ALLAMA IQBAL OPEN UNIVERSITY

USMAN NOTES
www.usmannotes.com

Name:

Registration No.:

Course Code: 801

Program: MSc. Economics

Semester: Autumn 2026

ALLAMA IQBAL OPEN UNIVERSITY, ISLAMABAD

Course: Introduction to Microeconomics (801) • Level: MSc. Economics

Semester: Autumn 2026

Assignment No. 1 (Units 1–5)

Total Marks: 100 • Pass Marks. 40

SOLVED ASSIGNMENT GUIDE

(For Study & Reference)

Note: This guide is prepared for study and reference purposes. Students are advised to read the relevant units of the course book, understand the concepts, and then write the assignment in their own words.

USMAN NOTES www.usmannotes.com

Question 1: What Are Economic Models? (20 Marks)

Q. 1 What are Economic Models? Explain.

Introduction

Economics studies a reality that is extraordinarily complicated. In a single market, millions of buyers and sellers, changing tastes, changing technology, government rules, weather, expectations and events in other countries all act at the same time. No human mind can follow all of them together, and no description that tried to include everything would explain anything. Economists therefore do what every science does: they simplify. They build an economic model — a deliberately simplified representation of a part of economic reality — which keeps the factors that matter most for the question in hand and sets the rest aside for the moment. Just as a map is useful precisely because it is not the territory itself, an economic model is useful precisely because it leaves things out. This answer explains what an economic model is, what it is made of, the main types of models used in microeconomics, how a model is

built and tested, and both the strengths and the limits of modelling, with examples from the theory of demand, supply and the firm.

Meaning of an Economic Model

An economic model is a simplified framework, expressed in words, diagrams, tables or mathematical equations, that shows how economic variables are related to one another and is used to explain past behaviour or to predict the effects of a change. A model is not a copy of reality; it is an argument in a compact form. It states: if these assumptions hold, and if this variable changes in this way while the others are held constant, then that result will follow. The famous phrase ceteris paribus — 'other things being equal' — is the working tool of every model, because it allows the economist to isolate one relationship at a time from the confusion of the real world.

A model therefore always contains two kinds of statements. The first are assumptions, which describe the setting: for example, that consumers aim to maximise satisfaction, that firms aim to maximise profit, that tastes and incomes are given, or that there are very many buyers and sellers. The second are conclusions or predictions, which are deduced logically from the assumptions: for example, that a rise in price, other things equal, will reduce the quantity demanded. The value of a model is judged not by whether its assumptions are literally true in every detail — they never are — but by whether its predictions are reasonably accurate and whether it illuminates the mechanism at work. A model that predicts well with simple assumptions is better than a complicated description that predicts nothing.

Elements of an Economic Model

Every model, whatever its form, is built from a small number of standard elements. First, there are variables, the magnitudes that the model talks about, such as price, quantity, income, cost and profit. Some variables are exogenous, determined outside the model and taken as given, such as the price of a raw material in a firm's cost model. Others are endogenous, determined inside the model, such as the quantity a firm chooses to produce. Secondly, there are parameters, the fixed numbers that describe behaviour or technology, such as the share of income spent on food or the responsiveness of demand to price. Thirdly, there are behavioural or functional relationships connecting the variables, often written as equations such as Qd = f(P, Y, Ps, T), which reads: quantity demanded depends on price, income, the prices of related goods and tastes. Fourthly, there is an equilibrium or solution concept, which states what it means for the model to 'settle' — for example, the price at which quantity demanded equals quantity supplied, so that no participant has a reason to change behaviour.

- Assumptions: the simplified conditions on which the model rests, for example rational behaviour and ceteris paribus.
- Variables: exogenous variables taken as given from outside, and endogenous variables explained within the model.
- Relationships: equations, schedules or curves that state how the variables affect one another.
- Predictions: the testable conclusions that follow, such as the direction in which price or quantity will move after a tax, a subsidy or a change in income.

Forms in Which Models Are Expressed

The same model can be expressed in four forms, and a good student should be able to move between them. The verbal form states the logic in words, as when we say that consumers buy more at lower prices because a lower price raises the purchasing power of income and makes the good cheaper relative to substitutes. The tabular form presents the relationship as a schedule, for example a demand schedule listing quantities demanded at successive prices. The geometrical form draws the relationship as a curve — the downward-sloping demand curve is a model in a picture — and it makes shifts, movements and equilibrium visible at a glance. The mathematical or algebraic form writes the relationship as equations, which is the most precise form and the one used at the MSc level, because equations allow exact solutions, comparisons and general proofs that words cannot deliver. None of these forms is 'the real model' by itself; the model is the underlying logic, and each form is a language for expressing it.

Types of Economic Models

Models are classified in several useful ways. By scope, a partial-equilibrium model studies one market in isolation, as Marshall's demand-and-supply model of a single commodity does, whereas a general-equilibrium model, in the tradition of Walras and Arrow–Debreu, studies all markets together and shows how they interact. By time, a static model compares equilibrium positions at a moment in time without explaining the path between them, a comparative-static model compares two equilibria before and after a change, and a dynamic model explicitly follows the adjustment over time, as in the cobweb model of agricultural prices. By behaviour, a positive model explains and predicts what is, without judging it, while a normative model evaluates what ought to be and is used in welfare economics. By construction, a theoretical

model is an abstract logical structure, while an empirical or econometric model estimates the relationships from real data so that the numbers, and not only the directions of change, can be predicted.

