

Manufacturing technician: study guide

Manufacturing, production, maintenance, and quality · suite apt-157-manufacturing-technician ·
generated 2026-09-15T15:37:53.587Z

Detail	Value
Title	Manufacturing technician: study guide
Generated	2026-09-15T15:37:53.587Z
Fixture/version	apt-157-manufacturing-technician
Sector	Manufacturing, production, maintenance, and quality
Guide version	v2

Private practice result. Not an official exam certificate, employer decision, hiring signal, admissions decision, or guaranteed outcome. Scores stay on this device unless you export them.

Answer keys and scoring logic stay server-side and are never included in any download or export.

How to use this guide

This file contains the whole study outline for this suite: every skill it draws on, the full lesson for each of those skills, worked examples, practice tips, a glossary, and where each piece of material comes from. Nothing here is a summary of a page you still have to visit.

1. Read the skills section end to end once, without timing yourself.
2. Work the guided practice mode for the suite, using the practice tips as a checklist.
3. Move to the mini-test only when guided practice feels unhurried.
4. Sit the full simulation last, once, in the conditions you expect on the day.

Practice attempts are stored on the device you used, never on an account. Exporting a result is the only way anything leaves that device.

Practice modes and durations

Practice modes for Manufacturing technician

Mode	Duration	What it is for
Guided practice	15 minutes	Untimed, with feedback after every item.
Mini-test	18 minutes	A short timed set for checking pace.
Full simulation	45 minutes	Full length and full time, in one sitting.

- Start guided practice:
<https://learn.novusstreamsolutions.com/aptitude/practice/new?bank=apt-157-manufacturing-technician&mode=guided>
- Start mini-test:
<https://learn.novusstreamsolutions.com/aptitude/practice/new?bank=apt-157-manufacturing-technician&mode=mini>
- Start full simulation:
<https://learn.novusstreamsolutions.com/aptitude/practice/new?bank=apt-157-manufacturing-technician&mode=full>

Skills covered, in full

This suite draws on 5 skill constructs. Each one below carries its complete lesson.

Technical reasoning

Applied technical principles, diagrams, tools, systems, measurements, and troubleshooting.

Technical reasoning is what sits above any single trade: reading a system diagram for what it actually does, isolating a fault by measurement instead of by guesswork, and taking the governing number off a drawing or a nameplate without importing assumptions. It is assessed in HVAC, instrumentation, mechatronics, process-operator and engineering-technician selection, and it is the construct that best predicts whether someone can diagnose an unfamiliar machine. Employers care about it because part-swapping is expensive and half-splitting is not. Practice sessions here stay on your device unless you choose to export them.

What you should be able to do after this lesson:

1. Trace a process or signal diagram end to end and state, for a given symptom, which components could physically cause it and which could not.
2. Isolate a fault by half-splitting a chain of stages, and say how many measurements a chain of a given length should take.
3. Convert a dimension with asymmetric tolerance into an acceptance window and decide whether a measured part passes, can be reworked, or is scrap.
4. Match a measuring instrument to a required resolution, and distinguish an instrument's resolution from its accuracy.
5. Describe a closed control loop as sensor, controller, setpoint, actuator and feedback, then design the one measurement that separates a sensor fault from an actuator fault.
6. Read the qualifier attached to a rated number (duty cycle, working load limit, working pressure, nominal versus maximum), and apply it correctly.

Worked examples and pitfalls

Half-splitting beats swapping parts: A conveyor will not stop when its photo-eye is blocked. The chain is: sensor, field cable, junction box, PLC input card, PLC program, output card, interposing relay, contactor. Eight places the signal can die. Swapping parts one at a time means four or five attempts on average, since the culprit is equally likely to sit anywhere in the eight, and every attempt costs a part. Half-splitting starts in the middle instead: watch the PLC input LED while a colleague blocks the beam. If it toggles, the sensor, field cable, junction box and input card are all proven good in a single observation, and eight candidates become four. Force the output in the PLC and watch the contactor: if it pulls in, the output card, interposing relay and contactor are good as well, so the fault is in the program logic rather than the hardware. Three checks resolve eight stages, because each check halves the remaining suspects. The trap is starting at whichever end is easiest to reach, which resolves one stage per check instead of half of them. The second trap is the sentence 'I replaced the sensor and it still fails', that proves only that the sensor was not the fault, at the price of a part and an hour.

Tolerance: in spec, or scrap?: A drawing calls a shaft 25.00 mm with a tolerance of plus 0.05 and minus 0.10. That is an asymmetric tolerance, so the acceptance window runs from 24.90 mm to 25.05 mm and the nominal is not at its centre. A part measuring 24.92 mm is inside the window and passes, even though it is below nominal: the most common wrong rejection on this style of item, made by anyone who silently reads the tolerance as plus or minus 0.05. A part at 25.06 mm fails by 0.01 mm, but it fails oversize, so material can still be removed and it is rework rather than scrap. A part at 24.85 mm fails undersize and there is no recovering it. One more layer the better items include: if you took that 25.06 reading on a caliper with 0.02 mm resolution, the reading is at the very limit of what the instrument can resolve, and the honest next step is

to re-measure with a micrometer before anyone scraps or reworks anything.

Reading a system diagram: what can actually cause this?: A tank fill line is drawn as supply, isolation valve V1, strainer, pump P1, check valve, control valve CV1, tank. A high-level switch LSH-1 is wired to close CV1. The reported symptom is that the tank overfilled. Work the path between the measurement and the element that stops flow: a CV1 that has stuck open, an LSH-1 that never actuated, and a broken wire in the LSH-1 loop are all consistent with the symptom. A blocked strainer is not: restricting the inlet reduces flow, and no amount of restriction causes an overflow. Nor is the check valve, whose function is to prevent reverse flow, not forward flow. Candidates pick the strainer because it is the component they know fouls in service, which is a memory of maintenance history rather than a reading of the diagram. The discipline that earns the mark is directional: a component can only be responsible if it lies on the causal path AND its failure mode pushes the system in the direction of the symptom.

