

Blueprint and plan reading: study guide

Construction and infrastructure · suite apt-178-blueprint-and-plan-reading · generated 2026-09-15T15:38:53.408Z

Detail	Value
Title	Blueprint and plan reading: study guide
Generated	2026-09-15T15:38:53.408Z
Fixture/version	apt-178-blueprint-and-plan-reading
Sector	Construction and infrastructure
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 Blueprint and plan reading

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-178-blueprint-and-plan-reading&mode=guided>
- Start mini-test:
<https://learn.novusstreamsolutions.com/aptitude/practice/new?bank=apt-178-blueprint-and-plan-reading&mode=mini>
- Start full simulation:
<https://learn.novusstreamsolutions.com/aptitude/practice/new?bank=apt-178-blueprint-and-plan-reading&mode=full>

Skills covered, in full

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

Spatial reasoning

Rotation, folding, views, orientation, maps, and three-dimensional visualization.

Spatial reasoning is the ability to turn an object in your head and be right about the result: to tell a rotation from a mirror image, fold a flat net into a solid and know which faces meet, read plan and elevation views back into a three-dimensional shape, and hold your own orientation steady while a route turns underneath you. It is a heavily weighted construct in pilot and aircrew selection, firefighter and military technical batteries, architecture and design admissions, and trades screening where a drawing has to be built. Much of it can be trained by replacing imagination with a checkable procedure. Practice is device-local unless you export it.

What you should be able to do after this lesson:

1. Distinguish a rotation from a reflection by checking whether the clockwise order of three marked features is preserved or reversed.
2. Read a cube net and name which faces end up opposite, using the rule that squares two apart in a straight strip are opposite.
3. Reconstruct a solid from plan, front and side views, count its unit cubes, and identify when the views leave the count undetermined.
4. Predict the hole pattern produced by unfolding a punched sheet, treating each unfold as a reflection across the fold line rather than a rotation.
5. Work with bearings and relative direction, converting between a map frame and a traveller frame without losing track of which is which.
6. Name the cross-section produced by cutting a cylinder, cone or cube on a given plane, including the oblique cuts people default to calling circles.

Worked examples and pitfalls

Rotation preserves handedness; reflection reverses it: A figure carries three distinguishable marks (say a red band, a blue band and a yellow band), which read red, blue, yellow going clockwise around its centre. Rotating that figure in the plane, by any angle, never changes that clockwise order: it still reads red, blue, yellow clockwise from whatever mark you start at. Mirroring it reverses the order, so a reflected copy reads red, yellow, blue clockwise. That single check answers the whole is-this-a-rotation-or-a-reflection family without imagining any motion. It works on letters, which is why no in-plane rotation of the letter R produces a backwards R, and no rotation of a lower-case b produces a d: b and d are mirror images across a vertical line, and turning one will never give you the other. Contrast b and q, which are related by a half turn and therefore genuinely are rotations of each other. The distinction is exactly what the clockwise-order test measures. In three dimensions the equivalent is handedness at a corner: pick three edges meeting at a vertex, and a rotation preserves whether they form a right-handed or left-handed set while a reflection swaps it. This is exactly why most Western dice, viewed at the corner where the 1, 2 and 3 faces meet, show those numbers running anticlockwise, and a die showing them clockwise is a mirror-image die rather than a differently rotated one. The one honest limitation: a figure with its own mirror line looks identical after reflection, so the test decides nothing for symmetric shapes, which is precisely why item writers use asymmetric ones.

Cube nets: two apart in a strip means opposite: A net made of six squares. Four of them, call them A, B, C and D, sit in a horizontal strip left to right. Square E sits directly above B, and square F sits directly below B. Fold it. The strip of four wraps around the cube as a band, so squares two apart in that strip end up on opposite faces: A is opposite C, and B is opposite D. E and F are the remaining pair, and they become the

top and bottom, opposite each other. The two rules that solve almost every net item are exactly these. Squares separated by one square in a straight line end up opposite. Squares sharing an edge on the net end up adjacent on the cube, and can never be opposite. Everything else is applying them carefully. A die adds one more constraint worth carrying: opposite faces of a standard die sum to seven, so 1 faces 6, 2 faces 5 and 3 faces 4, and a net showing 1 and 6 two apart in a strip is consistent while a net showing 1 and 2 two apart is not a standard die. The classic error on these items is trying to fold the net in your head all at once. Do not. Pick one square as the base, identify its opposite by the two-apart rule, then repeat for a second pair, and you have determined the cube in two cheap steps rather than one expensive act of visualisation.

