

ANSWER KEY — 31 JULY 2026

Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10
A	C	D	B	C	B	A	D	C	C
Q11	Q12	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20
B	A	D	D	B	A	C	B	D	A
Q21	Q22	Q23	Q24	Q25	Q26				
C	A	C	C	D	B				

LEGAL REASONING

Q1 A Separate legal personality (Salomon) EASY

The Salomon principle establishes that an incorporated company is a legal person entirely separate from its members, so the company's debts are its own and not the personal debts of the shareholders. Anita's liability, as a member, is limited to any amount that remains unpaid on her shares; she has not personally guaranteed the loans, so the bank cannot pursue her private assets for the company's default. That is exactly what option A states, and it is the correct answer. Option B is wrong because holding a majority of shares does not make a member a guarantor of the company's borrowings. Option C invents a rule that personal liability tracks the shareholding proportion, which the law of limited liability squarely rejects. Option D confuses control with identity: even total control does not destroy the company's separate personality. Hence (A) is correct: separate personality plus limited liability shields Anita.

Examiner's note — the trap: The trap is option D, which reasons that heavy control collapses the separate personality; but control is not identity, and Salomon shields even a dominant shareholder from the company's debts.

Q2 C Lifting the veil - facade to evade obligation MED

A court will pierce the corporate veil where a company is used as a mere facade or sham to evade an existing legal obligation. Here the trader is already bound by a valid non-compete covenant, and he forms a company purely to carry on the very business the covenant forbids, hoping the company's separate personality will insulate the competition. Because the company is a device to defeat a pre-existing obligation, the court may lift the veil and restrain the company as well as the trader, as option C states. Hence (C) is correct. Option A treats separate personality as absolute, which it is not. Option B wrongly asserts that every restraint of trade is void; reasonable covenants are enforceable. Option D fabricates an insolvency precondition that has nothing to do with the facade doctrine. The decisive feature is the evasion of an existing legal duty through a newly created shell.

Examiner's note — the trap: Option A is the tempting distractor because it correctly recites separate personality; the trap is treating that principle as absolute and ignoring the recognised facade-to-evade-obligation exception that governs these facts.

Q3 D Insurable interest and company property (Macaura) HARD

Following the reasoning in *Macaura v Northern Assurance Co*, a shareholder, even one who owns every share, has no insurable interest in the company's property, because that property belongs to the company as a separate legal person and not to the members. Gopal insured the timber in his own name, yet the timber was the company's asset; he therefore had no insurable interest in it and cannot recover on the policy. That is precisely option D, the correct answer. Option A commits the classic error of equating sole ownership of shares with ownership of the company's assets, which separate personality forbids. Option B invents a proportional-recovery rule unknown to the law of insurable interest. Option C makes recovery turn on Gopal's intention, but intention cannot create an insurable interest that does not exist in law. Hence (D) correctly denies recovery for want of insurable interest.

Examiner's note — the trap: Option A is the strongest lure, since owning all the shares feels like owning the timber; but the company's separate personality means the asset is the company's, so the sole shareholder has no insurable interest in it.

Q4 B One-person company as employer (Lee v Lee) HARD

On the authority of *Lee v Lee's Air Farming Ltd*, a company's separate personality allows it to enter into a contract of employment with the very person who owns and controls it. Rekha and the company are distinct legal persons, so the company could validly employ her as its chief pilot even though she was also its sole shareholder and governing director. She was therefore an 'employee', and her dependants can claim the workers' compensation, exactly as option B states. Hence (B) is correct. Option A denies that a sole controller can also be an employee, but separate personality makes precisely that relationship possible. Option C recharacterises her salary as a withdrawal of capital, which contradicts a genuine employment contract for services actually rendered as a pilot. Option D imposes an artificial condition requiring other unconnected employees, which the law does not demand. The two capacities - controller and employee - coexist because two separate persons are involved.

Examiner's note — the trap: Option A is the trap: it feels intuitively impossible to employ oneself, but the company is a separate person, so it can genuinely contract to employ its own sole owner-director.

Q5 C Groups of companies - single economic unit ELITE

The mere fact that a parent and its subsidiary form a single economic group does not entitle a court to treat them as one person; each company retains its own separate personality. The veil is pierced only in exceptional cases, such as where the subsidiary is a facade concealing the true position, and quite separately the parent may be liable on ordinary principles - for instance where the parent itself owed and breached a duty of care to the claimant. Option C captures this nuanced position and is the correct answer. Hence (C) is correct. Option A wrongly makes group membership automatic identity. Option B wrongly makes separate personality an absolute shield admitting no exception, ignoring both facade cases and the parent's own direct duties. Option D confines the claimant to the subsidiary's assets in every case, which overlooks the parent's possible independent liability. The key is that economic unity alone is not a ground for piercing.

