Course Companion

University Medicine

Every topic handout in one volume

exercises.lol

Contents

1. Topic 1 3
2. Topic 2 14
3. Topic 3 26
4. Topic 4 37
5. Topic 5 47
6. Topic 6 57
7. Topic 7 67
8. Topic 8 78
9. Topic 9 89
10. Topic 10 99
11. Topic 11 109
12. Topic 12 119
13. Topic 13 129

Topic 01

University Medicine

What this course is, and how to use it

This is the first course in University Medicine: **cell and molecular biology for medicine** . It is the preclinical base every later course stands on: biochemistry fills the cell, physiology makes it work, pathology explains when it breaks.

Scope limit, stated once: this is a study course, not a medical degree – no patients, no clinical placements; every case below is fictional and nothing here licenses or guides treatment.

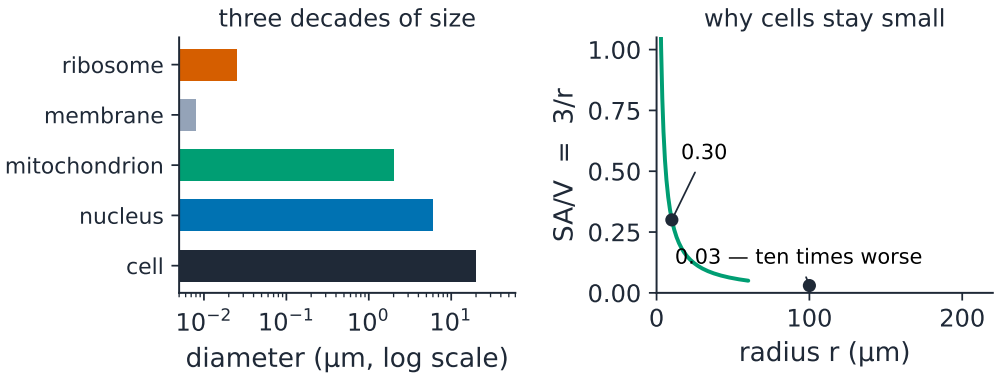
- Units 1.1–1.2 are the cell and its membranes: organisation, transport, and the electrical language of ions.
- Units 1.3–1.5 are the information line: replication, gene expression, and its regulation.
- Units 1.6–1.7 are decision-making: signalling between cells, and the cell's own life-and-death controls.

The medicine habit: **mechanism before fact**. Every number in this course is something you recompute – a Nernst potential, a compaction ratio, a mutation probability – and every clinical link names the mechanism it teaches.

Cell organisation and organelles

Why it matters. A eukaryotic cell is a factory of compartments: each organelle is a reaction vessel with the right pH, the right enzymes and the right surface. Medicine reads a cell the way an engineer reads a plant diagram – what lives where, and why.

The idea, with a picture. Organelles live on a size ladder spanning three orders of magnitude: a **ribosome** is 25 nm, a **mitochondrion** 2 µm, the **nucleus** 6 µm, the whole **cell** 20 µm. Function follows size: anything the whole cell must do quickly needs a membrane surface to do it on.



Big cells have proportionally little membrane – the surface-to-volume argument caps cell size

Definitions with numbers. The **nucleus** stores 2.2 m of DNA (6.4×10^9 bp $\times$ 0.34 nm/bp) inside a 6 µm envelope – a length about 4×10^5 times the nucleus's width. The **rough endoplasmic reticulum** is ribosome-studded protein export; the **Golgi** sorts and glycosylates; **lysosomes** keep pH ≈ 4.8 to dismantle cargo; **mitochondria** make ATP. **Ribosomes** translate mRNA (unit 1.4).

The derivation. Transport in and out scales with membrane **surface area** ($4\pi r^2$); demand – protein synthesis, ATP turnover – scales with **volume** ($\frac{4}{3}\pi r^3$). The ratio $SA/V = 3/r$: double the radius and every unit of cytoplasm gets half the membrane service. That single line is why cells stay small, why the ER folds so much surface into the cytoplasm, and why the intestinal lining is villi rather than flat.

> **Worked example 1 – the size ladder, verified.** A 20 µm hepatocyte: $SA/V = 3/10\ \mu m^{-1} = 0.3$. A 200 µm sphere (an oocyte): 0.03 – ten times worse ✓, which is why large cells are rare and membrane-folded. Cross-check the ladder: 20 µm / 25 nm = 800 – a ribosome is as small against a cell as a cell is against a 1.6 cm marble ✓.

> **Worked example 2 – clinical link: the lysosome that fails.** A (fictional) patient with progressive muscle weakness is found to lack a single lysosomal hydrolase. Mechanism, not name, is the answer: the substrate cannot be dismantled at pH 4.8, so it accumulates, the lysosome swells and the cell dies by being unable to recycle. One missing enzyme in one compartment = one storage disease: the organelle map is the diagnosis map.

Common mistakes. Treating "the cytoplasm" as one well-mixed bag – compartments differ in pH, calcium and redox state by design. Reading organelle counts as fixed (a liver cell has ~1000 mitochondria; a quiescent lymphocyte, few). And forgetting the membrane is where the cell meets the world.

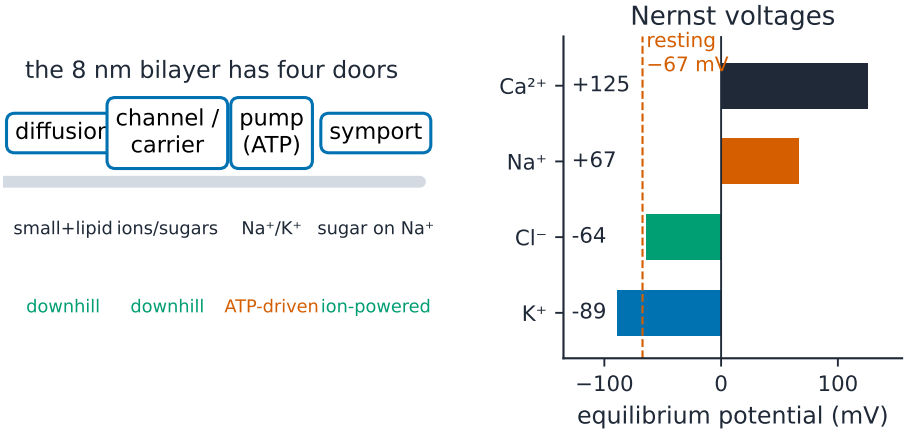
Summary. Organelles are reaction vessels on a 25 nm–20 µm ladder; $SA/V = 3/r$

caps cell size and explains folded membranes; 2.2 m of DNA packs into a 6 µm nucleus; compartment failure is disease.

Membranes and transport

Why it matters. Every drug target, every nerve impulse and every kidney's work reduces to one question: how does a molecule cross a lipid bilayer? The answer is always one of four doors.

The idea, with a picture. The **lipid bilayer** is a self-assembled 8 nm sheet impermeable to ions. Crossing is by: **simple diffusion** (small, lipophilic), **facilitated diffusion** through channels or carriers (down gradient, no ATP), **primary active transport** (ATP pumps, against gradient), and **secondary active transport** (an ion gradient drives something else uphill).



Four doors, four rules – and the equilibrium voltage each ion would need to stop moving

The Nernst equation. For ion X at 37 °C the **Nernst equation** gives the voltage at which the ion stops moving:

$$E_X = \frac{RT}{zF}\ln\frac{[X]_{out}}{[X]_{in}} = \frac{61.5\ mV}{z}\cdot\log_{10}\frac{[X]_{out}}{[X]_{in}}$$

Typical mammalian values – K^+ (140 in / 5 out): $E_K = 61.5\cdot\log_{10}(5/140) = -89.0$ mV ✓; Na^+ (12/145): +66.6 mV; Cl^- (10/110, $z=-1$): −64.0 mV; Ca^{2+} (100 nM/1.2 mM, $z=2$): +125.4 mV.

The derivation. At equilibrium the electrical force balances the chemical: $zFE = RT\ln([out]/[in])$ – that IS the Nernst line. With several ions permeable at once, replace the single ratio with a permeability-weighted sum (Goldman–Hodgkin–Katz): with $P_K : P_{Na} : P_{Cl} = 1 : 0.04 : 0.45$,

$$V_m = 61.5\log_{10}\frac{[K]_o + 0.04[Na]_o + 0.45[Cl]_i}{[K]_i + 0.04[Na]_i + 0.45[Cl]_o} = -67.3\ mV$$

✓ the **resting potential** – about 20 mV above E_K, because Na^+ leaks in and the Na^+/K^+-ATPase (3 Na^+ out, 2 K^+ in, 1 ATP) holds the gradients it needs.

> **Worked example 1 – hyperkalaemia arithmetic.** A (fictional) patient's extracellular K^+ doubles from 5 to 10 mM. New Nernst: E_K shifts by $61.5\log_{10}(10/5) = +18.5$ mV, from −89 to −71 mV ✓ – the depolarised E_K flattens the resting voltage, which is why raised potassium is a cardiac emergency: the mechanism is one logarithm.

> **Worked example 2 – clinical link: a broken chloride door.** The most common cystic-fibrosis variant deletes one phenylalanine (ΔF508) from a Cl^- channel, so the channel misfolds and never reaches the membrane. Follow the mechanism: Cl^- cannot leave → less salt reaches the airway surface, so less water follows it by osmosis → the surface liquid dehydrates → mucus cannot clear. One folding defect, one door, one disease – and small molecules that hold the fold open are the therapy rationale.

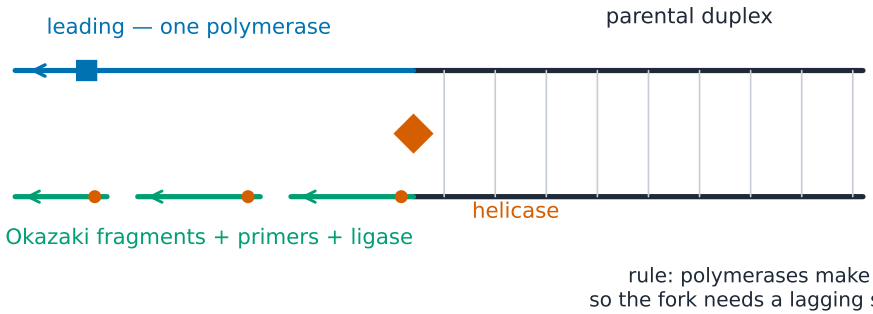
Common mistakes. Using the Nernst equation on non-equilibrium gradients (it is an equilibrium statement; GHK is the honest steady approximation). Forgetting z in the denominator (Ca^{2+} halves the slope). Calling channels "active" – only pumps burn ATP.

Summary. Four doors cross the bilayer; $E_X = (61.5/z)\log_{10}([out]/[in])$ mV at 37 °C; GHK with $P_K{:}P_{Na}{:}P_{Cl}$ = 1:0.04:0.45 gives −67 mV; 3 Na^+ out / 2 K^+ in per ATP keeps the battery charged.

DNA replication and repair

Why it matters. A cell must copy 6.4 billion letters with almost no mistakes, once per division, forever. The machinery that manages this – and the repair systems behind it – is where hereditary disease, mutagenesis and most oncology begin.

The idea, with a picture. The **replication fork** : helicase unzips, polymerases extend 5′→3′ only, so one strand is copied continuously (**leading**) and the other backwards in **Okazaki fragments** (**lagging**), each restarted by a primer.



One antiparallel rule forces the fork into a semiconservative machine with a lagging strand

Definitions with numbers. Rate: bacterial polymerase ~1000 nt/s; human ~50 nt/s per fork. Fidelity: raw polymerase error $\sim 10^{-5}$ per bp; **proofreading** buys $\times 100 \to 10^{-7}$; **mismatch repair** another $\times 100 \to 10^{-9}$ per bp ✓ – three systems deep, one billionth.

The derivation (the timing argument). Human S phase copies 6.4×10^9 bp in 8 h. A single origin with two forks at 50 bp/s would need $6.4 \times 10^9 / 100 = 6.4 \times 10^7$ s $\approx$ **2 years**. Replication works because ~30 000 origins fire in timed clusters: $6.4 \times 10^9 / (100 \times 3 \times 10^4) \approx 2100$ s ≈ 35 min of pure synthesis ✓, wrapped in chromatin time to fill S phase.

> **Worked example 1 – the fidelity budget.** After proofreading and mismatch repair the residual rate is $\sim 10^{-9}$/bp/division. For a 6.4×10^9 bp diploid genome: $6.4 \times 10^9 \times 10^{-9} \approx 6$ new mutations per daughter cell ✓ – the number oncology (unit 1.7) and genetics both lean on. Skip mismatch repair and the rate rises $\times 100$: hundreds of mutations per division, which is hereditary non-polyposis colon cancer's mechanism.

> **Worked example 2 – clinical link: sunlit skin without repair.** (Fictional patient) severe sunburn blisters after minutes outdoors, skin cancers by twenty. Mechanism: ultraviolet light fuses adjacent thymines into dimers; **nucleotide excision repair** is the system that cuts them out; without it every sunbeam is a mutagen with no auditor. The lesson generalises: each repair pathway owns one lesion class, and its absence names one syndrome.

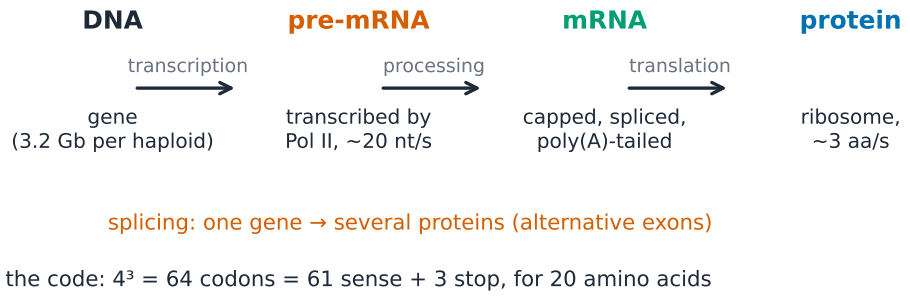
Common mistakes. Reading polymerases as reading 3′→5′ – they MOVE 3′→5′ along the template while MAKING 5′→3′. Calling S phase "copying once" without the origins argument. Quoting 10^{-9} as the raw polymerase rate (it is the post-repair rate).

Summary. Semiconservative, 5′→3′, one leading + one lagging strand; $10^{-5} \to 10^{-7} \to 10^{-9}$ fidelity ladder; ~30 000 origins turn 2 years into 8 hours; ~6 mutations survive per division; each repair pathway owns one lesion class.

Gene expression: transcription and translation

Why it matters. The gene is a recipe, not a machine: between DNA and protein stand RNA polymerase, processing, splicing and the ribosome – and most human disease mutations, it turns out, break the in-between steps, not the recipe itself.

The idea, with a picture. The **central dogma** with its eukaryotic detours: DNA → **pre-mRNA** mRNA (**transcription** by Pol II at ~20 nt/s) → 5′ cap + poly(A) + **splicing** (introns out, exons in) → mature mRNA → ribosome (**translation**) → protein (~3 amino acids/s). One gene, alternative exons, several proteins.



Transcription, processing, translation – three stages, three chances for regulation

Definitions with numbers. The genetic code: 64 **codons** = 61 sense + 3 stop; 20 amino acids from 4 bases in threes ($4^3 = 64$ ✓ – the arithmetic OF the code). Degeneracy averages ~3 codons per amino acid; AUG starts (methionine), UAA/UAG/UGA stop.

The derivation (the counting argument). Why triplets? Two bases give $4^2 = 16 < 20$: too few. Three give $4^3 = 64 > 20$: enough, with surplus spent on degeneracy – mostly at the third position, which is why third-base substitutions are so often silent. The code's redundancy is not waste; it is damage tolerance.

> **Worked example 1 – timing a protein.** **Preproinsulin** is 110 amino acids: translation at ~3 aa/s takes ~37 s; its 0.5 kb mRNA took ~25 s to transcribe at 20 nt/s. The numbers set expectations: making one insulin molecule costs about a minute of machinery, which is why secretion systems pre-build and store (unit 1.6 signalling triggers release, not synthesis).

Worked example 2 – clinical link: the splice that breaks. (Fictional pedigree) anaemia, and sequencing shows no coding mutation – the defect is a single base that creates a new splice site inside a β-globin intron. Mechanism: the spliceosome obeys consensus sequences; a new site lures it, a stretch of the intron is kept in the mRNA, the reading frame shifts and a premature stop follows. β-thalassaemia without a single amino-acid change: the dogma's middle step is the disease.

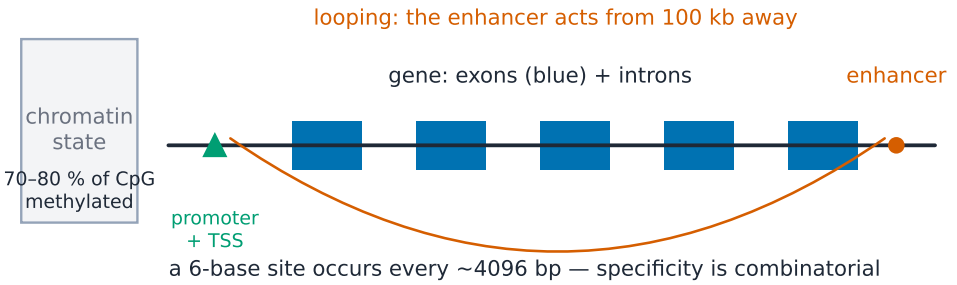
Common mistakes. Calling introns "junk" (regulatory and splicing-adjacent function, unit 1.5). Confusing the template strand with the coding sequence (the mRNA matches the coding strand, T→U). Expecting one gene = one protein (alternative splicing).

Summary. Pol II ~20 nt/s → processing (cap, poly(A), splice) → ribosome ~3 aa/s; $4^3 = 64$ codons for 20 amino acids with third-base degeneracy; one insulin molecule ≈ one minute; splice-site mutations break genes without touching their protein code.

Regulation of gene expression

Why it matters. Every one of your ~20 000 genes exists in nearly every cell; what makes a neuron a neuron is which genes are ON, when, and how much. Regulation is the actual genome – and the most common site of both evolution and disease.

The idea, with a picture. Control stacks: **transcription factors** bind promoters and **enhancers** (which can act from 100 kb away by DNA looping); **chromatin** state – histone marks and CpG **methylation** – decides access; **epigenetics** is the heritable part of that access without changing sequence.



Enhancer, promoter, transcription start, exons – and the chromatin gate in front of them

Definitions with numbers. Only ~1.5 % of the genome codes protein; regulatory sequence dwarfs it. 70–80 % of CpG dinucleotides in a differentiated cell are methylated – the switched-off baseline. The 2.2 m of DNA (unit 1.1) wraps nucleosomes (~200 bp each, about 7-fold shorter) folded into a ~30 nm fibre (about 40-fold); loops and further folding bring a mitotic chromosome near 10^4-fold.

The derivation (the combinatorics of specificity). A transcription factor recognises ~6–8 bases. Random 6-mers occur every ~4096 bp – about 800 000 times along each strand of the 3.2 Gb haploid genome: far too often for clean regulation. Specificity comes from COMBINATIONS: several factors binding adjacent sites cut false hits multiplicatively (4^{-6} to the power of cooperating sites) – the same logic that makes passwords long. Enhancers are cooperative logic gates, which is why regulation is network-shaped.

Worked example 1 – lactase persistence. The enzyme LCT falls after weaning in most of the world's population; northern Europeans carry a single base change ~14 kb UPSTREAM – inside an intron of a neighbouring gene, where it creates new enhancer binding. Mechanism chain: one regulatory base → one enhancer strengthened → LCT stays on into adulthood → milk digestible. A zero-protein-change trait that swept a continent: regulation evolves fastest.

Worked example 2 – clinical link: silencing gone wrong. (Fictional twin pair) identical genotypes, different severity: one carries methylation of the disease gene's promoter, the other does not. Same letters, different access – an **epigenetic** modifier of severity. The mechanism explains why "identical genes" never promises identical disease, and why methylation assays sit next to sequencers in the diagnostic laboratory.

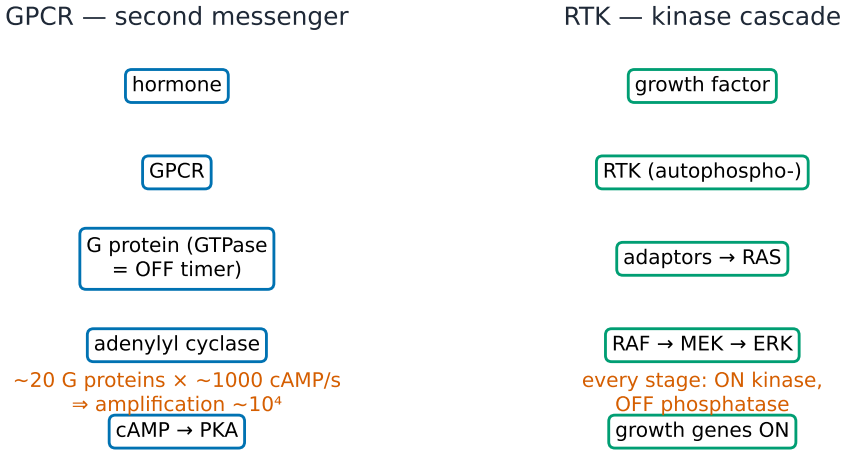
Common mistakes. Equating epigenetic with permanent (marks divide with the cell but can be reprogrammed). Treating enhancers as promoters-with-a-different-name (orientation- and distance-independent vs the fixed start site). Reading "gene ON/OFF" as binary – expression is graded and tissue-timed.

Summary. ~20 000 genes, ~1.5 % coding; enhancers are cooperative combinatorial switches acting by looping; methylation (70–80 % of CpG) and histone marks gate access heritably; one upstream base can rewire a phenotype.

Cell signalling

Why it matters. Cells are not self-sufficient: hormones, neurotransmitters and growth factors are instructions arriving at the membrane. Signal transduction is how a message at the surface becomes biochemistry inside – and roughly half of all drugs act at this layer.

The idea, with a picture. Two canonical cascades: a **GPCR** G (the hormone never enters; a G-protein switch relays to **second messengers** like cAMP) and an **RTK** (receptor phosphorylates itself; adaptor proteins pass the signal to RAS→MAPK growth commands). Both AMPLIFY: one surface event becomes thousands of intracellular events.



One molecule at the membrane becomes tens of thousands inside – amplification is the design

Definitions with numbers. Resting cAMP ~10^{-8} M; a maximal hormone pulse drives it to ~10^{-6} M – a hundredfold switch in seconds. One receptor can activate tens of G-proteins; each adenylyl cyclase makes ~1000 cAMP/s: amplification ~10^4 from hormone to second messenger ✓. GTPase activity of the Gα subunit is the built-in OFF timer (seconds).

The derivation (the amplification product). Signal output = $\prod$ (stage gains): hormone → receptor (1 : ~20 G-proteins) → cyclase (each ~1000 cAMP/s) → PKA (each phosphorylates many targets). The gains MULTIPLY, so the cascade is exponential in gain and, crucially, every stage is a control point: phosphodiesterase destroys cAMP, GTPase switches off Gα, phosphatases strip phosphates. Amplification and reset are the same machinery seen from two sides.

> **Worked example 1 – the cAMP ledger.** Basal 10 nM in a 2×10^{-12} L hepatocyte is 10^{-8} mol/L $\times\, 2\times10^{-12}$ L $\times\, N_A \approx 12\,000$ molecules ✓. One second of 10 activated cyclases at 1000/s would produce 10 000 cAMP – the entire resting pool is replaced in about a second, so the OFF enzymes matter exactly as much as the ON ones.

> **Worked example 2 – clinical link: the switch that cannot turn off.** (Fictional outbreak vignette) profuse secretory diarrhoea within hours. Mechanism: cholera toxin ADP-ribosylates the Gαs subunit, killing its GTPase – the OFF timer is disabled, adenylyl cyclase runs constitutively, cAMP stays at maximum, and the intestinal cell pumps Cl^- (unit 1.2's door) and water into the lumen. Therapy is fluid replacement while the toxin-modified proteins decay: you cannot inhibit a switch that is already welded on.

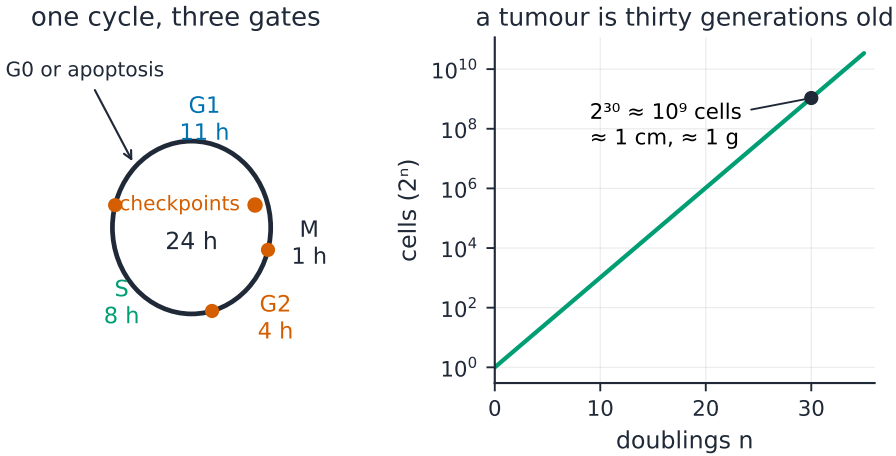
Common mistakes. Saying the hormone "enters the cell" (GPCR ligands do not; steroid hormones do – different mechanism entirely). Forgetting second messengers are destroyed continuously (PDE vs caffeine and other inhibitors). Reading kinase cascades as linear (feedback everywhere).

Summary. GPCR → G-protein → cAMP → PKA; RTK → RAS → MAPK; stage gains multiply to ~10^4; every ON stage has an OFF enzyme; cholera is the GTPase welded open.

Cell cycle, cell death and cancer biology

Why it matters. Division and death are the cell's most expensive decisions, gated by checkpoints. Cancer is not one disease but the accumulated failure of those gates – and every gate is a named molecular system you have already met (repair, signalling, regulation).

The idea, with a picture. The **cell cycle** : **G1** (growth, ~11 h) → **S** (DNA replication, ~8 h, unit 1.3) → **G2** (~4 h) → **M** (mitosis, ~1 h), with **checkpoints** at G1/S (the restriction point), G2/M and the spindle assembly check. Cells can leave to **G0** quiescence – or to **apoptosis** , programmed death, executed by caspases.



Three gates around one cycle – and thirty doublings from one cell to a visible tumour

Definitions with numbers. A typical cycling human cell: 24 h total (11 + 8 + 4 + 1) ✓. **p53** (guardian of the genome) pauses the cycle after DNA damage for repair – or commits the cell to apoptosis. Cancer's requirement: ~4–7 rate-limiting driver events, plus replicative immortality (telomerase) and escape from apoptosis.

The derivation (the doubling count). One cell doubling unchecked: after n divisions there are 2^n cells. A 1 cm^3 tumour masses about 1 g ≈ 10^9 cells: $2^{30} = 1.07\times10^9$ ✓ – a detectable tumour is already THIRTY generations old. The clinical implication is arithmetic: screening intervals and "why did it appear so fast" (it did not – it was silent for years) live on this exponential.

Worked example 1 – the driver-event lottery. Unit 1.3: ~6 mutations survive per division. Most hit passengers (vast non-coding space, unit 1.5); the chance any given mutation disables a specific gate is small – but 10^9 cells × decades buys many trials, each driver raising the rate further. Cancer incidence rising with age is this arithmetic: ~10^{16} divisions in a lifetime, each holding a lottery ticket.

Worked example 2 – clinical link: a family born one gate down. (Fictional pedigree) early sarcoma, breast cancer and a brain tumour across three generations – one inherited mutated p53 allele (Li–Fraumeni pattern). Mechanism: every cell starts one gate down; a single somatic hit on the second allele removes the checkpoint entirely. That is the two-hit logic of a **tumour suppressor gene** , and it ties the course together: repair (1.3), signalling (1.6) and regulation (1.5) are the gates cancer must pass.

Common mistakes. Treating apoptosis as passive death (it is active, programmed, ATP-dependent – necrosis is the messy one). Reading oncogenes and tumour suppressors as the same (gain-of-function vs loss-of-function; one hit vs two). Equating fast growth with fast cycles (usually fewer cells dying, not faster divisions).

Summary. 24 h cycle, three checkpoints, G0 and apoptosis as exits; p53 is the damage gate; $2^{30} \approx 10^9$ cells = 1 cm^3 = thirty silent doublings; 4–7 drivers plus immortality.

Topic 02

University Medicine

What this course is, and how to use it

This is the second course in University Medicine: **medical biochemistry** . MED01 built the cell; this course fills it – the molecules, the enzymes and the fuel flows that keep a human running, and the chemistry that clinical medicine reads every day.

Scope limit, stated once: a study course, not a medical degree; every case below is fictional and nothing here licenses or guides treatment.

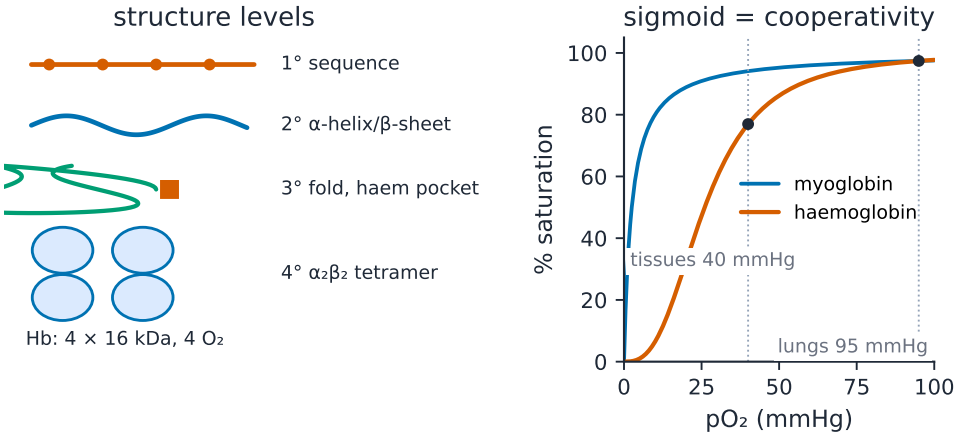
- Units 2.1–2.2 are the machinery: proteins and enzyme kinetics.
- Units 2.3–2.4 are energy: glycolysis and the oxidative engine.
- Units 2.5–2.7 are storage and traffic: glycogen, lipids, nitrogen.
- Unit 2.8 integrates: fed, fasting, and diabetes as one ledger.

The medicine habit: **mechanism before fact** – every enzyme constant, ATP count and fuel flow below is something you recompute, and every clinical link names the broken step.

Proteins and protein structure

Why it matters. Enzymes, transporters, antibodies, receptors – nearly everything a cell does, a protein does it. Structure is not decoration: one level builds on the last, and every level is a place a mutation can break.

The idea, with a picture. Four levels: **primary** (the amino-acid sequence, read N→C), **secondary** (α-helix, β-sheet, held by backbone H-bonds), **tertiary** (the full 3D fold, hydrophobic core), **quaternary** (subunit assembly).



Sequence, fold, 3D shape, assembly – each level is a place a mutation can break

Definitions with numbers. 20 amino acids; peptide bond ~1.33 Å. The reference protein is **haemoglobin** (Hb): 4 subunits ($\alpha_2\beta_2$), 4 haem groups, 64.5 kDa, one O_2 each – up to 4 O_2 per molecule.

The derivation (cooperativity). Myoglobin binds O_2 with a simple hyperbola; haemoglobin binds with a **sigmoid** S because each bound O_2 shifts the subunits from the low-affinity **T** state to the high-affinity **R** state – binding helps binding: **cooperativity**. The Hill coefficient $n_H \approx 2.8$ (max 4) says "almost fully cooperative". Physiological payoff: full loading in lungs ($pO_2 \approx 100$ mmHg), about 25 % unloading in resting tissue ($pO_2 \approx 40$ mmHg) and over 60 % in working muscle (≈ 20 mmHg), because the steep middle of the sigmoid sits exactly in that range.

> **Worked example 1 – the loading arithmetic.** Blood carries ~15 g Hb/dL; 64.5 kDa $\approx$ 64 500 g/mol: $15/64\,500 = 2.33 \times 10^{-4}$ mol/dL = 233 µmol/dL ×4 sites = 930 µmol O_2/dL $\approx$ 20.8 mL O_2/dL (22.4 mL per mmol) – the theoretical 1.39 mL O_2 per g Hb. The measured figure is 1.34 mL/g, 20.1 mL O_2/dL ✓, a little lower because some haemoglobin is in forms that cannot carry O_2.

> **Worked example 2 – clinical link: one charged bead.** (Fictional patient) HbS: one Glu→Val at position 6 of the β chain turns a charged surface patch hydrophobic; in the deoxy T state it docks into a pocket on a neighbour, polymerising the Hb and sickling the cell. Primary structure, one bead, changed quaternary behaviour – and vaso-occlusion follows. Sickle cell disease is a primary-structure disease with quaternary consequences.

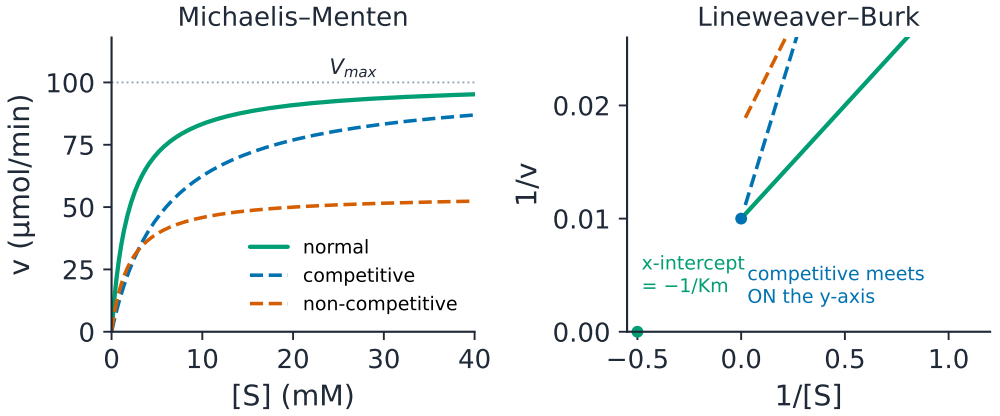
Common mistakes. Reading "denaturation" as breaking peptide bonds (it unfolds the fold; the sequence survives). Treating all mutations as equally severe (surface conservative changes can be silent; core or active-site ones cannot). Forgetting Hb's ligands (H^+, CO_2, 2,3-BPG all shift the curve – physiology lives on those shifts).

Summary. Sequence → fold → 3D → assembly; Hb is the reference tetramer with sigmoid, cooperative binding ($n_H \approx 2.8$); one bead (Glu6Val) can rewire quaternary behaviour.

Enzymes and enzyme kinetics

Why it matters. Nearly every diagnosis on a blood panel is an enzyme activity; nearly every drug is an enzyme inhibitor. Kinetics is how biochemistry becomes measurable – and how you will read the lab.

The idea, with a picture. An enzyme lowers the activation barrier and is unchanged by use. The rate curve is **Michaelis–Menten**: $v = \frac{V_{max}[S]}{K_m+[S]}$, where K_m is the substrate concentration at HALF-maximal velocity.



Saturation curves and their straight-line double-reciprocal transforms – the diagnostic pair

The derivation (why saturation). E + S ⇌ ES → E + P: at low $[S]$, almost all E is free, so $v \approx (V_{max}/K_m) \cdot [S]$ – first order in S. At high $[S]$, every E is in ES: v approaches V_{max} – zero order. The double-reciprocal (Lineweaver–Burk) line $1/v = (K_m/V_{max})\cdot(1/[S])+1/V_{max}$ turns the hyperbola into a straight line: y-intercept $1/V_{max}$, x-intercept $-1/K_m$.

> **Worked example 1 – reading data.** An enzyme gives $v = 33$ µmol/min at $[S] = 1$ mM and $v = 60$ at $[S] = 4$ mM. Write $v = V[S]/(K+[S])$ twice and divide: $33(K+1) = 15(K+4) \Rightarrow K_m = 1.5$ mM, $V_{max} = 33 \times 2.5 = 82.5$ µmol/min ✓. Check the half point: at $[S] = K_m = 1.5$: $v = 82.5 \times 1.5/3 = 41.25 = V_{max}/2$ ✓.

Worked example 2 – clinical link: the competitive blocker. (Fictional patient on a statin-like inhibitor) a drug mimics the substrate and occupies the active site (**competitive inhibition**): apparent K_m rises, V_{max} unchanged – overcome by more substrate, and the L-B lines intersect ON the y-axis. Non-competitive binding elsewhere lowers V_{max} and leaves K_m alone (lines meet ON the x-axis). Reading which line moved IS the mechanism.

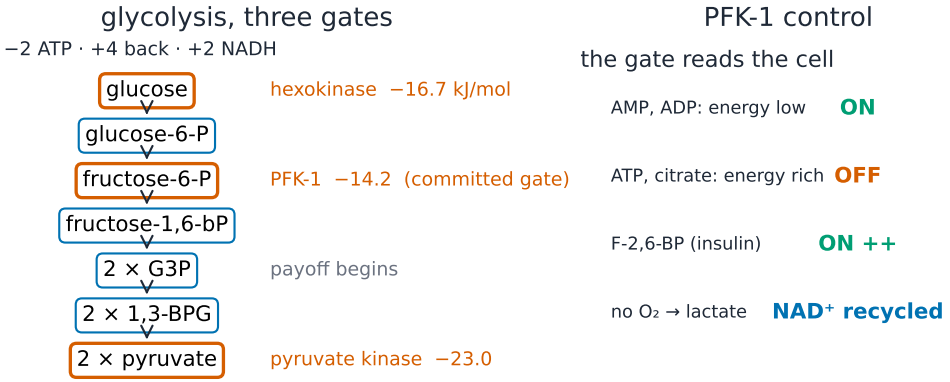
Common mistakes. Calling K_m an affinity (it is a concentration; small K_m happens to mean high apparent affinity). Fitting two points in your head without checking a third. Expecting inhibitors to change enzyme AMOUNT (induction does; inhibitors change activity).

Summary. $v = V_{max}[S]/(K_m + [S])$; $K_m = [S]$ at half-V_{max}; competitive: $K_m\uparrow$ V_{max} same; non-competitive: $V_{max}\downarrow$ K_m same; L-B intercepts identify which.

Bioenergetics and glycolysis

Why it matters. Cells pay for everything in ATP, and the currency exchange is exact: coupled reactions proceed only downhill in free energy. Glycolysis is the oldest, most conserved way to make ATP – and the metabolic anchor of every later unit.

The idea, with a picture. Coupling: an unfavourable reaction runs only when bolted to a favourable one (ΔG adds). Glycolysis splits one glucose (6 C) into two pyruvate (3 C): investment phase (2 ATP spent), payoff phase (4 ATP + 2 NADH made), **net 2 ATP + 2 NADH** ✓ – independent of oxygen.



