ALLAMA IQBAL OPEN UNIVERSITY

AIOU LIBRARY

Name:

Registration No.:

Course Code: 806

Program: MSc Economics

Semester: Autumn 2026

ALLAMA IQBAL OPEN UNIVERSITY, ISLAMABAD

Course: Advanced Macroeconomics (806)

Level: MSc Economics • Credit Hours: 03

Semester: Autumn 2026

Assignment No. 2 (Units 6–9) • Total Marks: 100 • Passing 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 and write the assignment in their own words.

Question 1: The Accelerator Model of Inventory Investment and the Flexible Accelerator Model of Capital Accumulation (20 Marks)

What is the relationship between the accelerator model of inventory investment and the flexible accelerator model of capital accumulation?

Introduction

Investment is the most volatile component of aggregate demand, and a large part of its volatility comes from a simple piece of arithmetic: firms hold stocks — of finished goods, materials and work in progress on the one hand, and of plant and machinery on the other — in some proportion to the output they expect to sell. When sales change, the desired stocks change with them, and the investment needed to move from the old stock to the new one can be many times larger than the change in sales itself. This is the accelerator principle. It appears in two closely related forms: the accelerator model of inventory investment, developed in the work of Lloyd Metzler, and the flexible accelerator model of capital accumulation, associated with Chenery, Goodwin and Koyck. This answer explains each model, and then shows that their relationship is one of family resemblance with an important twist: both are 'desired-stock'

theories driven by output, but the flexible accelerator adds the realistic feature of gradual adjustment, and the inventory model is best understood as the fast-adjusting special case of the same general stock-adjustment framework.

The Common Core: The Accelerator Principle

Both models begin by defining a desired (or optimal) stock proportional to output. For fixed capital, K* = vY, where v is the desired capital–output ratio. For inventories, N* = βY (or N* = β times expected sales), where β is the desired inventory–output (inventory–sales) ratio. Given the desired stock, investment is the change in the actual stock needed to reach it. In the naive or rigid accelerator, adjustment is assumed complete within the period, so net investment equals the change in the desired stock: I = vΔY for capital, and inventory investment I_n = βΔY for inventories. Two consequences follow immediately. Investment depends on the change in output, not its level, so a mere slowing of sales growth reduces investment even while sales are still rising; and because v and β multiply the change, investment fluctuates far more violently than output — the central fact about investment that both models exist to explain.

The Accelerator Model of Inventory Investment

Inventories are held because production and sales cannot be perfectly synchronised: firms produce in batches, must offer goods for immediate delivery, and need buffer stocks of materials against interruptions. Metzler's model assumes firms aim to end each period with inventories equal to a fixed fraction β of expected sales. Production in a period therefore equals expected sales plus the intended addition to inventories: firms produce for the market and for the shelf. The dynamics come from forecast errors. If sales turn out higher than expected, inventories are unexpectedly depleted; next period firms raise production both to serve

expected sales and to rebuild the stock, so output responds more than sales. If sales disappoint, unwanted stocks pile up, and firms cut production by more than the fall in sales to work the excess off. Inventory investment is thus procyclical and amplifying: it lengthens booms and deepens slumps, and in Metzler's hands it generates endogenous inventory cycles of a few years' duration. Empirically the model captures the striking fact that changes in inventories, though a small part of GDP, account for a large share of the variation in GDP during recessions.

The Flexible Accelerator Model of Capital Accumulation

Applied literally to fixed capital, the naive accelerator fails: it predicts implausibly violent swings in orders for plant and equipment, assumes firms can install any amount of capital instantly, and ignores spare capacity. The flexible accelerator repairs this by assuming partial adjustment. The firm closes only a fraction λ ($0 < \lambda < 1$) of the gap between desired and actual capital in each period:

$$I_t = \lambda(K^*_t - K_{t-1}) = \lambda(vY_t - K_{t-1}).$$

The reasons λ is less than one are concrete: delivery and construction lags for capital goods, rising costs of installing capital very quickly (adjustment costs), the irreversibility of investment which makes firms cautious, financing constraints, and the desire to be sure a rise in demand is permanent before committing to capacity. Substituting repeatedly backwards turns current investment into a distributed lag of past outputs: investment responds to a weighted average of the history of output changes, with weights declining geometrically (Koyck lags). The flexible accelerator therefore keeps the accelerator's engine — desired capital proportional to output — but adds a gearbox: demand changes feed into capital spending gradually and persistently, which matches the observed smooth, drawn-out pattern of fixed investment booms.

The Relationship Between the Two Models

The relationship can now be stated precisely. (1) Same structure: both are stock-adjustment models in which a desired stock is anchored to output by a technical or behavioural coefficient (v for capital, β for inventories) and investment is the flow that moves the actual stock towards the desired one. Inventory investment and fixed investment are not separate theories but two applications of one principle to two kinds of stock. (2) The flexible accelerator is the generalisation: the inventory model in its simple Metzler form assumes adjustment is completed within the period (λ close to 1), which is plausible for stocks that can be run down instantly by selling from the shelf or rebuilt by stepping up production. The flexible accelerator allows λ < 1, which is necessary for lumpy, slowly delivered fixed capital. Indeed, later versions of the inventory model itself (for example with production-smoothing motives and partial adjustment of stocks) adopt exactly the same flexible form: I_n = λ(βY − N_{t−1}). Seen this way, the flexible accelerator is the general stock-adjustment equation, the naive capital accelerator is its λ = 1 case, and the inventory accelerator sits near the high-λ end of the same family. (3) Different cycle roles follow from the different speeds: fast inventory adjustment creates short, sharp inventory cycles and amplifies turning points, while slow capital adjustment creates the longer, smoother investment cycles associated with booms in capacity. (4) Both models share the same limitations: the desired-stock ratios v and β are treated as constants, ignoring interest rates, relative prices, expectations and capacity utilisation, which later theories (neoclassical and q-theory for capital; buffer-stock and (S,s) models for inventories) set out to repair.