Three views, one solid, and the cubes you cannot see: A solid is built from unit cubes on a base one cube deep and three cubes wide. The column heights, left to right, are 3, 1 and 2. The front view is therefore a staircase profile of heights 3, 1, 2. The side view, looking along the row from the right, is a single column three cubes tall, because the tallest column hides the others behind it. The plan view is a three-by-one rectangle. Total cubes: 3 plus 1 plus 2, which is 6, and here the three views pin it down exactly because the plan view tells you the solid is only one cube deep. Now change one thing. Suppose the plan view is three by two, so the solid is two cubes deep. The front view can still read 3, 1, 2 while the count is anything from 6, one row of cubes with a hollow back, up to 12, a solid block matching the front profile in both rows. The views constrain the shape but do not determine it, and a well-written item will either supply the count or ask for the minimum and maximum rather than a single figure. That is the trap in this family: candidates who assume every hidden position is filled will over-count, and candidates who assume nothing is hidden will under-count. Read the plan view first to establish the footprint, then use the elevations to cap each column, then ask explicitly whether the item wants the minimum, the maximum, or a determined value.

Punched and unfolded, with coordinates: Take a square sheet 8 units across, with the origin at its centre so the corners are at plus or minus 4 in each direction. Fold the right half onto the left half across the vertical centre line, then fold the bottom half up onto the top half across the horizontal centre line. What is left visible is one quarter, the region where x runs from minus 4 to 0 and y runs from 0 to 4. Punch a single hole at the point minus 3, 3. Now unfold, one step at a time, and treat every unfold as a reflection. Undoing the horizontal fold reflects across the line y equals 0 and adds a hole at minus 3, minus 3. Undoing the vertical fold reflects both existing holes across x equals 0 and adds holes at 3, 3 and 3, minus 3. The finished sheet has four holes, at all four combinations of plus or minus 3. Two rules generalise from this. The number of holes is one punch multiplied by two for each fold the punch passed through, so two folds give four holes and three folds give eight, provided the punch went through every layer, which it does when the folds are through the middle. And the positions are mirror images across each fold line, never rotations. That second point is where the marks go: candidates who place the fourth hole by rotational symmetry get three holes right and one wrong, which is exactly the answer option the item includes.

Bearings, and the difference between the map and the walker: A walker leaves a gate, goes 300 m due north, then 400 m due east. The straight-line distance back to the gate is the hypotenuse of a 300 by 400 right triangle, which is 500 m: the 3-4-5 triangle, so no calculator is needed. The bearing from the gate to the walker is measured clockwise from north, and it is the angle whose tangent is 400 over 300, which is about 53 degrees, so the bearing is 053. The return bearing is that plus 180, which is 233. Bearings are always three digits and always clockwise from north, so 53 degrees is written 053 and a bearing of 350 is very slightly west of north, not south. Now the other half of this construct, which trips up more people than the trigonometry. You are facing south and you turn left. Which way are you now facing? East. Your left hand points east when you face south, and the standard wrong answer, west, comes from reading left as left-on-the-map while forgetting that the map is drawn north-up and you are not. The general fix is to state your heading before every turn and rotate your frame with the traveller: facing north, left is west; facing south, left is east; facing east, left is north; facing west, left is south. On route items, physically rotate the page so the direction of travel points up before working out a turn, and rotate it back only to read off the final compass direction.