Invest, pay off, and keep only three gates irreversible – hexokinase, PFK-1, pyruvate kinase

Definitions with numbers. The three essentially irreversible steps ($\Delta G^{\circ\prime}$ strongly negative): hexokinase (−16.7 kJ/mol), **PFK-1** (−14.2, the COMMITTED step and main control), pyruvate kinase (−31.4). Without O_2, NAD^+ is regenerated by **lactate dehydrogenase** : pyruvate + NADH → lactate (fermentation).

The derivation (why PFK-1 is the gate). Control theory says: gate the first committed, irreversible step. Hexokinase is pre-committed (other sugars use it); pyruvate kinase is too late (the investment is already made). PFK-1 is the first step unique to glycolysis – so **AMP and ADP activate it; ATP and citrate inhibit it**; fructose-2,6- bisphosphate (insulin signal) activates it powerfully. One enzyme sees the cell's whole energy state.

Worked example 1 – the ATP ledger. Per glucose: −2 (investment) +4 (payoff) = **net 2 ATP**; plus 2 NADH (cytosolic – worth 5 ATP via the malate shuttle or 3 via glycerol phosphate, unit 2.4). Anaerobically the 2 NADH are consumed re-reducing pyruvate to lactate: the anaerobic yield is exactly 2 ATP ✓.

Worked example 2 – clinical link: the gate that will not open. (Fictional patient) muscle cramps within seconds of sprint, no rise in blood lactate: muscle phosphorylase works (glycogen breaks down) but PFK-1 is deficient – sugar phosphates pile up before the gate, the payoff never comes, ATP falls and myoglobin spills. One gate, one exertional syndrome: the regulated step is the disease locus.

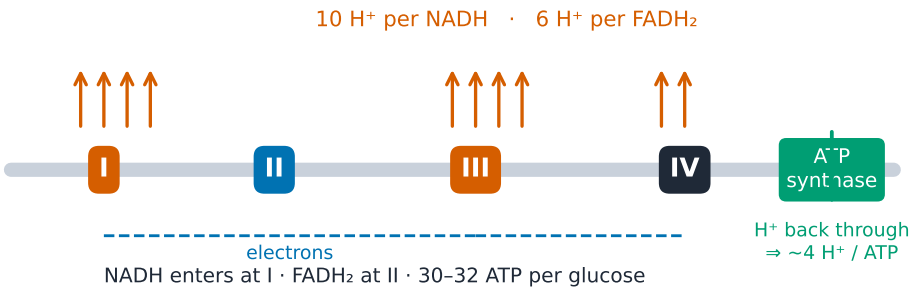
Common mistakes. Counting gross ATP where net is asked. Placing regulation at reversible steps (regulation lives at irreversible gates). Forgetting anaerobic NAD^+ regeneration is WHY lactate forms (not a waste product – a redox necessity).

Summary. Glucose → 2 pyruvate, net 2 ATP + 2 NADH, three irreversible gates with PFK-1 committed; AMP/ADP on, ATP/citrate off; lactate buys NAD^+ when O_2 is out.

TCA cycle and oxidative phosphorylation

Why it matters. Oxygen exists for one job in the body: to accept electrons at the end of the transport chain. The TCA cycle turns acetyl units into those electrons, and the ATP synthase turns their fall into ~90 % of the body's ATP. This unit is where energy metabolism becomes quantitative.

The idea, with a picture. The **TCA cycle** (Krebs): acetyl-CoA (2 C) + oxaloacetate (4 C) → citrate (6 C) → two decarbox- ylations return to oxaloacetate, yielding per acetyl: **3 NADH, 1 $FADH_2$, 1 GTP**. The **electron transport chain** (inner mitochondrial membrane): complexes I, III, IV pump H^+ out; the return flow through **ATP synthase** ATP makes ATP.



Electrons fall to oxygen; the gradient they build is the dam; ATP synthase is the turbine

The derivation (the modern ATP count). Per NADH: complexes I+III+IV pump 4+4+2 = **10 H^+**; per $FADH_2$ (entry at II): III+IV = **6 H^+**. ATP synthase needs ~3 H^+/ATP plus 1 H^+ to import Pi: **4 H^+/ATP**. So NADH ≈ 10/4 = **2.5 ATP**, $FADH_2$ ≈ 6/4 = **1.5 ATP** ✓ – and per glucose, counting every carrier: 2 (glycolysis ATP) + 2×2.5 (its NADH, malate shuttle) + 2×2.5 (pyruvate dehydrogenase NADH) + 2 GTP + 6×2.5 (TCA NADH) + 2×1.5 = **32 ATP** with the malate shuttle, **30** with glycerol phosphate (the older "36–38" used the rounded P/O of 3 and 2).

> **Worked example 1 – acetyl arithmetic.** One glucose → 2 pyruvate → 2 acetyl-CoA (2 NADH from pyruvate dehydrogenase). With pyruvate dehydrogenase, two TCA turns plus entry yield 2 GTP + 8 NADH + 2 $FADH_2$ ✓ = 2 + 20 + 3 = 25 ATP-equivalents – three quarters of the whole aerobic budget (25 of 32), all from electron pairs destined for O_2.

> **Worked example 2 – clinical link: the blocked turbine.** (Fictional overdose vignette) collapse with bright-red venous blood: cyanide binds complex IV, electrons cannot reach O_2, pumping stops, the gradient collapses and ATP synthase stops (it can even run backwards, spending ATP). The chemistry explains the signs: tissues cannot extract O_2 (venous blood stays red), lactate soars (glycolysis alone remains). Treatments exist to bypass or out-compete the block – mechanism first, antidote second.

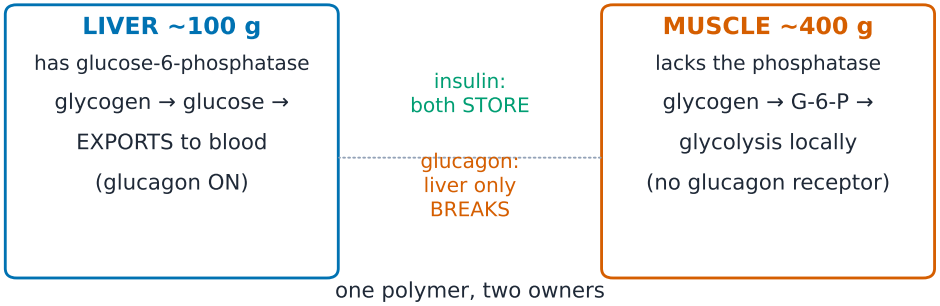
Common mistakes. Saying the cycle "makes ATP" directly (it makes GTP + reduced carriers; the chain makes the gradient). Counting CO_2 as from "the acetyl carbons of this turn" (they leave in LATER turns – the cycle mixes its carbons). Using P/O 3 and 2 without saying which convention.

Summary. Acetyl → 3 NADH + 1 $FADH_2$ + 1 GTP per turn; 10 H^+/NADH, 6 H^+/$FADH_2$, ~4 H^+/ATP ⇒ 2.5/1.5; 30–32 ATP per glucose (modern count, shuttle-dependent); complex IV is where oxygen earns its keep.

Glycogen and gluconeogenesis

Why it matters. The brain burns ~120 g glucose a day and stores none. Between meals the body lives on two glucose banks with different owners – and in fasting it must make glucose from scratch.

The idea, with a picture. Glycogen : branched glucose polymer ($\alpha1 \to 4$ chains, $\alpha1 \to 6$ branch every ~8–12 residues). Liver stores ~100 g – EXPORTS glucose to blood; muscle stores ~400 g – keeps it, burns it locally.



Same polymer, opposite jobs: the liver feeds the blood, the muscle feeds itself

Definitions with numbers. Glycogen phosphorylase cleaves $\alpha1 \to 4$ (phosphorolysis, no ATP spent); debranching enzyme handles the $\alpha1 \to 6$ stubs. **Gluconeogenesis** reverses glycolysis uphill around its three gates – at the cost of 4 ATP + 2 GTP

- 2 NADH per glucose (6 high-energy phosphates, vs glycolysis's net gain of 2).

The derivation (why two banks). Muscle lacks glucose-6-phosphatase: its glucose-6-P is trapped and feeds glycolysis – you cannot export what you cannot dephosphorylate. Liver HAS the phosphatase: glycogen → glucose-1-P → glucose-6-P → free glucose → blood. One enzyme's presence or absence assigns the organ's role; hormones aim it – **glucagon** (liver only) triggers breakdown, **insulin** triggers synthesis.

> **Worked example 1 – the overnight budget.** Liver glycogen ~100 g ≈ 400 kcal; the brain needs ~120 g/day ≈ 480 kcal/day. The bank covers roughly a day's brain use ✓ – after 12–24 h fasting, gluconeogenesis must take over (amino acids from muscle, glycerol from fat, lactate via the Cori cycle).

> **Worked example 2 – clinical link: the runner who cannot tap the bank.** (Fictional) cramps and dark urine after sprinting: muscle phosphorylase is deficient (McArdle pattern) – glycogen is stored but unreachable, ATP falls, myoglobin spills, and the expected lactate rise never comes. The lab IS the mechanism: absent lactate with maximal effort localises the break before the glycolytic gate.

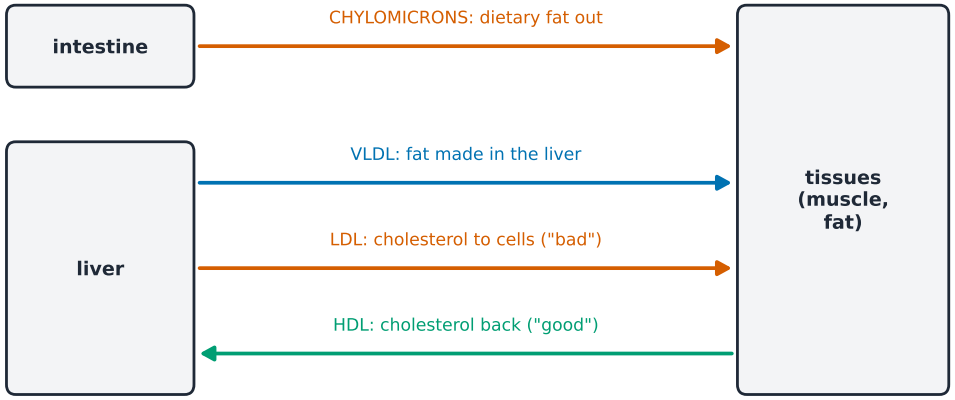
Common mistakes. Treating the two glycogen banks as one (muscle glycogen never becomes blood glucose). Counting gluconeogenesis as "glycolysis backwards for free" (it costs 6 high-energy phosphates). Placing glucagon receptors on muscle (there are none).

Summary. Liver ~100 g exports (has the phosphatase); muscle ~400 g keeps (lacks it); phosphorylase + debrancher break, synthase + branching enzyme build; glucagon/insulin set the sign; gluconeogenesis costs 4 ATP + 2 GTP + 2 NADH per glucose.

Lipid metabolism

Why it matters. Fat is the body's long-term energy warehouse – and its heaviest traffic problem: insoluble in water, lipids need a transport fleet whose failures are the world's most common disease.

The idea, with a picture. **β-oxidation** β in mitochondria: each cycle trims 2 C off a fatty acyl-CoA, yielding 1 acetyl-CoA + 1 $FADH_2$ + 1 NADH. Transport: **lipoproteins** – chylomicrons (intestine → tissues), VLDL (liver → tissues), LDL (cholesterol to tissues – the "bad" carrier), HDL (cholesterol back to liver – the "good" one).



Chylomicrons and VLDL deliver; LDL drops off; HDL returns – the freight map of fat

The derivation (the palmitate count). Palmitate (C16): β- oxidation needs $(16/2) - 1 = 7$ cycles → 8 acetyl-CoA + 7 $FADH_2$ + 7 NADH. Yield: each acetyl is worth 10 equivalents through the cycle ($3 \times 2.5 + 1.5 + 1$), so 8 acetyl give 80; the 7 $FADH_2$ give $7 \times 1.5 = 10.5$ and the 7 NADH give $7 \times 2.5 = 17.5$; minus 2 for activation: **106 ATP per palmitate** ✓ – against 30–32 per glucose: fat pays several-fold per carbon, because it is almost fully reduced.

> **Worked example 1 – ketones, the brain's diesel.** Sustained fasting: acetyl-CoA from β-oxidation outruns the TCA cycle (OAA diverted to gluconeogenesis), so the liver condenses it to **ketone bodies** (acetoacetate, β-hydroxybutyrate) and exports them – water-soluble, crossing to the brain, worth ~21.5 ATP per β-hydroxybutyrate. By day 3 ketones supply about a third of the brain's energy, and after several weeks up to two-thirds, sparing muscle protein ✓.

> **Worked example 2 – clinical link: the engine without a spark.** (Fictional infant) fasting hypoketotic hypoglycaemia: an acyl-CoA dehydrogenase is deficient, β-oxidation stalls, no ketones can be made AND no ATP to drive gluconeogenesis – the double failure explains BOTH lab findings. The diagnostic signature (low glucose WITH low ketones) is the mechanism read off a blood spot.

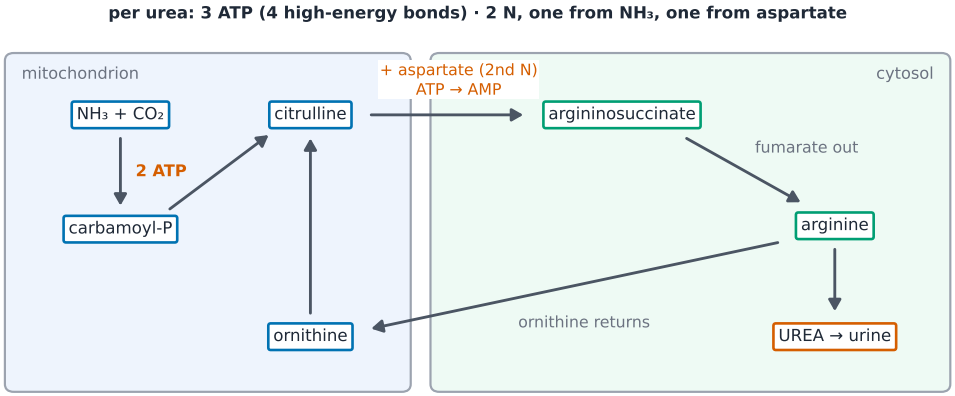
Common mistakes. Forgetting the activation cost (−2 equivalents) in ATP counts. Expecting the brain to burn long-chain fatty acids (it oxidises very little fatty acid – ketones are its substitute). Calling LDL "cholesterol" (it is the CARRIER particle).

Summary. β-oxidation: C16 → 8 acetyl-CoA + 7×($FADH_2$+NADH) → 106 ATP net; chylomicron/VLDL out, LDL deliver, HDL return; fasting makes the liver a ketone exporter.

Nitrogen metabolism and the urea cycle

Why it matters. Nitrogen cannot be stored: every amino acid not needed today must be deaminated today, and the ammonia released is neurotoxic at micromolar concentrations. Disposal is not optional – it is a five-step conveyor across two compartments.

The idea, with a picture. Amino groups funnel to **glutamate** (**transamination** : amino acid + α-ketoglutarate ⇌ keto-acid + glutamate), then oxidative deamination frees NH_3 in the liver mitochondrion. The **urea cycle** runs across mitochondrion and cytosol: NH_3 + CO_2 → carbamoyl phosphate → citrulline → (cytosol) arginino- succinate → arginine → **urea** + regenerated ornithine.



Five enzymes, two compartments, one water-soluble molecule carrying two nitrogens out

Definitions with numbers. Urea: 2 N per molecule (one from NH_3, one from aspartate), cost **3 ATP (4 high-energy bonds)** per urea. Normal blood ammonia is only ~10–35 µmol/L – tight, because the brain has no urea cycle.

The derivation (why the liver, why there). Ammonia detoxication must sit upstream of the systemic circulation: portal blood from the gut drains straight into the liver, which extracts amino acids and ammonia on first pass. The brain and muscle cannot afford the cycle's cost – they export nitrogen as **glutamine** (the ammonia taxi) and **alanine** (the carbon-and-nitrogen taxi) for the liver to process. Location IS the safety system.

> **Worked example 1 – the daily ledger.** A typical diet turns over ~70–100 g protein/day ≈ 11–16 g N; excreted almost entirely as urea (25–35 g/day). Cross-check: each urea carries 2 N (28 g of 60 g molecule = 47 % N by mass): 30 g urea × 0.47 ≈ 14 g N ✓ – the urine matches the diet.

> **Worked example 2 – clinical link: the blocked conveyor.** (Fictional newborn) lethargy then coma after first feeds, sky-high ammonia, orotic acid in urine: ornithine transcarbamylase deficiency – the cycle jams at step 2, carbamoyl phosphate spills into pyrimidine synthesis (hence the orotate). Mechanism-first management follows the block: limit nitrogen intake, scavenge ammonia, feed the alternatives – the biochemistry writes the protocol.

Common mistakes. Writing "urea removes ammonia" (it removes TWO nitrogens: NH_3 and aspartate's amino group). Expecting urea-cycle enzymes in brain (liver-only; that is why ammonia hits the brain). Forgetting the ATP cost (nitrogen disposal is expensive: 4 bonds).

Summary. Transamination funnels N to glutamate; the liver cycle (2 compartments, 5 enzymes) builds urea at 3 ATP each, exporting 2 N; glutamine and alanine are the taxis; hyperammonaemia with orotate points to OTC.

Integration of metabolism: fed, fasting and diabetes

Why it matters. Metabolism is one system with a switch. Insulin and glucagon are the lever, the organs are the gears, and diabetes is the switch jammed – the single most instructive integration in medicine.

The idea, with a picture. Fed state (insulin high): glucose into all cells, glycogen filled, extra carbon to fat. **Fasting state** (glucagon high): liver glycogenolysis (hours) → gluconeogenesis (days) → ketogenesis (brain adapts). One fuel ledger, three time regimes.

FED	insulin high — store	FASTING	glucagon high — release
liver	glycogen ↑, fat ↑	liver	glycogen → glucose out, then gluconeogenesis
muscle	glucose in, glycogen ↑	muscle	export amino acids
fat	take up, store TAG	fat	lipolysis → fatty acids
brain	glucose (120 g/day)	brain	glucose, then ketones

Liver arbitrates, muscle consumes or exports, fat stores or releases, brain demands – insulin/glucagon set the arrows

Definitions with numbers. Brain ~120 g glucose/day. Liver glycogen ~100 g (~a day). Ketones cover about a third of brain energy by day 3 of a fast, and up to two-thirds after several weeks. Diagnostic anchor: **HbA1c ≥ 6.5 %** indicates diabetes – haemoglobin glycation accumulates over the red cell's ~120-day life, integrating glucose like a tape recorder.

The derivation (the switch logic). One variable – is food arriving? – must set EVERY organ's arrows coherently. A hormone pair with opposite actions on the same enzymes (PFK-2, glycogen synthase, phosphorylase, lipase) achieves this: glucagon raises cAMP, so PKA phosphorylates them; insulin lowers cAMP (phosphodiesterase on) and switches on phosphatases, so the same enzymes lose their phosphates. The elegance is that ONE signal, not seven, keeps 4 organs from working at cross purposes – and diabetes shows what cross purposes cost.

Worked example 1 – the timeline of a fast. Hour 0 (meal ends): glucose falls. Hours 0–4: absorbed fuel used, insulin falls. Hours 4–12: liver glycogen drains. Hours 12–24: bank empty, gluconeogenesis ramps (muscle amino acids, glycerol, lactate). Days 1–3: protein sparing begins as ketones rise. Day 3: about a third of brain energy from ketones; weeks: up to two-thirds ✓ – each transition is the previous unit's mechanism, timed.

Worked example 2 – clinical link: the jammed switch, both ways. (Fictional patient) type 1 picture: no insulin, so despite high glucose the cells read FASTING – fat pours out, ketones flood (acidosis, Kussmaul breathing as the lungs compensate), muscle breaks down, urine loses glucose and water (osmotic diuresis, thirst). Every symptom is one switch read by four organs. Type 2 picture: insulin present but resisted – the same ledger, milder arrows, years instead of hours.

Common mistakes. Reading insulin as "lowers glucose" only (it sets the direction of ALL fuels). Treating brain glucose need as negotiable (ketones only partly substitute, and take days). Confusing HbA1c (3-month average) with a glucose reading (minutes).

Summary. Insulin/glucagon = one lever on all organs; fed → store, fast → glycogenolysis → gluconeogenesis → ketones on a ~24 h / 3 day clock; HbA1c ≥ 6.5 % integrates the tape; diabetic ketoacidosis is fasting metabolism running with high glucose.

Topic 03

University Medicine

What this course is, and how to use it

This is the third course in University Medicine: **human anatomy** – the map of the body. Every later clinical course reads this map: where a nerve runs, what a region contains, what one lesion silences.

Scope limit, stated once: a study course, not a medical degree; every case below is fictional, and nothing here licenses or guides treatment. English terms follow **Terminologia Anatomica**; Chinese terms follow the national standard for medical terms.

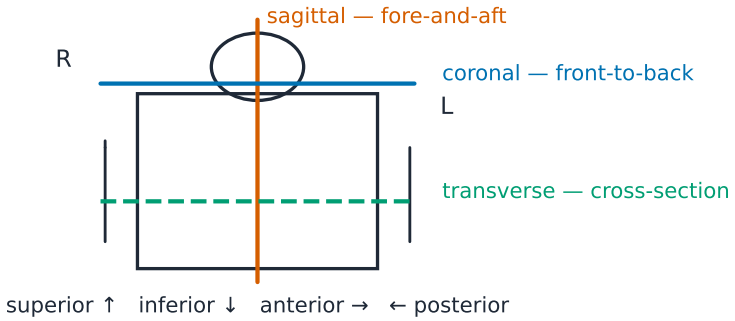
- Units 3.1–3.3 are the language and the limbs.
- Units 3.4–3.6 are the trunk cavities: thorax, abdomen, pelvis.
- Units 3.7–3.8 are head, neck and back.

The medicine habit: **structure explains symptom** – every nerve, vessel and space below is taught with the one lesion that exposes it, and every number (lengths, counts, levels) is stated so you can verify it.

Anatomical language, planes and imaging

Why it matters. Anatomy has a precise language because imprecision has consequences: a "mass on the left" means the patient's left, and a plane is a cut you can point to. Learn the grammar once and every later description is unambiguous.

The idea, with a picture. Position: **superior/inferior** / , **anterior/posterior** / , **proximal/distal** / . Planes: **sagittal** (fore-and-aft), **coronal** (front-and-back), **transverse** (cross-section).



Three planes, one grammar – every scan and section you will ever read is one of these

The derivation (why the grammar pays). Radiology reports are plane statements

("axial T2", "coronal CT") plus side statements. A CT slice is a transverse view displayed as if the patient's feet faced you: the patient's right appears on the LEFT of the image. That single convention, applied consistently, is why two readers agree about a lesion without seeing the patient.

> **Worked example 1 – the directional drill.** The elbow is distal to the shoulder and proximal to the wrist ✓; the sternum is anterior to the heart, which is anterior to the oesophagus ✓. Drill until instinct: each pair replaces a sentence of description.

> **Worked example 2 – clinical link: the wrong side.** (Fictional vignette) a tension pneumothorax is decompressed with a needle at "left, second intercostal space, mid-clavicular line". The grammar IS the safety system: side named first, landmark named in two dimensions, depth by space not centimetres – because "left" means the patient's left, always.

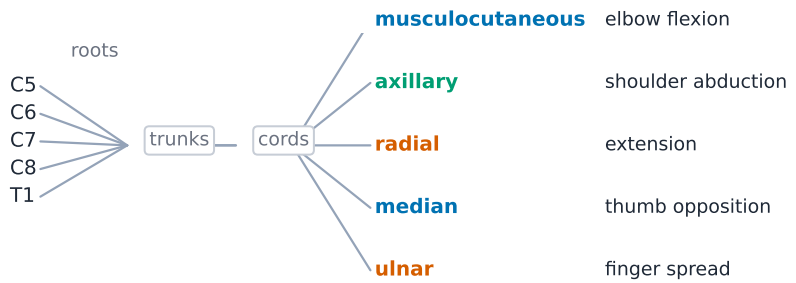
Common mistakes. Mixing observer and patient orientation (the X-ray convention above). Mixing the pairs' reference points (medial/lateral are measured from the midline; proximal/distal from where the limb attaches). Reading "coronal" as "crown down" – it is the plane OF the coronal suture, front-to-back.

Summary. Directional pairs, three planes, one display convention (patient's right on image left); the grammar makes every later description unambiguous.

Upper limb

Why it matters. The arm is the most examined region in preclinical anatomy because its wiring is a readable circuit: one plexus, five terminal nerves, each with a motor test and a sensory patch you can point to.

The idea, with a picture. Roots C5–T1 → trunks → cords → five terminal nerves: **musculocutaneous** (elbow flexion), **axillary** (shoulder abduction), **radial** (extension), **median** (precision grip), **ulnar** (fine finger control).



one nerve, one movement — lesions read as movement patterns

Five roots condense to five nerves – learn each with its motor test and sensory patch

The derivation (why the plexus condenses). Skin segments (**dermatomes**) C5–T1 must map onto a limb that rotates during development: the preaxial border (thumb side, C5–C6) and postaxial (little-finger side, C8–T1) explain why the dermatome strip spirals. The plexus is the remix: one nerve for one movement, drawing whatever roots that movement needs – so nerve lesions follow MOVEMENT patterns, not spinal-segment patterns. That distinction IS the diagnostic method.

> **Worked example 1 – the circuit read out.** Wrist drop = radial (Saturday-night palsy, spiral groove). Ape hand + loss of thumb opposition = median at the carpal tunnel. Claw hand = ulnar at the elbow ✓. Each is a one-line motor test: extend, oppose, spread.

> **Worked example 2 – clinical link: the pulse and the groove.** (Fictional patient) medial elbow pain radiating to the little finger, worse on phone-holding: the ulnar nerve passes BEHIND the medial epicondyle, in a groove you can palpate on yourself – its superficiality there is both the examination (Tinel's tap) and the entrapment. Anatomy you can feel on your own arm is anatomy you keep.

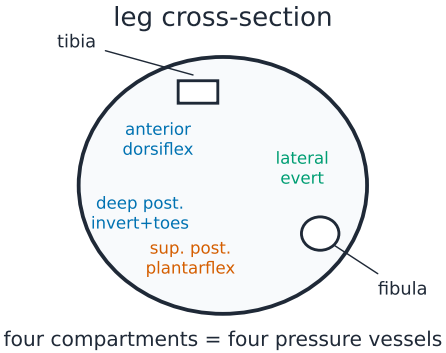
Common mistakes. Learning dermatomes as stripes without the rotation story (they spiral on limbs). Testing muscles by name instead of movement (test extension, not "triceps" first). Forgetting the plexus layers (roots/trunks/cords) exist to explain INJURY PATTERNS, not for their own sake.

Summary. C5–T1 → five terminal nerves; nerve lesions read as movement patterns (radial = extend, median = oppose, ulnar = spread); the ulnar groove is the palpable reminder that superficial anatomy is clinical anatomy.

Lower limb

Why it matters. The leg carries the body for a lifetime: its fascial compartments are pressure vessels, its nerves pass landmarks you can name, and its vessels are the highway between heart and foot.

The idea, with a picture. The thigh: **femoral** (anterior, knee extension), **obturator** (medial, adduction), **sciatic** (posterior → tibial + common peroneal). The leg has four fascial compartments – anterior (dorsiflexion), lateral (eversion), superficial posterior (plantarflexion), deep posterior (inversion + toes).



four compartments = four pressure vessels

the four pulse checkpoints

femoral	hip fold, mid-point quads/knee extension
popliteal	knee crease, deep behind the knee
posterior tibial	behind medial malleolus ankle, medial
dorsalis pedis	foot dorsum foot

Compartments are pressure vessels: each has its nerve, its muscles, its disaster

The derivation (why compartments matter). Fascia does not stretch: exercise swells muscle inside an unyielding sleeve, and pressure above capillary perfusion (~30 mmHg) shuts the bed down – **compartment syndrome** . The anatomy (closed sleeves) explains the physiology (pressure) explains the signs (pain on passive stretch, paraesthesia) explains the treatment (open the sleeve). Pulselessness comes late, if at all: a palpable pulse does not exclude compartment syndrome, because the pressure that kills capillaries is far below arterial pressure.

> **Worked example 1 – the nerve at the neck of the fibula.** The common peroneal nerve wraps the fibular neck, subcutaneous: foot drop with preserved inversion ✓ (tibial intact). The site is palpable – your own leg demonstrates why crossing legs or a tight cast strikes exactly this nerve.

> **Worked example 2 – clinical link: the failing vessel.** (Fictional patient) calf pain after 200 m walking, relieved by standing still: the same artery that serves every compartment narrows proximally – claudication distance is the perfusion budget of the whole limb read off one history. Distal pulses by name (femoral, popliteal, posterior tibial, dorsalis pedis) are the limb's four checkpoints.

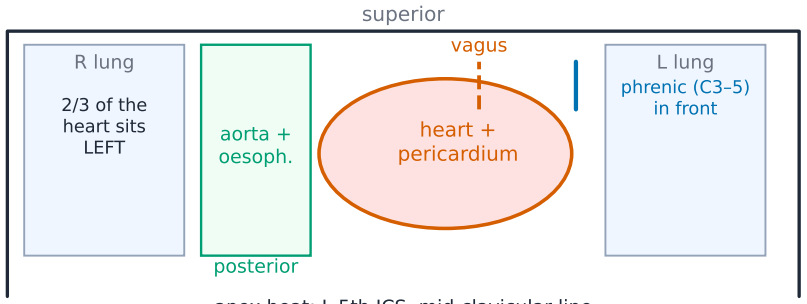
Common mistakes. Confusing dermatome and nerve-territory maps (they differ in limbs). Forgetting eversion (superficial peroneal) vs inversion (deep posterior/tibial) – the pair that localises a lesion. Testing compartments by poking (the sign is pain on PASSIVE stretch).

Summary. Femoral/obturator/sciatic above the knee; four leg compartments below, each with nerve + disaster; peroneal nerve at the fibular neck = foot drop with intact inversion; pulses at four named points.

Thorax

Why it matters. The thorax is the engine room: heart, great vessels, airways and the oesophagus packed in one split box, with a nerve and two pumps crossing between. Nearly every emergency reads this box.

The idea, with a picture. The **mediastinum** between the pleural cavities: superior (above the heart), anterior/middle/posterior around it – middle holds the heart and **pericardium** , posterior holds the descending aorta and oesophagus. The **phrenic nerve** (C3–5) runs anterior to the lung root (diaphragm motor); the **vagus** posterior (parasympathetic to viscera).



apex beat: L 5th ICS, mid-clavicular line
trachea bifurcates at the sternal angle (T4/5)

One box, three divisions – and the two nerves that take different routes past the lung root

The derivation (why the routes diverge). Embryology fixes adult anatomy: the heart descends dragging its pericardium; the phrenic (cervical origin, C3–5) follows the pericardium anteriorly, while the vagus follows the gut tube posteriorly. "Anterior = heart's nerve, posterior = gut's nerve" is a developmental fact that localises tumours: a lung-root mass anteriorly irritates the diaphragm (referred to the shoulder tip, C4 dermatome); posteriorly, it hoarsens the voice (left recurrent laryngeal under the aortic arch).

> **Worked example 1 – the levels that matter.** The trachea bifurcates at the sternal angle (T4/5) ✓; the aortic arch's three branches from right to left: brachiocephalic trunk, left common carotid, left subclavian ✓; the heart's apex beat sits at the left fifth intercostal space, mid-clavicular line ✓ – three numbers that anchor every chest examination.

> **Worked example 2 – clinical link: the tamponade.** (Fictional trauma) hypotension with engorged neck veins and muffled sounds: pericardium is a closed fibrous bag around the heart; 100–200 mL of blood inside it compresses the filling chambers – the anatomy (stiff bag, limited slack) IS the physiology (diastolic collapse) and the procedure (needle below the xiphoid, aimed at the left shoulder).

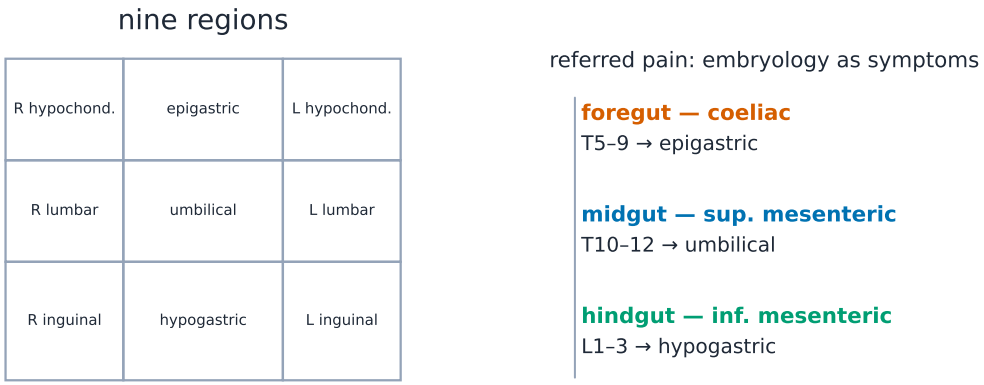
Common mistakes. Putting the heart mid-line (two-thirds left). Confusing phrenic and vagus at the root (anterior vs posterior). Forgetting the sympathetic chain on the necks of the ribs (the route of a chest-wall infection's spread and the landmark for a nerve block).

Summary. Mediastinum: middle = heart, posterior = aorta + oesophagus; phrenic (C3–5) anterior, vagus posterior; bifurcation, branch order and apex beat are the anchor numbers; a closed pericardium makes small volumes lethal.

Abdomen

Why it matters. The abdomen holds the metabolic organs in a nine-region grid, supplied by three unpaired arteries whose territories explain both **referred pain** and every surgical plane.

The idea, with a picture. Surface: nine regions (four quadrants for speed) – **epigastric**, **umbilical**, **hypogastric** columns flanked by hypochondriac, lumbar and inguinal regions. Visceral supply: **coeliac trunk** (foregut to the duodenum), **superior mesenteric** (midgut to the transverse colon), **inferior mesenteric** (hindgut).



Surface grid above, supply map below – one territory per embryological gut division

The derivation (why pain refers where it does). Viscera report to the spinal segments of their embryological origin, not their adult address: foregut → T5–9 (epigastric), midgut → T10–12 (umbilical), hindgut → L1–3 (hypogastric) ✓. Early, unlocalised visceral pain maps to the midline of that band; only when parieties are involved does it lateralise and sharpen. The referred-pain map is embryology read as symptoms.

> **Worked example 1 – the classic referrals.** (Fictional cases) the appendix: midgut → umbilical dull ache first, then right iliac sharpness as the parietal peritoneum joins in ✓; the gallbladder: foregut → epigastric ache first, then right upper quadrant, and right shoulder-tip pain once the diaphragm is irritated (phrenic, C3–5); the ureter: L1 → groin. One embryological sentence explains three bedside pictures.

> **Worked example 2 – clinical link: the bleeding territory.** (Fictional patient) painless bright-red rectal bleeding from internal haemorrhoids: the inferior mesenteric artery ends as the superior rectal artery, and the anal canal's supply and nerves split at the **pectinate line** – above it visceral (painless), below it somatic (exquisitely sensitive). Anatomy predicts which bleeds hurt and which do not.

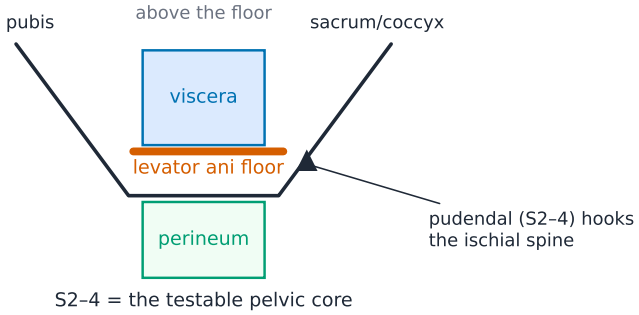
Common mistakes. Quadrants when regions are asked (and vice versa – say which system you use). Forgetting the gut's lymph and venous drainage follows the SAME territories (to the portal system, the coeliac nodes). Mixing renal (T12–L1, lumbar region) with gut pain.

Summary. Nine regions; coeliac/superior/inferior mesenteric = fore/mid/hindgut territories; referred pain T5–9/T10–12/L1–3 to epigastric/umbilical/hypogastric; the pectinate line splits sensation as well as supply.

Pelvis and perineum

Why it matters. The pelvis is the body's floor and its outlet: the bowl of bone, the sling of muscle, and the wiring that shares space with the birth canal and the rectum. Its anatomy explains continence, delivery and the most feared nerve injuries.

The idea, with a picture. Sagittal view: **pelvic floor** = **levator ani** + coccygeus, slung from pubis to coccyx; above it the viscera (bladder anterior, uterus/rectum behind), below it the perineum. The **pudendal nerve** (S2–4) curves around the ischial spine to supply the perineum – the landmark for a block.



A muscle floor carrying viscera; one nerve hooks the spine to reach the floor's outflow

The derivation (why the floor fails). The levator sling carries intra-abdominal pressure for decades; pelvic nerves (S2–4) drive both the bladder's detrusor (parasympathetic) and the external sphincter (somatic via pudendal). Continence is a pressure balance plus two neural switches – so childbirth's stretch (floor) or a spinal lesion (switches) breaks it by two different mechanisms, treated by two different logics.

> **Worked example 1 – the sacral segments read out.** S2–4 = the "pelvic core": detrusor contraction, sphincter control, perineal sensation, and the bulbocavernosus reflex arc ✓. One cone of cord, one exam: a lesion at S2–4 has a signature you can test at the bedside.

> **Worked example 2 – clinical link: delivery's two injuries.** (Fictional delivery) prolonged second stage: the fetal head compresses the pudendal nerve at the ischial spines → transient perineal numbness; a third-degree tear crosses the external sphincter → incontinence risk. Same event, two anatomies, two managements – nerve recovery vs sphincter repair.

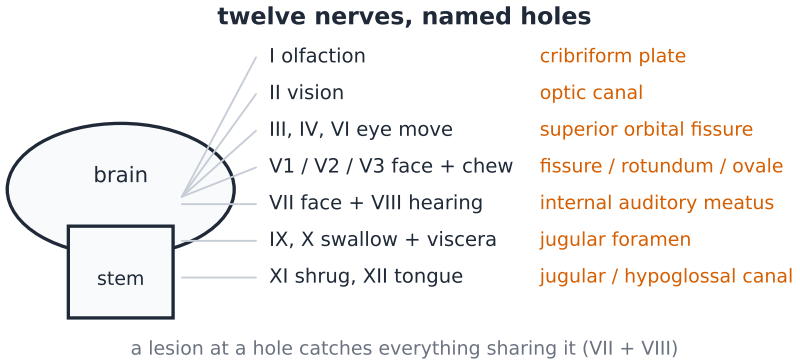
Common mistakes. Sexing the pelvis by one feature (the subpubic angle AND the greater sciatic notch shape both differ). Forgetting that the cavernous nerves for erection run beside the prostate (why prostate surgery puts them at risk). Putting the pelvic floor at the outlet (it is the SUPPORT; the perineum is the outlet region below).

Summary. Levator ani floor carrying viscera, perineum below; pudendal (S2–4) around the ischial spine; S2–4 is the testable pelvic core; delivery injures nerve or sphincter by different mechanisms.