Table 1: The two accelerator models compared

Feature	Inventory accelerator	Flexible accelerator (capital)

Stock concerned	Inventories of goods and materials (N)	Fixed capital stock (K)
Desired stock	$N^* = \beta \times$ expected sales/output	$K^* = v \times$ output (desired capital–output ratio)
Investment rule	Production = expected sales + desired stock change; errors drive dynamics	$I = \lambda(K^* - K_{-1})$, $0 < \lambda < 1$: partial adjustment
Speed of adjustment	Fast — within the period (λ near 1)	Slow — spread over many periods (distributed lag)
Main source of dynamics	Unintended inventory changes from forecast errors	Delivery lags, adjustment costs, caution, finance
Cycle implication	Short inventory cycles; amplifies turning points	Longer, smoother fixed-investment cycles

A Simple Numerical Illustration

Suppose sales rise by Rs 100 million and stay there, with $\beta = 0.25$ for inventories and $v = 3$ for capital. Desired inventories rise by Rs 25 million; the inventory accelerator concentrates essentially all of that Rs 25 million of inventory investment into the first period or two, after which inventory investment falls back to zero even though sales remain high — a sharp spike. Desired capital rises by Rs 300 million, but with $\lambda = 0.25$ the flexible accelerator produces net investment of Rs 75 million in the first period, about Rs 56 million in the second, Rs 42 million in the third, and so on — a smaller initial spike but a long tail, cumulating to the same Rs 300 million only gradually (before depreciation). One demand shock, two investment paths: the contrast in timing is exactly the difference between the two models.

Conclusion

The accelerator model of inventory investment and the flexible accelerator model of capital accumulation are siblings, not rivals. Both rest on the accelerator principle that a desired stock

is proportional to output and that investment is the flow which closes the gap between actual and desired stocks; both thereby explain why investment is far more volatile than output and why it depends on output's rate of change. They differ in the stock involved and, above all, in the speed of adjustment: inventories adjust almost completely within the period, so the inventory accelerator is effectively the $\lambda \approx 1$ case, whereas fixed capital adjusts only partially each period under the flexible accelerator's coefficient $\lambda < 1$, spreading investment over a distributed lag of past outputs. The flexible accelerator is thus the general stock-adjustment framework of which the inventory accelerator is the fast-adjusting special case.

Question 2: Tobin's Portfolio Approach to the Demand for Money (20 Marks)

Describe the Tobin's portfolio approach to demand for money in detail.

Introduction

Keynes explained the speculative demand for money by a stark all-or-nothing choice: expecting interest rates to rise or fall, each wealth holder puts everything into bonds or everything into money. James Tobin objected that this is not how people behave: real investors diversify, holding money and risky assets side by side, not because they are certain about interest rates but precisely because they are uncertain. In his 1958 paper 'Liquidity Preference as Behavior Towards Risk', Tobin rebuilt the demand for money as a problem in portfolio choice under risk, using the mean–variance analysis of his collaborator Harry Markowitz. The result — the portfolio approach — made the demand for money a branch of the general theory of asset demand, where money is held as the safe asset in a diversified portfolio. This answer sets out

the assumptions of the model, its analysis of risk averters and risk lovers, the determination of the optimum portfolio, the resulting demand function for money, and an assessment of the approach.

Assumptions and Setting

The wealth owner has a given stock of wealth W to allocate between two assets. Money is perfectly safe and liquid but yields no interest: its return is certain and equal to zero. Bonds (standing for all interest-bearing risky assets) pay an expected interest return r but their market price fluctuates, so holding them exposes the owner to an expected capital gain or loss g with mean zero and some risk, measured by the standard deviation σ of the return. The portfolio's expected return μ is the weighted average of the two asset returns, and its risk σ is proportional to the share placed in bonds: money contributes return but no risk, bonds contribute both. The investor therefore faces a trade-off along a straight opportunity line in (σ, μ) space, drawn in Figure 1: starting from the all-money point (zero risk, zero return), he can obtain higher expected return only by accepting more risk, at a rate set by the bond's return per unit of risk. Preferences are represented by indifference curves in the same space. Investors are divided into three types: risk averters, who require extra expected return as compensation for extra risk (their indifference curves slope upwards and are convex); risk lovers, who enjoy risk and will sacrifice return for it; and the risk-neutral, who care only about expected return. Tobin argued that most wealth holders are risk averters, and the main analysis concerns them.

