index over consumption given by the Constant Relative Risk Aversion (CRRA) function:

$$u(C) = \frac{C^{1-\theta}}{1-\theta}$$

where $\theta > 0$ measures the degree of risk aversion. (If $\theta = 0$, the agent is risk neutral. As $\theta \to \infty$ the agent becomes increasingly risk averse, such that in the limit, he will never accept risk. It turns out that the utility function we used above ($u = \log C$) is a special case of this utility function, when $\theta = 1$.) Find the optimal level of illegal criminal activity as a function of the risk aversion parameter θ. (It may help to consider a few different specific cases, e.g. $\theta = 0.5, \theta = 3$ and so on). What happens to the optimal level of criminal activity as $\theta \to \infty$?

6.3.3 Application: Optimal Insurance

In the application on criminal activity, we assumed (implicitly) that agents began on the certainty line (i.e. when they did not engage in criminal activity) and then voluntarily chose to incur some risk in order to take advantage of the favourable-bet. In this application, we consider a case where an agent originally finds herself facing a risky prospect and seeks to insure herself against the risk.

Suppose the agent's income is uncertain, and that she receives income y_g in the good state with probability π_g, and income y_b in the bad state, with probability π_b. (We represent this uncertain prospect by point A in the diagram below. For notational convenience, I will refer to the lottery over income as the risky prospect Y) The agent's expected utility is given by:

$$U(Y) = \pi_g u(y_g) + \pi_b u(y_b)$$

The indifference curve IC_0 represents all the consumption profiles which give the agent the same level of utility as the risky prospect Y. Since the agent is risk-averse, we know that the indifference will be convex shaped. Moreover, we know that when the indifference curve intersects the certainty line, it will have the same slope as the fair-odds line. Point B represents the **certainty equivalent** of the risky prospect Y. Recall - the certainty equivalent is the level of income which the agent receives with certainty which leaves her just as well off as if she faced the risky prospect Y.

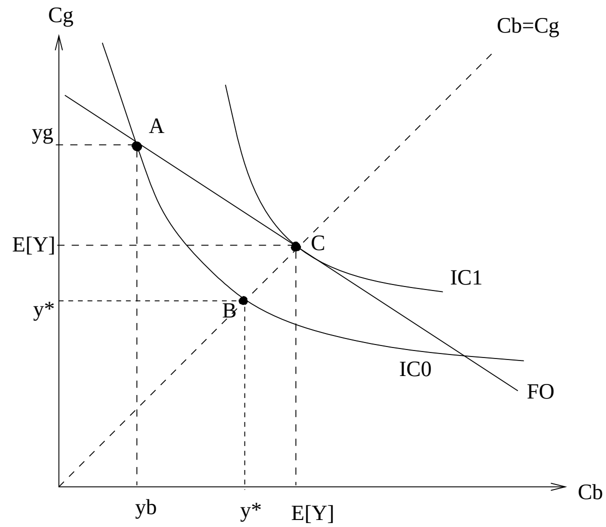

The expected value of the risky prospect Y is given by:

$$E[Y] = \pi_g y_g + \pi_b y_b$$

The fair-odds line is the set of all consumption profiles which give the agent the same expected consumption as under the risky prospect Y. The slope of the fair-odds line is given by:

$$slope\ FO = -\frac{\pi_b}{\pi_g}$$

Let point C be the intersection of the fair-odds line and the certainty line. By the definition of the expected value, it should be clear that point C corresponds to the situation when the agent receives $E[Y]$ for sure. We note that the certainty equivalent y^* is less than the expected value of the prospect $E[Y]$. This follows from the fact that the agent is risk averse —and so is willing to sacrifice income in order to avoid risk.

Insurance Suppose an insurance firm offers the agent the following insurance package: The agent can purchase insurance in \$1 increments. (She pays this amount prior to the state of the world being realised.) If the good state of the world is realised, then the agent forfeits the \$1. If the bad state of the world is realised, then the agent receives $\$r$ (i.e. the agent receives back her \$1 plus a further $\$(r-1)$.) By purchasing insurance, the agent can reduce the risk associated with her income. If so, she will receive less income in the good state (since she must pay the insurance cost), but she will receive more income in the bad state. Suppose she purchases x packages of insurance. Then, her consumption in the good and bad states is given by:

$$\begin{aligned} C_b &= y_b + (r-1)\,x \\ C_g &= y_g - x \end{aligned}$$

Combining these, gives the budget constraint associated with the insurance scheme. We have:

$$\begin{aligned} C_g &= y_g - \frac{C_b - y_b}{r-1} \\ \frac{1}{r-1}C_b + C_g &= \frac{1}{r-1}y_b + y_g \end{aligned} \tag{6.8}$$

The slope of the budget constraint is given by:

$$slope\ BC = -\frac{1}{r-1}$$

Actuarially Fair Insurance Consider the insurance scheme from the perspective of the insurer. The insurer gains \$1 with probability π_g (i.e. if the state is good, and the insurer does not need to compensate the agent) and loses \$ $(r-1)$ with probability π_b (if the bad state is realised). The expected payoff to the insurer is:

$$EV = 1 \cdot \pi_g - (r-1) \cdot \pi_b$$

We say that insurance is *actuarially fair*, if the expected payoff to the insurer is zero. (i.e. the gains during the good periods exactly offset the loses during the bad periods). The actuarially fair insurance package is given by:

$$\begin{aligned} EV &= 1 \cdot \pi_g - (r-1) \cdot \pi_b = 0 \\ &\Rightarrow\ r = \frac{\pi_g}{\pi_b} + 1 \end{aligned} \tag{6.9}$$

The budget line for an actuarially fair insurance package corresponds to the fair-odds line. (The reason why should be fairly intuitive. The fair-odds line represents all the consumption profiles in which the agent receives the same expected income. In an actuarially fair insurance scheme, the income which the agent forgoes in the good state is exactly offset by the compensation received in the bad state - so that expected income is unchanged). We further verify (by rearranging (6.9) that:

$$-\frac{\pi_b}{\pi_g} = -\frac{1}{r-1}$$

Recall, the slope of the budget constraint is given by $-\frac{1}{r-1}$, and the slope of the fair-odds line is given by $-\frac{\pi_b}{\pi_g}$. Since these are equal (and both line pass through point A), we confirm that the budget line for an actuarially fair insurance scheme coincides with the fair-odds line.

If a risk averse agent is offered fair insurance, then she will completely insure herself against the risk. (I.e. she will purchase sufficiently many 'units' of insurance, so that her consumption is on the certainty line - i.e. Point C on the above diagram). Why must this be true? We know that the optimal level of insurance must satisfy the tangency condition. Furthermore, we know that indifference curves are tangent to the fair-odds line along the certainty line. Hence - the optimal level of insurance must be the amount which alleviates the agent of all risk.

Unfair Insurance In reality, we know that insurers will never offer actuarially fair odds. (If they did, they would not be able to make any profit - let alone cover the costs of providing the insurance services, such as paying wages of employees.) Insurers will more likely offer unfair odds - i.e. those in which the expected return to the insurer is positive (or conversely, where the expected return to the agent is negative). Faced with an unfavourable bet, will a risk averse agent still choose to take out insurance?

The answer is yes! Consider the following diagram:

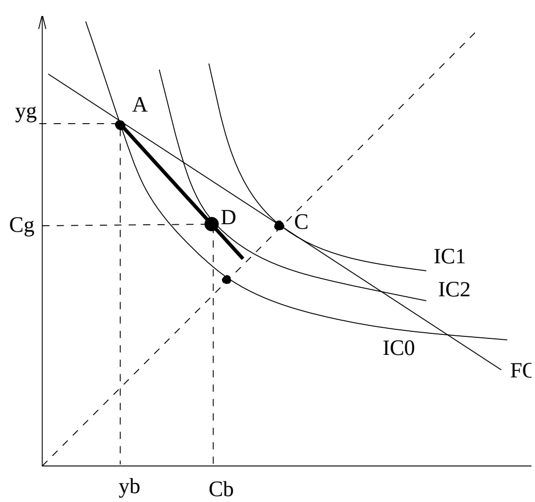

In the above diagram, the thick line is the budget constraint of the actuarially unfair insurance package, whilst the thin line is the fair odds line. (The budget constraint line lies below the fair-odds line, since the agent must forgo more consumption in the good state to guarantee a particular level of consumption in the bad state.) The optimal level of insurance is given by point D. The agent purchases $\$\,(y_g - C_g)$ worth of insurance. Note - at this point, the agent does not completely insure herself against risk. (How do we know this? At points along the certainty line, no indifference curve can be tangent to the budget line. Why? We know that, at points along the certainty line, every indifference curve has the same slope as the fair-odds line. But the slope of the budget line is steeper than the slope of the fair-odds line. Hence, we cannot have tangency along the certainty line.) It should be fairly intuitive that since the agent is getting an unfavourable deal, she will not be willing to completely insure herself against risk - however she will still be willing to partially do so.

But doesn't this contradict our previous claim that risk averse agents reject all unfavourable bets? No. We made that claim in the context of an agent who had the choice between keeping a certain level of wealth and opting to take an unfavourable bet. (Clearly in that case, the agent would prefer to keep her sure income). In this context, the agent faces a risky prospect. Even though she is offered an unfavourable bet, she will accept it to the extent that it helps reduce the risk she faces.)

It shouldn't come as any surprise that the willingness of an agent to accept unfair insurance depends on her degree of risk aversion. In the diagram below, two different agents face the same risky prospect Y, and are offered the same unfair insurance package. Point A represents the risky prospect Y. Point C is the situation where the agent chooses to completely insure herself against the risk. The individual in the left hand diagram is less risk averse than the agent in the right hand diagram, and so she chooses a lower level of insurance.

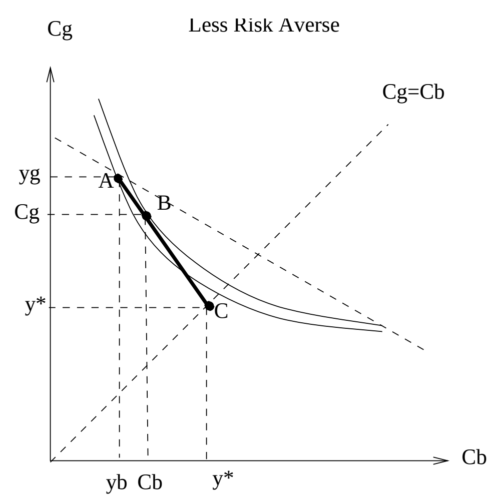

Exercise 5. Suppose the agent is perfectly risk averse. What do the agent's indifference curves look like? How much insurance will the agent purchase? Will she completely insure herself in the face of unfair insurance?

Example 30. Tom's utility function is given by $u(C) = \ln C$. Tom expects that there will be a recession next year with probability $\pi_b = 0.2$. If so, he will lose his job (and receive zero income). If not, then he keeps his job, which pays an annual salary of $100,000. An insurance firm offers to insure Tom against the risk of unemployment. Characterise an actuarially fair insurance plan. Suppose the insurer offers an insurance package, which pays Tom $3 if he loses his job, for each $1 packet of insurance purchased. How much insurance will Tom purchase?

Actuarially fair insurance satisfies:

$$\begin{aligned}
1 \cdot \pi_g - (r-1) \cdot \pi_b &= 0 \\
r &= \frac{\pi_g}{\pi_b} + 1 \\
&= \frac{0.8}{0.2} + 1 = 5
\end{aligned}$$

Hence, if the insurance package was fair, Joe would receive \$5 (if unemployed) for every \$1 packet of insurance purchased.

The marginal rate of substitution is given by:

$$\begin{aligned}
MRS &= -\frac{u'(C_b)}{u'(C_g)} \cdot \frac{\pi_b}{\pi_g} \\
&= -\frac{\frac{1}{C_b}}{\frac{1}{C_g}} \cdot \frac{\pi_b}{\pi_g} \\
&= -\frac{1}{4}\frac{C_g}{C_b}
\end{aligned}$$

The budget constraint is given by:

$$\begin{aligned}
\frac{1}{r-1}C_b + C_g &= \frac{1}{r-1}y_b + y_g \\
\frac{1}{2}C_b + C_g &= 100000
\end{aligned}$$

and so, the slope of the budget constraint is:

$$Slope\ BC = -\frac{1}{r-1} = -\frac{1}{2}$$

Then, the tangency condition implies:

$$\begin{aligned}
MRS &= slope\ BC \\
-\frac{1}{4}\frac{C_g}{C_b} &= -\frac{1}{2} \\
C_g &= 2C_b
\end{aligned}$$

Substituting this into the budget constraint, gives:

$$\begin{aligned}
\frac{1}{2}C_b + (2C_b) &= 100000 \\
C_b &= \$40000
\end{aligned}$$

and so:

$$\begin{aligned}
C_g &= 2C_b \\
&= \$80000
\end{aligned}$$

Tom pays \$20,000 in unemployment insurance. Since this pays out at a rate of 3 to 1, then if he finds himself unemployed, he receives a net payments of \$40,000 from the insurer (after payment for the insurance). $\square$

6.3.4 Optimal Investment Portfolios

In this section, we consider the question of how agents choose an investment portfolio amongst various risky instruments. When choosing amongst risky assets, the agent will be interested in:

- the expected return to the asset (μ)
- the risk associated with this return. (We usually measure risk using the standard deviation σ)

How do investors know the risk and expected return of assets? They might use historical data to estimate these. Let r_t be the return to the asset in period t, and suppose the agent has data on returns for the last T years. We can estimate:

$$\mu = \frac{1}{T}\sum_{t=1}^{T} r_t$$

$$\sigma = \sqrt{\frac{1}{T-1}\sum_{t=1}^{T}(r_t-\mu)^r}$$

Preferences over Risk and Return We assume that the investor has preferences over the levels of risk and return. We can represent these preferences (using indifference curves) on over the space of risk-return profiles. Consider a risk averse agent and an asset with risk and return as shown in the diagram below:

Assets in region II have a higher expected return and lower risk than asset A. Clearly, a risk averse agent strictly prefers these assets to A. Conversely, assets in region IV has a lower expected return and higher risk than asset A. A risk averse agent strictly prefers A to any asset in region IV. Hence, an indifference curve representing a risk averse agent's preferences over risk and return must be upward sloping (it must go through regions I and III). It should be apparent that risk-averse agents face a trade-off between accepting higher risk and receiving a higher expected payoff. The slope of the indifference curve represents the rate at which they are willing to make this trade-off.

We assume further that preferences over risk and return are convex —i.e. an agent prefers an asset with moderate risk and return to an asset with low risk but low returns or to an asset with high returns and high risk. If the agent faces an asset with low risk, he will be more willing to switch to a riskier asset with higher returns, than if the agent already faces high risk. Hence, the slope of the indifference curve is flat at first (when risk is low, more willing to accept more risk), but becomes increasingly steep as the level of risk increases. We have:

Figures/Uncertainty/lecture12f.pdf

For completeness, we characterise preferences over risk and return for risk neutral and risk loving agent. A risk neutral agent only cares about the expected return of an asset —he is indifferent between two assets with the same expected return, but different levels of risk. Hence, the indifference curves for a risk neutral agent are horizontal. The risk neutral agent prefers higher indifference curves to lower ones. We have:

Figures/Uncertainty/lecture12g.pdf

Return to the figure above which divides the risk-return space into quadrants. A risk-loving agent prefers assets with higher expected return, but also assets with higher levels of risk! Hence, the risk loving agent strictly prefers any asset in region III to asset A, and he prefers asset A to any asset in region I. Hence, a risk loving agent's indifference curve must be downward sloping (it must go through regions II and IV). Moreover, we assume his preferences over risk and return are convex. It follows that a risk loving agent's indifference map looks like:

Figures/Uncertainty/lecture12h.pdf

Efficient Portfolio Frontier Having characterised preferences, we now turn to the question of which risk-return combinations the investor can achieve. Suppose an investor can choose to invest his wealth in two assets A and B. The return to these asset r_A and r_B are uncertain (risky) prospects. The agent knows the expected returns (μ_A and μ_B) and the risks (σ_A any σ_B) associated with each asset. We assume that asset B has both higher risk ($\sigma_A < \sigma_B$) and higher expected return ($\mu_A < \mu_B$).

Figures/Uncertainty/lecture12i.pdf

The investor is not limited to choosing just between these two combinations of risk and return. Indeed, by dividing his portfolio between these two assets (in different proportions), the agent can generate any expected return between μ_A and μ_B. To see this more clearly, suppose a portfolio where the agent invests a fraction w_A of his wealth in Asset A, and a fraction w_B in Asset B. (Clearly we must have $w_A + w_B = 1$.) The return to the portfolio is given by $r_p = w_A r_A + w_B r_B$. What is the expected return to such a portfolio? We have:

$$\begin{aligned}
\mu_p &= E[r_p] \\
&= E[w_A r_A + w_B r_B] \\
&= w_A E[r_A] + w_B E[r_B] \\
&= w_A \mu_A + w_B \mu_B
\end{aligned}$$

which is a convex combination of the expected returns to each of the assets individually. Note, since $w_A = 1 - w_B$, this implies:

$$w_B = \frac{\mu_p - \mu_A}{\mu_B - \mu_A}$$

What is the risk associated with such a portfolio? We have:

$$\begin{aligned}
\sigma_p^2 &= Var\left[r_p\right] \\
&= Var\left[w_A r_A + w_B r_B\right] \\
&= Var\left[w_A r_A\right] + Var\left[w_B r_B\right] + 2Cov\left[w_A r_A, w_B r_B\right] \\
&= w_A^2 Var\left[r_A\right] + w_B^2 Var\left[r_B\right] + 2w_A w_B Cov\left[r_A, r_B\right] \\
&= w_A^2 \sigma_A^2 + w_B^2 \sigma_B^2 + 2w_A w_B \sigma_{AB} \\
&= w_A^2 \sigma_A^2 + w_B^2 \sigma_B^2 + 2w_A w_B \sigma_A \sigma_B \rho_{AB}
\end{aligned}$$

where $\sigma_{AB} = Cov\left[r_A, r_B\right]$ and $\rho_{AB} = Correl\left[r_A, r_B\right] = \frac{Cov[r_A, r_B]}{\sqrt{Var(r_A)}\sqrt{Var(r_B)}}$.

It should be clear that the risk associated with a portfolio depends upon the correlation between the returns to the assets in the portfolio. If the correlation is positive, then the level of risk will be higher, whilst if the correlation is negative, then the level of risk will be lower. (To see why, note that if correlation is positive, then if the return to asset A is below expectation, then it is likely that the return to asset B will also be below its expected level. Since the assets co-move, if the impact of a good/bad return on one asset is reinforced by a similarly good/bad return to the other. On the other hand, if the correlation is negative, then if the return to one asset is below expectation, then it is likely that the return to the other asset will be above expectation. This means that impact of a bad return on one asset is counteracted by a good return on the other. This reduces the overall risk.)

To make this relationship stark, we consider 3 cases: (i) perfect correlation ($\rho = 1$); (ii) no correlation ($\rho = 0$) and (iii) perfect negative correlation ($\rho = -1$).

Case 1: Perfect Correlation Suppose the assets are perfectly correlated, so that $\rho_{AB} = 1$. Then we have:

$$\begin{aligned}
\sigma_p^2 &= w_A^2 \sigma_A^2 + w_B^2 \sigma_B^2 + 2w_A w_B \sigma_A \sigma_B \\
&= \left(w_A \sigma_A + w_B \sigma_B\right)^2 \\
\sigma_p &= w_A \sigma_A + w_B \sigma_B \qquad (6.10)
\end{aligned}$$

which is simply a convex combination of the risks associated with each asset individually. Substituting (6.3.4) into (6.10), we have:

$$\sigma_p = \sigma_A + \frac{\sigma_B - \sigma_A}{\mu_B - \mu_A}\left(\mu_p - \mu_A\right)$$

which gives the relationship between portfolio return and portfolio risk. This is a linear equation passing through the points (μ_A, σ_A) and (μ_B, σ_B). We can plot the risk-return profiles for these various portfolios:

The line in the above diagram represents all the possible risk-return profiles that can be achieved by constructing a portfolio containing both assets A and B. Point C represents the risk-return pair for a portfolio containing 25% asset B and 75% asset A. Point D corresponds to the portfolio which is equally weighted between the two assets. Point E represents a portfolio compromising 90% of asset B and 10% of asset A. It should be clear that the greater the share of asset B in the portfolio, the larger is both the expected return and the risk.

Case 2: No Correlation Suppose the assets are not correlated, so that $\rho_{AB} = 0$. Then we have:

$$\begin{aligned} \sigma_p^2 &= w_A^2\sigma_A^2 + w_B^2\sigma_B^2 \\ \sigma_p &= \sqrt{w_A^2\sigma_A^2 + w_B^2\sigma_B^2} \end{aligned} \tag{6.11}$$

Substituting (6.3.4) into (6.11) gives an expression for the relationship between portfolio risk and expected return. We plot the risk-return profiles in the diagram below. The dotted line represents the portfolio where assets are perfectly correlated. We note that for each level of return, the portfolio risk (with no correlation) is smaller than in the perfect correlation case.

Case 3: Perfect Negative Correlation Suppose the assets are perfectly negatively correlated, so that $\rho_{AB} = -1$. Then we have:

$$\begin{aligned} \sigma_p^2 &= w_A^2\sigma_A^2 + w_B^2\sigma_B^2 - 2w_Aw_B\sigma_A\sigma_B \\ &= (w_A\sigma_A - w_B\sigma_B)^2 \\ \sigma_p &= |w_A\sigma_A - w_B\sigma_B| \end{aligned}$$

We plot the risk-return combinations in the usual way.

__Figures/Uncertainty/lecture12k.pdf

The risk-return frontier is given by the solid line. We note curiously, that although both asset A and asset B are risky prospects, there is a portfolio consisting of a specific combination of the assets for which there is zero risk. For such a portfolio, the effect of a positive variance to one of the assets will be exactly counteracted by a negative variance to the other asset.

The following points are worth noting:

Figures/Uncertainty/lecture12l.pdf

- A common objection to portfolios with negatively correlated assets is that positive earnings by one asset will be counteracted by negative earnings by the other asset. If so, how can a portfolio with negatively correlated assets make a positive return? This line of argument is fallacious. Each asset has a positive expected return. The correlation between the assets concerns whether the assets perform better or worse than their expected value. In the case of two negatively correlated assets, if one does slightly worse than its expected value (but still possibly makes a positive return), then the other asset will do slightly better than its expected value. Hence, the expected return to the investment will be positive, however the variance will be lower, because deviations in the return to one asset will be balanced by contrary deviations to the other asset.

- If assets are negatively correlated, the investor can achieve a higher expected at a lower level of risk. It is not always the case that the agent must accept higher risk in order to receive a higher return!

Summary In the above discussion, we characterised the efficient frontier for 3 special cases of assets —those in which there is perfect correlation (both positive and negative) and those for which there is no correlation. For generic assets (with some, but not perfect correlation), the frontier must between these. For example, if the correlation between assets A and B is $\rho = 0.3$, then the efficient frontier will lie between the $\rho = 0$ and $\rho = 1$ frontiers. Similarly, if the correlation between A and B is $\rho = -0.5$, then the efficient frontier will lie between the $\rho = -1$ and $\rho = 0$. We illustrate this in the following diagram (where the unbroken lines represent the base cases of $\rho = -1, \rho = 0$ and $\rho = 1$, whilst the dashed lines represent the frontiers for $\rho = -0.5$ and $\rho = 0.3$).

Optimal Portfolio Choice In the first section, we characterised the investors preferences over combinations of risk and return. In the next section, for any two assets, we characterised the risk-return frontier for various portfolio combinations of two assets. (We can think of the risk-return frontier as the "budget constraint" —it is the set of risk-return combinations that the investor can achieve given the available assets). The optimal portfolio is the portfolio on the efficient frontier which is on the highest indifference curve. As usual, the optimal portfolio must satisfy the tangency condition. We illustrate this diagrammatically:

Figures/Uncertainty/lecture12m.pdf

The above example illustrates the optimal choice of a risk averse agent choosing a portfolio amongst two assets which are somewhat negatively correlated. The optimal portfolio lines $\frac{2}{5}$ of the way along the frontier (from A to B) and so the optimal portfolio consists of 40%

asset B and 60% asset A. We note that, since the assets are negatively correlated, the agent is able to achieve a significantly larger expected return (than the return to asset A) at the cost of marginally higher risk.

It should be clear that risk neutral and risk loving agents will always choose to invest their entire wealth in asset B. (Recall, risk neutral agents are indifferent to risk, and maximise their utility by choosing the asset with the highest expected return. Recall further that risk loving agents like risk, and might even accept a lower expected return as compensation for higher risk.)

Part II

Intermediate Microeconomics

Chapter 8

Mathematical Tools

8.1 Multivariate Calculus

Let $x = (x_1, ..., x_n)$ be an n-ary vector, and consider the functions $y = f(x)$. For example, consider the bi-variate function: $f(x_1, x_2) = x_1^2 - x_1 \ln x_2$.

In single variable calculus, the derivative is constructed by taking small changes in the input (x) and asking how quickly the output $f(x)$ changed in response. Derivatives in the multivariate setting are defined analogously: we take a small change in *any* one of the inputs (x_i) and ask how quickly the output changes, assuming the other input is held fixed at its original level. We call this a **partial derivative**, which we denote $\frac{\partial f}{\partial x_i}$ or $f_i(x)$ where x_i is the variable that is allowed to change.

Naturally, rather than having a single derivative, we now have n partial derivatives —one for each component of the vector x. Formally, we have:

$$\frac{\partial f}{\partial x_i} = \lim_{h \to 0} \frac{f(x_1, ..., x_{i-1}, x_i + h, x_{i+1}, ..., x_n) - f(x_1, ..., x_i, ..., x_n)}{h}$$

We can compute partial derivatives in much the same way we computed derivatives in the single variable case, using the rules we noted above. When doing so, we treat all variables other than the one that we are taking the derivative with respect to as if they were constants.

Example 31. Consider the bivariate function $f(x_1, x_2) = x_1^2 + x_1 \ln x_2$. The partial derivatives are:

$$\begin{aligned} f_1(x_1, x_2) &= \frac{\partial f}{\partial x_1} = 2x_1 + \ln x_2 \\ f_2(x_1, x_2) &= \frac{\partial f}{\partial x_2} = \frac{x_1}{x_2} \end{aligned}$$

When computing the partial derivative w.r.t. x_1 we treat x_2 as if it were constant. So the function is effectively of the form: $x_1^2 + ax_1$, where a stands in for $\ln x_2$ (which we treat as a constant). Clearly the derivative will be of the form $2x_1 + a$. Similarly, when computing the partial derivative w.r.t. x_2, we treat x_1 as if it were constant. So the function is effectively of the form: $b + c \ln x_2$, where b and c stand in for x_1^2 and x_1 (which we treat as constants). Clearly the derivative will be of the form $\frac{c}{x_2}$.

We describe the derivatives as *partial derivatives* because they only show the (partial) effect of a change in the output —that coming from a change in a single input alone. In economics, often-times, several inputs will change simultaneously. The *total derivative* accounts for all of these changes, whereas the partial derivative takes each in isolation.

Just as we can find higher-order (second, third,...) derivatives of single-variable functions, so can we of multi-variate functions. In single variable calculus, the second derivative is the derivative of the derivative. For multi-variate, we know that that there are n first order partial derivatives to consider —and each of these can be differentiated with respect of each of n variables. I.e. we can take the (partial) derivative of $\frac{\partial f}{\partial x_1}$ w.r.t. x_1, or w.r.t. x_2,...., or w.r.t. x_n. Similarly, and we can take the (partial) derivative of $\frac{\partial f}{\partial x_n}$ w.r.t. x_1 or w.r.t. x_2,..., or w.r.t. x_n. Thus, there are n^2 second order partial derivatives:

$$\begin{aligned}
f_{11}(x_1, x_2) &= \frac{\partial^2 f}{\partial x_1^2} = \frac{\partial}{\partial x_1}\left(\frac{\partial f}{\partial x_1}\right) \\
f_{22}(x_1, x_2) &= \frac{\partial^2 f}{\partial x_2^2} = \frac{\partial}{\partial x_2}\left(\frac{\partial f}{\partial x_2}\right) \\
f_{12}(x_1, x_2) &= \frac{\partial^2 f}{\partial x_2 \partial x_1} = \frac{\partial}{\partial x_2}\left(\frac{\partial f}{\partial x_1}\right) \\
f_{21}(x_1, x_2) &= \frac{\partial^2 f}{\partial x_1 \partial x_2} = \frac{\partial}{\partial x_1}\left(\frac{\partial f}{\partial x_2}\right)
\end{aligned}$$

We refer to the first two as 'own' second-order partial derivatives, and the second two as 'cross' second-order partial derivatives. The reason for the nomenclature should be obvious.

[Young's Theorem] If the function f is 'sufficiently smooth' (technically, if f is twice continuously differentiable), then $f_{ij} = f_{ji}$. i.e.

$$\frac{\partial^2 f}{\partial x_i \partial x_j} = \frac{\partial^2 f}{\partial x_j \partial x_i}$$

Young's Theorem states that the order of differentiation does not matter. If we differentiate f w.r.t x_i first and then x_j, we get the same second-order partial derivative, as if we differentiated w.r.t. x_j first and then x_i.

Example 32. Continuing the previous example, we have: $f_{11} = 2$, $f_{12} = \frac{1}{x_2} = f_{21}$, and $f_{22} = -\frac{x_1}{x_2^2}$.

Exercise 6. For each of the following functions, find all of the first-order and second-order partial derivatives:

1. $f(x_1, x_2) = \frac{2x_1}{\ln x_2} + x_1x_2 + 5$
2. $g(x, y, z) = x^2 + 2xy - y^2 + 2\frac{y}{z} + z$

8.2 Unconstrained Optimization

We can find local optima of multivariate functions in much the same way we do with univariate functions. For a single variable function, we know that the derivative is typically zero at the optimum. Intuitively, if the derivative was non-zero, then the function can be increased (or decreased) further, by either taking a small change in the input. Thus, to be at the maximum, it must be that the derivative is zero —there are no further opportunities for improvements. In the multivariate case, the same logic applies, except we now need *all* of the partial derivatives to be zero. This guarantees that the function cannot be improved by making a small change to *any* of the inputs.

Definition 7. A *critical point* of a function $f(x_1, ..., x_n)$ is a point $(x_1, ..., x_n)$ at which $\frac{\partial f}{\partial x_i} = 0$ for each $i \in \{1, ..., n\}$.

We say the function satisfies the first order conditions at its critical points. As before, the critical points are candidate optimizers, but not every critical point is an optimizer, and not every optimizer s a critical point.

Lemma 6. *Let $f(x_1, x_2)$ be a bivariate function. Suppose (x_1^*, x_2^*) is a critical point of f. Then:*

- *(x_1^*, x_2^*) is a local minimum if: (i) $f_{11} > 0$ and $f_{22} > 0$, and (ii) $f_{11}f_{22} - f_{12}f_{21} > 0$, where the second-order partial derivatives are all evaluated at (x_1^*, x_2^*).*
- *(x_1^*, x_2^*) is a local maximum if: (i) $f_{11} < 0$ and $f_{22} < 0$, and (ii) $f_{11}f_{22} - f_{12}f_{21} > 0$, where the second-order partial derivatives are all evaluated at (x_1^*, x_2^*).*

Example 33. A firm can produce a quantity of output q by combining k units of capital and l units of labour. The firm's production function is $q = k^{0.25}l^{0.5}$. The firm can hire as much capital and labour as it desires. The per unit cost of capital is 10 and the per worker wage is 5. Firms can sell any quantity of output produced on the market at price 20. How much of each input should it hire? What quantity of output should it produce? What is the firm's optimal profit?

The firm's profit is $\pi = 20q - 10k - 5l$. Since the firm's output is determined by its production function, we have $\pi = 20k^{0.25}l^{0.5} - 10k - 5l$. Thus, the firm's profit is a function of two variables: k and l. The firm's problem is:

$$\max_{k,l} 20k^{0.25}l^{0.5} - 10k - 5l$$

The first order condtions are:

$$\begin{aligned}
\frac{\partial \pi}{\partial k} &= 5k^{-0.75}l^{0.5} - 10 = 0 \\
\frac{\partial \pi}{\partial k} &= 10k^{0.25}l^{-0.5} - 5 = 0
\end{aligned}$$

Combining these gives: $\frac{1}{2}\frac{l}{k} = 2$, and so $l = 4k$. Substituting this into the first condition gives:

$$\begin{aligned}
5k^{-}0.75(4k)^{0.5} &= 10 \\
2k^{-0.25} &= 2 \\
k^* &= 1
\end{aligned}$$

Then, since $l^* = 4k^*$, we have $l^* = 4$.

Next, we must check that the second-order conditions hold so that we have, indeed, found a maximum. Below are each of the second derivatives, evaluated at the critical point:

$$\begin{aligned}
\pi_{kk} &= -3.75k^{-1.75}l^{0.5} = -3.75(1)^{-1.75}(4)^{0.5} = -7.5 \\
\pi_{ll} &= -5k^{0.25}l^{-1.5} = -5(1)^{0.25}(4)^{-1.5} = -0.625 \\
\pi_{kl} &= 2.5k^{-0.75}l^{-0.5} = 2.5(1)^{-0.75}(4)^{-0.5} = 1.25 \\
\pi_{lk} &= 2.5k^{-0.75}l^{-0.5} = 2.5(1)^{-0.75}(4)^{-0.5} = 1.25
\end{aligned}$$

Hence, we verify that $\pi_{kk} < 0$, $\pi_{ll} < 0$ and $\pi_{kk}\pi_{ll} - \pi_{kl}\pi_{lk} = -7.5 \times -0.625 - 1.25 \times 1.25 = 3.125 > 0$. The second order conditions are satisfied. □

8.3 Constrained Optimization

Fundamentally, economics is the study of choice. Agents only have genuine choices to make when they face constraints that prevent them from getting everything they want. (Such constraints might include budget constraints, time constraints, technological constraints, institutional constraints etc.)

Example 34. Two canonical examples of constrained optimization problems:

1. *Utility Maximization subject to a budget constraint.* There is an agent who must choose quantities of goods x and y to purchase to maximize her utility subject to those goods being affordable. We have:

$$\max_{x,y} u(x,y) \text{ s.t. } p_x x + p_y y = I$$

2. *Profit Maximization subject to a technological constraint.* There is a firm that must choose quantities of inputs k and l to hire, and the quantity of output to produce q, subject to the output being technologically feasible given the inputs. We have:

$$\max_{q,k,l} pq - rk - wl \text{ s.t. } q = f(k,l)$$

8.3.1 The Method of Lagrange Multipliers

Consider a generic problem:

$$\max_{x_1,..,x_n} f(x_1,...,x_n) \text{ s.t. } g(x_1,...,x_n) = 0$$

$f(x_1,...,x_n)$ is the objective function, whilst $g(x_1,...,x_n)$ is the constraint function. (It should be clear that we can express each of the constraints in the above example in the form $g(\cdot) = 0$.)

Absent the constraint, we know what to do —find the critical points of f by setting all the first-order partial derivatives to zero. But these critical points may not satisfy the constraints (i.e. they may not be feasible). Our solution is to incorporate the constraint into the objective function, so that it is already taken into account when we take first order conditions.

Consider a new function:

$$\mathcal{L}(x_1,...,x_n,\lambda) = f(x_1,..,x_n) - \lambda\left(g(x_1,...,X_n)\right)$$

where λ is a new variable called the **Lagrange Multiplier**. The new function $\mathcal{L}$ is called the **Lagrangian**. It turns out that to find the constrained maximizer of the original objective function, it suffices to find the *unconstrained* maximizer of the Lagrangian. To see why, notice that the Lagrangian has the constraint incorporated into it. In fact, it is simply the objective function, augmented with a penalty whenever the constraint is not satisfied. The size of the penalty is given by the Lagrange multiplier.

Notice that the penalty term will be zero whenever the constraint is satisfied, since $g(x_1,...,x_n) = 0$. This ensures that, when the constraint is satisfied, the Lagrangian function simply coincides with the true objective function, so that maximizing the Lagrangian and constrained maximizing the objective function amount to the same thing. If the constraint is not satisfied, then the correction term $(\lambda g(x_1,...,x_n))$ will be positive, which causes the

Lagrangian to take a lower value than the objective. If we seek to maximize the Lagrangian then, we will be deterred from settling on values that violate the constraint.

But how do we choose the size of the penalty, λ? We let the calculus do it for us! Notice that $\frac{\partial \mathcal{L}}{\partial \lambda} = -g(x_1, ..., x_n)$. Hence, if we take the first order condition with respect to λ, we effectively get the condition $g(x_1, ..., x_n) = 0$, which guarantees that the constraint will be satisfied.

In summary, to do constrained optimization, we:

1. Form the Lagrangian function, which incorporates the constraint, in a way that penalizes the maximization whenever the constraint is violated.
2. Find the unconstrained optimum of the Lagrangian. To do so, we take first order conditions with respect to the original variables x_1,...,x_n, as well as the penalty variable λ (i.e. the Lagrange multiplier).

Example 35. Consider the following problem:

$$\max_{x_1,x_2} -x_1^2 - x_2^2 \,\text{s.t.}\, 2x_1 + x_2 = 5$$

Clearly, the unconstrained maximum is achieved when $\hat{x_1} = 0 = \hat{x_2}$ —you can easily check this by taking first order conditions —but this obviously does not satisfy the constraint.

The Lagrangian is:

$$\mathcal{L} = -x_1^2 - x_2^2 - \lambda(2x_1 + x_2 - 5)$$

Taking first order conditions, we have:

$$\frac{\partial \mathcal{L}}{\partial x_1} = -2x_1 - 2\lambda = 0$$
$$\frac{\partial \mathcal{L}}{\partial x_2} = -2x_2 - \lambda = 0$$
$$\frac{\partial \mathcal{L}}{\partial \lambda} = 5 - 2x_1 - x_2 = 0$$

We have a system of 3 equations in 3 unknowns. From the first two equations, we get:

$$\lambda = -x_1 = -2x_2$$

which implies that $x_1 = 2x_2$. Substituting this into the third equation give $5 - 2(2x_2) - x_2 = 0$, which implies $x_2 = 1$. Then, since $x + 1 = 2x_2$, it must be that $x_1 = 2$. Finally, since $\lambda = -x_1 = -2x_2$, we have $\lambda = -2$. The constrained optimum is $(x_1^*, x_2^*) = (2, 1)$.

Notice that the penalty $\lambda^* = -2$ is negative. Does this make sense? To get from the constrained optimum to the unconstrained optimum, both x_1 and x_2 would need to decrease, which would cause the constraint $(2x_1 + x_2 - 5)$ to become negative. If λ were positive, the penalty would be negative, and since the penalty is subtracted, this would cause the Lagrangian to be *larger* than the objective. Clearly this cannot be the case. This explains why $\lambda < 0$.

Example 36 (Firm's Choice cont.)**.** Return to the Example about the firm choosing which inputs to hire. Recall, the firm faces a production technology $q = k^{0.25}l^{0.5}$ and faces output price p and input prices r and w. The firm must choose the quantity of output to produce q and the quantities of each input k and l.

We can reformulate the firm's problem as a constrained optimization problem:

$$max_{q,k,l} pq - rk - wl \, \text{s.t.} \, k^{0.25}l^{0.5} = q$$

Before going any further, let us note crucially. There are several 'variables' in this problem. The variables q, k and l are *choice variables* —the agent chooses these to maximize his objective. The variables p, r and w are environmental variables —they describe market prices which the agent takes as given; he has no ability to choose these. We will often refer to these as *parameters* of the model. When solving for the optimum, you should find an expression for each of the choice variables in terms of the parameters.

The Lagrangian is:

$$\mathcal{L} = pq - rk - wl - \lambda\left(q - k^{0.25}l^{0.5}\right)$$

The first order conditions are:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial q} &= p - \lambda = 0 \\
\frac{\partial \mathcal{L}}{\partial k} &= -r + \lambda\left(0.25k^{-0.75}l^{0.5}\right) = 0 \\
\frac{\partial \mathcal{L}}{\partial l} &= -w + \lambda\left(0.5k^{0.25}l^{-0.5}\right) = 0 \\
\frac{\partial \mathcal{L}}{\partial \lambda} &= k^{0.25}l^{0.5} - q = 0
\end{aligned}$$

Using the first condition, we have $\lambda = p$. Substituting this into the second and third equations gives:

$$\begin{aligned}
\left(0.25k^{-0.75}l^{0.5}\right) &= \frac{r}{p} \\
\left(0.5k^{0.25}l^{-0.5}\right) &= \frac{w}{p}
\end{aligned}$$

This is a system of two equations in two variables, identical to what we had in the previous example. It follows that:

$$\begin{aligned}
\frac{0.25k^{-0.75}l^{0.5}}{0.5k^{0.25}l^{-0.5}} &= \frac{r}{w} \\
\frac{l}{2k} &= \frac{r}{w} \\
l = &= \frac{2r}{w}k
\end{aligned}$$

Now, substituting this back into equation (2) gives:

$$\begin{aligned}\left(0.25k^{-0.75}\left[\frac{2r}{w}k\right]^{0.5}\right) &= \frac{r}{p}\\ k^{-0.25} &= \frac{r}{p}\cdot\sqrt{\frac{w}{2r}}\\ k^* &= \frac{4p^4}{r^2w^2}\end{aligned}$$

Additionally, since $l^* = \frac{2r}{w}k^2$, then $l^* = \frac{8p}{rw^3}$.

Finally, substitute k^* and l^* into the last equation.

$$\begin{aligned}k^{0.25}l^{0.5} &= q\\ \left(\frac{4p^4}{r^2w^2}\right)^{0.25}\left(\frac{8p^4}{rw^3}\right)^{0.5} &= q\\ q &= \frac{4p^3}{rw^2}\end{aligned}$$

Hence the solution is $q^* = \frac{4p^3}{rw^2}$, $k^* = \frac{4p^4}{r^2w^2}$, $l^* = \frac{8p}{rw^3}$ and $\lambda^* = p$.

8.3.2 The Envelope Theorem

In the last example, the decision maker's optimal choice depended straight-forwardly on the parameters. In general, let $x = (x_1, ..., x_n)$ be the choice variables, and let $\alpha = (\alpha_1, ..., \alpha_p)$ be the parameters. The agent's generic decision problem is:

$$\max_x f(x;\alpha)\, s.t.\, g(x;\alpha) = 0$$

This setup allows for both the objective and constraint functions to depend on the parameters α. In the profit maximization problem, the parameters were $\alpha = (r, w, p)$, which entered the objective function, but not the constraint. In the utility maximization, the parameters were $\alpha = (p_x, p_y, I)$ which entered the constraint function, but not the objective. In general, the parameters may potentially affect either/both.

The Lagrangian is:

$$\mathcal{L}(x,\lambda;\alpha) = f(x;\alpha) - \lambda\left(g(x;\alpha)\right)$$

The first order conditions are:

$$\begin{aligned}\frac{\partial\mathcal{L}}{\partial x_i} &= \frac{\partial f(x;\alpha)}{\partial x_i} - \lambda\frac{\partial g(x;\alpha)}{\partial x_i} = 0 \text{ for each } i\\ \frac{\partial\mathcal{L}}{\partial\lambda} &= -g(x;\alpha) = 0\end{aligned}$$

Let $x^*(\alpha) = (x_1^*(\alpha), ..., x_n^*(\alpha))$ and $\lambda^*(\alpha)$ solve the first order conditions. Notice that these will generally be functions of the parameters.

The *value function* gives the value that the objective function achieves at the optimum. We have:

$$V(\alpha) = f(x^*(\alpha); \alpha)$$

Since the optimizers x^* are functions of the parameters, so will the value function.

Suppose there is a change in one of the parameters (e.g. suppose the wage increases). We would like to know how this affects the agent's value function. Suppose α_j changes. This will affect $V(\alpha)$ in two ways:

1. *Direct Effect*: To the extent that the objective function $f(x; \alpha)$ is responsive to α_j, when α_j changes, this will directly cause f (and thus V) to change.
2. *Indirect Effect*: When α_j changes, each of the optimizers $x_i^*(\alpha)$ will likely change, and this will indirectly cause V to change.

We see this in the following equation:

$$\frac{\partial V(\alpha)}{\partial \alpha_j} = \underbrace{\frac{\partial f(x^*(\alpha); \alpha)}{\partial \alpha_j}}_{\text{Direct Effect}} + \underbrace{\sum_{i=1}^{n} \frac{\partial f(x^*(\alpha); \alpha)}{x_i} \frac{\partial x_i^*(\alpha)}{\partial \alpha_j}}_{\text{Indirect Effect}}$$

Similarly, we know that the optima must satisfy the constraint, so that $g(x^*(\alpha); \alpha) = 0$. Totally differentiating w.r.t α_j gives:

$$\frac{\partial g(x^*(\alpha); \alpha)}{\partial \alpha_j} + \sum_{i=1}^{n} \frac{\partial g(x^*(\alpha); \alpha)}{x_i} \frac{\partial x_i^*(\alpha)}{\partial \alpha_j} = 0$$

where again, the first term is the direct effect of a changing α_j on the constraint (e.g. if the price of a good increases, fewer bundles are affordable), and the second term is the indirect effect of a changing α via its effect on the x_i^*s.

Now, since the first order conditions are satisfied at the optimum, we know that:

$$\frac{\partial f(x^*(\alpha); \alpha)}{\partial x_i} = \lambda^*(\alpha) \frac{\partial g(x^*(\alpha); \alpha)}{\partial x_i}$$

Substituting this into the expression for $\frac{\partial V(\alpha)}{\partial \alpha_j}$ gives:

$$\begin{aligned}
\frac{\partial V(\alpha)}{\partial \alpha_j} &= \frac{\partial f(x^*(\alpha); \alpha)}{\partial \alpha_j} + \lambda^*(\alpha) \sum_{i=1}^{n} \frac{\partial g(x^*(\alpha); \alpha)}{x_i} \frac{\partial x_i^*(\alpha)}{\partial \alpha_j} \\
&= \frac{\partial f(x^*(\alpha); \alpha)}{\partial \alpha_j} - \lambda^*(\alpha) \frac{\partial g(x^*(\alpha); \alpha)}{\partial \alpha_j} \\
&= \frac{\partial \mathcal{L}(x^*, \lambda^*; \alpha)}{\partial \alpha_j}
\end{aligned}$$

where the second line makes use of the result from the constraint that:

$$\frac{\partial g(x^*(\alpha);\alpha)}{\partial \alpha_j} = -\sum_{i=1}^{n} \frac{\partial g(x^*(\alpha);\alpha)}{x_i}\frac{\partial x_i^*(\alpha)}{\partial \alpha_j}$$

This result is known as the *Envelope Theorem.* It states that when a parameter α_j changes, its effect on the the value function only depends on its direct effect on the objective function $\frac{\partial f(x^*(\alpha);\alpha)}{\partial \alpha_j}$ and its direct effect on the constraint function $\frac{\partial g(x^*(\alpha);\alpha)}{\partial \alpha_j}$. The indirect effects (of the change in α_j causing the x^*'s to change, and those in turn causing V to change) all net out to zero. As we will see during this course, the Envelope Theorem is one of the most important and useful results in the study of economics.

One important insight that comes out of the Envelope Theorem is that it gives meaning to the Lagrange multiplier λ^*. To see this, modify the problem slightly: suppose we can write the constraint function as: $h(x;\alpha) = b$. (There is nothing fishy going on here; previously, we would have written $g(x;\alpha,b) = h(x;\alpha) - b = 0$.) For example, in the utility maximization problem, the constraint is $p_x x + p_y y = I$, and we can think of $\alpha = (p_x, p_y)$ and $b = I$. Then, applying the Envelope Theorem:

$$\begin{aligned}\frac{\partial V(\alpha,b)}{\partial b} &= \frac{\partial \mathcal{L}(x^*(\alpha,b),\lambda^*(\alpha,b);\alpha,b)}{\partial b}\\ &= \frac{\partial}{\partial b}\left[f(x^*(\alpha,b);\alpha) - \lambda^*(\alpha,b)(g(x^*(\alpha,b);\alpha) - b)\right]\\ &= \lambda^*(\alpha,b)\end{aligned}$$

Hence, the equilibrium value of the Lagrange multiplier can be interpreted as the amount by which the value function would increase if the constrained were 'loosened' by one unit. (E.g. in the utility maximization problem, it tells us how much more utility the agent could generate if he had one extra dollar to spend.) The Lagrange multiplier, thus, captures the *opportunity cost* or *shadow price* of the agent being subject to the constraint.