Table 1: Main types of economic models

Type	What it does	Example
Partial equilibrium	Studies one market holding other markets aside	Demand and supply for wheat
General equilibrium	Studies all markets and their interdependence together	Walrasian system of equations
Comparative static	Compares equilibrium before and after a change	Effect of a sales tax on price
Dynamic	Traces the path of adjustment over time	Cobweb model of farm prices
Econometric	Estimates relationships from observed data	Estimated demand function for tea

USMAN NOTES www.usmannotes.com

How a Model Is Built and Tested

Model-building follows a recognisable sequence. The economist first defines the problem precisely: what exactly is to be explained — a price, a quantity, a choice? Next, the essential variables are selected and simplifying assumptions are made explicit, because hidden assumptions are the commonest source of confusion in economics. Relationships between the variables are then specified, solution is derived — for example, by setting demand equal to supply or by maximising profit subject to a cost constraint — and predictions are drawn out. Finally, the predictions are confronted with evidence. If the facts repeatedly contradict the model, the assumptions are re-examined and the model is revised or replaced; if the facts broadly confirm it, the model earns a place in the economist's toolkit. This cycle of assumption,

deduction and testing is the scientific method applied to social life, and it is what separates a model from a mere opinion or a slogan.

Examples of Models in Microeconomics

Microeconomics is built from a family of interconnected models, several of which appear in this very assignment. The model of consumer behaviour uses preferences, indifference curves and a budget line to predict how demand responds to price and income changes, and the Hicksian theory of Question 2 is one refined version of it. The model of the firm uses a production function and cost curves to predict output and pricing decisions, as Questions 4 and 5 explore. The models of market structure — perfect competition, monopoly, monopolistic competition and oligopoly, which dominate Assignment No. 2 — are models in the full sense: each changes only the assumptions about the number of sellers, the nature of the product and the freedom of entry, and then deduces a different outcome for price, output and welfare. Recognising this shows why models matter: by changing one assumption at a time, the economist can see exactly which feature of reality is responsible for which result.

Importance and Limitations of Economic Models

Models are indispensable for at least four reasons. They bring clarity, forcing vague ideas into precise statements that can be checked. They bring economy of thought, allowing a trained economist to analyse a new situation quickly by recognising which standard model it resembles. They make prediction and policy possible: a government considering a subsidy, a minimum wage or a tax must rely on a model, explicit or implicit, of how buyers and sellers will respond. And they provide a common language through which economists in different countries can

communicate and criticise one another's reasoning. Without models, economics would be a heap of anecdotes about individual markets, with no general knowledge at all.

Their limitations must be stated with equal honesty. A model is only as good as its assumptions, and unrealistic assumptions — perfect information, no transport costs, perfectly rational agents — can mislead if the conclusions are applied mechanically to situations where the assumptions fail badly. Models can also create a false sense of precision when estimated numbers are presented without their margins of error, and they can omit things that are hard to measure, such as fairness, dignity or social norms, even though those things affect real behaviour. Finally, different models of the same phenomenon can give different advice, so the economist's judgement in choosing and adapting a model can never be replaced by the model itself. The correct attitude is neither worship nor dismissal: a model is a tool, to be used with knowledge of what it was built to do.

USMAN NOTES www.usmannotes.com

Conclusion

To conclude, an economic model is a simplified, logically consistent representation of an aspect of economic reality, built from assumptions, variables and relationships, and used to explain behaviour and predict the consequences of change. It may be expressed in words, tables, diagrams or equations, and it may be partial or general, static or dynamic, theoretical or empirical. Its worth lies not in photographic realism but in insight and predictive power. Mastering models — knowing what each one assumes, what it predicts and where it breaks down — is the central skill of the economist, and it underlies every other question in this course, from the Hicksian demand theory to the theory of the firm and the models of market structure.

Question 2: Hicksian Demand Theory (20 Marks)

Q. 2 Explain the Concept of Hicksian Demand Theory.

Introduction

The Hicksian demand theory, developed by Sir John Hicks, most fully in Value and Capital (1939) and in his later revision of demand theory, is one of the two great ordinal approaches to consumer behaviour, the other being the Slutsky approach and, in a different tradition, Samuelson's revealed preference theory. Its central problem is one that Marshall's cardinal utility theory had left unresolved: when the price of a good falls, the quantity demanded rises for two quite different reasons, and a satisfactory theory of demand must separate them. Hicks's achievement was to perform this separation using only ordinal preferences — indifference curves that rank bundles without measuring utility in numbers — and in doing so to produce a compensated demand function, now called the Hicksian demand function, which stands alongside the ordinary Marshallian demand function as a foundation of modern microeconomics and welfare economics.

The Problem: Splitting the Price Effect

When the price of a good falls, the total change in its quantity demanded is called the price effect. Hicks argued that this total effect is the sum of two components. The substitution effect is the change in demand that occurs because the good has become relatively cheaper than other goods: even a consumer whose real standard of living is held exactly constant will substitute the cheaper good for the now relatively dearer alternatives, so the substitution effect of a price fall on the quantity demanded is always positive. The income effect is the change in demand that occurs because the price fall has raised the consumer's real income:

with the same money income, the consumer can now afford bundles that were previously out of reach, and whether this leads him to buy more or less of the good depends on whether the good is normal or inferior. For a normal good the income effect reinforces the substitution effect; for an inferior good it works against it; and in the extreme Giffen case it outweighs it.

The Hicksian Compensation: Holding Utility Constant

The distinctive feature of Hicks's method is the way he isolates the substitution effect. He asks: after the price has fallen, by how much could we reduce the consumer's money income — 'compensate' him for the price change — so that he is left exactly as well off as he was before, neither better nor worse, resting on his original indifference curve? This adjustment is called the compensating variation in income. Once income has been reduced in this way, any remaining change in the quantity demanded cannot be due to a change in real welfare, because welfare, measured by the indifference curve, has been held constant; it can only be due to the change in relative prices. That remaining change is the pure Hicksian substitution effect. Removing the compensation — restoring the income that was taken away — then produces the income effect, as the consumer moves from the compensated position to the final equilibrium on a higher indifference curve.