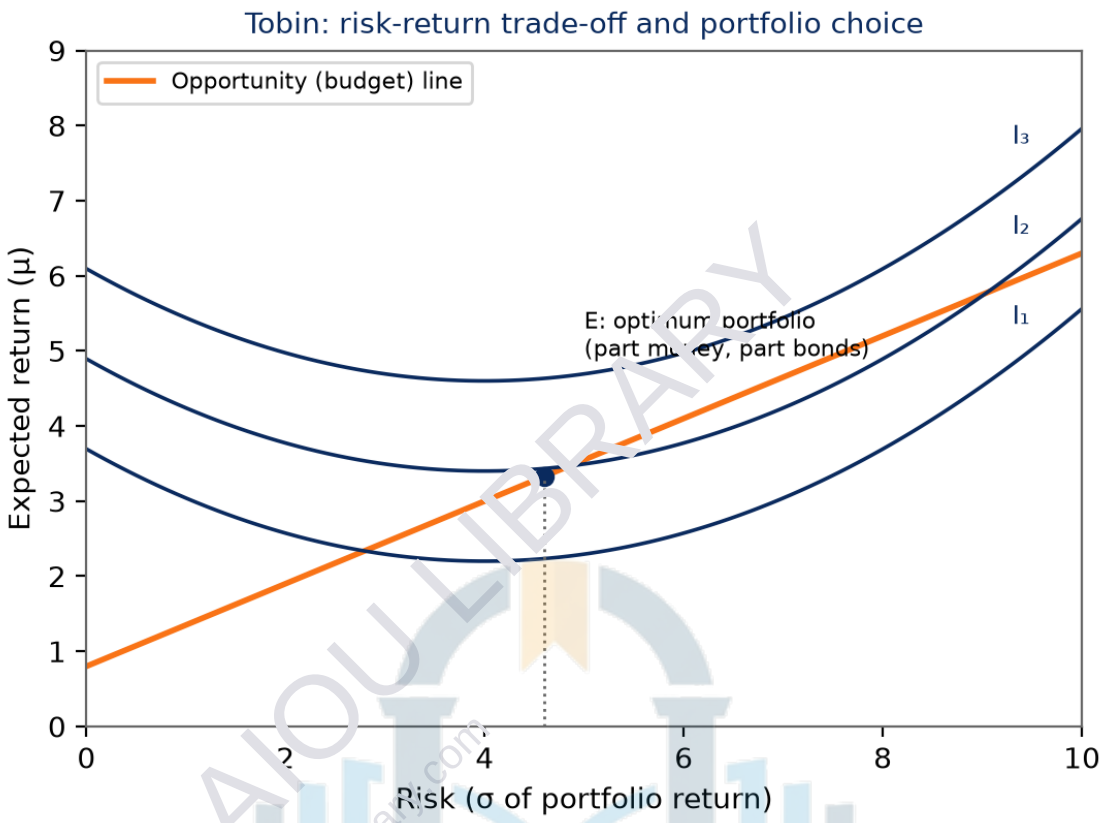

Figure 1: Tobin's portfolio balance. The opportunity line shows the expected return obtainable at each level of risk. The risk averter chooses the tangency point E on the highest attainable indifference curve — a diversified portfolio of money (safe) and bonds (risky).

The Optimum Portfolio and Diversification

The risk averter maximises satisfaction at the point E where his highest reachable indifference curve is tangent to the opportunity line. Except in limiting cases, E lies strictly between the two extremes: he holds some bonds, to earn return, and some money, to limit risk. Diversification is thus the rational response to uncertainty, not a sign of confused expectations — Tobin's central insight. Keynes's corner solutions reappear only as special cases: a risk lover runs to the risky end of the line, and an investor who expects strongly negative returns on bonds may hold only money. The chosen point divides wealth into a money share and a bond share, and anything that moves the opportunity line or the indifference map changes the demand for money.

What Determines the Demand for Money in This Approach

- The interest rate. A rise in the bond rate r rotates the opportunity line upwards — more return is now offered per unit of risk — and the tangency point normally moves towards more bonds and less money: the demand for money falls as the interest rate rises. This gives the familiar downward-sloping money demand curve, but derived from risk balancing rather than from a knife-edge speculative bet. (Tobin himself noted a subtlety: for some risk averters the income effect of a higher return can work the other way, since the same target return can now be reached with less risk; but for the community as a whole the substitution effect dominates and the aggregate demand for money is interest-inelastic in neither direction — it is simply negatively related to r.)
- Wealth. An increase in wealth expands attainable portfolios proportionally and, with the usual assumption that money is a 'normal' asset whose share does not fall as wealth grows, raises the amount of money held. The demand for money therefore rises with wealth — and hence, in the aggregate, with income, which connects the portfolio motive to the transactions story.
- Risk. If bonds become riskier (the variability of capital values increases), the opportunity line flattens: each unit of return now costs more risk. Risk averters shift towards money, and the demand for money rises at every interest rate. Financial crises are the vivid case: a spike in perceived risk produces a flight to liquidity that Keynes's model could describe only crudely.
- Expectations and other assets' returns. Expected inflation, expected returns on equities and property, and the riskiness of every alternative asset enter the same calculus; money demand is one equation in a system of asset demands that must add up to total wealth —

the feature that later made the approach the foundation of multi-asset macro models and of Friedman's restatement of the quantity theory.