Chapter 9

Producer Theory

9.1 Technology - the Production function

In this chapter, we study decision making by firms. We identify firms by their production technology —which describes the relationship between the inputs consumed in the production process and the output generated. Let $x = (x_1, ..., x_n)$ be a vector of n inputs, and let y be the quantity of output produced. A **Production Function** describes the maximum amount of the output which can be produced, given a set of inputs, if they are used in their most efficient way. We have:

$$y = f(x)$$

To build a theory that can make predictions, we need to put some structure on the behavior of production functions. We will add two assumptions. Let $x \geq x'$ indicate that $x_i \geq x'_i$ for each $i = 1, ..., n$.

Definition 8. The production technology is **monotone** if $x \geq x'$ implies $f(x) \geq f(x')$.

A production function is monotone if, increasing all inputs causes output to increase as well. An implication of monotonicity is that marginal products are non-negative. We have:

$$MP_i = \frac{\partial f(x)}{\partial x_i} \geq 0 \quad \forall i$$

Definition 9. The production technology is **convex** if

$$f(\lambda x + (1 - \lambda)x') \geq \lambda f(x) + (1 - \lambda)f(x') \;\; forall \lambda \in [0, 1]$$

A production function is convex if, the average of two input bundle can produce at least as much output as the least productive of the two original bundles. It states that the production function favors bundles that combine inputs in moderate quantities, rather than having large

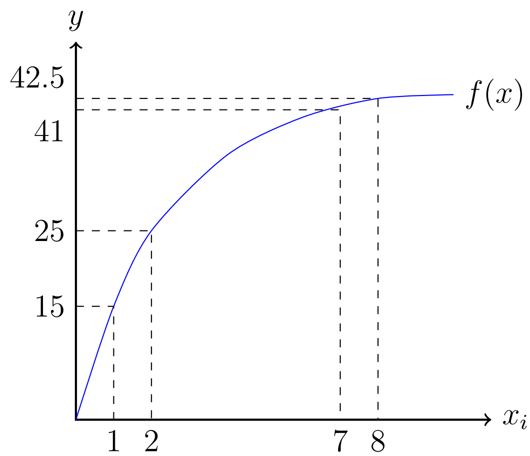

Figure 9.1: Law of Diminishing Marginal Returns

amounts of some inputs and small amounts of others. An implication of convexity is that the **law of diminishing marginal returns** —as we add more units of an input to the production process, holding constant the level of all other inputs, output increases, but at a smaller and smaller rate. Formally:

$$\frac{\partial MP_i(x)}{\partial x_i} = \frac{\partial^2 f(x)}{\partial x_i^2} < 0$$

Diminishing marginal returns implies that we cannot continue to increase output at a rapid rate by simply increasing the level of one input, ignoring the others. (For example, we cannot produce enough food to feed the world by simply adding more and more water to a single flower pot.) We illustrate the notion of diminishing returns in Figure 9.1. Notice that when the quantity of input x_i used in the production process increases from 1 to 2, the increase in output is significantly larger than the corresponding increase in output when the quantity of input x_i increases from 7 to 8. (The quantity of x_j is held constant throughout.)

We can represent a production technology using isoquants. An **isoquant** connects all input bundles that produce the same quantity of output. There are naturally many isoquants, one for each feasible output level.

Monotonicity and Convexity imply that isoquants satisfy the following properties:

- Isoquants are 'downward sloping'. An increase in the quantity of input i requires a decrease in the quantity of some other good (j, say) to ensure that overall output is unchanged. (Follows from montonicity.)
- Isoquants cannot intersect.
- 'Higher' isoquants (i.e. those further from the origin) represent input bundles producing larger outputs. (Again, follows by monotonicity.)

- Isoquants are 'convex' to the origin. (Follows by convexity.)

Definition 10. The Marginal Rate of Technical Substitution of input j with respect to input i, denoted $MRTS_{ij}$ is the quantity by which the firm must reduce input j after increasing input i by one unit, in order to keep output unchanged, holding the quantity of all other inputs constant.

$$MRTS_{ij} = -\frac{MP_i(x)}{MP_j(x)} = -\frac{\partial f(x) / \partial x_i}{\partial f(x) / \partial x_j}$$

To see that the formula must be true, consider a change in the quantity of input i, and suppose the Firm changes its utilization of input j to ensure output is unchanged, holding quantities of all other inputs constant. We can write $x_j(x_i)$ to capture that the quantity of input j is pinned down by the quantity of input i. Along an indifference curve, we know that: $f(x) = c$. Totally differentiating w.r.t x_i gives:

$$\begin{aligned} \frac{\partial f(x)}{\partial x_i} + \frac{\partial f(x)}{\partial x_j}\frac{dx_j(x_i)}{dx_i} &= 0 \\ \frac{dx_j(x_i)}{dx_i} &= -\frac{\partial f(x) / \partial x_i}{\partial f(x) / \partial x_i} \\ MRTS_{ij} &= \frac{MP_i}{MP_j} \end{aligned}$$

We also typically assume that the production function exhibits **decreasing returns to scale.** This means that if we increase the level of all inputs by some proportion then output will increase, but by a proportion less than the increase in inputs. I.e. if we double (or treble) the level of each input, output will increase, but will not double (or treble). Mathematically, we have:

$$f(\lambda x) < \lambda f(x) \quad \text{for any } \lambda > 0$$

We refer top this property as **diseconomies of scale**. (Traditionally we assume that up to a particular point, when a firm increases its scale, it can produce goods more efficiently by employing more specialised inputs and exploiting efficiencies that arise when producing goods in bulk. If so, the technology exhibits *increasing returns to scale* or *economies of scale.* However, economies of scale cannot continue to persist forever. At some point, the size of a firm becomes so large, that production becomes efficient - perhaps due to the growth of sinusoids. Hence - eventually - we assume that the technology will exhibit diseconomies of scale.

Comments:

1. When considering the firm's production decisions, we distinguish between the short-run and the long-run. The **short-run** is defined as the period short enough that the firm cannot vary all of the inputs used in the production process. (e.g. the size of a factory assembly line cannot be changed over-night. In the short run, the size of this

capital input is fixed.) The **long-run** is defined as the period long enough that the firm can vary all of its inputs. (e.g. continuing with the previous example, given sufficient time, the firm can increase the size of its assembly line to meet growing demand, or downsize a plant if needed.)

2. It is important to distinguish the concepts of *diminishing marginal product* and *diseconomies of scale*. Diminishing marginal product is a short run phenomenon. It says that if increasing quantities of a variable input are added to a fixed input, the marginal contribution of each additional variable input falls. Diseconomies of scale on the other hand, is a long run phenomenon. It says that even if we scale up all inputs in proportion, output wont increase by as much.

9.2 Cost Minimization

Our theory of the firm has two components. We must determine: (i) the quantity of output that the firm should produces, and (ii) the quantities of inputs that the firm should utilize in the production process. Ultimately we assume that firms are profit maximizers, and will produce the quantity that maximizes profit.

But there are multiple input combinations that produce this optimal output. Which combination should the firm choose? When contemplating producing *any* quantity of output y, we assume that firms choose the input combination that minimizes the cost of producing that output. To be clear —we are not saying that the firm's ultimate goal is to minimize costs; their goal is to maximize profits. Cost minimization is simply the criterion to choose amongst input combinations that achieve the same output. The cost minimizing input combination is the one that produces the desired output in the most efficient way.

Suppose a firm wishes to produce y units of output. Let $w = (w_1, ..., w_n)$ be a vector of input prices. We want to find the input bundle which minimises the cost of producing quantity y of output. (This amounts to find the the bundle on the isoquant corresponding to output quantity y, which lies on the lowest achievable isocost.) Formally, we have: $\min_x w \cdot x = \sum_i w_i x_i$ s.t. $y = f(x_1, ..., x_n)$.

The Lagrangian is:

$$\mathcal{L} = \sum_i w_i x_i - \lambda(f(x_1, .., x_n) - y)$$

The first order conditions are:

$$\frac{\partial \mathcal{L}}{\partial x_i} = w_i - \lambda \frac{\partial f(x)}{\partial x_i} = 0 \qquad \forall i$$
$$\frac{\partial \mathcal{L}}{\partial \lambda} = y - f(x_1, \ldots, x_n) = 0$$

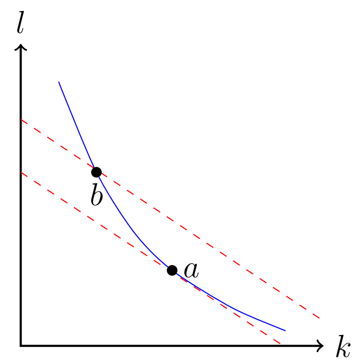

Figure 9.2: Optimal Input Choice

Applying the first equation to different inputs, we have:

$$\lambda = \frac{w_i}{MP_i} = \frac{w_j}{MP_j} \tag{9.1}$$

Two important insights:

1. The optimal input bundle must lie on the isoquant which corresponds to the quantity y which we seek to produce.(Why? If the bundle lies on a different isoquant, then we are either producing more output than we need to, or not producing enough to satisfy our needs.)

2. The optimal consumption bundle will satisfy the tangency condition:

$$\begin{aligned} slope\ of\ IQ &= slope\ of\ IC \\ MRTS_{ij} &= -\frac{MP_i}{MP_j} = -\frac{w_i}{w_j} \end{aligned}$$

   which follows from (9.1). (We can also prove this by contradiction. Suppose the optimal bundle is a, and at this point, the isoquant and isocost curves are not tangential. Then, at a, the curves must cross. But then, there is a point along the isoquant curve (b, say) which achieves a lower isocost line than a does. Since input bundles a and b both produce y units of output, and b is cheaper than a, then it cannot be optimal for a firm to choose input combination a. We have a contradiction.)

The tangency condition says that the rate at which producers are able to substitute between inputs and keep total output unchanged (i.e. $MRTS$) must be the same as the rate at which they are able to a substitute between inputs and keep total cost unchanged (i.e. slope of IC). If this is the case, then it is impossible for the firm to reallocate inputs in the production process in such a way as to reduce costs. Why? Suppose in $MRTS_{ij} = -3$ and $-\frac{w_i}{w_j} = -2$, and the firm decides to reallocate inputs by using 1 additional unit of i. This requires that it uses 3 fewer units of input j. But for a unit increase in x_i, reducing the amount of x_j by

2 units leaves costs unchanged - so reducing the amount of x_j by 3 units will cause costs to fall. Hence, it is possible for the firm to reallocate inputs in the production process in such a way as to reduce costs. But if so, the original input choice could not have been optimal. Hence, if $MRTS \neq slopeIC$, then there is always an opportunity for the firm to reallocate inputs in such a way as to reduce costs. It follows that, at the optimum, $MRTS = slopeIC$.

Example 37. Consider a firm with production function $f(x) = x_1^{0.5}x_2^{0.25}$. Suppose the firm seeks to produce some output y. The firm's problem is:

$$\max_{x_1,x_2} w_1x_1 + w_2x_2 \quad s.t. \quad x_1^{0.5}x_2^{0.25} = y$$

The Lagrangian is:

$$\mathcal{L} = w_1x_1 + w_2x_2 + \lambda\left(y - x_1^{0.5}x_2^{0.25}\right)$$

First order conditions are:

$$\begin{aligned}\frac{\partial\mathcal{L}}{\partial x_1} &= w_1 - 0.5\lambda x_1^{-0.5}x_2^{0.25} = 0\\ \frac{\partial\mathcal{L}}{\partial x_2} &= w_2 - 0.25\lambda x_1^{0.5}x_2^{-0.75} = 0\\ \frac{\partial\mathcal{L}}{\partial \lambda} &= y - x_1^{0.5}x_2^{0.25}\end{aligned}$$

Re-arranging the first two conditions in terms of λ gives:

$$\lambda = 2w_1\frac{x_1^{0.5}}{x_2^{0.25}} = 4w_2\frac{x_2^{0.75}}{x_1^{0.5}}$$
$$x_1 = 2\frac{w_2}{w_1}x_2$$

Substituting this into the third equation (i.e. production function) gives:

$$\left(2\frac{w_2}{w_1}x_2\right)^{0.5}x_2^{0.25} = y$$
$$\hat{x_2}(y, w_1, w_2) = \left(\frac{w_1}{2w_2}\right)^{\frac{2}{3}}y^{\frac{4}{3}}$$

which implies that:

$$\hat{x_1}(y, w_1, w_2) = \left(\frac{2w_2}{w_1}\right)^{\frac{1}{3}}y^{\frac{4}{3}}$$

Finally, substituting back into the expression for the Lagrange multiplier, we find:

$$\hat{\lambda}(y, w_1, w_2) = \left(2^{\frac{4}{3}}\right)w_1^{\frac{1}{3}}w_2^{\frac{2}{3}}y^{\frac{1}{3}}$$

Notice that the conditional factor demand functions are decreasing in the factor price, increasing in the price of the other factor and increasing in income. We say inputs are normal and substitutes in production.

9.2.1 Cost Function

The conditional factor demand functions give the input bundles which minimise the cost of producing a particular quantity of output y, given input prices w. The **cost function** specifies this minimised cost, for each potential output quantity.

Given the conditional factor demand functions $(\hat{x}_1 \ldots \hat{x}_n)$ it is easy to calculate the minimised cost. We have:

$$C\left(w, y\right) = \sum_i w_i \hat{x}_i(w, y)$$

Example 38. Continue with the previous example. The cost function is:

$$c(y, w_1, w_2) = w_1 \left(\frac{2w_2}{w_1}\right)^{\frac{1}{3}} y^{\frac{4}{3}} + w_2 \left(\frac{w_1}{2w_2}\right)^{\frac{2}{3}} y^{\frac{4}{3}} = 3 \cdot (2)^{-\frac{2}{3}} w_1^{\frac{2}{3}} w_2^{\frac{1}{3}} y^{\frac{4}{3}}$$

The cost function is increasing in input prices and in the desired output level.

Moreover, the marginal cost is:

$$MC = \frac{\partial c(y, w_1, w_2)}{\partial y} = \left(2^{\frac{4}{3}}\right) w_1^{\frac{1}{3}} w_2^{\frac{2}{3}} y^{\frac{1}{3}} = \hat{\lambda}(y, w_1, w_2)$$

The Lagrange multiplier encodes the firm's marginal cost! More on this below...

The cost function is a crucial ingredient in addressing the firm's profit maximization decision. After all — we cannot calculate the firm's profits without knowing its costs. Note that the cost function came from something more primitive —the firm's technology. We derive the firm's cost function as a consequence of optimal input choice.

Given this optimal construction, can we say anything about the behavior of the firm's marginal cost? Consider an increase in target output y. The existing input choices will no longer produce the desired output level; feasibility requires that they change. We can write the feasibility condition as:

$$f\left(\hat{x}_1(y, w), \ldots, \hat{x}_n(y, w)\right) = y$$

The conditional factor demands associated with output y should be such that, when combined, they produce y units of output. (Obviously!) Differentiating both sides of this feasibility condition by y gives:

$$\sum_i \frac{\partial f(\hat{x}(y, w)}{\partial x_i} \cdot \frac{\partial \hat{x}_i(y, w)}{\partial y} = 1$$

To interpret this, take each term under the summation sign. The second term in the pair tells us how much the demand for input i changes when target output goes up by 1. The first term tells us how much output will change if 1 more unit of the input is hired. The product of these gives the total amount output changes because demand for input i went up by the amount that it did. Summing over all inputs i gives the total change in output. Of course, this total change in output should be exactly 1.

Now, back to the cost function: $c(y, w) = \sum_i w_i \hat{x}_i(y, w)$. Differentiating this w.r.t. y, we have:

$$\begin{aligned}
\frac{\partial c(y,w)}{\partial y} &= \sum_i w_i \frac{\partial \hat{x}_i(y,w)}{\partial y} \\
&= \sum_i \hat{\lambda}(w,y) \frac{\partial f(\hat{x}(y,w))}{\partial x_i} \cdot \frac{\partial \hat{x}_i(y,w)}{\partial y} \\
&= \hat{\lambda}(y,w) \sum_i \frac{\partial f(\hat{x}(y,w))}{\partial x_i} \cdot \frac{\partial \hat{x}_i(y,w)}{\partial y} \\
&= \hat{\lambda}(y,w)
\end{aligned}$$

where in the second line, we use the first order conditions from the cost minimization problem to make a substitution, and in the 3rd line, we make use the property derived from the feasibility condition.

What do we find? The Lagrange multiplier reflects the marginal cost of production. This is a particular instantiation of a more general result known as the Envelope Theorem. (We will see it's application over and over again.) More generally, the Envelope Theorem tells us that the Lagrange multiplier is a measure of opportunity cost. It tells us how much the optimum would change if the constraint were tightened or loosened by one unit.

Going back to the first order conditions, we know that $\hat{\lambda} = \frac{w_i}{MP_i(\hat{x})}$ for each i. Taking these together gives:

$$\frac{\partial c(y,w)}{\partial y} = \frac{w_i}{\partial f(\hat{x}) / \partial x_i} = \frac{w_i}{MP_i(\hat{x})} \quad \forall i$$

The right hand side is the cost of producing one more unit of output by increasing input i along. (To see this, hiring 1 more unit of input i produces MP_i more output. So to produce 1 more unit of output, the firm would need to hire $1/MP_i$ more units of input i. Since each unit costs w_i, the fraction gives the total cost of purchasing this quantity.) The incremental cost of producing one extra unit is the same no matter which input (or which combination of inputs) we use to produce it!

9.3 Profit Maximization

Having derived the cost function, we can now analyse the firm's profit maximisation problem. Suppose the price of output is p. Then, the firm chooses the level of output y that maximises its profit:

$$\pi = py - c(y, w)$$

The optimal level of output y^* must satisfy the first order condition:

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

(Why must this be the case? Suppose y is profit maximising, but at y, $p \neq MC(y)$. For concreteness, suppose $p > MC$. Then, if the firm produces one additional unit of output, it will receive an additional revenue of p and incur an additional cost of MC. Since $p > MC$, producing this additional unit has the net effect of increasing the firm's profit. But if so, then the original output y could not have been profit maximising. Similarly, suppose $p < MC$. Then, the last unit of output produced generated additional revenue of p, but incurred an additional cost of MC. Since $p < MC$, the last unit produced generated a net loss for the firm. The firm could increase it's profits by decreasing it's output. But if so, then the original output y could not have been profit maximising. Hence, at the optimum, we must have $p = MC$.)

The firm's **supply function** $y^*(p, w)$ gives the profit maximising level of output as a function of the output and input prices. The **factor demand functions** $x_i^*(p, w)$ give the profit maximizing inputs that the firm employs in the production process, given output and input prices. Naturally, there must be a consistency between the factor demand functions and conditional factor demand functions. These must coincide when the firm chooses to produce the profit maximizing output level. Formally:

$$x^*(p, w) = \hat{x}(y^*(p, w), w)$$

Now, using our expression for marginal cost, we can re-write the profit maximizing condition as:

$$\begin{aligned} p &= MC \\ &= \frac{w_i}{MP_i(\hat{x}(y^*(p, w), w))} \quad \forall i \\ MP_i(x^*(p, w)) &= \frac{w_i}{p} \end{aligned}$$

Profit maximization requires the firm to hire inputs up to the point where the marginal product of each input is equal to its replacement cost (or real wage). To make sense of this, note that ${w_i}/{p}$ is the number of units of output that the firm needs to sell in order to cover the cost of hiring a unit of input i. Clearly, if the input produces more than its replacement cost, the firm earns a 'profit' by hiring that input. Similarly, if the input does not replace its cost, the firm makes a loss at the margin. At the optimum, the last unit of each input hired must just its hiring cost.

We could equivalently have written this condition as: $MRP_i(x^*(p, w)) = p \cdot MP_i(x^*(p, w)) = w_i$, where $MRP_i = p \cdot MP_i$ is the **marginal revenue product** of input i. The marginal revenue product is the additional revenue generated by hiring an additional input (which produces additional output that can be sold at price p). The intuition for this the same as above.

Example 39. Suppose the production function is given by $y = x_1^{0.5} x_2^{0.25}$. Given generic output and factor prices (p, w_1, w_2), find the firm's supply function. Verify that it is increasing in the output price p. What is the effect of an increase in factor prices w?

From the previous example, we know that the cost function is given by $C(w_1, w_2, y) = 3 \cdot (2)^{-\frac{2}{3}} w_1^{\frac{2}{3}} w_2^{\frac{1}{3}} y^{\frac{4}{3}}$. Then, the marginal cost function is:

$$\begin{aligned} MC(w_1, w_2, y) &= \frac{\partial C}{\partial y} \\ &= 2^{\frac{4}{3}} w_1^{\frac{2}{3}} w_2^{\frac{1}{3}} y^{\frac{1}{3}} \end{aligned}$$

Profit maximisation requires:

$$\begin{aligned} p &= MC(w_1, w_2, y) \\ p &= 2^{\frac{4}{3}} w_1^{\frac{2}{3}} w_2^{\frac{1}{3}} y^{\frac{1}{3}} \\ y^{\frac{1}{3}} &= \frac{p}{2^{\frac{4}{3}} w_1^{\frac{2}{3}} w_2^{\frac{1}{3}}} \\ y^* &= \frac{p^3}{16 w_1^2 w_2} \end{aligned}$$

Clearly, the supply function is increasing in p and decreasing in w_1 and w_2.

We calculate the factor demand function for the firm with production function given in the above example: $y = 2x_1^{0.5} x_2^{0.25}$. Recall, the marginal product functions are given by:

$$\begin{aligned} MP_1 &= x_1^{-0.5} x_2^{0.25} \\ MP_2 &= \frac{1}{2} x_1^{0.5} x_2^{-0.75} \end{aligned}$$

Applying condition (5.3), we have the following two equations:

$$x_1^{-0.5} x_2^{0.25} = \frac{w_1}{p} \tag{9.2}$$

$$\frac{1}{2} x_1^{0.5} x_2^{-0.75} = \frac{w_2}{p} \tag{9.3}$$

Rearranging (9.2) gives:

$$x_2 = \left(\frac{w_1}{p}\right)^4 x_1^2 \tag{9.4}$$

Substituting this into (9.3), we have:

$$\begin{aligned} \frac{1}{2} x_1^{0.5} \left[\left(\frac{w_1}{p}\right)^4 x_1^2\right]^{-0.75} &= \frac{w_2}{p} \\ \frac{1}{2} x_1^{-1} \left(\frac{w_1}{p}\right)^{-3} &= \frac{w_2}{p} \\ x_1^* &= \frac{p^4}{2 w_1^3 w_2} \end{aligned}$$

Then, by (9.4), we have:

$$\begin{aligned} x_2^* &= \left(\frac{w_1}{p}\right)^4 x_1^2 \\ &= \left(\frac{w_1}{p}\right)^4 \left[\frac{p^4}{2w_1^3 w_2}\right]^2 \\ &= \frac{p^4}{4w_1^2 w_2^2} \end{aligned}$$

We confirm that the factor demand functions are decreasing in their own factor prices, and increasing in the output price p (since - as we showed previously - these inputs are normal).

In the sub-section on cost minimisation, we derived the conditional factor demand functions $\hat{x}_i(w_1, w_2, y)$. How are these distinct from the factor demand functions $x_i^*(w_1, w_2, p)$? The conditional factor demand functions gave the input bundles that a firm would employ if it's objective was to produce some arbitrary level of output y, in the least cost way. This output level need not be the profit-maximising level - indeed, we do not take output prices into account when constructing the conditional factor demand functions. The conditional factor demand functions are the producer theory analogue of the compensated (Hicksian) demand curves. Recall the the compensated demand functions gave the optimal consumption bundles that a consumer would choose, if his goal was to achieve some arbitrary level of utility, and if affordability wasn't a consideration. The conditional factor demand functions give the optimal input bundle that a firm would choose, if its goal was to achieve some arbitrary level of output, and if profitability wasn't a consideration. To this extent, conditional factor demand functions - like the compensated demand functions - are a theoretical construct. What we actually observe are the factor demand functions.

9.4 The Profit Function

Definition 11. The **profit function** gives the firm's profit when it chooses the quantity of output and inputs optimally. We have:

$$\pi(p, w) = py^*(p, w) - \sum_i w_i x_i^*(p, w) = py^*(p, w) - c(y^*(p, w), w)$$

The profit function has several properties that contribute useful predictions about the firm's behavior.

Lemma 7. *(Hotelling's Lemma) The firm's supply function and factor demand functions can be recovered from the profit function as follows:*

- $y^*(p, w) = \frac{\partial \pi(p, w)}{\partial p}$.

- $x_i^*(p,w) = -\frac{\partial \pi(p,w)}{\partial w_i}$ *for each* i.

Proof. The proof makes use of the Envelope Theorem in much the same way as we used it when applied to the cost function. Recall, the profit function is $\pi(p,w) = py^*(p,w) - c(y^*(p,w),w)$. Differentiate this w.r.t. p, noting that p enters both directly into the firm's revenue, and indirectly through each instance of y^*. We have:

$$\frac{\partial \pi(p,w)}{\partial p} = y^*(p,w) + \underbrace{\left(p - \frac{\partial c(y^*,w)}{\partial y}\right)}_{=0} \cdot \frac{\partial y^*(p,w)}{\partial p} = y^*(p,w)$$

Similarly, differentiating w.r.t. w_i gives:

$$\frac{\partial \pi(p,w)}{\partial w_i} = -x_i^*(p,w) + \underbrace{\left(p - \frac{\partial c(y^*,w)}{\partial y}\right)}_{=0} \cdot \frac{\partial y^*(p,w)}{\partial w_i} = -x_i^*(p,w)$$

□

We will use this result in two important ways. First, Hotelling's Lemma makes a connection between the firm's profit and the producer's surplus. Recall, the producer's surplus is the area under the firm's supply curve that follows from a price change. Suppose the firm's profit changes from p_0 to p_1. By the fundamental theorem of calculus, the change in the firm's profit is:

$$\pi(p_1,w) - \pi(p_0,w) = \int_{p_0}^{p_1} \frac{\partial \pi(p,w)}{\partial p} dp = \int_{p_0}^{p_1} y^*(p,w)dp$$

which is precisely the producer surplus.

Similarly, the change in the firm's profit following a change in the price of input i (from w_{i0} to w_{i1} is the area under the associated factor demand function. We have:

$$\pi(p,w_1) - \pi(p,w_0) = \int_{w_{i0}}^{w_{i1}} \frac{\partial \pi(p,w)}{\partial p} dp = -\int_{w_{i0}}^{w_{i1}} x_i^*(p,w)dw_i$$

The second important result that follows from Hotelling's Lemma is about the direction of comparative statics. But, before we can present that result, an aside. We say a function is convex if it is "valley-shaped". We know that a convex function has non-negative straight second order partial derivatives.

Lemma 8. *The Profit function is convex. I.e.*

- $\frac{\partial^2 \pi(p,w)}{\partial p^2} \geq 0$
- $\frac{\partial^2 \pi(p,w)}{\partial w_i^2} \geq 0$ *for each* i

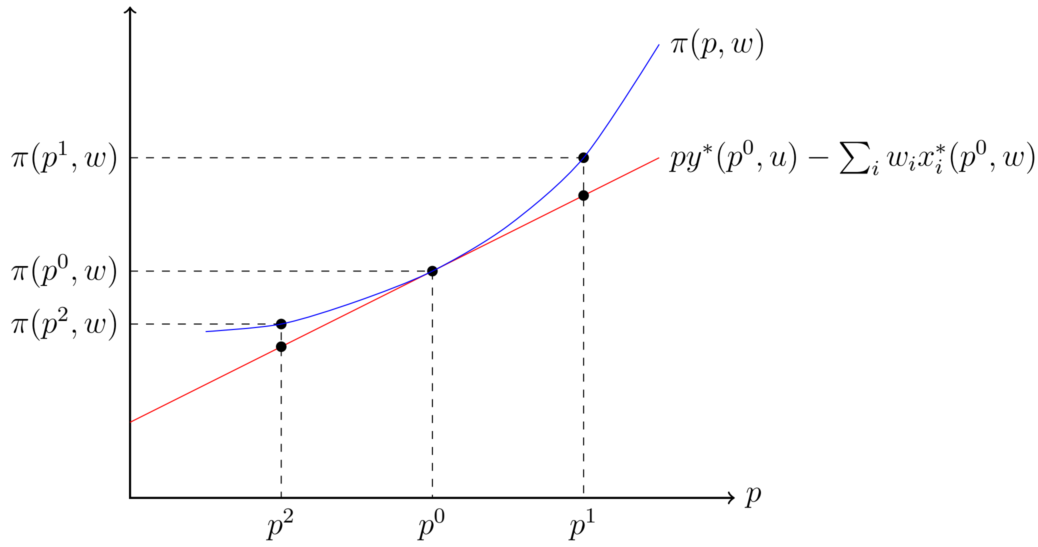

Figure 9.3: The profit function is convex

Rather than prove this result formally, I instead provide a heuristic argument. Consider Figure 9.3. We plot the output price on the horizontal axis and the firm's profit on the vertical. (We could have just as easily plotted an input price.) Throughout, assume input prices are fixed at some level w. Suppose at price p^0 the firm's true profit is $\pi(p^0, w)$, and that the optimal output and inputs are $y^*(p^0, w)$ and $x^*(p^0, w)$. Now consider changes in the output price. It is feasible for the firm to not change its production choice. Then, since the expression for profit is linear, the firm's new profit will be a linear function of the output price. This is represented by the red line in the diagram.

The red line represents a lower bound on the firm's profits. If it does *nothing* following a price change, it can still achieve this level of profit. Of course, the firm would probably benefit from re-allocating its input and output choices after a price change. So its true profit after reallocation will likely be above the red line; for example, the blue line. But this requires the blue line to be curved the way that it is represented in the diagram. It is valley-shaped. The profit function is convex.

Using this result, we can now do comparative statics. We have the following pair of results:

- The supply function is upward sloping. Firms produce more when prices are higher.

$$\frac{\partial y^*(p, w)}{\partial p} = \frac{\partial}{\partial p}\left(\frac{\partial \pi(p, w)}{\partial p}\right) = \frac{\partial^2 \pi(p, w)}{\partial p^2} \geq 0$$

- The factor demand functions are downward sloping. Firms hire fewer of an input when its price increases.

$$\frac{\partial x_i^*(p, w)}{\partial w_i} = \frac{\partial}{\partial w_i}\left(-\frac{\partial \pi(p, w)}{\partial w_i}\right) = -\frac{\partial^2 \pi(p, w)}{\partial w_i^2} \leq 0$$

Both of these results might have seemed intuitive, but it is nice to be able to show that they follow from the theory. And, since many of you probably struggled to follow some of the steps to get to this point, perhaps the results weren't really that intuitive....

Chapter 10

Consumer Theory

10.1 Utility Maximization

Suppose an agent can choose to purchase different quantities of n goods. Let $x = (x_1, ..., x_n)$ denote a bundle of goods, where x_i is the quantity of the i^{th} good. Let X denote the *consumption set* —the set of conceivable bundles the agent can consume. We typically take the consumption to be the non-negative orthant, $X = \mathbb{R}^n_+$.

10.1.1 Preferences & Utility

The agent has preferences over the various bundles in the consumption set. (She can make statements like 'I prefer bundle x to bundle y' or 'I am indifferent between bundles x and z'.) We represent the agent's preferences using a utility function $u(x)$.

Definition 12. A function $u : X \to \mathbb{R}$ *represents* an agent's preferences:

- Whenever the agent strictly prefers bundle x to bundle y, the function satisfies: $u(x) > u(y)$.
- Whenever the agent is indifferent between two bundles x and y, the function satisfies $u(x) = u(y)$

A utility function is simply a convenient way of representing an agent's preferences over bundles. The numbers that we use in the representation themselves have no innate meaning, except in so far as they convey the agent's rankings of different bundles.

An example may help illustrate the point:

Example 40. Suppose $n = 1$, so that there is only one type of good that the agent can consume. Suppose further that the agent dislikes the good, and that finds it increasingly bad to be given more of the good. The following utility functions all represent the agent's utility:

- $u_1(x) = -x$
- $u_2(x) = 5 - x^2$
- $u_3(x) = \frac{1}{1+x}$
- $u_4(x) = -\ln x$

All of these utility functions have the property that more is worse.

As the above exercise makes clear, there is nothing special about the numbers we assign, as long as they respect the appropriate ordering. Utility values do not have to be positive. If $u(x) = 2u(y)$ we cannot conclude that x is twice as preferred as y. The utility function does not capture strength of preference. Indeed, there is no unique (or even natural) scale to represent utility. And there is no hope of trying to measure it (e.g. using a 'hedonometer'). The only information it conveys is the ranking of bundles. Utility is an *ordinal* concept. Given this discussion, we have the following result:

Lemma 9. *Suppose u represents an agent's preferences and let $g(\cdot)$ be a strictly increasing function. Then $v(x) = g(u(x))$ also represents the agent's preferences.*

Notation: We write $x \geq y$ to mean that bundle x contains at least as much of each good than bundle y does, at at least more of some good. I.e. $x_i \geq y_i$ for each $i \in \{1, ..., n\}$, with strict inequality for at least on i.

Definition 13. We say that an agent's preferences are:

- *Monotone*, if whenever $x \geq y$, $u(x) > u(y)$.
- *Convex*, if $u(\lambda x + (1 - \lambda)y) > \min\{u(x), u(y)\}$, for any $\lambda \in (0, 1)$

Monotonicity captures the idea that 'more is better'. An immediate consequence of monotonicity is that $u_i = \frac{\partial u}{\partial x_i} > 0$ for each i.

Convexity captures the idea that moderation is preferred to extremes. The bundle $\lambda x + (1 - \lambda)y$ is a weighted average of bundles x and y. Convexity says that the average bundle cannot be worse than the least preferred of the two extreme bundles. (In the special case that the agent is indifferent between x and y, convexity says that the average bundle must be strictly preferred to either x or y.) The mathematical term for a function having the convexity property is '*quasi-concavity*'.

We can represent consumer preferences diagramatically using **indifference sets**. (When $n = 2$, we refer to indifference sets as 'indifference curves'; for $n > 2$, the geometry of the set is a hyper-surface; not a curve.) An indifference set contains all consumption bundles over which the consumer is indifferent —i.e. from which he derives the same utility. (Formally, an indifference set is a level set of the utility function.) Naturally, there are many indifference sets, each corresponding to a different level of utility. Taking all the indifference sets together gives an *indifference map*.

Indifference sets/curves have several properties that follow straight-forwardly from the properties of the utility function:

- Indifference curves are 'downward-sloping'. An increase in the quantity of good i requires a decrease in the quantity of some other good (j, say) to ensure that overall utility is unchanged. (Follows from mononoticity.)
- 'Higher' indifference sets (i.e. those further from the origin) represent more desirable bundles. (Again by monotonicity.)
- Indifference sets cannot intersect.
- Indifference sets are 'convex' to the origin. (Follows by convexity.)

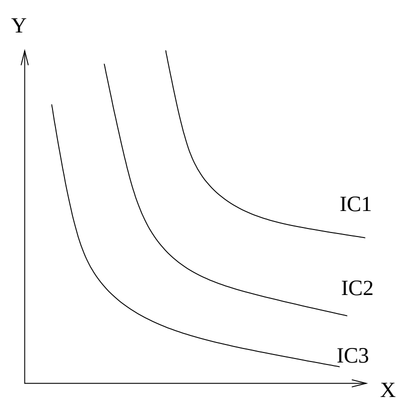

Definition 14. The Marginal Rate of Substitution of good j with respect to good i, denoted MRS_{ij} is the number of units of good j that the agent must forgo in order to remain indifferent after consuming one additional unity of good i, holding the quantity of all other goods constant.

$$MRS_{ij} = -\frac{MU_i(x)}{MU_j(x)} = -\frac{\frac{\partial u(x)}{\partial x_i}}{\frac{\partial u(x)}{\partial x_j}}$$

To see that the formula must be true, consider a change in the quantity of good i, and suppose the agent changes his consumption of good j to ensure utility is unchanged, holding quantities of all other goods constant. We can write $x_j(x_i)$ to capture that the quantity of good j is pinned down by the quantity of good i. Along an indifference curve, we know that: $u(x) = c$. Totally differentiating w.r.t x_i gives:

$$\begin{aligned} \frac{\partial u(x)}{\partial x_i} + \frac{\partial u(x)}{\partial x_j} \frac{dx_j(x_i)}{dx_i} &= 0 \\ \frac{dx_j(x_i)}{dx_i} &= -\frac{\partial u(x) / \partial x_i}{\partial u(x) / \partial x_i} \\ MRS_{ij} &= -\frac{MU_i}{MU_j} \end{aligned}$$

10.1.2 Budget Set

Consider a consumer with income y and suppose the prices of goods are given by the n-vector $p = (p_1, ..., p_n)$. The **budget set** is the set of all consumption bundles that the consumer can afford given his income, and given the prices of goods. The budget set is given by:

$$B(p, y) = \left\{ x \in X \mid \sum_{i=1}^{n} p_i x_i \leq y \right\}$$

The frontier of the budget set is known as the **budget constraint**. All consumption bundles along the budget line are just affordable - i.e. if the consumer purchases a bundle along this line, then he will exhaust his income. A point on the interior of the budget set represents a bundle where the consumer is not spending all of his income. Any point along the budget constraint must satisfy:

$$p \cdot x = \sum_{i=1}^{n} p_i x_i = y$$

10.1.3 Optimization

The agent chooses the bundle x^* to maximize her utility subject to her budget constraint. We have:

$$\max_x u(x) \ s.t. \ p \cdot x = y$$

Caution: In this problem the agent optimizes by choosing quantities of different goods $x_1, ..., x_n$, taking prices and income as given. It is not in the agent's power to choose prices, or (for the purpose of this model) to choose her income. Thus $x = (x_1, ..., x_n)$ are the choice variables, and $(p, y) = (p_1, ..., p_n, y)$ are the parameters.

The Lagrangian is:

$$\mathcal{L} = u(x) - \lambda(p \cdot x - y)$$

which, I re-write for clarity as:

$$\mathcal{L} = u(x_1, ..., x_n) - \lambda(p_1 x_1 + ... + p_n x_n - y)$$

We need to find the vector (x^*, λ^*) that maximizes the Lagrangian. The first order conditions are:

$$\frac{\partial \mathcal{L}}{\partial x_i} = \frac{\partial u(x)}{\partial x_i} - \lambda p_i = 0 \;\; \forall i = 1, .., n$$
$$\frac{\partial \mathcal{L}}{\partial \lambda} = y - p \cdot x = 0$$

This is a system of $n+1$ equations in $n+1$ variables (x_1,...,x_n & λ). We can solve this system simultaneously. Let $x^*(p, y) = (x_1^*(p, y), ..., x_n^*(p, y))$ and $\lambda^*(p, y)$ be the optimizers. We refer to the optimal consumption bundle as **Marshallian Demand** (after economist Lord Alfred Marshall).

Example 41. Suppose the consumer's utility function is given by $u(x_1, x_2) = -\frac{1}{x_1} - \frac{1}{x_2}$ (believe it or not, this satisfies the conditions for a utility function!). We have:

$$\begin{aligned} \max_{x_1, x_2} u(x_1, x_2) &= -\frac{1}{x_1} - \frac{1}{x_2} \\ \text{s.t. } p_1 x_1 + p_2 x_2 &= y \end{aligned}$$

We use the method of Lagrange Multipliers. We have:

$$\mathcal{L} = -\frac{1}{x_1} - \frac{1}{x_2} - \lambda(p_1 x_1 + p_2 x_2 - y)$$

The first order conditions are:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial x_1} &= \frac{1}{x_1^2} - \lambda p_1 = 0 \\ \frac{\partial \mathcal{L}}{\partial x_2} &= \frac{1}{x_2^2} - \lambda p_2 = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= y - p_1 x_1 - p_2 x_2 = 0 \end{aligned}$$

Combining the first two conditions, we have:

$$\begin{aligned} \lambda &= \frac{1}{x_1^2 p_1} = \frac{1}{x_2^2 p_2} \\ &\rightarrow \;\; x_2 = \sqrt{\frac{p_1}{p_2}} x_1 \end{aligned}$$

Substituting this into the third gives:

$$\begin{aligned} p_1 x_1 + p_2 \sqrt{\frac{p_1}{p_2}} x_1 &= y \\ x_1^*(p_1, p_2, y) &= \frac{y}{\sqrt{p_1}\left(\sqrt{p_1} + \sqrt{p_2}\right)} \end{aligned}$$

and then from the optimality condition:

$$x_2^*(p_1, p_2, y) = \frac{y}{\sqrt{p_2}\left(\sqrt{p_1}+\sqrt{p_2}\right)}$$

and:

$$\lambda^*(p, y) = \frac{(\sqrt{p_1}+\sqrt{p_2})^2}{y^2}$$

□

Definition 15. The **Indirect Utility Function** $V(p, y)$ gives the agent's utility under different prices and incomes, assuming that the agent chooses her consumption bundle optimally. We have:

$$V(p, y) = u(x^*(p, y))$$

The direct utility function $u(x)$ provides the agent's ranking over bundles —these are the things she actually cares about. The indirect utility function $V(p, y)$ instead provides the agent's ranking over budget sets. For different budget sets, different consumption bundles will be feasible, and this will affect how much utility the agent can purchase for herself. The agent doesn't instrinsically care about her budget set, but she does care instrumentally, in-so-far as the budget set affects which consumption bundles are affordable.

The indirect utility function is also relevant to policy-making. Although the agent doesn't control prices and incomes, and takes these as given, policy makers are able to control prices and incomes through taxes and transfer policies. But, of course, policy-makers can't actually make consumption choices for the agent. The indirect utility function describes for the policy-maker the agent's preferences over different policy-outcomes, understanding that these policies translate into utility through the consumption choices that the agent makes.

Recall: $V(p, y) = u(x^*(p, y)) - \lambda^*(p, y)(p \cdot x^*(p, y) - y)$. By the envelope theorem, we have:

$$\begin{aligned} \frac{\partial V(p, y)}{\partial p_i} &= -\lambda^*(p, y) x_i^*(p, y) < 0 \\ \frac{\partial V(p, y)}{\partial y} &= \lambda^*(p, y) > 0 \end{aligned}$$

Straight-forwardly, the agent's utility is decreasing in prices and increasing in income. Importantly, the envelope theorem gives meaning to the Lagrange multiplier λ. At the optimum, the Lagrange multiplier measures in utility units the value of loosening the budget constraint by \$1 —i.e. how much more utility the agent would have if she had one additional dollar to spend. In effect, λ is the shadow cost of the budget constraint.

With this understanding of the Lagrange multiplier in mind, return to the first order conditions. For each $i = 1, ..., n$, we have:

$$\frac{\partial u(x)}{\partial x_i} = p_i \frac{\partial V(p, y)}{\partial y}$$

The left-hand side of the FOC is the marginal benefit of increasing consumption of good i by one unit. The right-hand side is the marginal (opportunity) cost. [To see this, note that purchasing one more unit of good i has cost p_i. Making this choice, in effect, tightens the budget constraint by p_i. But we know that each dollar tightening of the budget constraint reduces utility by $\frac{\partial V}{\partial y}$.] The opportunity cost of purchasing one more unit of good i is the amount of utility that the agent would have otherwise been able to purchase if she had p_i more dollars to spend (optimally) on other goods. In particular, the opportunity cost is *not* the price of the good, but the utility forgone.

10.2 Expenditure Minimization

Although we take utility maximization and Marshallian demand to be reasonably descriptive of the agent's actual decision-making, it turns out that there is little we can say about the behavior of Marshallian Demand and Indirect Utility. Two issues are especially problematic:

- The theory makes no prediction about the slope of the Marshallian demand curve. It is possible that $\frac{\partial x_i^*(p,y)}{\partial p_i} > 0$ —i.e. that the 'Law of Demand' is violated. [Convince yourself that this is true graphically by drawing indifference curves appropriately.]
- The theory does not require that $\frac{\partial x_i^*}{\partial p_j}$ and $\frac{\partial x_j^*}{\partial p_i}$ have the same sign. We usually classify goods as substitutes or complements according to the sign of their cross-price effects. But if the cross-price effects don't have the same sign, then such a classification becomes impossible.

To make more useful theoretical predictions, we need a different approach. Consider the following hypothetical problem: The agent has some baseline utility level u_0 that she seeks to achieve. What is the minimum amount of income that the agent needs to achieve that utility, given prices p, if she chooses her consuption bundle optimally. Notice that the agent's income is not specified. She doesn't face a budget constraint. Rather she faces a utility constraint, and we ask how much income she would need to achieve it. The agent's problem is:

$$\min_x \; p \cdot x \; s.t. \; u(x) = u_0$$

Let μ denote the Lagrange multiplier (to distinguish it from the multiplier in the utility maximization problem). The Lagrangian is:

$$\mathcal{L} = p \cdot x - \mu(u(x) - u_0)$$

which, I re-write for clarity as:

$$\mathcal{L} = p_1 x_1 + ... + p_n x_n - \mu(u(x_1, ..., x_n) - u_0)$$

We need to find the vector (x^h, μ^h) that maximizes the Lagrangian. The first order conditions are:

$$\frac{\partial \mathcal{L}}{\partial x_i} = p_i - \mu \frac{\partial u(x)}{\partial x_i} = 0 \;\; \forall i = 1, .., n$$
$$\frac{\partial \mathcal{L}}{\partial \mu} = u_0 - u(x) = 0$$

This is a system of $n+1$ equations in $n+1$ variables (x_1,...,x_n & λ). We can solve this system simultaneously. Let $x^h(p, u_0) = \left(x_1^h(p, u_0), ..., x_n^h(p, u_0)\right)$ and $\mu^h(p, u_0)$ be the optimizers. We refer to the expenditure minimizing bundle as **Hicksian Demand** (after economist John Hicks) or *Compensated Demand.*

Example 42. Take the utility function from the previous example: $u(x_1, x_2) = -\frac{1}{x_1} - \frac{1}{x_2}$. We have:

$$\begin{aligned} &\min_{x_1, x_2} p_1 x_1 + p_2 x_2 \\ &\text{s.t.} \;\; -\frac{1}{x_1} - \frac{1}{x_2} \;\; = \;\; u_0 \end{aligned}$$

We use the method of Lagrange Multipliers. We have:

$$\mathcal{L} = p_1 x_1 + p_2 x_2 - \mu \left(-\frac{1}{x_1} - \frac{1}{x_2} - u_0 \right)$$

The first order conditions are:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial x_1} &= p_1 - \mu \frac{1}{x_1^2} = 0 \\ \frac{\partial \mathcal{L}}{\partial x_2} &= p_2 - \mu \frac{1}{x_2^2} = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= u_0 + \frac{1}{x_1} + \frac{1}{x_2} = 0 \end{aligned}$$

Combining the first two conditions, we have:

$$\begin{aligned} \mu &= x_1^2 p_1 = x_2^2 p_2 \\ &\rightarrow \;\; x_2 = \sqrt{\frac{p_1}{p_2}} x_1 \end{aligned}$$

Substituting this into the third gives:

$$\begin{aligned} -\frac{1}{x_1} - \frac{1}{\sqrt{\frac{p_1}{p_2}} x_1} &= u_0 \\ -\frac{1}{x_1} - \frac{\sqrt{p_2}}{\sqrt{p_1}} \frac{1}{x_1} &= u_0 \\ x_1^h(p_1, p_2, u_0) &= -\frac{\sqrt{p_1} + \sqrt{p_2}}{\sqrt{p_1} u_0} \end{aligned}$$

and then from the optimality condition:

$$x_2^h(p_1, p_2, u_0) = -\frac{\sqrt{p_1} + \sqrt{p_2}}{\sqrt{p_2} u_0}$$

and:

$$\mu^h(p_1, p_2, u_0) = \frac{\left(\sqrt{p_1} + \sqrt{p_2}\right)^2}{u_0^2}$$

□

Consider an agent who faces prices p and income y, and let the associated Marshallian demand be $x^*(p, y)$. Now, suppose the price of good i increases, and the new price vector is p'. The new Marshallian demand is $x^*(p', y)$. Since $\frac{\partial V}{\partial p_i} < 0$, we know that the agent's utility will decrease as a result of the price increase.