Head and neck

Why it matters. Twelve **cranial nerves** leave the skull through named holes; each hole and each nerve localises a lesion. Head and neck anatomy is the purest structure-localises-disease game in medicine.

The idea, with a picture. The twelve: I smell, II see, III/IV/VI move the eye (VI alone abducts it), V feels the face and chews, VII moves the face and tastes anterior, VIII hears and balances, IX/X swallow and talk (viscera), XI shrugs, XII moves the tongue. Foramina: II through the optic canal, V1 superior orbital fissure, V2 foramen rotundum, V3 foramen ovale, VII/VIII internal auditory meatus, X jugular foramen ✓.



Twelve nerves, named holes – a lesion's signs name the hole

The derivation (why nerves and holes pair up). The skull grew around the brain's outflow: each nerve claimed a channel, and the channel's location fixes what ELSE can be caught there (VII and VIII share the internal auditory meatus – a vestibular schwannoma deafens then paralyses the face in that order). Read the hole, and the overlapping signs follow.

> **Worked example 1 – the VII vs stroke drill.** (Fictional patient) sudden one-sided facial weakness: forehead SPARED = central (upper face has bilateral cortical input) vs forehead INVOLVED = peripheral VII (Bell's pattern) ✓. One muscle test (raise the eyebrows) splits the benign from the emergency.

> **Worked example 2 – clinical link: the apex at the lung tip.** (Fictional smoker) hand wasting, hoarseness, shoulder pain: an apical tumour climbs the cupola and hits, in sequence, the brachial plexus (C8–T1 hand), the recurrent laryngeal (hoarseness) and the sympathetic chain (a drooping eyelid and a small pupil – Horner's). Three systems, one anatomical corner; the corner explains the triad.

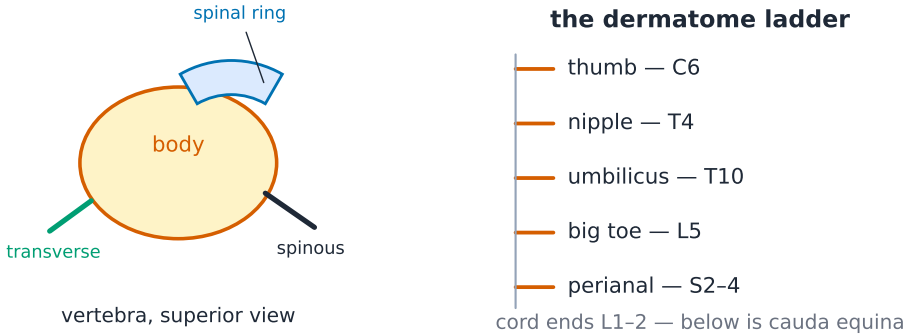
Common mistakes. Counting nerves without their holes (learn in pairs). Forgetting VI exits low and long – the false localiser in raised pressure. Splitting IX from X at the bedside: in the gag reflex IX carries the sensation and X the movement, while "aah" and the palate rise test X alone (the uvula swings away from the weak side).

Summary. Twelve nerves, named foramina; VII/VIII share a meatus; forehead-sparing splits central from peripheral VII; the lung apex is a crossroads of plexus, laryngeal nerve and sympathetic.

Back and vertebral column

Why it matters. The spine is a stack of 33 vertebrae shielding a cord, threaded by nerves at 31 pairs of levels, and held by a design that trades stability for movement at predictable cost – the site of most musculoskeletal pain and every cord emergency.

The idea, with a picture. 7 cervical, 12 thoracic, 5 lumbar, 5 sacral (fused), 4 coccygeal (fused) = 33 bones, 24 movable ✓. Curves: cervical and lumbar lordosis, thoracic and sacral kyphosis. The cord ends at L1–2 (adult); below that the **cauda equina** rides in lumbar cistern – the reason lumbar puncture is safe at L3/4.



One vertebra's ring of bone and ligament; the dermatome ladder that reads a cord level

The derivation (why levels matter twice). In embryos cord and canal match length; the spine grows faster, so adult roots run downward to their exit – the lumbar cord became the cauda equina, and the cone of cord (**conus medullaris**) sits at L1–2 ✓. Two consequences: a T10 bony injury damages the L1–2 cord segment (the gap between bony and cord level grows down the spine: about one segment in the neck, two in the upper thorax, three or more below T7), and below L2 the canal holds only roots (mass effect there is cauda equina syndrome, not cord).

> **Worked example 1 – the dermatome ladder.** Thumb C6, nipple T4, umbilicus T10, big toe L5, perianal S3–5 ✓ – five rungs that turn a sensory exam into a level read. Combined with the level gap, the ladder names both the segment injured and the bone to image.

> **Worked example 2 – clinical link: the emergency at the cord's end.** (Fictional patient) back pain, saddle numbness, retention: a disc or mass at L4–5 compresses roots of the cauda equina – the anatomical end-state (roots, not cord, S2–4 among them) predicts the exact signs (bladder, bowel, saddle), and decompression is measured in hours. The terminology (cauda equina vs cord) IS the urgency.

Common mistakes. Reading bony level as cord level (the gap grows down the spine). Forgetting the disc's posterior longitudinal ligament is strongest midline – discs herniate posterolaterally onto the NEXT lower root (L4/5 disc → L5 root) ✓. Counting sacral vertebrae as movable (5 fused, 4 fused).

Summary. 33/24 vertebrae, alternating curves; cord ends L1–2, so below is cauda equina; bony level ≠ neural level; five dermatome rungs read the level; saddle symptoms plus retention is the hours-count emergency.

Topic 04

University Medicine

What this course is, and how to use it

This is the fourth course in University Medicine: **histology and embryology** . MED03 gave the body's map; this course reads its fabric under the microscope and rewinds to how it was woven – the tissue-level structure that explains organ function, and the developmental sequence whose errors predict the anomalies.

Scope limit, stated once: a study course, not a medical degree; every case below is fictional and nothing here licenses or guides treatment.

- Units 4.1–4.3 are the tissues: preparation, the four basic types.
- Units 4.4–4.5 are organs read as tissue architecture.
- Units 4.6–4.7 are development: fertilisation to birth, and what failure looks like.

The medicine habit: **tissue architecture is organ function** – every layer, lumen and band below is taught with the physiology it enables and the anomaly its mis-making predicts.

Microscopy and tissue preparation

Why it matters. Histology is not "looking at pictures": the picture is manufactured, and every step of its manufacture can create or destroy the evidence. Reading a slide means knowing what the preparation did.

The idea, with a picture. The sequence: **fixation** (formalin cross-links proteins, preserving structure), **dehydration** (alcohols remove water), **embedding** (paraffin supports the tissue), **sectioning** (microtome, 3–5 µm sections), **staining** (haematoxylin–eosin: H&E, -).



Five steps, each with a failure mode – the artefact you can name is the artefact you can discount

Definitions with numbers. Section thickness 3–5 µm (a cell is 10–30 µm – one section slices through part of a cell). Resolution of the light microscope ~200 nm (half the wavelength of visible light): that limit, not the eye, decides what is visible. Magnification beyond ~1500× is empty – blur enlarged.

The derivation (why H&E works). Haematoxylin is basic and stains acidic structures **blue-purple** (nuclei: DNA and RNA are acids); eosin is acidic and stains basic structures **pink-red** (cytoplasm, collagen). The charge pairing is the whole logic: acid-base electrostatics, not magic – and special stains follow the same rule (PAS stains carbohydrates, so basement membranes and glycogen light up magenta).

> **Worked example 1 – the resolution arithmetic.** Two mitochondria 150 nm apart cannot be resolved (below the ~200 nm limit) ✓; a 2 µm bacterium and a 10 µm nucleus both can. The limit explains why electron microscopy (wavelength ~0.005 nm) exists: it buys ~1000× the resolving power, at the cost of killing the specimen.

> **Worked example 2 – clinical link: the frozen section.** (Fictional vignette) a surgeon waits mid-operation for a margin answer: the frozen-section technique skips dehydration/embedding – the tissue is frozen and cut within minutes, stained quickly, read at lower quality. Speed trades artefact: ice crystals tear the tissue, detail is blurred – but a same-hour answer beats a perfect tomorrow answer.

Common mistakes. Treating every gap as a real space (shrinkage during dehydration creates artificial spaces). Reading stain colour as intrinsic colour. Forgetting the section is one plane through a 3-D structure (a tube cut obliquely looks like an oval).

Summary. Fix–dehydrate–embed–section–stain, 3–5 µm sections; light microscopy resolves ~200 nm; H = acidic blue (nuclei), E = basic pink (cytoplasm); frozen sections trade quality for minutes.

Epithelium and connective tissue

Why it matters. Two of the four tissue types carry the body's whole division of labour: epithelium is the boundary workforce (barrier, transport, secretion), connective tissue is the infrastructure (support, transport highway, immune terrain). Organ function is mostly which epithelium sits on which connective tissue.

The idea, with a picture. Epithelium is classified by cell layer (simple one; stratified many) and cell shape (squamous , cuboidal , columnar): simple squamous for exchange (capillary, alveolus), simple cuboidal for transport (tubules), stratified squamous for abrasion (skin, oesophagus). **Connective tissue** = cells + **extracellular matrix** (fibres: collagen , elastic; ground substance), from loose (areolar) to dense regular (tendon) to specialised (bone, blood).

layer × shape

	squamous	cuboidal	columnar
simple	exchange capillary	transport tubule	absorption gut
stratified	abrasion skin	wet barrier ureter	protection conj.

connective spectrum

cells + matrix

loose areolar	cell wrap
dense regular	tendon
cartilage	joint
bone	scaffold
blood	plasma matrix

Layer × shape classifies every epithelium; fibre density and arrangement classifies every connective tissue

The derivation (structure follows function). Exchange needs minimal diffusion distance: one flat layer (simple squamous, ~0.5 µm) ✓. Abrasion defence needs replaceable depth: many layers (stratified squamous) ✓. Transport and secretion need machinery: tall cells packed with organelles (cuboidal/columnar). The classification grid is not taxonomy for its own sake – each cell of the grid predicts the underlying physiology.

> **Worked example 1 – reading a slide.** A section shows a free surface of tall cells with nuclei in a basal row, on a thin basement membrane, over loose matrix with small vessels: simple columnar epithelium on areolar connective tissue – a transport/absorption surface (intestine pattern). Every clue (apical surface, basal membrane, matrix) is a coordinate in the grid.

> **Worked example 2 – clinical link: the barrier that fails.** (Fictional biopsy) recurrent infections, steatorrhoea: the duodenal biopsy shows flattened villi – the absorptive simple columnar epithelium's surface area has collapsed. Malabsorption is not a "functional" mystery: the exchange architecture itself is gone, and the microscope reads the mechanism.

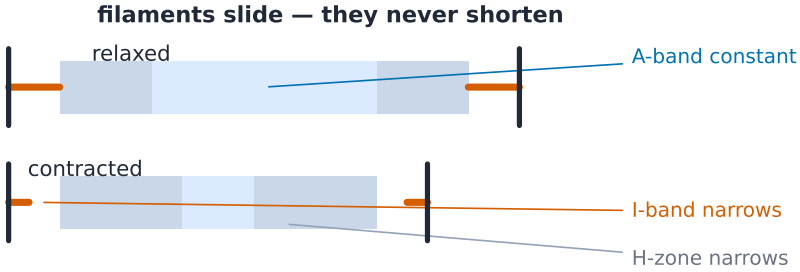
Common mistakes. Confusing the basement membrane (epithelium's anchor) with the whole matrix. Calling blood "not connective tissue" (it is: cells + matrix = plasma). Forgetting pseudostratified looks layered but IS one layer (all cells touch the basement membrane – respiratory epithelium).

Summary. Layer × shape classifies epithelium; each grid cell predicts function (exchange = simple squamous; abrasion = stratified); connective tissue = cells + collagen/elastic/ground substance, from areolar to tendon to bone and blood.

Muscle and nervous tissue

Why it matters. Movement and control are tissues with visible machinery: the sarcomere's banding IS the contraction mechanism, and the neuron's shape IS its connectivity. Histology here is physiology you can see.

The idea, with a picture. **Muscle** : skeletal (striated, voluntary, multinucleate at the periphery), cardiac (striated, involuntary, branched with intercalated discs), smooth (unstriated, involuntary, spindle cells). **Nervous tissue** : neurons (cell body, dendrites, axon) plus **glia** (astrocytes support, oligodendrocytes/Schwann cells myelinate, microglia scavenge).



A-band stays constant (myosin); the I-band and H-zone shrink as the filaments slide – the bands are the mechanism

The sarcomere derivation. A **sarcomere** runs Z-line to Z-line. The **A-band** (anisotropic, dark) is myosin's full length – it does NOT change in contraction. The **I-band** (isotropic, light) is actin-only; the **H-zone** is myosin-only within the A-band. As actin slides toward the centre, the I-band and H-zone NARROW and the Z-lines approach – filament lengths never change ✓. The banding you see is the sliding-filament model printed on the tissue.

> **Worked example 1 – identifying muscle in a section.** Peripheral nuclei + long parallel fibres + striations = skeletal ✓. Central nuclei, branching, intercalated discs = cardiac ✓. No striations, spindle cells, central single nuclei = smooth ✓. Three tissues, three fingerprints, no special stains needed.

> **Worked example 2 – clinical link: myelin read under the scope.** (Fictional nerve biopsy) a child with progressive weakness shows "onion bulbs" – concentric Schwann-cell processes around thinly myelinated axons: repeated de/remyelination made visible. The histology names the disease class (demyelinating neuropathy) because glia are as much "nervous tissue" as neurons.

Common mistakes. Expecting the A-band to shorten (it is myosin's constant length). Confusing intercalated discs with striations (discs are junctions between cells, dark lines across the fibre). Treating glia as packing material (they myelinate, buffer,

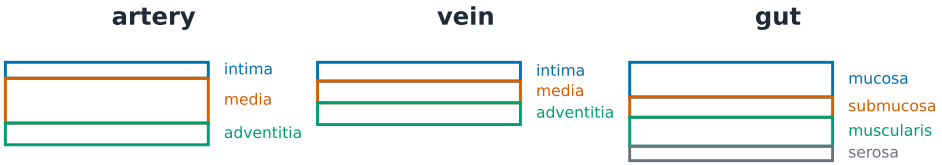
and scar).

Summary. Three muscle fingerprints (peripheral-nuclei striated / branched with discs / unstriated spindle); sarcomere: A constant, I and H narrow, filaments never shorten; glia are functional nervous tissue – myelin pathology is glial pathology.

Organ histology: circulatory, respiratory and digestive

Why it matters. An organ is a tissue arrangement solving a physiological problem – and the three great tube systems (vessels, airways, gut) wear their haemodynamics on their walls. Read the layers and you read the pressures.

The idea, with a picture. Vessels: **tunica intima** (endothelium), **media** (smooth muscle – thick in arteries, thin in veins), **adventitia** (collagen). Airway: the wall thins from bronchus to **alveolus** where the wall becomes simple squamous over capillaries – the blood–air barrier ~0.5 µm ✓. Gut: **mucosa** (epithelium + lamina propria + muscularis mucosae), **submucosa**, **muscularis externa** (inner circular, outer longitudinal), **serosa** or adventitia.



artery = mostly media (pulse) · vein = thin walls · gut = 4 coats

Three tubes, one grammar: lining, muscle coat, outer sheath – each layer named for its job

The derivation (layers track pressure and traffic). Arteries carry high-pressure pulse: a thick elastic/muscular media smooths it ✓. Veins carry low-pressure volume: thin walls, valves, big lumen ✓. The alveolus trades every defence for distance: no muscle, no glands, one flat cell plus its basement membrane and the capillary wall – exchange optimised to the exclusion of everything else ✓. The gut keeps a muscularis externa because its job is propulsion, plus a mucosa that is the body's most renewable frontline.

> **Worked example 1 – artery vs vein at a glance.** The vessel with the thick media nearly filling the wall and a round lumen is the artery; the floppy wide-lumen vessel with a thin wall is the vein ✓. The histological look IS the pressure record – an artery is "mostly media", a vein is "mostly lumen".

> **Worked example 2 – clinical link: the barrier that decides gas exchange.** (Fictional case) breathless with clear airways: the alveolar wall has thickened (fibrosis pattern) – the 0.5 µm barrier has become several µm, and diffusion length, not airflow, now limits oxygen. The histology explains why the inhaler cannot help: the problem is in the wall, not the tube.

Common mistakes. Calling every inner layer "mucosa" (vessels have intima; mucosa belongs to wet epithelial tubes). Forgetting the gut's two muscle coats run in opposite directions (circular squeezers, longitudinal chasers). Expecting alveoli to look like bronchi (they are the opposite: everything removed).

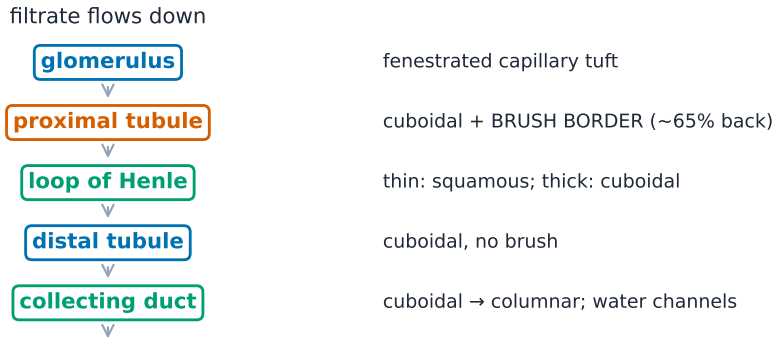
Summary. Vessel = intima/media/adventitia, media ∝ pressure; gut = mucosa/submucosa/muscularis externa/serosa; alveolus = simple squamous

- basement membrane + capillary, ~0.5 µm total; the layer list is the pressure/traffic profile.

Organ histology: urinary, endocrine, reproductive and skin

Why it matters. The kidney's tubule segments, the endocrine's vascular nets, the skin's stratified armour – each is a recognisable histological machine, and together they cover the urine, the hormones and the barrier: three systems whose diseases are read in biopsies every day.

The idea, with a picture. The **nephron**: glomerulus (fenestrated capillary tuft in Bowman's capsule) → proximal tubule (cuboidal with brush border) → loop of Henle (flat squamous thin limbs, cuboidal thick limb) → distal tubule (cuboidal, no brush) → collecting duct (cuboidal to columnar, with pale principal cells). Endocrine organs = cords or follicles wrapped in **fenestrated capillaries** (secreting into blood needs the shortest path). Skin = **stratified squamous keratinised** on dermis, with the epidermis turning over every ~4 weeks.



Each segment wears its job: brush border for reabsorption, flat for exchange, columnar for ducts

The derivation (epithelium predicts tubule function). The proximal tubule reabsorbs ~65 % of the filtrate: it needs surface area – hence a **brush border** (microvilli) that stains brightly ✓. The loop of Henle's thin limbs let water and ions move passively: flat squamous, minimal machinery ✓; its thick ascending limb pumps salt out and is cuboidal, packed with mitochondria ✓. The collecting duct conveys and fine-tunes: cuboidal to columnar, with water channels under hormonal control ✓. You can deduce a segment's physiology from one high-power field.

> **Worked example 1 – biopsy orientation.** A field shows a tuft of capillaries inside a capsule, beside a tubule with a fuzzy violet brush border: glomerulus plus proximal tubule ✓ – renal cortex. The signature pairing names the region without a map.

> **Worked example 2 – clinical link: the four-week frontline.** (Fictional burn vignette) the barrier is lost where the keratinising epithelium died: fluid loss and infection follow immediately, and healing depends on how much **basement membrane** survived, with the hair follicles and glands below it (they are where new epidermis grows from). One layer's survival decides the prognosis – histology as triage.

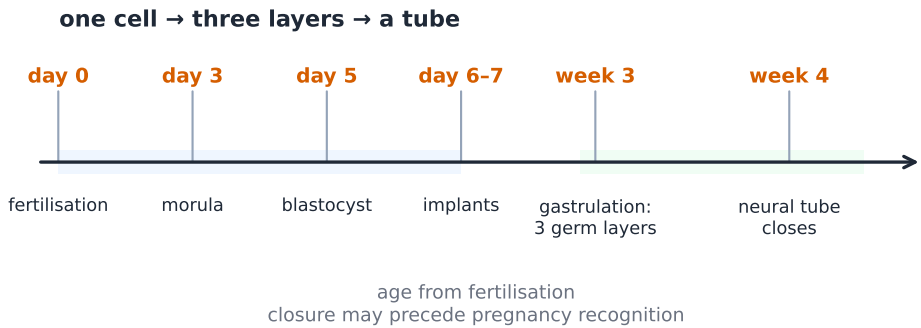
Common mistakes. Expecting every kidney tubule to look the same (the segments are the point). Forgetting endocrine cells sit directly ON capillaries (no ducts – that is what "endocrine" means). Calling the dermis "just connective tissue" (it is the vascular, sensory, structural half of skin).

Summary. Nephron segments read as epithelial signatures (brush = proximal; flat = thin loop; cuboidal to columnar = collecting); endocrine = secretory cells on fenestrated capillaries; skin = keratinised stratified squamous over vascular dermis, ~4-week turnover; basement membrane survival decides regeneration.

Early development: fertilisation to gastrulation

Why it matters. The first three weeks build the body's axes and layers from one cell – with a speed and precision that make them both the most vulnerable and the most instructive period of life. Every later organ trace begins here.

The idea, with a picture. Timeline: **fertilisation** (day 0) → **cleavage** → **morula** (~day 3) → **blastocyst** (~day 5, implants ~day 6–7) → **gastrulation** (week 3: forms the three germ layers) → **neurulation** (week 3–4: the neural tube closes).



Day 0 to week 3: one cell becomes three layers – the whole body is a mapping of these three

The derivation (why gastrulation is the pivot). The blastocyst is a ball of cells; the body is three layered **germ layers** with axes – **ectoderm** (nervous system, epidermis), **mesoderm** (muscle, bone, blood, kidney, heart), **endoderm** (gut, lungs' lining, liver, pancreas) ✓. Gastrulation is the reorganisation that CREATES the layers and the axes; everything after is specialisation. The germ-layer map is why a single week-3 field defect predicts a multi-organ syndrome: tissues are neighbours before they are organs.

> **Worked example 1 – the timeline arithmetic.** Fertilisation to implantation ≈ 6–7 days ✓; to gastrulation ≈ 15 days; to neural-tube closure completion ≈ day 28. Why folate matters periconceptionally: the neural tube closes in week 4 after conception – about two weeks after the first missed period, often before the pregnancy is known or the first antenatal visit – so the protective supplement must start before conception ✓.

> **Worked example 2 – clinical link: the tube that fails to close.** (Fictional case) elevated maternal alpha-fetoprotein, ultrasound shows an open spine: neural-tube closure failed at day 26–28. The defect level maps the deficit (lumbar → legs and bladder; cervical → quadriplegia) – development wrote the prognosis in anatomy.

Common mistakes. Mixing up the blastocyst's two parts (the inner cell mass becomes the embryo; the trophoblast becomes the placenta). Treating the germ layers as a list to memorise (they are a spatial map: tube-in-tube-in-body). Counting pregnancy weeks from conception (clinical weeks count from the last period: ~2 weeks ahead).

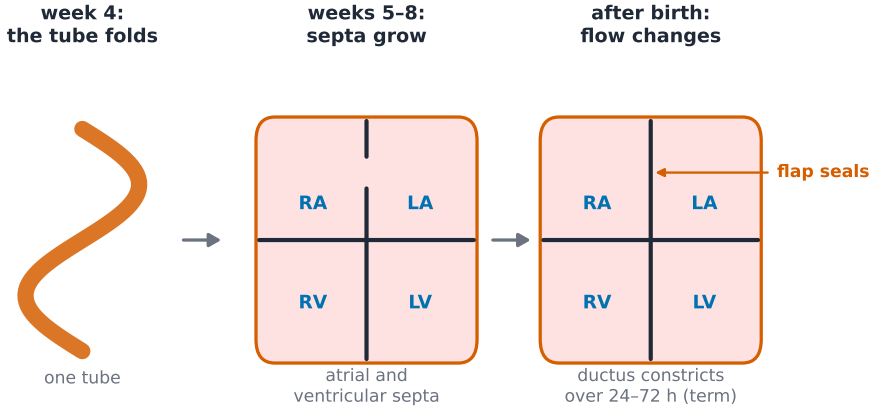
Summary. Day 0 fertilisation → day 5 blastocyst → week 3 gastrulation makes ectoderm/mesoderm/endoderm → week 4 neural tube; the germ-layer map predicts organ systems and syndrome patterns; closure is complete about two weeks after the

first missed period, so folate must start before conception.

Organogenesis and congenital anomalies

Why it matters. Weeks 4–8 turn three layers into recognisable organs – the most error-prone window in human life. Congenital anomaly reasoning is developmental logic read backwards: from a defect, infer the week and the process that failed.

The idea, with a picture. Heart **septation** : the heart tube folds (week 4), then septa close the atrial and ventricular holes – the atrial septum's **foramen ovale** must stay open in the fetus (high pulmonary resistance) and close functionally after birth ✓. Meanwhile: the palate's shelves fuse (weeks 7–10), the diaphragm closes (week 7), the limb buds pattern out (weeks 4–8).



The most error-prone construction project in the body: each septum has its week, and its failure has its name

The derivation (the timing predicts the defect). Organs form in overlapping windows; a teratogen or gene error hits whatever is mid-construction AT THAT WEEK. Weeks 3–4 (neural tube, heart tube): the classic malformation window; weeks 4–8 (septation, palate, limbs): the structural window; after week 8, growth and function mature rather than form ✓. Anomaly reasoning = matching defect to window: defects in organs that were forming in the same weeks (heart septa, palate, limbs) can share one cause, so a cluster is a clue, not a coincidence.

Worked example 1 – the septal arithmetic. The fetal circulation bypasses the lungs: right atrium → foramen ovale → left atrium, and pulmonary artery → ductus arteriosus → aorta ✓. At birth the first breath drops pulmonary resistance, left atrial pressure rises, and the foramen slams shut (functionally immediately, anatomically in months). The ductus arteriosus normally constricts over the first 24–72 hours in term infants; anatomical closure follows later. Patent foramen ovale (persistent in ~25 % of adults) usually matters only when something else forces a right-to-left shunt ✓.

Worked example 2 – clinical link: one cause, several systems. (Fictional child) cleft palate and a heart murmur from a ventricular septal defect. The palate and the ventricular septum were both being built in weeks 4–10, and some causes hit both at once: a teratogen in that window, or a genetic change in a signal both use (the 22q11 deletion is the classic one). So one anomaly is a reason to look for its neighbours: the doctor examining a child with a cleft always listens to the heart.

Common mistakes. Treating "congenital" as "genetic" (most anomalies are multifactorial; many are pure timing). Dating anomalies by birth findings (the error happened weeks earlier). Forgetting the placenta is fetal tissue (its vessels are part of fetal cardiovascular design).

Summary. Weeks 4–8 organogenesis: heart folds then septates, palate fuses, limbs pattern; defects date to their construction windows; the fetal foramen ovale and ductus must be open in utero and close functionally after birth; organs built in the same weeks can share one cause.

Topic 05

University Medicine

What this course is, and how to use it

This is the fifth course in University Medicine: **human physiology** – how the fabric built in MED03–04 actually runs. Physiology is quantitative medicine: pressures, volumes, clearances and concentrations that you compute, not admire.

Scope limit, stated once: a study course, not a medical degree; every case below is fictional and nothing here licenses or guides treatment.

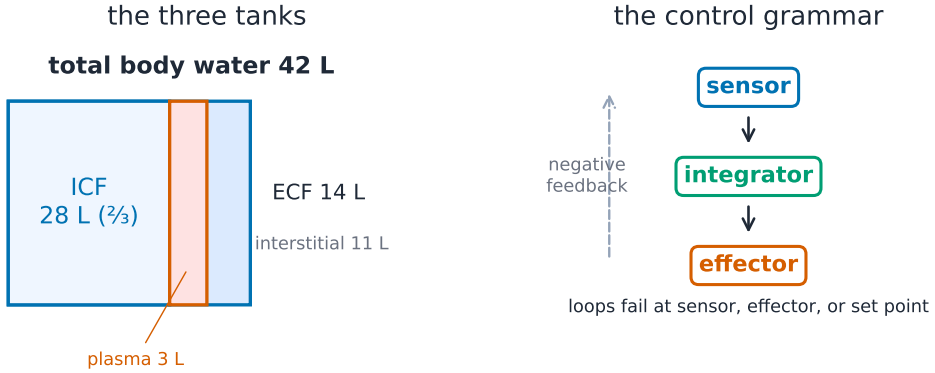
- Unit 5.1 is the frame: homeostasis as control.
- Units 5.2–5.5 are the great systems: excitation, circulation, respiration, the kidney.
- Units 5.6–5.7 complete the map: gut and hormones.

The medicine habit: **every claim is an equation you can evaluate** – cardiac output, alveolar gas, clearance, Henderson–Hasselbalch – and every number below is something you recompute.

Homeostasis and body fluids

Why it matters. The body is a controlled system: its whole design goal is keeping the internal environment constant while the outside swings. Learn the control frame once and every later system is an instance of it.

The idea, with a picture. Homeostasis = sensor → integrator → effector, in a negative-feedback loop. The body water: total ~42 L in a 70 kg adult – **intracellular fluid** 28 L (2/3), **extracellular fluid** 14 L (1/3: plasma 3 L + interstitial 11 L) ✓, separated by selective membranes.



42 litres in three tanks – and one feedback grammar every system speaks

The derivation (why the loop needs a set point). Negative feedback holds a variable near a set point only if the sensor can detect error AND effectors can push both ways. Every physiological control system in this course satisfies that template: baroreceptors sense pressure, the brainstem integrates, heart and vessels effect; chemoreceptors sense CO_2, the medulla integrates, ventilation effects. When you meet a new system, find its three boxes and its set point – the exam question and the clinical failure both live there.

> **Worked example 1 – the water budget arithmetic.** A 70 kg adult: 42 $L = 28$ ICF $+14$ ECF $= 28 + 3$ plasma $+11$ interstitial ✓. Gain 2 L of water: it distributes by compartment number of particles – ECF (with its sodium) takes $1/3 \approx 0.67$ L, ICF 1.33 L ✓. The osmolality maths (ECF tonicity governs the split) is the whole of water balance reasoning.

> **Worked example 2 – clinical link: the failed sensor.** (Fictional patient) a steadily low plasma sodium, yet the urine dilutes and concentrates normally around that lower level: the osmostat's set point has been reset low (the reset-osmostat form of SIADH) – the loop runs perfectly, holding the wrong target. Control loops fail three ways: sensor, effector, set point; diagnosis is deciding which.

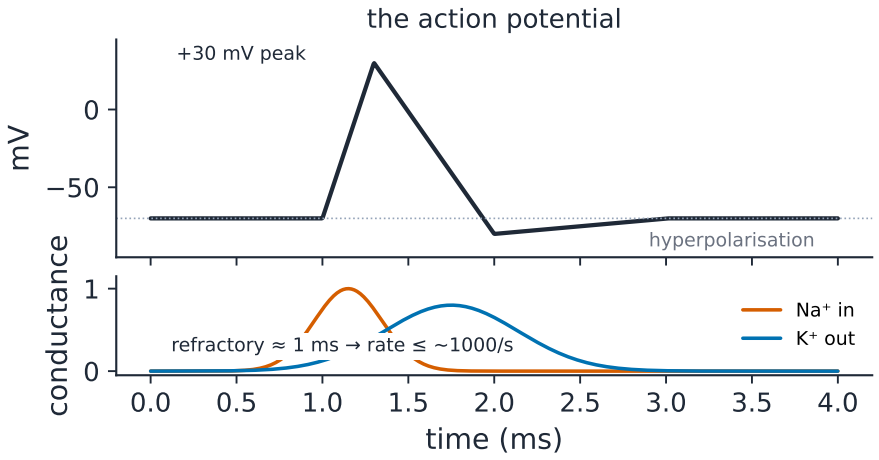
Common mistakes. Forgetting plasma is part of the ECF. Treating "feedback" as automatically negative (positive feedback exists – clotting, labour – and it drives an event to completion rather than holding a steady state). Counting water intake without metabolic water (~0.3 L/day) or losses through skin and breath (~0.9 L/day).

Summary. 42 L = 28 ICF + 14 ECF (3 plasma + 11 interstitial); sensor-integrator-effector around a set point; water distributes by osmolality; loops fail at sensor, effector or set point.

Excitable cells and muscle

Why it matters. Nerve and muscle speak one electrical language; the **action potential** is the sentence. Every rhythm you will ever reason about – heartbeat, breathing, peristalsis – is built from it.

The idea, with a picture. Rest: the permeability-weighted (GHK) membrane potential ≈ -70 mV (K^+ dominates, unit 1.2 of MED01). Excitation: voltage-gated Na^+ channels open → the upstroke overshoots to $+30$ mV → Na^+ inactivates, K^+ carries the repolarisation → a brief **hyperpolarisation**, then rest. The **absolute refractory period** (~1 ms in nerve) sets the maximum firing rate $\approx$ 1000/s ✓.



One sodium upstroke, one potassium repolarisation – and the refractory gap that limits the rate

The derivation (why the upstroke is self-reinforcing). Depolarisation opens Na^+ channels → Na^+ enters → further depolarisation: positive feedback to the spike. Repolarisation is negative feedback (K^+). The two loops in sequence explain the all-or-none shape and the threshold: below threshold, K^+ wins; above it, Na^+ wins ✓. Propagation is the same event regenerating along the membrane – myelin makes it jump (**saltatory conduction**), raising speed from ~1 m/s to up to ~100 m/s ✓.

> **Worked example 1 – from spike to twitch.** In skeletal muscle the spike releases Ca^{2+} from the sarcoplasmic reticulum; Ca^{2+} moves troponin, actin is freed, and the MED04 sarcomere contracts. Summation works because a single twitch (~100 ms) outlasts its refractory (~2 ms): twitches can add. At about 20 stimuli/s they fuse only partly (a bumpy, incomplete tetanus); at about 50/s and above they fuse into **tetanus** – smooth, maximal force ✓.

> **Worked example 2 – clinical link: the channel that fails.** (Fictional) periodic paralysis after a carbohydrate meal: insulin drives K^+ into cells, the extracellular K^+ falls, and muscle fibres with a faulty channel (a small inward leak) depolarise instead of hyperpolarising; their Na^+ channels sit inactivated and the fibres cannot fire – a few millivolts in the wrong direction silences a muscle. One ion, one channel, one sentence of mechanism.

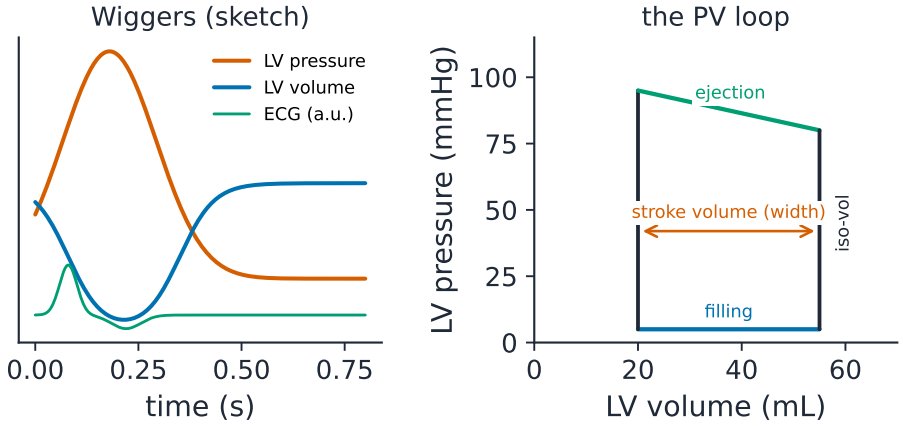
Common mistakes. Confusing refractory (no re-firing) with hyperpolarisation (harder to fire). Expecting local anaesthetics to block small pain fibres last (they go FIRST – rate dependence and small-fibre hit). Forgetting cardiac muscle is a functional syncytium via gap junctions – one long refractory prevents tetanus in the heart.

Summary. −70 mV rest; Na^+ upstroke (positive feedback), K^+ repolarisation; refractory sets the rate ceiling (~1000/s nerve); myelin → saltatory, up to ~100 m/s; Ca^{2+} couples spike to sarcomere; fast enough stimulation (~50/s) fuses twitches into tetanus – but not in the heart.

Cardiovascular physiology

Why it matters. The circulation is a pump, a set of pipes and a volume – the quantitative heart of medicine. Its three equations (output, resistance, pressure) run every shock, every failure and every blood-pressure clinic.

The idea, with a picture. Cardiac output $CO = HR \times SV$: $70 \times 70 = 4900$ mL/min ≈ 5 L/min ✓. Ohm's law of the circulation: $MAP - RAP = CO \times SVR$, so $MAP \approx CO \times SVR$ (RAP ≈ 0), where SVR is the **systemic vascular resistance** . The **pressure–volume loop** - of the left ventricle: filling → isovolumetric contraction → ejection → isovolumetric relaxation, its WIDTH the **stroke volume** , its AREA the external work ✓.



Pressure, volume and the ECG on one clock – and the same beat drawn as a loop whose width is the stroke volume

The derivation (Starling's law as a loop property). Starling's law Starling : increasing preload (filling pressure) stretches myocytes; more overlap-ready filaments eject more – the loop widens, up to the flat top of the curve ✓. Contractility tilts the whole loop (more ejection at the same filling); afterload resists opening the valve (higher pressure before ejection, smaller width). The loop lets you SEE the three determinants of output that the equation hides.

Worked example 1 – the arithmetic of shock. $MAP = CO \times SVR$: haemorrhage drops filling → CO 5 → 3 L/min; reflex vasoconstriction raises SVR by two-thirds; $MAP \approx 3 \times 1.67 \times SVR_{base} = 5 \times SVR_{base}$ – pressure fully preserved while FLOW to organs has fallen 40 % ✓: the cuff reads "compensated", the kidneys read the truth.

Worked example 2 – clinical link: the murmur's timing. (Fictional exam case) a systolic murmur after a myocardial infarct: a new valve leak adds a regurgitant volume – the loop's width splits into forward and backward ejection; the reflex raises HR to hold the 5 L/min forward. Timing (systole) plus loop (volume load) names the lesion class.