Examiner's note — the trap: Both A and B are extreme opposites designed to tempt; the elite trap is picking either the automatic-identity view (A) or the absolute-shield view (B), when the true rule sits in between and is stated only by option C.

Q6 B Unreasonable interference in nuisance MED

Private nuisance protects the use and enjoyment of land against unreasonable interference, and physical damage to neighbouring property caused by an emission is a paradigm actionable nuisance. Here the chemical works emits fumes that, over many months, actually damage the shrubs and discolour the paintwork of the adjoining house; this is a continuing interference causing real, tangible harm in a residential area, which the law readily treats as unreasonable. Option B states this correctly and is the answer. Hence (B) is correct. Option A wrongly claims fumes can never be a nuisance, contradicting settled authority. Option C invents a requirement that the works lack a licence or planning permission, but the presence of permissions is not decisive of nuisance liability. Option D asserts that public usefulness is a complete defence, yet it is well established that the utility of the defendant's activity is no defence to an action in private nuisance.

Examiner's note — the trap: Option D is the classic trap, invoking the works' usefulness to the public; but public benefit is not a defence to private nuisance, so it cannot defeat a claim for real damage to a neighbour's land.

Q7 A Abnormal sensitivity (Robinson v Kilvert) HARD

The law of nuisance judges the defendant's conduct by its effect on an ordinary occupier making ordinary use of land, and it disregards the abnormal sensitivity of the claimant or his property, as in *Robinson v Kilvert*. The modest, harmless rise in temperature would trouble no ordinary occupier; it damages only the claimant's exceptionally heat-sensitive specialist paper, a delicate use of land. Because the interference is measured against the ordinary occupier and not the sensitive one, the claim fails, which is exactly option A. Hence (A) is correct. Option B wrongly states that any physical damage to goods is always actionable, ignoring the sensitivity rule. Option C makes liability turn on the neighbour's knowledge of the paper, but knowledge does not convert a non-nuisance into a nuisance where the interference is harmless to ordinary use. Option D invents a bar on heat being a nuisance, which is not the law; heat can be a nuisance where it affects ordinary enjoyment.

Examiner's note — the trap: Option C is the seductive distractor, suggesting that the neighbour's knowledge of the delicate paper creates liability; but abnormal sensitivity is disregarded regardless of knowledge, since the ordinary-occupier standard governs.

Q8 D Locality and reasonableness in nuisance MED

In assessing whether an interference is unreasonable, the character of the locality is a central factor: what amounts to a nuisance in a quiet residential street may be entirely reasonable in an established industrial district. Here a new factory produces moderate noise, no worse than the surrounding works in a long-industrial area, so the locality strongly supports the conclusion that the noise is reasonable there. Option D captures this and is the correct answer. Hence (D) is correct. Option A wrongly claims locality is irrelevant and the standard uniform nationwide, contradicting the settled locality principle. Option B overstates matters by suggesting no nuisance action could ever be brought in an industrial area, however extreme the interference, which is false. Option C asserts that any measurable noise is automatically a nuisance everywhere, again ignoring that reasonableness is context-sensitive. The correct approach weighs the noise against the settled industrial character of the district.

Examiner's note — the trap: Option B tempts by over-reading the locality factor into a blanket immunity for industrial areas; but locality shapes reasonableness, it does not abolish nuisance liability where interference is truly excessive.

Q9 C Malice in nuisance (Christie v Davey) HARD

Malice - a deliberate intention to annoy - can convert an otherwise innocent activity into an actionable nuisance, as illustrated by *Christie v Davey* and *Hollywood Silver Fox Farm v Emmett*. The man's banging, whistling and shrieking against the party wall are done purely to disrupt his neighbour's lessons, and it is precisely this malicious purpose that makes the conduct unreasonable and therefore a nuisance, even though identical sounds made for an ordinary purpose might not be. Option C states this correctly and is the answer. Hence (C) is correct. Option A wrongly asserts that motive is always irrelevant to liability, but malice is exactly what tips the balance here. Option B demands proof of physical damage to the structure, yet nuisance protects comfortable enjoyment, not merely fabric. Option D invents a minimum duration of several months, but the malicious character, not the length of time, is decisive on these facts.

