

ANSWER KEY — 22 SEPTEMBER 2026

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

LEGAL REASONING

Q1 C s.3(d) Explanation - salts deemed same substance **HARD**

The Explanation to Section 3(d) of the Patents Act, 1970 deems salts of a known substance to be the same substance 'unless they differ significantly in properties with regard to efficacy'. Velortan citrate is a salt of the known compound velortan, so option A's argument that chemical distinctness makes it a new product is precisely what the Explanation shuts out. The only question is whether the salt shows enhanced efficacy. In *Novartis AG v Union of India* (2013) the Supreme Court held that, for a medicine, efficacy means therapeutic efficacy, and that properties such as better flow, thermodynamic stability and lower hygroscopicity, though useful in manufacture and storage, do not establish it. Faster dissolution in vats and a longer shelf life are exactly such manufacturing and storage advantages, and Pharmaxa has no data on how patients respond. The application therefore fails, which is option C. Option B equates shelf life with efficacy, which Novartis rejects. Option D misstates Section 3(e), which excludes only a mere admixture that aggregates the properties of its components.

Examiner's note — the trap: B is the trap: stability sounds like a real improvement, but Novartis confines efficacy for a medicine to therapeutic efficacy, so storage benefits cannot rescue the salt.

Q2 D Synergy - s.3(e) and s.3(d) combinations **ELITE**

Section 3(e) of the Patents Act, 1970 excludes only a mere admixture 'resulting only in the aggregation of the properties' of its components; the passage adds that a combination producing a synergistic effect, greater than the sum of the separate effects, is not a mere admixture. Kedar's trial data show exactly such synergy, so the 3(e) objection can be met. Under Section 3(d), the Explanation does list 'combinations' as derivatives deemed the same substance, but only 'unless they differ significantly in properties with regard to efficacy'; markedly greater pain relief is enhanced therapeutic efficacy in the Novartis sense, so the 3(d) objection can also be met. Option B ignores that proviso. Option C would bar every combination of known drugs, which the synergy principle contradicts. Option A makes the opposite error: the passage states that escaping Section 3 does not by itself make a claim patentable, because Section 2(1)(j) still requires novelty, an inventive step and industrial application. Only option D captures both halves.

Examiner's note — the trap: A is the elite trap: it gets the Section 3 analysis right but skips the final step, since clearing the exclusions is necessary but not sufficient for a patent.

Q3 B s.3(d) known process - new reactant exception **HARD**

Section 3(d) of the Patents Act, 1970 excludes 'the mere use of a known process, machine or apparatus unless such known process results in a new product or employs at least one new reactant'. The two exceptions are joined by 'or', so either one is enough to take a known process outside the bar. Here the fermentation process is known and its product, the vitamin, is also known, so the first exception does not help the firm. But the firm feeds the process a newly synthesised sugar that no one has used in it before, which is a new reactant. The second exception therefore applies, and Section 3(d) does not exclude the claim. As the passage stresses, escaping Section 3 is not the end: the process must still be new, involve an inventive step and be capable of industrial application under Section 2(1)(j). That is option B. Option A reads the two exceptions as if both were required. Option C ignores the new-reactant exception entirely, and option D misapplies Section 3(i), which concerns processes of treatment, not processes of manufacture.

Examiner's note — the trap: A is the trap: it correctly notes that the product is not new but treats that as fatal, when the statute is satisfied by a new reactant alone.

Q4 C Novartis - bioavailability needs research data **HARD**

A polymorph is expressly listed in the Explanation to Section 3(d) of the Patents Act, 1970, so it is deemed the same substance as the known compound unless it differs significantly in properties with regard to efficacy; option B is therefore wrong on the text. In *Novartis AG v Union of India* (2013), which also involved a crystalline form, the Supreme Court held that increased bioavailability is not by itself enough: the applicant must claim and establish, by research data, that the higher bioavailability results in enhanced therapeutic efficacy. Counsel's argument that more absorption 'must mean' the drug works better is exactly the inference Novartis refused to

presume. So the office should refuse unless Orbis produces that data, which is option C. Option A treats a pharmacokinetic figure as proof of efficacy. Option D overstates the law: Section 3(d) bars only a new form that does not enhance known efficacy, so a polymorph that is shown to be therapeutically superior can clear it.