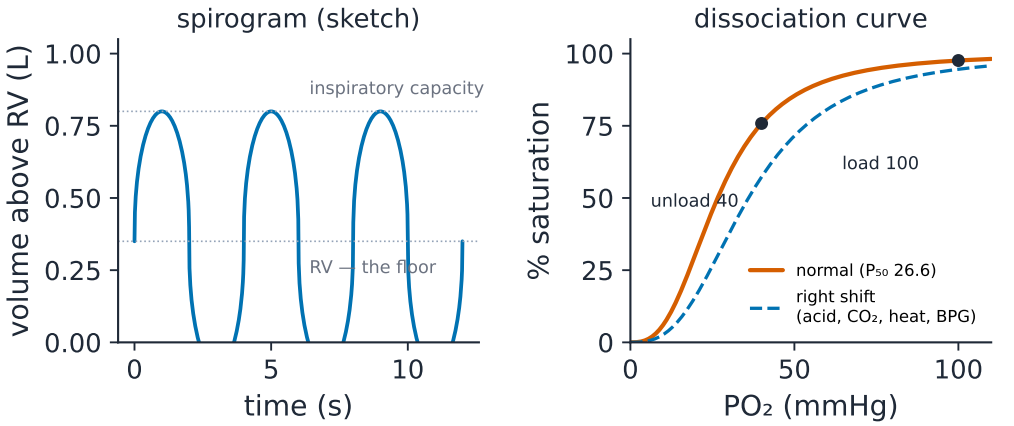
Common mistakes. Multiplying HR by an SV in mL and calling it L/min without converting. Reading MAP as "diastolic + a bit" ($MAP \approx DBP + \frac{1}{3}(SBP - DBP)$ ✓ 80 + 13 ≈ 93 mmHg). Forgetting the baroreceptor reflex senses STRETCH of the carotid sinus – pressure is inferred, and it adapts in 1–2 days.

Summary. $CO = HR \times SV$ (5 L/min); $MAP = CO \times SVR$; loop width = SV, area = work; preload widens (Starling), contractility tilts, afterload narrows; compensation masks flow – trust organ perfusion over the cuff.

Respiratory physiology

Why it matters. Gas exchange is pure applied physics with clinical arithmetic: one equation for alveolar gas, one curve for oxygen carriage, and one law of diffusion. Hypoxaemia reasoning is five-mechanism taxonomy.

The idea, with a picture. **The alveolar gas equation** $P_{AO2} = P_{IO2} - \frac{P_{ACO2}}{R}$: at sea level $P_{IO2} = 0.21 \times (760 - 47) = 150$ mmHg; with $P_{ACO2} = 40$, $R = 0.8$: $P_{AO2} = 150 - 50 = 100$ mmHg ✓. The **oxygen dissociation curve** : sigmoid, $P_{50} = 26.6$ mmHg; shifted right by acid, CO_2, temperature, 2,3-BPG ✓.



One clock of breaths, one curve that loads at 100 and unloads at 40 – and the four shifts that move it

The derivation (why the A–a gradient exists and matters). The alveolar equation says what the alveolus CONTAINS; the arterial blood shows what arrives. Their difference – the **A–a gradient** - , normally 5–15 mmHg (rising with age) – is the fingerprint of lung disease: shunt and V/Q mismatch widen it; hypoventilation and high altitude leave it normal while both still cause hypoxaemia ✓. One subtraction sorts the five causes of low oxygen into "lung" and "not lung".

Worked example 1 – the hypoventilation arithmetic. Opioid overdose (fictional): P_{ACO2} rises to 80 → $P_{AO2} = 150 - 80/0.8 = 150 - 100 = 50$ mmHg ✓: hypoxaemia with a NORMAL A–a gradient – the lung is fine, the bellows are not. Treatment follows the mechanism: ventilate.

Worked example 2 – clinical link: reading the curve at two organs. Lung capillary $P_{O2} \approx 100$ mmHg: saturation 97–98 % ✓. Tissue $P_{O2} \approx 40$: saturation falls to ~75 % – the curve's steep mid-portion releases about a quarter of the content on each pass through the tissue ✓. Anaemia lowers CONTENT at the same saturation: the saturation probe reads normal in severe anaemia – the machine measures binding, not carriage.

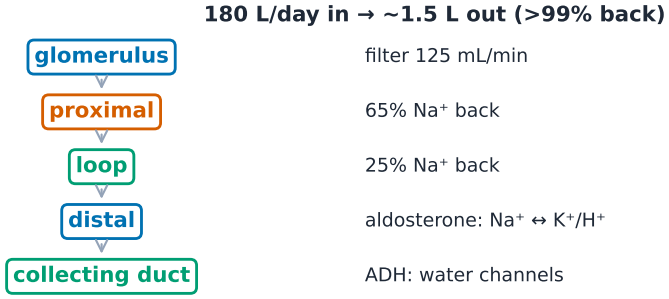
Common mistakes. Assuming $R = 1$ (it is about 0.8; with $R = 1$ the alveolar P_{O2} comes out 10 mmHg too high). Forgetting the water-vapour correction (47 mmHg) – 150, not 160, is the sea-level inspired number. Calling cyanosis an oxygen reading (it needs ~5 g/dL of deoxyhaemoglobin: invisible in anaemia, visible at high saturation in polycythaemia).

Summary. $P_{AO2} = 150 - P_{ACO2}/0.8$; A–a gradient separates lung (shunt, V/Q) from bellows/altitude causes; $P_{50} = 26.6$, right-shifted by acid/CO_2/heat/BPG; saturation ≠ content – anaemia deceives the probe.

Renal physiology and acid–base

Why it matters. The kidney is the body's chemist: it filters, reabsorbs, secretes and regulates. Its clearance arithmetic is the quantitative language of dosing, of kidney function and of every acid–base panel.

The idea, with a picture. Filtration : the **glomerular filtration rate** (GFR) ≈ 125 mL/min (~180 L/day filtered, ~1.5 L excreted – >99 % reabsorbed ✓). The nephron transports segment by segment (MED04's map, now with numbers): proximal 65 %, loop 25 % of the filtrate's Na^+; the distal nephron fine-tunes under **aldosterone** (Na^+ for K^+/H^+) and **ADH** (water channels).



Same segments as the histology, now carrying percentages and two hormonal taps

The derivation (clearance as a measurement). Clearance $C_x = \frac{U_x \cdot V}{P_x}$: the volume of plasma fully cleared of X per minute. Creatinine is freely filtered and barely reabsorbed, so its clearance APPROXIMATES GFR – e.g. $U = 1.2$ mg/mL, $V = 1.4$ mL/min, $P = 0.012$ mg/mL (1.2 mg/dL): $C = 1.2 \times 1.4 / 0.012 = 140$ mL/min (it slightly over-reads, because creatinine is also secreted) ✓. For a substance cleared entirely in one pass (para-aminohippuric acid), clearance = renal plasma flow (~660 mL/min).

> **Worked example 1 – the filtered-load arithmetic.** Plasma glucose 5 mmol/L × 125 mL/min = 625 µmol/min filtered ✓; the proximal transport maximum Tm ≈ 375 mg/min (~2 mmol/min per 1.73 m^2) – above a plasma glucose of ~11 mmol/L (200 mg/dL) glucose starts to appear in urine ✓. Tm/GFR alone predicts 2/0.125 = 16 mmol/L; the threshold is lower because nephrons differ and the weakest saturate first (splay) – Tm arithmetic plus spread, not mystery.

> **Worked example 2 – clinical link: the Henderson–Hasselbalch equation at the bedside.** The **Henderson–Hasselbalch equation** - : $pH = 6.1 + \log_{10} \frac{[HCO_3^-]}{0.03 \cdot P_{CO2}}$ ✓: a metabolic acidosis $[HCO_3^-] = 12$, compensated to pH 7.3 by $P_{CO_2} = 24$: $6.1 + \log(12/0.72) = 6.1 + 1.22 = 7.32$ ✓. One logarithm checks whether the compensation is appropriate (Winter's formula: expected $P_{CO_2} = 1.5 \times 12 + 8 = 26 \pm 2$ ✓) – mismatch means a second disorder hiding in the first.

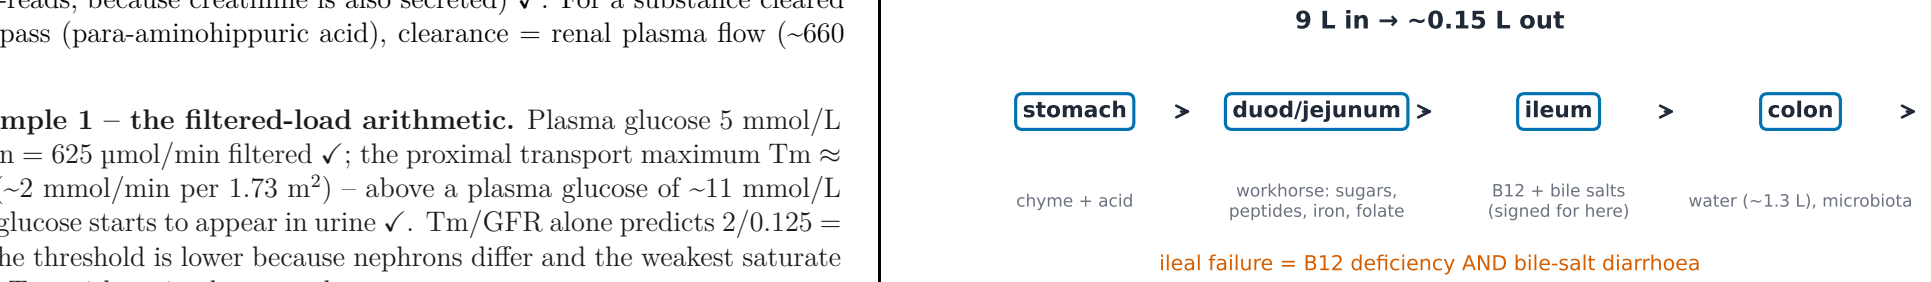
Common mistakes. Reading a "normal" creatinine as a normal GFR (creatinine $\propto$ 1/GFR, a hyperbola: losing the first half of GFR only doubles a low creatinine, which can still sit near the normal range; later losses send it soaring). Confusing clearance of a substance with its excretion RATE. Applying Henderson–Hasselbalch without checking the units of dissolved CO_2 (0.03 × mmHg).

Summary. GFR 125 mL/min, >99 % reabsorbed; proximal 65 % / loop 25 % Na^+; aldosterone and ADH tap the distal end; $C = UV/P$ estimates GFR (creatinine) or RPF (PAH); $pH = 6.1 + \log \frac{HCO_3}{0.03 P_{CO2}}$, with Winter's formula judging the compensation.

Gastrointestinal physiology

Why it matters. The gut is a conveyor with chemistry at every station: motility, secretion, digestion and absorption are staged along its length, and the staging – not any single step – is the teaching.

The idea, with a picture. The map by segment: mouth/stomach (mechanical + proteolysis begins), duodenum/jejunum (the absorption workhorse – iron, folate, most sugars and peptides), ileum (B12 with intrinsic factor, bile salts), colon (water, ~1.5 L/day of the 9 L handled, and microbiota) ✓. Control: the **enteric nervous system** (the "second brain", paced by the slow waves of the interstitial cells of Cajal) modulated by vagus and hormones (gastrin, secretin, CCK).



Nine litres in, 0.1 L out – every segment signs for its cargo

The derivation (why secretion is staged upstream of absorption). Fat digestion needs emulsification (bile) and a near-neutral pH for pancreatic lipase – so the stomach's acid is met by bicarbonate in the duodenum, and the peak of pancreatic

secretion arrives exactly where absorption needs it. Staging explains the diseases: pancreatic insufficiency (no enzymes: fats unabsorbed), ileal resection (B12 and bile salts lost: both deficiency AND a bile-salt diarrhoea as they reach the colon).

> **Worked example 1 – the nine-litre ledger.** In: food and drink 2, saliva 1.5, gastric 2.0, bile 0.5, pancreatic 1.5, intestinal 1.5 = 9 L/day ✓; out: ~8.9 L absorbed (small bowel 7.5, colon 1.4), 0.1–0.2 L in stool ✓. Diarrhoea arithmetic: cholera toxin keeps the gut secreting (MED01's welded switch) – litres of stool with an intact mucosa.

> **Worked example 2 – clinical link: the segment-specific failure.** (Fictional) after ileal resection: macrocytic anaemia (B12 needs the ileum-intrinsic factor complex) plus steatorrhoea (bile salts no longer recycled – the pool shrinks below the fat load). One segment's physiology predicts BOTH deficiencies and their timing (bile-salt pool: days to weeks; B12 stores: years).

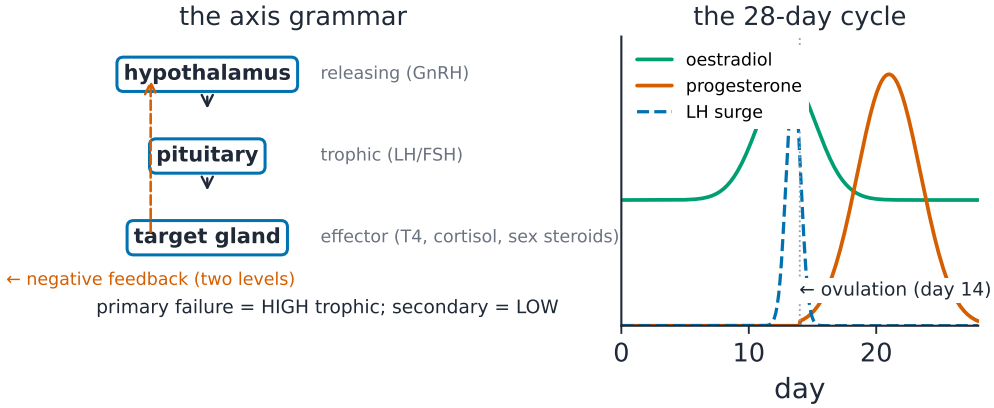
Common mistakes. Expecting the stomach to be the main absorber (alcohol and aspirin aside, it is not). Forgetting the migrating motor complex – the fasting "housekeeper" cycle that stops with feeding. Treating vomiting as a gastric event (it is medullary, with gastric output as the effector).

Summary. 9 L in / ~0.15 L out; jejunum does the bulk, ileum signs for B12 and bile salts, colon polishes water; enteric plexus under vagal/hormonal modulation; segment failure = segment-specific deficiencies.

Endocrine and reproductive physiology

Why it matters. Hormones are the body's slow control layer – signal by axis, act at a distance, feedback at the top. Every axis obeys one grammar, so learning the grammar once retires a dozen syndromes.

The idea, with a picture. The axis: **hypothalamus** (releasing hormone) → **pituitary** (trophic hormone) → **target gland** (effector hormone) → negative feedback to both higher levels ✓. In reproduction: GnRH → LH/FSH → gonads → sex steroids (plus inhibin selectively on FSH); the **menstrual cycle** is that axis in 28-day rhythm, ovulation triggered by the oestradiol-driven LH surge – positive feedback flipping the loop's sign once per cycle ✓.



One grammar: releasing → trophic → effector → feedback – and the one month-per-cycle sign flip that releases an egg

The derivation (why feedback order identifies the lesion). The axis fails at one level; the hormones tell you which: primary (target-gland) failure – low effector, HIGH trophic (feedback released); secondary (pituitary) – low trophic, low effector. One pattern pair, every axis ✓. E.g. (fictional labs) low free T4 with high TSH = the thyroid; low T4 with low TSH = the pituitary/hypothalamus.

> **Worked example 1 – the thyroid arithmetic.** The axis set point: TSH log-linearly tracks free T4 – a small T4 fall doubles TSH ✓, which is why TSH is the screening test: the axis amplifies the primary signal before symptoms exist. Feedback, not sensitivity of assay, is why.

> **Worked example 2 – clinical link: the cycle's two feedback signs.** (Fictional work-up) amenorrhoea with high LH: without ovulation there is little progesterone, and progesterone normally slows the hypothalamic GnRH pulses; faster pulses favour LH over FSH, so LH runs high (the PCOS-pattern logic). One axis diagram plus the sign rule explains the lab pattern and the endometrial risk (unopposed oestrogen).

Common mistakes. Reading single hormone levels without their feedback partner (an unpaired cortisol is nearly meaningless). Forgetting the posterior pituitary stores hypothalamic ADH (it is neural, not a gland fed by trophic hormones). Mixing the two ovarian follicle phases' hormones (oestrogen before, progesterone after the surge).

Summary. Axis: releasing → trophic → effector → two-level feedback; primary failure = high trophic, secondary = low; TSH's log-linear gain makes it the screen; ovulation is one sign flip per cycle – oestradiol's positive feedback driving the LH surge.

Topic 06

University Medicine

What this course is, and how to use it

This is the sixth course in University Medicine: **medical genetics** . MED01 explained the molecule; this course explains its inheritance – how variants travel through families and populations, and how the clinic reads them.

Scope limit, stated once: a study course, not a medical degree; every case below is fictional and nothing here licenses or guides treatment.

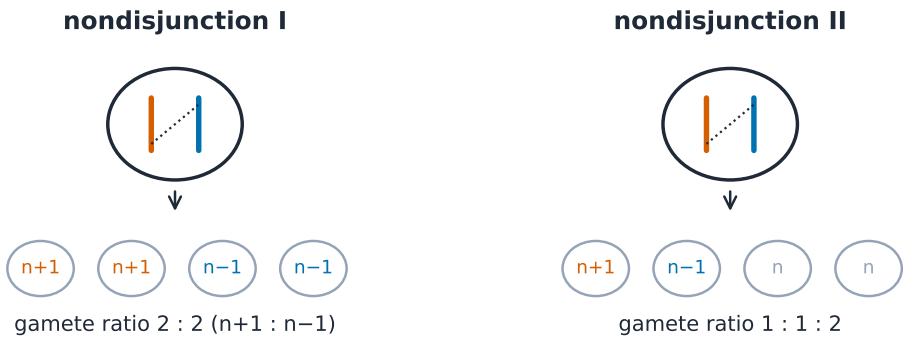
- Units 6.1–6.3 are the transmission rules: chromosomes, Mendel, and the exceptions.
- Units 6.4–6.5 are variation: mutation types and population frequencies.
- Units 6.6–6.7 are the clinic: testing, counselling, and the multifactorial majority.

The medicine habit: **every risk is a probability you compute** – segregation ratios, Hardy–Weinberg frequencies, Bayesian carrier updates – and every number below is something you recompute.

Chromosomes and karyotypes

Why it matters. The chromosome is the unit of transmission, and meiosis is the machine that moves it. When the machine misfires the consequence is a whole-chromosome dosage error – the most common survivable genetic conditions.

The idea, with a picture. The human **karyotype** : 46 chromosomes = 23 pairs (22 autosome pairs + XY/XX) ✓, counted at metaphase (MED04's mitosis). **Meiosis** : one duplication, two divisions → four haploid gametes. **Nondisjunction** – homologs (I) or sister chromatids (II) failing to separate – gives gametes with one too many or too few.



One mis-segregation, two windows (I vs II) – and different gamete ratios from each

The derivation (the 2:2 and the 1:1:2). Division-I nondisjunction: the homolog pair travels together → 2 gametes with n+1, 2 with n−1 – half and half ✓. Division-II: the homologs separate normally at I, then one sister pair sticks at II → 1 gamete n+1, 1 gamete n−1, 2 normal – a **1:1:2** ratio ✓. Down syndrome (trisomy 21, ~47,XX/XY,+21; incidence ≈ 1 in 700 births, rising with maternal age ✓) is the surviving example: about 95 % of cases are a free extra chromosome 21 from nondisjunction, about 90 % of those arise in the egg, and most of those at meiosis I.

> **Worked example 1 – writing a karyotype.** 46,XX = female, standard ✓; 47,XY,+21 = male, trisomy 21 ✓; 45,X = monosomy X (Turner, the only viable full monosomy ✓). The notation is a sentence: count, sex, change. Sex-chromosome aneuploidies (47,XXY; 47,XYY; 47,XXX) are milder than autosomal ones – X-inactivation and gene-poor Y buffer the dosage.

> **Worked example 2 – clinical link: the mosaic window.** (Fictional case) mild features, two cell lines 46/47,+21: nondisjunction occurred in an early mitotic division – a post-zygotic event. Mosaicism explains the milder picture (some tissue is euploid). It also warns the counsellor about the reverse case: a parent can carry trisomic cells only in the ovary or testis (**germline mosaicism**), so a condition can recur in a second child even when both parents' blood karyotypes are normal.

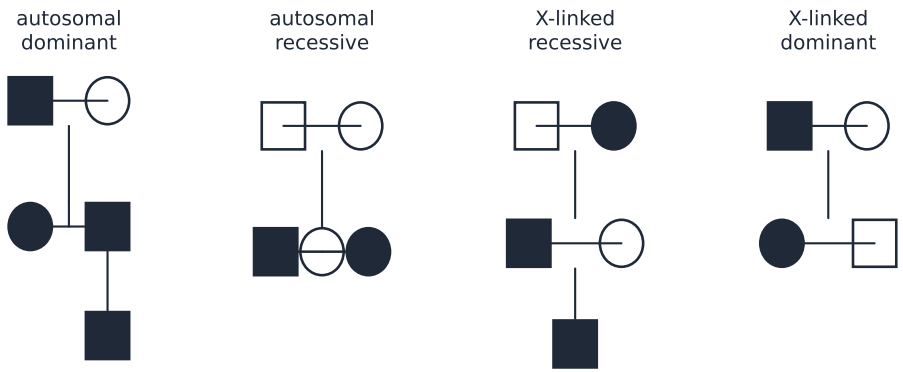
Common mistakes. Confusing meiosis-I and II windows (the gamete ratios differ: 2:2 vs 1:1:2). Calling Down syndrome "familial" (it is usually a de-novo meiotic error; recurrence risk is about 1 %, or the maternal-age risk if that is higher). Forgetting that 45,Y is never viable – X carries essential genes.

Summary. 46 = 23 pairs; meiosis I segregates homologs, II sister chromatids; nondisjunction I → 2×(n+1) and 2×(n−1); II → 1:1:2; trisomy 21 (~1/700) is the surviving autosomal example; 45,X the only viable monosomy; mosaicism softens and complicates.

Mendelian inheritance and pedigrees

Why it matters. Mendel's ratios are the default reasoning engine of genetic counselling: given a family, compute who is at risk. The pedigree is the data structure; the ratios are the arithmetic.

The idea, with a picture. **Pedigree** conventions: squares male, circles female; filled = affected; horizontal line = mating; double line = consanguinity. Four single-gene patterns: **autosomal dominant** (every generation, M:F equal, father-to-son possible; onset often adult), **autosomal recessive** (horizontal clustering, consanguinity, 1/4 of sibs), **X-linked recessive** X (no male-to-male, affected males related through females), **X-linked dominant** X (both sexes affected; an affected father passes it to ALL his daughters and to NONE of his sons ✓).



Four patterns, four signatures – read the squares and circles like a radiograph

The derivation (ratios from gametes). Aa × Aa: gametes A,a × A,a → AA, Aa, Aa, aa = 3:1 affected:unaffected ✓ (dominant); recessive **carriers** Aa × Aa → 1/4 affected. X-linked recessive: carrier mother X^AX^a × normal father → sons 1/2 affected, daughters 1/2 carriers ✓ – a Punnett grid per pattern, and every risk quote in counselling is one of these grids plus penetrance.

> **Worked example 1 – risk down a pedigree.** (Fictional) Huntington-like autosomal dominant: affected parent (Aa) × unaffected (aa) → each child 1/2 at risk ✓. The child of an untested, still-healthy 25-year-old in that family: $1/2 \times 1/2 = 1/4$ ✓, because at 25 the parent's own risk is still about 1/2. Age changes this by Bayes. Suppose (fictional figure) 20 % of carriers are still healthy at 60. For a healthy 60-year-old child of the affected parent: $P(\text{carrier} \mid \text{healthy at 60}) = \frac{\frac{1}{2}\times 0.2}{\frac{1}{2}\times 0.2+\frac{1}{2}\times 1} = \frac{1}{6}$ ✓. "Unaffected so far" is evidence, and it moves the risk from 1/2 to 1/6.

> **Worked example 2 – clinical link: the de-novo dominant.** (Fictional) achondroplasia in a child of average parents: ~80 % of cases are new mutations (paternal age effect) ✓; sib recurrence is near population risk, but the affected child's own children are at 1/2. Distinguishing inherited vs de-novo changes the whole counselling script.

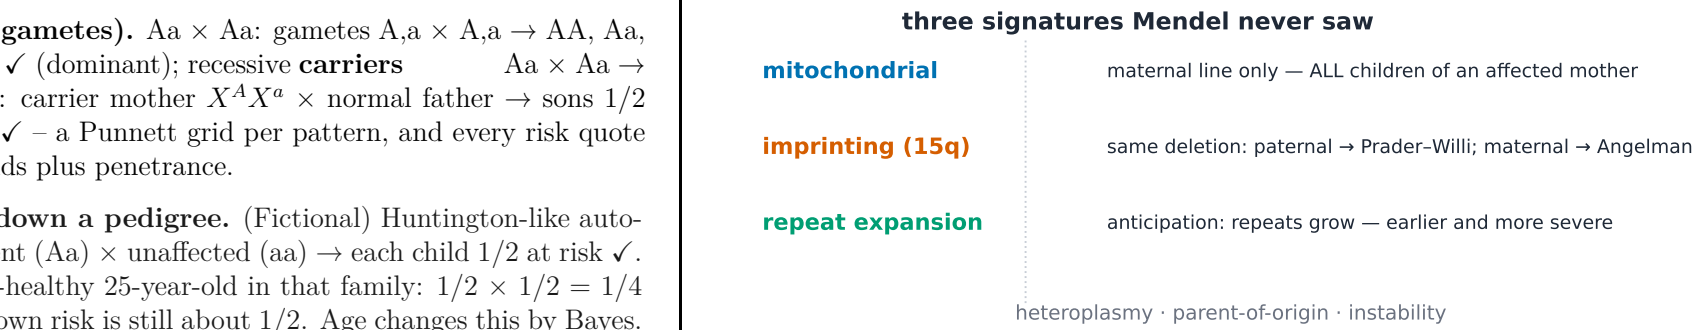
Common mistakes. Equating "dominant" with "common" (dominant diseases are often rare and often de-novo). Forgetting new mutations break the classic pedigree (no affected parents). Assuming small families show ratios (probabilities, not guarantees: two carrier parents have three unaffected children $(3/4)^3 \approx 42$ % of the time ✓, so a recessive condition can hide in a small family).

Summary. Four pedigree signatures; Punnett grids give 3:1, 1/4, 1/2-sons; age-of-onset and de-novo mutation modify raw ratios; risk = grid × pedigree evidence, updated by Bayes.

Non-Mendelian inheritance

Why it matters. The clinic's awkward families – maternal transmission, unequal expression, dosage surprises – are not Mendel failing; they are additional layers of transmission that every modern panel must interpret.

The idea, with a picture. Mitochondrial inheritance : mtDNA is maternally inherited only – all children of an affected mother may be affected, none of an affected father ✓; heteroplasmy (mixed mtDNA populations) makes severity a lottery. **Genomic imprinting** : expression depends on parental ORIGIN (Prader–Willi = paternal 15q deletion, Angelman = maternal 15q deletion – same region, opposite parental origin, opposite syndromes ✓). **Trinucleotide repeat expansion** : **anticipation** (earlier, more severe in successive generations) as repeats grow – Huntington CAG, fragile-X CGG.



Maternal lines, parent-of-origin, and expanding repeats – three signatures Mendel never saw

The derivation (why dosage origin matters). Imprinted genes are SILENCED by methylation (MED01's epigenetics) set in the germline and re-set each generation: only one parental copy works. Delete the WORKING copy → disease, so the same deletion on the paternal chromosome (no active copy, Prader–Willi) differs from the maternal (no active copy of a different gene, Angelman) ✓. Uniparental disomy (both copies from one parent) mimics a deletion: two silent copies.

> **Worked example 1 – the anticipation arithmetic.** (Fictional) fragile-X family: premutation (55–200 CGG) mothers transmit expanded (>200, full mutation) with expansion risk rising with repeat length ✓. Huntington: paternal transmission of large CAG (>60) often gives juvenile onset – the repeat is unstable, growing meiosis to meiosis: anticipation is measurable, not folklore.

> **Worked example 2 – clinical link: the maternal-only line.** (Fictional) a pedigree with affected mothers and children of both sexes, no male transmission over three generations: MELAS-pattern mitochondrial inheritance. Heteroplasmy explains the severity spread (the aunt deaf, the mother diabetic, the cousin encephalopathic) – one mutation, many phenotypes, because each child samples a different bottle-neck fraction of mutant mtDNA.

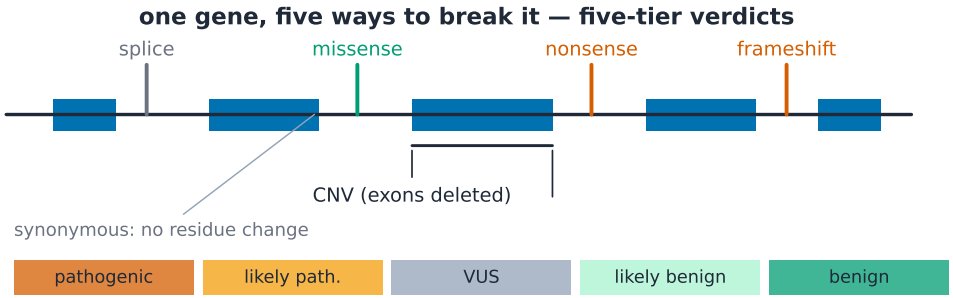
Common mistakes. Expecting mtDNA disease to affect every child of an affected mother equally (heteroplasmy). Confusing imprinting with dominant/recessive (both can be dominant-acting; the parent of origin is the clue). Treating anticipation as ascertainment bias (it is real meiotic instability).

Summary. mtDNA: maternal-only, heteroplasmy lottery; imprinting: same deletion, opposite parent, opposite syndrome (15q); repeat expansions: anticipation, premutation→full-mutation risk by length.

Mutation and variant interpretation

Why it matters. Sequencing found the variant; the clinic needs to know if it matters. Variant interpretation is the daily intellectual core of medical genetics – a graded evidence chain, not a yes/no.

The idea, with a picture. Variant types on a gene model: **synonymous** (no amino-acid change), **missense** (one residue), **nonsense** (stop; usually severe – truncation), **frameshift** (indel not a multiple of 3), **splice** , plus large **copy number variants** (CNVs: deletions/duplications of exons or whole genes). Grading: pathogenic, likely pathogenic, **variant of uncertain significance** (VUS), likely benign, benign – by population frequency, computational prediction, functional data, segregation, and the gene's established disease link ✓.



One gene, five ways to break it – and a five-tier verdict scale

The derivation (why frequency is the first filter). A variant common in healthy populations (e.g. >1 %) cannot cause a rare severe early disease: the arithmetic rules it out before any prediction algorithm runs. So allele frequency in population databases (gnomAD) dominates the first pass; predicted-deleterious-but-common variants are usually benign background ✓. Take a dominant disease affecting 1 person in 10 000 and a variant carried by 1 person in 100. Even if this one variant caused every case, only $\frac{1/10\,000}{1/100} = \frac{1}{100}$ of its carriers could be affected ✓ – so it cannot be a fully penetrant cause.

> **Worked example 1 – the VUS lecture.** (Fictional) a missense VUS in a breast-cancer gene found in an affected woman: the right answer is NOT to treat it as pathogenic. Cascade testing of relatives is not indicated on a VUS; the variant's interpretation may be upgraded later by segregation (does it track disease?) and accumulation of case data ✓. Counselling language: "we found a spelling change; we do not yet know if it means anything."

> **Worked example 2 – clinical link: the loss-of-function rule.** (Fictional) a nonsense variant early in a gene whose established mechanism is **haploinsufficiency** (one working copy is not enough – e.g. PAX6 in aniridia): a truncating variant in such a gene is strong evidence of pathogenicity. The rule: gene mechanism (haploinsufficiency vs dominant-negative vs recessive) predicts which variant TYPES matter in that gene.

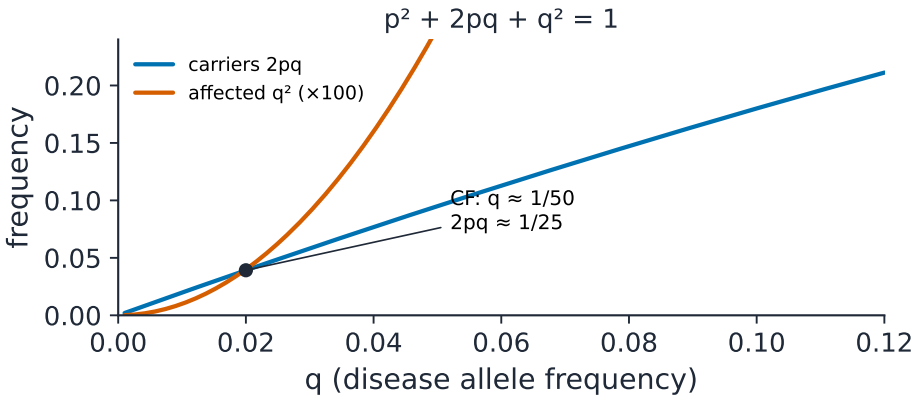
Common mistakes. Treating VUS as "probably pathogenic" (it is explicitly unknown – the commonest counselling failure). Applying dominant-negative logic to haploinsufficiency genes. Forgetting a variant can be pathogenic for one disease and benign for another (or recessive – a carrier without a second allele).

Summary. Five variant classes on a gene model; five-tier verdict scale; population frequency first (Bayes), then prediction, function, segregation; VUS stays VUS until evidence moves it; the gene's mechanism selects which variant types count.

Population genetics

Why it matters. Allele frequencies make carrier risks predictable before any family history – the arithmetic behind newborn screening, carrier panels and ethnic-prevalence counselling. One equation does the work.

The idea, with a picture. Hardy–Weinberg equilibrium - : $p^2 + 2pq + q^2 = 1$ – genotype frequencies from allele frequencies, stable across generations IF mating is random, the population large, and no selection/migration/mutation ✓.



One equation: allele frequencies in, genotype frequencies out – carrier risks included

The derivation (why the binomial appears). Random union of gametes is a binomial draw: two alleles, frequencies p and q, so P(AA) = p^2, P(Aa) = 2pq, P(aa) = q^2 ✓ – the 2 is the two carrier orders. Cystic fibrosis: $q^2 \approx 1/2500$ (Northern European) → q ≈ 1/50 → carriers 2pq ≈ 1/25 ✓ – the famous 4 % carrier rate from a disease rate. Sickle trait in West African populations ~1/4 carriers: q from 2pq, and heterozygote advantage (carriers resist malaria) explains why the allele stays common there although homozygotes are selected against.

> **Worked example 1 – carrier risk arithmetic.** A woman with no CF family history: carrier risk ≈ 1/25. Her partner screened negative (residual risk after a negative panel ≈ 1/300: a panel that finds 92 % of carriers leaves $\frac{(1/25)(0.08)}{1-(1/25)(0.92)} \approx \frac{1}{300}$): affected-child risk = 1/25 × 1/300 × 1/4 ≈ 1/30 000 ✓ – the whole carrier-screening industry is this one multiplication.

> **Worked example 2 – clinical link: the screened residual.** (Fictional) a couple where one is a known carrier and the partner tests negative: risk ≈ 1 × residual(1/300) × 1/4 ≈ 1/1200 ✓. Why not zero? Detection is imperfect (CNVs, rare promoter variants); the residual is gene- and panel-specific – counselling quotes it, never "no risk".

Common mistakes. Solving q = √(frequency) when given a CARRIER rate (carriers are 2pq, not q^2). Applying equilibrium where selection distorts (sickle regions, founder populations). Forgetting founder effects concentrate rare alleles (Tay–Sachs ~1/30 carriers in Ashkenazi Jews vs ~1/300 general ✓).

Summary. $p^2+2pq+q^2=1$; disease rate → q = √(disease) → carrier = 2pq; residual risk after negative screens is quoted, not zero; selection and founder effects are the standard deviations from equilibrium.

Genetic testing and counselling

Why it matters. A test result is information about a family, not just a patient – and each method sees a different layer of the genome. Choosing the test and counselling the result are the clinical skills.

The idea, with a picture. Methods: **karyotyping** (metaphase spread; sees ≥ ~5–10 Mb, balanced rearrangements), **microarray** (CMA: CNVs to ~50–100 kb; no balanced rearrangements ✓), **FISH** (targeted; rapid aneuploidy or microdeletion probe), **sequencing** (Sanger for one gene; NGS panels; exome/genome). Test choice by question: aneuploidy → karyotype/FISH/QF-PCR; developmental delay → CMA; known familial variant → targeted Sanger ✓.

the question chooses the window

karyotype	whole genome, ≥5–10 Mb	balanced rearrangements seen
microarray (CMA)	genome-wide CNVs ~50–100 kb	balanced changes MISSED
FISH	targeted, fast	one probe, one question
sequencing	base-level, chosen scope	Sanger → panel → exome/genome

Four windows on the genome, four resolutions – the question chooses the window

The derivation (why resolution trades scope). A karyotype sees the whole genome at low resolution (large, balanced, structural); sequencing sees base-level detail in the regions it covers; CMA covers the genome at intermediate resolution but MISSES balanced changes (no net dosage). No single test sees everything – the clinical question maps to the method that can answer it ✓. That mapping is the counsellor's first decision, before any consent form.

> **Worked example 1 – the cascade.** (Fictional) a familial pathogenic BRCA1 variant: the proband's positive result creates a TESTABLE risk (1/2 for each first-degree relative) – cascade testing finds carriers for screening and breaks the chain of uncertainty; relatives who test negative revert to population risk for that syndrome ✓. The pedigree multiplies the value of one result.

> **Worked example 2 – clinical link: the incidentalome.** (Fictional) a genome sequenced for developmental delay finds an adult-onset predisposition variant incidentally: the clinic must decide what to return. Consent before testing sets the menu; a VUS-heavy report is counselled as uncertain, not alarming – and children's results defer adult-onset disclosures where possible ✓.

Common mistakes. Ordering sequencing when a CNV is the likely mechanism (and

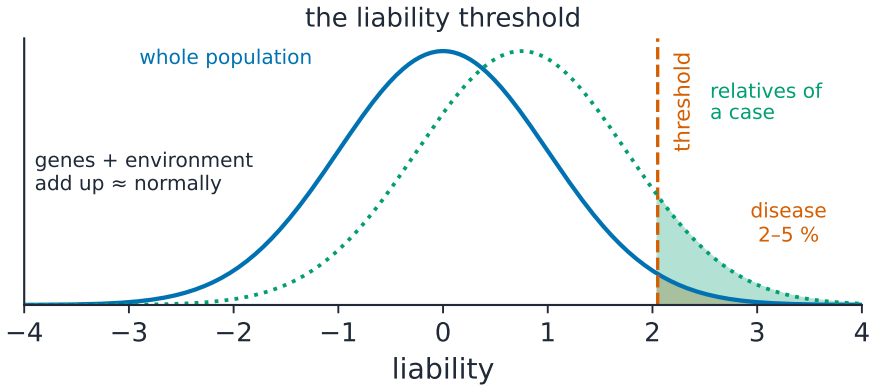
vice versa). Telling a relative who tests negative for the family variant they have "no risk" (they return to population risk, which is not zero). Forgetting pre-symptomatic testing of children is restricted to childhood-actionable conditions ✓.

Summary. Karyotype (large/balanced), CMA (CNV genome-wide), FISH (targeted fast), sequencing (base-level, chosen scope); the question picks the method; cascade testing multiplies one result; consent and the incidentalome are part of the test.

Multifactorial disease and pharmacogenomics

Why it matters. Most common disease – diabetes, hypertension, asthma, schizophrenia – and most drug responses are multifactorial: many genes of small effect plus environment. The clinic needs models that make "it runs in the family" quantitative.