Examiner's note — the trap: Option A is the trap, reciting the general idea that a landowner may use his land as he likes; but where the noise is made maliciously to annoy, motive becomes decisive and defeats that supposed freedom.

LOGICAL REASONING

Q10 **C** Assumption - no offsetting cause **MED**

The department concludes that switching from diesel to electric buses will make the city's air cleaner. This conclusion depends on an unstated premise: that generating the electricity to charge the buses does not add as much pollution to the city's air as the diesel buses removed. If the charging electricity were produced by dirty power plants within the city, the promised gain could be cancelled out. Option C states exactly this necessary assumption and is the correct answer. Hence (C) is correct. Option A concerns operating cost, which is irrelevant to the pollution claim. Option B concerns residents' approval, which does not bear on whether the air actually becomes cleaner. Option D notes adoption by other cities, which is at most supporting evidence and not something the conclusion logically requires. Only C plugs a gap that, if false, would break the argument, which is the hallmark of a genuine assumption.

Examiner's note — the trap: Option A distracts by sounding like a benefit of electric buses; but cheaper operation says nothing about air quality, whereas only C addresses the pollution the conclusion actually turns on.

Q11 **B** Assumption - negation test **HARD**

A genuine assumption is one whose falsity would collapse the argument. Apply the negation test to option B: if buses did not account for a large enough share of the city's air pollution, then removing even the dirtiest diesel buses would make no noticeable difference to overall air quality, and the department's promise of 'markedly cleaner' air would fail. Because negating B destroys the conclusion, B is an assumption the argument depends on, and it is the correct answer. Hence (B) is correct. Option A, about identical routes and timetables, is irrelevant to how much pollution is cut. Option C, about a battery warranty, concerns reliability, not air quality. Option D, about how long the diesel buses would otherwise have run, does not determine whether their removal cleans the air noticeably. Only B, when negated, undermines the core claim, which is exactly what the negation test is designed to reveal.

Examiner's note — the trap: Option D tempts because a longer remaining diesel life sounds relevant to the benefit; but the argument still needs buses to be a big pollution source, and only B fails the negation test by wiping out the conclusion.

Q12 **A** Assumption - bridging premise **MED**

The department leaps from 'the air will be cleaner' to 'residents will breathe far healthier air'. That inference assumes a bridge: that cleaner outdoor air will actually deliver a health benefit to residents rather than being cancelled by other harms. Without this link, cleaner air need not translate into healthier breathing, so the bridging assumption is essential. Option A states it and is the correct answer. Hence (A) is correct. Option B, that residents spend most of the day outdoors, is a much stronger and narrower claim than the argument needs; the conclusion can hold without residents being outdoors all day. Option C, about securing the funding, concerns whether the plan happens, not whether cleaner air is healthier. Option D, about the relative age of the diesel buses, is irrelevant to the health inference. Only A supplies the missing connection between cleaner air and better health that the conclusion assumes.

Examiner's note — the trap: Option B overreaches: it assumes residents are outdoors nearly all day, a far heavier claim than the argument requires; the modest bridging assumption the conclusion truly needs is stated by A.

Q13 **D** Assumption - hidden leakage **ELITE**

The subtlest hole in the argument is that the withdrawn diesel buses might not actually disappear from the city's air. If they are merely sold on and continue to be driven within the same city, the total quantity of diesel fumes over that city could stay much the same, and the promised clean-up would not materialise. The department therefore quietly assumes the old buses will not keep polluting the very same air, which is option D, the correct answer. Hence (D) is correct. Option A is far too strong, requiring that no resident has any pollution-sensitive illness, which the argument does not need. Option B, about reliability and punctuality, does not bear on air quality. Option C, that every diesel bus emits identical amounts, is unnecessary; the argument works whatever the spread of emissions. Only D identifies a hidden leakage that, if it occurred, would defeat the department's promise of cleaner city air.

Examiner's note — the trap: Options A, B and C are plausible-sounding but either too strong or off-target; the elite trap is overlooking the quiet displacement assumption in D - that the scrapped buses will not simply pollute the same city from new owners.