Examiner's note — the trap: D is attractive because it sounds strict, but Section 3(d) is a conditional bar, not an absolute one; a new form with proven therapeutic gain can qualify.

Q5 B s.2(7) Explanation - livelihood by self-employment MED

Section 2(7) of the Consumer Protection Act, 2019 excludes a person who obtains goods 'for any commercial purpose', but its Explanation provides that commercial purpose does not include use of goods bought and used by a person exclusively for earning his livelihood by means of self-employment. Harpreet drives the autorickshaw herself, employs no one, and it is her family's only income, so she falls squarely within the Explanation and remains a consumer. That is option B. Option A applies the dominant-purpose test of *Lilavati Kirtilal Mehta Medical Trust v Unique Shanti Developers* (2019) without the statutory carve-out: every self-employed livelihood earns money, which is exactly why the Explanation exists. Option C is contradicted by the definition itself, which expressly covers consideration 'paid or promised, or partly paid and partly promised, or under any system of deferred payment', so an instalment purchase qualifies. Option D invents an income ceiling that appears nowhere in the Act; the test is who uses the goods and for what, not how much they earn.

Examiner's note — the trap: A is the trap: it states the *Lilavati* profit-nexus test correctly, but the self-employment Explanation overrides that test for a person who works the goods herself.

Q6 D Pecuniary jurisdiction - value of consideration paid HARD

The passage states the rules notified in 2021 under the Consumer Protection Act, 2019: a District Commission hears complaints where the value of the goods or services paid as consideration does not exceed Rs 50 lakh, a State Commission from above Rs 50 lakh up to Rs 2 crore, and the National Commission above Rs 2 crore. It adds that the compensation claimed does not decide the forum. Manish paid Rs 45 lakh for the car, which is within the District Commission's limit, so the District Commission has pecuniary jurisdiction over the whole complaint, including the claim for compensation. That is option D. Option A adds the refund and compensation (45 lakh + 1.8 crore = 2.25 crore) and so uses the old 'amount claimed' approach that the rules displaced. Option B also looks at the compensation claimed rather than the consideration paid. Option C invents a split of one complaint between two forums; nothing in the rules divides a complaint by the heads of relief sought.

Examiner's note — the trap: A is the trap: adding up the claim feels natural and was the basis under the older regime, but the 2021 rules key jurisdiction only to the consideration paid.

Q7 A s.84(2) - express warranty, no-negligence no defence HARD

Under Chapter VI of the Consumer Protection Act, 2019 a product manufacturer is liable in a product liability action if the product 'does not conform to an express warranty', and Section 84(2) adds that he is liable even if he proves that he was not negligent or fraudulent in making the express warranty. The label 'Guaranteed shatterproof up to 150 degrees C' is an express warranty, and the kettle shattered at 100 degrees C, so it did not conform. The company's proof of careful testing and of an undetectable impurity goes only to negligence, which Section 84(2) makes irrelevant here. Tanvi's scalding is personal injury, which is 'harm'. The company is therefore liable, option A. Option B imports a negligence requirement that the provision expressly excludes. Option C misuses the idea of obvious dangers: a buyer relying on a shatterproof guarantee cannot be expected to know of a hidden flaw. Option D again shifts the inquiry to the standard of care, which is beside the point once the warranty is broken.

Examiner's note — the trap: B is the trap: in ordinary tort, proving due care defeats a claim, but Section 84(2) removes that defence where the product fails its own express warranty.