What a cut actually reveals: Cross-section items ask what shape a plane makes as it slices a solid, and the reflex answer of circle is wrong most of the time. A cylinder cut perpendicular to its axis gives a circle. Cut parallel to the axis, it gives a rectangle. Cut obliquely so the plane crosses only the curved surface, it gives an ellipse; cut obliquely so the plane also crosses one flat end, it gives an ellipse with one side sliced flat, a D-shape. A cone is richer, because its sections are the conic sections that give the family its name: perpendicular to the axis gives a circle; an oblique cut that crosses every side of the cone gives an ellipse; a cut parallel to one slanted side gives a parabola; a cut parallel to the axis gives a hyperbola branch; and a cut through the apex gives a triangle. A cube is the one that surprises people. A cut parallel to a face gives a square, obviously. A cut through three vertices that are mutually adjacent to one corner gives an equilateral triangle. And a cut through the midpoints of six edges, perpendicular to a long diagonal of the cube, gives a regular hexagon: a six-sided section from a six-faced solid with no curves anywhere. For a flat-faced solid, work these by counting faces rather than by visualising: the section of a convex polyhedron gains exactly one straight edge for every face the plane passes through, so a cut crossing six faces of a cube must close into a six-sided figure and a cut touching only three cannot have more than three sides. For curved solids the question becomes whether the plane meets the curved surface all the way round, which gives a closed circle or ellipse, or only partway, which leaves an open arc closed off by a straight chord.

How to practise this skill

- Replace visualisation with a rule wherever one exists. Two-apart-means-opposite on nets, and clockwise-order-preserved for rotations, both give a checkable answer in seconds and are reliable when mental imagery gets tired late in a section.
- Physically rotate the paper or the device on route and orientation items so the direction of travel points up, work out the turn, then rotate back to read the compass direction. Fighting the north-up convention in your head is the single largest source of left-right errors.
- On view and cube-count items, read the plan view first to fix the footprint, then use the elevations to cap each column, then decide explicitly whether the question wants a minimum, a maximum, or a determined count.
- Treat every unfold as a reflection across the fold line. Say the word reflection out loud as you draw each new hole; the rotational version of the pattern is a deliberately included wrong answer in almost every folding item.
- Build a physical reference once. Fold a paper cube net, mark the faces, and hold it: ten minutes with a real cube fixes the opposite-face rule more firmly than an hour of diagrams, and the memory transfers to items you have never seen.
- Practise untimed first and check whether your errors cluster in rotation, folding, views or orientation, because the four subskills train differently. Attempts are stored on this device only, so an honest error log costs nothing and is what tells you which drill to repeat.

Glossary

Chirality (handedness): The property that distinguishes a figure from its mirror image. Rotation preserves it and reflection reverses it, which is what makes the clockwise-order test decisive for asymmetric figures.

Net: A flat arrangement of faces that folds into a solid. On a cube net, squares two apart in a straight strip become opposite faces and squares sharing an edge become adjacent ones.

Plan view: The view looking straight down on an object. It fixes the footprint of a stacked-cube solid, which is why it is the first view to read.

Elevation: A view looking horizontally at an object, usually given as front and side. Each elevation caps the heights it shows but hides anything directly behind.

Orthographic projection: A drawing convention showing an object through separate parallel views with no perspective. Together the views constrain a solid but do not always determine it.

Bearing: A direction measured clockwise from north and written with three digits, so due east is 090 and a heading slightly west of north is 350. The reverse bearing differs by 180 degrees.

Egocentric versus allocentric frame: The difference between directions relative to a traveller who is turning and directions relative to a fixed map. Left means different compass headings in the two frames, which is why headings must be restated before every turn.

Cross-section: The shape revealed where a plane cuts a solid. It gains one straight edge for each flat face the plane crosses and one curve for each curved surface, which turns the question into a count.

Study this next

- Spatial reasoning skill hub: <https://learn.novusstreamsolutions.com/aptitude/skills/spatial-reasoning>
- Spatial reasoning practice bank: <https://learn.novusstreamsolutions.com/cognitive-skills/spatial-reasoning>
- Mechanical comprehension lesson:
<https://learn.novusstreamsolutions.com/aptitude/skills/mechanical-comprehension/lesson>
- Firefighter spatial reasoning suite:
<https://learn.novusstreamsolutions.com/aptitude/tests/firefighter-spatial-reasoning>
- Architecture and design aptitude suite:
<https://learn.novusstreamsolutions.com/aptitude/tests/architecture-and-design-aptitude>
- Spatial reasoning lesson on Novus Learn:
<https://learn.novusstreamsolutions.com/aptitude/skills/spatial-reasoning/lesson>

Where this material comes from

- Worked items written for Novus Learn from standard introductory geometry: cube-net opposite-face rules, orthographic view reconstruction, reflection across fold lines, right-triangle distances and bearings, and the conic and polyhedral cross-sections. The 3-4-5 distance, the 053 and 233 bearings, and the four hole coordinates were each computed and checked.
- Terminology checked against the public Wikipedia articles 'Net (polyhedron)', 'Multiview orthographic projection', 'Chirality', 'Bearing (navigation)' and 'Conic section'. Definitions only; every figure described 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.