Q14 **D** Strengthen - eliminate confounders **MED**

The researcher infers that walnuts improved heart health because the walnut group ended the year with better readings than the no-walnut group. That inference is strongest if the two groups were otherwise alike, so that the walnut habit was the main difference between them. Option D provides exactly this: closely matched diet, exercise and weight at the outset rules out obvious confounders and isolates the walnuts as the operative factor, which most strengthens the conclusion. Hence (D) is correct. Option A, about walnuts' long culinary history, is irrelevant to the causal claim. Option B, that volunteers enjoyed the habit, speaks to compliance, not to heart health. Option C, that other nuts may also help, does nothing to support walnuts specifically and could even dilute the claim. Only D closes the gap between correlation and causation by removing rival explanations, which is the classic way to strengthen a controlled comparison.

Examiner's note — the trap: Option B is attractive because enjoyment sounds positive; but liking the habit does not show walnuts caused the health gain, whereas matched groups in D remove the confounders that threaten the causal inference.

Q15 **B** Strengthen - reversibility test **HARD**

A powerful way to strengthen a causal claim is to show that removing the supposed cause reverses the effect. Option B does this: when the walnut group stopped eating walnuts, their cholesterol and blood pressure drifted back up towards earlier levels. This on-off pattern - benefit while eating walnuts, loss of benefit when they stop - strongly supports the claim that the walnuts, and not some background factor, produced the improvement. Option B is therefore the correct answer. Hence (B) is correct. Option A, about unequal group sizes, is at best a mild methodological quibble and does not support causation. Option C, that volunteers enjoyed the study, is irrelevant to whether walnuts caused the effect. Option D, about the researcher's earlier publications, speaks to her track record, not to this study's causal claim. Only B supplies the reversibility evidence that most convincingly ties the health gain to the walnuts themselves.

Examiner's note — the trap: Option D tempts by lending the researcher authority through past papers; but a strong publication record is not evidence about these walnuts, whereas the reversal on withdrawal in B directly supports causation.

Q16 A Strengthen - supply a mechanism **HARD**

An argument that X causes Y is strengthened when a plausible biological mechanism links the two. Option A supplies precisely that: walnuts are rich in a fatty acid that laboratory studies show directly lowers the kind of cholesterol linked to heart disease. A concrete pathway from walnut consumption to lower harmful cholesterol makes the observed heart benefit far more credible as a genuine causal effect rather than a coincidence. Option A is therefore the correct answer. Hence (A) is correct. Option B, that the walnut group was slightly younger, introduces a confounder and if anything weakens the claim. Option C, about cost and availability, is practical information irrelevant to causation. Option D, about volunteers feeling more energetic, describes a vague subjective impression that does not explain the cholesterol and blood-pressure findings. Only A offers a mechanism that explains how the walnuts could actually produce the measured cardiovascular improvement, which is what strengthening by mechanism requires.

Examiner's note — the trap: Option B looks relevant because age affects heart health; but a younger walnut group is a confounder that undercuts, not supports, the claim, whereas the biochemical pathway in A genuinely strengthens it.

Q17 C Strengthen - replication and generality **MED**

The recommendation is that adults generally should adopt the walnut habit, so it is strengthened by evidence that the benefit is not a one-off finding but holds broadly. Option C provides this: several large studies in different countries and age groups found the same heart benefit from a daily handful of walnuts. Replication across populations shows the effect is robust and generalisable, which directly supports a general recommendation. Option C is therefore the correct answer. Hence (C) is correct. Option A, that the study was confined to one small town, limits generality and weakens the recommendation. Option B, that rival researchers doubt the results, also weakens rather than strengthens. Option D, that the volunteers were the researcher's own friends and neighbours, points to a biased, unrepresentative sample and undermines the recommendation. Only C, by showing the benefit reproduces widely across countries and age groups, strengthens the case for advising adults in general to adopt the habit.

Examiner's note — the trap: Options A, B and D all quietly weaken the argument, so the trap is grabbing whichever sounds study-related; only C, replication across countries and ages, actually strengthens a recommendation aimed at adults generally.

QUANTITATIVE TECHNIQUES

Q18 B Simple average **EASY**

The simple average of the five turnout percentages is found by adding them and dividing by five. The turnout figures are 60, 72, 55, 66 and 68, which sum to 321. Dividing 321 by 5 gives 64.2 per cent, so option B is correct. Hence (B) is correct. Option A, about 62 per cent, understates the total, perhaps by mis-adding the figures. Option C, about 66.5 per cent, overstates it, as if a couple of the lower turnout values had been dropped. Option D, around 60 per cent, simply anchors on the smallest single value rather than averaging all five. The method is purely mechanical: sum the five percentages to 321, then divide by the count of 5, giving a clean simple average of 64.2 per cent for the district's five constituencies at the most recent poll.