Q8 C Product seller liability + definition of 'harm' ELITE

Two rules combine. First, a product seller who is not the manufacturer is liable under the Consumer Protection Act, 2019 in specified cases, one of which is where the identity of the manufacturer is not known. No one can trace who made the heater, so Deepa, as seller, is answerable; option A wrongly treats alteration or modification as the only ground of seller liability, when misuse, alteration or modification is in fact a defence available to the seller. Second, 'harm' in a product liability action includes personal injury and damage to property other than the product itself, but excludes harm caused to the product itself. Sunil's burns are personal injury and his curtains are other property, so both are recoverable; the burnt-out coil is harm to the product itself and is not. That is option C. Option B ignores the exclusion for the product itself. Option D goes too far the other way: the definition expressly covers damage to property other than the product, so the curtains are within it.

Examiner's note — the trap: B is the elite trap: every loss is causally linked to the wiring, but the statutory definition of harm carves out the defective product itself, whatever the causation.

Q9 A Failure-to-warn exception - expert-supervised products HARD

The Consumer Protection Act, 2019 provides that in a product liability action based on failure to warn, the manufacturer is not liable where the product was legally meant to be used or dispensed only by or under the supervision of an expert and the manufacturer employed reasonable means to give the warnings to such experts. Relifen may lawfully be dispensed only on prescription, and the manufacturer put a clear warning against driving in the prescribing information sent to every doctor and pharmacist. That is reasonable means of warning the experts. Anil's action is based solely on failure to warn, so the manufacturer is not liable; his grievance lies with

the doctor who did not pass the warning on. That is option A. Option B states the ordinary rule that the exception displaces. Option C misreads Section 84(2), which removes the no-negligence defence only for breach of an express warranty, not for warnings. Option D overstates the exception: prescription drugs are not excluded from product liability altogether, only from this failure-to-warn claim.

Examiner's note — the trap: C is the trap: it borrows Section 84(2)'s strict-liability language, but that provision is confined to express warranties and says nothing about warnings to experts.

LOGICAL REASONING

Q10 **D** Parallel - valid modus tollens **MED**

Argument I has the form: if a bill clears committee (P), it is debated on the floor (Q); the bill was not debated (not-Q); so it did not clear committee (not-P). This is modus tollens, and it is the one valid argument the manual refers to. Option D says a pharmacy may stock the vaccine only if it holds a licence. 'Only if' introduces a necessary condition, so the rule is: stocks (P) leads to licence (Q). The pharmacy has no licence (not-Q), so it does not stock the vaccine (not-P). That is the same valid structure, even though the wording ('only if' rather than 'any') looks different. Option A reverses it, affirming the consequent: being debated does not show the bill cleared committee. Option B denies the antecedent: a pharmacy that does not stock the vaccine might still be inspected. Option C weakens the rule to 'some', from which nothing follows about a particular pharmacy lacking a licence.

Examiner's note — the trap: A is the trap: it copies Argument I's subject matter almost word for word, but its logic runs backwards; parallel reasoning is about form, not topic.

Q11 **C** Parallel flaw - denying the antecedent **HARD**

Argument II reasons: every door-to-door campaigner (P) wins her ward (Q); Meera did not campaign door to door (not-P); so she will not win (not-Q). This denies the antecedent: the rule says door-to-door campaigning is sufficient for winning, not necessary, so Meera could win some other way. Option C has exactly this shape: installing panels (P) is sufficient for a rebate (Q); Ravi did not install panels (not-P); so no rebate (not-Q). He could qualify for a rebate on some other ground. Option A is valid modus tollens (not fined, so did not pay late). Option B is also valid: 'only tenants who signed may keep a pet' makes signing necessary for keeping a pet, so a pet-owner must have signed. Option D is flawed, but differently: it affirms the consequent, treating installing panels, a necessary condition for a rebate, as if it were sufficient. That is Argument IV's error, not Argument II's.

Examiner's note — the trap: D is the trap: it is also invalid and uses the same tenants-and-rebates story, but its flaw is affirming the consequent, which matches Argument IV rather than II.