Suppose we wish to bring the agent back to her original utility level. We would need to compensate her for the higher price she faces by giving her more income. The Hicksian demand gives the consumption bundle that the agent would choose after this compensation. (Note —although the agent achieves the same utility as before, he typically will not choose the original bundle.) This explains why Hicksian demand is also referred to as Compensated demand —it is the bundle chosen after the agent receives sufficient income compensation to achieve the original level of utility. As we shall see, Hicksian demand and income compensation are related to the income and substitution effects.

Definition 16. The **expenditure function** $e(p, u_0)$ gives the minimum amount of income needed for the agent to be able to achieve a desired utility level u_0 given prices p. We have:

$$e(p, u_0) = p \cdot x^h(p, u_0) = p_1 x_1^h(p, u_0) + ... + p_n x_n^h(p, u_0)$$

Lemma 10. *The expenditure function is concave. This implies that, for each i*

$$\frac{\partial^2 e(p, u_0)}{\partial p_i^2} \leq 0$$

Recall: $e(p, u_0) = p \cdot x^h(p, u_0) - u^h(p, u_0)(u(x^h(p, u_0)) - u_0)$. By the envelope theorem, we have:

Lemma 11 (Shephard's Lemma)**.** $\frac{\partial e(p,u_0)}{\partial p_i} = x_i^h(p, y) \geq 0$

If the agent was purchasing x_i^h units of good i, then a unit increase in the price of i will strain the agent's budget by x_i^h dollars. Absent income compensation, the agent will re-optimize and choose a different bundle that achieves a lower utility. To return him to the original level of utility, he will require x_i^h dollars.

There are two important results that follow from Shephard's Lemma:

1. Law of Demand
$$\frac{\partial x_i^h(p,u_0)}{\partial p_i} \leq 0$$

2. Slutsky Symmetry
$$\frac{\partial x_i^h(p,u_0)}{\partial p_j} = \frac{\partial x_j^h(p,u_0)}{\partial p_i}$$

Proof. To show (1), notice that:

$$\frac{\partial x_i^h(p,u_0)}{\partial p_i} = \frac{\partial}{\partial p_i}\frac{\partial e(p,u_0)}{\partial p_i} = \frac{\partial^2 e(p,u_0)}{\partial p_i^2} \leq 0$$

where the inequality follows from the concavity of $e(p,u_0)$. To show (2), notice that:

$$\frac{\partial x_i^h(p,u_0)}{\partial p_j} = \frac{\partial^2 e(p,u_0)}{\partial p_j \partial p_i} = \frac{\partial^2 e(p,u_0)}{\partial p_i \partial p_j} = \frac{\partial x_j^h(p,u_0)}{\partial p_i}$$

, where the middle equality follows from Young's Theorem. □

Our theory now provides two strong predictions about the nature of consumer choice. First, the law of demand says that Hicksian demand is downward sloping; when price increases, quantity demanded will fall. Second, cross-price effects are symmetric. An increase in the price of good j will affect the Hicksian demand for good i by exactly the same amount as an increase in the price of good i will affect the Hicksian demand for good j. An immediate implication is that the cross price effects have the same sign, so we can (without contradiction) classify goods as complements or substitutes. Given its nice properties, empiricists typically estimate Hicksian price effects (elasticities) rather than Marshallian ones.

But Hicksian demand is a hypothetical notion. We don't believe consumers actually behave this way. By contrast, we believe that Marshallian demand is descriptive of agent behavior, but our theory has no strong predictions about Marshallian demand. To be useful, we need a link between Marshallian demand and Hicksian demand that enables us to connect the actual (observable) choices of agents to the hypothetical (unobservable) ones that generate clear predictions.

10.3 Slutsky Equation

We begin by stating some straight-forward relationships between the utility maximization and expenditure minimization problems. Consider the indirect utility and expenditure functions.

- $V(p, e(p,u_0)) = u_0$

- $e(p, V(p, y)) = y$

Consider the first item. The left-hand side is the maximum utility that the agent can achieve at prices p with an amount of income which is the lowest income that achieves utility u_0. Clearly, this must be u_0. Similarly, take the second item. The left-hand side is smallest income needed (at prices p) to achieve the highest utility that can be achieved with income y at prices p. Again, this must be y.

We can state similar consistency properties for Marshallian and Hicksian demand:

- $x^*(p, e(p, u_0)) = x^h(p, u_0)$
- $x^h(p, V(p, y)) = x^*(p, y)$

Now, starting from the first bullet point: $x_i^h(p, u_0) = x_i^*(p, e(p, u_0))$. Differentiating w.r.t. p_j gives:

$$
\begin{aligned}
\frac{\partial x_i^h(p, u_0)}{\partial p_j} &= \frac{\partial x_i^*(p, e(p, u_0))}{\partial p_j} + \frac{\partial x_i^*(p, e(p, u_0))}{\partial y} \cdot \frac{\partial e(p, u_0)}{\partial p_j} \\
\frac{\partial x_i^h(p, u_0)}{\partial p_j} &= \frac{\partial x_i^*(p, e(p, u_0))}{\partial p_j} + x_j^h(p, u_0)\frac{\partial x_i^*(p, e(p, u_0))}{\partial y} \\
\frac{\partial x_i^h(p, V(p, y))}{\partial p_j} &= \frac{\partial x_i^*(p, y)}{\partial p_j} + x_j^h(p, V(p, y))\frac{\partial x_i^*(p, y)}{\partial y} \\
\frac{\partial x_i^h(p, V(p, y))}{\partial p_j} &= \frac{\partial x_i^*(p, y)}{\partial p_j} + x_j^*(p, y)\frac{\partial x_i^*(p, y)}{\partial y}
\end{aligned}
$$

where the penultimate line makes use of the assumption that $u_0 = V(p, y)$, so that $e(p, u_0) = e(p, V(p, y)) = y$. This is the famous **Slutsky Equation**. The left-hand side is the Hicksian price effect —it is unobserved, but we have strong predictions about it. The right-hand side is the sum of two term: a Marshallian price effect and a Marshallian income effect. Both of these are observable, although the theory provides no predictions about their behavior. The Slutsky equation provides a link between the unobservable and observable terms.

We often present the Slutsky equation in the following way:

$$
\underbrace{\frac{\partial x_i^*(p, y)}{\partial p_j}}_{\text{Total Effect}} = \underbrace{\frac{\partial x_i^h(p, V(p, y))}{\partial p_j}}_{\text{Substitution Effect}} \underbrace{- x_j^*(p, y)\frac{\partial x_i^*(p, y)}{\partial y}}_{\text{Income Effect}}
$$

The left-hand side is the 'total price effect'—it is how much Marshallian demand actually changes in response to a price change. This is what we actually observe. On the right-hand side, we decompose this total change into two distinct effects. The first effect is the 'substitution effect'. It asks how quantity demanded would change if after the change in prices, the agent was compensated with a corresponding change in income, to ensure that any change in demand was not caused by a change in the purchasing power of the agent's income (i.e. affordability). This substitution effect is simply the change in the Hicksian demand.

The second effect is the 'income effect'. It asks how much quantity demanded changed because the purchasing power of the agent's income fell. By Shephard's lemma, we know that to maintain her purchasing power after a unit increase in p_j, the agent would require an additional x_j dollars of income. Since she doesn't actually receive this, it is as if her income falls by x_j (at least in purchasing power terms). For each dollar increase in income, we know that demand will increase by $\frac{\partial x_i^*}{\partial y}$. Thus, since y effectively falls by x_j^*, demand will fall by $x_j \frac{\partial x_i^*}{\partial y}$.

If $j = i$, so that we are considering the own-price effect, we know that the substitution effect must be negative (since $\frac{\partial x_i^h}{\partial p_i} < 0$). However, the income effect may take either sign depending on the income elasticity. If i is a normal good, so that $\eta_i > 0$, then the income effect will be negative as well; the total effect will also be negative. By contrast, if i is an inferior good so that $\eta_i < 0$, then the income effect will be positive. The income and substitution effects move in opposite directions, and thus the overall direction is ambiguous. If the magnitude of the income is large enough, it may overwhelm the substitution effect, such that an increase in price generates an increase in quantity.

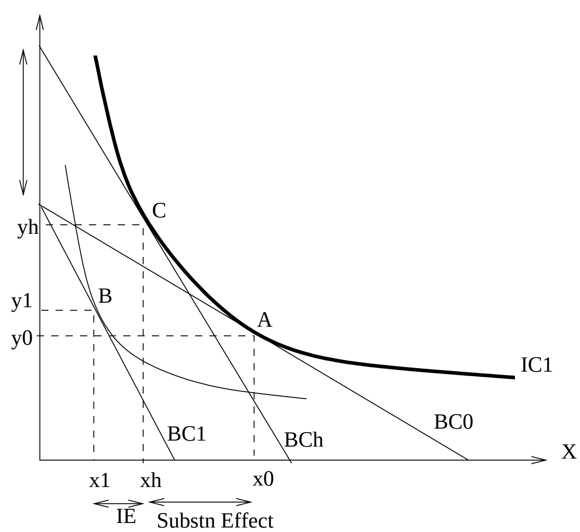

Suppose instead that $j \neq i$. We know by Slutsky symmetry that $\frac{\partial x_i^h}{\partial p_j} = \frac{\partial x_j^h}{\partial p_i}$. Moreover, the sign of these cross-price effects may either be positive (if the goods are substitutes) or negative (if they are complements). For concreteness, suppose they are substitutes. By the

Slutsky equation, the Marshallian cross-price effects are:

$$\frac{\partial x_i^*}{\partial p_j} = \frac{\partial x_i^h}{\partial p_j} - x_j^* \frac{\partial x_i^*}{\partial y}$$
$$\frac{\partial x_j^*}{\partial p_i} = \frac{\partial x_j^h}{\partial p_i} - x_i^* \frac{\partial x_j^*}{\partial y}$$

The sign of the Marshallian cross price effects depends on the size and sign of the income effect. For a normal good, the income effect moves in the opposite direction to the substitution effect. (The increase in p_j makes the agent poorer in real terms, which causes him to buy fewer normal goods.) The sign of the Marshallian cross-price effect will be ambiguous. Moreover, since the income effects can be different for different goods, it could be that the Marshallian cross price effect $\frac{\partial x_i^*}{\partial p_j} > 0$ but $\frac{\partial x_j^*}{\partial p_i} < 0$. Although we classify the goods as substitutes, an increase in the price of one may cause a decrease in the price of the other. But it is important to see that this is a result of the income effect overwhelming the substitution effect.

10.4 Elasticities

We are interested in quantifying the size of changes in demand following changes in prices or income.

Definition 17. The *(price) elasticity* of good i with respect to a change in the price of good j is:

$$\varepsilon_{ij} = \frac{\partial x_i}{\partial p_j} \cdot \frac{p_j}{x_i}$$

The *income elasticity* of good i is:

$$\eta_i = \frac{\partial x_i}{\partial y} \cdot \frac{y}{x_i}$$

You may be more familiar with elasticities defined as the percentage change in quantity divided by the percentage change in price or income, as appropriate. To see that our formal definition captures this, note that:

$$\frac{\%\Delta x_i}{\%\Delta p_j} = \frac{\Delta x_i}{x_i} \div \frac{\Delta p_j}{p_j} = \frac{\Delta x_i}{\Delta p_j} \cdot \frac{p_j}{x_i} \approx \frac{\partial x_i}{\partial p_j} \cdot \frac{p_j}{x_i} = \varepsilon_{ij}$$

Example 43. Return to our previous example, with $u(x_1, x_2) = -\frac{1}{x_1} - \frac{1}{x_2}$. We showed that the Marshallian demand for good 1 was given by:

$$x_1^*(p_1, p_2, y) = \frac{y}{\sqrt{p_1}(\sqrt{p_1} + \sqrt{p_2})}$$

Then, the Marshallian elasticities are:

$$\varepsilon_{11} = -\frac{y\left(1+\frac{\sqrt{p_2}}{2\sqrt{p_1}}\right)}{p_1(\sqrt{p_1}+\sqrt{p_2})^2} \cdot \frac{p_1}{\frac{y}{\sqrt{p_1}(\sqrt{p_1}+\sqrt{p_2}}} = -\frac{\sqrt{p_1}-\frac{1}{2}\sqrt{p_2}}{\sqrt{p_1}(\sqrt{p_1}+\sqrt{p_2})}$$

$$\varepsilon_{12} = -\frac{y\frac{\sqrt{p_1}}{2\sqrt{p_2}}}{p_1(\sqrt{p_1}+\sqrt{p_2})^2} \cdot \frac{p_1}{\frac{y}{\sqrt{p_1}(\sqrt{p_1}+\sqrt{p_2}}} = -\frac{1}{2\sqrt{p_2}(\sqrt{p_1}+\sqrt{p_2})}$$

$$\eta_1 = \frac{1}{\sqrt{p_1}(\sqrt{p_1}+\sqrt{p_2})} \cdot \frac{y}{\frac{y}{\sqrt{p_1}(\sqrt{p_1}+\sqrt{p_2})}} = 1$$

In the following subsections, we explore several properties of elasticities that follow from the theory.

10.4.1 Slutsky Equation in Elasticity Form

Let $s_i = \frac{p_i x_i^*(p,y)}{y}$ denote the (Marshallian) **budget share** of good i. It is the fraction of income that the agent spends on good i.

Return to the Slutsky equation. We seek to transform the price and income effects into elasticities:

$$\frac{\partial x_i^*}{\partial p_j} = \frac{\partial x_i^h}{\partial p_j} - x_j^* \frac{\partial x_i^*}{\partial y}$$

$$\frac{\partial x_i^*}{\partial p_j} \cdot \frac{p_j}{x_i^*} = \frac{\partial x_i^h}{\partial p_j} \cdot \frac{p_j}{x_i^*} - \frac{p_j x_j^*}{y} \frac{\partial x_i^*}{\partial y} \cdot \frac{y}{x_i^*}$$

$$\varepsilon_{ij} = \varepsilon_{ij}^h - s_j \eta_i$$

Notice that the income effect depends on the budget share of good j. The intuition is that for a 1% increase in the price of good j, the consequent loss in the purchasing power of income isn't 1%, but rather depends on the fraction of income devoted to purchasing good j. A 1% increase in p_j generates an s_j% decrease in the purchasing power of income, which translates into a $s_j\eta_i$ decrease in demand via the income effect.

10.4.2 Demand Homogeneity

Definition 18. A function $f(x_1, ..., x_n)$ is **homogeneous of degree** r if, for any $\lambda > 0$, $f(\lambda x_1, ..., \lambda x_n) = \lambda^r f(x_1, .., x_n)$.

Homogeneous functions have the property that if all inputs are scaled up in the same proportion, then output will also respond in a related way. For example, homogeneity of degree 1 implies that if all inputs are doubled (tripled, etc), output will exactly double (triple, etc).

[Euler's Theorem] Suppose $f(x)$ is homogeneous of degree r. Then:

$$x_1 \frac{\partial f(x)}{\partial x_1} + ... + x_n \frac{\partial f(x)}{\partial x_n} = r f(x)$$

Proposition 2. *Marshallian demand is homogeneous of degree* 0 *in* (p, y). *I.e.* $x^*(\lambda p, \lambda y) = x^*(p, y)$ *for all* $\lambda > 0$.

Corollary 1. *The indirect utility function is homogeneous of degree* 0 *in* (p, y). *I.e.* $V(\lambda p, \lambda y) = V(p, y)$ *for all* $\lambda > 0$.

If all prices and income are scaled up in proportion, the consumer's choice will not change. Intuitively, such a change in (p, y) leaves the budget set unchanged —anything, and only those things, that were previously affordable, will be affordable after the change. Thus, whichever bundle maximized utility under the old prices will also maximize utility under the scaled-up prices.

An implication of demand homogeneity is that an agent's choice does not depend on actual prices and income, but rather on relative prices (i.e. opportunity cost) and real income (purchasing power). Since the actual price levels do not matter, we are free to scale these up and down as we see fit, and this will not affect the consumer's decision. Often, for convenience, we set the price of some 'base' good equal to one, and then rescale all other prices and income relative to this good. This good is known as the **numeraire**.

Demand homogeneity also has implications for how price and income elasticities interact. Take any good i. Since $x_i^*(p, y)$ is homogeneous of degree 0, Euler's Theorem implies:

$$\sum_{j=1}^{n} p_j \frac{\partial x_i^*(p, y)}{\partial p_j} + y \frac{\partial x_i^*(p, y)}{\partial y} = 0$$

$$\sum_{j=1}^{n} \frac{\partial x_i^*(p, y)}{\partial p_j} \frac{p_j}{x_i^*(p, y)} + \frac{\partial x_i^*(p, y)}{\partial y} \frac{y}{x_i^*(p, y)} = 0$$

$$\sum_{i=1}^{n} \varepsilon_{ij} + \eta_i = 0$$

If we scale up p_j by 1%, the effect on the demand for good i is given (in percentage terms) by ε_{ij}. If all prices and income are scaled up by 1%, then the overall effect on demand will be the sum of all these elasticities. But we know the net effect must be zero.

Proposition 3. *Hicksian demand is homogeneous of degree* 0 *in* (p). *I.e.* $x^h(\lambda p, u_0) = x^h(p, u_0)$ *for all* $\lambda > 0$, *holding* u_0 *fixed.*

Corollary 2. *The expenditure function is homogeneous of degree* 1 *in* p. *I.e.* $e(\lambda p, u_0) = \lambda e(p, u_0)$ *for all* $\lambda > 0$.

A similar intuition holds for the Hicksian demand case. If all prices are scaled-up in proportion, then relative prices are unchanged, and so the cost minimizing bundle that achieves the desired utility level will be unchanged. Of course, now the cost of achieving this bundle will be higher by the amount of the scale factor. Hence, the expenditure function is homogeneous of degree 1.

An immediate of Euler's Theorem is:

$$\sum_{i=1}^{n} \varepsilon_{ij}^{h} = 0$$

Note —we could have also gotten to this result by using the Marshallian version and substituting using the Slutsky equation.

10.4.3 Engel & Cournot Aggregation

Since Marshallian demand satisfies the budget constraint, any changes in prices or income that affect the budget constraint in a regular way will also have a regular effect on Marshallian demand elasticities.

Take the budget constraint, and suppose there is a change in income. Differentiate w.r.t. y:

$$\begin{aligned}
\sum_{i=1}^{n} p_i x_i^*(p, y) &= y \\
\sum_{i=1}^{n} p_i \frac{\partial x_i^*(p, y)}{\partial y} &= 1 \\
\sum_{i=1}^{n} \frac{p_i x_i^*(p, y)}{y} \frac{\partial x_i^*(p, y)}{\partial y} \cdot \frac{y}{x_i^*(p, y)} &= 1 \\
\sum_{i=1}^{n} s_i \eta_i &= 1
\end{aligned}$$

This result is known as **Engel Aggregation**. It says that the budget-share weighted income elasticity must be one. Intuitively, if income goes up by 1%, then spending must also increase by 1%. But prices are held constant, and so the spending increase must be a consequence of increased demand. Engel Aggregation says that, on average, demand must increase by 1% —but, of course, it may increase by more for some goods, and by less (or even decrease) for other goods.

Now, consider the effect of a change in the price of some good, p_j say. Applying the same procedure:

$$\sum_{i=1}^{n} p_i x_i^*(p,y) = y$$

$$\sum_{i=1}^{n} p_i \frac{\partial x_i^*(p,y)}{\partial p_j} + x_j^*(p,y) = 0$$

$$\sum_{i=1}^{n} \frac{p_i x_i^*(p,y)}{y} \frac{\partial x_i^*(p,y)}{\partial p_j} \cdot \frac{p_j}{x_i^*(p,y)} = -\frac{p_j x_j^*(p,y)}{y}$$

$$\sum_{i=1}^{n} s_i \varepsilon_{ij} = -s_j$$

This result is known as **Cournot Aggregation**. It says that the budget-share weighted price elasticity (w.r.t p_j) is negative and equal in magnitude to the budget share of good j. Intuitively, if the price of good j goes up by 1%, and the agent was spending a fraction s_j of her budget on good j, then her budget tightens by s_j%. Cournot Aggregation says that, on average, demand must decrease by this amount (s_j%) —but, of course, it may decrease by more for some goods, and by less (or even increase) for other goods.

Expressing the Cournot Aggregation formula in terms of Hicksian elasticities is also instructive. Substituting the Slutsky equation gives:

$$\sum_{i=1}^{n} s_i \left(\varepsilon_{ij}^h - s_j \eta_i\right) = -s_j$$

$$\sum_{i=1}^{n} s_i \varepsilon_{ij}^h - s_j \sum_{i=1}^{n} s_i \eta_i = -s_j$$

$$\sum_{i=1}^{n} s_i \varepsilon_{ij}^h - s_j = -s_j$$

$$\sum_{i=1}^{n} s_i \varepsilon_{ij}^h = 0$$

where the third line makes use of Engel Aggregation. The Hicksian version says that the budget-share weighted compensated price elasticity (w.r.t p_j) must be zero. Even though p_j has risen 1%, the agent must still purchase a bundle that achieves the original utility level. On average, quantity demanded cannot change, since utility does not change. Of course, the compensated demand for some goods will fall, and the compensated demand for other goods will rise. (Indeed, since $\varepsilon_{jj} < 0$, it follows that the average *cross-price* elasticity must be positive. Goods are, on average, Hicksian substitutes.)

Chapter 11

Welfare & Consumer's Surplus

11.1 Measures of Welfare

In this chapter, we seek to quantify the welfare implications of a change in prices and income. A direct measure would be to simply ask how much the agent's utility changed. But, as we previously noted, the notion of utility is purely ordinal, and is not endowed with a standard scale or unit of measurement. It's not at all clear what a utility increase of 20 means, other than that utility increased.

Instead, we try to get at a numerical measure of welfare by asking questions like: 'How much would the agent have been willing to pay' to cause the change to come to pass? Or 'how much must we pay the consumer to compensate for the change?' Such an approach has the benefit of measuring welfare in dollars.

11.1.1 Money Metric Utility

Return to the expenditure function $e(p, u_0)$. Notice two things. First —it takes utility as an input; the higher the desired utility, the higher will be the neccessary expenditure. We saw this formally when we noted that $\frac{\partial e(p,u_0)}{\partial u_0} = \mu^h > 0$. Thus, the expenditure function is itself a measure of utility. Higher utility is associated with higher expenditure. Second, the expenditure function is measured in dollars —it has the desired unit of measurement.

Take an arbitrary price vector q.

Definition 19. The money metric utility function $m(u; q) = e(q, u)$ gives a measure of utility level u, as the amount of money needed to achieve that utility given reference prices q.

For example, suppose prices and incomes change from (p^1, y^1) to (p^2, y^2). We know that an agent choosing optimally will initially achieve utility $u_1 = V(p^1, y^1)$, and $u_2 = V(p^2, y^2)$ after the change. A dollar measure of this change in utility would be:

$$m(u_2; q) - m(u_1; q) = e(q; V(p^2, y^2)) - e(q; V(p^1, y^1))$$

for some (arbitrary) reference price vector q. Notice that our measure depends on the reference prices. The measure will be different for different q's. (This is to be expected; we cannot escape the original critique that utility has no objective scale for measurement.) The value of a dollar depends on the vector of prices that the agent faces. Notice also that, although the actual prices changed between the two scenarios, the reference price remained the same. It has to —otherwise we are not making a like-to-like comparison.

Although this approach works for any reference price level, there are some natural reference prices that we might focus on. Amongst them: what was the price level before the change took place? This would seem to be a natural reference point. Alternatively, what is the price level that the agent was anticipating, given the change?

11.1.2 Equivalence Variation

Definition 20. The *Equivalence Variation* (EV) is defined by:

$$EV = e(p^1; V(p^2, y^2)) - e(p^1; V(p^1, y^1))$$

The equivalence variation is the change in the money metric, using the original price level as the baseline reference. The equivalence variation asks 'how much money the agent would need (without the change) to achieve the same utility as she would have if the change in prices and income took place?'. An income change of EV is 'equivalent to' the anticipated change from (p^1, y^1) to (p^2, y^2). If the change hurts the agent, EV will be negative. We can think of EV as the maximum amount the agent would be *willing to pay* to avoid having the change take place. (For example, in a different context, the equivalence variation would be how much the agent would be willing to pay a blackmailer to avoid having harmful information revealed.)

In the following diagram, we plot consumption of a given good on the horizontal axis, and consumption of the numeraire good on the vertical axis. Since the price of the numeraire good is 1, buying m more units of the good is equivalent to buying m more dollars of the good. We can thus think of the vertical axis as being measured in dollars.

The diagram shows the EV associated with an increase in the price of the good along the horizontal axis. The agent is originally at point A. If the price increase occurs, she will choose bundle B. Instead, if there is no change in prices, but her income is reduced sufficiently to bring her utility to the new level, she will choose C. The EV is the necessary amount of income to be taken away. Since a price increase makes the agent worse off, the EV is negative.

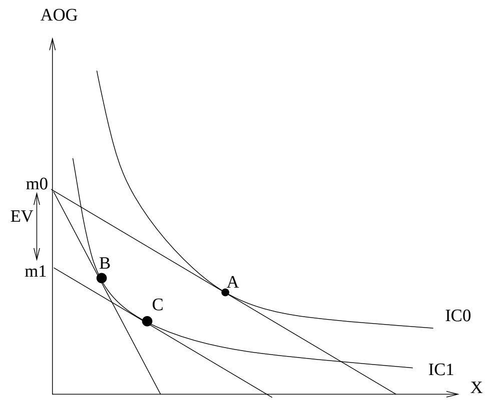

Note that we can re-write EV as follows:

$$\begin{aligned} EV &= e(p^1, u^2) - e(p^1, u^1) \\ &= e(p^1, u^2) - e(p^2, u^2) + e(p^2, u^2) - e(p^1, u^1) \\ &= - \left[e(p^2, u^2) - e(p^1, u^2) \right] + (y_2 - y_1) \end{aligned}$$

where we use the fact that $e(p^i, u^i) = e(p^i, V(p^i, y^i)) = y^i$.

Our formula for EV has been decomposed into two components: one relating to pure price changes and the other relating to pure income changes. Notice that the EV of a pure income change is just the income change itself. This should be obvious. with prices unchanged, what income change is equivalent to a given anticipated income change? Whatever was anticipated!

Now, consider a pure price effect, and for simplicity, suppose only the price of good j changes. (It is easy to extend to multiple price changes.) Recall, that $\frac{\partial e(p,u)}{\partial p_j} = x_j^h(p, u)$. We also know by the Fundamental Theorem of Calculus that for any (differentiable) function F:

$$F(b) - F(a) = \int_a^b \frac{dF(x)}{dx} dx$$

Hence:

$$\begin{aligned} EV &= - \left[e(p^2, u_2) - e(p^1, u_2) \right] \\ &= - \int_{p_j^1}^{p_j^2} \frac{\partial e(p, u_2)}{\partial p_j} dp_j \\ &= - \int_{p_j^1}^{p_j^2} x_j^h(p, u_2) dp_j \end{aligned}$$

which is the area between the Hicksian demand curve and the price axis. [More on this to follow.]

11.1.3 Compensating Variation

Definition 21. The *Compensating Variation* (CV) is defined by:

$$CV = e(p^2; V(p^2, y^2)) - e(p^2; V(p^1, y^1))$$

The compensating variation is the change in the money metric, using the new price level as the baseline reference. The compensating variation asks 'how much money the agent would need after the change to achieve her original level of utility?'. If the change hurts the agent, CV will be negative. The sign of CV is constructed to have the same direction as the change in utility. In a different context, the compensating variation would be the amount of damages awarded by a court in compensation for an irreversible act that harms the plaintiff.

We see the compensating variation in the following diagram.

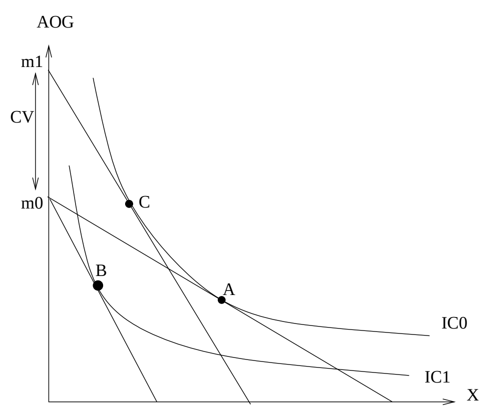

If we increase income from m_0 to m_1, then at the new price, the consumer can afford bundle C which lies on the original indifference curve IC_0. Notice that at point C, IC_0 is tangential to the new budget constraint. (This should be obvious, since the consumer will choose the new consumption bundle optimally, given his new level of income and the higher prices.) You should have noticed that the move from bundle A to bundle C, corresponds to the substitution effect. Since the price of good X is now higher, the consumer can achieve the same utility by consuming less of good X, and more of the composite bundle. The compensating variation is the additional income we had to give the consumer $(m_1 - m_0)$.

Note that we can re-write CV as follows:

$$\begin{aligned} CV &= e(p^2, u^2) - e(p^2, u^1) \\ &= e(p^2, u^2) - e(p^1, u^1) + e(p^1, u^1) - e(p^2, u^1) \\ &= -\left[e(p^2, u^1) - e(p^1, u^1)\right] + (y_2 - y_1) \end{aligned}$$

Similar to EV, our formula for CV has been decomposed into two components: one relating to pure price changes and the other relating to pure income changes. As before, notice that the CV of a pure income change is just the income change itself. This should be obvious. with prices unchanged, what income change compensates the agent for a given income change? Whatever was the actual income change!

Now, consider a pure price effect, and again for simplicity, suppose only the price of good j changes.

$$\begin{aligned} CV &= -\left[(p^2, u_1) - e(p^1, u_1)\right] \\ &= -\int_{p_j^1}^{p_j^2} \frac{\partial e(p, u_1)}{\partial p_j} dp_j \\ &= -\int_{p_j^1}^{p_j^2} x_j^h(p, u_1) dp_j \end{aligned}$$

which is the area between the Hicksian demand curve and the price axis. In fact, we had an analogous expression for EV, except that expression depended on Hicksian demand at the new utility level, whilst CV depended on the Hicksian demand at the old utility level.

11.2 Consumer's Surplus

11.2.1 Consumer's Surplus

The result that CV and EV can be represented as the area between the price axis and the Hicksian demand curve should have alerted you to another common measure of welfare: the consumer's surplus.

Definition 22. Suppose the price of good j changes from p_j^1 to p_j^2, holding all other prices constant. The *consumer's surplus* is defined by:

$$CS = -\int_{p_j^1}^{p_j^2} x^*(p, y) dp_j$$

The consumer's surplus is the area between the Marshallian demand curve and the price axis. This, again, is analogous to CV and EV, except now using Marshallian rather than Hicksian demand.

Is the consumer's surplus a valid measure of the agent's welfare? To answer this question, return to the Utility Maximization problem and the envelope theorem. Recall:

$$V(p, y) = u(x^*(p, y)) - \lambda^*(p, y)\left[p \cdot x^*(p, y) - y\right]$$

By the envelope theorem:

$$\frac{\partial V(p,y)}{\partial p_j} = -\lambda^*(p,y)x_j^*(p,y)$$
$$\frac{\partial V(p,y)}{\partial y} = \lambda^*(p,y)$$

Combining these gives the following result:

Proposition 4 (Roy's Identity)**.**

$$x_j^*(p,y) = -\frac{\frac{\partial V(p,y}{\partial p_j}}{\frac{\partial V(p,y)}{\partial y}}$$

Roy's Identity is analogous to Shephard's Lemma, but applied to Marshallian, rather than Hicksian, demand.

Now, using the same logic as above, after a change in the price of good j, the change in the agent's utility is:

$$V(p^2,y) - V(p^1,y) = \int_{p_j^1}^{p_j^2} \frac{\partial V(p,y)}{\partial p_j} dp_j = -\int_{p_j^1}^{p_j^2} \frac{\partial V(p,y)}{\partial y} x_j^*(p,y) dp_j$$

where the second inequality follows by Roy's Identity.

Now, if $\frac{\partial V(p,y)}{\partial y}$ is constant —and this is a big if —then we have:

$$V(p^2,y) - V(p^1,y) = -\frac{\partial V(p,y)}{\partial y} \int_{p_j^1}^{p_j^2} x_j^*(p,y) dp_j = \frac{\partial V(p,y)}{\partial y} \cdot CS$$

Hence, assuming the caveat holds, the change in the agent's utility is proportional to the consumer surplus. Why proportional? Notice that the change in utility is measured in 'utils', the consumer surplus is measured in dollars. The proportionality constant does the job of converting dollars into utils. Indeed, the proportionality constant is precisely $\frac{\partial V}{\partial y}$ —the utility value of one extra dollar. Since this is assumed constant, it is, in fact, the utility value of *every* dollar. Thus, assuming our caveat holds, the consumer surplus provides a dollar measure of the change in the agent's welfare.

But, note the importance of the caveat. If $\frac{\partial V}{\partial y}$ is non-constant, then CS is not a exact measure of the agent's welfare (although, in some cases, it may be a decent approximation).

11.2.2 Relationship to CV & EV

How are the compensating and equivalence variations and consumer surplus related to one another?

Consider the following diagram in which (as in the previous examples) the price of the horizontal good increases. The diagram draws both the compensating and equivalence variations, as well as the associated Marshallian and Hicksian demand functions.

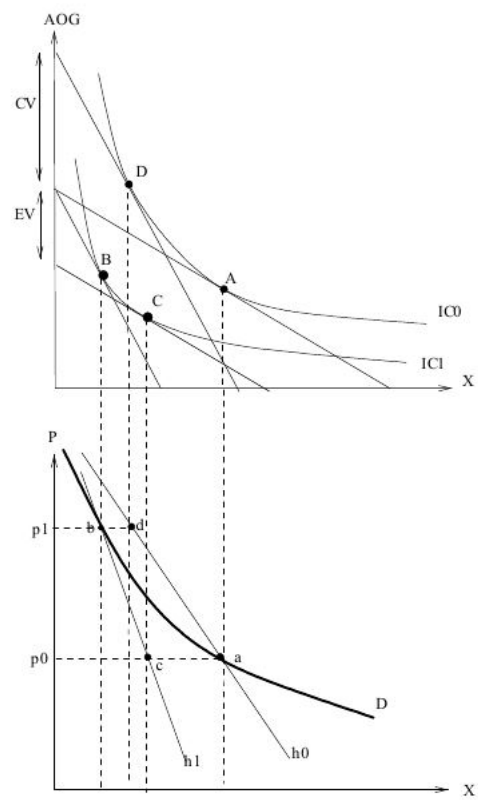

The Marshallian demand corresponds to the shift from A to B, as the price increases from p_0 to p_1. In so doing, the agent's utility decreases from u_0 to u_1. The equivalence variation requires that income be taken away from the agent sufficient to cause his new utility to be u_1. In this scenario, the agent will choose bundle C, and so the Hicksian demand at u_1 corresponds to the shift from C to B. The compensating variation requires that income be given to the agent sufficient to return him to utility u_0. In this scenario, the agent will choose bundle D, and so the Hicksian demand at u_0 corresponds to the shift from A to D. Notice that, as depicted in the diagram, both goods are normal goods.

We can now illustrate the compensating and equivalence variations and the consumer surplus using the demand curves in the lower panel:

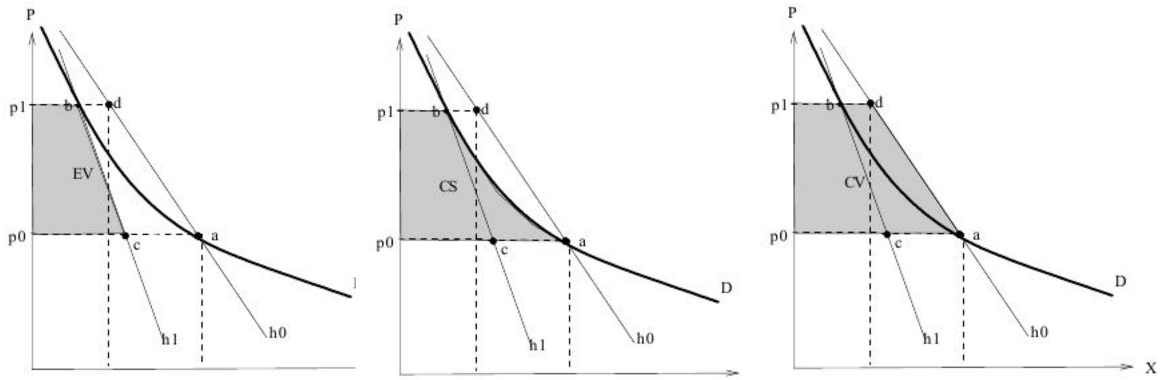

In the above diagrams, we noticed that $CV > EV$. Indeed, this was the case because the Hicksian demand associated with utility u_0 was higher at every price than the Hicksian demand associated with utility u_1. This property, in turn, followed from the fact that the good in question was a normal good. Recall, $u_0 > u_1$, and so the income associated with u_0 is higher than the income at u_1. But with higher income, an agent will consume more of a normal good (at every price level). The opposite would be true if the good were inferior.

We can also see this relationship by noting that the Hicksian demand curves, as drawn, are both steeper than the Marshallian demand curve. This follows immediately from the Slutsky equation. [To see this, note that:

$$\frac{\partial x_i^*(p,y)}{\partial p_i} = \frac{\partial x_i^h(p,V(p,y))}{\partial p_i} - x_i \frac{\partial x_i^*(p,y)}{\partial y} < \frac{\partial x_i^h(p,V(p,y))}{\partial p_i} < 0$$

Marshallian demand is more responsive to prices (in the sense of decreasing by a larger amount than Hicksian demand. Since we plot price on the vertical and quantity on the horizontal axis, this implies that Hicksian demand is steeper.] By the consistency properties, we also know that Marshallian and Hicksian demands must coincide at particular points; $x_i^*(p^0, y) = x_i^h(p^0, u_0)$ and $x_i^*(p^1, y) = x_i^h(p^1, u_1)$. This immediately tells us that c is to the left of a and d is to the right of b, and so $x_i^h(p, u_0)$ must be to the right of $x_i^h(p, u_1)$.

Note, again that the opposite would be true if the good was inferior, or if the price decreased rather than increased. In fact, we have the following relationship:

	Price Increase	Price Decrease
Normal Good	$CV > EV$	$EV > CV$
Inferior Good	$EV > CV$	$CV > EV$

Finally, we note that consumer surplus is bounded by the compensating and equivalence variations. (It turns out this result is only true for a single price change, and that with multiple price changes, CS may be very different from either CV or EV.) Thus, we can think of CS as an easily calculable approximation of CV and EV, even though it is not theoretically a correct measure of welfare.

11.2.3 Which Measure Should We Use?

We have two different measures of welfare - CV and EV. Which should we use? The answer depends in part on the context we wish to analyse. Consider the following example:

The government is concerned that the price of fresh fruit and vegetables is too high and that this is affecting the eating habits of low/middle income earners. The government proposes to subsidise the cost of fresh fruits and vegetables, and to pay for the subsidy by levying a lump-sum income tax. Then, the appropriate measure of welfare would be the compensating variation. The compensating variation is the highest income tax that would make the subsidy worthwhile to the consumer. If the lump-sum income tax is less than CV, the consumers are unambiguously better off. If the income tax is larger than CV, then consumers are worse off.

I illustrate these scenarios in the diagram below:

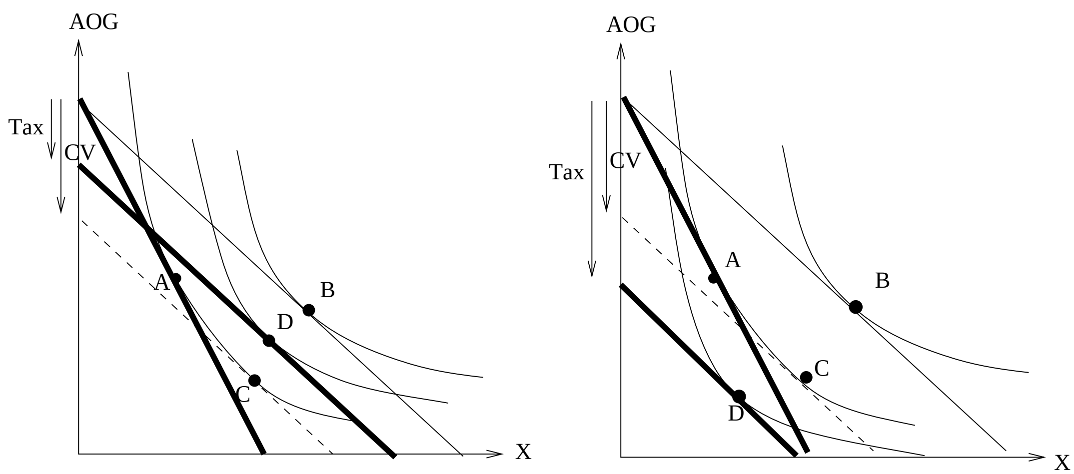

In the above diagram, the thick lines represent the actual budget constraints. The steep line is the budget constraint before the price decrease. The flatter line is the budget constraint after the change in price and income. (In the left hand panel, the magnitude of the income tax is small, relative to the right hand panel.) The thin line is the budget constraint if there was a price decrease (only) and no decrease in income. The dotted line is the budget constraint after the price change, and if just enough income is taken away from the consumer, so that they return to the original level of utility. In both diagrams, point A is the original consumption bundle; point B is the bundle that would've been chosen if there was only a price decrease, but not corresponding decrease in income; point C is the bundle that would be chosen at the new prices, if just enough income was taken away from the consumer to leave them as well off as they were prior to the price change; and point D is the actual new consumption bundle. Note that points A, B and C are identical in both panels - only point D differs. Clearly, if $tax > CV$ then the consumer is worse off, and consumes a bundle on a lower indifference curve. Conversely, if $tax < CV$, then the consumer is better off, and consumes a bundle on a higher indifference curve.

Suppose, on the other hand, the government decides to not decreases prices, but instead determines to supplement consumers' incomes, so that they can better afford fresh foods. Then, the appropriate measure of welfare would be the equivalence variation. The equivalence variation measures how much income the consumer would need to receive to make him as well off as he would've been had the subsidy been implemented. If the magnitude of the income supplement is larger than the EV, then clearly consumers are better off. Conversely, if the income supplement is smaller than EV, the consumer would've been better off with the subsidy.

11.3 Taxation/Subsidies and Excess Burden

We know that the government can affect consumer welfare by intervening in the market by providing price subsidies funded by lump-sum income taxes (or conversely by taxing goods and compensating consumers with income supplements). A natural question is this: can the government in such a way that consumers are better off AND there is no net cost to the government? If so, then there is a strong argument for government intervention in the market.

We define the **excess burden** or **deadweight loss** as the net cost to the government of implementing some policy, such that consumer welfare is unchanged. It is the difference between the compensating (or equivalent) variation associated with some policy and the cost to the government of that policy.

11.3.1 Example: School Vouchers

Consider the following example. The government wishes to increase participation in education. It can do this by subsidising the cost of education (and hence reduce the price from p_0 to p_1), with the subsidy being funded through a lump-sum income tax. Suppose the government levies the largest possible tax that does not make consumers worse off (relative to their utility prior to the implementation of the price subsidy and income tax) —this is the compensating variation. Consumers originally choose bundle A, and after the price and income changes, choose bundle B. The amount of revenue that the government raises through the lump-sum income tax is given by $CV = I_0 - I_1$. Note that the consumer's utility is unchanged.

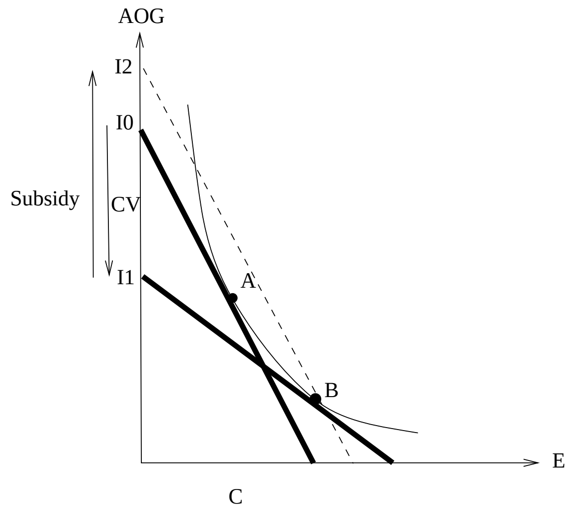

In the above diagram, the dashed line is parallel to the old budget line, but contains the new consumption bundle B. The cost of the subsidy to the government is given by $I_2 - I_1$. (This is the difference in intercepts of two budget lines that contain point B - one reflecting the old prices and one reflecting the new prices.) To see this formally, note that the formula for the new budget constraint is:

$$p_1 E_1 + y_1 = I_1$$

and the formula for the dotted line is:

$$p_0 E_1 + y_1 = I_2$$

where (E_1, y_1) represents the amount of education and all other goods after the policy change, and p_0 and p_1 represent the original and final prices of education. (Recall the price of all-other-goods is normalised to 1). Subtracting these two equations gives:

$$(p_0 - p_1) \times E_1 = I_2 - I_1$$

The left hand side of the expression is exactly the cost of the government subsidy. (The government provides the supplier with $(p_0 - p_1)$ for each unit produced, and the suppliers sell E_B units.) Hence, the cost of the subsidy can be represented by the distance between I_2 and I_1.

Note that the cost of the subsidy is more than the amount of revenue that the government can generate using income taxes. The excess burden of costs over revenues is given by $I_2 - I_0$. This illustrates a more general principle, that government policies that distort relative prices are less efficient than lump-sum transfers.

We could think of this problem slightly differently. Suppose the government, after implementing the subsidy, decided to raise taxes sufficiently to pay for the cost of the subsidy. Then, it must raise taxes by *more* than the compensating variation (since we have just shown that, if it raises taxes just by the level of CV, this will be insufficient to cover the entire cost of the subsidy). But doing so makes consumers worse-off than they were prior to the subsidy. Hence, it is not possible for the government to affect a policy through subsidies, which both maintains the welfare of consumers and keeps the government's budget in balance.

11.3.2 Excess Burden & Deadweight Loss

Consider an agent who originally faces price-income vector (p, y), and let $t = (t_1, ..., t_n)$ be a menu of commodity taxes. The post-tax price to the consumer is $p + t$, so that the price of good i is $p_t + t_i$. We allow for the possibility that $t_i < 0$, so that good i receives a subsidy.

The government's net revenue from such a policy is $R = \sum_{i=1}^{n} t_i x_i^*(p + t, y)$, and the agent's utility is $u = V(p + t, y)$. The equivalence variation is:

$$EV = -\sum_{I=1}^{n} \int_{p_i}^{p_i+t_i} x_i^h(q) dq_i$$

where $q = (q_1, ..., q_n)$ and $q_j = p_j + t_j$ for all $j \neq i$. The **excess burden** or **deadweight loss** is the sum of the government's revenue and the equivalence variation. (The former is the societal gain whilst the latter is the agent's loss from the policy.) We have:

$$\begin{aligned} DWL &= \sum_{i=1}^{n} \left[\int_{p_i}^{p_i+t_i} x_i^h(p + t, u) dq_i - \int_{p_i}^{p_i+t_i} x_i^h(q, u) dq_i \right] \\ &= \sum_{i=1}^{n} \int_{p_i}^{p_i+t_i} \left(x_i^h(p + t, u) - x_i^h(q, u) \right) dq_i \\ &\leq 0 \end{aligned}$$

where the first line makes use of the fact that $x_i^h(p+t, u) = x_i^*(p+t, y)$. To see why this expression is non-positive, recall that $\frac{\partial x_i^h}{\partial q_i} \leq 0$. Suppose $t_i > 0$ so that there is a genuine tax on good i. Then for $q_i \in (p_i, p_i+t_i)$, $x_i^h(p+t, u) \leq x_i^h(q, u)$, and so $\int_{p_i}^{p_i+t_i} \left(x_i^h(p + t, u) - x_i^h(q, u) \right) dq_i \leq 0$. Suppose, instead, that $t_i < 0$, so that there is a subsidy on good i. Then for $q_i \in (p_i + t_i, p_i)$, $x_i^h(p + t, u) \geq x_i^h(q, u)$. Then: $\int_{p_i+t_i}^{p_i} \left(x_i^h(p + t, u) - x_i^h(q, u) \right) dq_i \geq 0$, where we have reversed the integration bounds to make the integral positively oriented. Then $\int_{p_i}^{p_i+t_i} \left(x_i^h(p + t, u) - x_i^h(q, u) \right) dq_i \leq 0$. Hence the DWL is the sum of non-positive integrals, which must be non-positive.