Instrument choice: resolution is not accuracy: A steel rule resolves to roughly 0.5 mm. A vernier caliper marked 0.02 mm resolves to 0.02 mm. A 0 to 25 mm micrometer resolves 0.01 mm on the thimble, or 0.001 mm if it carries a vernier. A dial indicator reads 0.01 mm of relative movement but tells you nothing about absolute size without a reference. Asked to verify a 25.00 mm shaft with a tolerance of plus or minus 0.02 mm, the tolerance band is 0.04 mm wide: two divisions on that caliper, which is not enough to judge anything reliably. The workshop convention is that the instrument should resolve to about a tenth of the tolerance band, here 0.004 mm, so even the micrometer is marginal and comparison against gauge blocks is the defensible answer. The trap the item is built around is a digital display: showing four decimal places is a statement about resolution, not accuracy. An uncalibrated digital caliper will report 25.0000 mm with total confidence and be 0.03 mm out.

Closed loop: which element failed?: A room is meant to hold 21 degrees Celsius. A thermostat containing the sensor and the controller drives a valve on a radiator. The symptom: the room reaches 28 degrees Celsius and the valve stays open. Three explanations survive first inspection. The sensor reads low so the controller still believes the room is cold, the valve is mechanically jammed open, or the controller output has failed in the on state. One measurement separates them. Put an independent thermometer beside the thermostat. If the thermostat displays 17 degrees while the thermometer reads 28, the sensor is lying and everything downstream is behaving correctly. If the thermostat displays 28 and is still calling for heat, the sensor is fine and the fault is in the controller or the valve, which you then split by checking whether the valve actuator is being energised. The tempting non-answer is 'the room is too hot, so lower the setpoint'. That treats the symptom, and if the sensor reads seven degrees low the loop will simply settle seven degrees high again at the new setpoint.

Nameplates: read the qualifier, not just the number: A welding machine is rated 200 A at 40 percent duty cycle over a ten-minute period. That means four minutes of arc time and six minutes of cooling in every ten, at the full 200 A. It does not mean 40 percent of 200 A, and it does not mean 40 percent of an hour. Both wrong readings feel entirely natural, which is why they make good distractors. Run the machine continuously at 200 A and the thermal cut-out will open. The same discipline transfers across the whole trade: a hoist's working load limit is not its breaking load, a hose's working pressure is not its burst pressure, a motor's service factor describes a short-term overload allowance and not a continuous rating, and a pump curve's flow figure is quoted at a stated head. Whenever an item hands you a number in a table or on a plate, underline the qualifier printed next to it before you calculate anything; the wrong options are usually built by dropping exactly one qualifier.

How to practise this skill

- Trace every diagram with a pencil from input to output and name each block as you pass it. Technical items punish skimming far harder than they punish slow arithmetic.
- Rehearse half-splitting on systems you already know (a home network that has dropped out, a car that will not crank), and count the checks. The habit transfers to the test intact.

- Underline the qualifier beside every number a question supplies: per hour, at 20 degrees Celsius, at 40 percent duty, nominal, maximum. That single mark-up defuses most distractors.
- Convert the whole question to one unit system in one pass before calculating, rather than converting each intermediate result and accumulating rounding.
- Train yourself to state what a measurement PROVES rather than what it reads. '12 V present at the coil' proves the supply and the whole upstream path; it says nothing about whether the coil itself is good.
- Keep a running list of the schematic symbols and component names you personally keep getting wrong, and drill only those. Five focused minutes beats another full untargeted set, and your attempt history stays local to this device so the list is yours alone.

Glossary

Half-split fault finding: Testing at the midpoint of a chain of stages so that each measurement eliminates half the remaining suspects. A chain of eight stages resolves in about three checks rather than four swaps.

Tolerance band: The distance between the upper and lower acceptance limits of a dimension. A tolerance written as plus 0.05 and minus 0.10 has a 0.15 mm band that is not centred on the nominal.

Resolution: The smallest change an instrument can display. A four-decimal readout has fine resolution and may still be inaccurate if the instrument is out of calibration.

Accuracy: How close a reading is to the true value, established by calibration against a traceable standard. Independent of resolution, and the property that decides whether a part passes.

Interlock: A condition wired or programmed to inhibit an action until it is satisfied, such as a guard-door switch that prevents a motor start while the door is open.

Closed-loop control: A control arrangement where a sensor measures the process, a controller compares that measurement to a setpoint, and an actuator drives the process until the difference closes.

Duty cycle: The proportion of a stated period during which equipment may operate at a stated output, for example 40 percent of a ten-minute period, after which it must cool.

Schematic versus pictorial diagram: A schematic shows function and connection logic with no regard to physical layout; a pictorial or exploded view shows physical arrangement and assembly order but hides the logic.

Root cause: The condition whose removal stops a failure recurring, as opposed to the symptom, which is only what became visible.

Study this next

- Technical reasoning skill hub: <https://learn.novusstreamsolutions.com/aptitude/skills/technical-reasoning>
- HVAC and refrigeration suite: <https://learn.novusstreamsolutions.com/aptitude/tests/hvac-and-refrigeration>
- Instrumentation and control suite:
<https://learn.novusstreamsolutions.com/aptitude/tests/instrumentation-and-control>
- Engineering and technical design careers:
<https://learn.novusstreamsolutions.com/careers/families/engineering-and-technical-design>
- Encyclopedia search: troubleshooting: <https://learn.novusstreamsolutions.com/search?q=Troubleshooting>
- Technical reasoning lesson on Novus Learn:
<https://learn.novusstreamsolutions.com/aptitude/skills/technical-reasoning/lesson>

Where this material comes from

- Diagnostic and metrology examples written for Novus Learn from general maintenance practice: binary-search fault isolation, asymmetric dimensional tolerance, and the ten-to-one instrument selection convention.

- Terminology checked against the public Wikipedia articles 'Troubleshooting', 'Engineering tolerance', 'Accuracy and precision' and 'Control loop'. Definitions only; every scenario above is original.
- Novus Learn aptitude construct registry (catalog seed) for the construct scope and related suite mapping.
- Public educational framing only: not affiliated with any official exam board, publisher or employer. Ratings and limits described here are illustrative and never override the equipment documentation in front of you.

Educational preparation only. Novus Learn does not administer official exams and does not guarantee scores or hiring outcomes.

Data interpretation

Tables, charts, graphs, dashboards, trends, comparisons, and evidence-based conclusions.