Attention and concentration

Sustained focus, selective attention, and accuracy under time pressure.

Attention and concentration is the skill of still noticing on minute forty of a checking block as reliably as you did on minute two, and of pointing the noticing at the right thing when two things compete. It has three distinguishable parts - selective attention, sustained attention or vigilance, and divided attention - and assessments load them differently: cancellation and checking tasks load vigilance, conflict tasks load selection, and dispatch-style monitoring loads division. It is the construct that decides whether a records clerk, a dispatcher, an air-side controller or a quality inspector catches the one wrong digit in a shift. Practice here is device-local, with no account and nothing uploaded.

What you should be able to do after this lesson:

1. Count occurrences of a target in a dense string without double-counting or losing your place, using a fixed scan path rather than free looking.

2. Separate hits, misses and false alarms in your own results and work out whether you are set too eager or too cautious for the scoring rule in force.
3. Predict where in a long block your error rate will rise, and schedule micro-resets against that curve instead of pushing through it.
4. Recognise a transposition error, the class that survives a same-characters gist check and is therefore missed most often.
5. Explain why an automatic process such as reading interferes with a controlled one, and use that to anticipate which distractors will actually cost you time.
6. Run a two-stream monitoring task by alternating on a fixed cadence rather than attempting genuine simultaneity.

Worked examples and pitfalls

A cancellation count you can actually check: Count every 7 in this row: 4 7 1 7 7 3 9 7 2 8 7 5 7 6 0 7.

Working left to right and tapping once per hit: hits at the second, fourth, fifth, eighth, eleventh, thirteenth and sixteenth positions. That is seven sevens. Two error modes produce nearly all the wrong answers. The first is the adjacent pair - the 7 7 at positions four and five gets counted once, because the eye takes a repeated character as one perceptual object. The second is losing the place after the 9, where the visually similar 9 and 7 force a moment of re-checking and the scan restarts a character early, producing an over-count of eight. The defence for both is a fixed scan path with a physical anchor: a fingertip or cursor moving one character at a time, never jumping back. Re-scanning to check is the thing that creates the double-count, so if you must verify, verify by counting a second time from the RIGHT-hand end and comparing totals, never by re-reading part of the row.

Hits, misses and false alarms: the eager candidate loses: A checking batch contains 300 records, 40 of which are genuinely faulty. Candidate A flags 46 records and 36 of them are truly faulty. Hits 36, misses 40 minus 36 equals 4, false alarms 46 minus 36 equals 10. Candidate B flags 38 and 34 are truly faulty: hits 34, misses 6, false alarms 4. On hit rate alone A wins, 36 out of 40 which is 90 percent against B's 34 out of 40 which is 85 percent. Now apply the scoring rule that most checking tasks actually use, where a false alarm cancels a hit: A scores 36 minus 10 equals 26, B scores 34 minus 4 equals 30. B wins by four despite catching two fewer faults. This is why 'flag anything that looks odd' is bad advice on a scored checking task, and why the first thing to read on a checking item is whether wrong flags are penalised. Your response criterion - how much evidence you demand before flagging - is adjustable, and it should be set from the scoring rule, not from your temperament.

The transposition that passes every gist check: Compare these two lines and decide whether they match. Invoice 4820-7391-06. Invoice 4820-7931-06. They do not: the middle group reads 7391 in the first and 7931 in the second, with the 3 and the 9 swapped. Transpositions are the most-missed error class in record checking for a structural reason - the character SET is identical, the length is identical, the first and last characters of the group are identical, so every fast check the visual system runs comes back clean. Substitutions and omissions change the character inventory and get caught; transpositions do not. Two habits raise the catch rate. Read digits in fixed groups of two rather than as a whole number, so 73-91 against 79-31 becomes a mismatch at the first group instead of a subtle difference somewhere in a four-digit blur. And check groups in a deliberately non-natural order - last group, first group, middle group - because reading left to right lets the confirmation you built at the start carry you through the middle, which is exactly where the swap is usually planted.