If we focus on a single price increase, the deadweight loss is analogous to triangular area we are used to drawing, except it is defined over Hicksian demand, not Marshallian. <Draw Diagram>

This makes the general point that we saw in the school voucher example. If the government levies a tax on a good, then the revenue it raises will be smaller than the equivalence variation (which is the cost of compensating the consumer for the higher prices she faces). The reason is that, by the substitution effect, the higher price causes the agent to substitute away from the good, so the government earns revenue on a smaller tax base. Similarly, if the government levies a subsidy, then the cost to the government of the subsidy will be larger than the equivalent variation (which is the amount the agent would accept in lieu of the subsidy). The latter point makes the case starkly. If the point of the subsidy was to make the agent better off, the government could achieve this goal more cheaply by simply giving a cash transfer than by distorting prices. In either case, the net societal benefit is negative.

Technical Note: Given our expression for DWL, the marginal DWL from increasing t_j is:

$$\begin{aligned}
\frac{\partial DWL}{\partial t_j} &= \frac{\partial}{\partial t_j}\sum_{i=1}^{n}\int_{p_i}^{p_i+t_i}\left(x_i^h(p+t,u)-x_i^h(q,u)\right)dq_i \\
&= \sum_{i=1}^{n}\int_{p_i}^{p_i+t_i}\frac{\partial x_i^h(p+t,u)}{\partial p_j}dq_i \\
&= \sum_{i=1}^{n}t_i\frac{\partial x_i^h(p+t,u)}{\partial p_j} \\
&= \sum_{i=1}^{n}t_i\frac{\partial x_j^h(p+t,u)}{\partial p_i}
\end{aligned}$$

where the final line follows by Slutsky symmetry. A small increase in the tax on good i causes overall DWL to change by an amount that depends on the size of existing taxes on each good and the Hicksian price responsiveness of each good to the price of good i. We will see the relevance of this shortly.

11.4 Optimal (Ramsey) Taxation

Consider an agent who has preferences over n goods $(x_1, ..., x_n)$ and leisure L. The agent has T units of discretionary time, which she can spend either working or in leisure. (Let $z = T - L$ denote hours worked.) The prices of goods are $p_1, ..., p_n$ and the agent earns a wage w. Leisure is costless. The agent's problem is:

$$\max_{x_1,..,x_n,L} u(x_1, ..., x_n, L) \ s.t. \ p_1x_1 + ... + p_nx_n = w(T - L)$$

We can re-write this as:

$$\max_{x_1,..,x_n,L} u(x_1, ..., x_n, L) \ s.t. \ p_1x_1 + ... + p_nx_n + wL = wT$$

which makes clear that w is the (shadow) price of leisure and wT is the agent's (exogenous) full income.

There is a government that has an exogenous revenue requirement R.

11.4.1 First Best

Suppose the agent's leisure choice is observable and (thus) taxable. The government can levy proportional taxes τ_i on goods and τ_L on labor income, so that commodity prices are $(1 + \tau_i)p_i$ and the price of leisure is $(1 + \tau_L)w$.

Now, suppose the government sets a common tax rate of τ on all goods an leisure. Then the agent's problem becomes:

$$\max u(x,L)\ s.t.\ (1+\tau)p\cdot x+(1+\tau)wL=wT$$

or equivalently:

$$\max u(x,L)\ s.t.\ p\cdot x+wL=\frac{1}{1+\tau}wT$$

A uniform tax on all goods *and* leisure is equivalent to a lump-sum tax on the agent's exogenous full income (where a fraction $\frac{\tau}{1+\tau}$ is confiscated). Since this does not distort prices, the Hicksian demands are unaffected, and so we know there is no deadweight loss or excess burden. [Of course, the government will find it difficult to actually levy the lump-sum tax, so the uniform tax on commodities and leisure does the job in its place.] The government should set a uniform proportional tax rate of $\tau=\frac{R}{R+wT}$.

11.4.2 Second Best

What if leisure is unobservable to the government (since it is un-transacted) and so it cannot be taxed? The government could tax labor instead. Does this amount to the same thing?

Suppose the government sets a uniform tax rate τ on goods and labor. The agent's problem is:

$$\max u(x,L)\ s.t.\ (1+\tau)p\cdot x=(1+\tau)wz$$

Notice that the labor tax is actually a subsidy, since it increases the agent's income. Notice also that all the $(1+\tau)$'s cancel. A uniform tax collects no revenue. The revenue from commodity taxation is exactly offset by the subsidy to labor. The lesson is that if leisure cannot be taxed, the tax system has to distort prices if it has any hope of collecting revenue.

Return to our previous setup, where there are per-unit taxes t_i on goods and labor. The agent's budget constraint is:

$$\sum_{i=1}^{n}(p_i+t_i)x_i=(w+t_L)z$$

By demand homogeneity, we know that we can normalize the price of some good, and re-scale all other prices accordingly. By convention, we do this normalization by setting $t_L=0$. (But note —this does not mean we're ruling out labor taxation. Just that any system with labor taxation is equivalent to the normalized system without, after an appropriate re-scaling.) Also, for convenience, we denote $x_0=-z$ and $p_0=w$ so that the budget constraint is:

$$q\cdot x=q_0x_0+\sum_{i=1}^{n}q_ix_i=-wz+\sum_{i=1}^{n}q_ix_i=0$$

where $q_i=p_i+t_i$ and $q_0=w$ (since $t_0=0$). The agent's unearned income $y=0$.

The agent's problem is:

$$\max_{x_0,..,x_n} u(x_1, ..., x_n, T + x_0) \ s.t. q_0 x_0 + ... + q_n x_n = 0 = y$$

The government's problem is:

$$\max V(q, y) \ s.t. \sum_{i=1}^{n} (q_i - p_i) x_i^*(q, y) = R$$

Notice that the revenue constraint does not include x_0, since labor is untaxed by assumption. Taking producer prices p_i as given, choosing taxes $\{t_1, ..., t_n\}$ is equivalent to choosing post-tax prices $\{q_1, ..., q_n\}$. The Lagrangian is:

$$\mathcal{L} = V(q, y) + \mu \left[\sum_{i=1}^{n} (q_i - p_i) x_i^*(q, y) - R \right]$$

where $\mu > 0$ is the marginal welfare loss associated with increasing R by 1 unit.

The FOC for q_i is:

$$\frac{\partial V(q, y)}{\partial q_i} + \mu \left[x_i^*(q, y) + \sum_{j=1}^{n} t_j \frac{\partial x_j^*(q, y)}{\partial q_i} \right] = 0$$

The first term is the direct welfare effect on the consumer. The bracketed term consists of two parts. The first is the pre-substitution revenue effect; a unit increase in t_i increases revenue by 1 for each unit of good i consumed, and so revenue increases by x_i. The second term includes all substitution effects stemming from the rise in the price of good i.

Now, using Roy's Identity, substitute $\frac{\partial V(p,y)}{\partial q_i} = -\lambda x_i^*(q, y)$ where $\lambda = \frac{\partial V(q,y)}{\partial y}$ is the agent's (private) marginal utility of money. Furthermore, use the Slutsky equation to substitute $\frac{\partial x_j^*}{\partial q_i} = \frac{\partial x_j^h}{\partial q_i} - x_i^* \frac{\partial x_j^*}{\partial y}$. All told, we have:

$$-\lambda x_i^* + \mu \left[x_i^* + \sum_{j=1}^{n} t_j \left(\frac{\partial x_j^h}{\partial q_i} - x_i^* \frac{\partial x_j^*}{\partial y} \right) \right] = 0$$

Let

$$\alpha = \lambda + \mu \sum_{j=1}^{n} t_j \frac{\partial x_j^*}{\partial y}$$

be the social marginal utility of money. (The first term is the agent's private marginal utility of money. The term inside the summation is how much extra revenue the government would get if the agent had one extra dollar to spend. Thus, multiplying this by μ gives the additional utility benefit to the agent of having her tax burden reduced by this amount.) Notice that this is independent of i. Then, re-arranging the above gives:

$$\frac{1}{x_i^*}\sum_{j=1}^{n} t_j \frac{\partial x_j^h}{\partial q_i} = -\frac{\mu-\alpha}{\mu}$$

This is the classic Ramsey optimal taxation formula. Let us make sense of it. The numerator on the LHS is simply the marginal deadweight loss from raising the tax on good i by 1 unit. The denominator on the LHS is the (mechanical) additional revenue that the government receives by increasing the tax on good i by 1 unit (pre-substitution). The ratio is then the marginal deadweight loss per dollar revenue raised when taxing good i.

Next, notice that the RHS is a constant. Since there was nothing particular about the good i that we chose, this gives us the following result: *The marginal DWL per dollar of revenue raised must be the same for all taxable goods.* This should be intuitive. At the margin, if the marginal excess-burden (per unit revenue) was larger for good i than good j, then the government could reduce taxes on good i and increase taxes on good j in a way that kept revenue constant, and decreased overall DWL.

We can use Slutsky symmetry to re-write the Ramsey formula in the following way:

$$\frac{1}{x_i^*}\sum_{j=1}^{n} t_j \frac{\partial x_i^h}{\partial q_j} = -\frac{\mu-\alpha}{\mu}$$

This makes most intuitive sense when taxes are small. If so, then given a set of taxes $t_1, .., t_n$, the total change in the agent's (compensated) demand of good i is approximately $\Delta x_i^h \approx \sum_{j=1}^{n} t_j \frac{\partial x_i^h}{\partial q_j}$. Thus, the LHS is roughly the proportional change in the (compensated) demand for good i. Again, since the RHS is constant, optimal taxation requires that the demand for all goods change in proportion to one another. (Note —we call $\frac{\Delta x_i^h}{x_i}$ the *index of discouragement*. It is how much (relatively) taxation should depress demand for each good.)

We can re-write the Ramsey rule using elasticities rather than Hicksian price effects. We have:

$$\begin{aligned}
\frac{1}{x_i^*}\sum_{j=1}^{n} t_j \frac{\partial x_i^h}{\partial q_j} &= -\frac{\mu-\alpha}{\mu} \\
\sum_{j=1}^{n} \frac{t_j}{q_j}\cdot\frac{\partial x_i^h}{\partial q_j}\cdot\frac{q_j}{x_i^*} &= -\frac{\mu-\alpha}{\mu} \\
\sum_{j=1}^{n} \frac{t_j}{q_j}\cdot\varepsilon_{ij}^h &= -\frac{\mu-\alpha}{\mu}
\end{aligned}$$

where $\frac{t_j}{q_j} = \frac{t_j}{p_j + t_j}$ is the percentage of the post-tax price that is paid in taxes.

Recall, this equation holds for each good $i = 1, ..., n$, and so in fact, we have a system of n (linear) equations in n variables. We can write this in matrix form:

$$\begin{bmatrix} \varepsilon^h_{11} & \cdots & \varepsilon^h_{1n} \\ \vdots & \ddots & \vdots \\ \varepsilon^h_{n1} & \cdots & \varepsilon^h_{nn} \end{bmatrix} \begin{pmatrix} \frac{t_1}{q_1} \\ \vdots \\ \frac{t_n}{q_n} \end{pmatrix} = -\frac{\mu - \alpha}{\mu} \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}$$

The Ramsey rule is very general. I now consider some simple classic results. First, suppose that compensated cross-price effects are zero between all taxable goods; i.e. $\varepsilon^h_{ij} = 0$ for all $i \neq j$ and $i, j = 1, ..., n$. Then we have:

$$\frac{t_i}{q_i} = -\frac{\mu - \alpha}{\mu} \cdot \frac{1}{\varepsilon^h_{ii}}$$

This is the famous *inverse elasticity formula.* When taxable goods are Hicksian independent (so that there are no Hicksian cross price effects; i.e. an increase in the tax on good i does not cause DWL to change because of changes in the demand for good j)... the tax rate on each good should be inversely proportional to its (Hicksian) own-price elasticity.

But we know that, by demand Homogeneity, $\sum_{j=0}^{n} \varepsilon^h_{ij} = 0$. Thus, if by assumption $\varepsilon^h_{ij} = 0$ whenever $i \neq j$ and $j \neq 0$, then $\varepsilon^h_{i0} = -\varepsilon^h_{ii}$. [Recall —we only assumed that *taxable* goods were Hicksian independent.] Recall, the price of good 0 is the wage (i.e. the shadow price of leisure). Moreover, since $\varepsilon_{ii} < 0$ it must be that $\varepsilon_{i0} > 0$ —so that good i is a substitute for leisure.

The inverse elasticity formula becomes:

$$\frac{t_i}{q_i} = \frac{\mu - \alpha}{\mu} \cdot \frac{1}{\varepsilon^h_{i0}}$$

Thus, the tax rate on each good should be inversely proportional to its substitutability with leisure. Goods that are more complementary to leisure should have higher taxes.

In thinking about this insight, it is worth returning to the problem that motivated our second best analysis. If leisure were taxable, then optimal commodity taxation would be tantamount to lump-sum taxation and there would be no DWL. The reason why the optimal taxation problem becomes challenging is that leisure is not taxable, and this implies that to raise revenue, the government must levy price distorting taxes. We must then solve for the DWL minimizing tax profile. The Ramsey formula tells us that this optimal tax profile is constructed so as to indirectly tax leisure as best as possible by targeting goods that a complementary to leisure.

Chapter 12

General Equilibrium & Welfare

In the previous chapters, we have provided a descriptive theory of how goods and resources will be allocated within a market framework, given a vector of prices. In this chapter, we ask two big questions:

1. How should goods and resources be allocated?
2. Can the market achieve these desired allocations?

12.1 Setup

Consider the following environment: There are I agents, $i = 1, ..., I$. There are n goods $j = 1, ..., n$ and m inputs $k = 1, ..., m$. For simplicity we assume that each good is produced by a single (unique) firm, so that there are n firms in total. (We can easily generalize this assumption.) Let $x = (x_1, ..., x_n)$ denote a generic bundle of consumption goods, and $z = (z_1, ..., z_m)$ denote a generic bundle of inputs.

Agents have preferences over consumption bundles. Let $u_i(x)$ be a utility function that represents agent i's preferences over bundles. Similarly, each firm has a production technology represented by a production function $y_j = f_j(z)$, where y_j is the quantity of output produced by firm j, and z is the associated input bundle.

Each agent begins with an endowment of inputs $e^i = (e^i_1, ..., e^i_m)$. Notice that we use subscripts to denote inputs and superscripts to denote the agent who receives them. Let $\bar{e}_k = \sum_i e^i_k$ denote the aggregate endowment of input k. There is no initial endowment of goods. Consumption of final goods and services requires the agents to 'sell' resources to firms, who then engage in production, and 're-sell' final goods and services back to consumers.

Definition 23. An **allocation** specifies:

- A vector of inputs to be utilitized by each firm: $\{z^j\}_{j=1}^n = \{(z_1^j, \ldots, z_k^j)\}_{j=1}^n$.
- The quantity of goods to be produced by each firm: $y_1, \cdots, y_j$.
- A vector of goods consumed by each agent: $\{x^i\}_{i=1}^I = \{(x_1^i, \ldots, x_n^i)\}_{i=1}^I$.

An allocation specifies which firms get which resources, how much output is produced, and how that output is distributed amongst the agents in the economy.

Definition 24. An allocation is **feasible** if:

- All inputs are allocated: For each input k, $\sum_j z_k^j = \bar{e}_k$.
- All outputs are allocated: For each good j, $\sum_i x_j^i = y_j$.
- Outputs are technologically feasible: For each good/firm j: $y_j = f_j(z^j)$.

Our study of general equilibrium will be in two parts. We first begin by asking what types of allocations are 'desirable' or 'efficient'. To do so, we ask how a hypothetical benevolent social planner would choose to allocate inputs and outputs, if she could freely take from some and give to others. Allocations chosen by a benevolent social planner establish a benchmark against which to assess the performance of actual institutions, such as markets. In the second part of our analysis, we study the actual allocation that would arise in a competitive market, and compare it to the efficient benchmark.

12.2 Pareto Optimality

In this subsection, we step away from the market —you won't see any mention of prices —and answer the first question: which allocations are desirable?

Let a and b be two different allocations.

Definition 25. Allocation a **Pareto dominates** allocation b, if every agent is at least as well off under allocation a, and some agents are strictly better off. (Formally, $u_i(x_i(a), y_i(a)) \geq u_i(x_i(b), y_i(b))$ for each $i \in \{1, \ldots, I\}$, with strict inequality for at least one i.)

We say a is **Pareto superior** to b and that b is **Pareto inferior** to a. If an allocation is Pareto dominated, then it is possible to reallocate goods between agents in such a way as to make no agent worse off, and some agents better off.)

Definition 26. An allocation is **Pareto optimal** (or **Pareto efficient**) if there is no allocation that Pareto dominates it.

An allocation is Pareto optimal if it is impossible to reallocate goods in such a way that some agents are better off, without making other agents worse off.

Example 44. Suppose there are two goods: apples and bananas. Agent 1 enjoys bananas but has a strong preference for apples —she will willingly exchange any quantity of bananas for more apples. Agent 2 has the opposite preference —she enjoys apples, but has a strong preference for bananas. Suppose there are 10 apples and 20 bananas to allocate between the two agents.

- The allocation $(x_A^1, x_B^1) = (5, 5)$ and $(x_A^2, x_B^2) = (5, 10)$ is not Pareto optimal. E.g. it is Pareto dominated by the allocation $(x_A^1, x_B^1) = (5, 10)$ and $(x_A^2, x_B^2) = (5, 10)$. The original allocation was inefficient because it failed to allocated all possible goods.

- The allocation $(x_A^1, x_B^1) = (0, 20)$ and $(x_A^2, x_B^2) = (10, 0)$ is not Pareto optimal. It is Pareto dominated by $(x_A^1, x_B^1) = (10, 0)$ and $(x_A^2, x_B^2) = (0, 20)$.

- The allocation $(x_A^1, x_B^1) = (10, 5)$ and $(x_A^2, x_B^2) = (0, 15)$ is Pareto optimal.

- In fact, any allocation of the form: $(x_A^1, x_B^1) = (10, c)$ and $(x_A^2, x_B^2) = (0, 20 - c)$ or $(x_A^1, x_B^1) = (d, 0)$ and $(x_A^2, x_B^2) = (10 - d, 20)$ is Pareto optimal.

Comments

- A Pareto optimal allocation does not necessarily Pareto dominate all other allocations. (e.g. In the above example, consider the allocations $(x_A^1, x_B^1) = (8, 0)$ and $(x_A^2, x_B^2) = (1, 20)$, versus $(x_A^1, x_B^1) = (9, 2)$ and $(x_A^2, x_B^2) = (1, 18)$. We know that the first is Pareto optimal and the second is not. Yet, the first does not Pareto dominate the second.

- There may be many Pareto optimal allocations.

- Pareto optimality is an efficiency notion - it asks whether we have exploited every opportunity to reallocate goods in a way that is mutually beneficial to agents. However, it is not an equity notion. (It makes no comment about the distribution of goods between agents - as long as it is not possible to reallocated the goods in such a way that at least some agents are better off and no-one is worse off.) To this extent, we can think of the Pareto criterion as a minimal requirement of any notion of optimal social allocation. (Surely we wouldn't ever NOT want the Pareto condition to be satisfied.) In addition, we may wish to impose further criteria to attend to equity and/or other concerns.

How can we know if an allocation is Pareto optimal? Let's consider the social planner's (or Pareto Planner's) problem. Let $\gamma_1, \ldots, \gamma_I > 0$ be a set of 'Pareto weights', and suppose the

Pareto planner chooses the allocation to maximize:

$$\begin{aligned}
\max_{x^1,\ldots,x^I,z^1,\ldots,z^n} \quad & \sum_i \gamma_i u_i(x^i) \\
s.t. \quad & \sum_i x^i_j = f_j(z^j) \;\; \forall j = 1,\ldots,n \\
& \sum_j z^j_k = \overline{e}_k \;\; \forall k = 1,\ldots,m
\end{aligned}$$

I.e. the Pareto planner seeks to allocate resources and outputs in a feasible way that maximizes a weighted sum of the agent's utilities.

Lemma 12. *Suppose an allocation* $(\{x^i\},\{z^j\})$ *solves the Planner's problem for some set of Pareto weights* $\gamma_1,\ldots,\gamma_I$. *Then the allocation is Pareto optimal.*

To find the set of Pareto optima it suffices to solve the Planner's problem for various Pareto weights. Whatever solutions come out (as we vary the Pareto weights) will be Pareto optima. The intuition is straightforward. At a Pareto optimum, it is impossible to make some agents better off without making others worse off. But that condition must obviously be met if the planner is maximizing a weighted sum of agents' utilities. (Note —you may object that taking weighted sums of utilities is too simplistic, and that there are better ways of getting at social welfare. But we're not taking a stance on what the social welfare function should look like. We're simply arguing that if the planner were to maximize weighted sums of utilities, he would span out the set of Pareto optima.)

Let us write the Lagrangian for the Pareto planner's problem. Note that there are now $n+m$ constraints. Let λ^x_j denote the Lagrange multiplier associated with the constraint that the total quantity of good j allocated must equal the amount supplied by firm j. Similarly, let λ^z_k denote the Lagrange multiplier associated with the constraint that the total quantity of input k allocated must equal its aggregate endowment. Recall, these Lagrange multipliers can be interpretted as social opportunity costs.

$$\mathcal{L} = \sum_i \gamma_i u_i(x^i) - \sum_j \lambda^x_j \left(\sum_i x^i_j - f_j(z^j) \right) - \sum_k \lambda^z_k \left(\sum_j z^j_k - \overline{e}_k \right)$$

The first order conditions are:

$$\frac{\partial \mathcal{L}}{\partial x^i_j} = \gamma_i \frac{\partial u_i(x^i)}{\partial x_j} - \lambda^x_j = 0 \;\; \forall i, j \tag{12.1}$$

$$\frac{\partial \mathcal{L}}{\partial z^j_k} = \lambda^x_j \frac{\partial f_j(z^j)}{\partial z_k} - \lambda^z_k = 0 \;\; \forall j, k \tag{12.2}$$

From the first order conditions, we can show that Pareto optimality requires that 3 types of conditions be met:

1. The input mix for each firm must be optimal. (We cannot reallocate inputs between firms in such a way that each firm can produce (weakly) more.)

2. The output mix must be optimal for agents. (I.e. We cannot make all agents (weakly) better off by producing a different mix of goods (given finite inputs).

3. The allocation of outputs between the agents must be optimal. (I.e. given the outputs produced, we cannot make all agents (weakly) better off by re-allocating goods between them.)

12.2.1 Efficient Input Allocation

Start with the first type of condition: efficient input allocation. To see what is required here, take the second FOC. Since this must hold for any two inputs k and k', we have: $\lambda_j^x MP_k^j(z^j) = \lambda_k^z$ and $\lambda_j^x MP_{k'}^j(z^j) = \lambda_{k'}^z$. Combining these (substituting out for λ_j^x) gives:

$$\frac{MP_k^j(z^j)}{MP_{k'}^j(z^j)} = \frac{\lambda_k^z}{\lambda_{k'}^z}$$

Notice two things. First, the left hand side is the ratio of marginal products, which we know is the marginal rate of technical substitution for firm j. Second, the right hand side is independent of which firm we are analyzing. Hence, the MRTS must equal the same constant for all firms, and so the MRTS's must be equal for all firms. We have:

$$MRTS_{k,k'}^j(z^j) = -\frac{MP_k^j(z^j)}{MP_{k'}^j(z^j)} = -\frac{MP_k^{j'}(z^{j'})}{MP_{k'}^{j'}(z^{j'})} = MRTS_{k,k'}^{j'}(z^{j'})$$

But recall, the marginal rate of technical substitution is the amount of input k' that the firm can forgo in order to keep output unchanged if it hires one more unit of k. Pareto optimality requires that all firms be willing to trade-off inputs at the same rate.

We can represent this diagrammatically using the **Edgeworth box in Production**. (See Figure 12.1, below.) The dimensions of the Edgeworth box represent the total supply of two inputs $\overline{k}$ and $\overline{k'}$. The horizontal distances represent the quantities of k allocated to each firm, and the vertical distances represent the quantities of k'. The input allocation for firm j is measured from the bottom-left vertex O_j, whilst the allocation for firm j' is measured from the top-right vertex $O_{j'}$.

We can draw the isoquants for each firm within the Edgeworth box. Isoquants retain their usual properties — they are downward sloping, convex and non-intersecting. As usual, the isoquants which are further away from a firm's origin, represent input bundles producing a larger quantity of output.

Lemma 13. *An allocation of inputs is Pareto optimal if the isoquants of each firm going through that allocation are tangent to one another.*

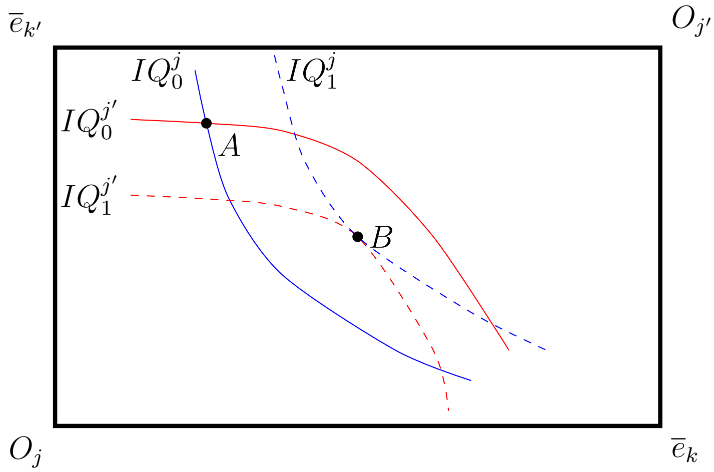

Figure 12.1: Edgeworth Box in Production

We prove this by contradiction. Consider an allocation where the isoquants are not tangent (such as A). Suppose we did a reallocation where we took 1 unit of input k from firm j' and gave it to firm j. Firm j could then forgo an amount of input k' equal to its MRTS, and keep output unchanged. Since j has steeper isoquant at A than j', it is clear that firm j can afford to give up more of input k' than firm j' needs to be given to compensate her for having the one unit of k taken away. So after fully compensating firm j', there are still some units of input k' remaining. These can either be allocated to firm j or to j' or it can be split between both firms. Either way, we have reallocated inputs in such a way as to (weakly) increase the quantities provided by both firms. Clearly, allowing the inefficient production plan could not be optimal.

It turns out that there are (infinitely) many Pareto optimal allocations of inputs. We refer to the set of Pareto optimal allocations as the **Contract Curve in Production**. At every point along the contract curve, there are isoquants for each firm that are tangent to each other. Note that the allocations O_j (where firm j receives nothing and firm j' receives all inputs) and $O_{j'}$ (where firm j' receives nothing, and j receives all inputs) must lie on the contract curve. (These allocations must be Pareto optimal, since it is impossible to increase the output of one firm without decreasing the output of the other). As we move along the production contract curve from O_j to $O_{j'}$, we move from optimal input allocations which favor production by firm j', to those which favour production by firm j.

Definition 27. The **Production Possibility Set** represents the set of output bundles that can be produced given different feasible allocations of inputs between the firms. The **Production Possibility Frontier** represents the set of **efficient** output bundles that are possible.

Along the production possibility frontier, if the quantity of some good increases, then the quantity of some other good must decrease. As we move along the contract curve in production, we span out the production possibility frontier.

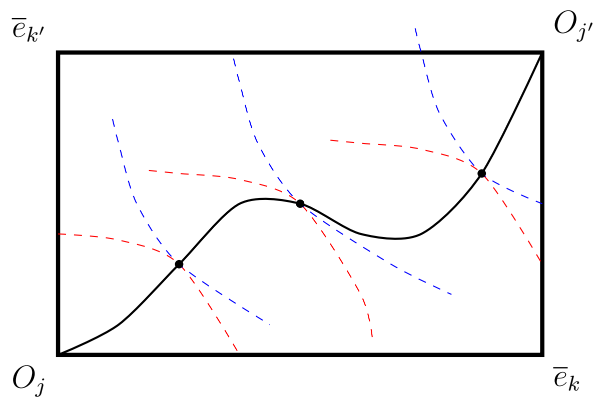

Figure 12.2: Contract Curve in Production

Suppose we wish to increase firm j's output by 1 unit and we do so by increasing the amount of k allocated to firm j. From the above discussion, we know that we must increase firm j's allocation of k by $\frac{1}{MP_k^j}$. Taking this amount of k away from firm j' causes its output to decrease by $\frac{MP_k^{j'}}{MP_k^j}$. If instead, we made the transformation by reallocating some other input k', firm 2's output would decrease by $\frac{MP_{k'}^{j'}}{MP_{k'}^j}$. But, along the contract curve in production, these ratios are equal. Hence, it doesn't matter which inputs we change to increase firm 1's output —in equilibrium, firm 2's output will decrease by the same amount.

Definition 28. The **marginal rate of transformation** is the quantity of good j' that must be forgone to increase the quantity of good j by 1 unit, assuming efficient input mix.

$$MRT_{j,j'} = -\frac{MP_k^{j'}}{MP_k^j}$$

for any input k.

The marginal rate of transformation is the slope of the production possibility frontier. See Figure 12.3.

12.2.2 Efficient Output Allocation

Now, consider any point along the Production Possibility Frontier (i.e. Contract Curve in Production) and suppose the Pareto Planner allocates inputs in the way that produces that output mix. How should the Planner then allocate the resulting output amongst the various agents?

To see what is required here, take the first FOC from the Planner's problem. Since this condition must hold for any two outputs j and j', we have: $\gamma_i MU_j^i(x^i) = \lambda_j^x$ and $\gamma_i MU_{j'}^i(x^i) =$

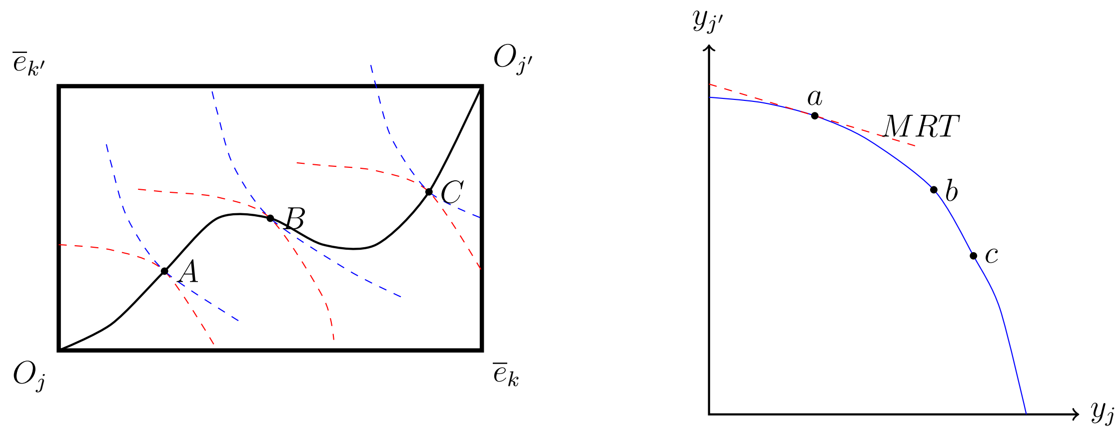

Figure 12.3: Production Possibility Frontier

$\lambda^x_{j'}$. Combining these (substituting out for γ_i) gives:

$$\frac{MU^i_j(x^i)}{MU^i_{j'}(x^i)} = \frac{\lambda^x_j}{\lambda^x_{j'}}$$

Notice two things. First, the left hand side is the ratio of marginal utilities, which we know is the marginal rate of substitution for agent i. Second, the right hand side is independent of which agent we are analyzing. Hence, the MRS must equal the same constant for all agents, and so the MRS's must be equal for all agents. We have:

$$MRS^i_{j,j'}(x^i) = -\frac{MU^i_j(x^i)}{MU^i_{j'}(x^i)} = -\frac{MU^{i'}_j(x^{i'})}{MU^{i'}_{j'}(x^{i'})} = MRS^{i'}_{j,j'}(x^{i'})$$

But recall, the marginal rate of substitution is the amount of good j' that the agent can forgo in order to keep utility unchanged if she consumes one more unit of good j. Pareto optimality requires that all agents be willing to trade-off outputs at the same rate.

As before, we can represent the set of all such possible allocations using an **Edgeworth Box in Consumption**. The dimensions of the box are $\overline{x} \times \overline{x'}$ - which represent the total supply of goods x and x' (given the point on the contract curve that was selected). Horizontal distances in the Edgeworth box represent the consumption of good x, whilst vertical distances represent the consumption of good x'. The consumption of agent i is measured from the bottom-left vertex O_i, whilst the consumption of agent i' is measured from the top-right vertex $O_{i'}$. See Figure 12.4.

We can represent each agent's preferences over allocations using indifference curves. Indifference curves retain their usual properties (downward sloping, convex and non-intersecting). As usual, indifference curves which are further away from an agent's origin, represent more preferred bundles. The analogy between the Edgeworth Box in Consumption and Production should be evident. The results are similarly analogous.

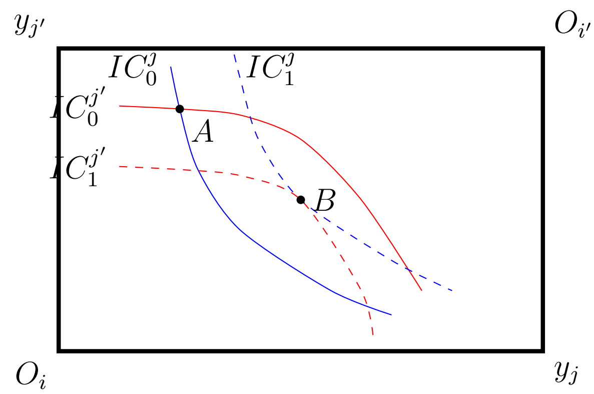

Figure 12.4: Edgeworth Box in Consumption

Lemma 14. *An allocation of goods is Pareto optimal if the indifference curves of each agent going through the allocation are tangent to one another.*

To understand this result, recall that the slope of an indifference curve (i.e. MRS) represents the rate at which agents are willing to forgo good j' in order to get an additional unit of good j such that they remain just as well off. Suppose $MRS^i_{j,j'} \neq MRS^{i'}_{j,j'}$. For concreteness, let $MRS^i = 3$ and $MRS^{i'} = 2$. Then, agent i is willing to forgo 3 units of good j' to get an additional unit of good j, whilst agent i' is only willing to forgo 2 units of good j'. Now, suppose we reallocate goods such that agent i receives one more unit of good j, and agent i' receives one fewer unit of j. i is willing to give up 3 units of j' in return for the additional unit of j. However, i' only requires an additional 2 units of j' to compensate him for the loss of good j. Hence, we can take 3 units of j' away from i, and give 2 of these units to i' - and leave both agents just as well off as before. But we still have one unit of j' remaining. We can either give this entirely to agent i, or entirely to agent i' - or share it between i and i' (in some proportion). Either way, it is always possible to make at least one agent better off without making any agent worse off. Such a reallocation will be possible whenever the MRS's are not the same.

It turns out that there are (infinitely) many Pareto optimal allocations of goods. We refer to the set of Pareto optimal allocations as the **Contract Curve**. At every point along the contract curve, there are indifference curves for each agent which are tangent to each other. We represent this in the following diagram. (The blue lines represent the indifference map of agent i whilst the red lines represent the indifference map of agent i'. The thick line is the contract curve). Note that the allocations O_i (where agent i receives nothing and agent i' receives everything) and $O_{i'}$ (where agent i' receives nothing, and i receives everything) must lie on the contract curve. (These allocations must be Pareto optimal, since it is impossible to make one agent better off without making another worse off). As we move along the contract curve from O_i to $O_{i'}$, we move from optimal allocations which favour agent i' to those which favour agent i. See Figure 12.5.

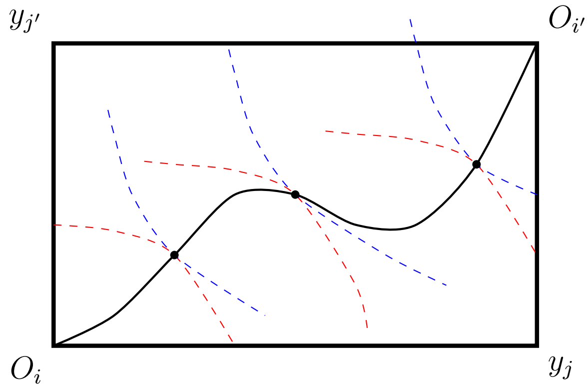

Figure 12.5: Contract Curve in Consumption

12.2.3 Efficient Output Mix

Suppose the planner allocates inputs to firms efficiently, so that it is impossible to reallocate them in a way that (weakly) increasing production by all firms. Given the output produced, suppose the planner also allocates these amongst agents efficiently, so that it is impossible to reallocated them in a way that makes all agents (weakly) better off. Are we done? There is an additional margin along which we might be able to improve things. It could be that the mix of outputs produced by firms does not match the output mix preferred by consumers. E.g. it could be that the agents have a strong preference for one good, but that inputs are allocated s.t. other goods are produced in high quantity instead.

To see what is required here, return to the first order conditions from the planner's problem. By the first order condition for inputs, for any two firms j and j', we have: $\lambda_j^x MP_k^j(z^j) = \lambda_k^z$ and $\lambda_{j'}^x MP_k^j(z^{j'}) = \lambda_k^z$. Combining these (substituting out for λ_k^z) gives:

$$\frac{\lambda_j^x}{\lambda_{j'}^x} \cdot \frac{MP_k^j(z^j)}{MP_k^{j'}(z^{j'})} = 1$$

Next, by the first order conditions for outputs, for any two goods j and j', we have: $\gamma_i MU_j^i(x^i) = \lambda_j^x$ and $\gamma_i MU_{j'}^i(x^i) = \lambda_{j'}^x$. Combining these (substituting out for γ_i) gives:

$$\frac{MU_j^i(x^i)}{MU_{j'}^i(x^i)} = \frac{\lambda_j^x}{\lambda_{j'}^x}$$

Notice that the same ratio of Lagrange multipliers is present in both equations. Solving out for this gives:

$$MRS = -\frac{MU_j^i(x^i)}{MU_{j'}^i(x^i)} = -\frac{MP_k^{j'}(z^{j'})}{MP_k^j(z^j)} = MRT$$

Pareto optimality requires that the marginal rate of transformation (i.e. the slope of the Pareto Frontier) coincide with the marginal rate of substitution for each agent. (Efficient

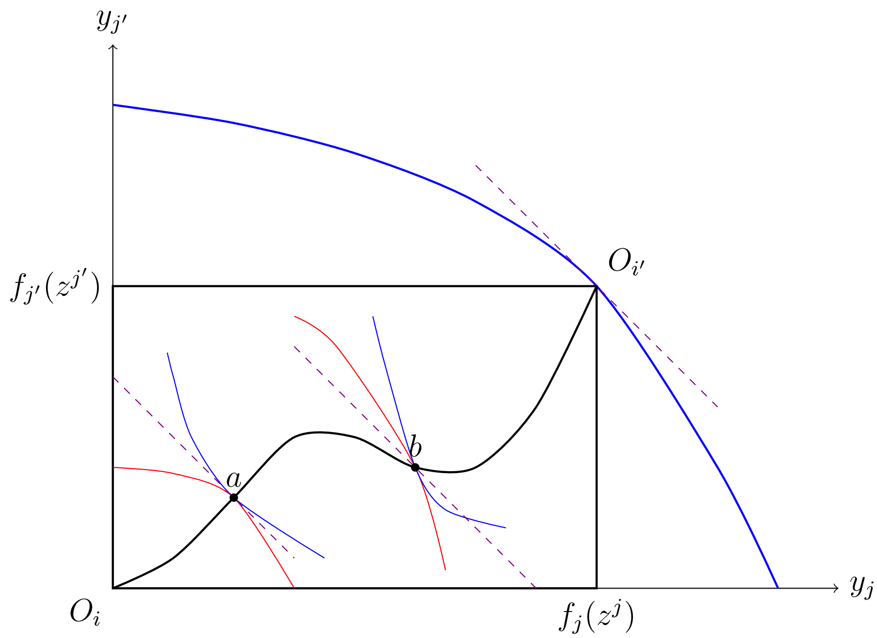

Figure 12.6: Efficient Output Mix

output allocation requires that all MRS's coincide, so it suffices that MRS=MRT for *some* agent.) Agents must be willing to trade off goods at the same rate that the market is able to transform one into the other, given the production processes of the different firms producing different goods. See Figure 12.6.

Intuitively, if this wasn't the case, there would exist a re-allocation of inputs that produces a different mix of outputs that is more preferred by the agents. For example, consider a point along the contract curve in consumption where $MRS^i_{j,j'} = -3$ and a point on the Pareto frontier where $MRT_{j,j'} = -2$. Then the agents would each be willing to give up 3 units of j' for one additional unit of j. But the PPF requires that they only forgo 2 units. So, we could allocate slightly more inputs to firm j, so that it produces one more unit of output; and do so by taking away inputs from firm j', causing its output to decrease by 2. We then give this additional unit of j to some consumer, who would willingly trade it for 3 units of j'. Of course, firm j' is producing 2 fewer units of j', which means there is still one unit of j' available. This can be allocated to any of the agents or shared amongst them. In any case, it is clear that we can make some agents better off without making any worse off.

In summary, an allocation is Pareto efficient if it satisfies:

1. Efficient Input Allocation
   - It is impossible to reallocate inputs in such a way as to increase output of some goods without decreasing output of others.

$$MRTS^j_{k,k'} = MRTS^{j'}_{k,k'} \;\; \forall k, k' \; \forall j, j'$$

2. Efficient Output Allocation

   - Taking the outputs (from efficient input allocation) as given, it is impossible to reallocate outputs in such a way as to increase utility of some agents without decreasing utility of others.

$$MRS^i_{j,j'} = MRS^{i'}_{j,j'} \ \forall j, j' \ \forall i, i'$$

3. Efficient Output Mix

   - It is impossible to efficiently re-allocate inputs to change the output mix, so that the resulting efficient output allocation improves the utility of some without reduces the utility of others.

$$MRS^i_{j,j'} = MRT_{j,j'} \ \forall j, j' \ \forall i$$

12.3 Competitive Equilibrium

In the previous section, we found the set of efficient allocations by asking how a hypothetical benevolent social planner would allocate inputs and outputs. One of the things that you will notice about our characterization of Pareto optima is that prices are nowhere to be seen. This is to be expected. Prices are a mechanism to help allocate resources when decisions are in a de-centralized manner. [I get to choose what I want, subject to my budget; and you get to choose what you want given your budget. But our choices are naturally directed by prices.] The Pareto planner had no need for such price signals — she was free to allocate goods however she saw fit, and was (assumed to be) perfectly informed, so she knew who wanted what and how much.

But, of course, the benevolent social planner doesn't exist. In this section, we ask whether we can achieve the same outcomes as the Pareto planner would choose, using the price mechanism (where each agent is free to make choices in their best interest). Such a 'decentralized' institution has the benefit of not needing a good, powerful and all knowing planner to make allocative choices. Naturally, prices will now feature prominently in our analysis.

Definition 29. A **competitive equilibrium** consists of (i) an allocation $(\{x^i\}, \{y_j\}, \{z^j\})$, and (ii) input and output prices w and p s.t.:

1. Each household i chooses x^i to maximize its utility subject to its budget constraint, taking prices as given:
$$\max_{x^i} u_i(x^i) \ \ s.t. p \cdot x^i = w \cdot e^i$$

2. Each firm j chooses (y^j, z^j) to maximize its profit subject to its technology, taking prices as given:
$$\max_{y_j, z^j} p_j y_j - w \cdot z^j \ \ s.t. y_j = f^j(z^j)$$

3. Prices are s.t. markets clear:

$$\sum_i x_j^i = y_j = f^j(z^j) \quad \forall j = 1, ..., n$$
$$\sum_j z_k^j = \bar{e}_k = \sum_i e_k^i \quad \forall k = 1, ..., m$$

In a competitive equilibrium, each household chooses the bundle to purchase to maximize its individual utility, taking its budget as given. (The household's income is derived from selling its endowment of inputs on competitive factor markets.) Similarly, each firm chooses the quantity of output to produce, and hires a bundle of inputs, in order to maximize its profits, taking its production technology as given. In both cases, households and firms take prices as given, and do not try to manipulate prices through their choices. Prices themselves are set by the 'invisible hand' of the market to ensure that markets clear (i.e. the demand equals supply).

We illustrate the basic idea of a competitive equilibrium in Figure 12.7, below. For simplicity, we consider an endowment economy, where agents i and i' are endowed with quantities of goods j and j'. Let the endowment be e. The agents are free to trade those goods at prices determined by the market. The purple line passing through e has slope equal to the relative prices, and represents the set of all bundles that can be achieved through some trade at market prices. Since each household maximizes utility subject to its budget constraint, it will work towards a trade that puts it on the highest possible indifference curve achievable along the purple line.

In the left panel, the prices are inconsistent with market clearance. The quantity of good j demanded exceeds its supply, whilst the quantity of good j' demand is less than its supply. There is an excess demand for j and over-supply of j'. Intuitively, the price of good j must be too low relative to j'. In the right panel, the market is equilibrium, and demand for each good is equal to its supply. (Moreover, the budget line is steeper, indicating that relative prices make good j less attractive.)

The right panel reveals an important insight. If households are each utility maximizing *and* if markets clear, it must be that the agents' indifference curves are tangent to the budget line at the same point. But this means the agents' indifference curves must be tangent to one another, which is the condition for Pareto optimality!

In fact, this result is true more generally. In a competitive equilibrium, we have the following:

1. Since each household i is utility maximizing, it must be that $MRS^i_{j,j'} = -\frac{p_j}{p_{j'}}$, for each i and each pair of goods j, j'. But this implies that $MRS^i_{j,j'} = MRS^{i'}_{j,j'}$ for all households and all goods. The competitive equilibrium allocation of goods lies along the contract curve in consumption.

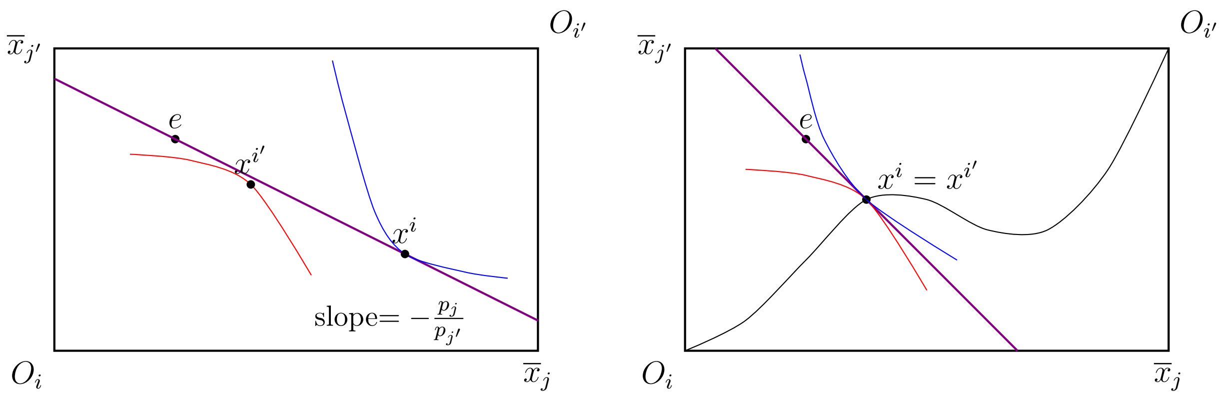

Figure 12.7: Optimal choices in an exchange economy given different prices. In the left panel, the relative prices are inconsistent with equilibrium. There is an excess demand for good j and an excess supply of good i. The right panel is consistent with equilibrium.