Data interpretation is the discipline of getting a correct number out of a table, chart or dashboard that was not built to make your question easy, and of saying so when the data cannot answer it at all. It dominates graduate and analyst screening, and it is the section where strong arithmetic still fails, because the marks are lost in the header row, the axis scale and the wording of the question. The same skill is the daily work of anyone who reports on a management pack, a clinical audit or a stock ledger. Practice is stored on this device only; there is no account and nothing is uploaded unless you export it.

What you should be able to do after this lesson:

1. Read a value correctly out of a table or chart including its units multiplier, footnotes and any 'excludes' or 'provisional' qualifier attached to the row.
2. Compute share of total, percentage change and percentage-point change from the same pair of cells, identify which of the three a question is asking for, and separate a movement in a rate from a movement in the underlying count.
3. Work with index numbers relative to a base year, including why a change of five index points is almost never a five percent change.
4. Join two tables on a shared key and produce a normalised figure (per head, per unit, per thousand) rather than comparing raw totals.
5. Recognise chart presentation effects (truncated axes, dual axes, cumulative versus periodic series), and answer from the numbers rather than from the visual impression.
6. Apply the 'cannot say' discipline: state precisely which extra fact would be needed before the question becomes answerable.

Worked examples and pitfalls

Read the header: (£000) changes every answer by a factor of a thousand: Table titled 'Regional revenue, year to March (£000)': North 1,240; South 986; East 1,455; West 719. Total = $1,240 + 986 + 1,455 + 719 = 4,400$, so the business turned over 4,400 thousand pounds, that is 4.4 million. East's share is $1,455 / 4,400 = 33.1$ percent. The whole set: North 28.2 percent, South 22.4, East 33.1, West 16.3, summing to 100. Two things go wrong here. The first is reading East's revenue as 1,455 pounds and then reporting a business with a total turnover of 4,400 pounds, which nobody notices because every option is scaled the same way, until the question asks for revenue in millions and only one option is right. The second is the comparison wording. East's share is $(1,455 - 719) / 4,400 = 16.7$ percentage points above West's, and East's revenue is $(1,455 - 719) / 719 = 102$ percent more than West's, that is slightly more than double. 'Sixteen point seven' and 'a hundred and two' both describe the same two cells honestly, and the question decides which one is correct. Note that the percentage-point figure must be computed from the unrounded shares rather than by subtracting the rounded ones, or the last digit will not survive.

One pair of rows, three correct increases: A complaints table: 2023, 120,000 orders, complaint rate 4.0 percent; 2024, 150,000 orders, complaint rate 5.0 percent. Three defensible answers to 'how much did

complaints increase?'. The rate rose by 1.0 percentage point. The rate rose by $(5.0 - 4.0) / 4.0 = 25$ percent in relative terms. And the count of complaints rose from $0.04 \times 120,000 = 4,800$ to $0.05 \times 150,000 = 7,500$, which is $(7,500 - 4,800) / 4,800 = 56.25$ percent. All three are arithmetically right; only one answers the question in front of you. The pattern to internalise is that a rate and a count move together only when the denominator is fixed, and here it is not. Order volume grew 25 percent as well. If the question is about customer experience, the rate is the honest figure; if it is about how many complaint handlers to hire, the count is. Test items usually ask for the one you would not have chosen.

Index numbers: five points is not five percent: A cost index with 2020 = 100 reads 104 in 2021, 111 in 2022 and 109 in 2023. From 2021 to 2023 the index rose 5 points, but the percentage change is $5 / 104 = 4.8$ percent, because the base for the comparison is 104, not 100. From 2022 to 2023 it fell 2 points, which is $-2 / 111 = -1.8$ percent, and note that costs fell even though the index remains 9 percent above the 2020 base. A level and a change are different claims. The only comparison where points and percent coincide is against the base year itself: 2020 to 2023 is 100 to 109, exactly plus 9 percent. Watch also for a rebased series, where a table switches to 2022 = 100 partway down; the two segments cannot be compared directly without converting one of them, and an item that quietly rebases is testing whether you read the column heading.

Joining two tables: totals and per-head figures disagree on purpose: Table 1, headcount by site: Leeds 84, Derby 47, Bristol 129. Table 2, absence days recorded in the same period: Leeds 630, Derby 300, Bristol 903. 'Which site has the worst absence problem?' On raw totals Bristol is worst at 903 days. Normalise per head and the ranking changes: Leeds $630 / 84 = 7.50$ days per employee, Bristol $903 / 129 = 7.00$, Derby $300 / 47 = 6.38$. Leeds is worst, Bristol is merely biggest. The organisation-wide figure is $1,833 / 260 = 7.05$ days per head, which is a useful reference line: Leeds is above it, the other two below. The general rule is that any comparison between units of different size demands a denominator, and the denominator has to come from the other table. Items are built so the raw-total answer and the per-head answer are both on the option list, and so the site with the biggest total is never the site with the highest rate.

The truncated axis: measure the numbers, not the bars: A quarterly satisfaction chart with a y-axis running from 78 to 82 shows bars at 79.2, 79.8, 80.4 and 81.1. Visually the last bar looks several times taller than the first, because only the top 4 points of a 100-point scale are drawn. The actual movement is $81.1 - 79.2 = 1.9$ points, which on the score's own scale is a relative rise of $1.9 / 79.2 = 2.4$ percent. If the question asks 'by approximately what percentage did satisfaction improve', the answer is about 2 percent, and the distractor built from the bar heights will be something like 40 or 400 percent. Related presentation effects to check before answering: a dual-axis chart where two series use different scales and appear to cross meaningfully when they do not; a cumulative series, where a flattening line still means the total is growing, just more slowly; and a logarithmic axis, where equal vertical distances are equal ratios rather than equal amounts. In every case the defence is the same. Find the printed numbers, or read the gridline values, and compute.