Graphically, the analysis uses an indifference-curve diagram with the good in question on the horizontal axis. The consumer starts in equilibrium where the original budget line touches the highest attainable indifference curve. When the price falls, the budget line swings outward around its intercept on the other good's axis, and the consumer reaches a higher indifference curve at a new tangency: the movement between the two tangencies is the total price effect. Hicks then draws a compensated budget line, parallel to the new budget line — because it faces

the new prices — but shifted inward until it just touches the original indifference curve. The tangency on the original curve with the compensated line shows how much of the good the consumer would buy at the new prices if his satisfaction were held at its original level. The movement along the original indifference curve to this point is the substitution effect, and the subsequent movement to the final equilibrium, when compensation is withdrawn, is the income effect.

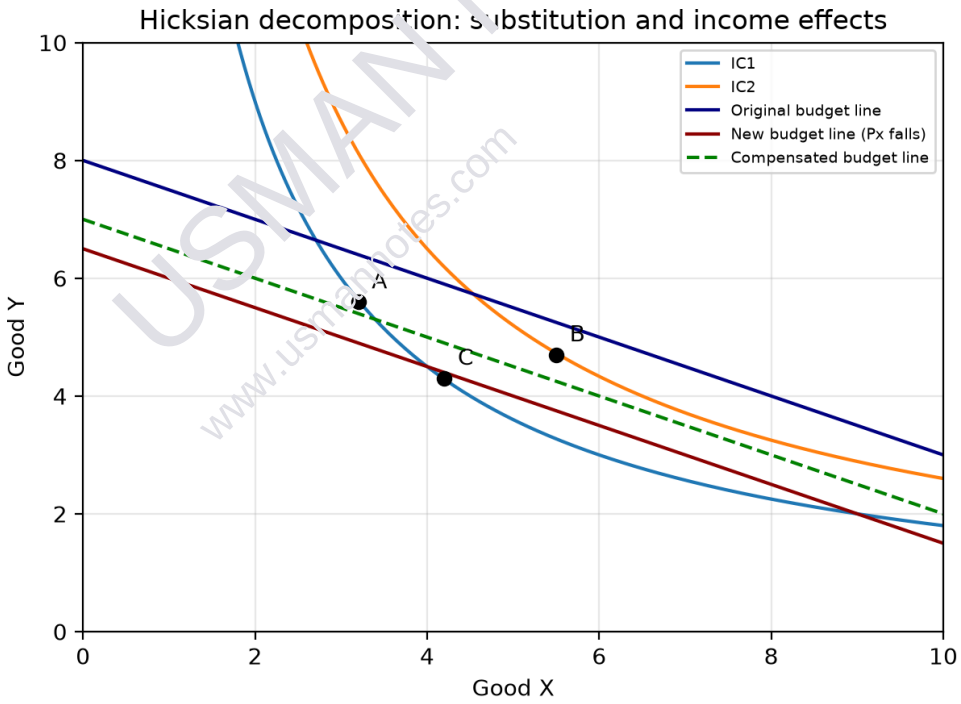


Figure 1: Hicksian decomposition — the movement from A to C along the original indifference curve is the substitution effect; the movement from C to B, when compensation is removed, is the income effect.

The Hicksian (Compensated) Demand Function

Repeating this thought experiment for every possible price produces the Hicksian demand function, written h(P, U), which gives the quantity demanded as a function of prices and a fixed level of utility, rather than of prices and a fixed money income as in the ordinary Marshallian function d(P, M). The Hicksian function answers the question: what is the cheapest way of

reaching a given level of satisfaction at given prices, and how much of each good does that cheapest bundle contain? It is therefore derived from the expenditure-minimisation problem — minimise total spending subject to reaching the target utility — just as the Marshallian function is derived from utility maximisation subject to the budget. The two problems are mirror images, known as the duality of consumer theory, and the expenditure function that records minimum spending is the bridge between them.

Because utility is held constant along a Hicksian demand curve, it contains the substitution effect only, and so it must always slope downward: there are no Giffen exceptions to the Hicksian law of demand. This is a major theoretical virtue. The ordinary Marshallian curve, by contrast, mixes both effects and can, in the Giffen case, slope upward. The Hicksian curve is also the theoretically correct basis for measuring changes in consumer welfare, because the compensating and equivalent variations that welfare economics uses are computed from it. Its practical weakness is equally clear: utility is not directly observable, so a Hicksian curve cannot simply be read off market data in the way an ordinary demand curve can be estimated; it is primarily a tool of theory and welfare measurement rather than of direct statistical estimation.

Table 2: Marshallian and Hicksian demand compared

Feature	Marshallian demand	Hicksian demand
Held constant	Money income	Utility (real satisfaction)
Derived from	Utility maximisation subject to budget	Expenditure minimisation subject to utility target
Contains	Substitution + income effects	Substitution effect only
Slope	Usually downward; Giffen exception possible	Always downward
Main use	Market demand, forecasting,	Welfare measurement, theory

	estimation	

Normal, Inferior and Giffen Goods in the Hicksian Analysis

The Hicksian decomposition classifies goods with precision. For a normal good, a price fall gives a positive substitution effect and a positive income effect, so quantity demanded unambiguously rises and the ordinary demand curve slopes downward. For an inferior good, the income effect is negative: as real income rises after the price fall, the consumer actually buys less of the inferior good, shifting part of his spending to better alternatives. Provided the substitution effect is stronger, which is the usual case, quantity demanded still rises on balance and the demand curve still slopes downward, but less steeply than it would without the offsetting income effect. For a Giffen good — a strongly inferior staple that absorbs a large share of a poor household's budget, the classic example being a basic staple food for very poor consumers — the negative income effect is stronger than the substitution effect, so a price fall actually reduces the quantity demanded and a price rise increases it. Hicks's framework shows exactly why this paradox is possible and why it is rare: it requires both strong inferiority and a large budget share.