2. Since each firm j is profit maximizing, it must be that $MRTS^j_{k,k'} = -\frac{w_k}{w'_k}$, for each j and each pair of inputs $k.k'$. But this implies that $MRTS^j_{k,k'} = MRTS^{j'}_{k,k'}$ for all firms and all inputs. The competitive equilibrium allocation of inputs lies along the contract curve in prodction.

3. Since each firm j is profit maximizing, it must be the $p_j MP^j_k = w_k$ for each j and each input k. Hence $\frac{p_j}{p_{j'}} \cdot \frac{MP^j_k}{MP^{j'}_k} = 1$ and so:

$$MRS^i_{j,j'} = -\frac{p_j}{p_{j'}} = -\frac{MP^{j'}_k}{MP^j_k} = MRT_{j,j'}$$

which implies that the output mix is efficient.

Taken together, we have the following famous result: (First Fundamental Theorem of Welfare Economics) Every competitive equilibrium allocation is Pareto optimal.

The first welfare theorem is a formal statement of the insight from Adam Smith that competitive markets, left to their own devices, will allocate inputs and goods in a way that maximizes the utilities market participants. It is important to recognize the caveats sitting behind this result:

1. Households and firms are price takers: their actions do not affect the prices of outputs and inputs. [This assumption is often reasonable in a market with many buyers and many sellers, each of whom is such a small part of the market, that any change in their behavior would have a negligible effect on the overall market price.]

2. There is no price discrimination. The logic above required that all agents were willing to trading-off goods/inputs at the market rate, and so were trading off goods at the same rate as one another. Of course, if different agents faced different prices, this logic would not hold.

3. There are no externalities. Agents' utilities depend on the bundle they consume, and not on the consumption of others.

We can also see this result by inspecting the optimization problems underlying the Pareto planner's choice and the competitive equilibrium. Writing these problems side-by-side, we have:

Pareto Optimality

Planner's Problem:

$$\max_{\{x^i\}.\{z^j\}} \sum_i \gamma_i u_i(x^i)$$
$$s.t. \sum_i x^i_j = f_j(z^k) \;\; \forall j$$
$$\sum_j z^j_k = \bar{e}_k \;\; \forall k$$

Competitive Equilibrium

1) Each household maximizes:

$$\max_{x^i} u_i(x^i) \;\; s.t. \;\; p \cdot x^i = w \cdot e^i$$

2) Each firm maximizes:

$$\max_{z^j} p_j f_j(z^j) - w \cdot z^j$$

3) Markets clear:

$$\sum_i x^i_j = y_j = f^j(z^j) \quad \forall j$$
$$\sum_j z^j_k = \bar{e}_k = \sum_i e^i_k \quad \forall k$$

Notice that the market clearance condition is present in both problems. If we now write the first order conditions side-by-side, we get:

Pareto Optimality

Planner's Problem FOCs:

$$\frac{\partial \mathcal{L}}{\partial x^i_j} = \gamma_i \frac{\partial u_i(x^i)}{\partial x_j} - \lambda^x_j = 0 \;\; \forall i, j$$
$$\frac{\partial \mathcal{L}}{\partial z^j_k} = \lambda^x_j \frac{\partial f_j(z^j)}{\partial z_k} - \lambda^z_k = 0 \;\; \forall j, k$$

Competitive Equilibrium

Competitive Equilibrium FOCs:

$$\frac{\partial \mathcal{L}^i}{\partial x^i_j} = \frac{\partial u_i(x^i)}{\partial x_j} - \mu_i p_j = 0 \;\; \forall i, j$$
$$\frac{\partial \mathcal{L}^j}{\partial z^j_k} = p_j \frac{\partial f_j(z^j)}{\partial z_k} - w_k = 0 \;\; \forall j, k$$

where μ_i is the Lagrange multiplier associated with household i's utility maximization problem. Notice that the first order conditions from both problems are very similar to one another. In fact, if we set $p_j = \lambda^x_j$ for each j, $w_k = \lambda^z_k$ for each k, and $\gamma_i = \frac{1}{\mu_i}$, then the two first order conditions exactly coincide. But if so, their solutions must coincide as well. So the competitive allocation will be Pareto optimal.

What does it mean to set $p_j = \lambda^x_j$ for all j? We are requiring the prices of goods in the market to reflect the social opportunity cost (to the Pareto planner) of allocating a unit of good j, and thus having one fewer unit to allocate to other agents. Prices in a competitive market reflect social opportunity costs. And since, agents will only purchase a good if they value it by more than the price, then everyone who ends up with a good must value it by more than its social opportunity cost. There are no mis-allocations which, by remedying, we could make everyone better off.

Similarly, by setting $w_k = \lambda_k^z$, we are requiring the prices of all inputs to reflect the social opportunity cost of allocating a unit of that input, and thus having one fewer unit to allocate to other firms. Again, prices reflect opportunity costs!

The first welfare theorem made a positive (i.e. predictive) comment about competitive equilibria — that they would coincide with *some* Pareto optimal allocation. We know, however, that there are many Pareto optima, and some are 'fairer' than others. What if the Pareto optimum implied by the competitive market is far from the Pareto optimum we wish we could implement? Does this justify government intervention in the market?

[Second Fundamental Theorem of Welfare] If agents' preferences and firms' technology satisfies certain regularity conditions, then any Pareto optimal allocation can be supported as a competitive equilibrium with transfers.

The second welfare theorem states that if we transfer wealth between individuals (or if we reallocate the original endowment between individuals), then we can force the competitive equilibrium to coincide with any of the Pareto optimal allocations we desire. The second welfare theorem opens the door for government involvement in the market. Whilst the competitive market is Pareto efficient —it may not generate our most favoured Pareto outcome. With a little tinkering - the Government can force the market to choose outcomes that are socially more desirable. (Note - the second welfare theorem does not prescribe that the government should be involved in affecting the prices of goods - this is the role of the market. However - by providing income transfers between individuals, the government can cause different prices to arise in the competitive equilibrium - and as such for a different Pareto allocation to be achieved.)

Chapter 13

Monopoly

In the analysis up to now, we have assumed that agents are price takers, i.e., the prices of goods are unaffected by their choices about how much to produce or consume. In this chapter, we consider decision-making under conditions of price-making power —where the decision-maker understands that their choices will affect the prices they pay/receive.

We being with the case of a monopoly; where the demand for the firm's good and market-wide demand are one and the same. If the firm wishes to increase output, to entice consumers to purchase the new output, it must reduce the price it charges (on all goods).

Let $p(q)$ denote the firm/market (inverse) demand function. We assume that $p'(q) < 0$, so that demand is downward sloping. Additionally, we make the assumption that the firm's revenue function $r(q) = p(q)q$ is concave (i.e. hill shaped). This requires:

$$r''(q) = 2p'(q) + qp''(q) < 0$$

Now, the first term is negative by the assumption of downward sloping demand, so this condition holds provided the second term is not too positive (i.e. the demand function is not too convex).

The monopolist has a cost function $c(q)$ that satisfies the standard assumptions: $c'(q) > 0$ (marginal cost is positive) and $c''(q) \geq 0$ (increasing marginal cost). The firm's profit is:

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

or alternatively

$$\pi = \max_p p \cdot q(p) - c(q(p))$$

where $q(p)$ is the demand function. Although the second formulation may seem more natural —firms choose prices —since the cost function is defined in terms of quantity, the firm formulation is often easier to work with.

The first order condition is:

$$p(q) + qp'(q) = c'(q)$$

Briefly, let us verify that the second order condition is satisfied, so that the FOC select a maximum. We have:

$$\pi'' = \underbrace{2p'(q) - qp''(q)}_{r''(q)} - c''(q)$$

which is the sum of two terms. The first term is negative (by the assumption that r is concave). The second term is also non-positive since $c' > 0$.

Return to the first order condition. Notice the difference between this expression and the profit-maximizing condition for a price-taking firm ($p = c'(q)$). For a price-taking firm, its revenue is $r(q) = pq$ and so the marginal revenue is simply the price ($r'(q) = p$). (This should be intuitive; each additional unit of output generates \$$p$ of revenue.)

By contrast, for the price-making monopolist, the revenue is $r(q) = p(q)q$. The marginal revenue, $r'(q) = p(q) + qp'(q)$, is the sum of two terms. The first term is the new revenue that the firm earns by selling one additional unit of output; this is simply the price of the marginal unit. (We call this as the *extensive margin*). The second term is the revenue lost on the previous units sold because the firm had to lower its price to sell the marginal unit. To sell one more unit, the demand curve requires that the price fall by $p'(q)$. Each of the q units previous sold now earn this many fewer dollars of revenue. (We call this the *intensive margin.*)

Notice that marginal revenue on the intensive margin is negative. Thus $r'(q) = p(q) - qp'(q) < p(q)$. When the firm increases output, the additional revenue earned is less the price commanded by the marginal unit. The marginal revenue curve lies below the demand curve.

Now, since profit maximization requires that the firm equate marginal revenue and marginal cost, this implies that, at the optimum $p(q^*) > c'(q^*)$. Recall —one of the conditions for Pareto optimality is that $p = MC$. Thus, the price-making power of a monopoly causes the resulting allocation to be inefficient. The rough intuition is as follows: efficiency requires that all opportunities for mutually beneficial gains from trade be exhausted. But if $p > c'(q)$, then the cost of producing the next unit of the good will be lower than the benefit to the marginal agent of consuming it. The monopoly output is below the socially efficient level.

The profit maximization condition reveals further insights if expressed in elasticity form. Recall, the elasticity of demand is $\varepsilon = \frac{\partial q}{\partial p} \cdot \frac{p}{q}$ and so $\frac{1}{\varepsilon} = \frac{\partial p}{\partial q} \cdot \frac{q}{p}$. First, note that:

$$MR = p(q) + qp'(q) = p(q)[1 + p'(q) \cdot \frac{q}{p(q)} = p(q)\left[1 + \frac{1}{\varepsilon}\right]$$

The firm's marginal revenue is intimately connected with the elasticity of demand. This should be intuitive. Suppose q increases by 1%. What is the effect on revenue? If prices were unchanged, the higher quantity would result in larger revenue. Put, by the demand curve, to sustain higher output, prices must fall, and this will cause revenue to fall. The overall effect is ambiguous, and depends on the relative sizes of the change in price and change in quantity —which is measured by the elasticity of demand.

In fact, we can say more. Marginal revenue is positive if $1+\frac{1}{\varepsilon} > 0$, which requires that $\varepsilon < -1$ (where we assume, by the law of demand that $\varepsilon < 0$). Hence, marginal revenue is positive if demand is elastic and negative otherwise. This, again is intuitive. If demand is elastic, then quantity changes by more (along the demand curve) than price, and so the quantity effect described in the previous paragraph will dominate. The opposite is true if demand is inelastic. We can re-state this result a different way: Suppose the monopolist increases its price (which causes quantity demanded to decrease). Will this cause revenue to increase or decrease? Revenue will increase as long as quantity falls by less (in percentage terms) than price increases —i.e. demand is inelastic. Thus, when demand is inelastic, increasing price increases revenue. By contrast, if demand is elastic, increasing price decreases revenue.

Return to the profit-maximization condition:

$$\begin{aligned} p(q) + p(q) \cdot p'(q) \cdot \frac{q}{p(q)} &= c'(q) \\ p(q)\left[1 + \frac{1}{\varepsilon}\right] &= c'(q) \\ p(q) &= \frac{\varepsilon}{\varepsilon + 1} c'(q) \end{aligned}$$

Several insights follow from this expression. First, a price-making firm sets its price by applying a mark-up over its marginal cost. This is seemingly descriptive of how many firms (especially smaller businesses) actually operate. It doesn't matter that the business owner may have no concept of calculus or first order conditions; profit maximization is equivalent to applying a simple rule of thumb: for each good, set prices by applying a markup to the item's cost. Moreover, the size of the markup is related to the elasticity of demand. The more elastic is demand, the lower will be the markup. [In fact, in the limit as $\varepsilon \to \infty$, so that the firm effectively becomes a price-taker, the markup shrinks to zero, so that the firm prices at marginal cost!]

Second, if the firm is profit maximizing, it must choose a quantity at which demand is *elastic*. (We can easily check that if demand were inelastic, then $1 + \frac{1}{\varepsilon} < 0$, which implies the firm should set a negative price, which obviously cannot be.) This seems somewhat counter-intuitive. We often associate monopolies with inelastic demand. And yet —if the firm did produce in the inelastic region of the demand curve, we know that by increasing price and decreasing quantity, it will increase revenue and decrease costs (since it is producing less), and this clearly increases profits. Thus, whenever demand is inelastic, the firm should exploit this by raising prices. If the firm has stopped raising prices, it must be that it has entered the elastic region of the demand curve. [Now, raising prices further causes revenue and costs to both decrease. It will stop raising prices if the former decreases by more than the latter.]

Re-arranging the profit-maximization condition gives:

$$\underbrace{\frac{p(q) - c'(q)}{p(q)}}_{\text{Lerner Index}} = -\frac{1}{\varepsilon}$$

The LHS defines the *Lerner Index*, which is the fraction of the output price that the firm earns as profit (at the margin). This is a measure of the firm's market power. The larger is the Lerner Index, the greater the ability of the firm to raise prices over marginal cost (in percentage terms). The expression reveals that market power is inversely related to the elasticity of demand. As demand becomes more elastic, the Lerner index (and thus markups) falls, such that in the limit as $\varepsilon \to \infty$ the $p(q) \to c'(q)$. (As demand becomes perfectly elastic, the monopolists behavior converges to that of a perfectly competitive firm.)

Example 45. A monopolist faces an (inverse) demand curve $p(q) = \alpha - \beta q$. For simplicity, suppose the firm has a constant marginal cost production technology $c(q) = cq + f$. [As should be evident, what's new in the study of monopoly is the shape of the demand curve. So, to focus attention on the role of demand, we often make the cost side super simple, to ensure that any insights we observe are not coming from funky cost behavior.] The firm's problem is:

$$\max_q (\alpha - \beta q)q - (cq + f)$$

. The first order condition is

$$\begin{aligned}(\alpha - \beta q) - \beta q - c = 0\\ q^m = \frac{\alpha - c}{2\beta}\end{aligned}$$

And $p^m = \frac{\alpha+c}{2}$.

Instead, suppose we found the optimum by applying the inverse elasticity rule. The elasticity of demand is $\varepsilon = -\frac{1}{\beta} \cdot \frac{p}{q}$. We have:

$$\begin{aligned}\frac{p - c}{p} &= \frac{\beta q}{p}\\ p - c &= \beta q\\ (\alpha - \beta q) - c &= \beta q\\ q^m &= \frac{\alpha - c}{2\beta}\end{aligned}$$

which coincides with our expression above.

13.1 Two-Part Tariff

Why does the monopolist raise prices over marginal cost? Although pricing at marginal cost is efficient in that it maximizes the volume of utility improving transactions, the firm itself makes a low profit, whilst the consumers earn a large consumer surplus. The firm would like to appropriate some of this consumer surplus. It can do so by raising the price. This has two opposing effects. On the one hand, raising the price means the firm appropriates more of the surplus of those consumers who continue to purchase the product. On the other hand, with a higher price, there are fewer trades and so the firm loses profit on the margin. The monopolist prices to optimize this trade-off.

In the preceding analysis we assumed that the only instrument that the monopolist has to extract profit/surplus is the price of the good. Now, suppose it has two instruments: it can choose a price as well as a fixed up-front fee (e.g. a membership fee). Can the monopolist do better with a two-part tariff?

Start with a simple case. Suppose all consumers are identical with individual demand function $q(p)$. At price p, each consumer earns a consumer surplus of $CS(p) = \int_p^\infty q(\pi)d\pi$ (since the consumer surplus is the area under the demand curve, given the quantity purchased). Hence, if the firm charges a price p, it can afford to charge a fixed fee of up to $CS(p)$. (Any fee larger than this will cause the net benefit to the consumer of purchasing the good to be negative.) Naturally, the firm will charge the largest fixed fee that doesn't deter consumption. The firm's problem is:

$$\max_p pq(p) - c(q(p)) + CS(p) = \max_p pq(p) - c(q(p)) + \int_p^\infty q(\pi)d\pi$$

The first order condition is:

$$\begin{aligned} q(p) + [p - c'(q(p))]q'(p) - q(p) &= 0 \\ [p - c'(q(p))]q'(p) &= 0 \\ p &= c'(q(p)) \end{aligned}$$

We return to marginal-cost pricing! What is the intuition? The reason for raising the price above marginal cost was to extract a larger fraction of the consumer surplus. But, now, the firm needn't use the price to do this, since it also has the up-front fee at its disposal. Since the fee is chosen to expropriate the entire consumer surplus, the firm should set the price that maximizes consumer surplus (without making a loss). Clearly this requires setting $p = c'(q(p))$.

Does this solve the problem? Not entirely. What if consumers are not homogeneous in their preferences? For simplicity, suppose there are two types of consumers. Low type preferences have individual demand functions $q_L(p)$, whilst high type preferences has demand function $q_H(p)$, with $q_H(p) > q_L(p)$. Suppose there are n_L and n_H low and high type agents, respectively, and let $n = n_L + n_H$ denote the total number of agents. Average individual market demand is $\overline{q}(p) = \frac{1}{n}[n_L q_L(p) + n_H q_H(p)]$ and total market demand is $n\overline{q}(p)$. For any price, let $CS_i = \int_p^\infty q_i(\pi)d\pi$ be the consumer surplus for a type $i \in \{L, H\}$ agent.

Suppose the firm sets price p. Then for any up-front fee $F \leq CS_L(p)$, both types of agents will buy the good. If the up-front fee is $F \in (CS_L(p), CS_H(p]$, then only the high type agent will purchase the good, and noone will purchase if the fee is $F > CS_H(p)$. I focus on situations where n_L is large enough that the firm will want to sell to both types and not price low types out of the market. Naturally, the fixed fee should be set at $F = CS_L(p)$. The firm's problem is:

$$\max_p \; n_L[CS_L(p) + pq_H(p) + CL(p)] + n_H[CS_L(p) + (p-c)q_H(p)] - c(n_Lq_L(p) + n_Hq_H(p))$$
$$= \max_p \; n\left[\int_p^\infty q_L(s)ds + p\overline{q}(p)\right] - c(n\overline{q}(p))$$

The first order condition is:

$$\begin{aligned} -nq_L(p) + n\overline{q}(p) + (p - c'(n\overline{q}(p)))n\overline{q}'(p) &= 0 \\ p - c'(n\overline{q}(p))\overline{q}'(p) &= -[\overline{q}(p) - q_L(p)] \\ \frac{p - c'(n\overline{q}(p))}{p} \cdot \overline{q}'(p)\frac{p}{\overline{q}(p)} &= -\frac{\overline{q}(p) - q_L(p)}{\overline{q}(p)} \\ \frac{p - MC}{p} &= -\frac{1}{\overline{\varepsilon}} \cdot \left[1 - \frac{q_L(p)}{\overline{q}(p)}\right] \end{aligned}$$

where $\overline{\varepsilon}$ is the average individual elasticity of demand. Note that $\left[1 - \frac{q_L(p)}{\overline{q}(p)}\right] \in (0,1)$ since $- < q_L(p) < \overline{q}(p)$. This expression is a modified version of the inverse elasticity rule. The left hand side is the Lerner index (i.e. the percentage markup). The right hand side is the inverse of the elasticity scaled by a factor less than 1. Thus, the mark-up will be lower than in the case where the firm's only instrument is the price. The firm still prices above marginal cost, but not as much.

What is the intuition? If consumers have homogeneous preferences, there is no need to raise prices above marginal cost; the monopolist instead uses the fee to extract all the consumer surplus. With heterogeneous agents, it can at most extract all of the consumer surplus from the lower-willingness-to-pay consumer, thus leaving some of the high-willingness-to-pay consumer's surplus on the table. The only way to extract this is, again, to raise prices above marginal cost, inviting a similar trade-off as before. The only difference is that much of the surplus is already extracted through the fee, so the benefit of raising prices is lower. Thus, the monopolist prices according to a scaled/dampened inverse elasticity rule. Moreover, the scale/dampening factor depends on the amount of fees that the monopolist forgoes (at the margin), relative to what it receives; which is a function of the ratio of the demand of the low-type to the average demand.

Example 46. There are n_L low type and n_H high type consumers, with $n = n_L + n_H$. A type $i \in \{L, H\}$ agent has demand function $q_i(p) = \frac{\alpha_i}{p^2}$ where $\alpha_H > \alpha_L$. Denote $\overline{\alpha} = \frac{n_H\alpha_H + n_L\alpha_L}{n_H + n_L}$. As before, the elasticity of demand (for both types) is $\varepsilon = -2$. Market demand is:

$$q(p) = n_Lq_L(p) + n_Hq_H(p) = \frac{n_L\alpha_L + n_H\alpha_H}{p^2} = \frac{n\overline{\alpha}}{p^2}$$

The firm has a constant marginal cost production technology. If the monopolist simply maximizes w.r.t. a price, we know that it applies a mark-up inversely proportional to the elasticity. We have:

$$p^m = \frac{\varepsilon}{\varepsilon+1}c = 2c$$

The monopoly output is:

$$q^m = q(p^m) = \frac{n\overline{\alpha}}{4c^2}$$

The firm's profit is:

$$\pi^m = (p^m - c)q^m = \frac{n\overline{\alpha}}{4c}$$

Next, suppose the monopolist chooses the optimal two-part tariff. If the price is p, the type i consumer's surplus is $CS_i(p) = \int_p^\infty q_i(s)ds = \frac{\alpha_i}{p}$. The firm maximizes:

$$\max_p \left\{ n\frac{\alpha_L}{p} + \frac{n\overline{\alpha}}{p^2}(p-c) \right\}$$

The first order condition is:

$$-\frac{n\alpha_L}{p^2} + \frac{n\overline{\alpha}}{p^2} - 2\frac{n\overline{\alpha}}{p^3}(p-c) = 0$$

$$1 - \frac{\alpha_L}{\overline{\alpha}} - 2\frac{p-c}{p} = 0$$

$$p\left(1 - \frac{\alpha_L}{\overline{\alpha}}\right) - 2(p-c) = 0$$

And so:

$$p^{2PT} = \left(\frac{1}{1+\frac{\alpha_L}{\overline{\alpha}}}\right) \cdot 2c$$

where the term in parenthesis is the dampening factor. Notice that the dampening factor depends on the ratio of α_L to average α. As $\alpha_L \to \overline{\alpha}$ (i.e. $\alpha_L \to \alpha_H$ and so the consumers become homogeneous), then the dampening factor converges to $\frac{1}{2}$ and so $p = MC$. As $\alpha_L \to 0$, then the monopolist cannot extract any consumer surplus from the low type, so the fixed fee is useless. $p = 2c$. There is no dampening.

13.2 Price Discrimination

In all of the above analysis, we have assumed that the monopolist must charge the same price and same fee to all agents. What if the monopolist was able to charge different prices to different agents? We call this price discrimination. We generally recognizes three types of price-discrimination: first, second and third.

First degree price discrimination (or perfect price discrimination) occurs when the firm can observe the preferences of each individual agent, and is able to extract the entire consumer surplus from each agent, by charging the appropriate fee. To successfully first-degree price-discriminate, the firm must have significant information about each potential consumer. In most cases, this is simply too informationally demanding. (But some examples remain: think of situations where consumers are forced to turn over information about their willingness/ability to pay in order to access a good. E.g. Colleges can give different amounts of financial aid to different students, and this is feasible because the college can require students to furnish significant information about their finances. Similarly, insurers can demand that potential customers furnish their medical and family histories, and charge different premiums accordingly.)

Second degree and third degree price discrimination take these informational constraints more seriously. Under second degree price discrimination, the seller doesn't know the willingness to pay of the buyer, but is able to structure contracts in a way that causes potential buyers to truthfully reveal their willingness to pay. We will study this further in the game theory section.

Third degree price discrimination occurs when the seller is able to put potential consumers into different categories (associated with willingness to pay) on the basis of observable characteristics. E.g. if men and women are believed to have different willingnesses to pay (on average), the seller may sell to men and women at different prices, gender being a (mostly) easily verifiable on inspection. Similarly, sellers may discriminate on the basis of age, student or disability status, citizenship or nationality, location, or time/day etc.

For simplicity, suppose agents can be classified into two clearly identifiable groups, $i \in \{1, 2\}$. Let $p_i(q)$ denote the (inverse) demand function for the good in group i. The firm is free to sell its good at different prices to consumers according to their group membership. The firm's problem is:

$$\max_{q_1, q_2} p_1(q_1)q_1 + p_2(q_2)q_2 - C(q_1 + q_2)$$

The first order conditions are:

$$\begin{aligned} p_1(q_1) + q_1 p_1'(q_1) &= c'(q_1 + q_2) \\ p_2(q_2) + q_2 p_1'(q_2) &= c'(q_1 + q_2) \end{aligned}$$

Notice that each condition has the form of the monopolist's standard profit maximization condition. To maximize overall profits, the firm should maximize profits in each sub-market.

Let us re-write the FOC in elasticity form. We have:

$$\begin{aligned} p_1 \left[1 + \frac{1}{\varepsilon_1}\right] &= c' \\ p_2 \left[1 + \frac{1}{\varepsilon_2}\right] &= c' \end{aligned}$$

Using the same intuition as above, the price in each sub-market should be a mark-up over marginal cost. It should be immediately apparent that the price will be higher in the sub-market that has less elastic demand. Less price-sensitive groups are charged the high price, whilst more price-sensitive groups receive a discount.

Price discrimination of this sort necessarily raises the monopolist's profit. Why? The firm is free to charge the same price in each sub-market. Since it doesn't, and it is profit-maximizing, it must be that its profits from price-discriminating are strictly larger than under a uniform price.

What about consumer welfare? Consumers in the high willingness to pay group are obviously made worse-off —they purchase fewer units at a higher price, and thus earn a smaller consumer surplus. By contrast, consumers in the low willingness to pay group are made better off —they purchase more units at a lower price, and thus earn a larger consumer surplus. The overall effect is ambiguous. Importantly, price discrimination often causes total quantity to increase, so market wide gains from trade will be larger. However, the firm is also appropriating a larger portion of this surplus.

Chapter 14

Normal Form Games

Game Theory is the study of decision making in an environment where outcomes are determined collectively by the choices of many agents. Contrary to environments we have studied so far, a decision maker cannot choose in isolation — she must be cognizant of the choices of other agents in making her own choice. We say such an environment is *strategic*.

Definition 30. A *normal form game* consists of:

- A set of players, $i \in I$
- A set of actions A_i for each $i \in I$
- Preferences over the set of action profiles, for each $i \in I$

Example 47. There are two players, $i \in \{1, 2\}$. Player 1 must choose an action from $A_1 = \{T, M, B\}$, and player 2 must choose an action from $A_2 = \{L, C, R\}$. Agents' preferences are shown in the following matrix, with agent 1's payoff indicated first in each pair:

	L	C	R
T	$9, 5$	$8, 6$	$1, 7$
M	$1, 3$	$2, 5$	$3, 6$
B	$2, 7$	$3, 6$	$2, 8$

14.1 Pure Strategies

Definition 31. A *pure strategy* for player i specifies an action $a_i \in A_i$ for player $i \in I$. A *strategy profile* $a = (a_1, .., a_n)$ a collection of strategies.

We often write $a = (a_i, a_{-i})$ to indicate the action chosen by player i, and the profile of actions chosen by all players other than i, a_{-i}.

Definition 32. A strategy for player i is a *dominant strategy*, if the player's payoff from choosing that strategy is at least as high as from choosing any other strategy, *no matter the strategies chosen by the other players*. I.e. $a_i \in A_i$ is a dominant strategy for player i if

$$u_i(a_i, a_{-i}) \geq u_i(a_i', a_{-i}) \ \forall a_i' \in A_i \text{ and } \forall a_{-i} \in A_{-i}$$

If player i has a dominant strategy, it is she should play it. She doesn't have to anticipate how the other players will play, because this strategy is best no matter what the other players do.

Example 48. Consider the example above. Player 1 does not have a dominant strategy. T is best for player 1 if player 2 chooses either L or C. But if player 2 chooses R, player 1 would rather chooses M. By contrast, Player 2 does have a dominant strategy; R is best for any choice of action by player 1. We should expect player 2 to choose R. Anticipating this, we should expect player 1 to choose M.

Notice that predicting player 1's strategy is harder than predicting player 2's. Player 2 doesn't need to worry about what player 1 is doing. By contrast, player 1's optimal strategy depends on her reasoning about what player 2 will do. Player 1's optimal strategy is more cognitively demanding than player 2's.

Definition 33. A strategy for player i is *strictly dominated* if there is some other strategy that achieves a strictly higher payoff, no matter what the other playhers do.

Players should not choose actions that are strictly dominated.

Example 49. Consider the following modified game:

	L	C	R
T	$9,5$	$8,6$	$1,7$
M	$1,3$	$2,5$	$0,2$
B	$2,7$	$1,6$	$4,8$

Neither player has a dominant strategy. (T is best for player 1 if player 2 chooses L or C, but not R. R is best for player 2 if player 1 chooses T or B, but not M.) Note, however, that M is strictly dominated by T for player 1. No matter what player 2 does, player 1 would do better to choose T than choose M. We can reason that player 1 will never choose M.

Is this sufficient to solve the game? Yes. Eliminating M as a possible strategy for player 1, R now becomes dominant for player 2 (amongst the remaining strategies). So player 2 should choose R. Then player 1 should choose B, anticipating that player 2's choice.

Notice, again, that finding reasoning to the equilibrium becomes more difficult. Player 2 needed to understand that player 1 would not choose M, in order to settle on playing R. Player 1 additionally needed to understand that player 2 would choose R (understanding that player 2 would reason that player 1 would not choose M because it is strictly dominated).

Can we solve all games by selecting dominant strategies and eliminating strictly dominated strategies?

Example 50. Consider the following, different modified game:

	L	C	R
T	$9,5$	$3,6$	$1,7$
M	$1,5$	$4,3$	$2,2$
B	$2,7$	$3,6$	$3,8$

Neither player has a dominant strategy. Moreover, no strategies are themselves strictly dominated. Our elimination procedures will not suffice. To find equilibrium strategies, each must anticipate the actions chosen by the others, and choose the best action accordingly. Moreover, this anticipation of others' strategies proceeds on the understanding that the other players themselves are reasoning in the same way (by anticipating and best responding).t work.

Elimination procedures are convenient because they use straightforward notions of sub-to rule out all but one action for each player. If elimination doesn't work, the players' strategic thinking becomes more c

Definition 34. A strategy profile $a = (a_1, ..., a_n)$ is a *Nash Equilibrium* if each player's action is optimal, given the profile of actions chosen by his opponents. I.e. for each player i

$$u_i(a_i, a_{-i}) \geq u_i(a_i', a_{-i}) \ \forall a_i' \in A_i$$

Notice that a Nash Equilibrium strategy is less demanding than a dominant strategy. A dominant strategy is best no matter what actions the opponents choose. A Nash Equilibrium strategy is merely best given what the opponents are rationally expected to do (although it might not be best for other choices that opponents are not anticipated to make). In a Nash Equilibrium, each agent is optimizing given what his opponents are doing, understanding that they themselves are optimizing given what he (and everyone else) is doing.

Example 51. In the previous example, (B, R) is pair of Nash Equilibrium strategies. Given that player 1 chooses, B, player 2 does best by choosing R, and given that player 2 is choosing R, player 1 does best by choosing B. By contrast, (M, C) is not a Nash equilibrium. Given that player 2 is choosing C, player 1 does best in choosing M. But since player 1 is choosing M, player 2 would rather play L instead.

We motivate Nash equilibrium in a variety of ways. First, we can think of the game-theorist recommending a profile of strategies for all the players to take. A strategy profile is a Nash equilibrium, if each player would want to follow the game-theorist's recommendation assuming all other players do so. Second, and equivalently, we can think of each player forming beliefs about the choices that the other players will make. A strategy profile is a Nash equilibrium if each player's beliefs are confirmed by the actions they take. Nash equilibrium is essentially a no regret condition. After observing the actions of the other players, no agent will want to change his action.

This second interpretation leads one to ask: 'how do the agents form beliefs about their opponents' behavior?'. (The same question manifests from the first interpretation, given that, in reality, the game-theorist is rarely present to whisper strategies into players' ears.) Players are assumed to form beliefs *rationally*. For example, in the above example, the players all understand that it cannot be that player 1 will choose M and player 2 will choose C. If player 2 is conjectured to choose C, it is understood that player 1 will want to choose M and so player 2 will want to switch and choose L instead. It is understood that the conjecture that player 2 chooses C is inconsistent. Rationality requires players to only hold conjectures that are not susceptible.

How do we find Nash Equilibria?

Definition 35. The best response function $b_i(a_{-i})$ selects the action(s) that are optimal for player i for every profile of actions chosen by her opponents. We have:

$$b_i(a_{-i}) = \arg\max_{a_i \in A_i} u_i(a_i, a_{-i})$$

The best response function is a menu that specifies which action player i should take for every possible profile of actions chosen by her opponents. A profile of actions is a Nash equilibrium if each player's action is a best response to actions chosen by the other players. Nash equilibrium requires mutual best response.

Example 52. Return to the example above. The best response functions are:

$$b_1(a_2) = \begin{cases} T & \text{if } a_2 = L \\ M & \text{if } a_2 = C \\ B & \text{if } a_2 = R \end{cases}$$

$$b_2(a_1) = \begin{cases} R & \text{if } a_1 = T \\ L & \text{if } a_1 = M \\ R & \text{if } a_1 = B \end{cases}$$

We can represent the best response functions by underlining the relevant entry in the payoff matrix:

	L	C	R
T	$\underline{9},5$	$3,6$	$1,\underline{7}$
M	$1,\underline{5}$	$\underline{4},3$	$2,2$
B	$2,7$	$3,6$	$\underline{3},\underline{8}$

Since Nash equilibrium requires mutual best response, we search for an action profile where each player is playing a best response. By inspection, the only Nash equilibrium is (B, R).

14.2 Canonical Games

The details of a game (i.e. what strategies are available, and what the payoffs are) will vary from context to context. However, it turns out that, in most contexts, the underling strategic incentives are analogous to those the appear in 5 canonical games. The canonical games themselves are stark, but in studying them, we learn about the nature of strategic interaction in more nuanced scenarios.

Before presenting the games, it is worth noting that the exact numbers specified in the payoff matrices are irrelevant except insofar as they determine which actions are more or less preferred. As usual with utilities, the numerical value is not informative about strength of preference —just the ranking of outcomes. Thus, two different payoff matrices that imply the same ranking over outcomes for each player capture the same strategic scenario.

14.2.1 The Prisoners' Dilemma

Two individuals are caught by the police in the commission of a minor crime. The police suspect them of also having committed a more serious crime, but do not have any firm evidence. The police offer to each agent the following deal: confess to the serious crime and turn state's witness against your friend; if you do, the prosecutor will recommend a more lenient sentence. If neither agent confesses, they will each get 1 year in jail for the minor crime. If one confesses and the other doesn't, the confessing agent is released without penalty (as reward for testifying against his friend), whilst the other gets 9 years in jail. If both confess, then neither testimony is necessary to convinct the other. But having admitted the crime, each gets a reduced sentence of 6 years in jail. The payoffs are:

	$Silent$	$Confess$
$Silent$	$-1,-1$	$-9,0$
$Confess$	$0,-9$	$-6,-6$

It should be clear that confessing is a dominant strategy for each player. (If your opponent confesses, you definitely want to confess and only get 6 years in jail, than not and get 9 years.

If your opponent stays silent, you'd rather confess and get off for free than stay silent and get a year in jail.) The Nash equilibrium is (C, C), and both agents go to jail for 6 years. Notice that, had both stayed silent, they would have only received one year in jail each. The equilibrium outcome is Pareto inefficient.

The Prisoners' Dilemma is famous because it showcases the (slightly counter-intuitive) result that, in the presence of strategic interaction, individually optimal decision making does not necessarily result in socially optimal outcomes. The first welfare theorem fails! There is scope for the government or regulation to direct the decision making of agents so as to make all agents better off.

The Prisoners' Dilemma dynamic is present in a number of strategic scenarios including: oligopolistic competition, public goods provision, (over)-utilization of public resources ('the tragedy of the commons'), arms races (in the military, as well as advertising/lobbying), team-work and free-riding, amongst many others. We will study several of these situations in future lectures.

14.2.2 Coordination Games

There are three versions of the coordination game; all similar, but each having a slightly different flavor. Let us consider each in turn:

Battle of the Sexes: There are two agents who must independently choose which leisure activity to enjoy one evening. Player 1 prefers going to the opera over a sporting event, whilst player 2 has the opposite preference. However, both players would rather enjoy an event together than separately. We can represent this strategic situation as:

	O	S
O	$2, 1$	$0, 0$
S	$0, 0$	$1, 2$

There are two Nash equilibria in the Battle of the Sexes. It is a Nash equilibrium for both agents to choose Opera and similarly an equilibrium for both agents to choose Sport. The agents wish to be coordinated with one another. The Battle of the Sexes captures the strategic dynamic in a variety of settings where agents wish to ensure that they are 'part of the crowd'. Examples include: adoption of a technology (e.g. QWERTY keyboard), driving on the right-side of the road, bank runs and revolutions (I want to leave my money in the bank if enough others do, but withdraw my money if others are withdrawing their money as well).

Chicken: There are two agents who are engaged in a test of mettle. Each gets into a car and drives towards the other at high speed. The first to swerve is revealed to be a 'chicken' and forever suffers the indignity of having lost his nerve. The player enjoys perpetual bragging rights. If both players do not swerve, then both die in head-on high speed collision. If both players swerve at the same time, then the whole exercise proves to be a dud. We represent this situation as:

	Sweve	*Not*
Swerve	$0, 0$	$-1, 1$
Not	$1, -1$	$-2, -2$

There are two Nash equilibria in the game of Chicken; $(Swerve, Not)$ and $(Not, Swerve)$. The agents want to be coordinated, however, unlike the Battle of the Sexes, the nature of the coordination is asymmetric. Each agent wants to swerve if and only if his opponent does not. Chicken captures the strategic dynamic in settings where agents want to be anti-coordinated. Examples include, public goods provision, reporting a crime (I want to report only if noone else does), voting, investing in R& D, etc.

Stag Hunt: There are two agents who are going hunting. Each can decide whether to hunt for stag or hare. Killing the stage requires the two agents to work together and triangulate; if they try alone, they will not succeed. By contrast, each agent can kill a hare by himself. Getting the stag is more valuable than a hare, which is in turn more valuable than not getting anything. We can represent this situation by:

	S	H
S	$3, 3$	$0, 1$
H	$1, 0$	$1, 1$

There are again two Nash equilibria in Stag Hut: (S, S) or (H, H). Like Battle of the Sexes, the equilibria are symmetric. What distinguishes Stag-Hunt from Battle of the Sexes is that the equilibria are Pareto-ranked in the former, but not the latter. In BoS, correctly coordinating on either outcome was socially optimal. Here, the socially efficient choice is for both agents to choose S, and this is an equilibrium, but it is also an equilibrium to coordinate on the socially inefficient choice (H, H). Stag Hunt reminds us that equilibrium need not coordinate us on the socially efficient choice.

All 3 games are coordination games with multiple equilibria. In the presence of multiple equilibria, the concept of Nash Equilibrium provides us with no guidance as to which one will prevail, or which is more likely. 'Refining' equilibria is it's own subfield of game theory, and one we will not turn our attention to. For our purposes, what is important to remember is that the concept of Nash equilibrium doesn't tell us how or why the players form beliefs about which equilibria they will coordinate around. Just that, the beliefs they form are consistent with mutual best response.

14.2.3 Matching Pennies

The final canonical game to consider is (anti)-Matching Pennies. There are two agents who must simultaneously each choose either H or T. Player 1 wins if the players match. Player 2 wins if they do not. The payoff matrix is:

	H	T
H	$1,-1$	$-1,1$
T	$-1,1$	$1,-1$

By inspection, there is no (pure) strategy profile that is a Nash Equilibrium. If the action profile results in player 1 winning, player 2 will wish she had chosen differently. Similarly, if player 2 wins, player 1 will wish she had chosen differently. There is no action profile that is optimal from the perspective of both players. Matching Pennies represents strategic situations that involve where at least one agent has an incentive to deceive the other. Examples include bluffing and bluff-calling (in poker, but also in negotiations more generally).

14.3 Mixed Strategies

Matching pennies is strategically similar to a familiar game —rock, paper, scissors. From playing this game, you should have an intuition that optimal play involves randomizing. Predictable play (i.e. choosing pure strategies) will very quickly result in a string of losses. When players randomize, rather than choosing an action to play, they determine to choose each action with some probability; their strategy is a lottery over the available actions.

Let A_i be the set of actions available to agent i. A *mixed strategy* is an assignment of probabilities α_i, such that $\alpha_i(a_i)$ is the probability that agent i plays actions $a_i \in A_i$. Naturally, we need $\alpha_i(a_i) \geq 0$ for each $a_i \in A_i$, and $\sum_{a_i \in A_i} \alpha_i(a_i) = 1$.

When players choose mixed strategies, they induce lotteries over final outcomes. Thus, the agents assess strategies according to the associated expected utility. Let $a = (a_1, ..., a_n)$ be an action profile, and let $\alpha = (\alpha_1, ..., \alpha_n)$ be a profile of mixed strategies. The expected utility for agent i is:

$$U_i(\alpha) = \sum_{(a_1,...,a_n)} \left(\prod_{j=1}^{n} \alpha_j(a_j) \right) u_i(a_1, ..., a_n)$$

Definition 36. A profile of mixed strategies $\alpha = (\alpha_1, ..., \alpha_n)$ is a (mixed strategy) Nash equilibrium if, for each i, $U_i(\alpha_i, \alpha_{-i}) \geq U_i(\alpha_i', \alpha_{-i})$ for every $\alpha_i \in \Delta(A_i)$.

A strategy profile is a (mixed strategy) Nash equilibrium if no agent can do better by choosing a different mixed strategy, given the mixed strategies chosen by all other players. As usual, if a mixed strategy profile is a Nash equilibrium, then the agents' strategies are mutual best responses.

Example 53. Consider the following variant of Matching Pennies, where the action labels and numerical payoffs have been changed, but the strategic incentives are unchanged.

	L	R
U	$1,-1$	$-1,2$
D	$-1,1$	$1,-1$

Suppose (p^*, q^*) is an equilibrium in which agent 1 chooses U with probability p and agent 2 chooses L with probability Q. The agents' expected utilities are:

$$\begin{aligned} E\left[u_1|p,q\right] &= pq - p\left(1-q\right) - \left(1-p\right)q + \left(1-p\right)\left(1-q\right) \\ &= 2p\left(2q-1\right) - 2q + 1 \end{aligned}$$

$$\begin{aligned} E\left[u_2|p,q\right] &= -pq + 2p\left(1-q\right) + \left(1-p\right)q - \left(1-p\right)\left(1-q\right) \\ &= q\left(2-5p\right) + 3p - 1 \end{aligned}$$

Clearly, $E\left[u_1\right]$ is increasing in p when $2q - 1 > 0$ (i.e. $q > \frac{1}{2}$), it is decreasing in p when $q < \frac{1}{2}$, and it is constant in p when $q = \frac{1}{2}$. Similarly, $E\left[u_2\right]$ is increasing in q for $p < \frac{2}{5}$, it is decreasing in q for $p > \frac{2}{5}$ and constant in q for $p = \frac{2}{5}$. The best response correspondences are:

$$p^*\left(q\right) = \begin{cases} 1 & q > \frac{1}{2} \\ [0,1] & q = \frac{1}{2} \\ 0 & q < \frac{1}{2} \end{cases}$$

and

$$q^*\left(p\right) = \begin{cases} 0 & p > \frac{2}{5} \\ [0,1] & p = \frac{2}{5} \\ 1 & p < \frac{2}{5} \end{cases}$$

By inspection, there is a unique pair of strategies $(p^*, q^*) = \left(\frac{2}{5}, \frac{1}{2}\right)$ which are mutual best responses. (To verify this, note that there cannot be an equilibrium with $p > \frac{2}{5}$. If so, then $q^* = 0$ and if so $p^* = 0 < \frac{2}{5}$. We similarly rule out equilibria with $p < \frac{2}{5}$, $q > \frac{1}{2}$ and $q < \frac{1}{2}$.) This constitutes the unique mixed strategy Nash Equilibrium.

The agents' payoffs are asymmetric and this causes the equilibrium mixing probabilities to be asymmetric. (You can easily verify that with the original specification of payoffs, $p^* = q^* = \frac{1}{2}$.) Notice that the asymmetry breaks towards player 2 benefiting more from choosing R than L. Nevertheless, player 2 mixes evenly between L and R. By contrast,. the payoffs to player 1 are purely symmetric, but player 1 chooses U with lower probability than D. Why? As the best response correspondences demonstrate, an agent only mixes if she is indifferent between her actions; if so, she has no direct interest in the rate at which she mixes. Any mixing probability will do. However, the mixing probability that she chooses will affect her opponent. Hence — each agent chooses her mixing probability, not for her own benefit, but to affect her opponent. In particular, in a totally mixed-strategy equilibrium, each player mixes at the rate which causes her opponent to be indifferent between his actions. (This explains why player 1 puts lower weight on U — since he needs to keep player 2 indifferent between choosing L and R, and player 2 strongly prefers R when player 1 is choosing L.)

Comments

- The best response to a mixed strategy is often a pure strategy!
- The agents are maximizing payoffs in an *ex ante* (rather than *ex post*) sense. For any given realization of outcomes following the randomization, some players may wish that a different outcome had prevailed. Nevertheless, before the outcome is realized, no player would want to change the probabilities with which they choose their actions.

Why focus on mixed strategies? As we argued, there are some games which have no pure strategy equilibria, but may have mixed strategies. In fact, to guarantee the existence of a Nash Equilibrium, it suffices to expand the strategy space to mixed strategies. (Of course, a pure strategy is a special case of a mixed strategy, where the induced lottery over outcomes is degenerate.)

[Nash's Existence Theorem] Every normal form game has an equilibrium (possibly in mixed strategies).

Chapter 15

Oligopoly

In previous chapters, we have studied decision-making under conditions of perfect competition and monopoly. In both cases, each individual firm could make its decision by simply maximizing its profit, taking demand as given. (For monopolists, this was because it was the only firm. In perfect competition, it is because the actions of any given firm do not affect the broader market conditions.) The study of oligopoly differs importantly in that with a few large firms, the decisions of each will impact the market conditions faced by the rest. Thus, each firm must makes its choice, anticipating the decisions made by the other firms.

There are two canonical approaches to modelling oligopoly: the quantity (Cournot) approach, and the price (Bertrand) approach. We begin by studying Cournot Oligopoly, then turn to Bertrand Oligopoly, and then return to see the connection between the two.

15.1 Cournot Oligopoly

There are n firms, $i = 1, ..., n$, that produce a homogeneous good. Firm i can produce a quantity of output q_i at cost $c_i(q_i)$. Let $q = (q_1, ..., q_n)$ denote the profile of quantities produced. We use (q_i, q_{-i}) to denote the quantity produced by firm i, and the vector of quantities produced by the fremaining firms. Additionally, we let $Q = \sum_{i=1}^{n} q_i$ denote the aggregate quantity supplied, and let $Q_{-i} = \sum_{j \neq i} q_j$ denote the quantity excluding firm i.

Let $p(Q)$ denote the (inverse) demand for the good in question. We assume that $p'(Q) < 0$, so that demand is downward sloping. Additionally, as with the monopoly model, we make the assumption that the market wide revenue function $r(Q) = p(Q)Q$ is concave, so that the first order conditions on r find maximum revenue, rather than a minimum. This requires:

$$r'(Q) = p'(Q) + Qp''(Q) < 0$$

Given the total market supply Q, a price $p(Q)$ emerges that causes the market to be in equilibrium; i.e. for demand to equal supply. Note, importantly, that firms do not choose their own prices; instead they contribute to the total market supply and then accept the market price that arises in equilibrium. This may be a more or less realistic assumption depending on the context. For example, in the global market for oil, it would seem reasonable. There is an active and transparent spot market for oil whose price quickly adjusts to changes in market demand and supply.

We consider a Nash equilibrium $q^* = (q_1^*, ..., q_n^*)$ of the Cournot game. In equilibrium, each firm chooses its quantity q_i taking as given the quantites that will be produced by all other firms, and the implications for the market price.