Conflict: why reading fights you: The classic demonstration is Stroop's, published in 1935: the word RED printed in blue ink, with the instruction to name the ink colour. Naming takes measurably longer, and errors go up, because reading a familiar word is automatic and cannot be switched off, so the automatic response has to be suppressed before the controlled one can be produced. The same conflict has a numeric version you can test on yourself in a second: how many characters are in the string 4 4 4? The answer is three, and

the digit 4 pulls at you the entire way. In an assessment this appears wherever the salient feature and the asked-for feature come apart - a chart where the tallest bar is not the answer to the question, a form where the highlighted field is not the one being verified, a row where the bold total is not what the stem requested. The practical move is to name the target feature out loud before you look - 'ink colour', 'character count', 'the value for March' - because pre-loading the target biases selection before the automatic reading response gets a chance to win.

Where the errors actually appear in a 45-minute block: Errors in a long checking block are not spread evenly. Mackworth's 1948 clock-watching study established the pattern that gives the effect its name: detection declines over a prolonged watch, with the sharpest deterioration early rather than at the very end. Practically, on a 45-minute self-timed checking block, expect your per-minute error rate in minutes 20 to 45 to run visibly above minutes 1 to 20 even though nothing about the material changed, and expect the subjective sense of effort to lag the actual decline, so it will not feel like you are getting worse. Two things work against it. Break the block into three fifteen-minute segments with a ten-second reset between them - look away, unfocus, breathe out - which costs thirty seconds of a 45-minute block, about one percent of the time, and buys back more than that in caught errors. And score your practice by segment rather than as one number, because a single overall accuracy figure hides exactly the information you need: whether your problem is skill, which shows as flat error rate, or endurance, which shows as a rising one.

Two streams: alternate on a cadence, do not try to merge: A dispatch-style monitoring task: keep a running total of the numbers announced on channel one while watching channel two for the code word AMBER. Genuine simultaneity is not available - the two tasks compete for the same control resource - so the choice is not whether to alternate but whether to alternate deliberately or accidentally. Deliberate looks like this: fix the arithmetic to a rhythm, updating the total only at each announcement and holding a single number between updates, which frees the gaps for channel two. If the total is 34 and the next announcement is 7, you spend under a second reaching 41 and then you are free again. Accidental looks like re-deriving the running total from the beginning because you did not trust it, which locks up the whole window and is when AMBER goes past unheard. The measurable failure of divided attention is almost never a failure to hear the target; it is a failure to have any spare capacity at the moment it arrived. The related phenomenon worth knowing is inattention blindness, illustrated by Simons and Chabris in 1999: an unexpected and perfectly visible event is missed entirely when attention is committed to a counting task.

How to practise this skill

- Fix your scan path before you start, and never re-scan backwards to double-check. Backward re-scanning is what produces both double-counts and lost places; if you need to verify a count, run it again from the other end and compare.
- Read the scoring rule before the first item. Whether false alarms are penalised changes the correct strategy completely, and a criterion set for the wrong rule costs more marks than slow reading does.
- Check numeric strings in fixed groups of two or three and in a deliberately shuffled group order. Transpositions are the errors that survive whole-string comparison, and they are the ones planted on purpose.
- Time your practice in segments and log accuracy per segment, not per session. A flat error rate means you need speed; a rising one means you need endurance and breaks, and the two problems have opposite fixes.
- Name the target feature out loud before each conflict item. Pre-loading the relevant dimension is the cheapest available defence against the salient-but-wrong answer.
- Practise with an actual distractor present rather than in silence. A test hall has movement, coughing and page-turning, and attention practice in perfect quiet trains a condition you will not meet.

Glossary

Selective attention: Prioritising one source or feature while suppressing competing ones. It is what fails when the visually salient element of a display captures the response instead of the requested one.

Sustained attention (vigilance): Maintaining detection performance over a long, low-event period. It is the capacity that checking, monitoring and inspection tasks are really sampling.

Vigilance decrement: The decline in detection performance over a prolonged watch, typically steepest early in the block. Mackworth's 1948 clock test is the standard public reference.

Divided attention: Handling two streams that compete for control. Performance is best modelled as fast alternation with a switch cost, not as genuine parallelism.