Merits and Criticisms of Hicks's Approach

The merits of the Hicksian theory are substantial. It is based on ordinal utility, which is a far weaker and more realistic assumption than Marshall's measurable, cardinal utility: consumers need only rank bundles, not attach numbers to satisfaction. It explains the law of demand more completely than Marshall, because it shows the two forces of which the price effect is composed instead of treating it as a single unexplained response. It resolves the Giffen paradox

within a single consistent framework, and it provides the rigorous foundation for modern welfare measures and for the Slutsky equation, which expresses the same decomposition in algebraic form. For these reasons the Hicksian apparatus, together with indifference-curve analysis, became the standard language of advanced consumer theory.

The criticisms should also be recorded. The compensating variation is a hypothetical construction: no actual consumer's income is adjusted when prices change, so the substitution effect Hicks isolates is not directly observed in behaviour, whereas Slutsky's alternative compensation — allowing the consumer still to afford the original bundle — is closer to observable data. Indifference curves themselves can be criticised as introspective constructions that are difficult to establish empirically for real consumers, a difficulty that led Samuelson to build the revealed preference theory on choices alone. Hicks himself, in later work, acknowledged weaknesses concerning the treatment of goods that must be bought in discrete units and cases where indifference curves have unusual shapes. These criticisms qualify the theory; they do not remove its central insight, which remains a permanent part of economics.

Conclusion

In conclusion, Hicksian demand theory explains consumer behaviour by holding utility constant through a compensating change in income, thereby separating the price effect into an always-positive substitution effect and an income effect whose sign depends on the nature of the good. From this construction flows the compensated Hicksian demand function, which always slopes downward, underpins welfare measurement through compensating and equivalent variation, and clarifies the cases of normal, inferior and Giffen goods within one ordinal framework. Its hypothetical compensation makes it less directly observable than the

Marshallian function, but as an analytical achievement — showing why demand behaves as it does, and not merely that it does — it is one of the landmarks of twentieth-century microeconomics.

Question 3: Determinants of Market Demand (20 Marks)

Q. 3 Explain the determinants of Market Demand in detail.

Introduction

Market demand is the total quantity of a good or service that all the buyers in a market, taken together, are willing and able to purchase at each possible price during a given period of time. It is obtained by adding together, horizontally, the demand of every individual buyer at each price. Two words in the definition deserve emphasis: willing and able. A desire unsupported by purchasing power is not demand in the economic sense, and purchasing power without willingness is not demand either. The law of demand tells us how quantity demanded responds to the good's own price, but price is only one influence among many. This answer explains each determinant of market demand in detail — the price of the good, consumer incomes, the prices of related goods, tastes and preferences, expectations, the number and composition of buyers, and several further factors such as credit, distribution of income and government policy — and distinguishes carefully between movements along a demand curve and shifts of the whole curve.

1. Price of the Good Itself

The most immediate determinant is the price of the good itself. The law of demand states that, other things equal, the quantity demanded of a good rises when its price falls and falls when its

price rises, producing the familiar downward-sloping demand curve. Two mechanisms explain this relationship: the substitution effect, because a lower price makes the good cheaper relative to alternatives and attracts buyers toward it, and the income effect, because a lower price raises the real purchasing power of every buyer's income. Changes in the good's own price cause a movement along the existing demand curve — an extension of demand when price falls and a contraction when it rises — rather than a shift of the curve. Every other determinant discussed below works differently: when one of them changes, demand changes at every price, and the entire curve shifts to the right for an increase in demand or to the left for a decrease.

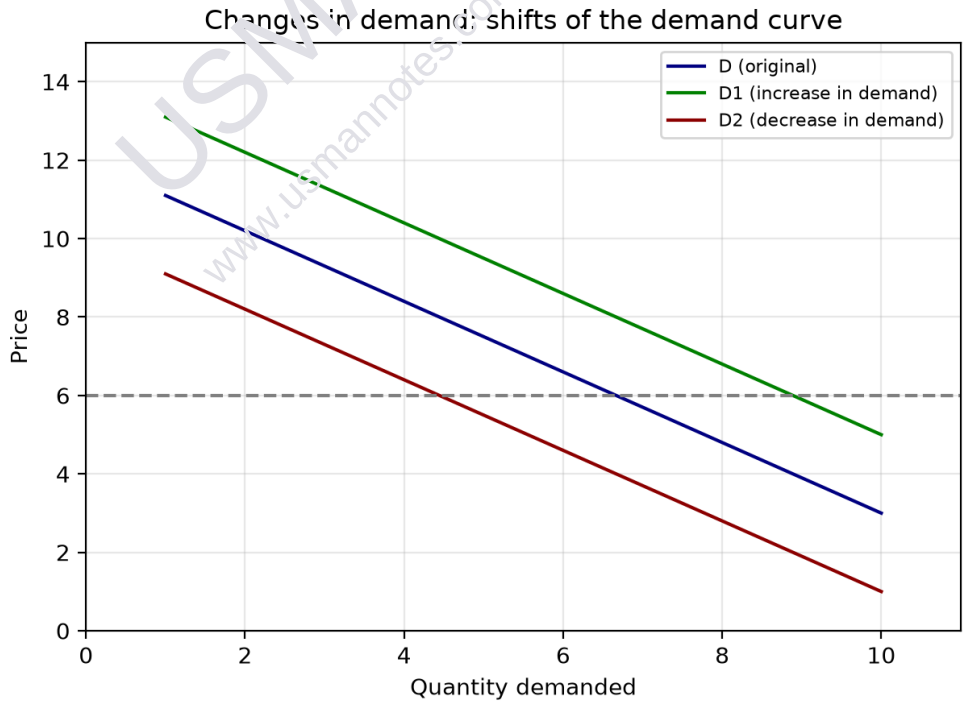


Figure 2: A change in the good's own price moves the buyer along a curve; a change in any other determinant shifts the whole curve — to D1 for an increase in demand, to D2 for a decrease.