Consider the best response function for firm i. Given a conjecture q_{-i} of other firms' quantities, firm i's problem is:

$$\max_{q_i} p\left(q_i + \sum_{j \neq i} q_j\right) q_i - c_i(q_i) = \max_{q_i} p(q_i + Q_{-i})q_i - c_i(q_i)$$

The first thing to note is that, since the behavior of other firms only affects firm i's profit through the price level, and this depends on total market supply, firm i's best response only depends on the sum of the other players' outputs —it is irrelevant how much each other firm contributed to this total.

The first order condition is:

$$p(Q) + p'(Q)q_i = c_i'(q_i)$$

which is analogous to the monopoly first order condition, except that the market price is not soley determined by firm i. In fact, we can define the residual demand for firm i, as the relationship between market price and firm i's quantity, taking the other firms' quantities as fixed. We have $p_i(q_i; Q_{-i}) = p(q_i + Q_i)$. Then, we can think of firm i as basically profit-maximizing as a monopolist, given her residual demand (taking the other firms' choices as given.) It is easily verified that the second order conditions are satisfied.

A Nash Equilibrium is a vector $q^* = (q_1^*, ..., q_n^*)$ (which implies a market supply Q^*), satisfying:

$$p(Q^*) + p'(Q^*)q_i^* = c_i'(q_i^*) \ \ \forall i$$

Example 54. Suppose $P = 100 - Q$ and let $c_i(q_i) = c_i q_i$, so that each firm faces a constant marginal cost (which possibly differs across firms). For notational convenience, let $\bar{c} = \frac{1}{n}\sum_i c_i$ be the average marginal cost. Firm i's best response function is the solution to:

$$\max_{q_i} \pi_i = (100 - q_i - Q_{-i})q_i - c_i q_i = (100 - Q_{-i} - c_i)q_i - q_i^2$$

By the first order condition:

$$(100 - Q_{-i} - c_i) - 2q_i = 0$$
$$q_i^*(Q_{-i}) = \frac{1}{2}(100 - Q_{-i} - c_i)$$

Notice that $\frac{dq_i^*}{dQ_{-i}} = -\frac{1}{2} \in (-1, 0)$. We will say more about this below. Now, in equilibrium, we have:

$$\begin{aligned} q_i^* &= \frac{1}{2}(100 - Q_{-i}^* - c_i) \\ \frac{1}{2}q_i^* &= \frac{1}{2}(100 - Q^* - c_i) \\ q_i^* &= (100 - Q^* - c_i) \end{aligned}$$

Summing over all firms gives:

$$\begin{aligned} Q^* = \sum_i q_i^* &= (100n - nQ^* - \sum_i c_i) \\ (n+1)Q^* &= n(100 - \bar{c}) \\ Q^* &= \frac{n}{n+1}(100 - \bar{c}) \end{aligned}$$

which gives the overall equilibrium market demand. Then,

$$p^* = 100 - Q^* = \frac{1}{n+1}100 + \frac{n}{n+1}\bar{c} = \frac{1}{n+1}(100 + \sum_i c_i)$$

which is the equilibrium market price. Notice that the equilibrium price is a weighted sum of the highest price that the firm can charge (100) and the average marginal cost ($\bar{c}$), which is lower. Moreover, the weight on the marginal cost term is increasing in the number of firms.

In equilibrium, each firm produces:

$$q_i^* = \frac{100}{n+1} + \frac{n}{n+1}\bar{c} - c_i$$

and earns a profit:

$$\pi_i^* = q_i^*(p^* - c_i) = \left(\frac{100 + n\bar{c}}{n+1} - c_i\right)^2$$

What do we notice. The equilibrium price P^* is decreasing in n, and the equilibrium market supply is increasing. Indeed, as $n \to \infty$, $P^* \to \bar{c}$, so that we return to (average) marginal cost pricing as the number of firms becomes large. We see that the market power to raise prices above marginal cost dissipates as more firms enter the market. By contrast, individual firm output q_i^* is decreasing in n. So as n increases, each firm produces less, but since there are more firms, more is produced in total. $\square$

Return to the first order condition in the general case.

$$p(q_i^*(Q_{-i}) + Q_{-i}) + p'(q_i^*(Q_{-i}) + Q_{-i})q_i^*(Q_{-i}) = c_i'(q_i^*(Q_{-i}))$$

Totally differentiating w.r.t Q_{-i} gives:

$$[p'(Q) + q_i^* p''(Q)]\left(1 + \frac{\partial q_i^*}{\partial Q_{-i}}\right) + [p'(Q) - c'(q_i^*)]\frac{\partial q_i^*}{\partial Q_{-i}} = 0$$

$$\frac{\partial q_i^*}{\partial Q_{-i}} = -\frac{p'(Q) + q_i^* p''(Q)}{(p'(Q) + q_i^* P''(Q)) + (p'(Q) - c'(q_i^*))}$$

(To see that $\frac{\partial q_i^*}{\partial Q_{-i}} \in (-1, 0)$, notice that numerator is negative (by the assumption that the revenue function is concave), and the denominator is negative, by the second order conditions. In fact, the denominator is *more negative*, because it is the sum of two negative terms, one of which is the same as the numerator. Hence, the denominator is larger in size than the numerator.)

What does this mean? Suppose other firms collectively increase output by 1 unit, causing the market price to fall. Then firm i will *reduce* its own output, creating an off-setting effect that pushes the price back up. However, firm i decreases its output by less than 1, so the it only partially offsets the price effect. (In the example, the firm would reduce its output by 0.5 for every unit increase by its opponents. Intuitively, the firm is willing to sacrifice some quantity to prop up the price, but not too much. We can think of this as an instance of 'Le Chatelier's Principle' —the inherent optimizing behavior of the system partially counteracts the imposed changed.

Notice, also, that when the other firm(s) increase their output, the net effect on firm i's profits is negative. Straightforwardly, firm i is seller a lower quantity at a lower price. Thus, when some firms increase their output, they impose an externality on other firms. The total costs of each firm's choice are not entirely experienced by the firm itself. Other firms share some of the burden.

15.1.1 Market Power

As a matter of public policy, we are often interested in understanding the connection between market structure and market power. For example, these considerations go to the heart of anti-trust policy, and forms the basis of regulatory oversight over mergers ad acquisition. The field of Law and Economics was born out of the realization by the legal community for the need for economic analysis in the analysis of such questions.

How should we measure market concentration? A naive answer might be to simply count the number of firms in the market. But, consider a market with one large firm that serves 90% of the market, and 10 small firms each serving %1. Such a market is effectively a monopoly; it stretches credulity to count the large and small firms equally.

A common measure is the Herfindahl-Hirschmann Index (HHI). Let $s_i = \frac{q_i}{Q}$ be the market share of firm i. The HHI is constructed as:

$$HHI = \sum_i s_i^2$$

HHI gives larger importance to big firms relative to small ones. HHI ranges from 0 to 1. In the case of monopoly (most concentrated), $HHI = 1$. A market with n equally sized firms has $HHI = \frac{1}{n}$. (Convince yourself that this is true.) If there are many (roughly) equally sized firms, $HHI \to 0$ as $n \to \infty$. We can think of $HHI = 0$ as the case of perfect competition.

Recall, if there are n equally sized firms, then $HHI = \frac{1}{n}$. Equivalently, $n = \frac{1}{HHI}$. In fact, even if the firms are not equally sized, we can think of inverse-HHI as a measure of the 'effective number of firms' in the market. (The effective number of firms need not be an integer.)

Example 55. Suppose there are four firms with market shares $s_1 = 0.7$, $s_2 = 0.1$, $s_3 = 0.1$ and $s_4 = 0.1$. Then the market concentration is $HHI = 0.7^2 + 0.1^2 + 0.1^2 + 0.1^2 = 0.52$. By contrast, if there were four equally sized firms were equally sized, then $HHI = 0.25^2 + 0.25^2 + 0.25^2 + 0.25^2 = 0.25$. Although both markets have four firms, there is a celar sense in which the first market is more concentrated, in the sense that it is dominated by a single firm. (Of course, the dominant firm still faces some competition, so the market is not as concentrated as a monopoly.) The effective number of firms in the first market is $n^{eff} = \frac{1}{0.52} = 1.92$. So, although there are four firms, the market is as concentrated as a market with two (roughly) equally sized firms. $\square$

The Department of Justice's Horizontal Merger Guidelines (2010) provide the following classification. A market is:

- *Competitive* if $HHI < 0.01$ (i.e. there are more than 100 effective firms).
- *Unconcentrated* if $HHI < 0.15$ (i.e. if there are more than 6.67 effective firms).
- *Moderately concentrated* if $HHI \in (0.15, 0.25)$ (i.e. if there are between 4 and 6.67 effective firms).
- *Highly concentrated* if $HHI > 0.25$ (i.e. if there are fewer than 4 effective firms).

The DOJs categorization of market concentration has implications for its role in overseeing mergers. (See DOJ Horizontal Merger Guidelines 2010.)

- A merger that results in an increase in HHI of less than 0.01 is unlikely to have adverse competitive effects, and typically does not require further analysis.
- Mergers resulting in unconcentrated markets are unlikely to have adverse competitive effects, and typically do not require further analysis.
- Mergers resulting in moderately concentrated markets that involve an HHI increase of more than 0.01 potentially raises significant competitive concerns and often warrant scrutiny.

- Mergers resulting in highly concentrated markets that involve an HHI increase of between 0.01 and 0.02 potentially raises significant competitive concerns and often warrant scrutiny. Those that increase HHI by more than 0.02 are presumed to likely increase market power. (The presumption may be rebutted with further evidence.)

HHI has many advantages as a measure of market concentration, including that it is easily calculated, based on relatively easily measured factors (market shares). But does it capture anything deep about market power? Or is it simply an arbitrary measure?

Return to the Cournot first order condition. We have:

$$\begin{aligned} p(Q^*) + p'(Q)q_i &= c_i'(q_i) \\ p(Q) + p'(Q)Q \cdot \frac{q_i}{Q} &= c_i'(q_i) \\ p(Q)\left[1 + \frac{1}{\varepsilon(Q)} \cdot \frac{q_i}{Q}\right] &= c_i'(q_i) \end{aligned}$$

or

$$L_i = \frac{p(Q^*) - c'(q_i^*)}{p(Q^*)} = \frac{1}{\varepsilon} \cdot \frac{q_i *}{Q^*} = \frac{1}{\varepsilon} \cdot s_i$$

The equilibrium price in the Cournot model follows a modified inverse elasticity formula. Each firm's Lerner Index —i.e. its markup over marginal cost —is the function of two terms: the inverse of the elasticity of demand, and the firm's market share. To make sense of this, recall that if the firm were a monopolist, it would receive a markup exactly equal to the inverse elasticity. When there are many firms, by contrast, its markup is dampened by its market share. For all firms, the markup will be lower than in the monopoly case; however, larger firms will earn a larger markup than smaller firms. This should be intuitive, since all firms receive the same price, the firms with lower marginal costs will have incentives to produce more output, and thus their market share should be higher.

The fact that the Lerner Index is proportional to market shares is noteworthy. It means that when thinking about market power, we should focus more on larger firms than smaller ones. But HHI does just that! Infact, suppose we calculate the average Lerner Index amongst firms in a given market, weighting firms by their market share. Then we have:

$$\overline{L} = \sum_i s_i L_i = \frac{1}{\varepsilon} \sum_i s_i^2 = \frac{1}{\varepsilon} \cdot HHI$$

The average Lerner Index (i.e. markup over marginal cost) in the market is directly proportional to HHI. This suggests that HHI really is a measure of market concentration that is meaningfully related to market power. We can re-write the above formula to make the relationship even more stark. Letting L^m denote the Lerner Index under monopoly, and recalling that $L^m = \frac{1}{\varepsilon}$, we then have:

$$HHI = \frac{\overline{L}}{L^m}$$

We know that markups and price-making power will be greatest under conditions of monopoly. HHI provides the exact link between market concentration and the average firm's market power, relative to a hypothetical monopolist. If $HHI = 0.8$, then the average firm in the market can markup prices to 80% of the level a monopolist would. As $HHI \to 0$, market power disappears, and firms price at marginal cost.

15.1.2 Collusion

A natural question is to compare outcomes under Cournot Oligopoly to those that would arise if the firms colluded, and choose quantities to maximize their joint profit. Intuitively, it must be that the firms can do at least as well by colluding, since at worst they can 'collude' by simply agreeing to produce what they would have otherwise produced in the Cournot equilibrium. (And, there is a strongly likelihood that they can each earn even larger profits, if collusion is successful.) We seek to understand why.

Suppose the firm's formed a cartel, and jointly made output choices for the group. The cartel will naturally want to maximize joint profits. (How it distributes these profits between the firms depends on their bargaining strengths, but evidently, each firm will get at least as much profit as they would have under the non-collusive arrangement.) The cartel's problem is:

$$\max_{q_1,\ldots,q_n} p\left(\sum_i q_i\right)\sum_i q_i - \sum_i c_i(q_i)$$

Taking first order condition w.r.t. q_i (for any firm i) gives:

$$p(Q) + p'(Q)Q = c_i'(q_i)$$

Let us interpret this formula. The left hand side is the market wide marginal revenue; it is how much total revenue amongst all firms changes if output by any firm increases by 1 unit. The right hand side is the marginal cost to firm i. The cartel allocates production to each firm so that, for each firm, the marginal revenue of its last unit of output to the entire cartel is equal to the marginal cost of production in *that* firm. Thus each firm's contribution to joint profits is optimal.

Contrast this to what happens in the non-collusive Cournot equilibrium. Each firm chooses the quantity for which its (private) marginal revenue and marginal cost coincide. But, we know that when one firm increases its output, it reduces the revenue of all other firms. In the non-collusive game, firms does not pay attention to cost they impose on other firms, whereas the cartel does. We can see this more writing the cartel first order condition in the following way:

$$\underbrace{p(Q) + p'(Q)q_i}_{\text{MR to firm } i} + \underbrace{p'(Q)Q_{-i}}_{\text{MR to other firms}} = \underbrace{c_i'(q_i)}_{\text{MC}}$$

The second term is negative, since $P'(Q) < 0$. In the non-collusive equilibrium, each firm chooses q to set the first term equal to the third term, ignoring the second term. Then, it must be from the joint perspective, the LHS is below the RHS. Each firm over-produces relative to the collusive optimum. The cartel will direct each firm to produce less output, which causes market supply to fall and the market price to rise.

But there is a problem. Suppose $q_1^c, ..., q_n^c$ are production plans that maximize joint profits —that solve the cartel's first order conditions. We have just shown that they do not solve any of the firm's individual first order conditions, taking as given the other firms' choices. The plan is not consistent with any firm's best response function. Given the production choices of other firms, each firm would want to raise its quantity above the cartel level. (The reason is exactly the opposite of why the cartel would want each firm to lower its output below the Cournot equilibrium level.) Each firm has an incentive to reneg, and produce more than the cartel instructs it to. [In a later chapter, we study how the cartel may enforce commitment.]

To recap. Although each firm acts in its best interest in the Cournot equilibrium (taking as given the other firm'), all firms earn lower profits that they would if they formed a cartel. The firms settle on a Pareto inferior outcome. Furthermore, even if the firms try to implement the Pareto optimal policy, they will not succeed, because each firm will be tempted to defect, by producing more than agreed. What canonical game form does this remind you of?

Example 56. Consider the above example with linear demand. Set $n = 2$, and suppose the firms have identical marginal costs $c_i = 25$. Then, the cournot equilibrium is $q_1^* = q_2^* = 25$, market supply is $Q^* = 50$ and the market price is $p^* = 50$. Each firm earns a profit of $\pi^* = 625$.

Now, let us solve for the cartel policy. The cartel maximizes:

$$\max_{q_1, q-2} [(100 - q_1 - q_2)q_1 - 25q_1] + [(100 - q_1 - q_2)q_2 - 25q_2]$$

Evidently, this is equivalent to maximizing:

$$\max_{Q} (100 - Q)Q - 25Q$$

Straightforwardly, the maximum occurs at $Q^c = 37.5$, and so each firm should produce $q_1^c = q_2^c = 18.75$. Immediately, we see that each firm over-produces relative to the cartel optimum. The cartel price is $p^c = 62.5$. The total cartel profits are 1406.25, and so each firm's fair share is 703.125.

Let us focus on these two focal quantities: $q^* = 18.75$ and $q^c = 25$. Suppose each firm must decide whether to produce the Cournot quantity or the collusive quantity. We have the following normal form game:

	18.75	25
18.75	$703.13, 703.13$	$525.94, 781.25$
25	$781.25, 525.94$	$625, 625$

It should be clear that choosing $q = 25$ is a dominant strategy, although it results in a Pareto inferior equilibrium allocation. This is precisely a Prisoners' Dilemma.

15.2 Bertrand Oligopoly

The Cournot model has a lot of interesting implications. But the lack of true price competition would seem to be limit its applicability. In this section, we consider an approach that directly models price competition. We begin with the stark case where there are no capacity constraints, and then relax this assumption.

15.2.1 No Capacity Constraints

There are n identical firms, each with an identical constant marginal cost technology $c(q) = cq$. Firms have unlimited capacity; in principle they could serve the entire market. The market demand function is $q(p)$. Each firm announces a price p_i. All consumers purchase from the firm with the lowest price. (Implicitly, quantity demanded is the quantity associated with the lowest price.) If there are multiple firms each having the lowest price, they equally split the market.

A crucial feature of the Bertrand model is its 'winner-takes-all' feature. The firms have strong incentives to slightly under-cut one another, because having just a slightly lower price might be enough to win the entire market. With this in mind, what do best response functions look like? Take a firm i.

$$p_i^*(p_{-i}) = \begin{cases} \min_{j\neq i} p_j - \epsilon & \text{if } \min_{j\neq i} p_j > c \\ c & \text{if } \min_{j\neq i} p_j \leq c \end{cases}$$

for some small ϵ. The intuition is obvious. If the lowest competing price is above the firm's marginal cost, then the firm can slightly under-cut, steal the entire market, and make a positive profit. (If the firm set that same minimum price, it would get a slightly higher price, but half or fewer customers.) By contrast, if the lowest price is at or below c, the best the firm can do is offer the lowest price consistent with not making a loss.

It should be evident that as long as the price is above marginal cost, there will be a strong incentive for firms to undercut one another, pushing the price further towards marginal cost. In a Nash equilibrium, it must be that all firms offer $p = c$.

It is worth pausing to consider just how stark this result is. In the Cournot model, as n increased, each firm's profit margin was squeezed, but it would still price above marginal cost. Marginal cost pricing only arose in the limit as the number of firms became infinitely (read very very) large. By contrast, in the Bertrand model, all it takes is 2 firms to compete with one another, and they will push the price all the way to MC. A little competition is sufficient to compete away all market power. Moreover, in equilibrium, both firms earn zero profits. Does this seem plausible?

What if the firms had constant but heterogeneous marginal costs. Let firm i have cost function $c_i(q_i) = c_i q_i$. How do things change? The best response functions are more-or-less as above (replacing the uniform c, with c_i). Each firm undercuts the lowest price until it reaches their marginal cost. But now, the implication is slightly different. Now, if there is a single firm whose marginal cost is lower than all the others, then the price will be competed down to (slightly under) the 2nd lowest marginal cost. The lowest-cost firm will win the entire market, and sell at a price above its marginal cost, thus earning a positive profit. The insight here is that there is a benefit to being a cost leader. [Of course, if the lowest cost is shared by multiple firms, they will compete the price all the way down to this cost, thereby competing away all profits.] Now, firms cannot freely choose their production technologies, but the logic of the cost-leader earning positive profits suggests a strong motive for firms to engage in R&D to reduce their cost of production below their competitors.

15.2.2 Capacity Constraints

An unrealistic feature of the baseline Bertrand model was that each individual firm had the capacity to serve the entire market. This provided each firm with out-sized incentives to undercut its opponents. In reality, firms are limited in how much they can produce (e.g. because of various fixed resources). If so, the benefits to undercutting may be quite muted, since the number of customers that are available to steal is small. Similarly, the cost to setting a high price may not be so high, because not all customers can go to the cheaper alternative.

Suppose firm i has capacity k_i, which is to say it can produce any quantity up to k_i at a constant marginal cost c. [For simplicity, let us return to uniform marginal cost.] Importantly, we assume that $K = \sum_i k_i < q(c)$. The total capacity of all firms is insufficient to meet demand if prices are at marginal cost. [The game theory becomes quite complicated if this condition is not met.]

Capacity constraints dampen the benefit of undercutting ones opponent. Let $p_K = q^{-1}(K)$. At price p_K, the demand for the good will be exactly equal to the maximum quantity the firms can jointly supply to the market. At this price, every firm is already producing at capacity. No firm has any incentive to lower its price further, since doing so does not steal away any additional consumer. [To be clear; *all* the consumers will want to purchase from that firm, but the firm will not be able to serve them all.]

Will the firms want to price above p_K? (If so, the firm will not be producing at capacity.) No! standard Bertrand logic still dictates that, then, cutting prices slightly will enable firms to boost their sales up to capacity, at little cost to margins. Thus, the equilibrium of the Bertrand game with capacity constraints is for firms to all coordinate on the price at which the market can just meet the demand. Under our assumption on K, this implies an equilibrium price above marginal cost.

We are now closer to a useful model. There is severe price competition between firms. This creates strong incentives for firms to undercut one another. But the incentive exists only when firms are producing below capacity. Once capacity is reached, there is no further reason to lower prices. When overall capacity is low, this can result in prices above marginal cost, consistent with strong price competition.

So far, we have taken the firm's production capacity as given. But, of course, firms choose their capacities, through their investments in resources (especially capital). And firms understand that their choice of capacity, along with the capacity choice of all other firms, determines the market equilibrium price. What capacity should the firm choose? Each firm chooses the capacity that maximizes its profits, taking as given the capacity choices of its competitors.

Let $k = (k_1, ..., k_n)$ be a profile of capacities, and let $K_{-i} = \sum_{j \neq i} k_j$ be the capacity of all firms except i. Let $p(q)$ denote the (inverse) demand function. Firm i's problem is:

$$\max_{k_i} \left[p(k_i + K_{-i}) - c\right] k_i$$

But this is a model of quantity selection, where the price arises endogenously as a function of the total quantity selected by all firms. In fact, it is precisely the Cournot model! Thus, there is a strong connection between the two approaches. Although we usually think of Bertrand competition as being more descriptively true, it is far too stark without capacity constraints. And since capacities are quantities that are chosen, the firms ultimately face a quantity choice captured by the Cournot model, that induces an equilibrium in the Bertrand game.

15.2.3 Heterogeneous Goods Bertrand Duopoly

Another strong assumption underlying the baseline Bertrand model is that firms produced homogeneous goods that are perfect substitutes for one another. Suppose, instead that the goods imperfect substitutes. Then, if one firm reduces its price below the other's, it will not necessarily steal all of their customers. This reduces the temptation to fight a price war all the way down to marginal cost.

Let's dig deeper. For simplicity, consider a market with 2 firms A and B, with demand functions:

$$q_A(p_A, p_B) = 12 - 2p_A + p_B$$
$$q_B(p_A, p_B) = 12 + p_A - 2p_B$$

Let us find the best response function for firm A. The firm's problem is:

$$\max_{p_A}(p_A - c)(12 - 2p_A + p_B) = \max_{p_A}\{-2p_A^2 + p_A(12 + p_B + 2c) - (12 + p_b)c\}$$

which is quadratic in p_A. Straightforwardly we have:

$$p_A^*(p_B) = 3 + \frac{1}{4}p_B + \frac{1}{2}c$$

Notice that the best response function is *increasing* in the other firm's price. We say the firms are *strategic complements*. When one firm increases its price, this creates an incentive for the other firm to increase its price as well. The incentive to undercut is muted. It should be clear that this will enable the firms to sustain an equilibrium above marginal cost.

Let us complete the equilibrium characterization. By symmetry:

$$p_B^*(p_A) = 3 + \frac{1}{4}p_A + \frac{1}{2}c$$

Hence:

$$\begin{aligned}
p_A^* &= 3 + \frac{1}{4}\left(3 + \frac{1}{4}p_A^* + \frac{1}{2}c\right) + \frac{1}{2}c \\
p_A^* &= \frac{15}{4} + \frac{1}{16}p_A^* + \frac{5}{8}c \\
p_A^* &= 4 + \frac{2}{3}c
\end{aligned}$$

and likewise $p_B^* = 4 + \frac{2}{3}c > c$.

Chapter 16

Extensive Form Games

16.1 Extensive Form Games

Normal form games had two requirements that we may find limiting: First is the requirement that players make decisions simultaneously. (Note —simultaneity is not to be interpreted literally; the players may decide at different times. What is important is that the players cannot observe their opponents' choices before making their own choice.) Second is the requirement that the strategic interaction be 'one-shot'; there are no repeated interactions. If we wish to incorporate these into our models, we need a model of sequential decision making. Sequentiality changes the strategic environment by allowing agents who move later in the game to condition their actions on the choices of the agents moving earlier. Strategies can take the form: "If my opponent chooses A, I will respond by choosing X; otherwise I will choose Y." Importantly, such 'contingent strategies' allow players to reward or punish their opponents on the basis of the actions they choose.

Example 57. Consider the following game: There are 2 players. Player 1 first chooses whether to put \$0, \$1 or \$3 into a pot. Player 2 then decides whether to match player 1's contribution or to take the money from the pot. If player 2 matches, then the game theorist adds a bonus to the pot: The bonus is \$2.5 if each player contributed \$1 and \$5 if each contributed \$3. At the end of the game, if there is money in the pot, player 1 gets back double what she put in, and player 2 gets whatever is left. [This game is meant to capture, in simple form, the interaction between a venture capitalist (player 1) who makes an initial investment into a project, and an entrepreneur (player 2), whose effort determines the success of the venture.]

We can represent this strategic situation using a game tree:

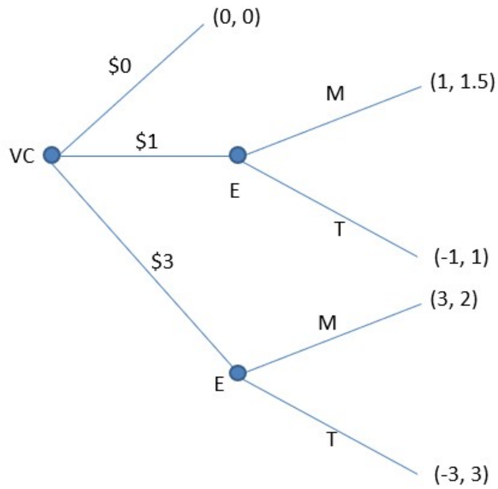

We refer to a path (sequence of actions) from the top of the tree down to any node in the tree as a history. In the above example, there are eight histories: $\emptyset$, \$0, \$1, \$3, $(\$1, M)$, $(\$1, T)$, $(\$3, M)$, and $(\$3, T)$. The histories after which there are no further actions are called terminal histories. Each terminal history is associated with a final outcome and payoffs. The histories after which some player has a choice to make are called *sub-histories*.

Definition 37. A *strategy* for player i specifies the action that player i will take for every sub-history at which he is required to take an action.

Note importantly — a strategy is a menu. It **must** specify the actions chosen after every sub-history, including those that will not arise along the (equilibrium) path of play. For example, in the Venture Capitalist example, a strategy must specify whether player 2 will match or take in both scenarios where player 2 has a choice; i.e. if player 1 choose \$1 or if she chooses \$3. Even if we are certain that player 1 would choose \$3, we must still specify what player 2 *would have done* had player 1 chosen \$1 instead. A strategy might be: 'T if player 1 contributes \$1 and M if player 1 contributes \$3'.

Definition 38. A profile of strategies is a Nash Equilibrium if no agent could do strictly better by changing her strategy, taking as given the strategy of all other players.

Example 58 (Venture Capitalist (cont.))**.** There is a unique Nash Equilibrium of the Venture capitalist game. Player 1 should contribute \$1. Player 2 should choose M if player 1 contributes \$1 and T if player 1 contributes \$3.

Let us check that this is a Nash Equilibrium. First consider player 1. Player 1 never wants to contribute if player 2 will take, and player 1 is better of contributing than not if player 2 matches. Hence, player 1 will not contribute \$3, but will be willing to contribute \$1. Next, consider player 2. Since player 1 contributes \$1, player 2 can either match and get 1.5 or take and get 1. Clearly matching is preferable.

Why did it matter what player 2 did in the sub-history where player 1 contributed \$3. Suppose we alter player 2's strategy so that she always chooses M. (This won't change what she does if player 1 contributes \$1.) If player 1's strategy is to contribute \$1, then player 2's new strategy is still optimal. However, player 1's choice is not optimal given player 2's strategy. Player 1 could now do better by contributing \$3, rather than \$1. And, of course, this cannot be an equilibrium, since if player 1 were contributing \$3, player 2 would want to take. Notice that, player 2's choice 'off-the equilibrium-path' does not affect her payoff, but it does affect the choice that player 1 would make, and therefore is strategically important.

Important —the optimal choice for players choosing earlier in the game depends on choices that other players will make later in the game. When choosing amongst her possible actions, player 1 needs to know what player 2 would do in every contingency, in order to choose the action that is best for her.

Example 59. Consider the following example, regarding the choice of a firm (the 'challenger') to enter an existing monopolistic market, and the choice of the monopolist (the 'incumbent') to retaliate by waging a price war or not.

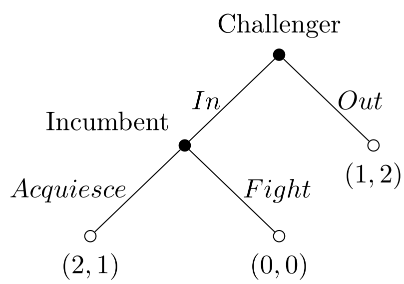

This game has two Nash Equilibria: $(Out, Fight)$ and $(In, Acquiesce)$. The first equilibrium is intuitive. The challenger chooses in. Given this choice, the incumbent's best response is to choose to acquiesce. This vindicates the challenger's choice to go in. By contrast, the second equilibrium is less intuitive. The incumbent threatens that he will fight if the incumbent goes in. This threat deters the challenger from entering. Given that the challenger stays out, the incumbent's decision to fight or not does not affect the outcome, and so there is no beneficial deviation for the incumbent.

Note, the importance of off-equilibrium-threats to sustaining equilibrium play. But are all threats made equal? In the above example, the incumbent may wish to posture and claim that he will fight if the challenger enters. But if the challenger calls his bluff and enters, the incumbent would not wish to follow through with his threat. The threat to fight is not *credible*. This demonstrates a problem with Nash Equilibrium — it permits agents to make wild threats about how they will behave off the equilibrium path, and these threats are believed by the other players. Although such behavior would be sub-optimal, as long as the game never goes down those paths, they player is never called-on to actually make a sub-optimal choice. Nevertheless, we would think that such threats are not realistic. The other players shouldn't believe the threat in the first place. This motivates the notion of *subgame-perfection*.

Definition 39. Let h be a non-terminal history of an extensive form game. The *subgame* following h is the part of the game that remains after history h has occurred.

We can think of a subgame as creating a new game, starting from any non-terminal node in the parent game, and then following the game from there onward.

Definition 40. A profile of strategies of an extensive form game is a *subgame perfect equilibrium* if there is no subgame in which some player can do strictly better by choosing a different strategy, taking as given the strategies of their opponents.

Subgame perfection requires that the agents' strategies be optimal (i.e. Nash equilibria), not only in the parent game, bit also in every subgame. One way to motivate this concept is to suppose that —whilst every player seeks to choose the equilibrium strategy, so that the game should follow the Nash equilibrium path —at every stage, a player may make a mistake (her hand may tremble when choosing her action) and send the game down an off-equilibrium path. Subgame perfection requires that the players continue to play optimal strategies, even after arriving at mistaken histories.

Off-equilibrium path threats in a subgame perfect equilibrium are 'credible' in that, should an agent call the threat-issuing player's 'bluff', the player will optimally carry through with the threat. In the incumbent-challenger example, above, there are two Nash equilibria, but only one —the one in which the challenger enters and the incumbent acquiesces —is subgame perfect. The other Nash equilibrium is not subgame perfect in that, if the challenger called the incumbent's bluff and entered, the incumbent would not want to follow throw with his threat of fighting a price war.

Lemma 15. *A subgame perfect equilibrium is a strategy profile that induces a Nash equilibrium in every subgame.*

Corollary 3. *Every subgame perfect equilibrium is a Nash equilibrium.*

The Corollary reminds us that Subgame Perfection is not a different solution concept to Nash Equilibrium — but rather an refinement of the Nash Equilibrium. In addition to the usual requirement of a Nash Equilibrium that strategies are optimal on the equilibrium path, subgame perfection insists that they are optimal off the equilibrium path as well. Hence every subgame perfect equilibrium will be a Nash equilibrium, but some Nash equilibria may not be subgame perfect.

To find the subgame perfect equilibria, we can use the method of **backward induction**. Start at the bottom of the tree. For each ultimate choice, select the action that is best for the player. Then progressively work back up the tree, assuming that subsequent play follows the path you have found to be optimal. Once you get to the top, you will have found the subgame perfect equilibrium.

Example 60. Return to the venture capitalist example. We indicate the subgame perfect equilibria diagrammatically on the game tree. The optimal actions at each stage are shown by bolding the associated edge.

Start at the bottom of the tree. In the sub-history after player 1 contributed \$1, player 2 does best to meet. So we bolden the branch corresponding to 'meet' (after \$1). In the sub-history after player 1 contributed \$3, player 2 does best to take. So we bolden the branch corresponding to 'take' (after \$3). Now, work up to the previous level. Player 1 must choose from the set $\{\$0, \$1, \$3\}$. He know foresees that player 2 will meet if he chooses \$1 and take if he contributes \$3. Comparing the associated payoffs, his optimal choice is to contribute \$1. We indicate this by boldening the branch corresponding to \$1. We are now at the top of the tree. The subgame perfect equilibrium strategies correspond to the bolded actions.

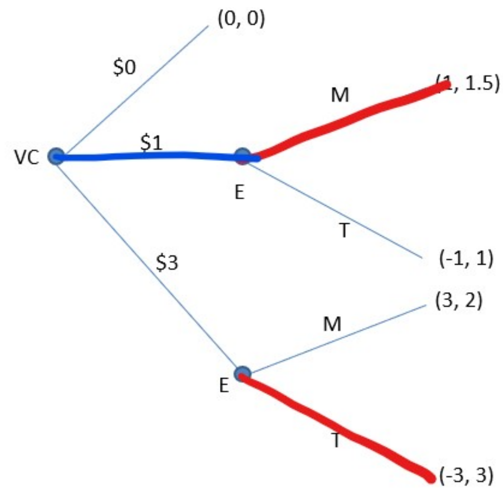

Example 61. Consider the following game (due to Osborne (2003))

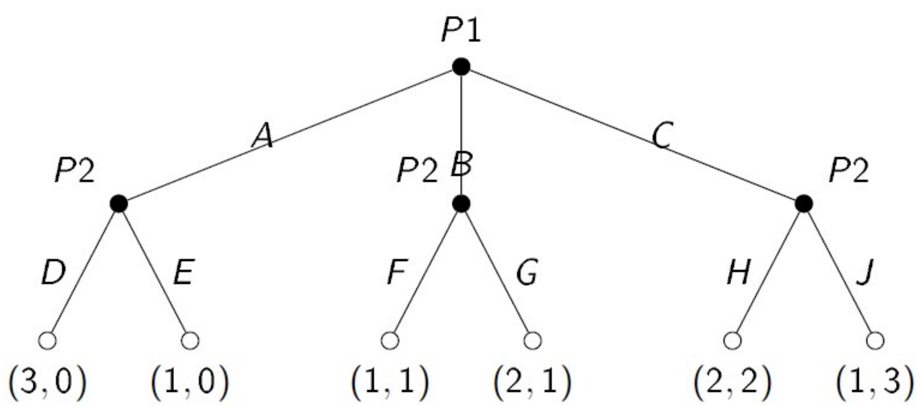

I denote a strategy profile by a pair (W, XYZ) where W is the action chosen by player 1, X is the action chosen by player 2 in history where player 1 chooses A, Y is the action by player 2 following player 1's choice of B, and Z is the action following C. This game has 6 subgame perfect equilibria: (A, DFJ),(A, DGJ), (A, EFJ),(B, EFJ),(C, EFJ), (B, EGJ). Recall, each of these is also a Nash Equilibrium. Additionally, there are 3 Nash equilibria which are not subgame perfect: (A, DFH), (A, DGH), (B, EGH).

16.2 Stackelberg Duopoly

The Cournot model of duopoly involved firms simultaneously choosing quantities, anticipating opponent choices (which affects the market price, and hence, profit). The Stackelberg problem has the same structure as Cournot, except that firms may choose outputs sequentially (or at the very least, one or some firms are known market leaders who are privileged to choose before the others). The sequential nature of the problem has distributional and efficiency implications, relative to the Cournot outcome.

Suppose there are two firms $i \in \{1,2\}$. Firm 1 is the leader and chooses its output q_1 first. Firm 2 then chooses its output q_2, after observing firm 1's output. Each firm has cost function $C_i(q_i)$. The market price is given by the inverse demand function $P(Q)$, where $Q = q_1 + q_2$ is market supply, and $P'(Q) < 0$.

How does the sequentiality change things? Let's solve the game by backward induction. Firm 2 chooses a quantity q_2 to maximize her profit, taking q_1 as given. [Note —unlike the case with simultaneous moves, where firm 2 merely conjectures firm 1's choice, in the sequential game, firm 1's choice is directly observable.] Firm 2's problem is:

$$\max_{q_2} P(q_1+q_2)q_2 - C_2(q_2)$$

The first order condition is:

$$P(Q) + q_2 P'(Q) - C_2'(q_2) = 0$$

Notice that this is exactly the same condition as in the Cournot case. Solving for q_2 gives the the best response function for firm 2, $q_2(q_1)$. So far, nothing has changed.

Now, consider firm 1's problem. Firm 1 chooses its quantity q_1 to maximize its profits, taking as given firm 2's strategy , but not taking a particular quantity as given. Firm 1 knows that firm 2's quantity is contingent on its choice q_1. Thus, rather than taking firm 2's quantity as given, it realizes that firm 2 will change its quantity depending on what firm 1 does, according to the best response function. Firm 1's problem is:

$$\max_{q_1} P(q_1+q_2(q_1))q_1 - C_1(q_1)$$

The first order condition is:

$$P(Q) + q_1 P'(Q)[1 + q_2'(q_1)] - C_1'(q_1) = 0$$

which is different from the usual Cournot first order condition. The condition is the sum of 3 terms. The first two terms together are firm 1's marginal revenue, whilst the third term is marginal cost. Focusing on the first two terms, the first term is revenue gained from selling one more unit (at price $P(Q)$). The second term is negative. It is the revenue lost on all the previous units sold because the price is now lower. In Cournot, when firm 1 increases quantity by 1 unit, total supply increases by 1 unit, and so the price falls by $P'(Q)$. In

Stackelberg, when firm 1 increases quantity by 1 unit, firm 2 best response by partially decreasing q_2, which means total supply increases by less than 1 unit. The overall effect is muted. This means price will fall by less, and so the loss in revenue for firm 1 will be lower. Firm 1 cam rely on firm 2 to mitigate the downside effect of its action.

Recall, $q_2'(q_1) \in (-1, 0)$ —i.e. when firm 1 increases its output, firm 2 will decreases its output, but less than one-for-one. Now, we can re-write firm 1's first order condition as:

$$P(Q) + q_1 P'(Q) - C_1'(q_1) = -q_1 P'(Q) q_2'(q_1)$$

In this form, the LHS is precisely the expression from the standard Cournot first order condition. But in standard Cournot, the RHS was zero. In the Stackelberg game, it is negative! (To see this, note that $P'(Q) < 0$ and $q_2'(q_1) < 0$.) Suppose firm 1 produced the Cournot output q_1^C. Firm 2 will best respond by producing its Cournot output q_2^C. If so, we know the LHS will be zero. But to satisfy the Stackelberg FOCs, the LHS should be negative at the optimum. [Put differently, the actual FOC is positive at the Cournot solution.] Firm 1 has an incentive to raise output. Hence $q_1^S > q_1^C$. Then, by firm 2's best response function, $q_2^S < q_2^C$. Additionally, since firm 2 decreases its output less than one-for-one, its must be that total output rises $Q^S > Q^C$, and so the Stackelberg price is lower than the Cournot price $P^S < P^C$. Since firm 2 produces less and receives a lower price, it must be that firm 2's profit has fallen $\pi_2^S < \pi_2^C$. Moreover, firm 1's Stackelberg profit must be higher than under Cournot, since it could have replicated the Cournot outcome, but chose not to, $\pi_1^S > \pi_1^C$.

The sequentiality of decision making in the Stackelberg setting provides a distinct advantage to the first mover. The first mover does not need to best respond to what the second mover does. Instead, the first mover can exploit the fact of this asymmetry by forcing firm 2 to carry the heavier burden of lowering quantity in order to keep prices from falling too low. In a sense, it is able to free-ride off firm 2.

Example 62 (Stackelberg Duopoly)**.** Suppose the (inverse) demand function is: $P = 100 - Q$ and firms have identical cost functions $c_i(q_i) = 40q_i$. We previously saw that the best response function for each firm in a Cournot duopoly is $q_i(q_{-i}) = 30 - \frac{1}{2}q_{-i}$. In equilibrium, each firm produces $q_i^C = 20$, so that the total output was $Q^C = 40$ and the market price is $P^C = 60$. Each firm earns a profit of $\pi^C = 400$.

Now consider the Stackelberg case. Firm 2's best response function remains $q_2 = 30 - \frac{1}{2}q_1$. Firm 1 chooses q_1 to solve:

$$\max_{q_1}(100 - q_1 - (30 - \frac{1}{2}q_1))q_1 - 40q_1 = \max_{q_1}(30 - \frac{1}{2}q_1)q_1$$

Straightforwardly, this solved by setting $q_1^S = 30$. Firm 2 best responds by setting $q_2^S = 15$. The equilibrium market quantity is $Q^S = 45$ and the equilibrium price is $P^S = 55$. Firm 1 earns profits of $\pi_1^S = 450$, whilst firm 2's profit is $\pi_2^S = 225$.

Notice that firm 1's quantity and profits are larger than in the Cournot duopoly, firm 2's quantity and profits are lower, the overall market supply is higher, whilst the market price is lower.

Before ending this section —a modelling comment. In our sequential game, we assumed that firm 1 would get to move first, and that once it chose its quantity, it wouldn't move again. Firm 1 is *committed* to its chosen quantity. In reality, it's not immediately clear why this should be the case, and why firm 1 wouldn't change its output having seen firm 2's. [Taking firm 2's output as given, firm 1 is not best responding!] Moreover, firm 2's behavior would be different if it anticipated firm 1 changing its mind. My goal here is not to investigate where the first-mover's commitment comes from. Rather, to note that, implicit in the preceding analysis was the (strong) assumption that the first mover had the ability to commit to policies, and that this commitment ability generated strong benefits for player 1.

16.3 Repeated Games

In the above sub-section, we showed that history dependent strategies, equilibrium strategies are affected by (credible) threats about the path of play off the equilibrium path. In this section, we turn our attention to repeated play of the simultaneous-move games from the first sub-section. We are particular interested in games whose Nash Equilibrium is Pareto dominated, such as the prisoner's dilemma; can repeated play with credible punishments cause agents to make efficient choices?

Example 63. We begin with a simple game, in which the following normal form game is repeated twice:

	L	M	R
L	$6,6$	$-1,7$	$-5,-5$
M	$7,-1$	$0,0$	$-5,-5$
R	$-5,-5$	$-5,-5$	$2,2$

The one-shot version of this game has two (pure strategy) Nash equilibria —(M, M) and (R, R). These equilibria are Pareto ranked — (R, R) Pareto dominates (M, M) —although neither is Pareto Optimal, since (L, L) gives strictly larger payoff to both agents.

How many equilibria exist in the two-round version of the game? Naturally, the second round reduces to a one-shot game, and so second period strategies must coordinate on one of these equilibria. Let h be any history after round 1. (There are 9 such histories: $h \in \{LL, LM, LR, ML, MM, MR, RL, RM, RR\}$.) After each of these histories, it is subgame perfect for both players to choose M or to both choose R. What about the first period? Clearly (static-game) Nash play is optimal in the first round, since it satisfies mutual best response. Hence there are 2^{10} SPE in which the agents choose Nash actions in each period. (There are two possible action pairs in the first period. For each of these, we must specify second round play following 9 potential histories. For each history, there are two possible action pairs that constitute a Nash Equilibrium.)

In addition, there are $2^6 = 64$ SPE in which agents choose non-Nash actions in the first stage. These strategies are of the form: choose L in period 1, and in the second period choose R after history (LL) and M for histories (L, M) and (M, L), and choose either R or M after any other history (there are 6 of these). How does the equilibrium work? The reason why (L, L) cannot be sustained as an equilibrium in the one-shot game is that players will be tempted to defect and choose M. This earns an additional payoff of 1. The idea is to reward players who stick to the plan and choose L and punish those who defect to M. The reward is (R, R) in round 2; the punishment is (M, M). The net second period cost of defection is 2 —which should outweigh the 1 unit gained in round 1 from defecting.

Important —take the subgame following any history of play in round 1. The payoffs in this subgame are exactly given by the payoff matrix. However, the payoffs for the overall game (starting from period 1) are not given by the entries in the matrix. These entries are just the payoffs for round 1 —but additional payoffs will follow in round 2. This is why we can sustain actions in period 1 that are not Nash-Equilibrium of the one-shot game; those actions are dominated in the one-period game, but not in the extended game.

16.3.1 Finitely Repeated Prisoners' Dilemma

Consider the following version of the Prisoner's Dilemma:

	C	D
C	$2, 2$	$0, 3$
D	$3, 0$	$1, 1$

Can repeated play generate cooperation in the Prisoners' Dilemma?

Lemma 16. *Suppose the Prisoners' Dilemma is repeated a finite number of times T. In the unique equilibrium, each player chooses D in every period, following every history of play.*

Proof. The proof is by backward induction. Base case: Begin with period T, and let h^T by any history of play up to period T. It is a dominant strategy for each agent to choose D in period T, no matter the history of play. So we know $a^T = (D, D)$. Inductive step: Suppose $a^\tau = (D, D)$ for all $\tau > t$. Now, consider the subgame at period t, following any history of play in the preceding $t - 1$ periods. Each agent knows that his action in period t will not affect his payoff in future periods (since both players will necessarily choose (D, D) in future periods). So each agent simply optimizes in period t as if he is playing a one-shot game. Hence $(a^t) = (D, D)$. $\square$

We can think of this result in two different ways:

1. Since the game ends in period T, the agents cannot sustain cooperation in period T by threatening to punish defection (since there are no further periods in which to meat out punishment). This makes period T play entirely predictable. Now consider period $T-1$. Agents may try to sustain cooperation in period $T-1$ by promising to cooperate in period T if the opponent cooperates (in period $T-1$), but switching to defection if the opponent defects. But such a threat is not credible, since it is known that agents will definitely choose D in period T. Lacking a credible threat to sustain cooperation, both agents will choose D in $T-1$. The same logic applies all the way forward to period 1. The fact that the game ends makes credible punishment impossible.

2. The agents know that they will eventually switch to choosing D towards the end of the game. Nevertheless they may wish to choose C at the start. However, each agent wants to be the one who defects first, and switches from C to D (because if not, she will bear the loss of cooperating whilst her opponent defects). If player 1 expects player 2 to switch at period t, he will want to switch at period $t-1$. In continuing to preempt each other, the players bring the switching time all the way forward to round 1.

The inability to sustain cooperation can be made stark by considering the following variant of payoffs:

	C	D
C	$100, 100$	$-1, 101$
D	$100, -1$	$0, 0$

Along the equilibrium path of play, both agents receive a discounted stream of utility of zero. If they cooperated in every period, they would have each received $100T$ (taking $\delta = 1$ for simplicity). The temptation to earn 101 in some period, rather than 100, by defecting on their opponent, causes both to earn 0!

16.3.2 Infinitely Repeated Prisoners' Dilemma

The lesson from finitely repeated PDs was that cooperation cannot be sustained when opportunities for punishment dry up. Furthermore, there is a contagion effect, where this lack of punishment at the end of the game, prevents punishment at *any* stage of the game. What if the game is repeated infinitely, so that the possibility of future punishments are always available? It should be clear that repetition of the static Nash actions remains an equilibrium strategy. (If I always expect my opponent to choose D, the best I can do is to always choose D.) However do other equilibria exist?