Hit, miss, false alarm, correct rejection: The four outcomes of any detection decision: flagging a real fault, missing one, flagging a clean record, and correctly leaving a clean record alone. Any honest accuracy claim needs at least the first three.

Response criterion: How much evidence you require before flagging. Shifting it trades misses against false alarms without changing your underlying sensitivity, which is why it must be set from the scoring rule.

Stroop interference: The slowing that occurs when an automatic response, such as reading a word, conflicts with the requested response, such as naming its ink colour. Named for Stroop's 1935 experiments.

Inattentional blindness: Failing to notice a fully visible, unexpected event because attention is committed elsewhere - the effect Simons and Chabris demonstrated in 1999 with a counting task.

Study this next

- Attention and concentration skill hub:
<https://learn.novusstreamsolutions.com/aptitude/skills/attention-concentration>
- Cognitive Skills Lab: attention practice:
<https://learn.novusstreamsolutions.com/cognitive-skills/attention-concentration>
- Emergency dispatch and 911 communications suite:
<https://learn.novusstreamsolutions.com/aptitude/tests/emergency-dispatch-and-911-communications>
- Error detection lesson: <https://learn.novusstreamsolutions.com/aptitude/skills/error-detection/lesson>
- Perceptual speed lesson: <https://learn.novusstreamsolutions.com/aptitude/skills/perceptual-speed/lesson>
- Attention and concentration lesson on Novus Learn:
<https://learn.novusstreamsolutions.com/aptitude/skills/attention-concentration/lesson>

Where this material comes from

- Character strings, invoice comparisons, batch counts and monitoring scenarios above are written for Novus Learn. All figures are original and no published test item is reproduced.
- Terminology and directions of effect follow widely published work: Mackworth (1948) on the vigilance decrement, Stroop (1935) on response conflict, Simons and Chabris (1999) on inattentional blindness, and the standard hit/miss/false-alarm framing from signal detection theory. Concepts only; no result figures are attributed to those studies.
- 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 not a clinical, diagnostic or attention-disorder screening measure.

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

Numerical reasoning

Arithmetic, fractions, percentages, ratios, rates, estimation, word problems, and number relationships.

Numerical reasoning is the ability to take a small set of supplied numbers (a price list, a staffing table, a fuel figure), and reach a defensible answer in around ninety seconds without a formula sheet. It is the most widely used quantitative section in graduate, management and public-sector screening, and it appears in banking, retail, armed forces and healthcare entry batteries alike. The arithmetic itself is deliberately ordinary: percentages, ratios, rates and averages. What is actually being measured is whether you pick the right operation quickly, keep the units straight, and answer the question that was asked rather than the one you started computing. 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. Move between fractions, decimals, percentages and ratios on sight, and pick whichever form makes a given item fastest (12.5 percent as one eighth, 33.3 percent as one third).
2. Apply percentage change in both directions, including recovering an original figure from a post-increase total, and combine successive changes multiplicatively rather than by adding them.
3. Split a total in a stated ratio by counting parts first, and work backwards from a stated difference between two shares.
4. Solve rate problems (speed, unit price, consumption, combined work rates) by carrying units through every intermediate line so the units of the answer prove the method.
5. Produce a one-significant-figure estimate before computing, and use it to eliminate the distractors built from the common wrong operations.
6. Read a multi-step stem and name, in one sentence, the single quantity the final question asks for before touching the numbers.

Worked examples and pitfalls

Reverse percentages: the item most candidates get backwards: A subscription price rose by 15 percent and now stands at 57.50 pounds. What was it before? The correct move is to divide, not subtract: the new price is 115 percent of the old, so the old price is $57.50 / 1.15 = 50.00$. Check it forwards: 15 percent of 50 is 7.50, and $50 + 7.50 = 57.50$. The tempting wrong answer takes 15 percent of the NEW figure, $0.15 \times 57.50 = 8.625$, and reports 48.88. That distractor is always on the option list because it is what most people do under time pressure, and it is wrong by 2.25 percent, small enough to look plausible. The same asymmetry drives the successive-change item: a price that rises 20 percent and then falls 20 percent does not return to where it started. Multiply the factors: $1.20 \times 0.80 = 0.96$, a net fall of 4 percent. Starting at 500 pounds you get 600 then 480, not 500. Percentages compose by multiplication; they never add.