The idea, with a picture. The **liability threshold model** : liability = genetic + environmental contributions (summing approximately normally); disease occurs when liability crosses a threshold ✓. Empiric recurrence risks: 2–5 % for most multifactorial malformations (vs 25 % recessive, 50 % dominant), RISING with (1) more affected relatives, (2) a more severe form in the affected person, and (3) an affected person of the LESS often affected sex ✓. That third rule is the **Carter effect** : pyloric stenosis is about four times commoner in boys, so an affected girl must carry more liability, and her relatives' risk is higher.



A normal curve of inherited+environmental liability; the tail past the threshold is the disease

Pharmacogenomics : genotype-guided prescribing – the best-established examples: CYP2C19 clopidogrel (poor metabolisers cannot activate the prodrug), TPMT/thiopurines (severe toxicity in deficient), HLA-B*57:01 abacavir hypersensitivity (screen before prescribing ✓).

The derivation (why risks rise with each case). If parents of an affected child sit above average liability, additional affected relatives push the FAMILY's inferred liability higher – the whole distribution shifts right, and more siblings cross the threshold: empiric 2–5 % for one affected rises toward 5–10 % with two ✓. This is why multifactorial counselling quotes empiric tables (population studies), not Mendel's ratios: many loci of small effect don't segregate 1:2:1.

> **Worked example 1 – the recurrence quote.** (Fictional) one child with neural-tube defect: empiric recurrence ~2–4 % (folate cuts it); two affected children: ~10 % ✓. The quote is from population tables, not a Punnett grid – and the folate advice (4 mg pre-conception) applies to EVERY subsequent pregnancy regardless of computed risk.

> **Worked example 2 – clinical link: the screened prescription.** (Fictional) a patient starting clopidogrel after a coronary stent: CYP2C19 poor-metaboliser genotype → an alternative antiplatelet drug from day one, since the prodrug is barely activated. About 15 % of East Asians are poor metabolisers (two loss-of-function alleles), against about 2 % of Europeans ✓. HLA-B*57:01 before abacavir: the test PREVENTS the hypersensitivity.

Common mistakes. Quoting 50 %/25 % to multifactorial families (those are Mendel's numbers; multifactorial is empiric 2–5 %). Adding "genetic" and "environmental" risk as percentages instead of liabilities. Treating a pharmacogenomic result as a disease diagnosis (it is a dosing instruction).

Summary. Liability (≈ normal) × threshold = multifactorial disease; empiric recurrence 2–5 %, rising with count and severity; folate regardless of risk; pharmacogenomics (CYP2C19, TPMT, HLA-B*57:01) turns genotype into prescriptions.

Topic 07

University Medicine

What this course is, and how to use it

This is the seventh course in University Medicine: **medical immunology** . The immune system is a distributed defence network with an identity problem: destroy everything foreign, spare everything self. Its failures in both directions are half of internal medicine.

Scope limit, stated once: a study course, not a medical degree; every case below is fictional and nothing here licenses or guides treatment.

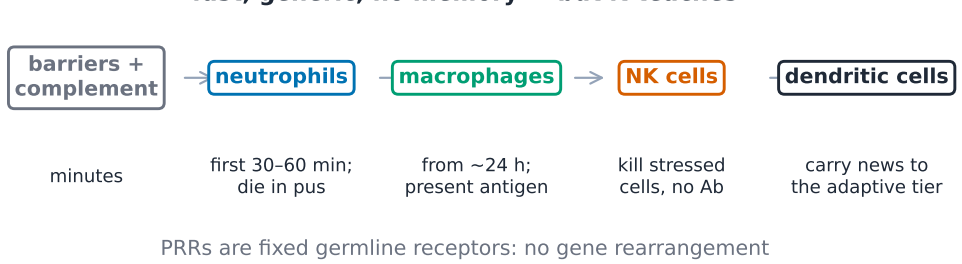
- Unit 7.1 is the fast response: innate immunity.
- Units 7.2–7.4 are the learned response: recognition, activation, memory.
- Units 7.5–7.7 are the failures: allergy, autoimmunity, deficiency, and the transplant boundary.

The medicine habit: **mechanism predicts lesion** – every defect below is traced to the cell or signal that failed, and every number (cell counts, response ratios, timelines) is something you recompute.

Innate immunity

Why it matters. Innate immunity is the response you are born with: minutes to hours, no learning, no memory. It holds the line while the adaptive response gears up – and its cells are the ones that DECIDE whether the adaptive response happens at all.

The idea, with a picture. Barriers (skin, acid, mucus) → the **complement** system (cascade that opsonises and lyses) → **phagocytes** : **neutrophils** (50–70 % of white cells, first to arrive ✓, die in hours), **macrophages** (tissue-resident, present antigen), **natural killer cells** (NK cells; kill virus-infected and stressed cells without antibodies) → **dendritic cells** (the bridge: sample, migrate, present).



Barriers, complement, phagocytes, NK – and the dendritic cell that carries the news to the adaptive system

The derivation (why innate cannot learn but can teach). Innate receptors (**pattern recognition receptors**, PRRs) bind conserved microbial patterns (PAMPs) – fixed germline receptors, no gene rearrangement ✓ (contrast 7.2). But dendritic cells turn an innate detection event into an adaptive instruction: without their co-stimulation, T cells ignore antigen (7.3's two-signal rule). Innate is fast because it is generic; adaptive is precise because it rearranges genes – and the dendritic cell is the main courier between them.

> **Worked example 1 – the arithmetic of an abscess.** A bacterial breach: neutrophils start arriving within the first hour (rolling, then squeezing between endothelial cells, **diapedesis**), dominate the first 6–24 hours and die in the pus; monocytes, which become macrophages, take over at 24–48 hours ✓. The marrow makes about 10^{11} neutrophils a day in health – about 70 million a minute ($10^{11} \div 1440 \approx 7 \times 10^{7}$ ✓) – and can raise that several-fold in severe infection.

> **Worked example 2 – clinical link: the count that warns.** (Fictional) fever with neutrophil count 0.3×10^{9}/L (normal 2–7): chemotherapy has stripped the first responders; the risk is overwhelming sepsis within hours – febrile neutropenia is an emergency precisely because the innate tier is gone while the adaptive tier is too slow.

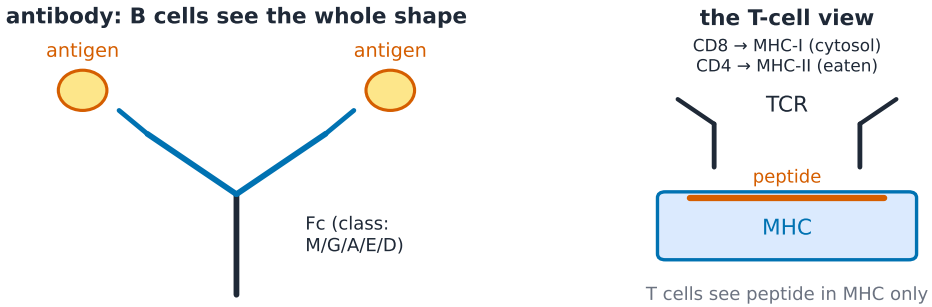
Common mistakes. Calling NK cells adaptive (they are innate lymphocytes – no rearranged receptor). Expecting complement to be an antibody-only system (the alternative and lectin pathways activate without antibody ✓). Forgetting eosinophils/basophils (parasites and allergy, 7.5).

Summary. Barriers → complement → neutrophils (fast, suicidal) → macrophages (sustain, present) → NK (innate killing) → dendritic cells (the courier to 7.2–7.3); PRRs are fixed receptors – no memory, but full teaching rights.

Antigen recognition: antibodies, receptors and MHC

Why it matters. Recognition is the immune system's core engineering: gene segments are shuffled to build receptors that can bind anything – and MHC molecules display samples of what is inside every cell for inspection. Both ideas explain vaccination, transplantation and most immunological tests.

The idea, with a picture. The **antibody** (immunoglobulin, Ig): two heavy + two light chains; V regions bind antigen, C regions define class (IgM first, IgG memory/transfer, IgA mucosa, IgE allergy, IgD receptor). **T-cell receptors** T (TCRs) on T cells see antigen ONLY as peptides held by **major histocompatibility complex** (MHC) molecules on cell surfaces – MHC-I (nearly all cells, displays cytosolic samples, CD8 inspects) and MHC-II (professional presenters, displays endocytosed samples, CD4 inspects) ✓.



The B cell sees the whole shape; the T cell sees a peptide in an MHC groove – two recognition chemistries

The derivation (how near-infinite diversity is built). Ig genes exist as segment libraries: V, D and J segments are recombined (combinatorial joining), junctions are trimmed and random nucleotides added (junctional diversity). Fewer than 200 working segments give about 2×10^6 heavy–light combinations by joining alone, and junctional diversity takes the possible total past 10^{11} ✓; one person carries perhaps 10^7–10^8 different clones at any time. Each B cell expresses ONE specificity, and antigen picks which clones grow (**clonal selection**); T cells undergo the same rearrangement in the thymus. The cost: the system cannot pre-know self – hence tolerance checkpoints (7.6).

> **Worked example 1 – the class switch arithmetic.** IgM pentamer: 10 binding sites but low affinity; after class switching and somatic hypermutation (in the germinal centre), IgG affinity rises 10- to 1000-fold over weeks to months ✓ – the titre rise a vaccine series exploits. This **affinity maturation** is why dose 2 given 4–8 weeks later yields higher-avidity antibody than dose 1 (7.4's curve).

> **Worked example 2 – clinical link: the mismatched 3 of 6.** (Fictional transplant) HLA-A, -B, -DR are the typed loci; each parent contributes one haplotype, so siblings have a 25 % chance of a full match (2 identical haplotypes: $1/2 \times 1/2$) ✓. The more mismatches, the higher the risk of rejection (7.7): T cells see mismatched MHC directly as foreign – MHC is, by design, the most polymorphic gene family in the body.

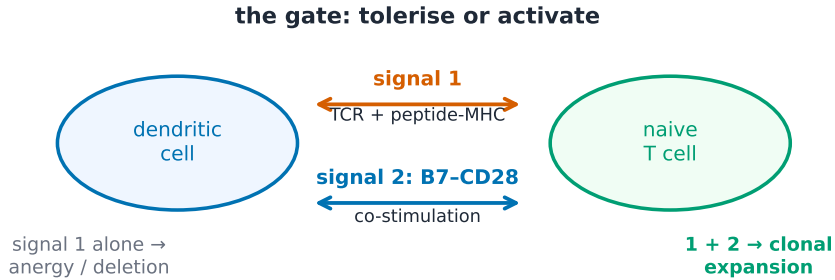
Common mistakes. Saying antibodies kill (they neutralise and mark – complement and phagocytes kill). Expecting T cells to see free antigen (they never do – peptide + MHC only ✓). Forgetting MHC-I displays what a cell is MAKING (cytosol) and MHC-II what it has EATEN (endosome).

Summary. Two recognition chemistries: Ig/BCR (free antigen, whole shape) and TCR (peptide in MHC); V(D)J recombination + junctional diversity → over 10^{11} possible receptors; IgM→IgG class switch with 10- to 1000-fold affinity maturation; MHC-I/CD8 cytosolic, MHC-II/CD4 endocytosed; sibling full-match 1/4.

Adaptive responses: T and B cells

Why it matters. Adaptive responses are the system's precision tier: clonal expansion, division of labour, and regulation – and their failure modes (7.5–7.7) all trace to signals described here.

The idea, with a picture. T cells: CD4 "helpers" (macrophage activation, B-cell help, cytokine orchestration) vs CD8 "killers" (cytotoxic granules, Fas–FasL). **B cells**: antibody factories (plasma cells: thousands of molecules per second ✓) and memory. **The two-signal rule**: antigen binding alone tolerises; full activation needs signal 1 (TCR + peptide-MHC) + signal 2 (**co-stimulation**, e.g. B7–CD28) from an activated dendritic cell ✓.



Signal 1 alone tolerises; signal 1 + co-stimulation activate – the decision gate of every adaptive response

The derivation (why the gate exists). A receptor library this large MUST bind some self – so evolution put the decision at activation: signal 1 alone (a quiet cell presenting self) → anergy or deletion; signal 1 + signal 2 (danger-PAMPs driving B7

expression) → clonal expansion ✓. The gate explains vaccination (adjuvants PROVIDE signal 2 – that is their whole job), autoimmunity (bypassed when tissue damage makes self look dangerous) and transplant rejection (injury at harvest supplies signal 2).

> **Worked example 1 – the clonal arithmetic.** Only about 1 naive T cell in 10^5–10^6 recognises a given viral peptide. Once activated it divides every ~6 hours; after 3 days, $2^{12} \approx 4000$ daughters; after 5 days ~$2^{20} \approx 10^6$ ✓ – the proliferation that makes an adaptive response 4–7 days late but overwhelming. CD8 expansion during acute infection can be >1000-fold of the specific population: 10^4-fold or more.

> **Worked example 2 – clinical link: help and the broken bridge.** (Fictional) X-linked hyper-IgM: CD40L is missing, so T cells cannot give B cells the CD40 signal – IgM accumulates but class switching fails; the patient suffers pyogenic infections (no IgG) AND opportunistic ones (no macrophage activation). One molecule, both arms of the T-helper role severed: helper signals are wiring, not decorations.

Common mistakes. Equating "helper" with gentle (TH1 activates macrophages to kill intracellular bugs – aggressive help). Forgetting **regulatory T cells** T (Tregs, the fourth population: suppressor gate, 7.6's tolerance). Timing errors: primary response peaks day 7–14; a positive test on day 1 is unlikely to be the acute antibody ✓.

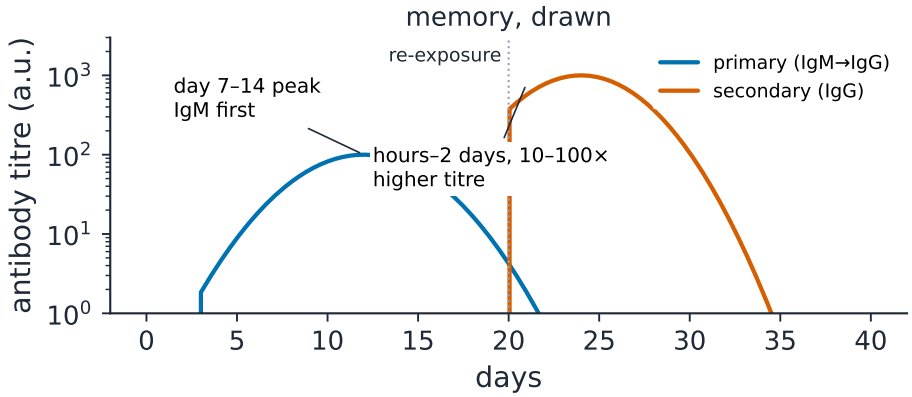
Summary. CD4 helps, CD8 kills, B secretes and remembers; signal 1

- co-stimulation (B7–CD28) is the gate that tolerises or activates; clonal expansion from one cell to 10^6 in days; CD40L–CD40 bridges T help to B switching; adjuvant = signal 2.

Immunological memory and vaccination

Why it matters. Immune memory is the adaptive system's dividend: a second encounter is faster, higher and better-targeted. Vaccination is memory engineered deliberately – the single most cost-effective medical intervention ever measured.

The idea, with a picture. The **primary** response: lag 4–7 days, IgM then IgG, peak ~day 10–14, low titre. The **secondary** response: lag hours–2 days, IgG immediately, titre 10–100× higher, higher affinity ✓ – memory B and T cells wait at higher frequency and in better locations (lymph node memory, tissue-resident memory).



First exposure: slow, IgM-first, modest. Second: immediate, IgG, 10–100× – memory, drawn

The derivation (why memory changes the curve). Memory cells skip the naive-to-activation gate partially (higher sensitivity to antigen, faster division) and exist at 10–100× the naive frequency ✓; affinity maturation pre-tuned in the first response's germinal centres. **Herd immunity** follows from the same curve: transmission breaks when the susceptible fraction falls below $1/R_0$, so the immune fraction must exceed $1 - 1/R_0$ – for measles ($R_0 \approx 12$–18) that is ~92–94 % immune ✓; for influenza ($R_0 \approx 1.3$) ~25 %. One equation explains why measles needs near-universal coverage.

> **Worked example 1 – reading a serology.** (Fictional) day 4 post-rash IgM positive / IgG negative; day 14 IgM positive / IgG up 4-fold or more: acute infection, seroconversion ✓. IgG only, stable titre: past infection or vaccination. The class and the rise ARE the timeline.

> **Worked example 2 – clinical link: the schedule's logic.** (Fictional clinic) hepatitis B vaccine at 0, 1, 6 months: dose 1 primes (only about a third to a half respond), dose 2 boosts most, dose 3 matures affinity and catches slow responders – the 3-dose arc mirrors the primary/secondary curve with the third dose as a settled secondary peak ✓. Live vaccines (MMR) replicate gently: one dose immunises ~95 %, the second catches the remainder.

Common mistakes. Expecting the first vaccine dose to protect (it primes; protection is measurable after dose 2 in multi-dose series). Equating antibody titre with T-cell memory (both exist; titres wane while protection persists via memory B and T cells). Interpreting a single IgG titre without a baseline (only the RISE or known exposure dates make it diagnostic).

Summary. Primary: day 7–14, IgM→IgG, modest. Secondary: hours–2 days, IgG, 10–100×. Memory cells: higher frequency, better location, pre-matured affinity; herd

threshold $1 - 1/R_0$ (measles ~95 %); schedules are the curve, engineered.

Hypersensitivity

Why it matters. The immune system can overreact in four distinct ways, together called **hypersensitivity** – each with its own timescale and cell types. Naming the type predicts the test, the treatment and the timeline: allergy clinics are this taxonomy applied.

The idea, with a picture. **Type I**: IgE + mast cells, IMMEDIATE (minutes) – allergy, anaphylaxis. **Type II**: IgG against a CELL or tissue, minutes–hours – transfusion, autoimmune haemolysis. **Type III**: immune COMPLEX deposition, hours–days – serum sickness, glomerulonephritis. **Type IV**: T-cell mediated, 48–72 hours – contact dermatitis, tuberculosis skin test ✓.

the clock names the mechanism

I — immediate	IgE + mast cells	minutes	anaphylaxis
II — cytotoxic	IgG on cells	hours	haemolysis
III — complexes	immune complex deposit	days	serum sickness
IV — delayed	T cells	48–72 h	contact dermatitis

Antibody on mast cells, antibody on cells, antibody complexes, T cells – four mechanisms, four clocks

The derivation (why the clock tells you the mechanism). Type I is limited by pre-formed IgE waiting on mast cells → seconds. Type II needs antibody to reach the target and recruit effectors → hours. Type III needs complexes to form AND deposit → days. Type IV needs T-cell recruitment and macrophage activation → 48–72 hours ✓. The delay is the biology: each later type has one more transport/assembly step.

> **Worked example 1 – the anaphylaxis arithmetic.** (Fictional) bee sting, previous mild reaction; now urticaria + wheeze + hypotension within 8 minutes: type I, adrenaline intramuscular immediately ✓ (the antihistamine is adjunct; the adrenaline is the decision). The earlier sting sensitised: it made IgE, which now sits on mast cells; this sting cross-links it. Type I needs a previous exposure, which is why a first sting cannot cause anaphylaxis by this route.

> **Worked example 2 – clinical link: the day-3 itch.** (Fictional) nickel watch strap: eczematous, itchy plaque at 48–72 hours exactly where metal touched – type IV, T cells meeting hapten-presenting Langerhans cells in the skin. Patch testing reads at 48 h and 72 h BECAUSE the mechanism needs that long ✓; antihistamines do nothing (no histamine-first step), topical steroids do (they quiet T cells).

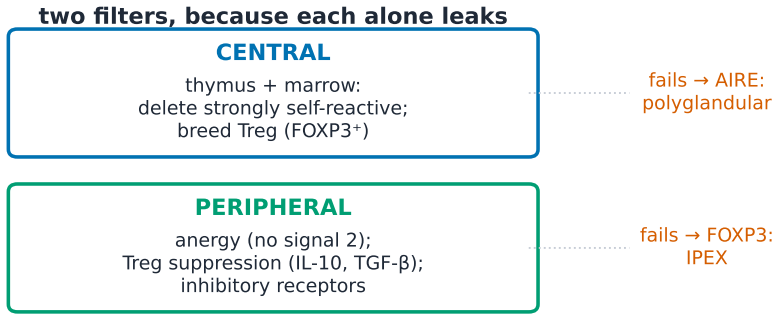
Common mistakes. Treating all allergy as type I (contact dermatitis is type IV – antihistamines useless). Forgetting type II includes drug-induced cytopenias (the drug coats the cell; IgG clears it). Expecting skin-prick tests to diagnose type IV (they test IgE, type I only).

Summary. I = IgE/mast cells/minutes; II = IgG on cells/hours; III = complexes/days; IV = T cells/48–72 h. The clock names the mechanism; the mechanism names the treatment.

Autoimmunity and tolerance

Why it matters. The immune system's hardest problem is self/not-self at the molecular level, and it is solved by checkpoints, not by perfection. When checkpoints fail, the result is autoimmunity – lifelong, organ-specific or systemic – and each disease is a named checkpoint failure.

The idea, with a picture. **Central tolerance**: in the thymus and marrow, strongly self-reactive developing lymphocytes are DELETED (negative selection); the thymus also breeds Tregs (regulatory T cells, FOXP3$^+$). **Peripheral tolerance**: anergy (no co-stimulation), suppression (Treg cytokines IL-10/TGF-β), regulation by inhibitory receptors ✓.



Two great filters – delete in the organs, suppress in the tissues – and the diseases named when each fails

The derivation (why checkpoints, not a perfect library). A receptor library built by random rearrangement cannot pre-exclude self; so the system over-generates then deletes/suppresses ✓ – two filters because each alone leaks. The check-point map predicts disease: AIRE failure → autoimmune polyglandular syndrome (thymic self-

display broken); FOXP3 failure → IPEX (Treg-less: allergy, enteropathy, endocrinopathy); and blocking the checkpoints CTLA-4 or PD-1 in cancer immunotherapy causes autoimmune side effects (colitis, thyroiditis) by design ✓.

> **Worked example 1 – the twin arithmetic.** Twin studies put type-1 diabetes concordance at ~30–50 % in identical twins ✓: genes load the gun (HLA-DR3/DR4), environment (viral trigger, diet) and stochastic T-cell selection pull the trigger. Genetic HLA associations imply recognition thresholds, not determinism.

> **Worked example 2 – clinical link: molecular mimicry's bill.** (Fictional) sore throat → 2 weeks later migratory arthritis and carditis: anti-streptococcal antibodies cross-reacting with heart and joint tissue (**molecular mimicry** ; rheumatic fever) ✓ – a type II attack on self by a response built for the microbe. The ASO titre RISE confirms the trigger; penicillin treats the strep, anti-inflammation treats the crossfire.

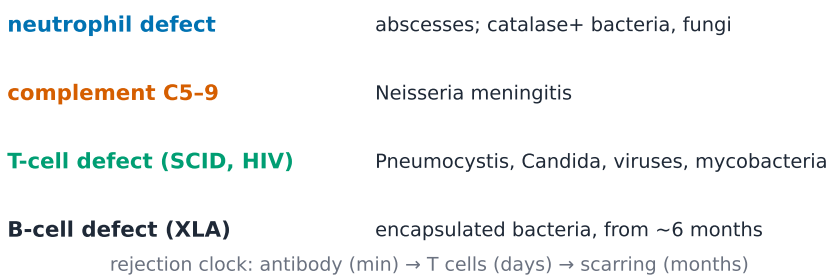
Common mistakes. Reading ANA positivity as lupus (~5–15 % of healthy adults are low-positive; titer and pattern matter ✓). Equating autoimmunity with immunodeficiency's opposite (they overlap: one mechanism's failure often exposes another's). Forgetting organ-specific (thyroid, T1D) vs systemic (SLE) frames.

Summary. Central: delete + breed Treg. Peripheral: anergy, suppression. Failure genes name diseases (AIRE, FOXP3); twin concordance ~30–50 % = genes + environment; mimicry turns infection into self-attack.

Immunodeficiency and transplantation

Why it matters. The immune system's two edges meet here: too little (**immunodeficiency**) → infection by organism and timing; too much (or mismatched) → rejection. Both are read from the same cell map built in this course.

The idea, with a picture. The defect-to-infection map: **neutrophil** defects → catalase-positive bacteria and fungi (abscesses, Aspergillus); **complement (C5–9)** → Neisseria meningitis ✓; **T-cell (SCID, HIV)** → opportunists (Pneumocystis, Candida, viruses, mycobacteria); **B-cell (X-linked agammaglobulinaemia)** → recurrent sinopulmonary infection with encapsulated bacteria after maternal IgG wanes (~6 months) ✓. **Transplant rejection**: hyperacute (pre-formed anti-blood-group or anti-HLA antibody, minutes–hours), acute (T-cell, days–weeks), chronic (antibody + scarring, months–years).



Which arm failed predicts which organism flourishes – the map is the diagnosis

The derivation (why the organism names the defect). Each immune arm specialises: neutrophils clear extracellular bacteria/fungi; complement punches holes in Neisseria's thin outer membrane, and the membrane attack complex is the body's main defence against it ✓; T cells police intracellular and opportunistic life; IgG opsonises encapsulated bacteria. A recurring organism at the wrong age is a pointer to the failed arm – the map is applied cell biology.

> **Worked example 1 – the age-of-onset clock.** (Fictional infant boy) recurrent pneumonia and ear infections from 6 months (maternal IgG gone), very low immunoglobulins, almost no B cells, tiny tonsils: X-linked agammaglobulinaemia ✓. T-cell defects present EARLIER (thrush, diarrhoea, failure to thrive in the first months); neutrophil defects can show in the first weeks (delayed separation of the umbilical cord in leukocyte adhesion deficiency). Age at first presentation is the first clue to the arm.

> **Worked example 2 – clinical link: matching and the drug list.** (Fictional transplant) ABO-compatible, 3/6 HLA-mismatched kidney: acute rejection is blunted by calcineurin inhibitors (tacrolimus blocks the calcium–calcineurin path from the T-cell receptor, signal 1, so IL-2 is not made), belatacept blocks B7–CD28 (signal 2), sirolimus blocks mTOR downstream of the IL-2 receptor, and mycophenolate starves dividing lymphocytes of purines ✓. Each immunosuppressive turns down a named step from 7.3: the drug list is the cell map, priced.

Common mistakes. Forgetting HIV is an acquired T-cell defect (the course's most common cause of opportunistic infection). Timing errors: hyperacute rejection = antibody (immediate, surgical), acute = T cells (days), chronic = fibrosis (months). Expecting vaccination to work in severe combined defects (live vaccines are CONTRAINDICATED ✓).

Summary. The defect names the organism: neutrophils → catalase-positive/fungi; C5–9 → Neisseria; T cells → opportunists; B cells → encapsulated after maternal IgG wanes. Rejection clock: antibody/T cell/scarring. Immunosuppressives = 7.3 signals

turned down; live vaccines contraindicated in severe T defects.

Topic 08

University Medicine

What this course is, and how to use it

This is the eighth course in University Medicine: **medical microbiology and infection** . MED07 built the defences; this course names the attackers – bacteria, viruses, fungi and parasites – and the laboratory methods that tell them apart.

Scope limit, stated once: a study course, not a medical degree; every case below is fictional and nothing here licenses or guides treatment.

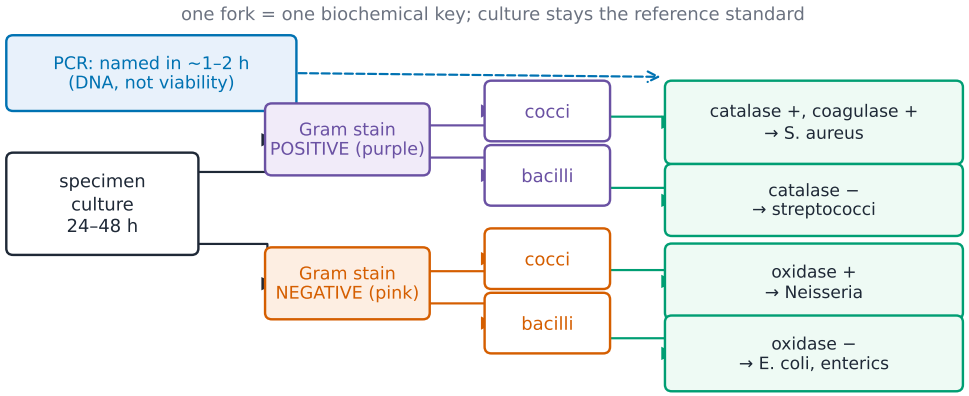
- Units 8.1–8.4 are the organisms: classification, bacteria, viruses, fungi and parasites.
- Units 8.5–8.7 are the interaction: transmission, antimicrobials and prevention.

The medicine habit: **the organism's biology predicts its disease** – every mechanism below is traced from structure or lifecycle to the clinical picture, and every number (doubling times, **incubation periods** , R_0) is something you recompute.

Classification and the diagnostic laboratory

Why it matters. Before any therapy question stands the naming question: what is growing? The diagnostic laboratory is a decision tree turned into instruments, and reading its outputs is a clinical skill.

The idea, with a picture. The tree: specimen → culture on solid media → **Gram stain** (positive purple / negative pink) → shape (cocci, bacilli) → catalase/oxidase, then biochemistry, serology, antigen tests and **PCR** for the fast lane ✓. Culture remains the reference standard: most bacteria grow visibly in 24–48 h; fastidious or slow organisms demand special media and patience.



Gram stain first, then shape, then biochemistry – the tree a microbiology laboratory walks every day

The derivation (why Gram staining separates). The stain is chemistry, not colour preference: crystal violet–iodine complexes are TRAPPED by the thick **peptidoglycan** wall of Gram-positive cells (many layers, **teichoic acids** ✓), while Gram-negative walls (1–2 peptidoglycan layers plus an outer membrane) lose the complex at the decolouriser step and take up the pink counterstain. The same chemistry predicts each group's drug susceptibilities (8.6): what stains, targets.

> **Worked example 1 – reading a report.** (Fictional) blood culture at 48 h: Gram-positive cocci in clusters; catalase positive, coagulase positive. The tree walks: clusters → staphylococci; catalase + → not streptococci; coagulase + → *Staphylococcus aureus*, not the coagulase-negative skin flora ✓. Each step is a fork with one biochemical key – the report's three lines are three forks.

> **Worked example 2 – clinical link: the fast lane.** (Fictional) suspected meningitis: PCR of cerebrospinal fluid can name *S. pneumoniae* or *N. meningitidis* within ~1–2 h, while culture needs 24–48 h ✓. The trade: PCR tells you what DNA is present, not what is alive and drug-sensitive – culture plus susceptibilities still follow. Naming fast, acting informed.

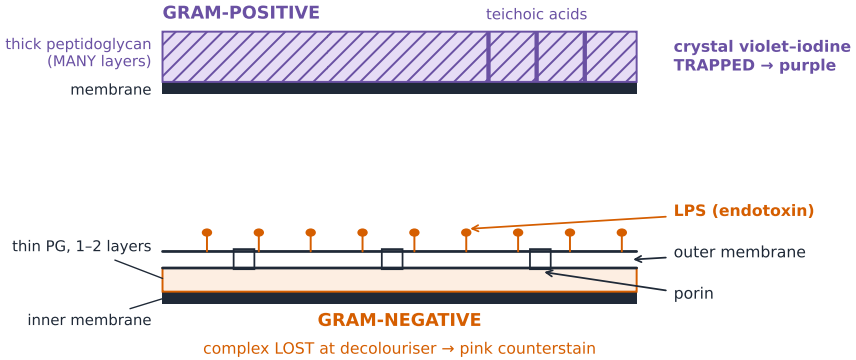
Common mistakes. Reading a Gram stain as a species (it narrows to a shape-and-colour class). Expecting PCR to replace culture (viable organism and susceptibilities come only from culture). Forgetting contaminant biology (coagulase-negative staphylococci in ONE of two bottles are more likely skin flora than pathogen ✓).

Summary. Gram stain (thick-wall purple vs thin-wall pink) → shape → biochemistry → serology/PCR; culture 24–48 h remains the reference; PCR names in ~1–2 h but only culture proves viability; report reading = walking the tree fork by fork.

Bacteria: structure, growth and pathogenesis

Why it matters. A bacterium's structure decides both its disease and its therapy: the wall sets the Gram class, the toxins set the picture, and the growth rate sets the pace of everything.

The idea, with a picture. Gram-positive: thick peptidoglycan with teichoic acids, no outer membrane. Gram-negative: thin peptidoglycan PLUS an outer membrane whose **lipopolysaccharide** (LPS, endotoxin) faces outward, with porins ✓. Common armament: **capsules** (defeat phagocytosis – MED07's opsonisation gap), pili/fimbriae (adhesion), flagella (motility), **exotoxins** (secreted proteins), and the growth curve: lag → log → stationary → death.



The wall chemistry that the Gram stain sees – and that drugs target

The derivation (why doubling time is destiny). $N = N_0 \cdot 2^{t/g}$ with **generation time** *g*. *E. coli* doubles in ~20 min: in 10 h that is 30 generations, $2^{30} \approx 10^9$ cells from one founder ✓. The tuberculosis bacterium doubles in ~20 h: 10 h is half a generation. Fast growers cause acute 24-hour disease and visible cultures at 24–48 h; slow growers cause weeks-long disease and demand weeks of culture (8.1). The arithmetic is the whole clinical pace.

> **Worked example 1 – toxin classes, named.** (Fictional teaching set) Toxin A blocks protein synthesis in a wound infection; toxin B lyses red cells on a plate: the first behaves like diphtheria's fragment-A enzyme (an ADP-ribosyltransferase that halts elongation), the second like the β-haemolysin family. Mechanism explains the picture: synthesis blockers kill with minimal local destruction (distant organ failure), membrane attackers destroy locally ✓.

Worked example 2 – clinical link: the capsule read. (Fictional) sputum Gram stain: Gram-positive lancet-shaped diplococci, each with a clear halo. The halo IS the capsule (*S. pneumoniae*) ✓ – the structure that defeats phagocytosis, and the reason people without a spleen are at risk of overwhelming pneumococcal disease. The stain names the virulence factor.

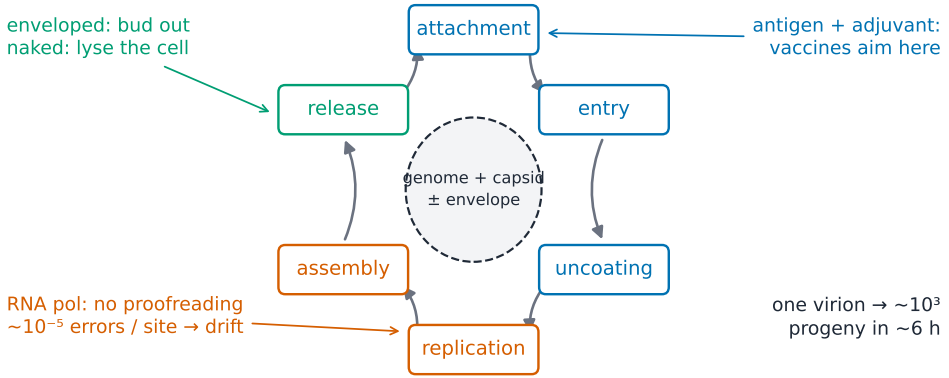
Common mistakes. Calling LPS an exotoxin (it is a wall component released at lysis – endotoxin ✓). Treating all fast growers as pathogens (the gut's commensal *E. coli* doubles at the same speed). Reading absence of growth as absence of organism (fastidious organisms need their media and their days).

Summary. Wall: thick Gram-positive vs LPS-membrane Gram-negative; virulence: capsule, adhesins, exotoxins; growth $N = N_0 2^{t/g}$ – a 20-min doubler makes 10^9 in 10 h, a 20-h doubler makes half a generation; toxin mechanism predicts the clinical picture; the capsule is visible on the stain.

Viruses

Why it matters. Viruses are genomes in coats – obligate intracellular parasites whose lifecycle IS the disease. The cycle tells you why antivirals must target replication steps and why vaccines (7.4) work so well against them.

The idea, with a picture. Structure: nucleic acid (DNA or RNA) + **capsid** (symmetric protein shell) ± **envelope** (host-derived membrane studded with viral glycoproteins) ✓. The cycle: attachment → entry → uncoating → replication of genome and proteins → assembly → release. Enveloped viruses bud out (host membrane kept continuous); naked viruses lyse the cell.



Attachment, entry, uncoating, replication, assembly, release – every antiviral target and every vaccine antigen sits on this cycle

The derivation (why mutation follows genome chemistry). RNA polymerases lack proofreading: influenza's mutation rate is ~10^{-5} substitutions per site per replication, against ~10^{-8}–10^{-6} for DNA viruses, whose polymerases proofread – tens to a thousand times lower ✓. Hence **antigenic drift** year by year and updated vaccines. DNA viruses proofread: stable antigens, stable vaccines. Segmented genomes (influenza: 8 segments) enable **antigenic shift** – reassortment between human and animal strains – the mechanism behind pandemic influenza.

Worked example 1 – the burst arithmetic. (Fictional) one influenza virion infects a respiratory cell; ~6 h later ~10^3 progeny bud off ✓; two rounds (12 h) give 10^6. An incubation period of 1–2 days is only ~4–8 cell rounds: viral arithmetic reaches symptoms faster than any culture.

Worked example 2 – clinical link: reading the season's chart. (Fictional public-health data) winter spike of fever and myalgia, incubation 1–2 days, attack rates ~20 % in offices: the pattern of influenza A (short incubation, drift-matched vaccine) – compare the 2–3-week incubation of a DNA virus like varicella, whose outbreaks build over weeks because each generation of cases takes longer. The epidemic curve's slope encodes the virus's cycle speed ✓.

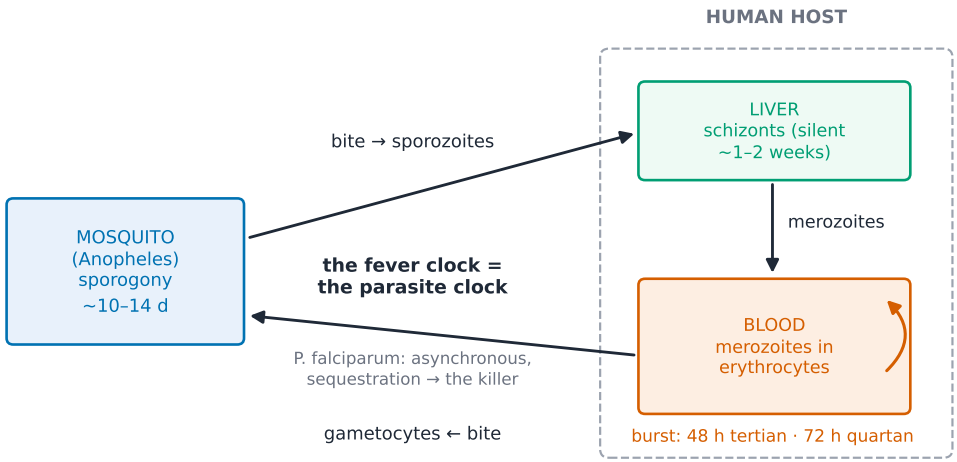
Common mistakes. Treating all RNA viruses as fragile (the envelope decides environmental survival: naked RNA viruses like poliovirus are hardier). Expecting antibacterials to act on viruses (their targets are 8.6's bacteria-specific structures). Forgetting that enveloped viruses carry host membrane – a bud is a disguise as well as an exit.