Yes! I demonstrate some of these equilibria by example. Before doing so, we need to know how to find SPE in infinite games. (We know to use backward induction for finite games, but this method cannot work in an infinite horizon.

Proposition 5 (One-shot deviation principle.). *A strategy profile of an infinitely repeated game is a subgame perfect equilibrium if and only if no player can gain by changing her action after any history, given both the strategies of the other players, and the remainder of her own strategy.*

Example 64 (Grim Trigger Strategies). Consider the following strategy for player i: $s_i(\phi) = C$ and for every $t = 1, 2, ...$

$$s_i\left(a^1, .., a^t\right) = \begin{cases} C & \text{if } a^\tau = (C, C) \text{ for every } \tau \le t \\ D & \text{otherwise} \end{cases}$$

This is the grim-trigger strategy. Agents begin by cooperating, but the immediately and forever switch to defect after either agent chooses D. We call this the *grim trigger*, because the punishment is severe. Once a player's trust is lost, it can never be regained.

We can think of the strategy has having one of two states — $\mathcal{C}$ and $\mathcal{D}$. The strategy prescribes that each agent should choose C when the state is $\mathcal{C}$, and choose D when the state is $\mathcal{D}$. The strategy must also prescribe the dynamics of state transitions. We have that the original state is $\mathcal{C}$, that $\mathcal{D}$ is an absorbing state, and that the state transitions from $\mathcal{C}$ to $\mathcal{D}$ if some action pair other than (C, C) was chosen in the previous round.

Is it an equilibrium for both agents to choose such a strategy? We must check if there is a favorable deviation to either agent. Recall, by the one-shot deviation principle, it suffices to check for one period deviations. If the agents follow the strategy, then C is chosen by both agents in every period, and so the payoffs are:

$$\sum_{t=0}^{\infty} \delta^{t-1} u(C, C) = \sum_{t=0}^{\infty} \delta^t (2) = \frac{2}{1-\delta}$$

Suppose some agent deviates in period t and chooses D. Then the path of play will be (C, C) until period $t-1$, (D, C) in period t, and (D, D) for every future period. The agent's lifetime utility is:

$$\sum_{\tau=0}^{t-1} \delta^\tau (2) + \delta^t (3) + \sum_{\tau=t+1}^{\infty} \delta^\tau (1) = 2\frac{1-\delta^t}{1-\delta} + 3\delta^t + \frac{\delta^{t+1}}{1-\delta}$$

Such a deviation is not favorable provided that:

$$\begin{aligned} 2 &\ge 2\left(1-\delta^t\right) + 3(1-\delta)\delta^t + \delta^{t+1} \\ 0 &\ge \delta^t - 2\delta^{t+1} \\ \delta &\ge \frac{1}{2} \end{aligned}$$

Hence, there is an equilibrium in which agents always cooperate, sustained by the threat of eternal punishment, provided that $\delta \ge \frac{1}{2}$. We shown that we can sustain cooperation if agents are sufficiently patient — i.e. if they put sufficient weight on the future.

Why can we sustain cooperation in the infinite game, but not in the finite game? In the finite game, the game eventually ends. When it does, the ability to punish disappears, and this undermines threats to punish in previous periods. By contrast, in the infinite game, there always exist additional periods in the future in which to punish deviations. The threat of punish will enforce cooperation provided that the agents are sufficiently patient (i.e. they care enough about the future).

Are there other types of trigger strategies? Yes. Below are other common strategies:

1. Limited Punishment Strategies: This strategy is similar to the grim trigger, but is less draconian, because the opponent is forgiven after sufficiently many periods of punishment.

   The game begins in state P_0 and in this state, the agent chooses C. The state remains in state P_0 unless the opponent chooses D. If so, the state transitions to punishment state P_1 where the agent chooses D. The state then automatically transitions to state P_2 (and then again to state P_3 and then back to state P_0) regardless of the agents' choices. In punishment states P_2 and P_3, the agent chooses action D. Hence, following a defection by the opponent, the agent automatically defects for 3 periods, and then reverts back to cooperating.

2. Tit-for-tat. In this strategy, the agent begins by choosing C, and then simply copies what his opponent chose in the previous round. The strategy is:

$$s_i(\phi) = C$$

$$s_i\left(a^1, ..., a^t\right) = \begin{cases} C & \text{if } a^t{=}(\cdot, C) \\ D & \text{otherwise} \end{cases}$$

Chapter 17

Public Goods & Externalities

17.1 Public Goods

Most of our analysis to this point has focused on *private goods* —those who consumption only affects a single agent. Now we consider *public goods*. First some terminology.

We say a good is *non-rival* if its consumption by one agent does not reduce its availability for consumption by others. E.g. a film screening is non-rival. The fact that I'm watching the film doesn't mean there is less film for you to watch (up to the point where the cinema fills up!). By contrast, an apple is clearly rival. If I consume the apple, then it is not available for you to consume. We say a good is *non-excludable* if it is impossible to prevent people from consuming it. A film screening in an enclosed theatre is clearly excludable —you can prevent people (non-payers in particular) from entering. A fireworks display, by contrast, is non-excludable. You cannot prevent non-payers from looking up and enjoying the fireworks show!

We say a good is *public* if it is both non-rival and non-excludable. Examples of public goods include: police protection and national defense, lighthouses and streetlights, free-to-air television and radio broadcasts etc. Not all non-private goods are public. Some goods are non-rival but excludable, e.g. cable (or subscription streaming service) television —we refer to such goods as *club goods*. By contrast, other goods are rival but non-excludable, e.g. street parking. It is important to note that there are a range of goods that are commonly referred to as 'public goods', typically because they are provided by governments, such as healthcare and education, but which are quite clearly rival and excludable, and so are technically private goods (even if they are socially very valuable).

17.1.1 Efficient Provision of a Discrete Public Good

Start with a simple example. There are two agents $i \in \{1, 2\}$ and two goods —a private good x and a public good G. Agent i has income y_i, and must determine how to allocate this between the public and private good. Let g_i be the allocation to the public good, and $x_i = y_i - g_i$ be private consumption. Agents have standard preferences over the two goods $u_i(G, x_i)$ where u_i is strictly increasing in both components.

First, suppose the public good is a fixed object that is either provided or not. (Later we will consider choosing amongst different quantities of the public good.) Suppose it costs c to provide the public good, so that:

$$G = \begin{cases} 1 & \text{if } g_1 + g_2 \geq c \\ 0 & \text{if } g_1 + g_2 < c \end{cases}$$

Should the public good be provided? For each consumer i, let r_i denote the maximum she would be willing to pay in order that the public good be provided. We call r the *reservation price*. r_i satisfies:

$$u_i(1, y_i - r_i) = u_i(0, y_i)$$

Now, let (g_1, g_2) be a pair of contributions satisfying $g_1 +_g 2 \geq c$. If it is desirable for both agents to make this contribution, then:

$$u_i(1, y_i - g_i) > u_i(0, y) = u_i(1, y_i - r_i)$$

which implies that $y_i - g_i > y_i - r_i$ for each i (since u is strictly increasing in both components). Hence $g_i < r_i$. Combining this equation for both agents:

$$r_1 + r_2 > g_1 + g_2 \geq c$$

Thus, if the sum of the willingnesses-to-pay amongst members of the public exceeds the cost of provision, then it is efficient to provide the public good (in the sense that there is a way to allocate the burden of paying for the good in a way that makes everyone better off from doing so rather than not). Note the contrast to private goods, for which, efficiency requires provision provided that the individual consumer's willingness to pay is larger than the cost of provision. In the case of public goods, it could be that the willingness to pay of each agent is (well) below the cost of provision. However, since the good is non-rival and so can be enjoyed by many, it may still be efficient to provide as long as the joint willingness to pay exceeds the cost of provision.

17.1.2 Equilibrium Provision of a Discrete Public Good

Suppose $r_1 + r_2 > c$ so that provision is efficient. Will the good be provided by the marketplace? The answer is often no! We consider two possibilities: (1) the good is provided privately by the market; and (2) the good is provided by the government, based on the voting behavior of the public.

First, consider the private market case (where, implicitly, there is no coordination amongst the consumers). For tractability, suppose $r_i = 100$ for each $i \in \{1, 2\}$ and that $c = 150$. Each agent must decide whether to purchase (individually) or not. We can represent the consumers' payoffs in the following matrix:

	Buy	Don't Buy
Buy	$25, 25$	$-50, 100$
Don't Buy	$100, -50$	$0, 0$

Clearly the equilibrium is for neither agent to buy, even though we know that both could be made better-off if each made a moderate contribution. Naturally, the problem here is one of coordination.

Second, consider the voting case. Suppose there are 3 agents with $r_1 = 90$, $r_2 = 30$ and $r_3 = 30$, with $c = 99$. Then $r_1 + r_2 + r_3 = 150 > c$ and so provision is efficient. Suppose in the voting scheme, each individual understands that they will be liable for a $1/3$ share (through taxation). Then the per-person cost is 33, which is above the reservation cost of 2 out of 3 consumers. A majority will note for the good not to be provided.

17.1.3 Efficient Provision of a Continuous Public Good

Now suppose the public good can be provided in any (non-negative) amount. Again, for simplicity, suppose there are two agents, although the results generalize easily. If $g_1 + g_2$ is contributed to the public good, then $G = g_1 + g_2$ is provided, and agent i receives utility $u_i(G, x_i) = u_i(g_1 + g_2, y_i - g_i)$.

We argued in Econ 201 that if an allocation was *Pareto efficient*, it must maximize a weighted sum of the agent utilities. Thus, the efficient provision of public goods must satisfy:

$$\max_{g_1, g_2} \gamma_1 u_1(g_1 + g_2, y_1 - g_1) + \gamma_2 u_2(g_1 + g_2, y_2 - g_2)$$

The first order conditions are:

$$\begin{aligned} \gamma_1 \frac{\partial u_1(G, x_1)}{\partial G} + \gamma_2 \frac{\partial u_2(G, x_2)}{\partial G} &= \gamma_1 \frac{\partial u_1(G, x_1)}{\partial x_1} \\ \gamma_1 \frac{\partial u_1(G, x_1)}{\partial G} + \gamma_2 \frac{\partial u_2(G, x_2)}{\partial G} &= \gamma_2 \frac{\partial u_2(G, x_2)}{\partial x_2} \end{aligned}$$

It must be that $\gamma_1 \frac{\partial u_1(G,x_1)}{\partial x_1} = \gamma_2 \frac{\partial u_2(G,x_2)}{\partial x_2}$. Taking either equation and dividing the LHS by the RHS gives:

$$\begin{aligned} \frac{\frac{\partial u_1(G,x_1)}{\partial G}}{\frac{\partial u_1(G,x_1)}{\partial x_1}} + \frac{\frac{\partial u_2(G,x_2)}{\partial G}}{\frac{\partial u_2(G,x_2)}{\partial x_2}} &= 1 \\ MRS^1_{Gx} + MRS^2_{Gx} &= 1 \end{aligned}$$

Recall, the marginal rate of substitution is the amount of the private good that each agent is willing to give up to get one more unit of the public good. Efficiency requires that the sum total of these willingnesses-to-forgo should coincide with the amount of private consumption that must be forgone (which in our setup is 1). The analogy to discrete case should be obvious.

Example 65. Suppose the agents have identical Cobb Douglas preferences $u_i(G, x_i) = \alpha \ln G + (1-\alpha)\ln x_i$, and let the incomes be y_1 and y_2 (not necessarily the same). The marginal rate of substitution for agent i is:

$$MRS^i = \frac{\frac{\partial u_1(G,x_1)}{\partial G}}{\frac{\partial u_1(G,x_1)}{\partial x_1}} = \frac{\frac{\alpha}{G}}{\frac{1-\alpha}{x_i}} = \frac{\alpha}{1-\alpha} \cdot \frac{y_i - g_i}{g_1 + g_2}$$

Efficiency requires:

$$\begin{aligned} MRS_1 + MRS_2 &= 1 \\ \frac{\alpha}{1-\alpha} \cdot \frac{y_1 - g_1}{g_1 + g_2} + \frac{\alpha}{1-\alpha} \cdot \frac{y_2 - g_2}{g_1 + g_2} &= 1 \\ \frac{\alpha}{1-\alpha} \cdot \frac{y_1 + y_2 - g_1 - g_2}{g_1 + g_2} &= 1 \\ \alpha(y_1 + y_2) - \alpha G &= (1-\alpha)G \\ G^{eff} &= \alpha(y_1 + y_2) \end{aligned}$$

17.1.4 Equilibrium Private Provision of Continuous Public Good

Now that we've found the efficient quantity, what level of public good will actually be provided in a private market? Suppose each agent i voluntarily and independently contributes g_i towards the provision of the public good. Then $G = g_1 + g_2$ units of the public good are provided. Each player makes their contribution to maximize their utility, taking as given their opponent's contribution. Naturally, we solve for a Nash equilibrium.

Player i's problem is:

$$\max_{g_i} u_i(g_i + g_{-i}, y_i - g_i)$$

The first order condition is:

$$\frac{\partial u_i(G, x_i)}{\partial G} = \frac{\partial u_i(G, x_i)}{\partial x_i}$$

Let's interpret the expression. The LHS is the marginal benefit to agent i of increasing g_i by 1 unit, taking as given g_{-i}. The RHS is the marginal (opportunity) cost to agent i, in terms of x_i forgone. Each agent chooses their contribution level g_i to maximize *their* net benefit. Contrast this to the first order condition that defines the efficient output:

$$\frac{\partial u_i(G, x_i)}{\partial G} + {\color{red}\frac{\gamma_{-i}}{\gamma_i} \frac{\frac{\partial u_{-i}(G,x_{-i})}{\partial G}}{\frac{\partial u_i(G,x_i)}{\partial G}}} = \frac{\partial u_i(G, x_i)}{\partial x_i}$$

The term in red is the 'external benefit' that the other agent receives (normalized to the same unit of utility as agent i) when agent i increases her contribution by 1 unit. The social planner naturally takes this into account, but with private provision, the agent herself does not. Naturally, we should expect this to generate an under-provision of the public good.

We see this by re-writing the FOC in terms of marginal rates of substitution. We have $MRS_i = 1$ for each i, which implies $MRS_1 + MRS_2 = 2$. As before, the LHS is the marginal social benefit to increasing G by one unit. Notice that the condition has become more demanding. The public good will be provided only up to the point that its marginal social benefit is twice its marginal cost (whereas efficiency requires its provision as long as $MSB > MC$). The private market will under-provide the public good.

Example 66. Continue the previous example with Cobb-Douglas preferences. Agent i maximizes $u_i(G, x_i) = \alpha \ln(g_i + g_{-i}) + (1-\alpha)\ln(y_i - g_i)$, taking g_{-i} as given. The first order condition implies:

$$\begin{aligned} \frac{\alpha}{g_i + g_{-i}} + \frac{1-\alpha}{y_i - g_i} &= 0 \\ (1-\alpha)(g_i + g_{-i}) &= \alpha(y_i - g_i) \\ g_i^*(g_{-i}; y_i) &= \alpha y_i - (1-\alpha) g_{-i} \end{aligned}$$

This is agent i's best response function. Notice that g_i is decreasing in g_{-i}. There is a free-riding problem! As her opponent increases her contribution, each agent has an incentive to reduce her contribution and free-ride off her opponent. This results in a decreased overall provision of the public good.

What is the equilibrium provision? Combining the best response functions:

$$\begin{aligned} g_1^* &= \alpha y_1 - (1-\alpha) g_2^* \\ g_2^* &= \alpha y_2 - (1-\alpha) g_1^* \end{aligned}$$

which implies that: $g_1^* + g_2^* = \alpha(y_1 + y_2) - (1-\alpha)(g_1^* + g_2^*)$, and so:

$$G^* = g_1^* + g_2^* = \frac{\alpha}{2-\alpha}(y_1 + y_2) < \alpha(y_1 + y_2) = G^{eff}$$

17.1.5 Public Provision of a Continuous Public Good through Voting

Now, consider public provision through voting. To make the problem more interesting, suppose there are n agents, $i = 1, ..., n$. Suppose it has been exogenously determined that agent i will pay for a share $s_i \in [0, 1]$ of whatever quantity of the public good is provided, such that $\sum_i s_i = 1$. (Think about s_i as representing marginal tax rates that are determined independently of the level of public good spending.) The equilibrium level of public provision G_v is the one for which a majority of voters cannot be found who would either increase or decrease it.

First, let us find the optimal level of public goods provision for a given voter. Voter i's ideal G solves: $\max_G u_i(G, y_i - s_i G)$. The first order condition implies:

$$\begin{aligned} \frac{\partial u_i(G, x_i)}{\partial G} &= s_i \frac{\partial u_i(G, x_i)}{\partial x_i} \\ MRS^i &= s_i \end{aligned}$$

The intuition is straight-forward. Each agent wants to provide public goods up to the point where their wililingness-to-pay (MRS^i) is equal to *their share* (s_i) of the cost of provision. Let G^i be the optimal provision of public goods for agent i, and assume (naturally) that agent i prefers allocations less and less, the further is the quantity from G^i.

Let m index the median voter —the one who has the middle value of G^i. Thus, half the voters have an ideal public good provision below G^m and half have an ideal level above G^m. Suppose the government proposes $G < G^m$. Then a majority exists (i.e. the median and everyone who prefers even larger G) that would vote to raise G to G^m. Suppose instead that the government proposes $G > G^m$. Again, a majority (the median and everyone who prefers lower provision) would vote to lower G to G^m. Thus, the voting equilibrium will select $G_v = G^m$, and so G_v is characterized by the condition:

$$MRS^m = s_m$$

Is this efficient? Recall, the efficient quantity is given by: $\sum_i MRS^i = 1$, which we can re-write as:

$$\frac{1}{n}\sum_i MRS^i = \frac{1}{n}$$

Thus, efficiency requires that the average MRS is equal to the average contribution $\frac{1}{n}$. But depending on the distribution of incomes and contributions shares $\{s_i\}$, the G_v that solves the voting equilibrium may either be higher or lower than the efficient level.

Why the ambiguity? With voting there are two sources of inefficiency that push in opposite directions: First, as with the private market case, individual agents ignore the benefit to others of public goods provision, and this will cause public goods to be under-provided. Second, with government provision, any given agent doesn't pay for the cost of the public good, but only a fraction of the cost —forcing the remainder of the costs onto the rest of society. This makes public goods seemingly cheaper, and so will result in an over-provision (all else equal). The two effects taken together may either result in an under- or over-provision, depending on the size of each effect.

Example 67. Continue with the Cobb-Douglas example. Each agent i has preferences $u_i(G, x_i) = \alpha \ln G + (1-\alpha)\ln(y_i - s_i G)$. Analogous to results above, the efficient level of public goods is

$$G^{eff} = \alpha \sum_i y_i = n\alpha\overline{y}$$

, where $\overline{y}$ is the average income. Now, each individual agent's ideal level of public provision is given by the first order condition;

$$\frac{\alpha}{G} - s_i \frac{1-\alpha}{y_i - s_i G} = 0$$

, which implies:

$$G^i = \alpha \frac{y_i}{s_i} = (n\alpha\overline{y}) \cdot \frac{y_i}{\overline{y}} \cdot \frac{1}{ns_i} = G^{eff} \cdot \left(\frac{y_i}{\overline{y}} \cdot \frac{1}{ns_i}\right)$$

Hence $G_v > G^{eff}$ provided that:

$$\frac{y_m}{\overline{y}} > ns_m$$

or equivalently:

$$\frac{y_m}{s_m} > \frac{\overline{y}}{\frac{1}{n}}$$

Recall $G^i = \alpha\frac{y_i}{s_i}$, and so each agent's ideal level of public goods provision is proportional to $\frac{y_i}{s_i}$. All else equal, we know that richer agents will want more of the public good than poorer agents (because of diminishing marginal utility from the private good). Hence, as income increases, so will demand for public goods. Similarly, as s_m increases, the agent must internalize more of the cost of financing the public goods, and this will cause demand for public goods to decrease. For Cobb-Douglas preferences, what matters is the ratio.

Now, the efficient level G^{eff} is proportional to that ratio for the average agent: $\frac{\overline{y}}{\frac{1}{n}}$. Hence, the voting equilibrium level will be higher than the efficient level if the median voter's ratio is larger than this ratio, and vice versa.

17.2 Externalities

An important ingredient to the First Welfare Theorem was that each agent's utility only depended on their allocation, and was unaffected by what was given to others (except through the social resources constraint). But we know, in many instances, that the consumption of some directly affects the utility of others. We refer to such consequences to third parties as *externalities.*

In general, externalities may either be positive (e.g. education, immunisation) or negative (e.g. pollution). In this section we will explore the relevant concepts through the canonical example of a firm whose production process emits pollution.

Example There are two firms. Firm 1 produces output x which it sells in a competitive market at price p. The production of x involves direct costs $c(x)$ to firm 1, and also imposes external costs $e(x)$ on firm 2. Both costs functions ($c(x)$ and $e(x)$) are assumed increasing and convex in x, so that the marginal cost and marginal external cost are positive and increasing in the level of output. The profits of the two firms are:

$$\pi_1 = \max_x px - c(x)$$
$$\pi_2 = -e(x)$$

Clearly, the profit maximizing output is given by $p = c'(x^{max})$ (i.e. standard marginal cost pricing).

To determine the efficient level of output, we ask what output would be produced if the two firms merged so as to internalize the externality. (Under this scenario, there are no longer an externalities, and so the first welfare theorem will apply.) If so, the joint firm would choose output to maximize:

$$\pi = \max_x px - c(x) - e(x)$$

The first order condition is:

$$p = c'(x^{eff}) + e'(x^{eff})$$

Clearly $x^{max} > x^{eff}$ and so, in the presence of a negative externality, there will be an over-provision in the private market. (To see this, note that that since $e' > 0$, $c'(x^{max}) = p = c'(x^{eff}) + e'(x^{eff}) > c'(x^{eff})$. Then, since $c'' > 0$, $c'(x^{max}) > c'(x^{eff})$ implies $x^{max} > x^{eff}$.)

17.2.1 Solutions to the Problem of Externalities

Here we consider two common solutions: (1) Pigouvian taxes and (2) Missing markets.

Pigouvian Taxes. Suppose the government imposes a per-unit tax on firm 1's output. Firm 1's first order condition becomes:

$$p = c'(x) + t$$

and so, setting $t = e'(x^{eff})$ will cause firm 1 to naturally make the socially efficient choice. The tax causes the firm to *internalize the externality.* Of course, to implement such a scheme, the government would need to know the externality function, which might be quite informationally demanding.

Missing Markets. Under the missing markets view, the problem of externalities arises because firm 2 values the pollution created by firm 1, but has no way to influence it. We could solve this problem by explicitly adding a market for pollution in which firm 2 can participate.

Consider a market for pollution. For simplicity, suppose that producing x units of output generates x units of pollution. (As we will see below, it is not at all important that the relationship be one-for-one.) Let r be the price of pollution. For each unit of pollution produced, firm 1 must find a buyer to sell it to; it can't simply dump it in the environment. The natural buyer will be firm 2, since it is the other actor affected by the pollution. Let x_1 be firm 1's production (and thus its supply of pollution). Let x_2 be firm 2's demand for pollution. We have:

$$\pi_1 = \max_{x_1} px_1 + rx_1 - c(x_1)$$
$$\pi_2 = \max_{x_2} -rx_2 - e(x_2)$$

The first order conditions are:

$$p + r = c'(x_1)$$
$$0 = -r - e'(x_2)$$

The first equation defines the supply of pollution (formally, $x^s = [c']^{-1}(p+r)$) and the second equation defines the demand for pollution $x^d = [e']^{-1}(-r)$. Since, in equilibrium, we must have $x_1 = x_2$, then the first order condition simplifies to:

$$p = c'(x) + e'(x)$$

which is the condition for efficiency.

The basic idea is that equilibrium in the market for pollution will causes the pollution to be priced exactly so that it causes the firm to internalize the externality it creates. Notice that $r^* = -e'(x^*) < 0$ —which makes sense; pollution is a bad. The buyer will not want to pay for it, but rather needs to be paid to accept the pollution. The buyer (firm 2), will be willing to buy pollution as long as the payment received ($-r$ per unit) is larger than the cost it faces in 'enjoying' the bad ($e'(x)$).

To relax the one-for-one assumption, modify the problem slightly. Suppose firm 1 can produce x units of output and y units of pollution at cost $c(x, y)$. If there is no market for pollution, then the firm's maximization problem is:

$$\max_{x,y} px - c(x, y)$$

which implies first order conditions:

$$\begin{aligned} p &= \frac{\partial c(x, y)}{\partial x} \\ 0 &= \frac{\partial c(x, y)}{\partial y} \end{aligned}$$

Firm 1 chooses output such that price equals marginal cost, and chooses the level of pollution that minimizes the cost of production (for the desired output level). Since pollution is not priced, the firm will pollute as long as doing so reduces costs.

Now, suppose there is a market for pollution, and that each unit of pollution costs r. For each unit of pollution produced, the firm must sell it on the market —it cannot simply dump it in the environment. Naturally, firm 2 (who experiences the pollution) will be the 'buyer'. Now, the firms' profits are:

$$\begin{aligned} \pi_1 &= \max_{x,y_1} px + ry_1 - c(x, y_1) \\ \pi_2 &= \max_{y_2} -ry_2 - e(y_2) \end{aligned}$$

The first order conditions are:

$$\begin{aligned} p &= \frac{\partial c(x, y_1)}{\partial x} \\ r &= \frac{\partial c(x, y_1)}{\partial y} \\ -r &= \frac{\partial e(y_2)}{\partial y} \end{aligned}$$

Now, in equilibrium, it must be that $y_1 = y_2$. Hence, we have: $p = \frac{c(x,y))}{\partial x}$ and $\frac{c(x,y))}{\partial y} = -\frac{e(y))}{\partial y}$ —which is exactly the condition for efficiency. Notice —the second condition states that marginal cost to firm 1 of increasing pollution should coincide with the marginal cost to firm 2 of accepting it. (What reallocation could take place if this weren't true?)

Creating markets for externalities seems like a promising way forward. The problem is that these markets are often very thin. For example, in the above example, the market consisted of just two participants —a single buyer and single seller —which makes the assumption of a competitive market rather dubious.

17.2.2 Truth-telling Mechanism

(This subsection is drawn from Varian 2010.) We previously noted that significant problem with the Pigouvian tax approach is that it requires the government to know what the efficient tax is —an informational requirement that is probably too demanding. But, it might be less unreasonable to believe that the firms themselves know the level of externalities, and thus the efficient tax rate. Can the government design a mechanism to get the firms to reveal this ideal tax rate?

Yes. The procedure is in two stages:

1. Announcement Stage: Firms $i \in \{1, 2\}$ each (independently) announce a Pigouvian tax t_i (not necessarily truthfully).
2. Choice Stage: If firm 1 produces a positive quantity x, it pays a tax $t_2 x$, and firm 2 receives compensation of $t_1 x$. Additionally, each firm pays a penalty that depends on the difference between their announced rates.

The form of the penalty is inconsequential, except that it is zero if $t_1 = t_2$ and positive otherwise. For tractability, suppose the penalty is $(t_1 - t_2)^2$. Hence, the payoffs are:

$$\pi_1 = \max_x px - c(x) - t_2 x - (t_1 - t_2)^2$$
$$\pi_2 = -e(x) + t_1 x - (t_1 - t_2)^2$$

It turns out the this mechanism is truth-revealing. Let us solve the for the sub-game perfect equilibrium, using backward induction. In the second stage only firm 1 has a choice. The first order condition is:

$$p = c'(x) + t_2$$

which is precisely the output that results from a pigouvian tax of t_2. Let $x(t_2)$ be the resulting output. Note that $x'(t_2) < 0$.

What about the first stage? Each firm chooses the tax to announce, taking as given their opponent's announcement, and the equilibrium second stage decision. Firm 1's decision is easy. t_1 only enters her payoff through the penalty; she should set $t_1^*(t_2) = t_2$.

Firm 2's decision is less straight-forward. The problem is:

$$\max_{t_2} -e(x(t_2)) + t_1 x(t_2) - (t_1 - t_2)^2$$

The first order condition is:

$$[t_1 - e'(x(t_2))]x'(t_2) - 2(t_1 - t_2) = 0$$

Now, since $t_1 = t_2$, and since $x'(t_2) < 0$, this simplifies to:

$$t_2 - e'(x(t_2)) = 0$$

But from firm 2's first order condition from stage 2, we know that $t_2 = p - c'(x(t_2))$. Thus, firm 2 will choose t_2 to ensure:

$$p = c'(x(t_2)) + e'(x(t_2))$$

which is the condition for efficiency.

What is the intuition? For firm 1, it is straight-forward. It wants to match firm 2's announcement to avoid the penalty. What about firm 2? If firm 1's announces t_1 higher than the efficient tax, then firm 2 will wants to set a low tax, to entice firm 1 to produce high x, which will result in a high compensation (which more than makes up for the cost of pollution). But, knowing this, firm 1 will not make a high announcement. Similarly, if firm 1 announces t_1 lower than the efficient level, then firm 2 will announce a very high t_2 to discourage firm 1 from producing, since firm 2 will not be sufficiently compensated. But this will encourage firm 1 to raiser her announcement. The only tax rate that is equilibrium consistent is the efficient one, since it *just* compensates firm 2.

Chapter 18

Games with Incomplete Information

18.1 Normal Form 'Bayesian' Games

In all the games considered to this point, we have assumed that the agents have perfect information about the payoffs of all agents (both themselves and their opponents). In reality, there may be considerable uncertainty about this. For example:

1. Firms may not know each others' cost functions.
2. A seller's willingness to part with a good may depend on its quality, which the buyer might not know.

To capture this uncertainty, we assume that each agent can identify its different potential realizations, and associates each with a 'state of the world'. For example, firm 1 might know that firm 2 has a constant marginal cost technology, and that the marginal cost can take one of 3 values: $\{c_L, c_M, c_H\}$, where $c_H > c_M > c_L > 0$. We can think of there being 3 'states': $\{L, M, H\}$, each associated with a different cost function for firm 2. Firm 1's uncertainty, then, is captured by its beliefs over the likelihood that each of the 3 states is actually the true state. E.g. firm 1 may have beliefs $\Pr(L) = 0.6$, $\Pr(M) = 0.3$ and $\Pr(H) = 0.1$, in which case, it doesn't know the true state, but it believes that it is twice as likely that L is the state as M, which, in turn, is three times as likely to be the state as H. [Note —this conceptualization of uncertainty requires that the agents know or at least have beliefs about what the possibilities are. The agent cannot say: 'I have no earthly idea what my opponent's cost function looks like.']

How does the agent assign probabilities to different states? We allow for the possibility that agents receive a (noisy) signal about the state of the world, that guides their belief formation. For example, suppose the agents are playing Battle of Sexes, and the relevant states are $\{Rainy, Sunny\}$. The payoffs from the different options (opera and sports) are dependent on the weather on the evening of the event. The players must decide which event to attend in the morning. Before making their decision, each player might observe some signal about the likely evening weather (e.g. is it cloudy? are people carrying umbrellas?) The players form beliefs about the likelihood of rain accordingly. Different players may observe different signals and form different beliefs.

An implication of players receiving signals is that the same player might, depending on which signal they receive, choose different strategies. We often refer to a player's signal as their 'type', and expect that strategies will be 'type-dependent'.

With this discussion in mind, let us define a Normal Form Bayesian Game.

Definition 41. A *Bayesian Game with Imperfect Information* consists of:

- a set of players, $i \in I$.
- a set of actions A_i for each player $i \in I$
- a set of states, $\omega \in \Omega$.
- a set of signals T_i that agent i may receive, and a signal function $\tau_i(\omega)$ that associates a signal with each state. (We sometimes refer to an agent's signal as her type.)
- a belief function (probability distribution) $\pi_i(\omega|t_i)$ for each player, about the states consistent with the signal.
- a state dependent payoff function $u_i(a, \omega)$ for each player $i \in I$, which depends on the action profile $a = (a_1, .., a_I)$

Example 68. Consider a variant of the battle of the sexes game (due to Osborne (2003)) in which player 1 is unsure whether player 2 wants to meet her or to avoid her. The game is summarized in the game tableau below:

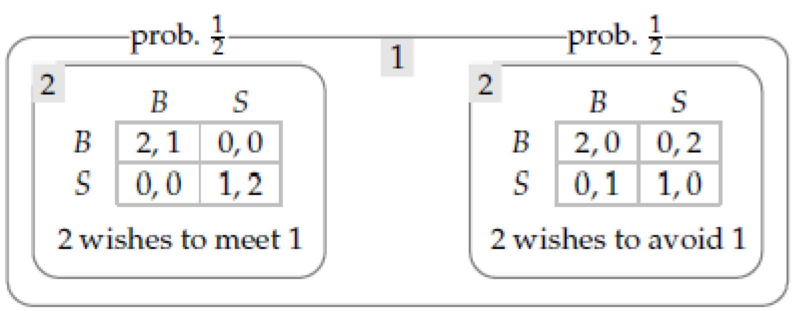

The main features of the game are as follows: There are two states of the world: $\Omega = \{meet, avoid\}$. Player 1 receives signal z and has signal function: $\tau_1(meet) = z = \tau_1(avoid)$, which captures the idea that the signal is uninformative for player 1. Player 2 receives one of two signals, m and v; her signal function satisfies: $\tau_2(meet) = m$ and $\tau_2(avoid) = v$, which captures the idea that her signal perfectly reveals the state (i.e. player 2 knows her preferences!). Player 1's beliefs satisfy $\pi_1(meet) = \frac{1}{2}$ and $\pi_1(avoid) = \frac{1}{2}$. Player 2's beliefs satisfy: $\pi_2(meet) = 1$, $\pi_2(avoid) = 0$ and $\pi_2(avoid) = 1$, $\pi_2(meet) = 0$.

Definition 42. A *pure strategy* for player i is a function $a_i : T_i \to A_i$ that assigns an action to each signal (information set).

Definition 43. A profile of strategies $(a_1, .., a_I)$ is a *Bayes Nash Equilibrium* if every type of every player maximizes his expected utility, given the strategies of every other player (and associated types) and the players beliefs about the state of the world. Formally, a strategy profile $(a_1^*(t_1), ..., a_I^*(t_I))$ is a Bayes Nash Equilibrium if, for each type $t_i \in T_i$ of player i:

$$\sum_{\omega \in \Omega} \pi_i (\omega | t_i) \, u_i \left(a_i^* (t_i), a_{-i}^* (\tau_{-i} (\omega)), \omega \right) \geq \sum_{\omega \in \Omega} \pi_i (\omega | t_i) \, u_i \left(a', a_{-i}^* (\tau_{-i} (\omega)), \omega \right)$$

for every $a' \in T_i$.

Example 69. In the Battle of the Sexes variant, the unique Bayes Nash Equilibrium is given by the strategy profile: $a_1^* (z) = B$ and $a_2^* (m) = B$ and $a_2^* (v) = S$. To verify that this is an equilibrium, we check that neither player would want to change their actions, given the equilibrium action of the other. Since player 1 chooses B, player 2 would choose B after receiving message m (since this gives her utility 1, which is preferable to utility 0 received if she chose S), and she would choose S after receiving message v. Expecting player 2 to behave in this way, and given his beliefs that each state is equally likely, player 1's expected utility from choosing B is $\frac{1}{2} \times 2 + \frac{1}{2} \times 0 = 1$, whilst his expected utility from choosing S is $\frac{1}{2} \times 0 + \frac{1}{2} \times 1 = \frac{1}{2}$. Clearly B is the better choice.

In fact, it is optimal for player 1 to choose B no matter what probability he assigns to the two states (as long as he puts some probability weight on the *avoid* state). By contrast, in the standard (complete) information game, there were two Nash Equilibria —one in which the players coordinated on S. Introducing uncertainty can help with equilibrium selection and coordination.

Example 70. Contagion. Consider the following game (also due to Osborne (2003)):

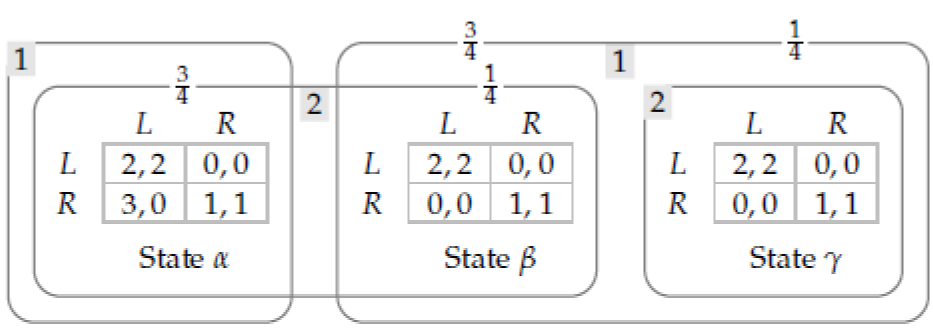

In this game, agents are uncertain about each other's knowledge. There are 3 states —α, β, γ. Player 1's preferences are the same in states β and γ, but differ in state α. Player 1 knows his preferences, so he knows whether the state is α or not. (But note, he cannot distinguish states β and γ). Player 2's preferences are identical in all states. She can distinguish state γ, but not states α and β from each other. If the true state is γ, both players will know which game they are playing. By contrast, if the state is α or β, player 1 will know the game form, but player 2 will not. To formalize the information sets and signalling technology, suppose $T_1 = \{a\neg a\}$, where $\pi_1(\alpha \,|\, a) = 1$ and $\pi_1(\alpha \,|\, \neg a) = 0$ with $\pi_1(\beta \,|\, \neg a) = \frac{3}{4}$ and $\pi_1(\gamma \,|\, \neg a) = \frac{1}{4}$. Similarly, let $T_2 = \{c, \neg c\}$, where $\pi_2(\gamma \,|\, c) = 1$ and $\pi_2(\gamma \,|\, \neg c) = 0$ with $\pi_2(\alpha \,|\, \neg c) = \frac{3}{4}$ and $\pi_2(\beta \,|\, \neg c) = \frac{1}{4}$.

The unique Nash equilibrium of this game is for both types of both players to choose R. To see this, begin by considering the type a of player 1 (i.e. player 1 knows the state is α). It is a dominant strategy for player 1 to choose R. Hence, we know that $a_1^*(a) = R$. Next, consider a type $\neg c$ of player 2 (i.e. the type who knows the state is either α or β, but cannot distinguish these). Player 2's expected utility if she chooses L is either 0 (if player 1 chooses R in state β) or $\frac{1}{2}$ (if player 1 chooses L in state β). Hence, choosing L gives expected utility of at most $\frac{1}{2}$. By contrast, if player 2 chooses R, her expected utility is at least $\frac{3}{4}$ (and potentially 1 if player 1 chooses R in state β). Hence, it is a dominant strategy for player 2 to choose R when her signal is $\neg c$. We have: $a_2^*(\neg c) = R$.

Now, consider a type $\neg a$ of player 1. He cannot distinguish states β and γ, but knows that if the state is β, that his opponent will choose R. Hence, choosing R gives expected utility of at least $\frac{3}{4}$, whilst choosing L gives expected utility of at most $\frac{1}{2}$. Hence $a_2^*(\neg a) = R$.

Then finally, the c type of player 2 chooses R, knowing that her opponent will always choose R in states β and γ.

What is interesting about this example is that, when the state is γ, both players know that they are playing a coordination game. In the world with certainty, there would be two equilibria —one where both agents chose L and one where they both chose R. However, adding uncertainty, even though all agents know the game payoffs, only one equilibrium survives. Why? Player 1 and player 2 both know what the payoffs are. [Player 2 knows the state is γ. Player 1 knows it is either β or γ. In either case, the payoffs are the same.] Furthermore, player 2 knows that player 1 knows what the payoffs are. [Player 2 knows that since the state is γ, player 1 must have received the signal $\neg a$.] However, player 1 doesn't know whether player 2 knows what the payoffs are. [Player 1 believes the state is β with probability $\frac{3}{4}$. But when the state is β, player 1 knows that player 2 will receive signal $\neg c$, which means player 2 might believe that the true state is α, which has different payoffs.] Moreover, although player 2 knows the true payoffs, and player 2 knows that player 1 knows the true payoffs, player 2 understands that player 1 doesn't know whether player 2 knows the true payoffs. This uncertainty causes both players to play 'conservatively'. Player 1 has a strict incentive to choose R just in case player 1 thinks the true state is α. But, understanding this, player 2 chooses R as well —even though, given the payoffs as understood by both players, (L, L) would also be an equilibrium. This example illustrates the important role of uncertainty and beliefs, not just in terms of what each player directly knows about the game, but what each player believes the other player knows about the game, what each player believes the other player be lives that she believes about the game, and so on.

Example 71. (Providing a Public Good) Consider a polity with n people. There is a public good which will be provided if at least one member of the group pays the cost c. Let v_i be agent i's valuation of the public good. Valuations are private, so each agent only knows their own valuation, but not the valuations of the other agents. However, it is known that each agent's valuation is an i.i.d draw from a distribution, and that $F(v)$ is the probability that $v_i < v$ for each v. We assume $v \in [\underline{v}, \overline{v}]$, so that $F(\underline{v}) = 0$ and $F(\overline{v}) = 1$. Furthermore, $c < \overline{v}$ and so $F(c) < 1$. (It is not the case that all agents' valuations are definitely below the cost.)

We say a strategy profile is symmetric if all similarly situated agents employ the same strategy. Let us characterize the symmetric Bayesian Nash equilibrium of this game.

Take a given player i. First, since each of the other agents decisions' to contribute or not depends on their valuation, and these are unknown to player i, from player i's perspective, it is uncertain whether *any* of her opponents will provide the public good. Let P_i be the probability that she assigns to none of her opponents contributing. Player i's utility from contributing is $u_i(C) = v_i - c$ and her expected utility from not contributing is $u_i(N) = (1 - P_i)v_i$. Thus, she should contribute provided that:

$$\begin{aligned} u_i(C) &> u_i(N) \\ v_i - c &> (1 - P_i)v_i \\ v_i &> \frac{c_i}{P_i} \end{aligned}$$

I.e. she should employ a threshold strategy in which she contributes only if her valuation lies above some threshold.

Now, in a symmetric equilibrium, all agents should use the same threshold. Let this threshold be v^*. The probability that each agent does not contribute is $F(v^*)$, and so the probability that none of the other agents contributes is $P_i = F(v^*)^{n-1}$. Thus we have:

$$v^* F(v^*)^{n-1} = c$$

Since $F()^{n-1} \leq 1$, it follows that $v^* \geq c$. Moreover, if $v^* = c$, then we need $F(v^*) = F(c) = 1$, which violates our assumption that $F(c) < 1$. Hence $v^* > c$. Each player's threshold for contributing exceeds the cost. Thus it is possible for every agent to have a valuation exceeding the cost but lower than the threshold, which results in the public good not being provided, although each agent would be individually willing to pay for it. Why? Although each agent is willing to pay for the good, they would rather free-ride of other agents.

How is the threshold affected by the number of players n? Let v_n^* be the threshold with n players, so that $v_n^* F(v_n^*)^{n-1} = c$. Take $m > n$. Since $F(v_n^*) < 1$, then $F(v_n^*)^{m-1} < F(v_n^*)^{n-1}$, and so $v_n(v_n^*)^{m-1} < c$, and so $v_m^* > v_n^*$. As n increases, so does the threshold, so each agent becomes less likely to contribute, *ceteris paribus*. The incentive to free ride becomes stronger as the number agents.

18.2 Extensive Form Games

Definition 44. An *extensive form game* with imperfect information (and chance moves) consists of:

- a set of players, $i \in I$
- a set of terminal histories.
- a function (the player function) that assigns either a player or "chance" to every history that is a proper sub-history of some terminal history.
- a function that assigns to each history that the player function assigns to chance, a probability distribution over the actions available after that history (where such a distribution is independent of all other such distributions).
- for each player i, an information partition of the set of histories assigned to that player by the player function, such that for every history h in any given partition, the set $A(h)$ of actions is the same.
- for each player, preferences over the set of lotteries over the terminal outcomes.

Example 72. Card Game. Each of two players begins by putting a dollar in the pot. Player 1 is then dealt a card that is equally likely to be *High* or *Low*; she observes her card, but player 2 does not. Player 1 may *see* or *raise*. If she *sees*, she shows her card to player 2. If her card is *High*, she takes the money in the pot, and if it is *Low*, player 2 takes the money; in both cases the game ends. If she *raises*, she adds a dollar to the pot and player 2 chooses whether to *pass* or *meet*. If player 2 *meets*, she adds a dollar to the pot, and player 1 shows her card (with the same outcome as before). If she *passes*, player 1 takes the money in the pot.

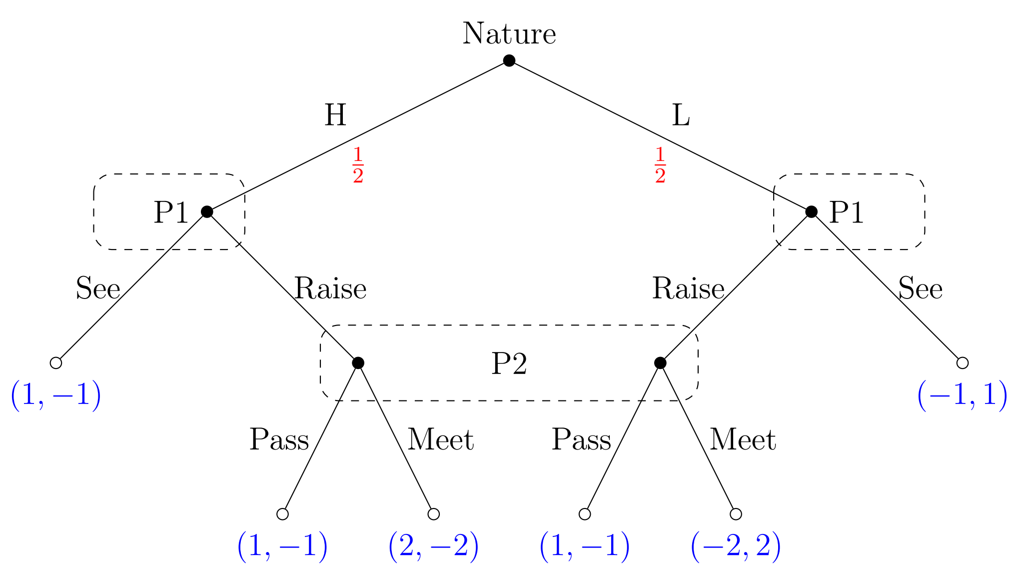

- The set of terminal histories is: $H = \{\phi, H, L, (H, See), (L, See), (H, Raise, Pass), (L, Raise, Pass),$
- The player function is: $P(h)$, where $P(\phi) =$ "*chance*", $P(H) = P(L) = 1$, and $P(H, Raise,) = P(L, Raise) = 2$.
- The chance moves: $\Pr[H|\phi] = \frac{1}{2} = \Pr[L|\phi]$.
- The information partitions: For player 1: $\{H\}$ and $\{L\}$. For player 2: $\{(H, Raise), (L, Raise)\}$.

Definition 45. A *behavioral strategy* for player i in an extensive form game, is a function that assigns to each of i's information sets I_i a probability distribution over the actions in $A(I_i)$, with the property that each probability distribution is independent of every other distribution.

Definition 46. A *belief system* in an extensive form game is a function that assigns to each information set, a probability distribution over the histories in that information set. An *assessment* is a pair (β, μ) consisting of a profile of behavioral strategies and a belief system.

Definition 47. An assessment (β, μ) is a *weak sequential equilibrium* if it satisfies the following conditions:

1. Sequential Rationality: Each player's strategy is optimal in the part of the game that follows each of her information sets, given the strategy profile and her belief about the history in the information set that has occurred. Precisely, for each player i and each information set I_i of player i, player i's expected payoff to the probability distribution $O_{I_i}(\beta, \mu)$ over terminal histories generated by her belief μ_i at I_i and the behavior prescribed subsequently by the strategy profile β is at least as large as her expected payoff to the probability distribution $O_{I_i}((\gamma_i, \beta_{-i}), \mu)$ generated by her belief μ_i at I_i and the behavior prescribed subsequently by the strategy profile (γ_i, β_{-i}), for each of her behavioral strategies γ_i.
2. Weak consistency of beliefs with strategies: For every information set I_i reached with positive probability given the strategy profile β, the probability assigned to by the belief system to each history h^* in I_i is consistent with Bayes' Rule.