Cannot say: revenue is not profit: A product table shows units sold and total revenue. Product P: 4,200 units, 71,400 pounds. Product Q: 1,800 units, 41,400 pounds. Average selling price is $71,400 / 4,200 = 17.00$ for P and $41,400 / 1,800 = 23.00$ for Q, so Q earns more per unit while P earns more in total. Now the statement to evaluate: 'P is more profitable than Q.' The correct response is cannot say. Profit needs cost, and the table has no cost column; a product with a 17 pound price and a 16 pound unit cost is less profitable than one priced at 23 with a cost of 9, and nothing here rules that out. Contrast with 'Q generated more revenue per unit than P', which the table fully supports and which is true. The habit worth building is to finish every cannot-say judgement with the missing input named out loud ('cannot say, because unit cost is not given') because that forces you to distinguish a genuinely unanswerable item from one you simply have not worked hard enough on.

How to practise this skill

- Read the title, the units line, the row and column headers and any footnote before you look at a single value. Roughly the first fifteen seconds of an item should contain no arithmetic at all, and that fifteen

seconds is what prevents the thousand-fold and percentage-point errors.

- For every item, write down which of the three quantities is wanted (share of total, relative change, or change in percentage points), before computing. Most wrong answers on this construct are correct arithmetic applied to the wrong quantity.
- Whenever two groups differ in size, ask what the denominator should be. If a question compares sites, teams, countries or periods of unequal length and you have not divided by something, you are almost certainly answering the wrong question.
- Practise the cannot-say items separately and force yourself to name the missing variable each time. Candidates who train only on computational items reliably over-answer inference statements under time pressure.
- Do not redraw or re-scale charts in your head. Locate the gridline values or the data labels, and if neither exists, interpolate between two labelled gridlines and state the bound rather than guessing a point value.
- Time yourself per item rather than per section. Data interpretation sets share a stimulus, so the first item costs the reading time and the rest should be fast; if item four takes as long as item one, you did not build a mental map of the table.

Glossary

Units multiplier: A scaling note in a table title or column header, such as (£000), (millions) or (per 1,000 population). It applies to every value in scope and is the single most common source of order-of-magnitude errors.

Index number: A series rescaled so a chosen base period equals 100. Changes between two non-base periods must be divided by the earlier value, so a movement in index points is not a percentage change except when measured from the base.

Rebasing: Restating an index against a new base period. Segments of a series with different bases cannot be compared directly, and a table that rebases partway down is testing whether you read the headings.

Truncated axis: A chart whose value axis does not start at zero, which exaggerates the apparent size of differences between bars or points. Legitimate for showing small movements in a large quantity, misleading if read as area or height.

Cumulative series: A line showing a running total rather than each period's value. It can only go up or stay flat, so a flattening cumulative line means the periodic figure is falling, not that the total is.

Weighted average: An average in which each value is multiplied by the size of the group it represents. Averaging two group percentages directly is only correct when the groups are the same size, which in these tables they rarely are.

Normalisation: Dividing a raw figure by an exposure measure (headcount, units sold, population, days open), so groups of different size can be compared. The denominator usually lives in a second table.

Cannot say: The verdict when a statement is neither supported nor contradicted by the data supplied. A correct cannot-say answer can always be defended by naming the specific missing variable.

Study this next

- Data interpretation skill hub: <https://learn.novusstreamsolutions.com/aptitude/skills/data-interpretation>
- Data interpretation practice in the Cognitive Skills Lab:
<https://learn.novusstreamsolutions.com/cognitive-skills/data-interpretation>
- Numerical reasoning lesson:
<https://learn.novusstreamsolutions.com/aptitude/skills/numerical-reasoning/lesson>
- Financial data interpretation suite:
<https://learn.novusstreamsolutions.com/aptitude/tests/financial-data-interpretation>
- Scientific data interpretation suite:

<https://learn.novusstreamsolutions.com/aptitude/tests/scientific-data-interpretation>

- Data interpretation lesson on Novus Learn:

<https://learn.novusstreamsolutions.com/aptitude/skills/data-interpretation/lesson>

Where this material comes from

- Every table, index series and chart described above was constructed for Novus Learn, and each figure was verified by recomputing the totals and the reverse calculation.
- Definitions of index numbers, rebasing and weighted averages cross-checked against standard public references such as the Wikipedia articles 'Index (economics)' and 'Weighted arithmetic mean'. Terminology only; no data or item text is taken from any source.
- Novus Learn aptitude construct registry (catalog seed) for construct scope and suite mapping.
- Public educational framing only: not affiliated with any official exam board, publisher or employer, and no copyrighted test item is reproduced.

Educational preparation only. Novus Learn does not administer official exams and does not guarantee scores or hiring outcomes.

Mechanical comprehension

Forces, motion, gears, pulleys, levers, fluids, pressure, and basic machines.

Mechanical comprehension is the ability to predict what a physical system will do (which way a gear turns, how hard you have to pull, how far the load actually rises) from forces, moments and the simple machines, rather than from a memorised formula sheet. It carries real weight in apprenticeship entry batteries, military technical selection, and screening for maintenance, machine operation and rigging roles. The same reasoning is daily work on site: sizing a jack, choosing a block-and-tackle arrangement, deciding why a belt slips under load but not at idle. Practice here is device-local: no account, and nothing leaves this device unless you export it.

What you should be able to do after this lesson:

1. Balance moments about a pivot (effort x effort arm = load x load arm) and name whether a given tool is a first-, second- or third-class lever.
2. Work out the speed, torque and rotation direction of a driven gear from tooth counts, including what an idler gear does and does not change.
3. Count the rope parts that support a moving block to get the mechanical advantage of a pulley system, then state how much rope must be pulled per metre of lift.
4. Apply pressure = force / area to a hydraulic jack and predict the piston-travel penalty that comes with the force gain.
5. Predict rpm across a belt or chain drive from pulley or sprocket diameters, and say when slip breaks the prediction.
6. State the trade-off that governs every simple machine: multiplied force always costs distance, and friction always makes the delivered advantage smaller than the ideal one.