2. Income of Consumers

The level and change of consumers' income is one of the most powerful determinants of market demand. For a normal good, demand rises as income rises: families whose incomes

improve buy better food, larger houses, more travel and more education, so the demand curve shifts to the right. For an inferior good, demand falls as income rises, because consumers switch away from the cheaper, lower-quality alternative they previously had to accept: demand for the cheapest unbranded staples or for second-hand goods often behaves in this way. The sensitivity of demand to income is measured by the income elasticity of demand, and it differs greatly between goods. Necessities such as basic food grains have low income elasticity — demand grows more slowly than income, a regularity known as Engel's law — while luxuries such as cars, air travel and restaurant meals have high income elasticity, which is why their markets expand rapidly in a growing economy and contract sharply in a recession. Expected future income matters as well as current income: households confident about their future earnings spend and borrow more freely today.

3. Prices of Related Goods: Substitutes and Complements

Demand for a good depends on the prices of goods related to it, in two opposite ways. Substitute goods are rivals in consumption — tea and coffee, butter and margarine, one brand of mobile phone and another. When the price of a substitute rises, buyers switch toward the good in question, so its market demand increases and its curve shifts right; when the price of a substitute falls, demand for the good falls. Complementary goods, by contrast, are used together — cars and petrol, printers and ink cartridges, bread and butter in some uses. When the price of a complement rises, the joint activity becomes more expensive, so demand for the good falls; when the price of a complement falls, demand for the good rises. The strength of these cross-effects, measured by the cross elasticity of demand, tells firms and policymakers how closely goods compete: a high positive cross elasticity marks close substitutes, which is

essential information for competition policy and for any firm deciding whether a rival's price cut must be matched.

4. Tastes, Preferences and Social Influences

Behind every demand curve lie the tastes and preferences of buyers, and when tastes change, demand changes at every price. Tastes are shaped by habit and culture, by advertising and marketing, by fashion, by health and environmental concerns, and by the behaviour of other consumers. A successful advertising campaign, a celebrity endorsement or a social-media trend shifts demand to the right without any change in price or income; a health scare about a food product, or growing concern about sugar and plastic packaging, shifts it sharply to the left. Two social effects deserve special mention. The bandwagon effect makes people demand a good partly because others are buying it, as with fashionable clothing or popular phone models, while the snob and Veblen effects make exclusivity itself attractive, so that some luxury goods are demanded more because their high price signals status. Tastes also change systematically with age, education and urbanisation, which is why long-run market demand for services, private education and processed food grows as societies develop.

5. Expectations About the Future

Buyers are forward-looking, so expectations about future prices, incomes and availability shift current demand. If consumers expect the price of a good to rise soon — because a tax increase has been announced, a harvest has failed or a currency depreciation is expected to raise import prices — they buy more now, stocking up while the good is cheaper, and current demand shifts to the right. If they expect prices to fall, as buyers often do before a new model of a phone or car is launched, they postpone purchases and current demand shifts left. Expectations about

future income work similarly: fear of job losses makes households cut discretionary spending and increase precautionary saving, reducing demand across many markets at once, while confidence about the future raises demand, especially for goods bought on instalments. In asset and housing markets, expectations can become self-fulfilling for a time, as expected price rises draw in speculative demand that pushes prices up further.

6. Number of Buyers and the Composition of the Population

Because market demand is the sum of individual demands, anything that changes the number of buyers changes market demand directly. Population growth increases the demand for virtually all goods, especially necessities, housing, education and health care. Migration and urbanisation concentrate that demand geographically and change its character, raising demand for urban services and rented housing. Equally important is the composition of the population: an ageing population increases the demand for medicines, health care and retirement services while reducing demand for toys and primary schooling, whereas a youthful population does the opposite. Changes in household structure matter too — a rise in the number of households, through smaller family size or more young people living independently, increases demand for housing, furniture and appliances even if total population is unchanged. This is why serious market forecasting begins with demography, not with price.

7. Distribution of Income, Credit and Other Factors

Several further determinants complete the picture. The distribution of income matters because different income groups spend on different goods: a transfer of income toward poorer households, whose spending is concentrated on necessities, raises the demand for basic food and clothing and lowers the relative demand for luxuries, even when total national income is

unchanged. The availability and cost of consumer credit is another major factor: cheaper loans and instalment facilities raise the demand for durable goods such as cars, motorcycles and appliances, while tight credit and high interest rates depress it, as monetary policy intends. Season and weather shift demand predictably — for cold drinks, fans, warm clothing and ice cream — and festivals and occasions create temporary surges. Government policy acts as a determinant as well: taxes and subsidies change effective prices, but public health campaigns, bans, quality standards and public provision of a service also shift private demand directly. Finally, the prices of assets and the general price level influence demand through wealth effects and through households' attempts to protect the real value of their money.

Table 2. Summary of the determinants of market demand

Determinant	Change	Effect on demand for the good
Own price	Price falls	Quantity demanded rises (movement along the curve)
Income	Income rises	Rises for a normal good; falls for an inferior good
Price of a substitute	Substitute becomes dearer	Demand rises (curve shifts right)
Price of a complement	Complement becomes dearer	Demand falls (curve shifts left)
Tastes	Good becomes more fashionable / healthier image	Demand rises
Expectations	Future price expected to rise	Current demand rises
Number of buyers	Population / households increase	Market demand rises

Conclusion

To conclude, market demand is determined jointly by the price of the good itself, which produces movements along the demand curve, and by a set of background determinants —

incomes, the prices of substitutes and complements, tastes and social influences, expectations, the number and composition of buyers, the distribution of income, credit conditions, season and government policy — each of which shifts the entire curve. A complete analysis of any real market must work through these determinants systematically rather than attributing every change in sales to price alone. Firms use this understanding for forecasting and pricing, and governments use it when predicting the effects of taxes, subsidies and macroeconomic policy on particular industries. The demand curve, properly understood, is not a fixed line but a summary of a changing economic and social situation.