Example 73 (Card Game (cont.))**.** Solve the game backward. Let $I_2 = \{(H, Raise), (L, Raise)\}$ denote player 2's information set when it is her turn to move. Let $\mu_2(H|I_2) \in [0, 1]$ denote probability she assigns to the the card being High. Player 2's expected utility from choosing Pass is -1, and her expected utility from choosing Meet is $2 - 4\mu$. Hence: $\beta_2(meet|I_2) = 1$ if $\mu_2(I_2) < \frac{3}{4}$, $\beta_2(meet|I_2) = 0$ if $\mu_2(I_2) > \frac{3}{4}$ and $\beta_2(meet|I_2) \in [0, 1]$ if $\mu_2(I_2) = \frac{3}{4}$.

Now consider Player 1, and history $I_1 = \{H\}$. If he sees, his payoff is 1. If he raises, his payoff is $\beta_2(meet|I_2) \cdot 2 + (1 - \beta_2(meet|I_2)) \cdot 1 = \beta_2(meet|I_2) \cdot +1$. Hence, $\beta_1(raise|H) = 1$ if $\beta_2(meet|I_2) > 0$ (i.e. if $\mu_2(I_2) \leq \frac{3}{4}$) and $\beta_1(raise|H) \in [0, 1]$ if $\mu_2 > \frac{3}{4}$. Now consider history $I_1 = \{L\}$. If he sees, his payoff is -1. If he raises, his payoff is $\beta_2(meet|I_2) \cdot (-2) + (1 - \beta_2(meet|I_2)) \cdot 1 = 1 - 3\beta_2(meet|I_2)$. Hence, $\beta_1(raise|L) = 1$ if $\beta_2(meet|I_2) < \frac{2}{3}$ (i.e. if $\mu_2 \geq \frac{3}{4}$) and $\beta_1(raise|L) = 0$ if $\beta_2(meet|I_2) > \frac{2}{3}$ (i.e. if $\mu_2 \leq \frac{3}{4}$) and $\beta_1(raise|L) \in [0,1]$ if $\beta_2(meet|I_2) = \frac{2}{3}$ (i.e. if $\mu_2 = \frac{3}{4}$). Sequential rationality implies:

$$\mu_2(H|I_2) = \frac{\frac{1}{2}\beta_1(raise|H)}{\frac{1}{2}\beta_1(raise|H) + \frac{1}{2}\beta_1(raise|L)}$$

How does this hang together? Suppose $\mu_2(I_2) < \frac{3}{4}$. Then $\beta_2(meet|I_2) = 1$, which implies $\beta_1(raise|H) = 1$ and $\beta_1(raise|L) = 0$. Then, by Bayes' Rule, $\mu_2(H|I_2) = 1 \not< \frac{3}{4}$. Contradiction. Suppose $\mu_2(I_2) > \frac{3}{4}$. Then $\beta_2(meet|I_2) = 0$, which implies $\beta_1(raise|H) \in [0,1]$ and $\beta_1(raise|L) = 1$. Then, by Bayes' Rule, $\mu_2(H|I_2) = \frac{\beta_1(raise|H)}{\beta_1(raise|H)+1} \leq \frac{1}{2} \not> \frac{3}{4}$. Contradiction. Finally, suppose $\mu_2(I_2) = \frac{3}{4}$. Then $\beta_2(meet|I_2) \in [0,1]$. There are several cases to consider:

- If $\beta_2(meet|I_2) = 0$, then (as before) $\mu_2(H|I_2) \leq \frac{1}{2} \neq \frac{3}{4}$. Hence this cannot be the case.
- Suppose $\beta_2(meet|I_2) \in \left(0, \frac{2}{3}\right)$. Then $\beta_1(raise|H) = 1$ and $\beta_1(raise|L) = 1$, and so $\mu_2(H|I_2) = \frac{1}{2} \neq \frac{3}{4}$. Contradiction.
- Suppose $\beta_2(meet|I_2) \in \left(\frac{2}{3}, 1\right]$. Then $\beta_1(raise|H) = 1$ and $\beta_1(raise|L) = 0$, and so $\mu_2(H|I_2) = 1 \neq \frac{3}{4}$. Contradiction.
- Suppose $\beta_2(meet|I_2) = \frac{2}{3}$. Then $\beta_1(raise|H) = 1$ and $\beta_1(raise|L) \in [0,1]$, and so $\mu_2(H|I_2) = \frac{1}{1+\beta_1(raise|L)} = \frac{3}{4}$ provided that $\beta_1(raise|L) = \frac{1}{3}$.

Hence, in sequential equilibrium, player 1 will always raise after history H, raise with probability $\frac{1}{3}$ after history L and see with the complementary probability. Player 2 holds beliefs $\mu_2(H|raise) = \frac{3}{4}$, and meets with probability $\frac{2}{3}$, and passes with the complementary probability. The equilibrium expected payoffs for player 1 are: $E[u_1(see|L)] = -1$ and $E[u_1(raise|L)] = \frac{2}{3}(-2) + \frac{1}{3}(1) = -1$ and $E[u_1(raise|H)] = \frac{2}{3}(2) + \frac{1}{3}(1) = \frac{5}{3} > 1 = E[u_1(see|H)]$. (This confirms that it is optimal to always raise following history H, and to mix following history L.) The equilibrium expected payoffs for player 2 are: $E[u_2(meet|I_2)] = \frac{3}{4}(-2) + \frac{1}{4}(2) = -1 = E[u_2(see|I_2)]$, which confirms that mixing is optimal.

Chapter 19

Asymmetric Information

So far, we have studied market failures arising from market power (monopoly, oligopoly), externalities, and public goods. There is a third major source of market failure —asymmetric information. We usually divide asymmetric information problems into two types: problems arising from hidden information and problems arising from hidden actions. With hidden information, the more informed party has an incentive to use their informational advantage to gain a larger share of the surplus. Understanding this, the less informed party will behave more 'conservatively', including by not engaging in surplus creating transactions, which it otherwise would have done if it were symmetrically informed. This leads to the problem of *adverse selection*. When there is hidden information, the parties may wish to take actions to have the information revealed to mitigate the adverse selection. The nature of these actions depends on which party is taking the action. When the informed party takes an action that credibly communicates her information, we say there is *signalling*. By contrast, when the uninformed party takes an action designed to entice the informed party to reveal her information, we say there is *screening*.

Hidden action is common in Principal-Agent interactions, where a principal seeks to engage an agent to perform some task on her behalf. Often, the principal cannot perfectly monitor the agent. Then, if the agent's preferences differ from the principal's, the agent may behave differently from how the principal had intended. This is the problem of *moral hazard*. Parties can mitigate moral hazard problems by engaging in optimal contracting.

19.1 Adverse Selection

Example 74 (The Market for Lemons)**.** Consider a buyer on the used car market who meets a seller. Suppose a car may either be high quality or low quality. Let v_i be the value to the buyer of a quality-i type car with $v_H > v_L$ and let $r_i < v_i$ be the reservation price of a typical seller of quality-i type car, where $i \in \{L, H\}$. Assume that the seller knows the quality of the car, but the buyer doesn't —but has prior belief that the probability of being high quality is π. Assume there is buyer-competition, so the buyer will pay his full valuation to the seller.

What happens if the quality is known? The buyer will pay v_H for a high-quality car and v_L for a low quality car. Trades will always take place, regardless of quality. This is efficient since $v_i > r_i$. Regardless of quality, the trade is always social welfare enhancing, because for each quality type, the buyer values the car more than the seller.

Now let quality be unknown by buyer. Suppose the buyer offers $\pi v_H + (1-\pi)\,v_L$ —which is the average value of the car. A low-type seller will always accept the buyer's offer. But a high-type seller will reject the offer if $\pi v_H + (1-\pi)\,v_L < r_H$ (which will be true if $\pi < \frac{r_H - v_L}{v_H - v_L}$). If so, the buyer's offer will only be accepted by the low type. This means that (i) only the low quality of car ever gets sold (which is inefficient), and (ii) the buyer will have over-paid for the car. Obviously, this cannot be an equilibrium. In equilibrium, the buyer, knowing that only the low-quality seller will deal with him, will only offer to pay v_L, and this deal will proceed. But this definitely precludes a sale with the high quality seller.

This example illustrates the problem of *adverse selection.* Adverse selection occurs when one side of the market is less informed about some relevant characteristic (quality, cost, value etc.) than the other. There is *hidden information.* The strategic incentives for different agents to enter/exit the market result in inefficient outcomes. Since the uninformed agent cannot distinguish good and bad quality products, he can only afford to offer a price equal to the expected value. But this is often too low for high-quality sellers which forces high quality sellers out of the market. In equilibrium, only low quality products are sold. Bad quality goods drive out the good quality ones.

A particularly important example of adverse selection is in the market for healthcare. Suppose (as the Affordable Care Act requires), that insurers cannot charge different prices for different agents based on their risk status (e.g. based on pre-existing conditions). The best insurers can do is charge a premium equal to the expected medical costs incurred by the agent with the average risk status in the insured pool. But agents know (better) their risk status, and at this 'average' price, low risk ('healthy') agents might find insurance to be not worthwhile. So they will drop out of the market. But this makes the average agent remaining in the market riskier, which will cause insurers to raise premiums. With higher premiums, some more 'healthy' people will drop out of the market, causing the insured pool to become riskier, and so on. In extreme cases, this dynamic can cause insurance premiums to rise so high that only the most unhealthy individuals find it optimal to purchase insurance. All other individuals are driven out. We say there is a 'death spiral'.

Example 75 (Death spiral)**.** A buyer wants to purchase a good from a seller. Let v be the quality of the good, and suppose, at every quality level, the seller demands a price at least $r = kv$, where $k < 1$. The seller knows the quality, but the buyer believes it is drawn from a distribution $v \sim U\,[0,1]$. Since $k < 1$, it should be that at every quality level, the buyer values the good more than the seller. So if there were complete information, trades would always take place.

What happens with asymmetric information. Suppose the buyer offers a price P. The buyer knows that this price will be accepted by the seller if the quality of her good is $v \leq \frac{P}{k}$. Hence, conditional on the offer being accepted, buyer knows that the average quality of the good that he receives will be: $E[v \,|\, v < \frac{P}{k}] = \frac{P}{2k}$. Is it worthwhile to the buyer to make such a trade?

Suppose $k \leq \frac{1}{2}$. Then $E[v \,|\, v < \frac{P}{k}] \geq P$, so the buyer will on average be better off. This is true no matter the price. So the best that the buyer can do is offer $P = k$. This guarantees that every seller will sell to the buyer, and the expected value of the good received is $\frac{1}{2} > P$.

Suppose instead that $k \in \left(\frac{1}{2}, 1\right)$. Then, for every $P > 0$, $E[v \,|\, v < \frac{P}{k}] < P$. So any positive price that the buyer offers, she will have over-paid on average if the sale goes ahead. The only price consistent with optimal behavior is $P = 0$. Trade only occurs when the seller has the lowest quality good. We say the market 'completely unravels'. [One way to see this is to argue step-wise. Suppose the buyer originally assumes all sellers will part with the good. Then he will offer at most $P_0 = \frac{1}{2}$. But at this price, the buyer then reasons that only sellers with quality $v \leq \frac{1}{2k}$ will sell. So the buyer offers at most the average price in this sub-market, which is $P_1 = \frac{1}{4k}$. But at this lower price, the buyer reasons that only sellers with quality $v \leq \frac{1}{4k^2}$ will sell. So the buyer offers at most the average price in this sub-sub-market, which is $P_2 = \frac{1}{8k^2}$. This process of reasoning continues *ad infinitum*. At the n^{th} round, the buyer will offer at most $P_n = \frac{1}{2\cdot(2k)^n}$. Since $2k > 1$, its clear that this maximum offer is getting smaller in each round, and that $P_n \to 0$ as $n \to \infty$.]

19.2 Screening

In the adverse selection problem, we might think we could restore efficiency if the informed party could reveal information to the uninformed party. However there is no guarantee that the informed agent will be truthful - in the above case, a low-quality car owner has an incentive to claim to have a high quality car, in order to get a better price. 'Screening' occurs when the uninformed party creates incentives that entice the informed party to truthfully reveal information. The basic idea is to offer different versions of the product or contract —with different versions being more or less optimal for different types of sellers. (Examples of screening contracts include second degree price discrimination in product markets, different levels of deductibles in insurance etc.)

19.2.1 A Model of Second Degree Price Discrimination

There is a monopolist with constant marginal cost technology $c(q) = cq$. There are two types of potential buyers —high and low. The monopolist does not know the buyer's type, but knows that each buyer has a probability π of being the high type. A type $i \in \{H, L\}$ buyer has utility $u_i(q, T) = \alpha_i q - \frac{1}{2}q^2 - T$, where $\alpha_H > \alpha_L > c$, and T is the total cost to the buyer of purchasing q units of the good.

Suppose the unit price of the good is p. Then, the agent's problem is:

$$\max_q \alpha_i q - \frac{1}{2}q^2 - pq$$

Taking the first order conditions, we find the agent's (Marshallian) demand function is $q_i(p) = \alpha_i - p$. So the agents' utility functions imply linear demands with different intercepts. (This is identical to the demand function we studied we looked at third degree price discrimination.)

If the firm could third degree price discriminate, we know that it would charge a different price to each type of buyer. [In fact, if it could price discriminate *and* charge a two-part tariff, it would set a different fixed fee for different types, but then sell all units at marginal cost.] Here, the monopolist cannot distinguish the types, so third degree price discrimination is not possible. And the optimal two-part tariff will not work, since all agents would profess to being the low type in order to pay the lower fixed fee. Rather than setting a unit price and letting the agent choose how many units to purchase, suppose the monopolist instead creates 'packages' (q, T) where it offers to sell q units at total package price T. [This implies a per-unit price of $\frac{T}{q}$ for this package, but different packages may have different implied per-unit prices.] Since there are two types of buyers, it suffices for the monopolist to offer two types of packages —one aimed at each type.

As a benchmark, consider the problem under full information. What packages would the monopolist charge? For a type i buyer, the monopolist seeks to maximize:

$$\max_{q,T} \; T - cq \; s.t. \; \alpha_i q - \frac{1}{2}q^2 - T \geq 0$$

The constraint is known as the *participation constraint* or *individual rationality (IR)* constraint. It states that the buyer must be made no worse off by buying her intended package, rather than not, or she would simply not purchase. Naturally, the monopolist will want to set T as high as possible without violating the constraint. So $T = \alpha_i q - \frac{1}{2}q^2$ —the firm sets a package price that extracts all of the buyer's consumer surplus.

What quantity will it sell? Substituting the constraint into the objective gives:

$$\max_q \alpha_i q - \frac{1}{2}q^2 - cq$$

i.e. the monopolist chooses q to maximize the joint surplus between buyer and seller. [This should be obvious, since the monopolist then goes on to extract all of the surplus which would otherwise accrue to the buyer. So the seller wants to maximize the total surplus.] Solving the first order condition gives: $\hat{q}_i = \alpha_i - c$. It is as if the seller prices at marginal cost, and the extracts the full consumer surplus through the package fee. The fee for each type would be $\hat{T}_i = \frac{1}{2}(\alpha_i - c)^2$.

19.2.2 Separating Equilibrium

Now consider the screening contracts with hidden information. The principal offers two contracts (q_H, T_H) and (q_L, T_L), targeted at the respective agents. However, the seller cannot force the buyer to choose the intended package. They must willingly do so (and in so doing, they reveal their private information about their type). Thus, in addition to the participation constraint, the contracts must satisfy the *incentive compatibility constraint*:

$$\alpha_i q_i - \frac{1}{2}q_i^2 - T_i \geq \alpha_i q_{-i} - q_{-i}^2 - T_{-i}$$

The incentive compatibility constraint states that a type i buyer should do at least as well purchasing the package intended for her type, than purchasing the other package. We say, neither type has an incentive to 'imitate' the other.

The monopolist's problem is:

$$\begin{array}{rl} \max\limits_{q_H, q_L, T_H, T_L} & \pi(T_H - cq_H) + (1-\pi)(T_L - cq_L) \\ & \text{subject to} \\ (IR_H) & \alpha_H q_H - \frac{1}{2}q_H^2 - T_H \geq 0 \\ (IR_L) & \alpha_L q_L - \frac{1}{2}q_L^2 - T_L \geq 0 \\ (IC_H) & \alpha_H q_H - \frac{1}{2}q_H^2 - T_H \geq \alpha_H q_L - \frac{1}{2}q_L^2 - T_L \\ (IC_L) & \alpha_L q_L - \frac{1}{2}q_L^2 - T_L \geq \alpha_L q_H - \frac{1}{2}q_H^2 - T_H \end{array}$$

For each agent only one of these constraints will bind. (To see this, take a H-type buyer, and notice that the LHS terms of IR_H and IC_H are the same, but the RHS terms are different. Generically, one of the RHS terms will be larger than the other. Since both inequalities must be satisfied in equilibrium, it must be that the more demanding inequality is satisfied. Note —it cannot be that both inequalities were slack. If it were, then the monopolist must be leaving money on the table, and thus not profit maximizing.)

Which constraints bind. Intuitively, the high type wants to buy more and is willing to pay a higher price than the low type. So, in equilibrium, the low type will never do better to imitate the high type. By contrast, since the monopolist will charge the low type a lower price, the high type might want to imitate the low type, in order to get the lower price. Thus, the participation constraint will bind for the low type, and the incentive compatibility constraint will bind for the high type. [It is not difficult to rigorously show that this must be the case, by simply guessing other pairs of binding constraints, and showing that you get a contradiction. For example, it cannot be that IR_H and IR_L both bind. If IR_L binds, so that $\alpha_L q_L - \frac{1}{2}q_L^2 - T_L = 0$, then since $\alpha_H > \alpha_L$, it must be that $\alpha_H q_L - \frac{1}{2}q_L^2 - T_L > 0$. So type-$H$ strictly prefers type L's bundle to not participating. But incentive compatibility states that H must (weakly) prefer her own bundle to L's bundle. Thus, H must strictly prefer her own bundle to not participating, which means that IR_H cannot be binding.]

Using IR_L, we have: $T_L = \alpha_L q_L - \frac{1}{2}q_L^2$. Using IC_H, we have:

$$T_H = T_L + \alpha_H(q_H - q_L) - \frac{1}{2}(q_H^2 - q_L^2) = \underbrace{\alpha_H q_H - \frac{1}{2}q_H^2}_{\text{Consumer Surplus}} - \underbrace{(\alpha_H - \alpha_L)q_L}_{\text{Info Rent}}$$

The monopolist is able to extract the full consumer surplus from the low type, but not from the high type. If it tried to fully expropriate the high type, H would find it better to mimic the low type. Hence, the monopolist must 'leave some money on the table' for the high type —just enough that the high type prefers to reveal herself as the high type. We call this an *informational rent.* The uninformed party must pay the informed party to reveal her information! Substituting these into the firm's objective problem gives:

$$\max_{q_H, q_L} \pi[\alpha_H q_H - \frac{1}{2}q_H^2 - cq_H - (\alpha_H - \alpha_L)q_L] + (1-\pi)[\alpha_L q_L - \frac{1}{2}q_L^2 - cq_L]$$

The first order conditions are:

$$\frac{\partial \pi}{\partial q_H} = \pi(\alpha_H - q_H - c) = 0$$
$$\frac{\partial \pi}{\partial q_L} = (1-\pi)(\alpha_L - q_L - c) - \pi(\alpha_H - \alpha_L) = 0$$

Solving the FOCs gives: $q_H^* = \alpha_H - c$ and

$$q_L^* = \begin{cases} \alpha_L - c - \frac{\pi}{1-\pi}(\alpha_H - \alpha_L) & \pi \leq \frac{\alpha_L - c}{\alpha_H - c} \\ 0 & \pi > \frac{\alpha_L - c}{\alpha_H - c} \end{cases}$$

where, we acknowledge the possibility of a corner solution at $q_L = 0$ if π is high enough.

What do we notice? The monopolist produces the efficient quantity for the high type, but an inefficiently low quantity for the low type. We say there is downward distortion of the low type. As before, the monopolist full extracts the low type's consumer surplus, but leaves the high type with an informational rent. Why is there downward distortion of the low-type's quantity? First, we know that the low type will not seek to imitate the high type, in equilibrium, so there is no need to distort the high type's package. However, we know that the high type might seek to imitate the low type, and to prevent this, the monopolist must give the high type an informational rent. The rent will be higher the more attractive it is for the high type to imitate the low type. [As we showed previously, this benefit from pretending to be a low type is proportional q_L.] Thus, the monopolist can reduce the informational rents it pays to the high type by 'damaging' the low-type package, thus making it less attractive. Of course, by 'damaging' the low-type package, the monopolist earns lower profits from the low type. Thus there is a trade-off. Damaging benefits the monopolist to the extent the eventual buyer is a high type, but hurts the monopolist if the eventual buyer is a low type. Naturally, the rate at which the monopolist trades these off depends on the likelihoods ($\frac{\pi}{1-\pi}$) of meeting a high versus low type agent. The higher is the likelihood of meeting a high type agent, the more damage the monopolist will be willing to inflict. In fact, if π is high enough, the monopolist will simply not sell to the low type, and thus not need to provide any informational rents to the high type.

There is an important broader insight here. We see 'versioning' in a variety of different product markets. For example, airlines have 'business', 'premium economy' and 'economy' versions of what is largely the same product. The business and premium economy versions have more space and more comfort, whereas economy is designed to be unpleasant. Airlines don't necessarily squeeze economy seats close together to fit more seats onto the plane. (They are simultaneously giving more room to the more expensive versions, so that overall, there may actually be fewer seats on the flight.) The point is to make the cheaper option less attractive to wealthy travelers, thus encouraging them to buy the higher cost versions. Economy is not bad because the airline seeks to punish lower income travelers or because it cannot afford to provide better comfort. It is uncomfortable in order to entice travellers to pay for upgrades.

This insight was first noted by Jules Dupruit (a 19th Century French economist) who wrote:

> 'It is not because of the few thousand francs which would have to be spent to put a roof over the third-class seats that some company or other has open carriages with wooden benches.... What the company is trying to do is prevent the passengers who can pay the second-class fare from traveling third-class; it hits the poor, not because it wants to hurt them, but to frighten the rich.... And it is again for the same reason that the companies, having proved almost cruel to third-class passengers and mean to second-class ones, becomes lavish in dealing with first-class passengers. Having refused the poor what is necessary, they give the rich what is superfluous.'

19.2.3 Pooling Equilibrium

The separating equilibrium is useful because it allows the monopolist to separate the types and get more output from the high-type. However, separating the types is costly for the principal. Might the principal simply be better off not separating the types? Whilst this doesn't necessarily get the efficient output, it saves on informational rents.

A pooling contract is one where the monopolist offers the same contract to both potential buyers. (Obviously the IC constraints no longer matter, since there is no longer an opportunity to imitate the other type.) Of course, the monopolist must still satisfy the participation constraints.

The monopolist's problem is:

$$\max_{q_P, T_P} \; T_P - cq_P$$
$$\text{subject to}$$
$$(IR_H) \quad \alpha_H q_P - \frac{1}{2}q_P^2 - T_P \geq 0$$
$$(IR_L) \quad \alpha_L q_P - \frac{1}{2}q_P^2 - T_P \geq 0$$

Since both constraints must be satisfied and $\alpha_H > \alpha_L$, it should be clear that (IR_L) will bind and (IR_H) will not. Hence $T_P = \alpha_L q_P - \frac{1}{2}q_P^2$. Substituting this into the objective function reduces the problem to:

$$\max_{q_P} \alpha_L q_P - \frac{1}{2}q_P^2 - cq_P$$

which coincides with the monopolist's full information problem when facing the low type. We have: $q_P^* = \hat{q}_L = \alpha_L - c$ and $T_P = \hat{T}_L = \frac{1}{2}(\alpha_L^2 - c^2)$. The firm's profit is: $\Pi_P = \frac{1}{2}(a_L - c)^2$. The firm produces the efficient quantity for the low type and fully expropriates all of the low type's consumer surplus. The quantity is inefficiently low for the high type, but the high type retains a positive surplus.

19.2.4 Other Applications

- Principal is buyer and agent is seller with unknown costs. (Baron and Myerson, 1982). Low cost firm is undistorted and receives informational rent. High cost firm's quantity is downward distorted, and receives no rents.

- Principal is lender, agent is borrower with unknown productive. (Freixas and Laffont, 1990). High productivity borrower is undistorted in credit market and earns an informational rent. Low productivity borrower is rationed in lending, but receives a lower interest rate.

19.3 Signalling

Screening occurs when the uninformed party creates incentives for the informed party to reveal their information. As we saw, this essentially requires the uninformed party to 'pay' the informed party. Now we consider the alternative approach where the informed party voluntarily reveals their information to the uninformed party. The informed party must do so in a way that is *credible* —since talk is cheap, it is not enough to say that you are the high-type. *Signalling* occurs when the high type takes a costly action that the low type would never find it optimal to do, thus revealing her type.

19.3.1 A Model of Education (Spence 1973)

Consider the following environment: There are two players: a worker and an employer. The worker's marginal product θ may either be high or low: $\theta_H > \theta_L$. The employer doesn't know the worker's productivity, but knows the prior probability π that the worker is a high type. The employer must decide how much to pay the worker. Labor markets are competitive, so the employer will pay the worker her (expected) marginal product.

The worker can signal her quality by investing in education. The cost of acquiring education is $c\,(e,\theta) = \frac{e}{\theta}$. (More generally, we take $c\,(e,\theta)$ such that c is increasing and convex in e for each θ, $c\,(e,\theta_H) < c\,(e,\theta_L)$ for every e, and $c_e\,(e,\theta_H) < c_e\,(e,\theta_L)$ for every e.) Hence, for both types of workers, each extra unit of education acquired is costlier than the previous unit, and the cost and marginal cost for the high type are lower than for the low type, at any level of education. The agent's utility function is $u\,(w,e;\theta) = w - c\,(e,\theta)$. Note importantly that education does not improve productivity —it is simply a tool used to demonstrate an agent's type, since it is less costly for high types can acquire education than low types.

A strategy for the worker is a pair (e_H, e_L), which specifies the level of education chosen by each type. A strategy for the employer is a function $w(e)$ that assigns a wage to be paid for every possible level of educational attainment. Naturally, $w(e) = E[\theta|e] = \Pr(\theta = \theta_H)\theta_H + (1 - \Pr(\theta = \theta_H))\theta_L$. Thus the wage function is directly related to the employer's beliefs about the worker's type. A *Perfect Bayesian Equilibrium* is a triple $(e_H, e_L, w(\cdot)$ such that:

- Each type of agent does strictly better choosing their equilibrium level of education than any other level, given the employer's wage function. I.e

$$w(e_i) - \frac{e_i}{\theta_i} \geq w(e') - \frac{e}{\theta_i}\ \forall e' \geq 0$$

- The employer must have correct beliefs about the equilibrium actions of the agents. Thus:
  - If $e_H \neq e_L$, the employer's beliefs should be $\Pr(\theta = \theta_H \,|\, e_H) = 1$ and $\Pr(\theta = \theta_H \,|\, e_L) = 0$. Thus $w(e_H) = \theta_H$ and $w(e_L) = \theta_L$.
  - If $e_H = e_L = \hat{e}$, the employer's beliefs should satisfy: $\Pr(\theta = \theta_H \,|\, e = \hat{e}) = \pi$, and so $w(\hat{e}) = \pi\theta_H + (1 - \pi)\theta_L$.
  - The employer is free to assign any beliefs for any e not chosen in equilibrium (i.e. 'off the equilibrium path'.)

19.3.2 Separating Equilibria

Suppose there is a separating equilibrium with effort levels (e_H, e_L), with $e_H \neq e_L$. In equilibrium, it must be that $w(e_H) = \theta_H$ and $w(e_L) = \theta_L$. These education choices must be optimal for each agent type. First, take the low-type agent. Incentive compatibility implies:

$$\theta_L - \frac{e_L}{\theta_L} \geq w(e) - \frac{e}{\theta_L} \;\; \forall e$$

In particular, this must be true for $e = 0$, and so $\theta_L - \frac{e_L}{\theta_L} \geq w(0)$. Now $w(e) \in [\theta_L, \theta_H]$ for every e, so $w(0) \geq \theta_L$. Hence $\theta_L - \frac{e_L}{\theta_L} \geq \theta_L$, which can only be true if $e_L = 0$. Hence, $w(0) = \theta_L$. Since the low-type agent will be discovered in equilibrium, and education is costly, the low type will do best by not educating. Of course, in equilibrium, the low type must not have an incentive to imitate the high type. Incentive compatibility for type L implies:

$$\theta_L \geq \theta_H - \frac{e_H}{\theta_L}$$

and so $e_H \geq \theta_L(\theta_H - \theta_L)$. This is the minimal level of education that is sufficiently costly that the low type would rather reveal herself as a low-type than try to mimic the high type. It is costly enough to credibly signal that the agent is a high type.

Incentive compatibility for type-H is:

$$\theta_H - \frac{e_H}{\theta_H} \geq w(e) - \frac{e}{\theta_H}$$

for any other e. In particular, incentive compatibility must hold for $e = 0$. (I.e. the high-type must not want to imitate the low-type.) Hence, $\theta_H - \frac{e_H}{\theta_H} \geq \theta_L$, and so $e_H \leq \theta_H(\theta_H - \theta_L)$. Hence $e_H \in [\theta_L(\theta_H - \theta_L), \theta_H(\theta_H - \theta_L)]$, and $w(e_H) = \theta_H$.

Are there any other profitable deviations for the players? In equilibrium, there shouldn't be. To ensure this, we need that $w(e) \leq \theta_L - \frac{e}{\theta_L}$ *and* that $w(e) \leq \theta_H - \frac{e_H - e}{\theta_H}$. (The former ensure that the low type has no profitable deviation, and the latter ensures that the high type doesn't. It suffices to set $w(e) = \theta_L$ whenever $e < e_H$ and $w(e) = \theta_H$ whenever $e \geq e_H$. (Also, the employer could se $w(e) = \theta_L$ for all $e \neq e_H$.)

We can sustain a continuum of separating equilibria with $e_L = 0$ and $e_H \in [\theta_L(\theta_H - \theta_L), \theta_H(\theta_H - \theta_L)]$. Amongst these equilibria, the most efficient is the one in which $e_L = 0$ and $e_H = \theta_L(\theta_H - \theta_L)$ —the low-type does not engage in costly education, and the high-type educates the minimum amount to separate himself from the low type. Notice that, with signalling contracts, the 'good' type must pay to reveal his type, whereas with screening contracts, the 'good' type was paid to reveal his information!

19.3.3 Pooling Equilibria

Suppose there is a pooling equilibrium, where $e(\theta_H) = e(\theta_L) = e_P$. Then, necessarily, $\mu(e_P) = \pi$ and as such, $w(e_P) = \pi(\theta_H - \theta_L) + \theta_L$. Again, these education choices must be optimal for each type. Incentive compatibility for the low type implies: $\pi(\theta_H - \theta_L) + \theta_L - \frac{e_P}{\theta_L} \geq w(e) - \frac{e}{\theta_L}$, for every e. In particular, this must hold true for $e = 0$. Hence, $\pi(\theta_H - \theta_L) + \theta_L - \frac{e_P}{\theta_L} \geq w(0) \geq \theta_L$, and so $e_P \leq \pi\theta_L(\theta_H - \theta_L)$.

Are there any other profitable deviations for the players? Since the deviations are off the equilibrium path, we can assign beliefs to the principal however we like. It suffices to take $w(e) = \theta_L$ whenever $e < \pi\theta_L(\theta_H - \theta_L)$ and $w(e) = \pi\theta_H + (1-\pi)\theta_L$ for any $e \geq \pi\theta_L(\theta_H - \theta_L)$. As above, this implies that there is never an incentive to choose $e \in (0, e_P)$ or $e > e_P$. We've already checked that the L-type agent will not deviate from e_P to 0. It is simple to check that the same is true for the high-type. Hence, there are a continuum of pooling equilibrium, with $e_P \in [0, \pi\theta_L(\theta_H - \theta_L)]$. Of course, the most efficient pooling equilibrium has $e_P = 0$.

Which equilibrium is better? In the most efficient pooling equilibria, the agent utilities are $\left(u_H^P, u_L^P\right) = (\pi\theta_H + (1-\pi)\theta_L, \pi\theta_H + (1-\pi)\theta_L)$. In the most efficient separating equilibrium, the agent utilities are: $\left(u_H^S, u_L^S\right) = \left(\theta_H - \frac{\theta_L(\theta_H - \theta_L)}{\theta_H}, \theta_L\right)$.

Obviously the pooling equilibrium is better for the low-type agent. It is also better for the high-type agent if:

$$\begin{aligned} \pi(\theta_H - \theta_L) + \theta_L &\geq \theta_H - \frac{\theta_L}{\theta_H}(\theta_H - \theta_L) \\ \pi &\geq 1 - \frac{\theta_L}{\theta_H} \end{aligned}$$

Again this is intuitive. In the separating equilibrium, the high-type agent has to expend costly resources to signal his high-type. If low-types are abundant, then average income is low, and so there is a strong benefit to separating. But if low-types are few, then the average income is already high. The cost of signalling is larger than the gain it produces.

19.3.4 Other Applications

- Warranties: Agent is seller. Principal is buyer. A warranty is a costly action for the seller — they incur a cost whenever the product fails. Obviously it is costlier for the low-quality seller than the high quality one. Warranties may enable separation. (But it may or may not be efficient for sellers to carry out repairs –> pooling)

- Advertising: Agent is seller. Principal is buyer. Advertisements are costly, but send information about product quality. A larger campaign is needed to convince public of value of poor quality product.

- Corporate Finance: Agent is a risk averse entrepreneur with risky project. Principal is the 'market'. Holding equity in the firm is costly for the risk averse entrepeneur — but sends a signal about the project quality. An entrepreneur with a bad quality project would want to hold less equity than an entrepreneur with high quality project. Enables separation, but results in inefficient risk bearing.

- Exchange Rates: Agent is LDC government with poor reputation. Principal is the 'market'. Government has private information about their commitment to not devalue. Dollarizing is a costly signal, but more costly for less committed governments. Enables separation, but potentially inefficient defaults.

- Incomplete Contracts: Principal and Agent are both potential partners in a joint venture. Sender is better informed about his commitment to relationship than partner. Leaving contract incomplete (e.g. not signing pre-nup) is costly for committed partner (it may cause inefficient outcomes in future contingencies); but it is even more costly for the uncommitted partner. Enables separation, but results in inefficient outcomes *ex post.*

19.4 Moral Hazard

Adverse selection was a consequence of hidden information — the informational asymmetry resulted in inefficient allocations of goods (both in the sense of net-loss transactions occurring and value-creating transactions not taking place). The challenge with adverse selection (i.e. signalling/screening) is to provide the correct incentives for the informed party to credibly transfer her information. Moral Hazard, by contrast, is a consequence of hidden action. Consider a situation where an agent's choice of action affects some other person (the principal). If the principal could directly observe the agent's action, then he could directly contract for the agent to perform the desired action. By contrast, if the principal cannot perfectly observe the agent's actions, then she can only create incentives/rewards/punishments for the agent that depend on observable outcomes that are (statistically) correlated with the agent's action.

For example, in the workforce, entrepreneurs would like to write labor contracts that reward high effort and punish low effort. However, since owners can typically only imperfectly monitor their workers, this may not be possible. Instead, they can reward good or bad outcomes, even if sometimes, high effort results in bad outcomes and low effort results in good outcomes.

19.4.1 Simple Model

There are two players — a principal and agent. The agent's preferences are given by $w - c(e)$ where e is the effort level of the agent, w is the agent's wealth and c exhibits increasing marginal cost. The agent has an outside option of not working that generates utility $\bar{u}$.

For simplicity, suppose the agent can only exert one of two types of effort: e_L and e_H, with $e_L < e_H$. The agent's effort affects the likelihood that a project will succeed. Let p_L and p_H denote the probability of success given the different effort levels, and let x_S and x_F denote the payoff to the principal if the project succeeds and fails (respectively), with $x_S > x_F$. For simplicity, assume the principal has risk neutral preferences.

The principal's expected utility as a function of the agent's effort is:

$$\pi(e) = \begin{cases} p_H (x_S - x_F) + x_F & e = e_H \\ p_L (x_S - x_F) + x_F & e = e_L \end{cases}$$

The benefit to the principal of having the agent exert high effort is $\Delta\pi = (p_H - p_L)(x_S - x_F)$. The net cost of exerting high effort is $\Delta c = c(e_H) - c(e_L)$. Suppose that $\Delta\pi > \Delta c$, so that exerting high effort is efficient.

Suppose the principal want to entice the agent to exert high effort. He offers the agent a contract with outcome-dependent wage (w_S, w_F). The agent will exert high effort provided that this contract satisfies both the IC and IR constraints. The principal's problem is:

$$\begin{array}{ll} \max_{w_S,w_F} & p_H(x_S - w_S) + (1-p_H)(x_F - w_F) \\ \text{s.t.} & p_H w_S + (1-p_H) w_F - c(e_H) \geq p_L w_S + (1-p_L) w_F - c(e_L) \\ & p_H w_S + (1-p_H) w_F - c(e_H) \geq \bar{u} \end{array}$$

Notice that we can re-write the principal's problem in the following way:

$$\begin{array}{ll} \max_{w_S,w_F} & p_H x_S + (1-p_H)x_F - p_H(w_S - w_F) - w_F \\ \text{s.t.} & (p_H - p_L)(w_S - w_F) \geq c(e_H) - c(e_L) \\ & p_H(w_S - w_F) + w_F \geq c(e_H) + \bar{u} \end{array}$$

The incentive compatibility constraint depends only on the wage differential $w_S - w_F$. The individual rationality constraint and the principal's payoff both depend on w_F and the wage differential $w_S - w_F$. Intuitively, the principal will want to keep both w_F and $w_S - w_F$ low. But both of these will need to be high enough to incentivize the agent.

In equilibrium, it must be that both the IC and IR constraints bind. [To see this, note that, if the IC constraint did not bind, the principal could reduce $\Delta w = w_S - w_F$ a little without violating the constraint. This would generate higher expected profits for the principal without causing the agent to reduce her effort. Similarly, if IR did not bind, then (assuming IC binds, as we've shown it must), the principal could reduce w_F without causing the agent to reduce her effort.]

We have two equations in two variables (w_F, w_S) (or equivalently, $(w_F, \Delta w)$. We can solve these simultaneously. From the IC constraint:

$$\Delta w = \frac{c(e_H) - c(e_L)}{p_H - p_L}$$

Substituting this into the IR constraint gives:

$$\begin{aligned} w_F &= c(e_H) + \overline{u} - p_H \Delta w \\ w_F &= c(e_H) + \overline{u} - \frac{p_H}{p_H - p_L}(c(e_H) - c(e_L)) \\ &= \overline{u} - \frac{p_H c(e_L) - p_L c(e_H)}{p_H - p_L} \end{aligned}$$

Finally:

$$\begin{aligned} w_S &= w_F + \frac{c(e_H) - c(e_L)}{p_H - p_L} \\ &= \bar{u} + \frac{(1-p_L)e(c_H) - (1-p_H)e(c_L)}{p_H - p_L} \end{aligned}$$

Comments

This contract is optimal if $\bar{u} \leq p_H x_S + (1-p_H)\, x_F - c\,(e_H)$.

Proof. The optimal contract that induces high effort solves:

$$\begin{aligned} \begin{bmatrix} p_H - p_L & -(p_H - p_L) \\ p_H & (1-p_H) \end{bmatrix} \begin{bmatrix} w_S \\ w_F \end{bmatrix} &= \begin{bmatrix} \Delta c \\ \bar{u} + c_H \end{bmatrix} \\ \begin{bmatrix} w_S \\ w_F \end{bmatrix} &= \frac{1}{p_H - p_L} \begin{bmatrix} (1-p_H) & (p_H - p_L) \\ -p_H & (p_H - p_L) \end{bmatrix} \begin{bmatrix} \Delta c \\ \bar{u} - c_H \end{bmatrix} \\ &= \begin{bmatrix} \bar{u} + \frac{1-p_L}{p_H - p_L} c_H - \frac{1-p_H}{p_H - p_L} c_L \\ \bar{u} - \frac{p_L}{p_H - p_L} c_H + \frac{p_H}{p_H - p_L} c_L \end{bmatrix} \end{aligned}$$

Like-wise, the optimal contract that induces low effort is: $(w_F', w_S') = \left(\bar{u} + c_L - \frac{p_L}{p_H - p_L}\Delta c, \bar{u} + c_L + \frac{1-p_L}{p_H - p_L}\Delta c\right)$. The expected profits are: $\pi_H = p_H x_S + (1-p_H)\, x_F - \bar{u} - c_H$ $\pi_L = p_L x_S + (1-p_L)\, x_F - \bar{u} - c_L$. By assumption $\pi_H > \pi_L$ since $(p_H - p_L)(x_S - x_F) > c_H - c_L$. Hence, the firm always prefers inducing high effort over low effort provided that this generates positive profits. This is guaranteed by the condition: $\bar{u} \leq p_H x_S + (1-p_H)\, x_F - c\,(e_H)$. □

Note that the average wage paid to the worker is simply the reservation utility plus the cost of exerting high effort. The firm retains all remaining profits.

In the above equilibrium, the agent's effort was not observable. The principal could only condition his contract on the observable outcome, rather than observable effort. Since outcomes and effort are correlated a contract inducing high effort exists. What would the contract look like if the firm could observe effort? The firm can now condition the contract on both effort and outcomes, and does so to maximize profit. For each $e \in \{H, L\}$, we have:

$$\begin{aligned} \min_{w_{eS}, w_{eF}} \quad & \{p_e w_{eS} + (1-p_e)\, w_{eF}\} \\ \text{s.t.} \quad & p_e w_{eS} + (1-p_e)\, w_{eF} - c\,(e) \geq \bar{u} \end{aligned}$$

Clearly, the participation constraint must bind. Hence we have: $p_e w_{eS} + (1-p_e)\, w_{eF} = \bar{u} + c\,(e)$. Notice that state-dependent wages are not crucial in and of themselves — all that matters is that the worker is remunerated for his effort and reservation utility on average. It suffices to make wages constant (non-outcome contingent), with $w_e = \bar{u} + c\,(e)$.

On average, the worker is paid the same amount when effort is observable or unobservable. However, the worker's remuneration is stochastic in the unobservable case — and this creates the incentive to exert effort. The worker is forced to internalize some of the randomness that the firm faces in its profit stream. By contrast, in the observable case, the worker receives a constant wage, and the firm internalizes all the randomness in payoffs.

19.5 Continuum of Payoffs

We switch to the setting of the owner of a firm who is trying to incentivize his manager to exert appropriate effort. Continue to assume that the manager (formerly the worker) can choose between two effort levels — $e \in \{e_L, e_H\}$. Let $\pi \in [\underline{\pi}, \bar{\pi}]$ denote the payoff to the project. Suppose π is stochastically related to e, with $f(\pi|e)$ being the conditional probability density function of π, given e. Assume that the distribution of π conditional upon high effort e_H first-order stochastically dominates the distribution of π conditional upon low effort e_L. Formally,

$$F(\pi|e_H) \leq F(\pi|e_L)$$

for all $\pi \in [\underline{\pi}, \bar{\pi}]$, where $F(\cdot|e)$ is the conditional cumulative distribution function. In particular, this implies that $E[\pi|e = e_H] \geq E[\pi|e = e_L]$. (We proved in Section 2 that $E[X] > E[Y]$ if X first order stochastic dominates Y.)

The manager is an expected utility maximizer with utility $u(w, e) = v(w) - c(e)$ where $c(e_H) > c(e_L)$. We sometimes specialize to the case: $v(w) = w$ — which implies that the manager is risk neutral.

19.5.1 Optimal Contract with Effort

Suppose the owner's can observe the manager's effort. The owner offers the manager a contract that specifies the manager's effort $e \in \{e_L, e_H\}$ and his wage payment as a function of observed profits. We have:

$$\begin{array}{ll} \max_{e\in\{e_L,e_H\},w(\pi)} & \displaystyle\int (\pi - w(\pi)) f(\pi|e)\, d\pi \\ \text{s.t.} & \displaystyle\int v(w(\pi)) f(\pi|e)\, d\pi - c(e) \geq \bar{u} \end{array}$$

We solve this problem in two steps. First, find the optimal wage contract for each effort level e. Then, we determine which is the idea effort choice. Note that maximizing profits is equivalent to minimizing the expected value of the owner's compensation scheme. Forming the Lagrangian, with multiplier λ, we have:

$$\mathcal{L} = \int w(\pi) f(\pi|e)\, d\pi - \lambda \left[\int v(w(\pi)) f(\pi|e)\, d\pi - c(e) - \bar{u} \right]$$

The FOC is:

$$\begin{aligned} f(\pi|e) - \lambda v'(w) f(\pi|e) &= 0 \\ v'(w(\pi)) &= \frac{1}{\lambda} \end{aligned}$$

If the manager is strictly risk averse (so that $v'(w)$ is strictly decreasing in w), and immediate implication is that the optimal compensation scheme $w(\pi)$ is constant. (This is simply a risk sharing result — analogous to our results from the section on efficiency in asset markets.) Given that the contract explicitly dictate's the manager's effort choice, and there is no problem with providing incentives, the risk-neutral owner should fully insure the risk-averse manager against any risk in his income stream. By the constraint, we have:

$$\begin{aligned} v(w^*) - c(e) &= \bar{u} \\ w^* &= v^{-1}[\bar{u} + c(e)] \end{aligned}$$

which is exactly the result we obtained above. By contrast, if the manager is risk neutral (so that $v'(w) = 1$), then any compensation scheme that satsifies the constraint (i.e. that gives the agent an expected wage of $\bar{u} + c(e)$) is optimal.

Now, consider the optimal choice of e. The owner solves:

$$\max_{e \in \{e_L, e_H\}} \Pi(e) = \int \pi f(\pi|e)\, d\pi - v^{-1}(\bar{u} + c(e))$$

and the optimal e^* is found by a simple comparison of these.

19.5.2 Optimal Contract without Effort

Now suppose the owner cannot observe the manager's effort level. As above we saw in the simple example, to incentivize the manager to exert high effort, the owner must relate the manager's pay to the realization of profits — which results in a stochastic compensation scheme. To highlight this point (and focus on the work-incentives, unencumbered by the risk-sharing), we focus on the case of a risk neutral manager. In the principal agent problem with unobservable managerial effort and a risk-neutral manager, an optimal contract generates the same effort choice and expected utilities for the manager and owner as when effort is observable.

Proof. The idea of the proof is as follows: Hidden action imposes an additional constraint upon the owner — it is now subject to the manager's IC constraint as well. Clearly the owner cannot do better when effort is unobservable, than when it is observable. Hence, if we can find some compensation schedule that is as good for the principal as in the case with observable effort, then it must be optimal (for the principal).

Consider the compensation schedule: $w(\pi) = \pi - \alpha$, where α is some constant. Essentially, the owner "sells the firm" to the manager, and retains a fixed income stream α. If the manager accepts this contract, he will choose the effort level that maximizes provided that:

$$\begin{aligned} \int w(\pi) f(\pi|e)\, d\pi - c(e) &= \int \pi f(\pi|e)\, d\pi - \alpha - c(e) \\ &= \Pi(e) + \bar{u} - \alpha \end{aligned}$$

Since $\bar{u}$ and α are constant, then the effort level that maximizes this, exactly coincides with the effort level that maximizes the owner's profit. Hence, the same effort choice will be made. Further, the manager will accept the contract provided that:

$$\int \pi f(\pi|e)\, d\pi - \alpha - c(e) \geq \bar{u}$$

Let α^* be the α for which this expression holds with equality. Clearly, this is the largest α which is consistent with the agent accepting the contract. Moreover, since the owner's profit is α, this maximizes the firm's profit. We have:

$$\Pi = \int \pi f(\pi|e)\, d\pi - c(e) - \bar{u}$$

which is precisely the firm's profit in the case with observable effort. $\square$