Merits of the Approach

The portfolio approach repaired the weakest joint in Keynesian liquidity preference. It explains observed diversification instead of predicting implausible corners; it derives an interest-elastic demand for money from rational, risk-averse behaviour without requiring each holder to forecast interest rates confidently; it explains why money is held even when interest rates are high and expected to be stable (as a risk-reducing asset, money is never dominated); and it embeds money in the balance sheet so that monetary policy can be analysed as changing the relative supplies of assets that the public must be induced to hold — the germ of the portfolio-balance models of open-economy macroeconomics. It also unified the theory of money demand with the theory of finance, where mean–variance portfolio selection became the standard tool.

Criticisms and Limitations

The approach has real limitations. Two assets are a drastic simplification: the real menu includes deposits, bills, bonds of many maturities, equities and real assets, with liquidity differing by degree. Mean–variance preferences are tractable but special — they are exactly right only for quadratic utility or normally distributed returns, and they treat gains and losses symmetrically although people demonstrably dislike losses more. The model is static and partial: it takes wealth and expected returns as given and says little about the transactions role of money that Baumol and Tobin himself modelled separately, and the individual response of money demand to the interest rate is not unambiguous once income effects are admitted.

Finally, in modern financial systems the relevant ‘safe asset’ is usually an interest-bearing deposit, so the stark money-versus-bonds contrast needs restating in terms of relative risks and returns across the whole spectrum. None of this overturns the core insight; it marks out where later work had to go.

Conclusion

Tobin’s portfolio approach recasts the demand for money as one side of a risk–return balancing act. Wealth holders divide their wealth between safe, unrewarding money and risky, rewarding bonds so as to reach the tangency of their highest indifference curve with the opportunity line, and they therefore normally hold both. From that single construction flow the main determinants of money demand: it falls when interest rates rise, rises with wealth, and rises when alternative assets become riskier or their expected returns fall. By grounding liquidity preference in rational diversification, Tobin removed the all-or-nothing fragility of Keynes’s speculative motive and founded the modern treatment of money as one asset among many in the public’s portfolio.

Question 3: Inflation Expectations, the Shifting Phillips Curve and Automatic Adjustment to Shocks (20 Marks)

Explain how the ability of inflation expectations to shift the Phillips curve helps the economy to adjust, automatically, to aggregate supply and demand shocks.

Introduction

The original Phillips curve offered policy-makers a stable menu: choose a point combining lower unemployment with higher inflation, or lower inflation with higher unemployment. The

experience of the 1970s destroyed that menu — inflation and unemployment rose together — and the explanation, supplied independently by Friedman and Phelps, was that the curve shifts with expected inflation. Far from being a mere complication, this shifting property is what allows a market economy to absorb shocks and return by itself to its natural rate of unemployment without permanent policy intervention. This answer states the expectations-augmented Phillips curve, and then traces the automatic adjustment mechanism for a demand shock and for a supply shock, showing in each case how moving expectations do the work that a fixed-curve world would have left undone.

The Expectations-Augmented Phillips Curve

The modern Phillips curve is written:

$$\pi = \pi^e - \beta(u - u_n) + v, \quad \beta > 0,$$

where π is actual inflation, π^e is expected inflation, u is the unemployment rate, u_n is the natural rate, and v captures supply shocks. For a given π^e there is a downward-sloping short-run Phillips curve (SRPC): unemployment below u_n generates inflation above expectations. But the curve is indexed by π^e: when expected inflation rises, the whole SRPC shifts upwards by the same amount — at every unemployment rate, inflation is higher, because wage bargainers and price setters build the expected inflation into nominal settlements in order to protect real wages and margins. When π^e falls, the curve shifts down. In the long run, expectations are fulfilled ($\pi = \pi^e$, $v = 0$), so $u = u_n$ whatever the inflation rate: the long-run Phillips curve (LRPC) is vertical at the natural rate, as drawn in Figure 2. Adjustment, in this framework, is simply the economy moving between short-run curves as expectations revise, until it lands back on the LRPC.