Ratio splits: count the parts before you divide: A 4,830 pound budget is split between three teams in the ratio 3:5:6. Add the parts first: $3 + 5 + 6 = 14$. One part is $4,830 / 14 = 345$. The shares are $3 \times 345 = 1,035$, $5 \times 345 = 1,725$ and $6 \times 345 = 2,070$, and they sum back to 4,830, which is the check you should always run. Two classic failures. The first is dividing by 3 because there are three teams. That gives 1,610 each and ignores the ratio entirely. The second is treating 5 as a fraction and computing five sixths or five fourteenths of something other than the total. The more interesting version of this item gives you a difference instead of the total: 'the second team receives 690 pounds more than the first, what is the whole budget?' The difference between the shares is $5 - 3 = 2$ parts, so one part is $690 / 2 = 345$, and the budget is $14 \times 345 = 4,830$. Same number, reached from the other end, and the arithmetic is trivial once you have made the parts explicit.

Rates and units: 2 h 15 min is 2.25, not 2.15: A delivery van covers 174 km in 2 hours 15 minutes. Average speed is distance over time, and the time must be in hours: 15 minutes is $15/60 = 0.25$ h, so $174 / 2.25 = 77.3$ km/h. Type 2.15 into the calculator instead and you get 80.9 km/h: an error of roughly 4.7 percent that

sits comfortably inside the plausible range and will match one of the options. Now extend it. The van consumes 8.6 litres per 100 km, so the trip needs $174 \times 8.6 / 100 = 14.964$ litres, and at 1.48 pounds per litre that is $14.964 \times 1.48 = 22.15$ pounds. Notice the units doing the work: (km) $\times$ (L / 100 km) leaves litres; (L) $\times$ (pounds / L) leaves pounds. If your intermediate line has km still attached at the end, you have divided when you should have multiplied. Write the unit next to every number and the method checks itself.

Unit pricing: normalise before you compare: Three pack sizes of the same fluid. Pack A: 750 ml for 3.60 pounds. Pack B: 2 litres for 9.40. Pack C: 1.5 litres for 7.20. Convert all three to price per litre. A: $3.60 / 0.75 = 4.80$ per litre. B: $9.40 / 2 = 4.70$. C: $7.20 / 1.5 = 4.80$. So B is cheapest by 10p a litre, and A and C are identical despite looking like different deals. The 'bigger pack is cheaper' heuristic happens to hold here but is not a rule, and test writers know it. Half of these items are built specifically so the largest pack loses. Now add the multibuy that these questions love: pack A is on three-for-two. Three packs give 2.25 litres for the price of two, 7.20 pounds, which is $7.20 / 2.25 = 3.20$ per litre and beats everything. The trap in the multibuy version is dividing by the number of packs paid for rather than the volume received.

Combined work rates: add rates, never times: Printer A completes a 3,000-page run in 50 minutes; printer B completes the same run in 75 minutes. Running together, how long? Convert to rates: A prints $3,000 / 50 = 60$ pages per minute, B prints $3,000 / 75 = 40$ pages per minute, and together they print 100 pages per minute, so the job takes $3,000 / 100 = 30$ minutes. Verify by counting output: in 30 minutes A produces 1,800 pages and B produces 1,200, which is 3,000 exactly. The two wrong answers you will see on the option list are 62.5 minutes (the average of 50 and 75) and 125 minutes (the sum). Both are impossible on inspection: two machines working together must finish faster than the faster machine alone, so any answer above 50 minutes is wrong before you compute anything. The general form is $1 / (1/50 + 1/75)$, and it is worth being able to write that line directly, but the sanity bound catches the error faster than the algebra does.

Estimate first, then compute: killing distractors in ten seconds: 'A department spent 847,300 pounds in 2024. In 2025 the budget fell by 12.5 percent. What was the 2025 spend?' Options: 741,387.50 / 953,212.50 / 105,912.50 / 762,570. Estimate before anything else: 12.5 percent is exactly one eighth, one eighth of roughly 850,000 is roughly 106,000, so the answer is roughly 744,000. That single line eliminates three options. 953,212.50 applied the change upwards. 105,912.50 is the size of the reduction, not the resulting spend: the 'answered the wrong question' distractor, and the most commonly selected wrong option on items of this shape. 762,570 is a 10 percent cut, planted for anyone who misread the rate. Exact working: $847,300 \times 7/8 = 5,931,100 / 8 = 741,387.50$. Doing it as a fraction avoids the decimal multiplication altogether. Learn the fraction equivalents cold (12.5 percent is $1/8$, 16.7 percent is $1/6$, 37.5 percent is $3/8$, 62.5 percent is $5/8$) because a fraction turns most percentage items into one division.