Summary. Genome (DNA/RNA) + capsid ± envelope; six-step cycle; no proofreading → ~10^{-5} errors per site → drift; 8 segments → shift; one virion → 10^3 in ~6 h; incubation period = the number of rounds.

Fungi and parasites

Why it matters. Eukaryotic pathogens behave differently from bacteria and viruses: dimorphic fungi switch with temperature, protozoa and helminths run complex life cycles through vectors and hosts. These cycles decide geography, diagnosis and whose blood film you are reading.

The idea, with a picture. Fungi: yeasts (single cells, budding), moulds (hyphae), **dimorphic fungi** (mould in the environment at 25 °C, yeast in tissue at 37 °C ✓) – opportunists exploit the immunocompromised host along MED07's defect map (T-cell failure → *Pneumocystis*, *Candida*; neutrophil failure → *Aspergillus*). Parasites: protozoa (malaria, amoebae) and helminths (worms). The **malaria life cycle**: mosquito injects **sporozoites** → liver schizonts → erythrocytic **merozoites** cycling in red cells → **gametocytes** taken up by the next mosquito ✓.



Liver stage, blood stage, gametocyte stage – each stage names a drug class and a vaccine target

The derivation (why the fever is periodic). Synchronous erythrocytic bursting every 48 h (*P. vivax*, *P. ovale* – tertian) or 72 h (*P. malariae* – quartan) ✓ produces the classic periodic fevers: the parasitaemia's clock IS the symptom's clock. *P. falciparum* synchronises poorly – fevers less regular – but cytoadherence of infected red cells (sequestration) makes it the killer: the cycle explains both the textbook pattern and the exception's danger.

> **Worked example 1 – the film read.** (Fictional) blood films from a returned traveller: the thick film finds the parasites; the thin film shows ring forms in ~3 % of erythrocytes plus crescent-shaped gametocytes → *P. falciparum* ✓. The 3 % parasitaemia and the banana gametocyte – morphology names the species and grades severity in one look.

> **Worked example 2 – clinical link: dimorphism as a geography lesson.** (Fictional) a farmer's pneumonia after dust exposure; culture at 25 °C grows a mould, at 37 °C a yeast: a dimorphic fungus (the histoplasma pattern) ✓ – the fungus's thermostat switch doubles as its identity card: the same organism is a mould in the environment and a yeast in you.

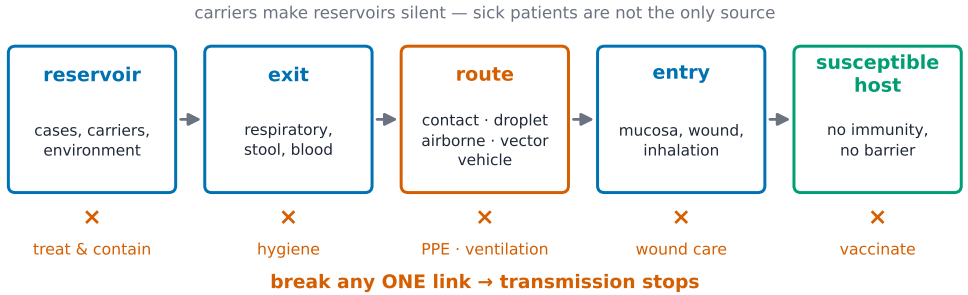
Common mistakes. Treating fungi like bacteria (their eukaryotic machinery narrows selective-toxicity options). Expecting all parasites to show in blood films (most need stool, serology or tissue). Forgetting that opportunistic mycoses follow MED07's arms: which defect, which fungus.

Summary. Yeasts/moulds/dimorphs (25 °C mould, 37 °C yeast); opportunists follow the immunodeficiency map; malaria: sporozoite → liver → 48-h blood cycle → gametocyte; fever periodicity = the parasite's clock; falciparum's asynchrony + sequestration = the killer.

Host–pathogen interaction and transmission

Why it matters. An infection is a negotiation between a population of organisms and a host – the outcome (colonisation, carrier state or disease) and the spread between hosts follow rules you can compute.

The idea, with a picture. The **chain of infection** : reservoir → exit → transmission route (contact / droplet / airborne / vector / vehicle) → entry → susceptible host ✓; break any link and transmission stops. The **basic reproduction number** R_0: average secondary cases from one case in a fully susceptible population – measles ~12–18, influenza ~1.3, ancestral SARS-CoV-2 ~2.5–3 ✓ (MED07's herd threshold $1 - 1/R_0$ builds on this).



Reservoir → exit → route → entry → susceptible host: break any one link, transmission ends

The derivation (why R_0 and the serial interval set epidemic speed). Early in an epidemic, cases multiply by R_0 every serial interval (measles ~10–12 days, influenza ~2–3 days) ✓. Doubling time $t_d \approx \mathrm{SI} \times \log 2 / \log R_0$: influenza ($R_0 \approx 1.3$, SI 2–3 d) doubles in ~5–8 days; measles ($R_0 \approx 15$, SI ~11 d) would double in ~2.5–3 days in a susceptible population. The same equation explains why isolation slows spread (it lowers the effective R) and why short-interval outbreaks outrun contact tracing.

> **Worked example 1 – the attack-rate arithmetic.** (Fictional wedding) 120 guests ate the buffet, 54 developed gastroenteritis within 48 h: **attack rate** = 54/120 = 45 % ✓; stratified by dish, 85 % of seafood eaters vs 10 % of non-eaters fell ill → the vehicle is named. One fraction, stratified once, is an outbreak investigation's first page.

Worked example 2 – clinical link: breaking the chain, link by link. (Fictional ward outbreak of diarrhoeal illness) the reservoir (cases plus contaminated environment) → exit (stool) → route (hands and surfaces) → entry (mouth) → susceptible hosts (ward patients). Hand hygiene breaks the route; cohorting breaks the susceptible-host link; each intervention maps onto a named link of the chain – the diagram is the outbreak plan, drawn ✓.

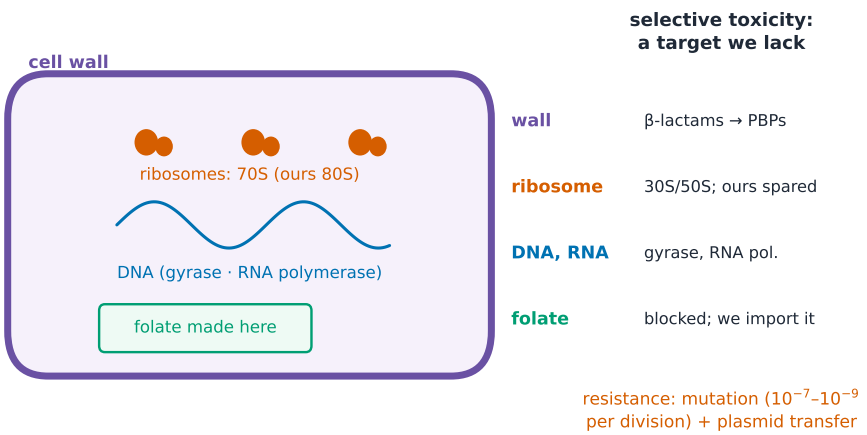
Common mistakes. Equating R_0 with severity or case-fatality (transmissibility and virulence are different axes). Using R_0 (intrinsic) where the effective R (behaviour, immunity) is the operative number. Forgetting carriers: reservoirs need not be sick – chronic typhoid carriers drive transmission ✓.

Summary. Chain of five links; R_0 sets epidemic pace jointly with the serial interval ($t_d \approx \mathrm{SI} \cdot \log 2 / \log R_0$); stratified attack rates name vehicles; carriers make reservoirs silent; interventions = links broken.

Antimicrobials and resistance

Why it matters. Every antibacterial exploits a difference between bacterial and human cells; every resistance mechanism is the bacterium closing that difference. Mechanism-first thinking (no regimens, per the course plan) explains both the drugs and the laboratory's susceptibility report.

The idea, with a picture. Selective toxicity needs bacteria-specific targets: the wall (β-lactams bind penicillin-binding proteins, blocking peptidoglycan cross-linking ✓), the ribosome (bacterial 30S/50S subunits vs human 40S/60S), nucleic- acid enzymes (DNA gyrase, RNA polymerase), and folate synthesis (bacteria make it; humans import it) ✓.



Wall, ribosome, DNA enzymes, folate – four target classes on one cell

The derivation (why resistance is inevitable arithmetic). Mutation supplies resistance at ~10^{-7}–10^{-9} per cell per generation; an overnight broth culture holds ~10^9 cells per mL, so every culture carries mutants against any single drug ✓. Combination approaches exist BECAUSE two simultaneous independent mutations are the product of probabilities (~10^{-14}–10^{-18}). **Horizontal gene transfer** (plasmids, transposons) then spreads the wins between species – chemistry gave the targets; population genetics sets their shelf life.

Worked example 1 – reading a susceptibility report. (Fictional) *S. aureus* from an abscess: β-lactamase positive, susceptible to flucloxacillin. The lab's two lines trace to 8.2's wall: β-lactamase chemically OPENS the β-lactam ring (a hydrolysis), while an altered target (PBP2a) is what makes a true methicillin-resistant strain – different mechanisms, different drugs, both named by the report ✓.

Worked example 2 – clinical link: the stewardship logic. (Fictional ward) rising resistance in unit infections tracked to broad spectrum use: narrow the drug to the reported susceptibility and the selective pressure on the ward's flora falls. The mechanism chain: drug pressure → selection of the resistant fraction (the 10^{-7} arithmetic) → stewardship narrows the pressure. No regimens – the principle IS the mechanism ✓.

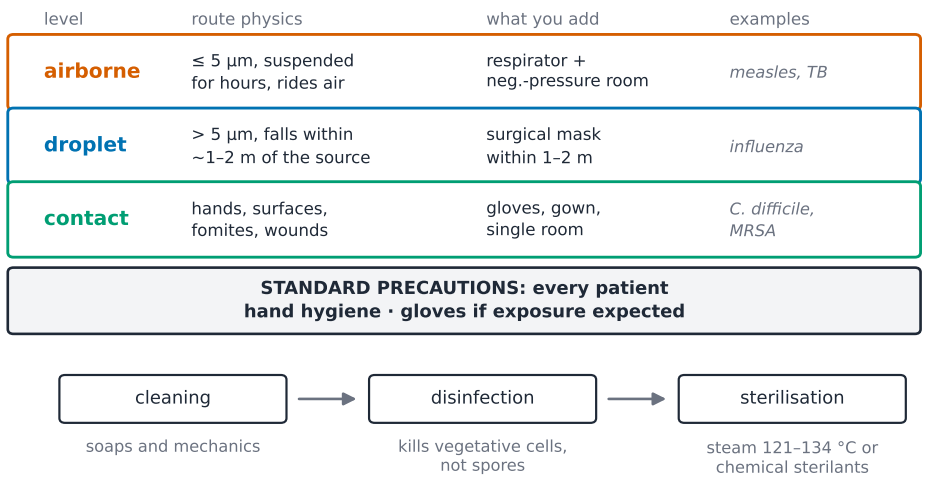
Common mistakes. Treating resistance as all-or-nothing (susceptibility is graded and reported in categories). Expecting new targets to defeat resistance permanently (the 10^{-9} clock restarts for every new molecule). Forgetting that host immunity is the drug's partner (the defences of MED07 do most of the clearing).

Summary. Targets: wall (PBPs), ribosomes, gyrase/polymerase, folate – all bacteria-specific; resistance arises at 10^{-7}–10^{-9} per cell per generation plus horizontal transfer; combination and stewardship exist because of the probability arithmetic; reports name the mechanism, not just the organism.

Infection prevention and control

Why it matters. Prevention is applied transmission knowledge: every precaution level, sterilisation method and screening programme is a link of 8.5's chain, deliberately cut. The cheapest infection is the one that never happens.

The idea, with a picture. Precaution levels: standard precautions (every patient: hand hygiene, gloves for anticipated exposure) plus transmission-based additions – contact (gowns, single room), droplet (masks within ~1–2 m), airborne (fitted respirator, negative-pressure room) ✓. The decontamination ladder: cleaning → disinfection → **sterilisation** (steam 121–134 °C, or chemical for heat-labile instruments) ✓.



Standard for everyone; add the level that matches the route – contact, droplet or airborne

The derivation (why the route sets the level). Droplets (> 5 µm) fall within ~1–2 m: a mask at that distance breaks the chain ✓. Airborne particles (≤ 5 µm) stay suspended for hours and ride air currents: only a fitted respirator plus room ventilation breaks that link. Match the level to the PHYSICS of the route: over-isolation wastes, under-isolation leaves the chain open.

> **Worked example 1 – the vaccination arithmetic.** (Fictional clinic) staff influenza vaccine uptake rises from 40 % to 90 %: if every vaccinated worker were protected, the ward's susceptible-staff fraction would fall from 60 % to 10 % – a six-fold drop (with a vaccine that protects 60 %, it falls from 76 % to 46 %), on top of each vaccinated worker's own reduced risk ✓. Herd thinking (7.4) applied to the staff who meet every reservoir: prevention is a population calculation.

> **Worked example 2 – clinical link: the outbreak clock.** (Fictional) a surgical-site infection cluster presenting on post-operative days 5–7: the timing names the acquisition window (the theatre or the first 48 h of aftercare), the organism's incubation arithmetic (8.2) dates the breach, and the response follows the chain – reservoir (case or environment), route (contact), entry (the wound). One clock, the whole investigation ✓.

Common mistakes. Treating standard precautions as optional when the patient looks well (the reservoir is often silent). Using airborne precautions for droplet organisms (measles and tuberculosis are airborne; influenza is managed with droplet precautions ✓). Confusing disinfection with sterilisation (spores survive disinfection).

Summary. Standard always; add contact/droplet/airborne to match the route's physics (> 5 µm falls in ~1–2 m, ≤ 5 µm drifts for hours); sterilisation for invasive instruments; vaccination is herd arithmetic; the incubation clock dates every breach.

Topic 09

University Medicine

What this course is, and how to use it

This is the ninth course in University Medicine: **neuroscience** . MED03 named the structures and MED05 gave the cell's electrical grammar; this course traces the circuits and turns them around: given a deficit, find the lesion. Prerequisites: MED03 (anatomy), MED05 (physiology).

Scope limit, stated once: a study course, not a medical degree; every case below is fictional and nothing here licenses or guides treatment.

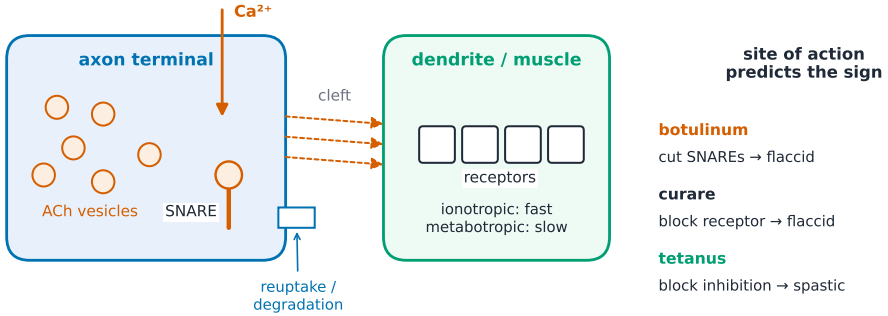
- Units 9.1–9.4 are the wiring: synapses, sensory, motor and special- sense pathways.
- Units 9.5–9.7 are the systems built on it: autonomic outflow, higher functions, and lesion localisation.

The neuroscience habit: **the pathway predicts the lesion** – every tract below is learned in two directions, forward (trace it) and backward (localise along it), and every number is something you recompute.

Neurons, synapses and neurotransmitters

Why it matters. The **synapse** is the nervous system's decision point – and the target of most neuroactive drugs and toxins. Predicting an effect at the synapse is applied cell biology.

The idea, with a picture. Transmission in six steps: action potential reaches the terminal → **voltage-gated Ca^{2+} channels** open → Ca^{2+} influx triggers SNARE-mediated vesicle fusion → transmitter released into the cleft → binds postsynaptic receptors → terminated by **reuptake** or degradation ✓.



Six steps, each a drug or toxin target – botulinum before the cleft, curare after it

The derivation (why transmitter chemistry predicts the effect). The postsynaptic receptor decides the sign: **ionotropic receptors** open channels fast (ms): nicotinic ACh, GABA-A (Cl^- influx → inhibition), glutamate AMPA/NMDA (Na^+/Ca^{2+} → excitation) ✓; **metabotropic receptors** act through G-proteins over hundreds of milliseconds to seconds (muscarinic, adrenergic, most dopamine and serotonin receptors). The receptor, not the transmitter, sets the sign: ACh excites skeletal muscle through nicotinic receptors but slows the heart through M2 receptors. In the adult brain glutamate acts almost only on excitatory receptors and GABA/glycine on inhibitory ones ✓.

> **Worked example 1 – the NMJ arithmetic.** (Fictional lab) one vesicle releases ~5,000–10,000 ACh molecules; one nerve impulse releases ~50–300 vesicles, so $50 \times 5{,}000 = 2.5 \times 10^5$ to $300 \times 10{,}000 = 3 \times 10^6$ molecules ✓. The end-plate potential this makes is about 3–5 times larger than threshold needs – the safety factor – which is why weakness appears only when about 70–80 % of receptors are blocked ✓.

> **Worked example 2 – clinical link: three toxins, three steps.** (Fictional teaching set) botulinum cleaves SNAREs PREsynaptically → no release → flaccid paralysis; curare blocks the nicotinic receptor POSTsynaptically → flaccid paralysis with normal release; tetanus toxin travels up the motor axon to the spinal cord and blocks release from inhibitory interneurons (glycine and GABA) → unopposed excitation → spastic paralysis ✓. Site of action predicts the sign of the paralysis.

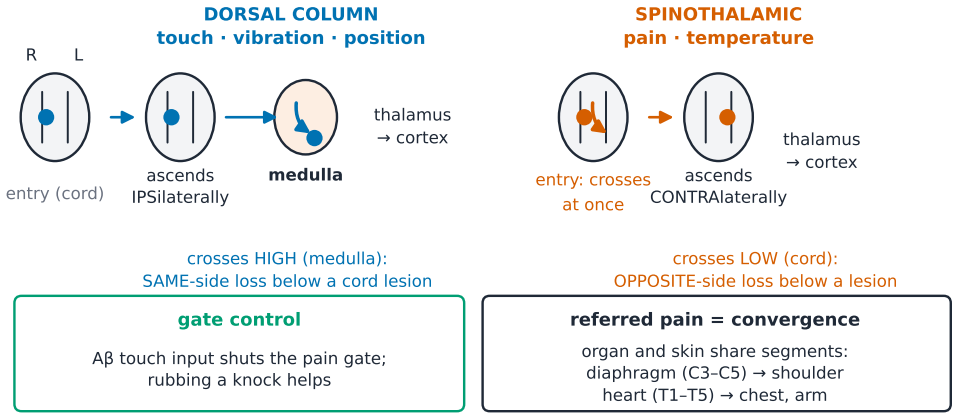
Common mistakes. Calling ACh simply excitatory (the receptor decides: nicotinic excites muscle, M2 slows the heart). Expecting reuptake blockers to change synthesis (SSRIs block the presynaptic serotonin TRANSPORTER, a step after release). Forgetting Ca^{2+} is the trigger, not the fuse: no Ca^{2+} influx, no release, however depolarised the terminal.

Summary. AP → Ca^{2+} → SNARE fusion → cleft → receptor → reuptake or degradation; ionotropic fast (channel), metabotropic slow (G-protein); glutamate excites, GABA/glycine inhibit; NMJ ~5–10k molecules/vesicle; toxin site predicts flaccid (pre- or post-synaptic block) vs spastic (disinhibition).

Somatosensory and pain pathways

Why it matters. Two ascending systems carry the body's sensory news, and they cross at different levels – a sensory deficit's PATTERN names its tract and its side. This is the daily grammar of the neurological examination.

The idea, with a picture. The **dorsal column** pathway (fine touch, vibration, proprioception): enters the cord, ascends IPSilaterally, crosses in the MEDULLA ✓. The **spinothalamic tract** (pain, temperature): crosses within 1–2 segments of entry in the CORD, then ascends contralaterally ✓.



One crosses in the cord, one in the medulla – the level of crossing is the level of the lesion's signature

The derivation (why the crossing levels differ clinically). Because dorsal-column fibres cross high, a CORD lesion numbs fine touch BELOW the lesion on the SAME side; because spinothalamic fibres cross low, the same lesion numbs pain on the OPPOSITE side ✓ – the Brown-Séquard mirror (9.7). **Gate control** adds the modulation: large Aβ fibres activate an inhibitory interneuron that closes the gate on C-fibre pain traffic – rubbing a knock works because Aβ traffic arrives ✓.

> **Worked example 1 – tracing a hemisection.** (Fictional case) left hemicord injury at T10: ipsilateral (left) loss of vibration and proprioception below T10, contralateral (right) loss of pain and temperature starting 1–2 segments below, plus left UMN signs below – each sign traced to one tract and its crossing ✓.

> **Worked example 2 – clinical link: referred pain by convergence.** (Fictional) diaphragmatic irritation felt at the shoulder tip, a **referred pain** : both C3–C5 afferents and diaphragmatic visceral fibres converge on the same spinal segments, and the brain misattributes the visceral source to the somatic map – the same arithmetic that refers cardiac pain to the T1–T5 chest and arm ✓.

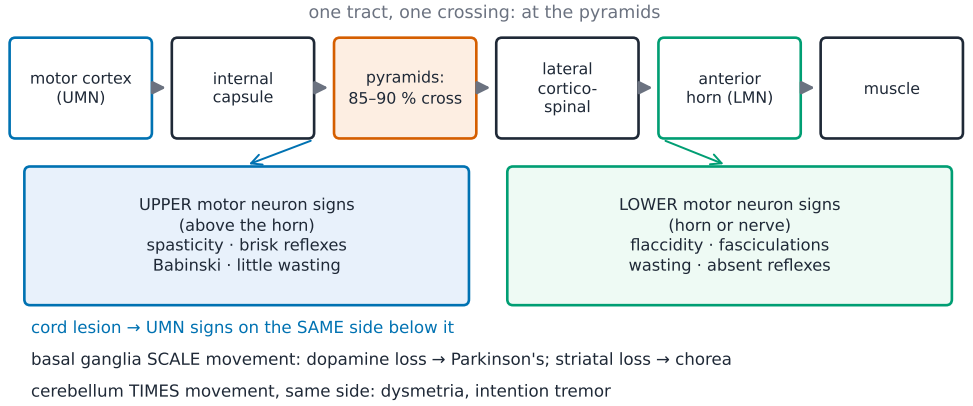
Common mistakes. Tracing pain pathways as uncrossed (they cross in the cord, at entry). Expecting a spinal LEVEL on the face (facial sensation rides the trigeminal system to the brainstem, not the cord). Putting the pain level at the lesion (spinothalamic fibres ascend 1–2 segments before crossing, so the pain level sits 1–2 segments below).

Summary. Dorsal column: ipsilateral to the medulla, then crossed – fine touch/vibration/proprioception. Spinothalamic: crosses in the cord within 1–2 segments – pain/temperature. Cord lesion → ipsilateral fine touch, contralateral pain; Aβ gates C-fibre traffic; convergence refers visceral pain to the somatic map.

Motor systems

Why it matters. Voluntary movement is a two-neuron system with computers attached – and each level fails with its own signature. **Upper motor neuron** (UMN) vs **lower motor neuron** (LMN) signs are the first fork of every motor examination.

The idea, with a picture. The **corticospinal tract** : precentral gyrus (Betz cells, layer V) → posterior limb of internal capsule → cerebral peduncle → pyramid, where ~85–90 % cross (the **pyramidal decussation**) to the lateral corticospinal tract ✓ → anterior horn lower motor neuron (LMN), which exits to the muscle.



One crossing at the pyramids, two neuron classes – upper-motor signs above the anterior horn, lower-motor signs at it

The derivation (why the two neuron classes fail differently). The UMN is a modulatory input; its loss releases the LMN pool from inhibition → spasticity, hyperreflexia, a **Babinski sign** (an upgoing big toe), minimal atrophy ✓. The LMN IS the final common path; its loss removes the drive → flaccidity, fasciculations, atrophy, areflexia ✓. The computers tune the loop: basal ganglia scale movement amplitude (dopamine balance; Parkinson = too little → bradykinesia, rigidity, rest tremor; Huntington = striatal loss → chorea), cerebellum times it (ipsilateral incoordination: dysmetria, **intention tremor** , dysdiadochokinesia) ✓.

> **Worked example 1 – one lesion, two sign sets.** (Fictional) left-sided cervical cord compression at C5: at the level, LMN signs in C5 muscles (biceps wasting, absent biceps jerk); below it, UMN signs. On which side? The lateral corticospinal tract has ALREADY crossed at the medulla, so the UMN signs appear on the SAME side (left) below C5 ✓: the crossing level is written into the anatomy.

Worked example 2 – clinical link: the tremor fork. (Fictional) a rest tremor that improves with movement plus bradykinesia points to the basal ganglia (Parkinson pattern); a tremor that appears ONLY with intention plus dysmetria points to the cerebellum ✓. Same word – tremor – two circuits, two examinations, two prognoses.

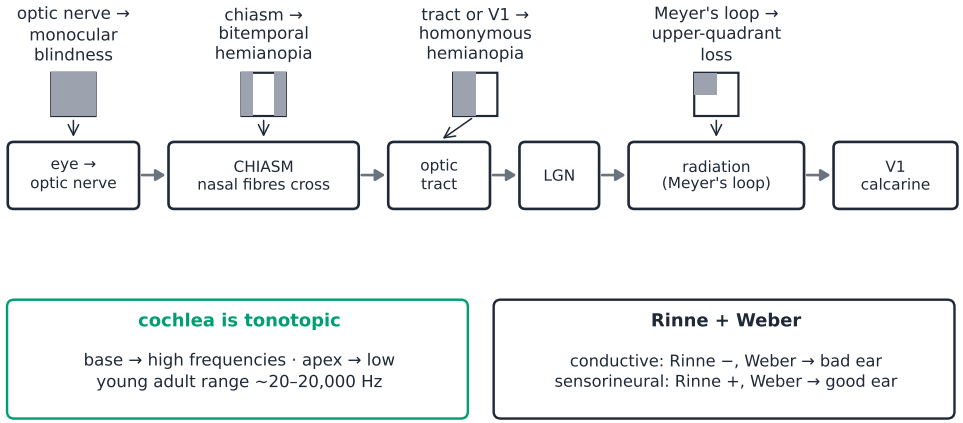
Common mistakes. Expecting corticospinal signs contralateral below a CORD lesion (it crossed at the medulla – contralateral is above the foramen magnum only). Equating spasticity with strength (spasticity is tone, not power). Forgetting the cerebellum is ipsilateral (its tracts double-cross).

Summary. Precentral → internal capsule → pyramids (~85–90 % decussate) → anterior horn → muscle; UMN = spastic, hyperreflexic, Babinski; LMN = flaccid, wasting, areflexic; cord lesions give ipsilateral UMN below (already crossed); basal ganglia scale amplitude, cerebellum times movement (its tracts double-cross to the same side).

Special senses

Why it matters. Vision and hearing are mapped pathways: because the map is preserved end to end, a field defect or an audiogram names the lesion's site with surveyor's precision.

The idea, with a picture. The visual pathway: photoreceptors → retinal ganglion cells → **optic nerve** → **optic chiasm** (nasal fibres cross) → optic tract → lateral geniculate nucleus → optic radiation → primary visual cortex (V1, calcarine sulcus) ✓.



Each lesion site stamps a signature field defect – nerve, chiasm, tract or radiation

The derivation (why lesions stamp signatures). Nasal retinal fibres cross at the chiasm and each carries the TEMPORAL visual field: an optic nerve lesion → monocular blindness; a chiasmal lesion → bitemporal **hemianopia** (loss of half a visual field); a tract/radiation/V1 lesion → contralateral homonymous hemianopia ✓; a temporal-lobe radiation lesion (Meyer's loop, inferior fibres) → contralateral superior quadrantanopia ✓. The map is the lesion detector.

Worked example 1 – reading the fields. (Fictional) driver misses the left half of both visual fields after a stroke: a LEFT homonymous hemianopia (named after the lost half-field) → lesion behind the chiasm on the RIGHT, contralateral to the lost fields ✓. Each word of the field name encodes a side and a level.

Worked example 2 – clinical link: the cochlear map. (Fictional audiogram) air conduction lost, bone conduction preserved (**Rinne test** negative, Weber lateralises to the affected ear) → conductive loss (middle ear); air still better than bone (Rinne positive) with both reduced, and Weber lateralising to the BETTER ear → sensorineural (cochlea/nerve) ✓. The cochlea is laid out by pitch, its **tonotopy** – base transduces high frequencies, apex low (young adult range ~20–20,000 Hz) – so the audiogram's SHAPE localises within the cochlea too. Vestibular twins: three canals in orthogonal planes report rotation; otolith organs report linear acceleration and tilt; nystagmus is the reflex made visible.

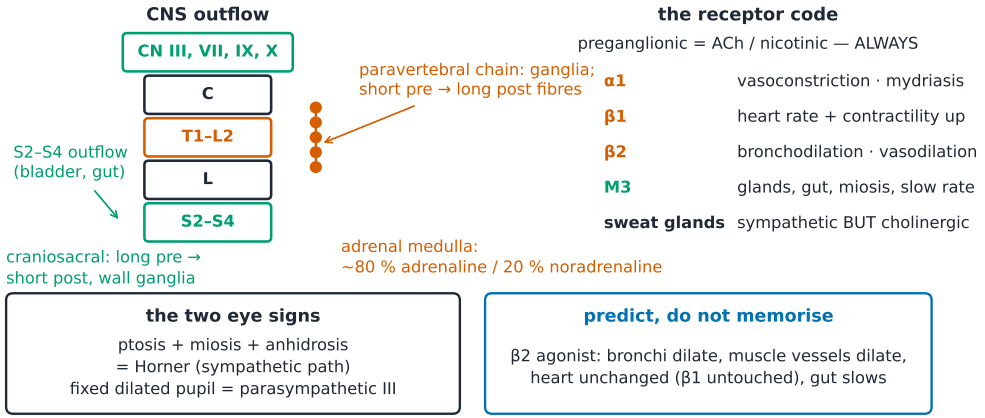
Common mistakes. Calling a homonymous hemianopia monocular (it is binocular – same half of EACH field). Reading Rinne without Weber (the pair, not either alone, separates conductive from sensorineural). Forgetting the temporal field projects to the NASAL retina (the crossing that makes chiasmal lesions bitemporal).

Summary. Nerve → chiasm (nasal crosses) → tract → LGN → radiation → V1; nerve = monocular, chiasm = bitemporal, behind = contralateral homonymous, Meyer's loop = superior quadrantanopia; Rinne+Weber split conductive from sensorineural; cochlea tonotopic (base high, apex low, ~20–20,000 Hz).

Autonomic nervous system

Why it matters. The autonomic system runs the viscera through two opposed limbs, and its receptor map is how organ responses – and autonomic drug effects – are predicted rather than memorised.

The idea, with a picture. The **sympathetic nervous system**: thoracolumbar outflow T1–L2, short pre-/long postganglionic fibres, ganglia in the paravertebral chain; the adrenal medulla is a modified ganglion releasing ~80 % adrenaline / 20 % noradrenaline ✓. The **parasympathetic nervous system**: craniosacral outflow (CN III, VII, IX, X and S2–S4), long pre-/short postganglionic fibres, ganglia in or near the wall ✓.



T1–L2 versus craniosacral – outflow anatomy is the receptor map drawn on the body

The derivation (why receptors predict organ responses). The autonomic code: all preganglionic fibres are ACh/nicotinic; sympathetic postganglionic is noradrenaline (except sweat glands, ACh), parasympathetic postganglionic is ACh/muscarinic ✓. Organ response = receptor expressed: $\alpha 1 \rightarrow$ vasoconstriction, mydriasis; $\beta 1 \rightarrow$ heart rate and contractility up; $\beta 2 \rightarrow$ bronchodilation, vasodilation; M2 $\rightarrow$ heart rate down; M3 $\rightarrow$ glands, gut motility, miosis, bronchoconstriction ✓. Predict the eye, the airways, the gut – one receptor table, every organ.

> **Worked example 1 – one agonist, four organs.** (Fictional teaching set) a selective $\beta 2$ agonist: bronchi dilate (asthma logic), skeletal-muscle vessels dilate, skeletal muscle shows a fine tremor, and the heart speeds a little (a reflex to the vasodilation plus some cardiac $\beta 2$ receptors) ✓ – each line read off the receptor table, no memorisation.

> **Worked example 2 – clinical link: the lesion that droops.** (Fictional) **Horner syndrome** – ptosis, miosis, anhidrosis – on one side: a sympathetic PATHWAY lesion (anywhere eye-to-spine) with intact parasympathetics ✓. The opposite pupil sign – a fixed dilated pupil with a drooping lid and an eye turned down and out – would implicate the parasympathetic fibres in the III nerve: two drooping lids, two systems.

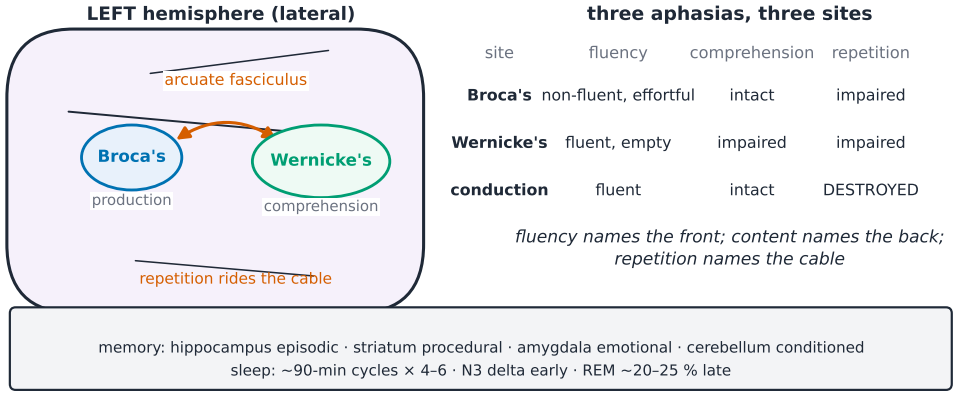
Common mistakes. Sympathetic = accelerates everything (it CONSTRICTS gut vessels and skin vessels; the response set is receptor-written). Forgetting sweat glands are sympathetic but cholinergic (the classic exception). Assuming the adrenal medulla is innervated postganglionically (it is a ganglion modified to secrete into blood).

Summary. Sympathetic T1–L2 (chain, short pre; adrenal medulla ~80/20 adrenaline/noradrenaline); parasympathetic craniosacral (III, VII, IX, X, S2–S4; wall ganglia). Preganglionic = ACh/nicotinic always; $\alpha 1$ constricts, $\beta 1$ drives heart, $\beta 2$ dilates, M2 slows the heart, M3 secretes; Horner = sympathetic lesion; fixed dilated pupil = III.

Higher functions: sleep, memory and language

Why it matters. Sleep, memory and language are network functions – their architecture predicts their disorders, and their localisation is clinical linguistics and electrophysiology applied.

The idea, with a picture. Sleep cycles ~90 minutes, four to six per night: N1 (transition), N2 (sleep spindles, K complexes), N3 **slow-wave sleep** (delta 0.5–4 Hz, growth hormone, hardest to wake), REM (atonia, vivid dreams, ~20–25 % of the night) ✓. Memory systems: hippocampus $\rightarrow$ episodic; striatum $\rightarrow$ procedural/habit; amygdala $\rightarrow$ emotional tagging; cerebellum $\rightarrow$ conditioned motor learning ✓. Language: **Broca's area** (inferior frontal, usually left) – production; **Wernicke's area** (superior temporal) – comprehension; the **arcuate fasciculus** connects them for repetition ✓.



Broca produces, Wernicke comprehends, the arcuate fasciculus repeats – three aphasias, three sites

The derivation (why lesion sites name the aphasias). Broca lesion $\rightarrow$ non-fluent, effortful speech with intact comprehension and intact understanding of own errors; Wernicke lesion $\rightarrow$ fluent, empty speech with impaired comprehension; arcuate lesion $\rightarrow$ **conduction aphasia**: fluent speech, good comprehension, REPETITION destroyed (the cable between them cut) ✓. Fluency names the front; content names the back; repetition names the cable.

> **Worked example 1 – the night's arithmetic.** (Fictional sleep study) 8 hours, five cycles: N3 dominates the first two (slow-wave pressure dissipates), REM episodes lengthen toward morning – early-night sleep loss costs growth hormone, late-night loss costs REM ✓. The architecture, not just the total, is the physiology.

Worked example 2 – clinical link: HM's lesson. (Classic teaching case, no identification needed) bilateral removal of the medial temporal lobes, including most of both hippocampi → profound ANTEROGRADE amnesia with intact procedural learning, while most remote memories survived: the hippocampus writes new episodic memories, the striatum still learns skills ✓ – dissociation is the experiment that separates memory systems.

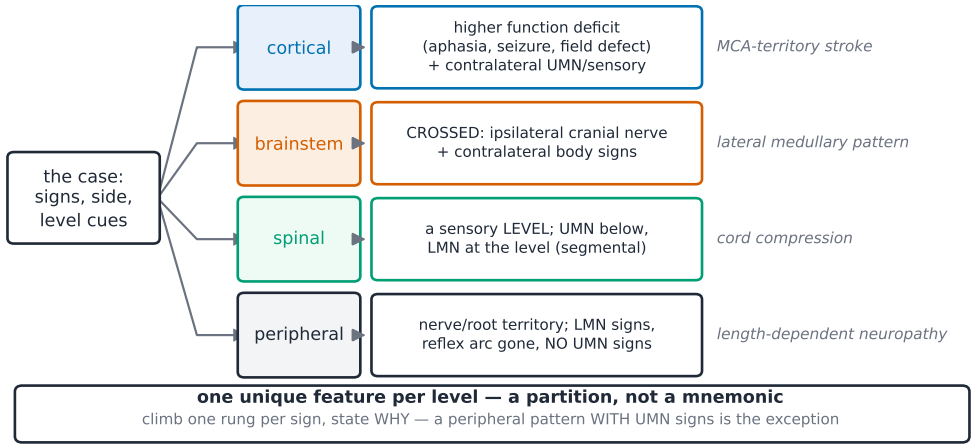
Common mistakes. Thinking dreams prove REM (they occur in NREM too; vivid recall is REM). Expecting hippocampal lesions to erase OLD memories (consolidated memories live in cortex; the hippocampus writes new ones). Calling all language deficit "aphasia of the left" without checking fluency and repetition first.

Summary. 90-min cycles × 4–6: N3 early (delta, GH), REM late (~20–25 %, atonia); hippocampus episodic, striatum procedural, amygdala emotional, cerebellum conditioned; Broca non-fluent, Wernicke fluent- empty, arcuate = repetition; HM: no new episodic memory, skills intact.

Lesion localisation

Why it matters. Everything above converges here: given a fictional case's signs, name the lesion's level and side – the skill this course was built to teach, one justified step at a time.