Worked examples and pitfalls

Levers: balance the moments, then name the class: A wheelbarrow carries 60 kg whose centre of mass sits 0.4 m from the wheel axle, and you grip the handles 1.2 m from that same axle. Take moments about the axle: lift force x 1.2 m = 60 kgf x 0.4 m, so the lift force is $60 \times 0.4 / 1.2 = 20$ kgf, about 196 N. The mechanical advantage is simply the arm ratio, $1.2 / 0.4 = 3$. Two traps live in this item. The first is measuring the load arm from your hands instead of from the fulcrum; the arm is always the perpendicular distance to the pivot, and the pivot here is the wheel contact, not the barrow body. The second is calling it a first-class lever

because that is the one everyone pictures. A wheelbarrow is second class: fulcrum at one end, load in the middle, effort at the far end, which is why its mechanical advantage is always greater than one. Compare tweezers or a pair of tongs, where the effort sits between fulcrum and load: that is third class, mechanical advantage below one, and you are deliberately trading force away to buy speed and control at the tip. Your own forearm lifting a weight is the same arrangement, which is why a 5 kg dumbbell loads the biceps far more than 5 kg.

Gear trains: tooth counts set speed, meshes set direction: A 12-tooth driver turning at 300 rpm meshes directly with a 36-tooth gear. The ratio is driven teeth over driver teeth, $36/12 = 3:1$, so the output turns at $300/3 = 100$ rpm and, ignoring friction, carries about three times the torque. One external mesh reverses rotation, so the output turns opposite to the driver. Now drop a 20-tooth idler between them. Step it through: $300 \times 12/20 = 180$ rpm at the idler, then $180 \times 20/36 = 100$ rpm at the output. The overall ratio is unchanged at 3:1, the idler's tooth count cancels, but there are now two external meshes, two reversals, so the output turns the SAME way as the driver. That is the whole reason idlers are fitted. The tempting wrong answer treats the idler as another reduction stage and reports 180 rpm or some product of both ratios. The check that never fails: only the first and last gear in a simple train affect the ratio, and the direction depends on whether the number of external meshes is odd (reversed) or even (same). An internal or ring mesh, as in a planetary set, does not reverse at all.

Pulleys: count the rope parts that carry the load: A 200 kg load, about 1,962 N. Hung from a single pulley bolted to a beam, the pulley only changes the direction you pull; both rope parts still meet at a fixed axle, the mechanical advantage is 1, and you pull the full 200 kgf. Hang the pulley on the load instead, with one rope end anchored above and the other in your hands, and two rope parts now support the moving block: mechanical advantage 2, effort about 100 kgf or 981 N, but you must pull 2 m of rope for every 1 m the load rises. Build a tackle with four parts supporting the moving block and the effort falls to $200/4 = 50$ kgf, about 490 N, at the cost of 4 m of rope pulled per metre of lift. The trap is counting sheaves instead of supporting parts. A three-sheave arrangement gives three parts if the dead end is made off to the fixed block and four if it is made off to the moving block, and the answer differs by 33 percent. The second trap is treating the figure as delivered force: real sheaves lose a few percent each to bearing and rope friction, so quoted mechanical advantage is the ideal velocity ratio, and the effort you actually feel is higher.

Hydraulics: pressure is shared, force and travel are traded: A jack has a 2 square centimetre input piston and a 50 square centimetre output ram. Push the small piston with 100 N and the pressure in the fluid is $100 / 2 = 50$ N per square centimetre, which is 500 kPa. Pascal's principle says every part of the confined fluid sees that same 500 kPa, so the large ram feels $50 \text{ N/cm}^2 \times 50 \text{ cm}^2 = 2,500$ N. Mechanical advantage 25. Nothing is free: fluid is effectively incompressible, so volume in equals volume out. A 25 cm stroke on the small piston moves $2 \times 25 = 50$ cubic centimetres, and 50 cubic centimetres spread over a 50 square centimetre ram raises it 1 cm. Twenty-five centimetres of pumping buys one centimetre of lift: exactly the force-times-distance bargain a lever strikes. Candidates lose this item by assuming both pistons travel the same distance, or by concluding the jack manufactures energy. A related item asks about a tank: the pressure at the bottom depends on the height of fluid above and its density, not on the width of the vessel, so a narrow 3 m standpipe reads the same bottom pressure as a 3 m deep swimming pool.

Belts and chains: same belt speed, different rpm: A motor pulley 100 mm in diameter runs at 1,450 rpm and drives a 250 mm pulley through an open V-belt. The quantity that is genuinely shared is belt speed, not rpm, so the driven pulley turns at $1,450 \times 100/250 = 580$ rpm. Note the direction of the ratio: with gears you divide by driven teeth, with belts you divide by driven diameter, and the arithmetic looks the same only because both are proportional to circumference. An open belt drives both pulleys the same way; a crossed belt reverses the driven pulley, which is the entire point of the classic crossed-belt diagram item. The trap is slip. A flat or V-belt under an overload slips, so the driven speed falls below 580 rpm while the motor holds 1,450: the symptom of a glazed belt or a weak tensioner. A chain or a toothed timing belt physically cannot slip a tooth without damage, which is why camshaft and indexing drives use them and why a belt-driven answer

and a chain-driven answer to the same question can legitimately differ.

Springs in series and parallel: the answer most people reverse: Two identical coil springs, each 20 N/mm. Mount them side by side, both carrying the same load, and they act in parallel: stiffness adds to 40 N/mm, so a 400 N load compresses the pair $400/40 = 10$ mm. Now stack them end to end, in series. Each spring carries the whole 400 N, force is not split down a chain, so each deflects $400/20 = 20$ mm and the total is 40 mm. The combined stiffness is $400/40 = 10$ N/mm, half of one spring. Most candidates reverse the two because 'series' sounds like the case where things add. Two reliable checks: in parallel the arrangement is always stiffer than one spring, and in series it is always softer than the softest spring. That is also why suspension designers add a helper spring in parallel to raise rate under load, and why a long slender bolt clamps more forgivingly than a short stubby one carrying the same preload.

How to practise this skill

- Sketch the free body and mark the pivot before writing a single number. Mechanical items are lost by choosing the wrong fulcrum far more often than by mis-multiplying.
- Fix the three lever classes to a tool you can picture (pliers and a see-saw for first class, a wheelbarrow and a nutcracker for second, tweezers and your own forearm for third), and say the class out loud before answering.
- Sanity-check every answer against the work rule: if a machine multiplies force it must divide distance. An answer where the effort both moves less and is smaller than the load is wrong somewhere.
- Convert units before you multiply, not after. A moment computed with millimetres on one side and metres on the other is off by a thousand and looks perfectly plausible.
- For any direction question, walk the mesh or the belt with a finger: each external gear pair reverses, an internal or ring mesh does not, an open belt keeps direction and a crossed belt flips it.
- Work untimed until you can explain the reasoning aloud, then add the clock. Attempts are recorded on this device only, so a messy first pass through gear trains costs you nothing but time.