How to practise this skill

- Memorise the fraction-to-percentage table up to eighths and twelfths. Almost every 'nasty' percentage on these tests is a clean fraction in disguise, and the fraction route is two or three times faster than decimal multiplication.
- Write the unit beside every intermediate number, including the ones you keep in your head. Most numerical errors on timed sections are unit errors, not arithmetic errors, and units are the only cheap way to catch them.
- Rehearse reverse percentages as their own drill until dividing by 1.15 feels more natural than subtracting 15 percent. It is a single, identifiable technique that accounts for a disproportionate share of lost marks.
- Before selecting, re-read the last clause of the stem out loud. 'By how much did it fall' and 'what was it after the fall' have different answers, and both are always on the option list.
- Keep an error log with four columns (misread the question, wrong operation, unit slip, arithmetic slip), and tally which one you actually commit. Three sessions is usually enough to show that one column dominates, and that is the column to train.
- Work untimed until your method is stable, then add the clock in stages: generous, then realistic, then 20

percent tighter than the real test. Attempts stay on this device, so a slow first pass costs nothing.

Glossary

Percentage point: The arithmetic difference between two percentages. A rate moving from 4 percent to 5 percent rises by one percentage point but by 25 percent in relative terms; questions specify which they want and the two answers are both on the option list.

Reverse percentage: Recovering an original amount from a figure that already includes a percentage change, by dividing by the multiplier (1.15 for a 15 percent rise, 0.88 for a 12 percent fall) rather than subtracting the percentage from the new figure.

Multiplier (decimal factor): The single number a quantity is multiplied by to apply a percentage change: 1.20 for plus 20 percent, 0.80 for minus 20 percent. Successive changes are handled by multiplying the factors, which is why plus 20 then minus 20 gives 0.96, not 1.

Ratio part: One unit of a ratio split. Dividing the total by the sum of the ratio terms gives the value of one part; every share and every difference between shares is then a whole number of parts.

Unit rate: A quantity expressed per one of something: price per litre, pages per minute, litres per 100 km. Comparisons are only valid once every option has been converted to the same unit rate.

Combined rate: The sum of two or more individual rates working simultaneously. Times cannot be added or averaged; only rates add, and the combined time is the total work divided by the combined rate.

Significant figure estimate: Rounding every input to one meaningful digit to get an approximate answer in a few seconds. Its purpose is not accuracy but elimination: it removes any option that is off by a factor or has the sign of the change wrong.

Distractor: A wrong option written deliberately to match a specific predictable error: subtracting instead of dividing, reporting the change instead of the result, averaging instead of combining rates. Recognising the pattern is often faster than recomputing.

Study this next

- Numerical reasoning skill hub: <https://learn.novusstreamsolutions.com/aptitude/skills/numerical-reasoning>
- Numerical reasoning practice in the Cognitive Skills Lab:
<https://learn.novusstreamsolutions.com/cognitive-skills/numerical-reasoning>
- Data interpretation lesson:
<https://learn.novusstreamsolutions.com/aptitude/skills/data-interpretation/lesson>
- Office numeracy suite: <https://learn.novusstreamsolutions.com/aptitude/tests/office-numeracy>
- Finance, accounting and insurance careers:
<https://learn.novusstreamsolutions.com/careers/families/finance-accounting-and-insurance>
- Numerical reasoning lesson on Novus Learn:
<https://learn.novusstreamsolutions.com/aptitude/skills/numerical-reasoning/lesson>

Where this material comes from

- All worked figures above were written and checked for Novus Learn. Every division, ratio split and rate calculation in this lesson was verified by recomputing it in the reverse direction.
- Conventions only (percentage point, unit rate, significant figures) cross-checked against standard public references such as the Wikipedia articles 'Percentage', 'Ratio' and 'Rate (mathematics)'. No item text is drawn from any published test.
- 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.

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.

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/blueprint-and-plan-reading/study-outline>