Q12 **B** Parallel - 'most + most' quantifier flaw **ELITE**

Argument III chains two 'most' statements: most A are B, and most B are C, so most A are C. This is invalid, because the A's that are B may fall largely among the minority of B's that are not C. For example, 60 of 100 attendees may improve, but if 1,000 members improve overall and the 60 attendees are among the 400 improvers who win nothing, the conclusion fails. Option B reproduces the exact form: most cyclists wear helmets, most helmet-wearers avoid injury, so most cyclists avoid injury. Option C replaces the second 'most' with 'all', which makes the argument valid: every helmeted cyclist avoids injury, and helmeted cyclists are most cyclists. Option D is invalid too, but it uses 'some' throughout, which is a different quantifier pattern. Option A has a different structure altogether: it moves from a claim about most of a group to a claim about a subset (partners) of that group. Only B matches both the quantifiers and the chain.

Examiner's note — the trap: D is the elite trap: it shares the chained, invalid shape, but a parallel answer must keep the same quantifiers, and 'some' is not 'most'.

Q13 **A** Parallel flaw - necessary treated as sufficient **MED**

In Argument IV, 'only students who have paid the fee may vote' makes paying the fee a necessary condition for voting: vote leads to paid. The argument then observes that Arjun has paid and concludes that he may vote, treating the necessary condition as if it were sufficient. He might still be barred for some other reason, such as not being enrolled. Option A repeats this exactly: 'may renew only if she has cleared her dues' makes clearing dues necessary for renewal; Kavya has cleared her dues; so she may renew. Clearing dues alone does not guarantee renewal. Option B is valid: lacking a necessary condition, Kavya cannot renew. Option D is also valid modus tollens on a sufficient-condition rule. Option C is invalid, but it denies the antecedent of a sufficient-condition rule, which is the error in Argument II, not Argument IV. Only option A mistakes a necessary condition for a sufficient one.

Examiner's note — the trap: C is the trap: it is flawed and looks similar, but it commits Argument II's error (denying the antecedent), not Argument IV's (affirming a necessary condition).

Q14 **B** Role - rival explanation raised then rebutted **MED**

S3 reports that critics point to two other changes in the same six months, the new metro line and the nationwide rise in fuel prices, 'either of which could have kept cars away'. These are alternative causes that would explain the fall recorded in S2 without crediting the charge. The argument does not stop there: S4 derives a testable consequence of these alternatives, S5 shows the consequence did not occur, and S6 concludes that the charge, 'and not the metro or the price of fuel', accounts for most of the fall. So S3 presents rival explanations that the argument then tests and largely rejects, which is option B. Option A is wrong because the conclusion is in S6. Option C reverses S3's function: it cuts against the corporation's reading of the figures rather than reinforcing it. Option D is wrong because the author does not concede that the rival explanations succeed; S6 expressly rejects them as the main cause, while allowing that the corporation's figures alone were not enough.

Examiner's note — the trap: D is the trap: the author does take S3 seriously, but taking an objection seriously is not conceding it; S6 rejects the rival causes as the main explanation.

Q15 A Bold-face pair - evidence vs conclusion HARD

S2 is the corporation's evidence: 26 per cent fewer cars crossing into the district than a year earlier. S6 opens with 'The corporation's figures alone could not have settled the matter', which is an explicit judgment that S2 was insufficient on its own, because of the rival explanations in S3. S6 then concludes that 'the charge ... accounts for most of the fall', which is exactly the claim S2 was cited to support in S1. So S6 both qualifies S2 and endorses its conclusion, now resting on the extra support of S4 and S5. That is option A. Option B overstates: S6 does not call the figures unreliable, and it still relies on the fall they record, since that is the fall being explained. Option C ignores that S6 adds the causal attribution and the rejection of the rival causes. Option D contradicts the word 'alone' in S6: the author expressly denies that S2 was decisive by itself.

Examiner's note — the trap: B is the trap: 'could not have settled the matter' sounds like a rejection, but the author still accepts the figures; they were insufficient, not wrong.