Question 4: Deriving the Cost Function from the Production Function (20 Marks)

Q. 4 Derive the Cost Function from the Production Function.

Introduction

The production function and the cost function are two views of the same underlying technology, seen from opposite sides. The production function looks at the firm from the engineering side: it states the maximum output that can be produced from each combination of inputs. The cost function looks at the firm from the financial side: it states the minimum total cost of producing each level of output, given the prices of inputs. The second is derived from the first by adding one piece of information — input prices — and one piece of behaviour — cost minimisation. This answer carries out that derivation step by step: it defines both functions, shows how the least-cost input combination is found for a given output, converts the resulting expansion path into a total cost curve, works a complete numerical example with a

Cobb–Douglas production function, and then draws the average and marginal cost curves from the total cost function, noting the duality between production and cost throughout.

The Two Starting Points: Production Function and Input Prices

Write the production function in general form as Q = f(L, K), where Q is output, L is labour and K is capital, and f embodies the best available technology. The function tells us what is technically possible, but by itself it says nothing about cost, because the same output can usually be produced by many different input combinations — more labour with less capital, or more capital with less labour — and these combinations do not cost the same. To choose among them we need the prices of the inputs: the wage rate w for labour and the rental price of capital r. With those prices, the cost of any input combination is given by the isocost equation C = wL + rK, which is a straight line in input space whose slope is the ratio of input prices, −w/r. The firm's problem is now well defined: for each output level, find the point on the corresponding isoquant that lies on the lowest possible isocost line. That is the whole logic of the derivation; everything that follows is its execution.

Step One: Cost Minimisation for a Given Output

Formally, the firm minimises C = wL + rK subject to the constraint f(L, K) = Q, for the target output Q. Using the method of Lagrange, the first-order conditions for an interior optimum require the marginal rate of technical substitution between the inputs to equal the ratio of their prices: MRTS = MPL / MPK = w / r. In words, the firm should employ inputs so that the extra output obtained from the last rupee spent on labour equals the extra output from the last rupee spent on capital; if one rupee bought more output on labour than on capital, total cost could be reduced by switching spending toward labour until equality is restored. Solving the

tangency condition together with the production constraint gives the conditional (or derived) factor demands: the quantities of labour and capital the cost-minimising firm will employ, expressed as functions of output and input prices, L*(Q, w, r) and K*(Q, w, r). They are called conditional because they are conditional on the output to be produced; unlike ordinary factor demands, they already embody the minimisation.

Step Two: From the Expansion Path to the Total Cost Function

As the target output rises, the cost-minimising tangency moves from isoquant to isoquant, and the line joining all these tangency points is the firm's expansion path, which shows how the optimal input combination grows with output at constant input prices. Now substitute the conditional factor demands back into the cost equation: C(Q) = w·L*(Q, w, r) + r·K*(Q, w, r). The result is the cost function: minimum total cost expressed as a function of output (for given input prices and technology). Adding fixed costs — the costs of inputs that do not vary with output in the short run — gives the short-run total cost function, TC = FC + VC(Q). From TC, the whole family of unit cost curves follows by simple operations: average cost AC = TC/Q, average variable cost AVC = VC/Q, and marginal cost MC = dTC/dQ, the increase in total cost from one more unit of output. Nothing new has been assumed at this stage; the cost curves are simply the production function read through the lens of input prices and minimising behaviour, which is why cost curves have the shapes they do — the law of diminishing returns in production reappears, mirrored, as rising marginal cost.

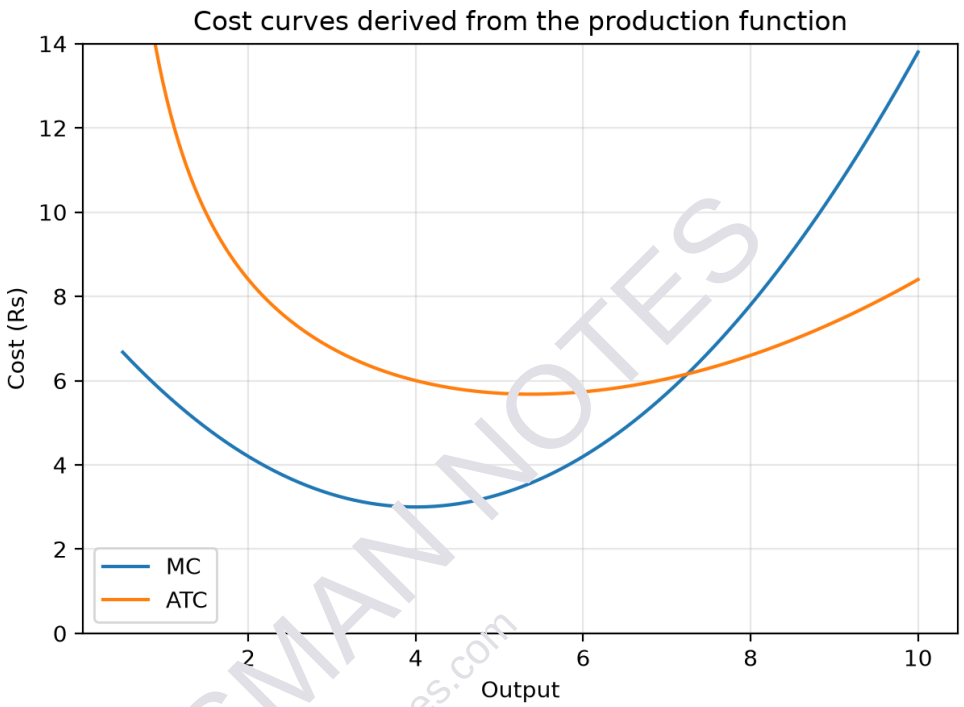


Figure 3: The marginal and average cost curves obtained from the total cost function derived in the worked example below.