Glossary

Moment (torque): A force multiplied by the perpendicular distance from its line of action to the pivot, in newton-metres. A body is in rotational balance when clockwise moments equal anticlockwise moments about the same point.

Fulcrum: The pivot a lever turns about. Its position relative to the load and the effort defines the lever class and therefore whether the tool multiplies force or multiplies movement.

Mechanical advantage: Output force divided by input force. Ideal mechanical advantage is the geometric figure; actual mechanical advantage is what you get after friction, and it is always lower.

Velocity ratio: Distance moved by the effort divided by distance moved by the load. In a lossless machine it equals the ideal mechanical advantage, which is why force gain always costs travel.

Gear ratio: Driven teeth divided by driver teeth. A ratio above one reduces speed and multiplies torque; only the first and last gear of a simple train affect it.

Idler gear: A gear placed between driver and driven purely to add a mesh. It reverses the output direction and bridges a distance without altering the overall ratio.

Pascal's principle: Pressure applied to a confined fluid is transmitted undiminished to every part of it, which is what lets a small piston at high pressure produce a large force on a bigger ram.

Head: The height of a fluid column above a point. Pressure at depth is density x gravity x height and is independent of the vessel's width or shape.

Study this next

- Mechanical comprehension skill hub:
<https://learn.novusstreamsolutions.com/aptitude/skills/mechanical-comprehension>
- Millwright and industrial mechanic suite:
<https://learn.novusstreamsolutions.com/aptitude/tests/millwright-and-industrial-mechanic>
- Spatial reasoning lesson: <https://learn.novusstreamsolutions.com/aptitude/skills/spatial-reasoning/lesson>
- Skilled trades, maintenance and repair careers:
<https://learn.novusstreamsolutions.com/careers/families/skilled-trades-maintenance-and-repair>
- Encyclopedia search: simple machine:
<https://learn.novusstreamsolutions.com/search?q=Simple%20machine>
- Mechanical comprehension lesson on Novus Learn:
<https://learn.novusstreamsolutions.com/aptitude/skills/mechanical-comprehension/lesson>

Where this material comes from

- Worked examples written for Novus Learn from standard introductory mechanics: moments about a pivot, gear ratios from tooth counts, rope-part counting for tackle, and Pascal's principle.
- Terminology checked against the public Wikipedia articles 'Simple machine', 'Lever', 'Gear train' and 'Pascal's law'. Definitions only; every number and item above is original.
- Novus Learn aptitude construct registry (catalog seed) for the construct scope and related suite mapping.
- Public educational framing only: not affiliated with any official exam board, publisher or employer, and no copyrighted test item is reproduced.

Educational preparation only. Novus Learn does not administer official exams and does not guarantee scores or hiring outcomes.

Safety judgment

Hazard recognition, safe sequencing, escalation, and risk controls.

Practise recognizing hazards and choosing a safe, proportionate response: protect people, control immediate exposure, use the correct reporting path, and verify that the control is effective.

What you should be able to do after this lesson:

1. Separate a hazard from the likelihood, severity, and exposure that shape its risk.
2. Choose an immediate control that stays within the person's training and authority.
3. Recognize when work should pause and when a supervisor, emergency process, or qualified specialist is needed.

Worked examples and pitfalls

Worked scenario: a damaged machine guard: A guard is loose before a scheduled run. Do not operate the machine or improvise a repair beyond your authorization. Keep people away, isolate or label the equipment only as procedure permits, report the defect to the responsible person, and wait for an approved inspection or repair. A deadline does not remove the hazard.

Risk-triage questions: Ask: what can cause harm, who is exposed now, how severe could the outcome be, what control is available, and who has authority to apply it? The safest answer is not always the most dramatic option; it is the option that controls the real exposure without creating a new hazard.

How to practise this skill

- Treat warnings, permits, isolation rules, and personal protective equipment as parts of a system, not interchangeable shortcuts.
- Do not ask an untrained person to investigate a hazard simply because they are nearby.
- In real work, site procedures and qualified safety direction override any general preparation heuristic.

Glossary

Hazard: A source or situation with the potential to cause harm.

Risk: A judgement about possible harm that considers likelihood, severity, and exposure.

Control: A measure that removes a hazard or reduces exposure to it under an approved process.

Study this next

- Occupational safety and health:
<https://learn.novusstreamsolutions.com/search?q=Occupational%20safety%20and%20health>
- Hazard: <https://learn.novusstreamsolutions.com/search?q=Hazard>
- Risk assessment: <https://learn.novusstreamsolutions.com/search?q=Risk%20assessment>
- Safety judgment lesson on Novus Learn:
<https://learn.novusstreamsolutions.com/aptitude/skills/safety-judgement/lesson>

Where this material comes from

- Novus Learn original safety-critical, hazard-recognition, and operational suite scenarios.
- Novus educational framework: identify the hazard, protect people, control exposure, report, and verify.

Educational preparation only. Novus Learn does not administer official exams and does not guarantee scores or hiring outcomes.

Problem solving

Selecting methods, combining information, troubleshooting, and reaching practical solutions.

Problem solving is the construct that asks you to choose a method, not just execute one: to combine a table with a rule, decide whether an estimate settles the question or an exact calculation is needed, narrow a fault by halving the search space, and check the answer against the constraint that actually binds. It shows up across operations, logistics, manufacturing, technician and analyst selection, and in the situational sections of public-safety batteries. It is also the construct where the marking rewards a defensible route as much as a number. Everything you practise stays on this device unless you export it.

What you should be able to do after this lesson:

1. Combine a pricing or specification table with a written brief to reach an answer neither source contains on its own, and find the break-even point where the better option changes.
2. Narrow a fault on a chain of components by bisection, and state the worst-case number of tests before you begin.
3. Solve a working-backwards problem from a known end state, then verify by running the sequence forwards.
4. Decide whether an order-of-magnitude estimate settles a feasibility question or whether the margin is close enough to require exact arithmetic.
5. Identify the binding constraint in a resource-allocation problem and explain why a per-unit heuristic on the non-binding resource gives the wrong answer.
6. Separate a root cause from a symptom, and describe the test that would confirm the cause rather than

merely correlate with it.