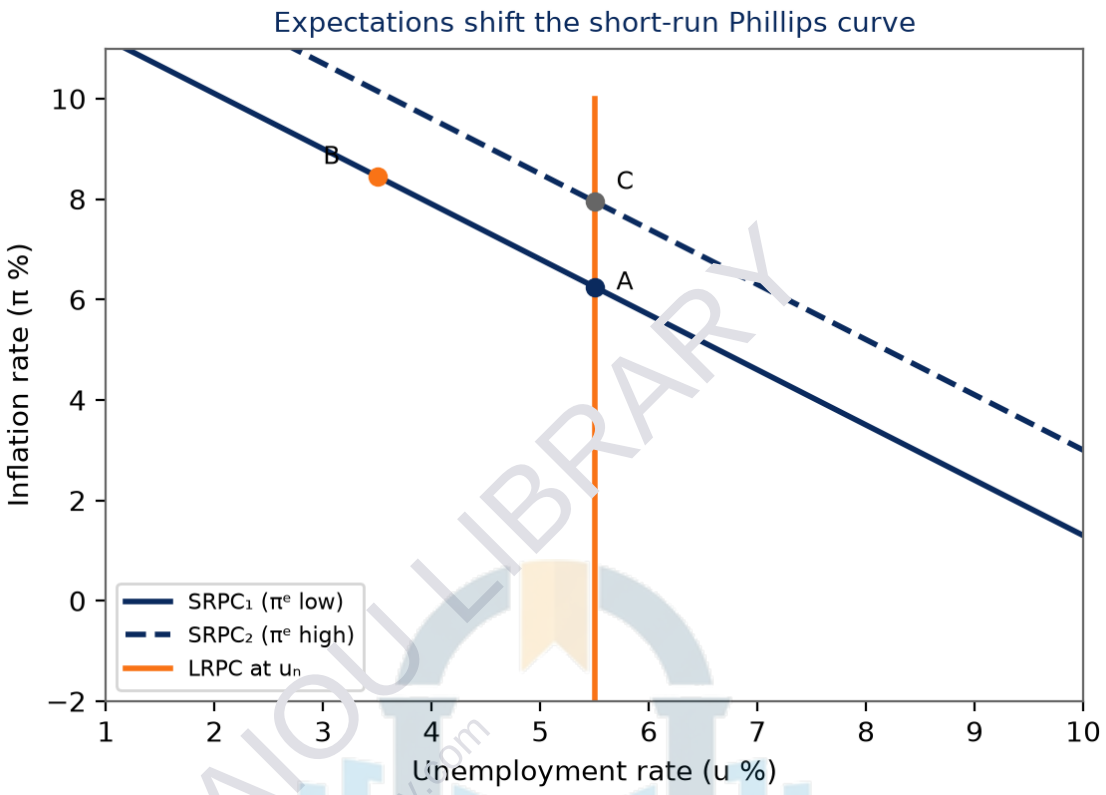

Figure 2: The expectations-augmented Phillips curve. A demand expansion moves the economy A → B along $SRPC_1$; as expectations catch up, SRPC shifts up to $SRPC_2$ and the economy settles at C — back at the natural rate u_n, with permanently higher inflation.

Adjustment to a Demand Shock

Consider first a positive aggregate demand shock — say a monetary expansion — hitting an economy resting at point A in Figure 2, with $u = u_n$ and inflation equal to expected inflation. Spending rises, firms expand output and hiring, and unemployment falls below u_n to point B. Labour markets tighten, nominal wages are bid up faster than expected, firms pass the costs on, and actual inflation rises above π^e. This is the only phase in which the shock has real effects, and it lasts exactly as long as expectations lag behind reality. Wage earners and firms now revise π^e upwards towards the new actual inflation. As they do, the SRPC shifts up: settlements now embody higher expected inflation, so at any given unemployment rate inflation is higher, and the stimulus needed to hold unemployment below u_n grows. The economy slides up and to

the right along successive SRPCs until, at point C, unemployment is back at u_n and expectations are again correct — at a permanently higher inflation rate. The return of unemployment to its natural rate is automatic: no policy reversal was required, because expectation adjustment itself progressively neutralised the real effect of the nominal shock. A negative demand shock unwinds symmetrically: unemployment rises above u_n, inflation falls below expectations, π^e is revised down, the SRPC shifts down, and the economy returns to u_n at lower inflation — although, as discussed below, this downward adjustment is typically slower and more painful.

Adjustment to a Supply Shock

A supply shock enters through v. An adverse shock — a sharp rise in imported oil prices, a harvest failure — shifts the SRPC upwards immediately: at unchanged expectations and any given unemployment rate, inflation is higher, and if demand policy does nothing unemployment also rises as firms cut output in the face of higher costs: stagflation. Now the shifting mechanism goes to work in reverse gear. The shock raises the price level but, if it is a one-off change in relative prices rather than a continuing process, actual inflation falls back once the price-level jump has passed. Because actual inflation then runs below the expectations formed during the shock, π^e is gradually revised downwards, the SRPC shifts back down, and unemployment returns to u_n without any permanent increase in inflation. The economy, in other words, distinguishes a temporary supply disturbance from a permanent change in the inflation regime, and expectation adjustment performs that distinction automatically. The crucial condition is that policy does not 'validate' the shock: if the central bank expands demand to fight the unemployment, actual inflation stays high, expectations ratchet up, the SRPC stays shifted, and what began as a one-off price-level shock is converted into permanently

higher inflation — the 1970s error. A favourable supply shock runs the film backwards: the SRPC shifts down, inflation falls at unchanged unemployment, expectations follow down, and the gain is locked in.

Why the Automatic Mechanism Matters — and Its Limits

The economic significance of the mechanism is that stability does not depend on continuous fine-tuning. Because expectations chase actual inflation and shift the SRPC, demand shocks have only temporary real effects and one-off supply shocks have only temporary inflation effects; the economy is self-correcting towards the natural rate, and policy's proper long-run task reduces to choosing the inflation rate at which that return happens. The mechanism also explains the costs of disinflation: to reduce entrenched inflation, π^e must be dragged down, and with adaptive expectations that requires a period of unemployment above u_n — the sacrifice ratio. Two limits deserve emphasis. First, speed: with backward-looking expectations, adjustment is gradual, and 'automatic' does not mean 'fast' or 'painless'; credible policy announcements can speed it up by moving π^e directly, which is why central-bank credibility matters. Second, hysteresis: if long unemployment spells erode skills and detach workers from the labour force, u_n itself rises during the adjustment, and the vertical LRPC shifts — the return destination has moved. Both limits qualify, but neither removes, the central self-adjusting property.