Examiner's note — the trap: The trap here is anchoring on one low value like Ashok Nagar's 60 per cent (option D); the simple average must combine all five figures, which lands at 64.2, not 60.

Q19 D Average of registered voters **MED**

To find the average number of registered voters per constituency, add the registered-voter counts and divide by five. The counts, in thousands, are 120, 150, 100, 180 and 150, which sum to 700. Dividing 700 by 5 gives 140 thousand, so option D is correct. Hence (D) is correct. Option A, exactly 150 thousand, mistakes the most common single value (which appears twice) for the average, ignoring the smaller constituencies. Option B, about 120 thousand, anchors on one of the lower values rather than averaging. Option C, around 160 thousand, overstates the mean, as if the largest constituency dominated. The correct procedure aggregates all five registered-voter figures to 700 thousand and divides by the number of constituencies, five, producing an average of exactly 140 thousand registered voters per seat across the district.

Examiner's note — the trap: Option A is the trap: 150 thousand appears twice and feels like the typical value, but the mode is not the mean, and averaging all five counts gives 140 thousand.

Q20 A Weighted (overall) turnout **HARD**

The district's overall turnout is total votes cast divided by total registered voters, not the simple average of the percentages, because the constituencies differ in size. Votes cast (in thousands) are $120 \times 0.60 = 72$, $150 \times 0.72 = 108$, $100 \times 0.55 = 55$, $180 \times 0.66 = 118.8$ and $150 \times 0.68 = 102$, summing to 455.8. Total registered voters are 700 thousand. So overall turnout is 455.8 divided by 700, which is about 65.1 per cent, making option A correct. Hence (A) is correct. Option B, 64.2 per cent, is the simple average of the percentages and is the intended trap,

since it ignores constituency size. Option C, about 68 per cent, and option D, around 70.3 per cent, both overstate the weighted figure. Because the larger constituencies (Danapur and Bela) carry more weight, the size-weighted turnout of 65.1 per cent differs from the unweighted 64.2 per cent, which is the whole point of the question.

Examiner's note — the trap: Option B, 64.2 per cent, is the deliberate trap: it is the simple average of the five percentages, but overall turnout must weight each constituency by its registered voters, giving 65.1 per cent.

Q21 C Missing value from a new average ELITE

Adding a sixth constituency, Fatehpur, raises the simple average turnout of all six constituencies to exactly 65 per cent. With six constituencies averaging 65, the total of the six turnout figures must be $6 \times 65 = 390$. The five original turnout figures sum to $60 + 72 + 55 + 66 + 68 = 321$. Fatehpur's turnout is therefore 390 minus 321, which equals 69 per cent, so option C is correct. Hence (C) is correct. Option A, about 65 per cent, wrongly assumes the new value equals the new average, which would only be true if the old average were already 65. Option B, roughly 64 per cent, mistakes Fatehpur's turnout for the old five-seat average. Option D, about 75 per cent, overshoots by overcorrecting. The reliable method is to convert the target average into a required total (390), subtract the known total of the existing five (321), and read off the missing value, giving 69 per cent.

Examiner's note — the trap: Option A is the elite trap: students assume the added value equals the new average of 65; but since the old average was 64.2, Fatehpur must exceed 65 to pull the mean up, landing at 69.

Q22 A Column total from a table EASY

The total budgeted expenditure for this year is the sum of the five 'this year' allocations. Reading the column, the figures in crores are 4,500 for Education, 3,000 for Health, 6,000 for Infrastructure, 1,500 for Agriculture and 3,000 for Administration. Adding these gives $4,500 + 3,000 + 6,000 + 1,500 + 3,000 = 18,000$, so this year's total is Rs 18,000 crore, making option A correct. Hence (A) is correct. Option B, Rs 15,500 crore, is in fact last year's total and is the intended trap for anyone who reads the wrong column. Option C, about Rs 16,500 crore, and option D, around Rs 19,200 crore, are simply mis-additions of the correct column. The task is a straightforward column sum: identify the five 'this year' entries and add them carefully to reach the overall budget of Rs 18,000 crore for the year under review.

Examiner's note — the trap: Option B, Rs 15,500 crore, is the trap: it is last year's total, so reading across to the wrong column yields it; this year's five figures add to Rs 18,000 crore.