The idea, with a picture. The decision map: **cortical** – higher function deficit (aphasia, seizures, field defects) + contralateral UMN/sensory signs; **brainstem** – CROSSED findings (ipsilateral cranial nerve + contralateral body); **spinal** – a sensory LEVEL with UMN signs below and LMN signs at the level; **peripheral** – LMN/sensory loss in a nerve, root or dermatome distribution, reflex arc gone, no UMN signs ✓.



Cortical, brainstem, spinal or peripheral – the pattern forks, each fork names a level

The derivation (why the map is complete). Each level owns a feature no lower level has: cortex owns higher functions; brainstem owns cranial nerves AND the crossing of tracts (hence ipsilateral face with contralateral body); cord owns the LEVEL (segmental dermatomes) with UMN below; periphery owns individual nerves and roots with their specific territories ✓. One unique feature per level – the map is a partition, not a mnemonic.

Worked example 1 – localise in four steps. (Fictional) right facial weakness (lower face) + right hemiparesis + aphasia: aphasia → left cortex; lower-face-only facial weakness with sparing of the brow → UMN VII → lesion above the VII nucleus; face and body weak on the SAME side (not crossed) → above the brainstem; right-sided weakness → left hemisphere → left middle cerebral artery territory ✓. Each sign climbed one rung; the localisation is the climb, restated.

Worked example 2 – clinical link: the level that moves. (Fictional) burning pain in both feet, absent ankle jerks, normal plantar responses, loss of pain to the knees bilaterally: peripheral pattern (LMN + pain loss without UMN) → a length-dependent polyneuropathy; contrast a cord-compression picture (a LEVEL on the trunk, UMN below) ✓. Symmetry, length dependence and the reflex arc separate periphery from cord.

Common mistakes. Localising from a single sign (the map works on PATTERNS; one sign has many addresses). Forgetting brainstem lesions cross (the single most useful exam clue). Skipping the general examination (wasting and fasciculations together with brisk reflexes and no sensory loss is motor neuron disease – UMN and LMN signs at once, the exception the map must allow for).

Summary. Cortex: higher function + contralateral signs. Brainstem: crossed findings. Cord: sensory level, UMN below, LMN at level. Periphery: nerve/root territories, LMN, no UMN. Four levels, four signatures; localise step by justified step, and say WHY at each rung.

Topic 10

University Medicine

What this course is, and how to use it

This is the tenth course in University Medicine: **pathology** – the study of disease at the level of cells, tissues and organs. MED04 gave the normal histology and MED05 the physiology; this course explains what goes wrong, from reversible injury to the structured report that carries the answer.

Scope limit, stated once: a study course, not a medical degree; every case below is fictional, diagrams are schematic (no surgical photographs) and nothing here licenses or guides treatment.

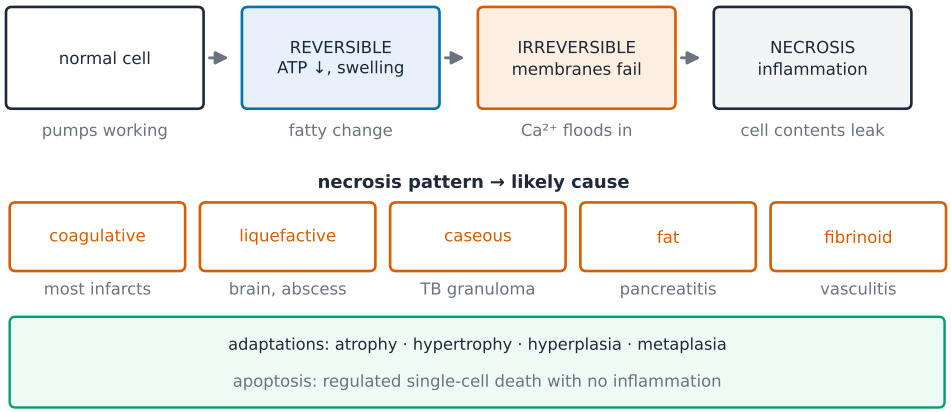
- Units 10.1–10.4 are general pathology: injury, inflammation, haemodynamics, **neoplasia**.
- Units 10.5–10.7 are the systems and the laboratory: cardiorespiratory and gut/liver/kidney disease, then the report itself.

The pathology habit: **mechanism predicts finding** – every lesion below is traced from its mechanism to the finding it produces, and every number is something you recompute.

Cell injury, adaptation and death

Why it matters. All disease begins as cell stress: whether a cell recovers or dies decides everything downstream – and the PATTERN of death names the cause.

The idea, with a picture. Reversible injury: hypoxia/ischaemia → ATP falls → Na^+/K^+ pump fails → cell and organelle swelling, fatty change ✓. Irreversible: membrane damage (phospholipase activation), Ca^{2+} influx, mitochondrial damage – the point of no return ✓.



ATP depletion swells the cell; membrane damage kills it – and the death pattern names the cause

The derivation (why necrosis patterns name causes). Dead cells keep their outlines when PROTEOLYSIS is slow – coagulative necrosis in solid organs after ischaemia ✓; the brain's lysosome-rich tissue liquefies – liquefactive; TB debris with granulomas gives cheese-like caseous; lipase on fat cells gives fatty necrosis (pancreas); immune deposits damage vessel walls – fibrinoid ✓. **Apoptosis** is regulated deletion: no inflammation, single cells, caspases.

> **Worked example 1 – the adaptation quartet.** (Fictional teaching set) disuse → atrophy (smaller cells); hypertension → hypertrophy (bigger cells, no new fibres in cardiac muscle); hormones/liver → hyperplasia (more cells); smoking → bronchial **metaplasia** (squamous replaces columnar – one differentiation traded for another) ✓.

> **Worked example 2 – clinical link: the reversible window.** (Fictional) ischaemic kidney tubule cells at 20 minutes: swollen, ATP low, Na^+ in – restore perfusion and they recover. At hours: membranes leak cytosolic enzymes such as LDH into the blood, where a laboratory detects them, and Ca^{2+} has entered – beyond the point of no return ✓. The enzyme curve is the cell's biography, readable from the blood.

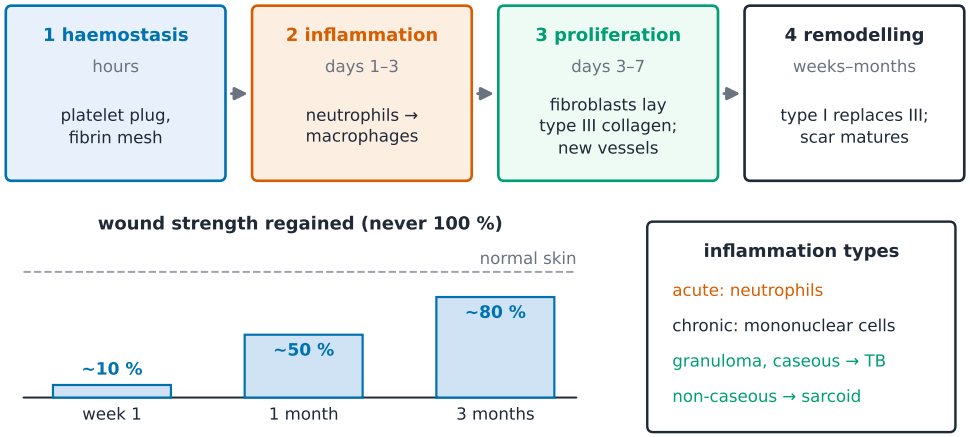
Common mistakes. Calling apoptosis a kind of necrosis (opposite mechanisms: regulated deletion vs enzymatic dissolution with inflammation). Reading **caseous necrosis** as specific to TB (it is granuloma-associated, classically TB). Forgetting fatty change is reversible (an early sign of injury, not of death).

Summary. Low ATP → swelling (reversible); membrane failure + Ca^{2+} → death (irreversible); necrosis patterns: coagulative (ischaemia, solid organs), liquefactive (brain, abscess), caseous (TB), fat (pancreas), fibrinoid (vessels); adaptations: atrophy, hypertrophy, hyperplasia, metaplasia.

Inflammation and repair

Why it matters. Inflammation is the body's response to injury – mediated by the same cells MED07 introduced. Its duration and cell type decide whether you heal, scar, or wall off a pathogen.

The idea, with a picture. Acute: minutes–hours, **neutrophils** , exudate, the cardinal signs (rubor, calor, tumor, dolor, functio laesa) ✓. Chronic: days–onward, mononuclear cells (macrophages, lymphocytes), scarring – and **granulomas** : aggregated macrophages, caseous (TB) or non-caseous (sarcoid) ✓.



Haemostasis → inflammation → proliferation → remodelling: each phase has its cells and its clock

The derivation (why wounds gain strength on a schedule). Repair runs a fixed programme: haemostasis (platelet plug, hours) → inflammation (neutrophils then macrophages, days 1–3) → proliferation: **granulation tissue** (fibroblasts, type III collagen, new vessels, days 3–7) → remodelling (type I replaces III, months – wound strength ~10 % at 1 week, rising fast over the next month and levelling at ~70–80 % from about 3 months, never 100 %) ✓. This schedule is why dehiscence risk peaks early and why sutures hold what collagen cannot yet.

> **Worked example 1 – dating a wound.** (Fictional post-mortem teaching case) neutrophils dominate day 1–3; granulation tissue and type III collagen appear by day 3–7; by week 2 macrophages recede and collagen bridges – each histological field carries a date ✓.

> **Worked example 2 – clinical link: the granuloma fork.** (Fictional biopsy) lung nodule with caseating granulomas → think mycobacteria (TB); non-caseating with no organism → sarcoid pattern ✓. The presence of necrosis inside the granuloma is the fork – one finding, two investigative worlds.

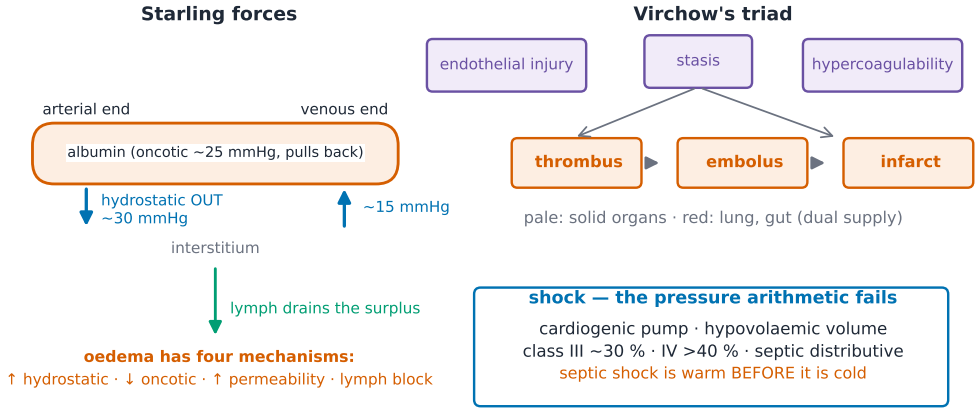
Common mistakes. Defining chronic inflammation by time alone (it is defined by its cells – macrophages, lymphocytes, plasma cells – with tissue destruction and repair going on together). Calling granulation tissue "granuloma" (granulation = repair vessels + collagen; granuloma = aggregated macrophages). Forgetting first-intention (clean approximated wound) heals faster than second-intention (open, granulating).

Summary. Acute: neutrophils, exudate, cardinal signs; chronic: mononuclear + scarring; granulomas caseous (TB) vs non-caseous (sarcoid); healing: clot (h) → inflammation (d 1–3) → granulation tissue (d 3–7) → remodelling to type I (~70–80 % strength from 3 months).

Haemodynamic disorders

Why it matters. Fluids and clots: **oedema** , thrombosis, **embolism** , **infarction** and shock are the final common paths of half of internal medicine – and all are arithmetic on pressures.

The idea, with a picture. Starling forces: capillary hydrostatic pressure pushes out (~30 mmHg arterial end, ~15 venous); plasma **oncotic pressure** (albumin, ~25 mmHg) pulls back; the small net surplus returns as lymph ✓. **Virchow's triad** : endothelial injury, stasis, hypercoagulability → thrombosis → embolism → infarction ✓.



Push out, pull back, drain the surplus – and the three forces that turn blood to clot

The derivation (why oedema has four main mechanisms). Net filtration = (capillary − tissue hydrostatic) − (plasma − tissue oncotic), and oedema collects when filtration outruns lymph drainage. So it arises when hydrostatic pressure rises (heart failure, with the kidney's sodium and water retention adding volume), oncotic pressure falls (liver, kidney, malnutrition), permeability rises (inflammation, exudate) or lymph is blocked (after surgery, filarial) ✓ – almost every clinical oedema is one of these four, and the mechanism names the treatment logic. Shock follows the same

arithmetic: cardiogenic (pump), hypovolaemic (volume; class III 30–40 % blood loss, class IV >40 %), distributive (septic: vasodilated capacity).

> **Worked example 1 – reading the pressure pair.** (Fictional) right heart failure: venous hydrostatic rises → dependent pitting oedema; **nephrotic syndrome** : albumin lost (≥3.5 g/day) → oncotic falls → periorbital oedema ✓. Same finding, opposite mechanism – the distribution and the blood albumin separate them.

> **Worked example 2 – clinical link: triad to infarct.** (Fictional) long flight (stasis) + post-op state (hypercoagulability) → calf thrombus → embolises → pulmonary artery occlusion → infarct or sudden death ✓. Each step of the chain has its own test: D-dimer, ultrasound, CT.

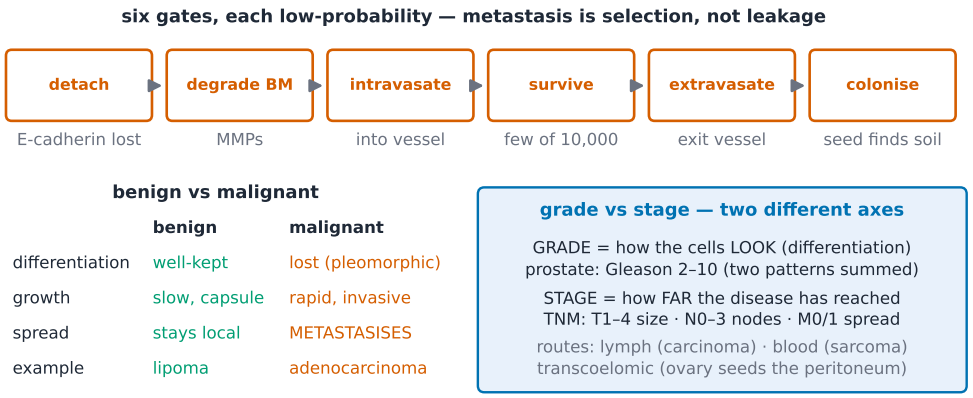
Common mistakes. Calling all oedema the same disease (four mechanisms, four arithmetic causes). Forgetting infarcts are usually pale (solid organs) but red in lung, gut and testis (loose tissue, dual supply) – and septic infarcts become abscesses. Equating shock with low pressure alone (septic shock is warm and vasodilated before it is cold).

Summary. Starling: out (~30/15) − back (~25) = lymph; oedema = ↑ hydrostatic, ↓ oncotic, ↑ permeability, lymph block; Virchow: injury + stasis + hypercoagulability → thrombus → embolus → infarct; shock: cardiogenic, hypovolaemic (class III 30–40 %), distributive.

Neoplasia

Why it matters. Neoplasia is regulated growth lost – and the vocabulary (benign, malignant, grade, stage) decides every prognosis conversation. The steps of spread explain every metastasis.

The idea, with a picture. Benign: well-differentiated, slow, encapsulated, local ✓. Malignant: pleomorphism, rapid, invasive, **metastasis** – the defining property ✓. Invasion in steps: growth → detach (E-cadherin lost) → degrade basement membrane (MMPs) → intravasate → survive the bloodstream (fewer than 1 in 10,000 cells do) → arrest → extravasate → colonise a hospitable soil ✓.



Detach, degrade, enter, survive, exit, colonise – six gates, and the last one is the soil

The derivation (why grade and stage are different axes). Grade = how abnormal the cells LOOK (well/moderate/poorly differentiated; prostate's Gleason score sums the two commonest patterns, 6–10 in modern practice, reported as Grade Groups 1–5) – a tissue property ✓. Stage = how FAR the disease has reached (TNM: T1–T4 size/invasion, N0–N3 nodes, M0/M1 metastasis) – an anatomy property ✓. Prognosis follows stage first, grade second; treatment follows both.

> **Worked example 1 – the arithmetic of spread.** (Fictional teaching set) a 1 cm tumour (about 10^9 cells) can shed ~10^6 cells a day into the blood, yet fewer than 1 in 10,000 survive to colonise – metastasis is a cascade of low-probability gates, not leakage ✓. That is why "positive circulation cells" alone do not equal metastases.

> **Worked example 2 – clinical link: routes predict addresses.** (Fictional) carcinomas travel lymphatic-first (breast → axillary nodes); sarcomas go haematogenous, and so do renal cell carcinoma (into the renal vein) and hepatocellular carcinoma (into the hepatic and portal veins), with the lung a common target; prostate cancer reaches the vertebrae through the valveless vertebral veins; ovarian carcinoma seeds the peritoneum (transcoelomic) ✓. The route explains the destination – and the first scan is chosen accordingly.

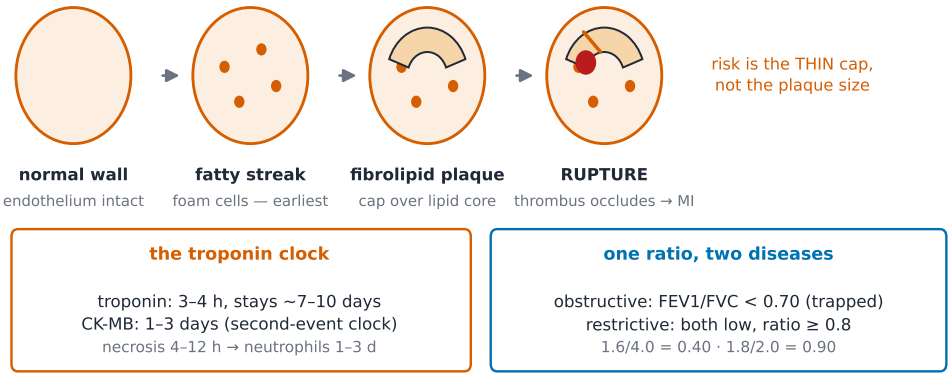
Common mistakes. Using "metastasis" loosely (it defines malignancy – benign tumours do not metastasise). Mixing grade with stage (look vs reach). Assuming larger always means worse (stage T is size/invasion, but N and M outrank it).

Summary. Benign: differentiated, slow, local; malignant: atypia, invasion, metastasis; cascade: detach (E-cadherin) → MMPs → intravasation → survival (under 1 in 10,000) → colonisation (seed and soil); grade = appearance (Gleason 6–10, Grade Groups 1–5), stage = extent (TNM); routes: lymphatic (carcinoma), blood (sarcoma, kidney, liver), transcoelomic (ovary).

Cardiovascular and respiratory pathology

Why it matters. Atherosclerosis and its rupture are the world's leading cause of death; airflow limitation is its respiratory mirror. Both are read off timelines and ratios.

The idea, with a picture. The plaque sequence: endothelial injury (LDL entry) → **fatty streak** (foam cells, the earliest lesion) → fibrolipid plaque (fibrous cap over lipid core) → **complicated plaque**: rupture → thrombosis → occlusion → infarction ✓.



Foam cells to fibrous cap to rupture – and the troponin clock that dates the event

The derivation (why troponin is the clock). Myocardial cells dying release **troponin**: blood rises 3–4 h after injury, stays elevated ~7–10 days ✓; CK-MB rises and falls faster (~1–3 days), flagging a SECOND event troponin cannot separate. The infarct itself runs a histological timeline: coagulative necrosis from 4–12 h, neutrophils 1–3 days, macrophages 3–7 days, granulation tissue from about 7–10 days, a firm scar by about 2 months – coagulative necrosis again, from 10.1's rulebook.

> **Worked example 1 – reading the spirometry pair.** (Fictional lab) FEV1 1.6 L, FVC 4.0 L: ratio 0.40 < 0.70 → obstructive (asthma/COPD pattern) ✓. FEV1 1.8, FVC 2.0: ratio 0.90 with both volumes low → restrictive (fibrosis) ✓. Obstruction traps air (FEV1 falls more); restriction shrinks both (ratio preserved or raised).

> **Worked example 2 – clinical link: the ratio as mechanism.** (Fictional) emphysema destroys alveolar walls: elastic recoil falls, airways collapse on expiration – FEV1 collapses with them, FVC less so, ratio falls ✓. Fibrosis stiffens the lung: both volumes shrink but flow per volume is preserved, ratio ≥0.8. One ratio, two disease families, two mechanism chains.

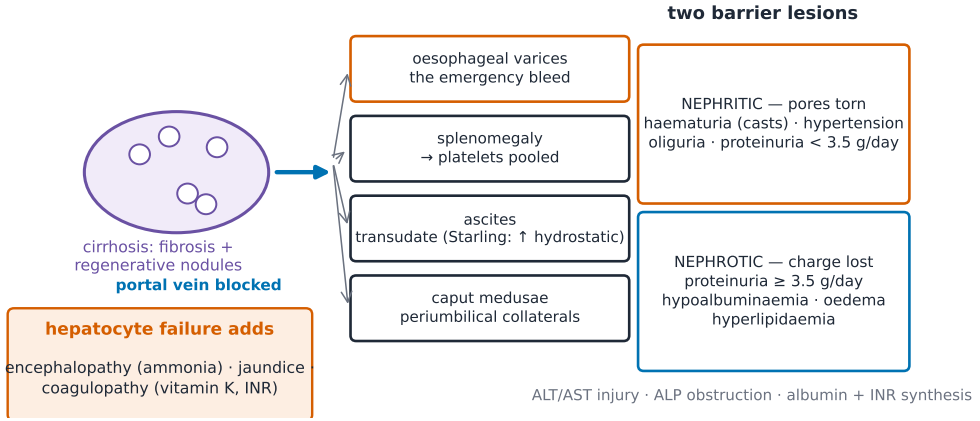
Common mistakes. Reading plaque size as risk (vulnerable thin-cap plaques rupture; heavily calcified stable ones often do not). Using troponin as an old-event marker beyond its window (use it to diagnose the acute; troponin persists ~7–10 days). Equating COPD with chronic bronchitis (one umbrella, several phenotypes).

Summary. Injury → LDL → foam-cell fatty streak → fibrolipid plaque → rupture → thrombosis → infarct; troponin 3–4 h, ~7–10 d; CK-MB 1–3 d (second-event clock); FEV1/FVC <0.70 = obstruction, both low with preserved ratio = restriction.

Gastrointestinal, liver and renal pathology

Why it matters. Cirrhosis, nephritis and nephrosis are pattern diseases: one organ's structure fails, and the consequences cascade through pressures and pores. Reading their blood tests is applied pathophysiology.

The idea, with a picture. Cirrhosis: bridging fibrosis + regenerative nodules → **portal hypertension** → oesophageal varices, splenomegaly (→ thrombocytopenia), ascites, caput medusae; hepatocyte failure adds encephalopathy (ammonia), jaundice, coagulopathy (vitamin K, INR) ✓.



One blocked vein, five collateral consequences – each with its own emergency

The derivation (why nephritic and nephrotic are different holes). Nephritic: inflammation RIPS pores → red cells and protein leak modestly: haematuria, hypertension, oliguria, proteinuria <3.5 g/day ✓. Nephrotic: the podocytes and their slit diaphragms are damaged (in minimal change disease the foot processes flatten), so the barrier's size and charge selectivity fail → albumin floods through: proteinuria ≥3.5 g/day, hypoalbuminaemia, oedema, hyperlipidaemia ✓. Two patterns, two barrier lesions, two workups.

> **Worked example 1 – the liver panel decoded.** (Fictional) ALT/AST high with ALP normal = hepatocellular (injury pattern); ALP high with transaminases low-normal = cholestatic (obstruction pattern); albumin + INR = synthetic function – a failing liver makes less of everything ✓. Three questions, one panel.

Worked example 2 – clinical link: the urine that grades the barrier. (Fictional) frothy urine + periorbital oedema + albumin <3 g/dL + protein 4.2 g/day: nephrotic ✓. Coca-Cola urine with blood pressure up and red-cell casts: nephritic ✓. The dipstick and the microscope separate the two barrier lesions before any biopsy.

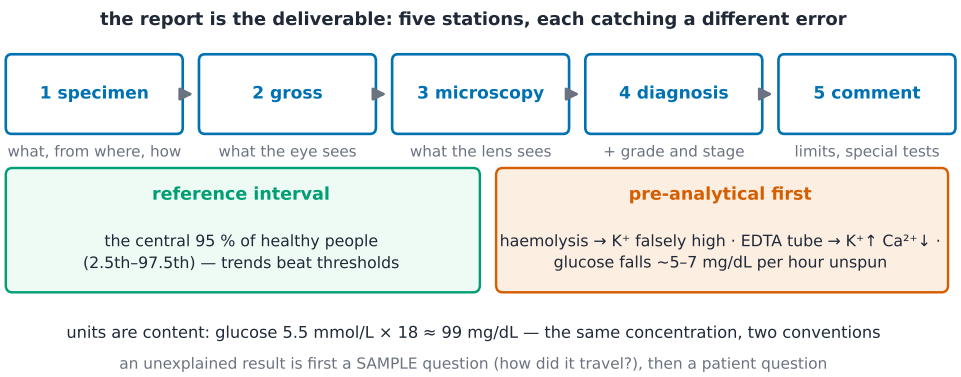
Common mistakes. Attributing varices to gastritis (they are portal pressure's escape valves). Forgetting cirrhosis's thrombocytopenia is splenic sequestration (a haemodynamic cause, not marrow failure). Reading "proteinuria" without quantity (3.5 g/day is the nephrotic line).

Summary. Cirrhosis = fibrosis + nodules → portal HTN: varices, splenomegaly, ascites, caput medusae; failure adds encephalopathy (NH_3), jaundice, coagulopathy. ALT/AST = injury, ALP = obstruction, albumin/INR = synthesis. Nephritic: torn pores (blood, HTN, <3.5 g); nephrotic: injured podocytes (≥3.5 g/day, hypoalbuminaemia, oedema, lipids).

Laboratory medicine and the pathology report

Why it matters. The laboratory is where mechanism becomes number – and where numbers are made wrong by tubes, time and temperature. The report is the deliverable of this whole course.

The idea, with a picture. A **reference interval** holds ~95 % of healthy people (2.5th–97.5th percentile); a structured report runs: specimen → gross description → microscopy → diagnosis (+ grade/stage) – plus comment ✓.



Specimen, gross, microscopy, diagnosis, comment – five stations, each catching a different error

The derivation (why pre-analytical errors dominate). The test measures the SAMPLE, not the patient: **haemolysis** bursts red cells (K^+ falsely high); EDTA tube wrong for chemistry (K^+ high, Ca^{2+} low); glucose falls ~5–7 % per hour in an unseparated tube (cells keep eating); delay and heat denature ✓. An unexplained result is first a sample question, then a patient question.

Worked example 1 – the unit trap. (Fictional) glucose 5.5 mmol/L vs 99 mg/dL – the same concentration, two conventions; multiply mmol/L by 18 for mg/dL ✓. A "high" result in the wrong units is a phone call, not a diagnosis.

Worked example 2 – clinical link: three impossible results. (Fictional) K^+ 8.2 with a calm ECG and a recently shaken tube: haemolysis until proven otherwise. Glucose "2.8" drawn after a delayed transport: pre-analytical consumption. Calcium "1.9" in an EDTA tube: tube error, redraw ✓. Each is caught by asking HOW the sample travelled before WHY the patient is ill.

Common mistakes. Treating the interval as normal/abnormal (5 % of healthy people, 2.5 % on each side, sit outside by definition; trends and prior probability decide). Skipping the specimen description in the report (it anchors the diagnosis). Believing the analyser cannot lie (it measures the tube's contents faithfully).

Summary. Interval = central 95 % of healthy; report stations: specimen, gross, microscopy, diagnosis, comment; pre-analytical first: haemolysis (K^+), wrong tube (EDTA: $K^+\uparrow$ $Ca^{2+}\downarrow$), glucose −5–7 % per hour unspun; units: glucose mmol/L ×18 = mg/dL.

Topic 11

University Medicine

What this course is, and how to use it

This is the eleventh course in University Medicine: **pharmacology** – what the body does to a drug (kinetics) and what a drug does to the body (dynamics). MED02 supplied the receptor chemistry, MED05 the physiology, MED08 the antimicrobial targets.

Scope limit, stated once: a study course, not a medical degree; every case is fictional, every number is an illustrative teaching value inside a labelled calculation, drug names are international nonproprietary names (INN), and nothing here is prescribing guidance or a regimen.

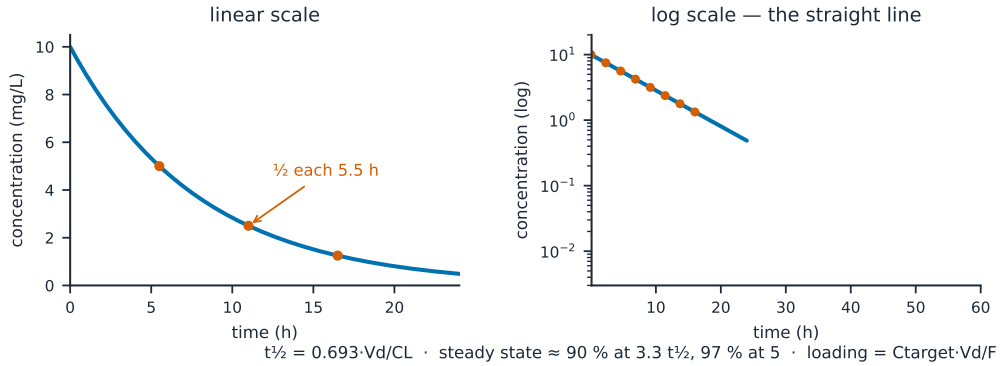
- Units 11.1–11.2 are the engine: pharmacokinetics and pharmacodynamics.
- Units 11.3–11.6 are systems: autonomic, cardiovascular/renal, antimicrobial/anticancer, CNS and analgesia.
- Unit 11.7 is safety: adverse reactions, interactions, error prevention.

The pharmacology habit: **mechanism predicts effect, arithmetic predicts exposure** – every curve below is read in both directions, and every half-life is something you recompute.

Pharmacokinetics

Why it matters. Kinetics turns a dose into a concentration – and concentration, not dose, is what the body sees. **Half-life** , clearance and volume are the three levers.

The idea, with a picture. Volume of distribution V_d (the apparent volume a drug "occupies"), **clearance** CL (volume cleared per time), linked by $t_{1/2} = 0.693 \cdot V_d/CL$ ✓; absorption with **bioavailability** F (fraction reaching circulation); **first-pass metabolism** cuts oral F ✓.



Half-life is the log-linear slope; steady state arrives after ~4–5 half-lives, loading doses skip the wait

The derivation (why 4–5 half-lives reach steady state). Repeated dosing adds drug each interval while clearance removes a fixed fraction: the gap closes exponentially toward **steady state** – after 3.3 half-lives ~90 % of steady state, after 5 half-lives ~97 % ✓. A **loading dose** fills the tank at once: $D_L = C_{target} \cdot V_d/F$ (illustrative calculation: target 10 mg/L, V_d 40 L, F 1.0 → 400 mg – a teaching value, not a prescription) ✓.

> **Worked example 1 – the three levers (illustrative values).** V_d 40 L, CL 5 L/h: $t_{1/2} = 0.693 \times 40/5 = 5.5$ h ✓. Halve clearance (liver disease): $t_{1/2}$ doubles to 11.1 h – the same dose now accumulates; steady state is later and higher ✓.

> **Worked example 2 – clinical link: the first-pass toll.** (Fictional teaching value) a drug with oral $F = 0.6$: of a labelled 100 mg oral dose, 60 mg reaches the circulation – the liver extracts 40 % on the first pass ✓. Portal failure or a bypass raises F and the same tablet behaves like a bigger one: kinetics, not pharmacology, changed the effect.

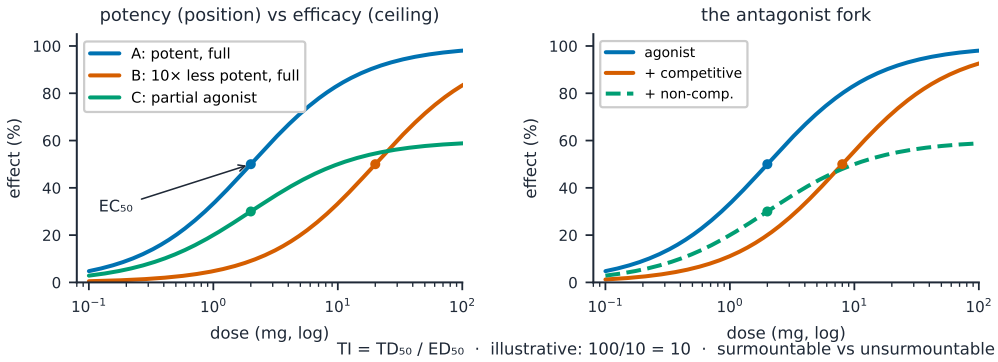
Common mistakes. Equating half-life with duration of effect (effect depends on concentration vs the drug's target range). Reading V_d as a real volume (it is a proportionality constant – about 3 L means the drug stays in plasma, 500 L means tissue binding). Forgetting maintenance dosing aims at $CL \cdot C_{ss}$.

Summary. $t_{1/2} = 0.693\, V_d/CL$; steady state ≈ 90 % at 3.3 $t_{1/2}$, 97 % at 5; loading $C_{target}V_d/F$; oral F = fraction past the first pass; clearance sets maintenance, volume sets the load.

Pharmacodynamics

Why it matters. Dynamics is the grammar of every drug chart: what a receptor does with a molecule – potency, efficacy, and the antagonists that undo both.

The idea, with a picture. The dose–response curve (log dose): position along the x-axis = **potency** (EC_{50}); ceiling = **efficacy** (E_{max}) ✓. Agonists reach E_{max}; **partial agonists** reach a lower ceiling (and displace full agonists); competitive antagonists shift the curve RIGHT in parallel; non-competitive lower E_{max} ✓.



Potency is position, efficacy is ceiling – antagonists shift right or cut the ceiling

The derivation (why the therapeutic index is a ratio). **Therapeutic index** $TI = TD_{50}/ED_{50}$ – the toxic dose over the effective dose, each read from its own cumulative curve (illustrative: TD_{50} 100 mg, ED_{50} 10 mg → TI 10) ✓. A drug with TI 2 asks for care; TI 100 forgives. Narrow-TI drugs (the ones needing level monitoring) are those whose toxic and effective curves overlap.

> **Worked example 1 – reading a curve pair.** (Fictional curves) drug A: EC_{50} 2 mg, E_{max} 100 %; drug B: EC_{50} 20 mg, E_{max} 100 % – B is 10× less potent but equally efficacious; drug C: EC_{50} 2 mg, E_{max} 60 % – a partial agonist, potent but capped ✓. Three claims, three curve features.

> **Worked example 2 – clinical link: the antagonist fork.** (Fictional) a competitive antagonist co-administered: the curve shifts right – the agonist's effect can be overcome by concentration (surmountable). A non-competitive antagonist cuts E_{max} – no concentration wins ✓. The overdose logic follows: one is a tug-of-war, the other removes the rope.

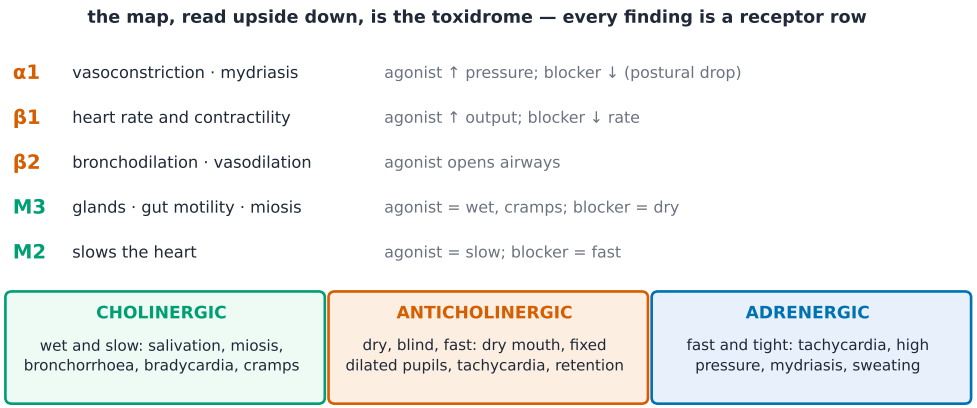
Common mistakes. Calling a potent drug a strong one (potency is dose-position; efficacy is the ceiling). Expecting a partial agonist to behave like a weak full agonist in combination (it COMPETES – lowering the mixture's ceiling). Reading TI from a single curve (it needs two: effect and toxicity).

Summary. Potency = EC_{50} position; efficacy = E_{max}; partial agonist = lower ceiling that displaces; competitive → parallel right shift (surmountable); non-competitive → lower ceiling; $TI = TD_{50}/ED_{50}$ (illustrative 100/10 = 10).

Autonomic pharmacology

Why it matters. The autonomic receptor map (MED09 §5) is the most-prescribed territory in medicine – every agonist and blocker is a prediction exercise, and the **toxidromes** are its emergencies.

The idea, with a picture. The code: $\alpha 1$ constricts vessels and dilates pupils; $\beta 1$ drives the heart; $\beta 2$ opens airways; M3 drives secretion, constricts the pupil and speeds the gut, while M2 slows the heart ✓. Agonists push the receptor; blockers remove its tone.



One map, five receptor rows – every autonomic drug is an arrow into it

The derivation (why toxidromes are the map upside down). Cholinergic excess = everything M stands for: salivation, lacrimation, miosis, bronchorrhoea, bradycardia, gut cramps and diarrhoea ✓. Anticholinergic = M blocked: dry mouth, dilated fixed pupils, tachycardia, urinary retention, hot dry skin ✓. Adrenergic excess = α+β on: tachycardia, hypertension, mydriasis, sweating. The syndrome is the arithmetic of the map.

> **Worked example 1 – predict four drugs.** (Fictional INN teaching set) a $\beta 1$ blocker lowers heart rate and contractility (cardio-selective logic); an $\alpha 1$ blocker lowers blood pressure and can cause postural drop; a $\beta 2$ agonist opens airways with little cardiac effect; an M blocker dries secretions and blurs vision ✓ – four predictions straight off the table.

> **Worked example 2 – clinical link: two overdoses, two maps.** (Fictional) patient A: pinpoint pupils, sweating, cramps and diarrhoea – cholinergic; patient B: dry, dilated, fast – anticholinergic ✓. Each finding is one receptor row; the collection names the class before any test returns.