Conclusion

Inflation expectations shift the short-run Phillips curve one-for-one, and this shifting is the economy's automatic stabiliser at the level of the price–unemployment relation. After a demand shock, the economy moves along the SRPC only until expectations catch up; the curve

then shifts, progressively cancelling the shock's real effects and returning unemployment to its natural rate at a new inflation rate. After a temporary supply shock, the curve shifts adversely but shifts back as actual inflation subsides and expectations revise down — provided policy does not accommodate the shock into a permanent inflation. The same mechanism that Friedman and Phelps used to refute the permanent trade-off thus also explains how a decentralised economy heals itself after disturbances, and why the long-run Phillips curve is vertical at the natural rate.

Question 4: The Harrod–Domar Model: A Critical Analysis and Its Relevance for Developing Countries (20 Marks)

Critically analyze the Harrod-Domar model and its relevance for developing countries.

Introduction

The Harrod–Domar model, developed independently by Roy Harrod (1939) and Evsey Domar (1946), was the first rigorous attempt to turn Keynes's short-run apparatus towards the long-run problem of growth. By asking what rate of growth of income is required to keep a growing capital stock fully employed, it produced one famous equation — the warranted growth rate equals the saving ratio divided by the capital–output ratio — and one famous anxiety: that this balanced path is a knife-edge. The model became the workhorse of post-war development planning and the intellectual basis of foreign-aid calculations. This answer states the model and its assumptions, works a numerical example, sets out the main criticisms, and then judges its relevance for developing countries, where it has been used most and where its defects show most clearly.

The Model

Saving is a fixed proportion of income, S = sY. Investment creates demand through the multiplier and simultaneously creates capacity: Domar's 'dual character' of investment. Capacity grows according to the productivity of capital, σ = 1/v, where v is the fixed incremental capital–output ratio (ICOR): a capital addition I can support extra output I/v. For the new capacity to be exactly absorbed by new demand, income must grow at the rate at which investment itself generates income; equilibrium between the demand and supply effects requires the warranted rate of growth:

$g_w = s / v.$

Harrod distinguished three rates: the actual rate g, the warranted rate g_w (the growth that leaves entrepreneurs content with their investment decisions), and the natural rate g_n (growth of the labour force plus labour productivity). If actual growth exceeds the warranted rate, capacity is strained, firms invest still more, and growth accelerates away; if it falls short, excess capacity discourages investment and the economy spirals down. Balanced, self-sustaining growth therefore requires $g = g_w = g_n$ simultaneously — a triple coincidence between parameters set by savers (s), technology (v), and demography (g_n) that nothing in the model brings about. This instability is the knife-edge.

A Numerical Illustration

Take a developing economy saving 24 per cent of income with an ICOR of 4. The warranted growth rate is $g_w = 0.24/4 = 6$ per cent per year. If population grows at 2.5 per cent and there is no productivity growth, the natural rate is only 2.5 per cent; since $g_w > g_n$, the economy's saving effort outruns its ability to employ the growing labour force with the capital being

created — capital accumulation proceeds but unemployment of labour worsens. Conversely, if saving were only 12 per cent, g_w = 3 per cent, and raising growth to 6 per cent would seem to require doubling saving or halving the ICOR — the arithmetic that drove the 'financing gap' aid doctrine: aid fills the gap between domestic saving and the investment needed for a target growth rate. The example shows both the model's planning appeal (targets reduce to two measurable parameters) and its rigidity (growth can be raised only by s or v, nothing else).

Table 1: Harrod–Domar arithmetic for selected saving rates (ICOR v = 4)

Saving rate (s)	Warranted growth g = s/v	Per-capita growth if population grows 2.5%
12%	3.0%	0.5%
16%	4.0%	1.5%
20%	5.0%	2.5%
24%	6.0%	3.5%
28%	7.0%	4.5%

Critical Analysis

- Fixed proportions. The model assumes capital and labour combine in rigid proportions, so v is a constant of nature. Solow showed that once substitution between factors is allowed, the capital–output ratio adjusts and the knife-edge disappears: growth converges to the natural rate, and the model's instability is an artefact of its most unrealistic assumption.
- Empirical instability of v. In practice the ICOR varies widely across countries, sectors and decades, and rises when investment is badly allocated; a parameter that is supposed to carry the whole theory is neither constant nor mainly technological — it reflects institutions, project selection and efficiency, which the model cannot discuss.

- No technical progress or human capital. With labour-augmenting progress absent, long-run per-capita growth is chained to demography; the model has nothing to say about education, skills, innovation or technology transfer, which are the substance of development.
- Saving is not the only constraint, and may not be binding. The model assumes saving automatically becomes productive investment through an efficient financial system. In many economies the binding constraints are foreign exchange (the two-gap extension of Chenery), infrastructure, governance, demand, or entrepreneurship; forced saving drives that finance wasteful projects lower effective growth.
- Closed-economy and no-government simplifications. Trade, aid, debt, fiscal policy and relative prices are absent although for developing countries these are first-order determinants of what investment is feasible and profitable.
- Knife-edge exaggeration. Real economies do not explode or collapse when growth deviates from a warranted path; prices, wages, interest rates and utilisation rates adjust. The instability result dramatises a genuine point — that demand for the output of new capacity is not guaranteed — but as literal dynamics it is too strong, as Solow's stable alternative demonstrates.