A Worked Example: Cobb–Douglas Production Function

Let the production function be Q = L^0.5 K^0.5, a Cobb–Douglas function with constant returns to scale, and let input prices be w = Rs 25 per unit of labour and r = Rs 100 per unit of capital. The marginal products are MPL = 0.5 L^−0.5 K^0.5 and MPK = 0.5 L^0.5 K^−0.5, so the MRTS, MPL/MPK, equals K/L. Cost minimisation requires MRTS = w/r, that is K/L = 25/100 = 1/4, so the expansion path is K = L/4: the firm always uses four units of labour for every unit of capital, because labour is the relatively cheap input. Substituting into the production function, Q = L^0.5 (L/4)^0.5 = L/2, so the conditional factor demands are L* = 2Q and K* = Q/2. Total cost is therefore C = wL* + rK* = 25(2Q) + 100(Q/2) = 50Q + 50Q = 100Q. The cost function is TC = 100Q: with constant returns to scale and constant input prices, total cost is proportional to output, average cost equals marginal cost at Rs 100 per unit at every output, and the AC and

MC curves coincide in a horizontal line. If fixed costs of, say, Rs 1,000 must also be paid, the short-run function becomes TC = 1,000 + 100Q, so AC = 1,000/Q + 100 falls steadily toward Rs 100 while MC remains Rs 100 — exactly the familiar behaviour of a firm with spare capacity and constant marginal cost.

Table 4: Cost schedule from TC = 1,000 + 100Q

Output Q	Fixed cost	Variable cost (100Q)	Total cost	Average cost	Marginal cost
0	1,000	0	1,000	—	—
10	1,000	1,000	2,000	200	100
20	1,000	2,000	3,000	150	100
50	1,000	5,000	6,000	120	100
100	1,000	10,000	11,000	110	100

USMAN NOTES www.usmannotes.com

The example also shows how changes in input prices rotate the expansion path and shift the cost function. If the wage rose to Rs 100, equal to the rental rate, cost minimisation would require K = L, the conditional demands would become L* = K* = Q, and total cost would rise to TC = 200Q: the same technology now produces every output at double the unit cost. This is the sense in which the cost function 'contains' the production function plus prices — a dearer input both raises cost and induces substitution toward the cheaper input along each isoquant, and the derived cost function records the final, optimised result of that substitution.

Short Run, Long Run and Duality

The derivation differs between the short run and the long run in one respect: which inputs the firm is free to vary. In the long run all inputs are variable, the firm can always reach the expansion path, and the long-run cost function is the lower envelope of the possibilities. In the

short run capital is fixed at some level $\bar{K}$, so the firm can only vary labour: the required labour is found by inverting the short-run production function $Q = f(L, \bar{K})$, and the short-run cost function, $TC = wL(Q) + r\bar{K}$, necessarily gives costs at least as high as the long-run function for the same output, because a constrained minimisation can never beat an unconstrained one. This relationship — long-run average cost as the envelope of short-run average cost curves — is a direct consequence of deriving both from the same production function. Economists call the whole correspondence duality: for well-behaved technologies, the cost function contains all the information in the production function, and properties such as returns to scale translate exactly — increasing returns appear as falling long-run average cost, constant returns as constant average cost, and decreasing returns as rising average cost.

Conclusion

To conclude, the cost function is derived from the production function in three steps: take input prices from the market, find for each output the least-cost input combination where the marginal rate of technical substitution equals the input-price ratio, and substitute those conditional factor demands into the cost equation $C = wL + rK$. The total cost function that results generates the average, variable and marginal cost curves used throughout the theory of the firm, and the worked Cobb–Douglas example, in which $Q = L^{0.5} K^{0.5}$ with $w = 25$ and $r = 100$ yields $TC = 100Q$, shows the derivation producing a complete, internally consistent cost schedule. Production theory and cost theory are therefore not two subjects but one: the cost function is the production function expressed in money, after the firm has done its minimising work.

Question 5: The Concept of Variable Costs (20 Marks)

Q. 5 Explain the concept of Variable Costs in detail.

Introduction

Every firm's total cost divides into two great classes according to how the cost behaves when output changes. Fixed costs are those that must be paid even if nothing is produced, and do not change with the level of output in the short run. Variable costs are the opposite class: costs that rise and fall with the quantity produced. Understanding variable costs is not a bookkeeping nicety; it is the key to the firm's most important short-run decisions — whether to produce at all, how much to produce, whether to accept an extra order, and at what price a product can profitably be sold. This answer defines variable costs precisely, gives their main examples, explains how total, average and marginal variable cost behave and why, distinguishes the short-run from the long-run view, connects variable costs with the laws of returns, and shows their role in pricing, break-even analysis and the shutdown decision.

Meaning of Variable Costs

Variable costs are those costs which vary directly with the level of output: they increase as more is produced, decrease as less is produced, and are zero, or close to zero, when output is zero. They are also called prime costs or direct costs in some accounting traditions, because they attach directly to the act of producing each unit. The defining test is behaviour, not the name of the item: a cost is variable if, and to the extent that, it changes when output changes within the period under consideration. Total variable cost, TVC, is therefore a function of output, TVC = TVC(Q), and it is drawn as a curve rising from the origin. Two related concepts are built on it. Average variable cost, AVC = TVC/Q, is the variable cost of a typical unit, and

marginal cost is dominated in the short run by variable cost, since MC = ΔTVC/ΔQ when fixed costs do not change. Together with fixed costs, variable costs make up total cost: TC = TFC + TVC, an identity that underlies the whole short-run theory of the firm.