Q23 C Share of total MED

Infrastructure's share of this year's budget is its allocation divided by the total. Infrastructure receives Rs 6,000 crore, and the total budget this year is Rs 18,000 crore. So the share is 6,000 divided by 18,000, which equals one-third, or about 33.3 per cent, making option C correct. Hence (C) is correct. Option A, exactly 25 per cent, would correspond to a quarter of the budget, which Infrastructure exceeds. Option B, about 30 per cent, is a near-miss that undershoots the true one-third. Option D, approximately 40 per cent, overstates Infrastructure's slice. The reliable approach is to divide the single head's allocation by the correct total of Rs 18,000 crore, which cleanly reduces 6,000 over 18,000 to one-third. Recognising the fraction reduces the arithmetic to a simple ratio and confirms that Infrastructure commands almost exactly a third of the entire budget for the year.

Examiner's note — the trap: Option B, about 30 per cent, is the near-miss trap for those who estimate rather than divide; 6,000 over 18,000 is exactly one-third, or 33.3 per cent, not 30.

Q24 C Largest percentage increase HARD

Percentage increase for each head is (this year minus last year) divided by last year. Education: $500/4,000 = 12.5$ per cent. Health: $500/2,500 = 20$ per cent. Infrastructure: $1,000/5,000 = 20$ per cent. Agriculture: $300/1,200 = 25$ per cent. Administration: $200/2,800$ is about 7.1 per cent. The largest percentage increase is Agriculture at 25 per cent, so option C is correct. Hence (C) is correct. Option A points to Infrastructure's rise of 1,000 crore, which is the largest absolute increase but only a 20 per cent proportional rise, and is the main trap. Option B cites Health's 500-crore rise, again an absolute figure, worth only 20 per cent. Option D correctly computes Education at 12.5 per cent but that is far from the largest. Because percentage change is proportional to the base, Agriculture's smaller absolute rise on a low base of 1,200 yields the biggest percentage jump.

Examiner's note — the trap: Option A is the classic trap: Infrastructure has the largest absolute increase of 1,000 crore, but percentage change is proportional, and Agriculture's 300 on a base of 1,200 is a larger 25 per cent.

Q25 D Overall percentage growth HARD

The total budget grew from last year's Rs 15,500 crore to this year's Rs 18,000 crore. The increase is 18,000 minus 15,500, which is 2,500 crore. Percentage growth is 2,500 divided by 15,500, times 100, which is about 16.1 per cent, so option D is correct. Hence (D) is correct. Option A, about 13.9 per cent, wrongly divides the 2,500 increase by this year's total of 18,000 rather than by last year's base. Option B, roughly 25 per cent, overstates the rise substantially. Option C, approximately 2.5 per cent, mistakes the raw increase of 2,500 crore for the percentage itself, dropping the division by the base entirely. The correct method takes the change of 2,500 crore over the previous-year base of 15,500 crore, giving a growth of about 16.1 per cent for the total budget across the two years.

Examiner's note — the trap: Option A is the base-selection trap: dividing the 2,500 rise by this year's 18,000 gives 13.9 per cent, but percentage growth must be taken on last year's base of 15,500, giving 16.1 per cent.

Q26 B Solve for a required allocation ELITE

Next year the total budget stays at Rs 18,000 crore, and Education plus Health together must make up exactly half of it. Half of 18,000 is 9,000 crore, so Education and Health combined must total 9,000 crore. Education is held at 4,500 crore, so Health must be 9,000 minus 4,500, which is 4,500 crore, making option B correct. Hence (B) is correct. Option A, Rs 9,000 crore, is the combined half itself, not Health's individual share, and is the principal trap for those who forget to subtract Education. Option C, Rs 6,000 crore, matches Infrastructure's figure but has no basis in this calculation. Option D, about Rs 5,250 crore, is an unmotivated intermediate value. The method is to convert 'half of the budget' into 9,000 crore, then subtract the fixed Education allocation of 4,500 crore, leaving Health at exactly 4,500 crore for next year.

Examiner's note — the trap: Option A, Rs 9,000 crore, is the elite trap: it is the combined half of the budget, but the question asks only for Health, so Education's 4,500 must be subtracted, leaving 4,500.

RAPID FIRE — ANSWERS

1. Which landmark case established that a company is a separate legal person from its members? — **Salomon v Salomon & Co Ltd (1897).**
2. In company law, what does 'lifting the corporate veil' mean? — **Looking past the company's separate personality to the persons behind it, in exceptional cases such as fraud or facade.**