Q16 D Bold-face pair - rival prediction disconfirmed ELITE

S4 is conditional: 'If those two changes, rather than the charge, were responsible for the fall', then ring-road traffic 'should have fallen by a similar proportion'. The antecedent is the critics' hypothesis from S3, so S4 derives a prediction from the rival view, using the ring road as a comparison case exposed to the same metro and fuel effects but not to the charge. S5 reports that ring-road counts instead rose by 9 per cent, which is inconsistent with that prediction and so tells against the critics' hypothesis. That is option D. Option B is the tempting mirror image, but the prediction in S4 belongs to the critics, not the corporation, and S5 disconfirms rather than confirms it. Option A is wrong because the main conclusion is in S6; S4 is an intermediate conditional. Option C misreads S4 as a concession: it is a test the author sets up precisely in order to show that it is failed.

Examiner's note — the trap: B is the elite trap: it correctly sees prediction-plus-observation but attributes the prediction to the wrong hypothesis and reverses the observation's effect.

Q17 C Method - comparison case to eliminate rivals HARD

The argument accepts the corporation's evidence (S2), identifies rival explanations (S3), and then tests them by looking at a comparison case: the uncharged ring road, which the passage says is used by the same commuters and equally exposed to fuel prices (S4). If the metro and fuel prices were the cause, the ring road should also show a fall; it rose by 9 per cent instead (S5). The author therefore concludes that the charge accounts for most of the fall (S6). That is supporting a causal claim by using a comparison case to test and discount competing explanations, which is option C. Option A describes a mere time-order argument, which the passage never relies on. Option B names a kind of attack the passage does not make: the period is not called unrepresentative. Option D describes the critics' position in S3, not the argument's overall method, which ultimately rejects the two alternative causes.

Examiner's note — the trap: D is the trap: it accurately describes S3 in isolation, but a method question asks about the whole argument, which answers that objection rather than resting on it.

QUANTITATIVE TECHNIQUES

Q18 A Volume of cylinder + cone HARD

Formula: volume of wheat = volume of cylinder + volume of cone = $\pi r^2 h + (1/3) \pi r^2 h_c = \pi r^2 (h + h_c/3)$. Substitution: $\pi r^2 = (22/7) \times 3.5 \times 3.5 = 38.5$ square metres. Cylinder volume = $38.5 \times 10 = 385$ cubic metres. Cone volume = $(1/3) \times 38.5 \times 1.2 = 15.4$ cubic metres. Total = $385 + 15.4 = 400.4$ cubic metres. Mass = $400.4 \times 0.75 = 300.3$ tonnes = 300,300 kg. Number of 50 kg bags = $300,300 / 50 = 6,006$. Result: 6,006 bags, option A. The distractors come from standard slips. Counting only the cylinder gives $385 \times 0.75 \times 20 = 5,775$ (option B). Forgetting the one-third in the cone formula treats the roof as a cylinder: $(385 + 46.2) \times 15 = 6,468$ (option C). Using two-thirds instead of one-third gives $(385 + 30.8) \times 15 = 6,237$ (option D). Note that the slant height, $\sqrt{3.5^2 + 1.2^2} = 3.7$ m, is needed only for surface area, not volume.

Examiner's note — the trap: B is the trap for students who read 'cylindrical wall' and stop; the passage says the wheat also fills the roof cone, which adds 15.4 cubic metres.

Q19 C Trapezium area x flow rate HARD

Formula: flow per hour = cross-sectional area x speed; time = volume needed / diverted flow per hour. Substitution: area of the trapezium = $(1/2) \times (2.4 + 1.6) \times 1.5 = 3$ square metres. Speed = 3 km/h = 3,000 m/h, so the canal carries $3 \times 3,000 = 9,000$ cubic metres per hour. One-quarter is diverted: $9,000 / 4 = 2,250$ cubic metres per hour. Volume needed = area of block x depth = $45 \times 10,000 \times 0.08 = 36,000$ cubic metres. Time = $36,000 / 2,250 = 16$ hours. Result: 16 hours, option C. Treating the channel as a rectangle using the surface width gives $2.4 \times 1.5 = 3.6$ square metres, a diverted flow of 2,700 cubic metres per hour and 13 hours 20 minutes (option A). Using the bed width gives 2.4 square metres and 20 hours (option B). Forgetting that only one-quarter of the flow is diverted gives $36,000 / 9,000 = 4$ hours (option D). Converting km/h to m/h and hectares to square metres before dividing avoids unit errors.