Examples of Variable Costs

The clearest examples come from physical production. Raw materials and components are the classic variable cost: a furniture maker uses twice as much wood for twenty chairs as for ten, and a bakery's flour, sugar and yeast bills follow its output of bread. Direct labour is variable when workers are hired by the hour, by the day or by the piece, as in much of Pakistan's garment and stitching industry where payment is per piece produced. Power and fuel used in production — electricity for machines, gas for furnaces, diesel for delivery vehicles per trip — vary with the running of the plant, unlike the fixed monthly connection charges. Packaging materials, consumable tools, sales commission paid per unit sold, freight charged per shipment and royalties paid per unit are also variable. In services the same logic applies: a clinic's medicines and disposable supplies per patient, a restaurant's ingredients per meal and a transport company's fuel per kilometre are variable costs, whereas rent, insurance, permanent salaries and licence fees are fixed — at least within the short run.

Table 5: Variable and fixed costs compared

Feature	Variable costs	Fixed costs
Behaviour with output	Rise and fall with output	Unchanged over a range of output
At zero output	Zero or near zero	Still incurred in full
Examples	Raw materials, piece-rate wages, fuel, packaging	Rent, insurance, permanent salaries, depreciation
Per-unit behaviour	AVC first falls, then rises (U-	AFC falls continuously as output

	shaped)	spreads it
Relevance	Decide whether and how much to produce in the short run	Recovered only in the long run / at break-even

Behaviour of Total and Average Variable Cost: The Laws of Returns

Total variable cost does not rise at a constant rate, and the pattern of its rise is dictated by the laws of returns operating through the production function. In the first stage of production, as the variable input is added to the fixed plant, the law of increasing returns (or diminishing costs) operates: each additional worker or unit of material is used more efficiently because specialisation becomes possible and the fixed equipment is better utilised, so total variable cost rises less than proportionately with output, and average variable cost falls. In the second stage, the law of constant returns gives a range where TVC rises roughly proportionately and AVC is at or near its minimum. In the third stage, the law of diminishing returns takes command: the fixed plant becomes crowded, additional variable inputs add less and less extra output, more input — and therefore more cost — is needed for each additional unit, TVC rises more than proportionately, and AVC turns upward. This is why the TVC curve has its characteristic inverse-S shape — first concave, then convex — and why the AVC curve is U-shaped, with marginal cost cutting it at its minimum point.

A simple schedule makes the behaviour concrete. Suppose a small workshop has the variable costs shown in a typical textbook pattern: producing the first few units, average variable cost falls from, say, Rs 120 to Rs 90 as workers specialise and machines stop standing idle; it reaches its minimum around the output where the plant is used at its designed intensity; beyond that, overtime wages, machine breakdowns, congestion and wastage push AVC up to Rs 110, Rs 140

and beyond. Marginal cost follows the same story even more sharply, falling first and then rising steeply. The managerial lesson is direct: the cheapest unit cost is not obtained by producing as little as possible, nor as much as possible, but at the output where the fixed plant and the variable inputs are combined in their best proportions.

Short Run and Long Run: All Costs Become Variable

The distinction between fixed and variable costs is a short-run distinction, and its time dimension must be understood. In the long run — a period long enough for the firm to change its plant size, renegotiate its leases, replace its machinery and enter or leave the industry — there are no fixed costs at all: every cost is variable, because every input can be adjusted. Rent is fixed only while the lease runs; salaries are fixed only while contracts run; even the factory building can be sold or extended. Some costs are also only partly variable in practice, called semi-variable or mixed costs: an electricity bill with a fixed monthly charge plus a per-unit rate, or a maintenance contract plus per-repair charges, must be split into its fixed and variable components before analysis. Costs can also rise in steps — one supervisor handles up to twenty workers, the twenty-first requires a second supervisor — so that a 'fixed' cost jumps at certain output levels. These refinements do not weaken the concept; they remind the analyst to define the period and the relevant range of output before classifying any cost.

Role of Variable Costs in Business Decisions

Variable costs govern the firm's key short-run decisions. The shutdown decision is the most famous: in the short run, a firm should continue producing only if its revenue covers its variable costs, that is, if price is at least equal to minimum average variable cost. If price falls below that point, every unit produced adds more to cost than to revenue, and the firm loses less by

shutting down and bearing only its fixed costs; if price is above AVC, producing contributes something toward fixed costs and is better than closing. Pricing and order decisions use the same logic: an extra export order or a special bulk order can profitably be accepted at any price above its marginal (variable) cost, provided normal sales are not disturbed, even if that price is below full average cost. Break-even analysis is built entirely on the split: the break-even output is Q = TFC / (P − AVC), showing how the contribution of each unit — price minus variable cost — accumulates to cover fixed costs and then to create profit. Cost control also focuses on variable costs, because they are the costs that respond immediately to efficiency in material use, labour productivity and wastage.

Conclusion

In conclusion, variable costs are those costs that vary with the level of output — raw materials, direct and piece-rate labour, production fuel and power, packaging and commission being the standard examples — and they are summarised by the total, average and marginal variable cost functions derived from the production process. Their pattern of behaviour, first falling and then rising per unit, reflects the laws of increasing and diminishing returns, and in the long run the fixed–variable distinction disappears as all inputs become adjustable. Above all, variable costs are decision costs: they determine the shutdown point, the profitability of extra orders, the break-even output and the contribution of each unit sold. A firm that knows its variable costs accurately knows the floor below which production destroys value; a firm that does not is pricing and producing in the dark.

USMAN NOTES www.usmannotes.com