Worked examples and pitfalls

Two price structures and the distance where they cross: A delivery of 55 km. Courier A charges 8.00 base plus 0.45 per km. Courier B charges 20.00 flat for the first 40 km, then 0.90 per km beyond that. Which is cheaper? Courier A costs 8 plus 0.45 times 55, which is 8 plus 24.75, giving 32.75. Courier B costs 20 plus 0.90 times 15, which is 20 plus 13.50, giving 33.50. Courier A wins by 0.75. Now the question the item is really testing: is A always cheaper? At 45 km, A costs 8 plus 20.25 which is 28.25, while B costs 20 plus 4.50 which is 24.50. B wins comfortably. So the answer flips somewhere between, and finding where is one line of algebra. Above 40 km, B costs 20 plus 0.9 times distance minus 40, which simplifies to $0.9d - 16$. Set that equal to A's 8 plus 0.45d: 8 plus 0.45d equals 0.9d minus 16, so 24 equals 0.45d, so d is about 53.3 km. Below 53.3 km B is cheaper; above it A is. Check at the crossover: A is 8 plus 24 which is 32.00, and B is 48 minus 16 which is also 32.00. The general lesson is that whichever option has the lower per-unit rate always wins eventually, regardless of the base charges, and the base charges only decide where eventually starts. A single quoted distance never answers a which-is-cheaper question for a fleet.

Halving the search space beats walking the chain: Forty sensors sit on a single daisy-chained cable and exactly one connection is broken, cutting off everything downstream. How many tests do you need in the worst case? Walking the chain from one end and testing each sensor in turn takes up to 40 tests and 20 on average. Test the midpoint instead. If sensor 20 responds, the break is in the upper half and you have eliminated 20 candidates in one measurement; if it does not, the break is below and you have eliminated the other 20. Repeat on the surviving half: 40 becomes 20, then 10, then 5, then 3, then 2, then 1. Six tests, worst case, because each test halves what is left and two to the power six is 64, comfortably more than 40. The saving grows as the problem grows: 1,000 candidates need only ten tests. Two conditions have to hold for bisection to be valid, and stating them is part of the answer. The fault must be monotone, meaning everything on one side of the break behaves differently from everything on the other; and a single test at any point must tell you which side you are on. When there are two independent breaks, or when a test is only meaningful at the ends, bisection does not apply and a different strategy is needed. Candidates who reach for bisection reflexively on a problem that fails those conditions lose more than they save.

Working backwards from the number you were given: A team started a quarter with an unknown budget. In week one it spent half of it. In week two it spent 400 of what remained. In week three it spent a third of what remained after that. It finished with 1,200. How much did it start with? Attempting this forwards means carrying an unknown through three operations. Backwards it is arithmetic. After week three, 1,200 is what is left having spent a third, so 1,200 represents two thirds of the week-three opening balance, which was therefore 1,800. Before week two's spend of 400, the balance was 1,800 plus 400, which is 2,200. That 2,200 is what remained after spending half, so the starting budget was 4,400. Verify forwards, always: 4,400 less half is 2,200; less 400 is 1,800; less a third of 1,800, which is 600, leaves 1,200. Correct. The mechanical rule is to invert each operation and apply the inversions in reverse order. The inverse of spending a third is dividing by two thirds, not multiplying by three, and that specific slip is the most common wrong answer in this family. Working backwards is the right tool whenever the end state is known exactly and the operations are individually invertible, which covers most budget, mixture and journey problems phrased as how much did it start with.

When an estimate is enough, and when it is not: A maintenance window is three hours. Two technicians must service 1,850 units, each taking about 4.5 minutes, plus a shared 20-minute setup and 15-minute teardown. Feasible? Estimate first: 1,850 units at 4.5 minutes is 8,325 technician-minutes; split between two people that is about 4,163 minutes, which is roughly 69 hours. The window is 3 hours. The answer is no by a factor of more than twenty, and no refinement of the setup and teardown figures could change it. Precision here would be wasted effort. Now the same problem with 90 units. Ninety at 4.5 minutes is 405 technician-minutes, halved to 202.5 minutes, plus 35 minutes of setup and teardown, giving 237.5 minutes.

That is 3 hours and 57 and a half minutes against a 3-hour window. Still no, but only just, and now every assumption matters: whether the technicians can genuinely work in parallel, whether setup is shared or duplicated, whether 4.5 minutes is a mean or a best case. The judgement being assessed is knowing which regime you are in. If your rough figure misses the target by an order of magnitude, stop and answer. If it lands within a factor of two, the estimate has not settled anything and you must do the exact arithmetic and name your assumptions.

The binding constraint decides, and it is rarely the obvious one: A van carries at most 900 kg and at most 6 cubic metres. Pallet type X weighs 150 kg, occupies 0.8 cubic metres and is worth 200. Pallet type Y weighs 90 kg, occupies 1.4 cubic metres and is worth 260. What mix maximises value? The instinctive heuristic is value per kilogram: X gives 200 over 150, about 1.33, while Y gives 260 over 90, about 2.89, so load Y. That is wrong, and the reason is that weight is not what runs out. Value per cubic metre tells the opposite story: X gives 250 and Y gives about 186. Enumerate the feasible whole-pallet mixes. Six X uses 900 kg and 4.8 cubic metres, worth 1,200. Five X and one Y uses 840 kg and 5.4 cubic metres, worth 1,260. Four X and two Y uses 780 kg and exactly 6.0 cubic metres, worth 1,320. Three X and three Y needs 6.6 cubic metres and does not fit. Two X and three Y uses 570 kg and 5.8 cubic metres, worth 1,180. Four Y alone uses 5.6 cubic metres and is worth 1,040. The best mix is four X and two Y at 1,320. Look at what binds it: volume is used to the last cubic metre while 120 kg of payload goes unused. Once volume is identified as the binding constraint, X's superior value per cubic metre explains the X-heavy answer, and the spare weight explains why two Y still get on board. The transferable move is to compute the usage of every constraint at your proposed answer and see which one is exhausted; a per-unit ratio computed against a non-binding resource is worse than no heuristic at all.