Examiner's note — the trap: A catches students who multiply the top width by the depth; the trapezium's area uses the mean of the two parallel sides.

Q20 B Excavation, bulking and spreading ELITE

Formula: rise in level = volume of earth spread / area over which it is spread. Substitution: volume dug = $20 \times 16 \times 4 = 1,280$ cubic metres. Dug earth bulks by 20 per cent, so the volume spread = $1,280 \times 1.2 = 1,536$ cubic metres. The earth is spread over the rest of the plot, excluding the tank's own footprint: $100 \times 80 - 20 \times 16 = 8,000 - 320 = 7,680$ square metres. Rise = $1,536 / 7,680 = 0.2 \text{ m} = 20 \text{ cm}$. Result: 20 cm, option B. Ignoring the bulking gives $1,280 / 7,680 = 16.67 \text{ cm}$ (option A). Spreading the bulked earth over the whole 8,000 square metres, as if the tank's area could be covered, gives $1,536 / 8,000 = 19.2 \text{ cm}$ (option C). Making both errors gives $1,280 / 8,000 = 16 \text{ cm}$ (option D). The two adjustments, bulking and excluding the tank's footprint, each push the answer up, and a careful solver must apply both to reach 20 cm.

Examiner's note — the trap: C is the elite trap: it remembers the 20 per cent bulking but spreads the earth over the tank's own footprint, which cannot receive it.

Q21 **D** Rectangle minus circle **MED**

Formula: area not reached = area of rectangle - area of circle swept by the arm. Substitution: rectangle = $210 \times 140 = 29,400$ square metres. The arm of 70 m sweeps a circle of radius 70 m: area = $(22/7) \times 70 \times 70 = 15,400$ square metres. Because the field is 140 m wide, the circle of diameter 140 m just fits across the width and lies wholly inside the field. Area not reached = $29,400 - 15,400 = 14,000$ square metres. Converting at 1 hectare = 10,000 square metres gives 1.40 hectares. Result: 1.40 hectares, option D. Option A, 1.54 hectares, is the area the sprinkler does reach. Option B, 0.42 hectares, is what a solver gets by subtracting the circle from a 140 m square ($19,600 - 15,400 = 4,200$) and forgetting the extra 70 m of length. Option C, 2.94 hectares, is the whole field. Checking that the circle fits inside the field is the step that confirms the simple subtraction is valid.

Examiner's note — the trap: B is the trap: the four corners of the inscribed square are the usual textbook answer, but this field is a rectangle and has two further 70 m strips beyond the circle.

Q22 **C** Average speed with wind - harmonic mean **MED**

Formula: average speed = total distance / total time; ground speed = airspeed $\pm$ wind speed. Substitution: going out into a 12 km/h headwind, ground speed = $60 - 12 = 48 \text{ km/h}$, so time out = $12 / 48 = 0.25 \text{ h} = 15 \text{ minutes}$. Returning with the wind, ground speed = $60 + 12 = 72 \text{ km/h}$, so time back = $12 / 72 = 1/6 \text{ h} = 10 \text{ minutes}$. Total flying time = 25 minutes = $5/12 \text{ h}$; total distance = 24 km. Average = $24 / (5/12) = 57.6 \text{ km/h}$. Result: 57.6 km/h, option C. As a check, for equal distances the average is the harmonic mean: $2 \times 48 \times 72 / (48 + 72) = 6,912 / 120 = 57.6$. Option A assumes the headwind and tailwind cancel, but the drone spends longer in the slow leg. Option B includes the 3-minute hover ($24 / (28/60) = 51.4$), which the question excludes. Option D is simply the outbound ground speed.

Examiner's note — the trap: A is the trap: wind gains and losses look symmetric, but the slower leg takes more time, so a round trip in wind is always slower than in still air.