Relevance for Developing Countries

Despite these defects, the model spoke directly to the situation of post-colonial developing countries, and its influence was correspondingly deep. Capital genuinely was their scarcest factor: poor economies saved little, financial markets were rudimentary, and the model identified the proximate bottleneck — the rate of capital formation — around which the Five-Year Plans of Pakistan, India and many others were organised. Its simple arithmetic gave

planners targets for saving, investment and aid that finance ministries still compute: the World Bank and planning commissions used financing-gap versions of Harrod–Domar for decades, and the model underwrites the still-sound insight that sustained growth requires a high and efficiently used investment rate — East Asia's success was built on saving and investment rates far above those the model would have been given in 1950. Its warnings are equally relevant in reverse: where states treated s/v as a mechanical recipe, raising investment by decree while ignoring the efficiency term v, the result was white-elephant projects, rising ICORs and debt without growth — a pattern visible in several aid-financed industrialisation drives. The model is therefore best used in developing countries as a consistency framework and a first approximation — a check that saving, investment and growth targets cohere — supplemented by the things it omits: human capital, technology, institutions, export demand and the quality of investment allocation.

Conclusion

The Harrod–Domar model reduces growth to g = s/v and balances the economy on a knife-edge where actual, warranted and natural growth coincide. Critically, its fixed proportions, constant capital–output ratio, silence on technology and human capital, and assumption that saving automatically becomes productive investment are all seriously unrealistic, and Solow's neoclassical model removed its instability result by relaxing the first of them. For developing countries the model remains relevant in the limited but real sense that capital scarcity and low saving were, and in many places still are, binding constraints, and its arithmetic still disciplines plans and aid calculations. Used as a mechanical growth recipe it misleads and has misled; used

as a framework insisting that growth must be financed by real saving and validated by demand, it retains a permanent place in development economics.

Question 5: What Determines the Steady-State Growth Rate of Per Capita Output? Short-Run Influences on Growth (20 Marks)

What factors determine the growth rate of steady-state per capita output? Are there other factors that could affect the growth rate of output in the short run?

Introduction

The question is the central question of the Solow growth model, and its answer is one of the most surprising results in economics: in the long run — in the steady state — an economy's output per head can grow only as fast as technical progress, and neither saving more nor accumulating capital faster can raise that long-run growth rate, although both raise the level of income. In the short run, by contrast, almost everything matters: saving, population growth, demand conditions and the distance from the steady state all change the growth rate temporarily. This answer derives the steady-state result from the Solow apparatus, distinguishes carefully between level effects and growth effects, and then lists the factors that move growth in the short run, including the transitional dynamics that dominate the experience of real economies.

The Solow Framework and the Steady State

Output is produced with capital and effective labour: Y = F(K, AL), where A is labour-augmenting technology growing at the exogenous rate g, and the labour force grows at rate n. Writing k for capital per effective worker and y = f(k) for output per effective worker, capital accumulates

according to $\dot{k} = sf(k) - (n + g + \delta)k$: saving per effective worker, $sf(k)$, must cover the investment needed to equip new workers, to keep up with technical progress and to replace depreciation, $(n + g + \delta)k$. The steady state k^* is the point in Figure 3 where the two are equal. At k^*, capital and output per effective worker are constant — that is the definition of the steady state — while output per actual worker, $Y/L = A \cdot y$, grows at exactly the rate of growth of A, namely g. Aggregate output grows at $n + g$. Hence the model's headline answer: the steady-state growth rate of per capita output is determined solely by the (exogenous) rate of technological progress, g.