Cause, symptom, and the test that tells them apart: A pump trips out twice a shift. The obvious fix is to reset it, which works for about four hours. Ask why once and you find the thermal overload relay is tripping. Ask again and the motor is running hot. Again and the bearing is running hot. Again and the grease has dried out. Again and the lubrication interval was set for a duty cycle far lighter than the pump has actually been running since the line was rebalanced last year. Each level supports a different fix: resetting the trip costs nothing and lasts four hours; replacing the relay costs a part and lasts until the bearing seizes; regreasing lasts weeks; changing the lubrication schedule to match the real duty cycle is the one that stops the fault recurring. Stop asking why when the next answer stops being something you can act on, or when it leaves the boundary of the system you can change. There is one discipline that separates this from storytelling. A correlation is a lead, not a cause: the fact that the trips started after the line was rebalanced is suggestive, not proof. The confirming test is to intervene and observe: restore the original duty cycle, or apply the corrected greasing interval to this pump and not to the identical one on the parallel line, and see which one trips. If a proposed cause cannot be tested by changing it, treat it as a hypothesis and say so rather than closing the investigation.

How to practise this skill

- Write the question you are actually being asked at the top of your working, in your own words. Problem-solving items usually supply more information than the question needs, and the commonest failure is a correct calculation of something nobody asked for.
- Before any comparison of two options, ask where they cross. One quoted quantity never settles which is cheaper or faster, and the break-even is usually a single line of algebra away.
- State the worst-case number of steps before starting a search or a fault trace. If bisection applies, say so and say why; if the fault is not monotone or a test only works at the ends, say that instead and pick a different method.
- Estimate before you compute, then decide explicitly which regime you are in. Missing the target by an order of magnitude ends the question; landing within a factor of two means the estimate proved nothing and the exact numbers are now compulsory.
- Finish every allocation or capacity problem by listing how much of each constraint your answer consumes.

The constraint you exhaust is the one that explains the answer, and the one you do not is where a per-unit heuristic will mislead you.

- Always verify a backwards solution by running it forwards. Attempts stay on this device, so the extra minute costs nothing and catches the inverted-fraction error that makes this family's most attractive wrong answer.

Glossary

Break-even point: The value at which two options cost or take the same. Below it one option wins and above it the other does, which is why a single quoted case never answers a comparison question.

Binding constraint: The limit that is fully consumed by the chosen solution. Identifying it explains the answer and reveals which per-unit ratios are relevant and which are noise.

Slack: The unused portion of a constraint at a proposed solution. Spare capacity in one resource is what lets a mix include items that look inefficient on the binding resource.

Bisection: Halving a search space with each test. It reduces a search of n candidates to about $\log$ base two of n tests, but only where the fault is monotone along the chain.

Working backwards: Solving from a known end state by inverting each operation in reverse order. The reliable method whenever the final value is given and every step is individually invertible.

Order-of-magnitude estimate: A rough calculation intended to settle feasibility rather than produce a figure. It is decisive when the answer is off by a factor of ten and useless when the margin is within a factor of two.

Root cause: The condition whose removal stops the fault recurring, as opposed to a symptom whose treatment suppresses it temporarily. A candidate cause that cannot be tested by changing it remains a hypothesis.

Monotone fault: A fault where everything on one side of a break behaves differently from everything on the other. The precondition that makes bisection valid, and the assumption most often left unstated.

Study this next

- Problem solving skill hub: <https://learn.novusstreamsolutions.com/aptitude/skills/problem-solving>
- Problem solving practice bank: <https://learn.novusstreamsolutions.com/cognitive-skills/problem-solving>
- Diagrammatic reasoning lesson: <https://learn.novusstreamsolutions.com/aptitude/skills/diagrammatic-reasoning/lesson>
- Maintenance technician suite: <https://learn.novusstreamsolutions.com/aptitude/tests/maintenance-technician>
- Supply chain planner suite: <https://learn.novusstreamsolutions.com/aptitude/tests/supply-chain-planner>
- Problem solving lesson on Novus Learn: <https://learn.novusstreamsolutions.com/aptitude/skills/problem-solving/lesson>

Where this material comes from

- Every figure above was computed and checked for Novus Learn: the courier break-even at about 53.3 km, the six-test bisection bound for 40 candidates, the 4,400 starting budget verified forwards, and the four-X-two-Y loading enumerated across all feasible whole-pallet mixes.
- Terminology checked against the public Wikipedia articles 'Binary search algorithm', 'Break-even', 'Shadow price' and 'Root cause analysis'. Definitions only; every scenario above is original.
- Novus Learn aptitude construct registry (catalog seed) for the construct scope and related suite mapping.
- Public educational framing only: not affiliated with any official exam board, publisher or employer, and no copyrighted test item is reproduced.

Educational preparation only. Novus Learn does not administer official exams and does not guarantee scores or hiring outcomes.

Recommended preparation

- Practice technical reasoning:
<https://learn.novusstreamsolutions.com/aptitude/skills/technical-reasoning/lesson>
- Build numerical reasoning:
<https://learn.novusstreamsolutions.com/aptitude/skills/numerical-reasoning/lesson>
- Practice safety judgement: <https://learn.novusstreamsolutions.com/aptitude/skills/safety-judgement/lesson>

Answer keys, scoring and privacy

Answer keys and scoring logic stay server-side and are never included in any download or export.

Formal suite answer keys and scoring logic stay on the server and are not part of any download, in any format. Downloadable keys exist only for the open, untimed practice material (the practice packs on the puzzles, cognitive-skills and reasoning practice lab pages), where the answers are already public teaching content.

Novus Learn needs no account. Your practice history lives in this browser's storage on this device, and is sent to a server only if you create an account and switch on backup. This file contains no attempt, session or result link, so it is safe to share.

Private practice result. Not an official exam certificate, employer decision, hiring signal, admissions decision, or guaranteed outcome. Scores stay on this device unless you export them.

Answer keys and scoring logic stay server-side and are never included in any download or export.

Source: <https://learn.novusstreamsolutions.com/aptitude/tests/manufacturing-technician/study-outline>