Q23 **B** Relative speed with a head start **HARD**

Formula: when two bodies approach each other, time to meet = gap / (sum of speeds). Substitution: drone A flies alone from 10:00 to 10:10, covering $54 \times (10/60) = 9 \text{ km}$. The remaining gap at 10:10 = $45 - 9 = 36 \text{ km}$. Closing speed = $54 + 36 = 90 \text{ km/h}$, so time to meet = $36 / 90 = 0.4 \text{ h} = 24 \text{ minutes}$ after 10:10, that is at 10:34. Distance of A from the hub = $9 + 54 \times 0.4 = 9 + 21.6 = 30.6 \text{ km}$. Check from B's side: B covers $36 \times 0.4 = 14.4 \text{ km}$ from the warehouse, and $45 - 14.4 = 30.6 \text{ km}$. Result: at 10:34, 30.6 km from the hub, option B. Option A assumes both drones left at 10:00 ($45 / 90 = 0.5 \text{ h}$, $54 \times 0.5 = 27 \text{ km}$). Option C runs A's clock for 36 minutes from 10:00 while using B's schedule, a mismatch. Option D gets the time right but forgets A's 9 km head start when computing the distance.

Examiner's note — the trap: D is the trap: the meeting time is correct, but the 9 km that A flew before B took off must be added to A's distance from the hub.

Q24 **A** Parallel data links - combined throughput **ELITE**

Formula: when parts of one job are split so that both finish together, time = total work / combined rate. Substitution: file size = $3.6 \text{ GB} = 3,600 \text{ MB} = 3,600 \times 8 = 28,800 \text{ megabits}$. Useful rates: $100 \times 0.9 = 90 \text{ Mbps}$ and $60 \times 0.9 = 54 \text{ Mbps}$, combined 144 Mbps. Time = $28,800 / 144 = 200 \text{ seconds} = 3 \text{ minutes } 20 \text{ seconds}$. Check: the fast link carries $90 \times 200 = 18,000 \text{ megabits}$ and the slow link $54 \times 200 = 10,800 \text{ megabits}$, which add to 28,800, so both finish together. Result: 3 minutes 20 seconds, option A. Option B, 180 seconds, uses the rated 160 Mbps and ignores the 90 per cent efficiency. Option C, 320 seconds, sends the whole file over the fast link alone ($28,800 / 90$). Option D, about 267 seconds, splits the file equally, so the slow link takes $14,400 / 54 \text{ seconds}$ and finishes last. The key step is converting bytes to bits: forgetting the factor of 8 would give an answer eight times too small.

Examiner's note — the trap: D is the elite trap: an equal split looks fair, but it leaves the slower link as the bottleneck; the best split is proportional to each link's useful speed.

Q25 **D** Round-trip range against wind **HARD**

Formula: time out + time back = available time, with ground speed = airspeed $\pm$ wind. Substitution: out against a 10 km/h headwind, ground speed = $40 - 10 = 30 \text{ km/h}$; back with the wind, $40 + 10 = 50 \text{ km/h}$. Let the distance be $d \text{ km}$. Then $d/30 + d/50 = 30 \text{ minutes} = 0.5 \text{ h}$. Taking the LCM, $(5d + 3d) / 150 = 0.5$, so $8d = 75$ and $d = 9.375 \text{ km}$. Result: 9.375 km, option D. Check: $9.375 / 30 = 0.3125 \text{ h}$ (18.75 min) out and $9.375 / 50 = 0.1875 \text{ h}$ (11.25 min) back, a total of 30 minutes. Option A, 10 km, is the range in still air (40×0.25), which assumes that the wind's losses and gains cancel. Option B, 7.5 km, gives each leg 15 minutes at the headwind speed of 30 km/h. Option C, 12.5 km, does the same at the tailwind speed of 50 km/h. Splitting the time equally between the two legs is wrong because the slower leg must take longer.

Examiner's note — the trap: A is the trap: intuition says the tailwind gives back what the headwind takes, but in time terms wind always shrinks the out-and-back range.