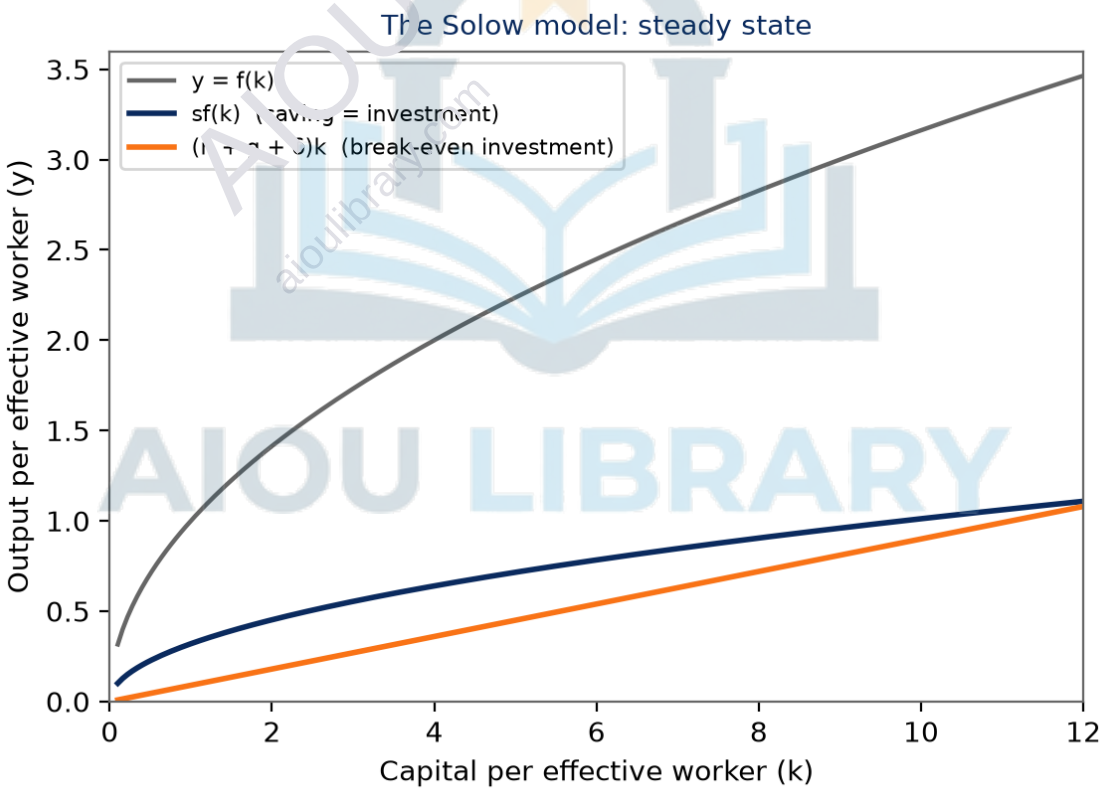

Figure 3: The Solow diagram. The economy rests where saving per effective worker sf(k) equals break-even investment (n + g + δ)k. Per capita output at that point grows only at the rate of technical progress, g.

Level Effects versus Growth Effects

The parameters often confused with growth determinants — the saving rate s, population growth n and depreciation δ — in fact determine levels, not the steady-state growth rate. A higher saving rate shifts sf(k) upwards, raises k* and raises the steady-state level of output per head (a level effect), and while the economy travels to the new steady state growth is temporarily above g; but once there, per capita growth falls back to g. Higher population growth swells the break-even line, lowers k* and lowers the steady-state level of income per head, again without changing its long-run growth rate. Technology is doubly privileged: its level A determines where the steady-state path lies, and its growth rate g determines how fast the whole path rises. This distinction explains why two economies with identical saving behaviour can have permanently different incomes (different A) and why the phenomenal investment rates of East Asia produced decades of rapid growth that nonetheless slowed as those economies neared their steady states.

Table 1: Level effects and growth effects in the Solow steady state

Factor	Effect on steady-state level of output per head	Effect on steady-state growth of output per head
Saving rate (s) ↑	Rises (higher k*)	None (only transitional growth while k rises)
Population growth (n) ↑	Falls (capital dilution)	None
Depreciation (δ) ↑	Falls	None
Technology level (A) ↑	Rises proportionally	None
Rate of technical progress (g) ↑	Path shifts up over time	Rises one-for-one — the sole long-run growth determinant

Factors Affecting Growth in the Short Run

Outside the steady state, the growth rate is rich in determinants, and most observed growth episodes are precisely such transitional or short-run phenomena. First, the distance from the steady state: an economy below its k* — after a war, a disaster, or at the start of a catch-up process — has a marginal product of capital far above the break-even requirement, so each unit of saving yields rapid capital deepening and growth well above g; this conditional convergence is the model's explanation of post-war German and Japanese miracles and of fast-growing poor economies, and of why growth slows as it approaches k*. Second, changes in the saving and investment rate: raising s lifts the economy onto a higher path, generating a transitional growth burst that can last decades given the slow speed of convergence (roughly a few per cent of the remaining gap per year). Third, population and labour-force changes — demographic transitions, migration, rising female participation, longer working lives through education and health — change transitional capital widening needs and the effective labour input directly. Fourth, human capital: in the augmented Solow model of Mankiw, Romer and Weil, investment in education behaves like investment in physical capital, so schooling drives powerful transitional growth and explains much of the cross-country income variation that the basic model leaves to A. Fifth, demand-side and cyclical factors entirely outside the supply-side model: in the short run output is demand-determined, so monetary and fiscal policy, export booms, terms-of-trade movements, financial crises and confidence shocks all move the growth rate for quarters or years, even though they leave the steady state untouched. Sixth, one-off technology events — adoption of an existing technology backlog, a resource discovery, institutional reform that raises total factor productivity — shift A's level and produce

transitional growth spurts. Finally, endogenous growth theory (Romer, Lucas, Aghion–Howitt) adds the important qualification that if knowledge accumulation has increasing returns or is driven by profit-seeking R&D and education, policy and saving can influence even the long-run growth rate — g ceases to be manna from heaven and becomes partly a choice variable.

Conclusion

In the Solow steady state, output per effective worker is constant, so output per head grows at the rate of technological progress g — and only g determines that growth rate. Saving, population growth and depreciation determine the level of the steady-state path, not its slope; confusing the two is the commonest error in popular growth discussion. In the short run, however, growth is affected by a long list of additional factors: the gap between actual and steady-state capital (convergence), changes in saving and investment, demographic change, human-capital accumulation, technology adoption and institutional reform, and demand-side shocks that dominate at business-cycle horizons. Since real economies spend much of their lives in transition rather than in the steady state, these short-run determinants are not a footnote to the theory but the main story of observed growth — with the steady-state result standing behind them as the discipline that such bursts must eventually subside to the rate set by technical progress.