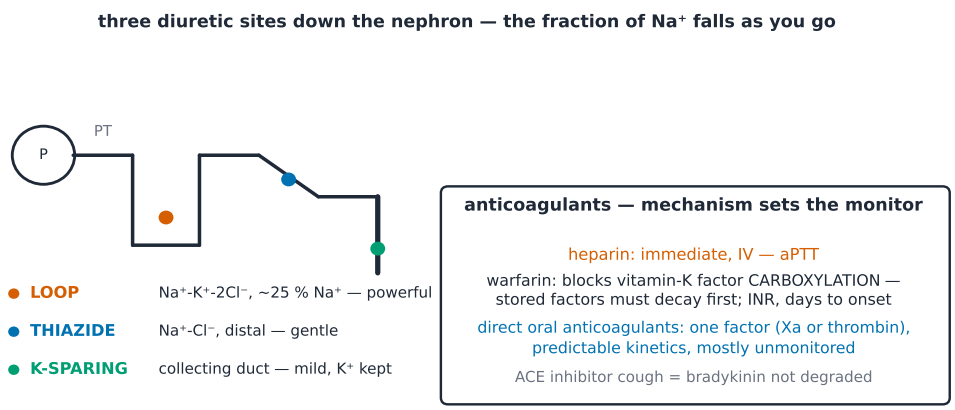
Common mistakes. Expecting β-blockade to lower blood pressure instantly (output falls first; the full antihypertensive effect takes weeks of remodelling). Forgetting muscarinic versus nicotinic (M effects dominate toxidromes; N is ganglia and the NMJ). Treating sweating as an adrenergic sign only (sweat glands are sympathetic but cholinergic, at muscarinic receptors – the MED09 exception, which is why the anticholinergic patient is hot and dry).

Summary. α1 constrict/dilate pupil · β1 heart · β2 airway · M3 secrete/constrict pupil/speed gut, M2 slow heart; agonist = push, blocker = silence; cholinergic = wet and slow; anticholinergic = dry, blind, fast; adrenergic = fast and tight; predictions come off the map, not memory.

Cardiovascular and renal drugs

Why it matters. Hypertension, heart failure and thrombosis are mechanism-rich: each drug class is a named pressure-regulating or clotting step turned down, and the nephron is the **diuretic** map.

The idea, with a picture. Antihypertensives by mechanism: diuretics (volume), **ACE inhibitors** ACE (block angiotensin II formation; the bradykinin cough), angiotensin receptor blockers, calcium-channel blockers (vasodilation), β1 blockers (output and renin) ✓. Diuretics by nephron site: loop (Na^+-K^+-$2Cl^-$, thick ascending, ~25 % of Na^+), thiazide (Na^+-Cl^-, distal tubule), potassium-sparing (collecting duct) ✓.



Loop, thiazide and potassium-sparing act at three descending sites – the fraction reabsorbed shrinks as you go

The derivation (why anticoagulants monitor differently). Unfractionated heparin acts immediately and is monitored with the aPTT; **warfarin** blocks vitamin-K clotting-factor synthesis – existing factors must decay first, so onset lags days and the INR measures the effect, not the dose ✓; direct oral **anticoagulants** target one factor (Xa or thrombin) with predictable kinetics, mostly unmonitored. The monitoring test is a property of the mechanism.

> **Worked example 1 – site to electrolyte.** (Fictional) loop diuretic at the thick ascending limb: powerful Na^+ loss, K^+ loss (downstream flow), and Ca^{2+} loss; a potassium-sparing agent at the collecting duct: mild, K^+-retaining ✓. Each site's chemistry predicts its electrolyte fingerprint.

> **Worked example 2 – clinical link: the cough that isn't infection.** (Fictional) on an ACE inhibitor, a dry persistent cough: bradykinin no longer degraded in that pathway ✓ – mechanism names the side effect, and the class switch (to an ARB) follows from the mechanism, not trial and error.

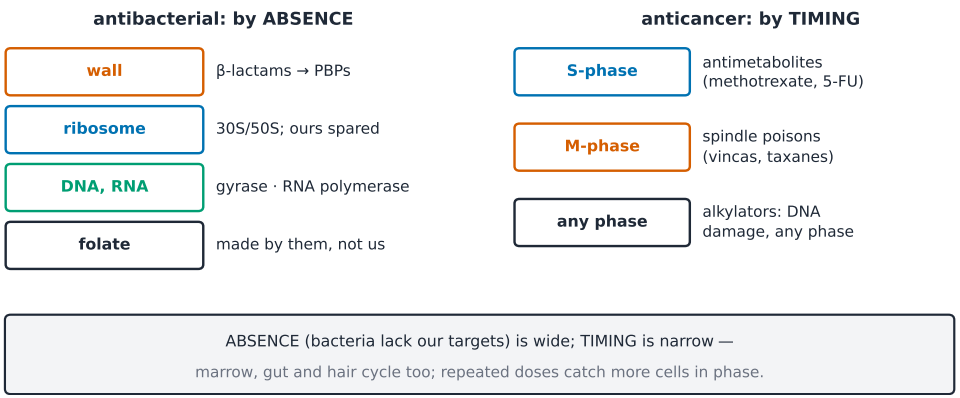
Common mistakes. Placing thiazides in the loop (they act distally – weaker, slower, Ca^{2+}-retaining). Expecting warfarin to act on factors already made (it blocks CARBOXYLATION of new ones). Confusing anticoagulants (clots) with antiplatelets (platelet plugs – different pipeline, different monitoring).

Summary. Volume/ACE/ARB/CCB/β1 = the antihypertensive mechanisms; loop 25 %, thiazide distal, K-sparing collecting duct; heparin→aPTT now; warfarin→INR after factor decay; DOACs one factor, predictable; mechanism names the side effect.

Antimicrobial and anticancer drugs

Why it matters. The two attack maps share one logic – **selective toxicity** against cells that differ from ours – and the plan defers to MED08 §6 for the antibacterial targets.

The idea, with a picture. Antibacterial targets (MED08): wall (β-lactams), ribosome, nucleic-acid enzymes, folate ✓. Anticancer: **cell-cycle-specific** agents hit cycling cells – S-phase (antimetabolites: methotrexate, fluorouracil), M-phase (vinca alkaloids, taxanes) – while alkylators damage DNA regardless of cycle ✓.



One target map across both fields: bacteria-specific targets left, cell-cycle phases right

The derivation (why selectivity differs between the fields). Antibacterial selectivity rides on structures human cells LACK (wall, different ribosome) – a wide margin. Anticancer selectivity rides on TIMING (cycling cells take the hit) – a narrow margin, because marrow, gut and hair cycle too ✓. The margin predicts the side-effect profile before any list is memorised.

> **Worked example 1 – the arithmetic of resistance (from MED08).** β-lactamase cleaves the ring (enzyme); altered PBP2a is target reshaping; the 10^{-7}–10^{-9} per-cell mutation supply explains why combinations outlast single agents ✓ – the pharmacology of failure is arithmetic.

> **Worked example 2 – clinical link: schedule is mechanism.** (Fictional teaching schedule) an S-phase antimetabolite given as a short bolus misses cells not in S at that moment – the same total dose fractionated over time meets more cells in their vulnerable phase; an M-phase vinca depends on cells arriving at mitosis ✓. Cycle specificity dictates scheduling logic.

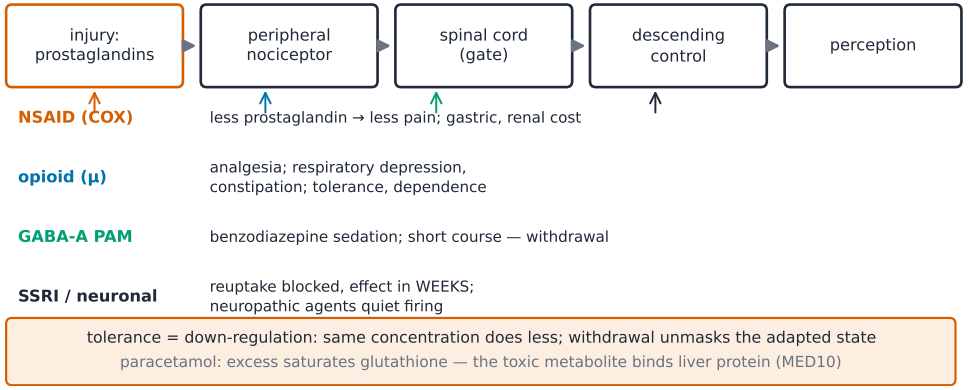
Common mistakes. Expecting anticancer "selectivity" to spare normal cycling tissue (marrow and mucosa cycle – the predictable toxicities). Treating resistance as a drug failure only (it is evolution in a dish: pressure selects). Calling antimetabolites non-specific (they are S-phase SPECIFIC).

Summary. Antibacterial: wall/ribosome/DNA/folate (MED08). Anticancer: S-phase antimetabolites, M-phase vinca/taxanes, alkylators cycle-independent; selectivity by ABSENCE (bacteria) vs TIMING (cancer); schedule follows the cycle.

CNS drugs and analgesia

Why it matters. CNS pharmacology is receptor physiology named: opioids, sedatives and antidepressants each sit at a known synapse – and **tolerance** and **dependence** are the learning the synapse does.

The idea, with a picture. **Opioids** act at μ receptors: analgesia with respiratory depression, constipation, miosis; tolerance and dependence develop ✓. **NSAIDs** block cyclo-oxygenase (COX-1 housekeeping/COX-2 inflammatory): less prostaglandin → less pain and inflammation, but gastric and renal cost ✓. Benzodiazepines are positive allosteric modulators at GABA-A; SSRIs block serotonin reuptake (weeks to full effect – receptor adaptation, not the synapse alone).



Peripheral COX, spinal modulation, descending opioid control – the pain pathway with its drug sites marked

The derivation (why tolerance is physiology, not weakness). Chronic receptor activation down-regulates and desensitises: the same concentration does less – larger doses restore effect (tolerance), and abrupt withdrawal unmasks the adapted state (dependence) ✓. This is why benzodiazepines are short-course and why opioids' effective dose climbs; the WHO analgesic ladder (1986, revised 1996) is the classic framework naming pain severity to drug class.

> **Worked example 1 – three pains, three sites.** (Fictional) prostaglandin-mediated inflammation – a NSAID at the peripheral COX site; severe post-operative pain – an opioid at μ receptors (with the respiratory price); neuropathic burning – neither class alone; agents acting on neuronal excitability ✓ – matching the site to the drug, not the severity alone.

> **Worked example 2 – clinical link: the paracetamol lesson.** (Fictional teaching case) an overdose of paracetamol saturates the safe conjugation pathways and uses up glutathione, so the toxic metabolite (NAPQI) binds liver proteins – the MED10 necrosis rulebook applies ✓. Mechanism names both the antidote logic (acetylcysteine restores glutathione) and the organ at risk.

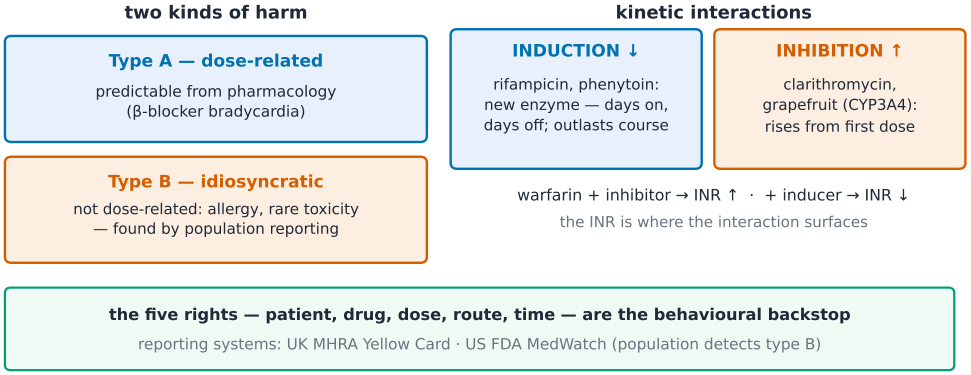
Common mistakes. Equating analgesic classes by strength (they act at different SITES – combinations are additive by mechanism). Expecting SSRIs to work in days (reuptake blockade is immediate; clinical effect needs weeks of adaptation). Forgetting constipation is the opioid effect that barely tolerises.

Summary. Opioid μ: analgesia + respiratory depression + constipation, tolerance/dependence; NSAID COX-1/2: periphery, gastric/renal cost; benzodiazepine GABA-A PAM; SSRI reuptake, weeks; tolerance = down-regulation; paracetamol = glutathione and liver.

Adverse reactions, interactions and safe-medication principles

Why it matters. Every drug chart is an experiment with **adverse drug reactions** , interactions and human error – the last unit is where the course's mechanisms become safety habits.

The idea, with a picture. Type A reactions: DOSE-RELATED, predictable from pharmacology (β-blocker bradycardia); type B: idiosyncratic, not dose-related (allergy, rare toxicity) ✓. Interactions: **enzyme induction** (rifampicin, phenytoin → a co-medication's levels fall) versus **enzyme inhibition** (clarithromycin, grapefruit at CYP3A4 → levels rise) ✓.



Induction lowers exposure, inhibition raises it – and the ADR types split by dose-relatedness

The derivation (why induction and inhibition move levels for days). Induction builds new enzyme – days to build, days to unwind; inhibition occupies existing enzyme – begins with the first co-dose ✓. So a new inhibitor raises a substrate's level within doses (warfarin effect climbs with it), while an inducer's effect arrives late and outlasts its stop date. Timing is the mechanism made visible.

> **Worked example 1 – one interaction, three questions.** (Fictional) warfarin + a new CYP inhibitor: INR rises (more warfarin effect); warfarin + rifampicin: INR falls (induction clears it); either change is detectable by the INR's rhythm, which is why monitoring interval follows interaction risk ✓.

> **Worked example 2 – clinical link: the five rights and the report.** (Fictional ward) a near-miss caught at the bedside: the five rights – right patient, drug, dose, route, time – are the behavioural backstop for every mechanism above; national reporting systems (the UK MHRA Yellow Card scheme; the US FDA MedWatch) exist because type B reactions are found only by population reporting ✓ – dated jurisdiction-bound statements, as the plan requires.

Common mistakes. Blaming type B reactions on dose (they are idiosyncratic – dose neither predicts nor prevents them). Forgetting the inducer's tail (its effect on enzyme outlasts the course). Treating interactions as pharmacological trivia – most are kinetic, and the INR/aPTT/level is where they surface.

Summary. Type A = dose-related, type B = idiosyncratic; induction → levels fall (slow on, slow off), inhibition → levels rise (fast on); five rights as the behavioural backstop; Yellow Card (UK) / MedWatch (US) for population detection of type B.

Topic 12

University Medicine

What this course is, and how to use it

This is the twelfth course in University Medicine: **epidemiology** , **biostatistics** and **evidence-based medicine** (EBM) – the toolkit that turns clinical questions into numbers and numbers into decisions. It aligns with the reviewed UNI-BUS-05 statistics material where the mathematics overlaps.

Scope limit, stated once: a study course, not a medical degree; every dataset below is synthetic or openly licensed, every abstract appraised is fictional, and nothing here licenses or guides treatment.

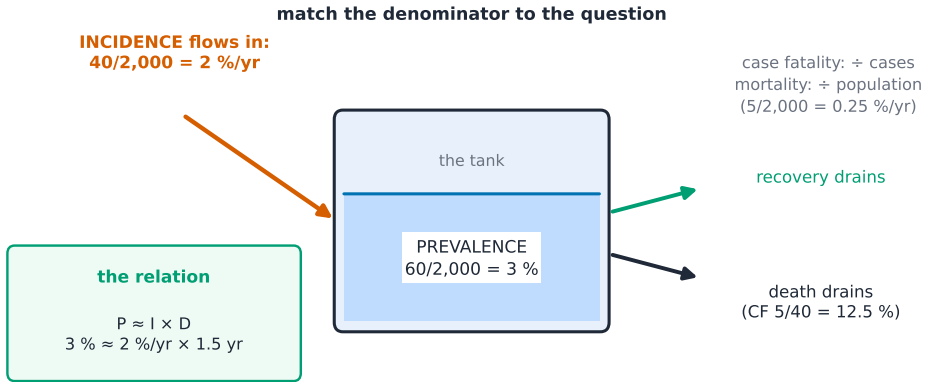
- Units 12.1–12.3 measure and compare: disease frequency, study designs, bias and causation.
- Units 12.4–12.6 quantify: diagnostic tests, inference, trials and reviews.
- Unit 12.7 applies: appraisal, certainty, and the patient in front of you.

The EBM habit: **ask in PICO, measure in risks and ratios, decide with intervals not rituals** – every calculation below is synthetic, machine-checked, and something you recompute.

Measuring disease: incidence and prevalence

Why it matters. Before any comparison, disease must be counted – and the two classic counts answer different questions: how fast new cases appear, and how many exist at once.

The idea, with a picture. Incidence = new cases ÷ people at risk ÷ time; **prevalence** = existing cases ÷ population at an instant ✓. They are linked: prevalence ≈ incidence × duration – long diseases pile up, brief ones vanish between surveys.



Incidence fills the tank; cure and mortality drain it – prevalence is the water level

The derivation (why standardisation exists). Crude rates mix two effects: disease risk and population age structure. A district with more older residents can show a higher crude rate with LOWER age-specific rates ✓. Direct standardisation applies each group's age-specific rates to one standard population and sums – the artificial total is comparable because the age mix is now identical.

> **Worked example 1 – a synthetic town.** 2,000 residents, one year, 40 new cases of a disease → incidence 40/2,000 = 2 % per year ✓. On survey day 60 people have the disease → point prevalence 60/2,000 = 3 % ✓. Check the link: 3 % ≈ 2 %/year × 1.5 years average duration – the numbers are a consistent system.

> **Worked example 2 – clinical link: mortality's three faces.** (Fictional dataset) of the 40 new cases, 5 died within the year: **case fatality** 5/40 = 12.5 % ✓; the yearly deaths from this cause in the whole town were 5/2,000 = 0.25 % (cause-specific mortality rate); and if the town recorded 20 deaths from all causes that year, this disease caused 5/20 = 25 % of them (proportional mortality). Same deaths, three denominators, three different questions.

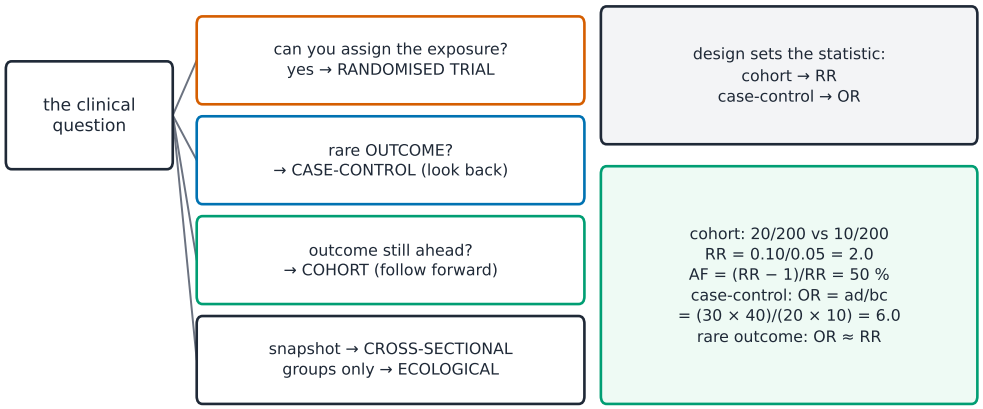
Common mistakes. Calling prevalence a rate (it is a proportion – no time unit). Comparing crude rates across populations with different age structures (standardise first). Reading case fatality as mortality (the denominator is cases, not people).

Summary. Incidence = new/at-risk/time; prevalence = existing/population; P ≈ I × D; standardisation fixes the age mix; case fatality (cases), cause-specific mortality (population), proportional mortality (all deaths) – match the denominator to the question.

Study designs

Why it matters. Every clinical question has designs that answer it well and designs that cannot – choosing the design is choosing which biases you will fight.

The idea, with a picture. The ladder: randomised trial (experiment) → cohort (follow exposed and unexposed forward) → case-control (look backward from outcome) → cross-sectional (one snapshot) → ecological (groups, not individuals) ✓.



Rare disease or outcome → case-control; exposure measured ahead of a common outcome → cohort; both questions at once → trial

The derivation (why RR belongs to cohorts and OR to case-control). A cohort samples by EXPOSURE, so risks in each arm are estimable: **relative risk** $RR = \frac{a/(a+b)}{c/(c+d)}$ ✓. A case-control samples by OUTCOME, so the "risk" in each arm is an artefact of the sampling ratio – only the **odds ratio** $OR = \frac{a/c}{b/d} = \frac{ad}{bc}$ is valid ✓. The design dictates the statistic, not preference.

> **Worked example 1 – one synthetic cohort.** Exposed 200, 20 events; unexposed 200, 10 events: risks 0.10 vs 0.05 → $RR = 2.0$ ✓, attributable fraction among exposed $(RR-1)/RR = 50\ \%$ ✓ – the exposed group's events are half "because of" the exposure under the assumptions.

> **Worked example 2 – clinical link: the case-control mirror.** (Fictional) 30 of 40 cases exposed; 20 of 60 controls exposed: $OR = \frac{30\times40}{10\times20} = 6.0$ ✓. A rare outcome makes OR approximate RR – one number, two designs, and the sampling decides which is legal.

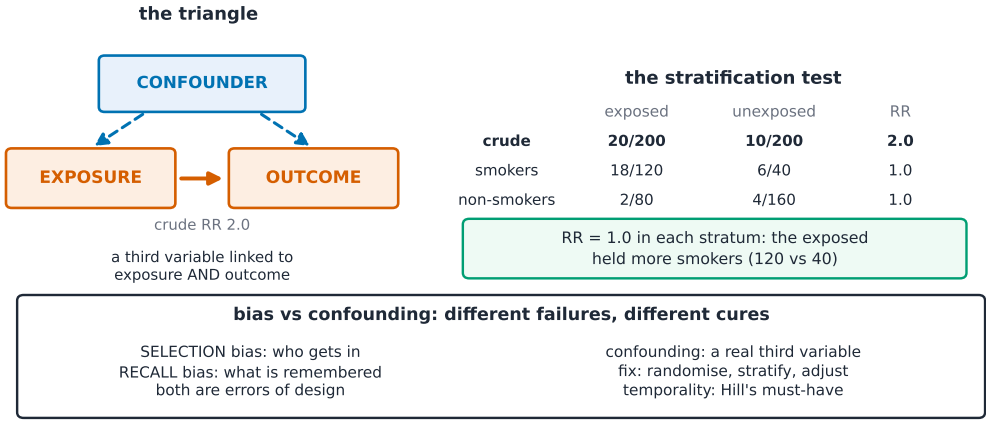
Common mistakes. Computing RR from a case-control study (the denominators are chosen by the researcher, not found). Reading a cross-sectional odds ratio as causal (exposure and outcome are simultaneous). Equating "prospective" with "better" (a prospective cohort still confounds; only randomisation balances the unknowns).

Summary. Trial → cohort (RR) → case-control (OR) → cross-sectional → ecological; $RR = \frac{a/(a+b)}{c/(c+d)}$, $OR = ad/bc$; cohort samples by exposure, case-control by outcome; rare-disease OR ≈ RR; attributable fraction $(RR-1)/RR$.

Bias, confounding and causation

Why it matters. Numbers lie in three ways: systematically (bias), through a third variable (confounding), and by association without cause. Naming the failure mode is half of critical appraisal.

The idea, with a picture. **Bias** is systematic error in design or measurement: **selection bias** (who gets in), information/recall bias (what is remembered or recorded). **Confounding** is a third variable linked to both exposure and outcome that pretends to be the cause ✓.



The confounder sits on the back path – stratify and the spurious association collapses

The derivation (why stratification is the proof). If a third variable is the true driver, the association should vanish inside strata where it is constant: crude RR 2.0, but within non-smokers RR 1.0 and within smokers RR 1.0 → smoking was the confounder ✓. If the association persists within strata (say 2.0 and 1.9), the exposure keeps its own signal. Randomisation prevents confounding prospectively by balancing even the unknowns ✓.

Worked example 1 – the stratified table. (Synthetic) crude: exposed 20/200 events vs unexposed 10/200 → RR 2.0. Stratify by smoking: non-smokers 6/100 vs 3/100 (RR 2.0 ✓ inside), smokers 14/100 vs 7/100 (RR 2.0) – the association survives **stratification** : smoking was NOT the confounder here; the exposure's signal stands. Now give the same crude table a different split – exposed: 120 smokers (18 events) and 80 non-smokers (2 events); unexposed: 40 smokers (6 events) and 160 non-smokers (4 events). Smokers' risk is 15 % in both arms and non-smokers' 2.5 % in both, so RR is 1.0 inside each stratum: the exposed group simply held more smokers, and smoking WAS the confounder ✓. One crude table, two verdicts – only the strata can tell them apart.

Worked example 2 – clinical link: the appraisal questions. (Fictional abstract) "coffee drinkers had twice the risk": ask who was studied (selection), how exposure was measured (recall – patients with illness remember differently), and what a confounder would need to look like (associated with coffee AND with the outcome, like smoking). Bradford Hill's viewpoints – temporality first, then strength, consistency, dose-response, plausibility – turn a headline into a checklist ✓.

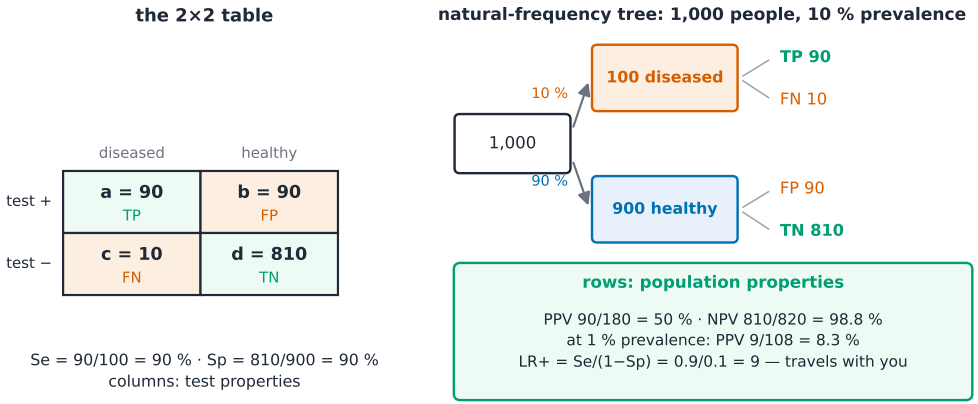
Common mistakes. Treating confounding and measurement bias as one problem (confounding is a mixing of effects through a third variable, cured by randomisation, stratification or adjustment; selection and information bias are errors in who is studied and how things are measured, which no analysis fully repairs). Expecting big data to cure confounding (more precision around a biased estimate). Demanding every Hill viewpoint (they are viewpoints, not a scorecard; temporality is the only non-negotiable).

Summary. Bias = systematic error (selection, information/recall); confounder = third variable on both paths; the stratification test collapses the association if confounded; randomisation balances the unknowns; Hill's viewpoints with temporality first.

Diagnostic tests

Why it matters. A test's numbers are not properties of the test alone – sensitivity and specificity belong to the test, but **predictive values** belong to the population using it. Misreading this is the commonest quantitative error in clinical reasoning.

The idea, with a picture. The 2×2: **sensitivity** = a/(a+c) (diseased, test positive); **specificity** = d/(b+d) (healthy, test negative); **PPV** = a/(a+b), **NPV** = d/(c+d) – and the natural-frequency tree shows why prevalence moves the predictive values ✓.



Sensitivity and specificity live in the columns; predictive values live in the rows – prevalence sets the row widths

The derivation (why prevalence drives PPV). PPV = $\frac{Se \times P}{Se \times P + (1-Sp)(1-P)}$ ✓. With Se = Sp = 90 %: at prevalence 10 %, PPV = $\frac{0.9 \times 0.1}{0.9 \times 0.1 + 0.1 \times 0.9} = 0.5$ ✓ – even an excellent test is a coin-flip in a low-prevalence clinic. The false-positive arm scales with the healthy majority.

Worked example 1 – the natural-frequency tree (synthetic). 1,000 people, prevalence 10 %: 100 diseased → 90 true positives, 10 false negatives; 900 healthy → 810 true negatives, 90 false positives ✓. PPV = 90/180 = 50 %; NPV = 810/820 = 98.8 % ✓. Same test at prevalence 1 %: 9 TP vs 99 FP → PPV 9/108 = 8.3 % ✓ – one tree, two worlds.

Worked example 2 – clinical link: the likelihood ratio travels. $LR+ = Se/(1-Sp) = 0.9/0.1 = 9$ ✓: a positive result multiplies the pre-test odds by 9 whatever the clinic's prevalence. Pre-test odds 1:9 (10 %) → post-test 9:9 = 1:1 → probability 50 % – the same PPV, derived in one line. **Likelihood ratios** (LRs) are the portable form of the test; predictive values are local.

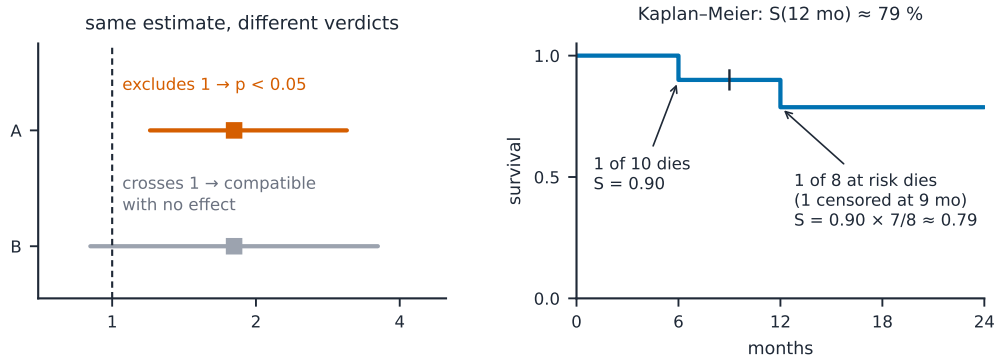
Common mistakes. Quoting sensitivity as "the test is 90 % right" (it is right about 90 % of the DISEASED – say of what). Applying a textbook PPV to a different-prevalence clinic. Confusing the two columns (sensitivity reads down the diseased column; specificity down the healthy column).

Summary. Se = a/(a+c), Sp = d/(b+d) – test properties; PPV/NPV – population properties, move with prevalence; natural-frequency trees make it visible; $LR+ = Se/(1-Sp)$ multiplies odds portably (0.9/0.1 = 9); 50 % PPV at 10 % prevalence despite Se = Sp = 90 %.

Statistical inference for clinical data

Why it matters. Inference is how a sample speaks about a population – and misread p-values and confidence intervals are the signature errors of the modern literature.

The idea, with a picture. A 95 % **confidence interval** is the range of true values compatible with the data at that level; a **p-value** P is the probability of data at least this extreme IF the null hypothesis were true ✓. The CI summarises size and precision; the p-value only summarises surprise.



The interval carries the estimate's size and precision; the survival curve steps down at events and drops at censoring gaps

The derivation (why the test follows the data type). Counts by group → χ^2; two means (normal-ish) → t-test (paired when each row is measured twice); several means → ANOVA; skewed/small-n → rank (non-parametric) alternatives; two continuous measurements → correlation (Pearson, or Spearman for ranks) ✓. The measurement scale decides the test – deciding after looking at results invites the wrong question.

> **Worked example 1 – reading one interval.** (Fictional abstract) RR 1.8, 95 % CI 1.2–3.1: the interval excludes 1, so the data are incompatible with "no effect" at this level (p < 0.05 follows); the width says the estimate is imprecise but consistently above 1 ✓. An RR 1.8 with CI 0.9–3.6 says the opposite – same point estimate, different verdict, and only the interval tells you.

> **Worked example 2 – clinical link: the Kaplan–Meier arithmetic.** (Synthetic) 10 patients enrolled; at 6 months 1 event among 10 at risk: $S = 9/10 = 0.90$; at 12 months 1 event among the 8 still at risk (one censored): $S = 0.90 \times 7/8 = 0.7875 \approx 79$ % ✓. Median survival is the time the curve crosses 50 %; censoring removes people from the denominator, not from the estimate's logic.

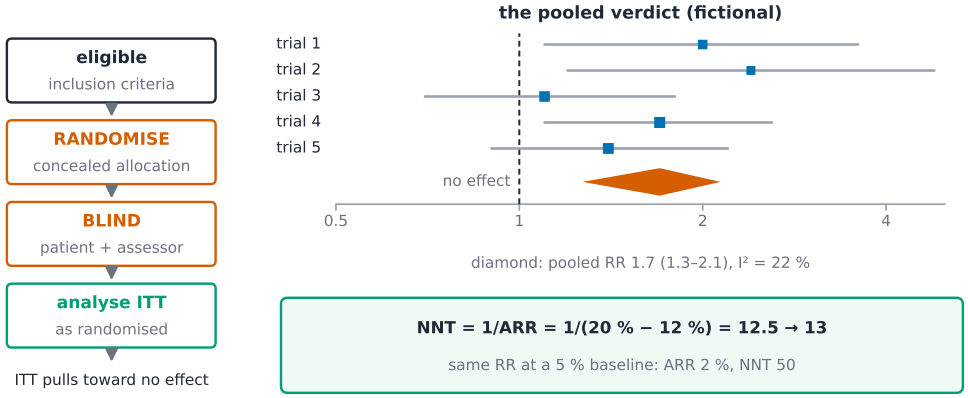
Common mistakes. "p > 0.05 means no difference" (it means the data are compatible with no difference – absence of evidence ≠ evidence of absence). Reading the CI as a 95 % probability the truth lies inside (the interval is a property of the procedure, applied to this sample's estimate). Choosing the test after seeing which gives p < 0.05.

Summary. CI = compatible truths, carries size and precision; p-value = P(data this extreme | null); never "proven" or "no effect" from p alone; χ^2/t/ANOVA/rank/correlation chosen by data type; the **Kaplan–Meier curve** - multiplies (1 − events/at-risk); median = 50 % crossing.

Clinical trials and systematic reviews

Why it matters. Trials are the machinery of proof and reviews are its assembly – the vocabulary (randomisation, blinding, ITT, NNT) is what every abstract and every forest plot assumes you know.

The idea, with a picture. A trial's spine: eligibility → randomisation (assignment by chance, allocation concealed) → blinding (participant, clinician, outcome assessor) → analysis by **intention to treat** (analyse as randomised, everyone counted) ✓. Reviews pool into a **forest plot** : squares (size = weight), lines (CI), diamond (pooled), and **heterogeneity** I^2 ✓.



Randomise, blind, analyse as randomised – then pool: squares weighted by precision, the diamond the pooled verdict

The derivation (why NNT is the clinical size of an effect). NNT $= 1/ARR$: control event rate 20 %, treated 12 % → ARR 8 % → NNT = 1/0.08 = 12.5 → 13 people treated for one to benefit ✓. NNT converts an abstract ratio into a workload – and its denominator (ARR) shrinks as baseline risk falls, which is why trial results do not transplant to lower-risk clinics unchanged.

Worked example 1 – the ITT logic. (Fictional trial) 20 % control vs 12 % treated: ARR 8 %, NNT 13 ✓. If 30 % of the treated arm stopped their assignment early, ITT still counts them as treated – diluting the apparent effect toward the null – because the question is the POLICY of assignment, not the biology of the adherent ✓. Per-protocol answers a different question and inherits adherence bias.

Worked example 2 – clinical link: reading the diamond. (Fictional forest) five trials, individual RRs 1.1–2.4 with wide CIs, pooled RR 1.7 (1.3–2.1), $I^2 = 22$ % ✓: the diamond sits right of 1 and its interval excludes 1, with low heterogeneity – a consistent effect across settings. The same diamond with $I^2 = 78$ % would prompt "why do these trials disagree?" before any pooling is trusted.

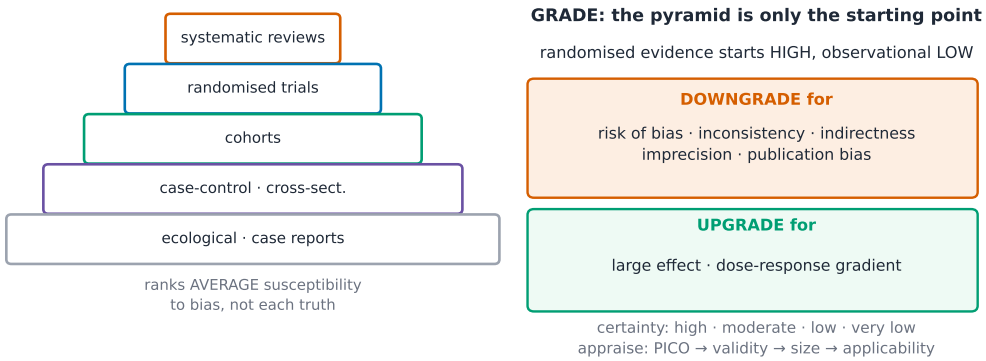
Common mistakes. Equating randomisation with blinding (chance balances groups; blindness balances expectations – a trial can have either alone). Reading per-protocol as the honest analysis (it breaks the randomisation's protection). Treating the pooled diamond as gospel despite high heterogeneity.

Summary. Randomise (concealed) → blind → analyse as randomised (ITT, dilutes toward null); NNT = 1/ARR (20 % → 12 %: ARR 8 %, NNT 13); forest: squares weighted, diamond pooled, I^2 disagreement; low-baseline clinics shrink ARR and inflate NNT.

Critical appraisal and applying evidence

Why it matters. Evidence becomes medicine only at the bedside: appraised for validity, sized for importance, and matched to the patient in front of you. This unit is the course's synthesis.

The idea, with a picture. The appraisal spine: **PICO** (population, intervention, comparison, outcome) framing → validity (were the results biased?) → size (how large, how precise?) → applicability (does my patient resemble the studied population?) ✓. Certainty grading (**GRADE**): randomised evidence starts high and observational evidence low; downgrade for risk of bias, inconsistency, indirectness, imprecision or publication bias; upgrade for large effects or a dose-response gradient ✓.



The pyramid ranks average susceptibility to bias, not the truth of each study – a flawed trial can mislead more than a good cohort

The derivation (why the hierarchy is a prior, not a verdict). Designs differ in their TYPICAL bias control: randomisation balances confounders measured and unmeasured; cohorts carry confounding; case-control adds recall bias ✓. But the hierarchy grades the design class, while the verdict belongs to the study in hand – a small, unblinded, per-protocol-analysed trial may deserve less trust than a large, pre-registered cohort with propensity adjustment. GRADE exists precisely because design class is only the starting point.

Worked example 1 – a fictional abstract in five questions. (Fictional) "Randomised, double-blind, n = 400; RR 0.75 (0.60–0.94)": validity – randomised and blinded ✓; size – 25 % relative reduction, interval excludes 1; local applicability – were patients like mine enrolled (age, severity, co-medications)?; outcomes – patient-relevant or surrogate?; certainty – any downgrade features? Five questions, one structured verdict ✓.

Worked example 2 – clinical link: transplanting the NNT. (Fictional) a trial in severe disease (baseline 20 %, ARR 8 %, NNT 13); my patient's clinic runs at 5 % baseline: with the same relative risk reduction the ARR scales to $8\,\% \times 5/20 = 2\,\%$ → NNT 50 ✓. The relative effect may travel; the absolute burden – the workload and the harm-benefit arithmetic – is local.

Common mistakes. Appraising only for validity (a valid but tiny effect is still useless: size and precision decide importance). Grading by journal or design label alone (GRADE starts there and moves with the evidence's flaws). Applying group averages to the patient whose defining feature was an exclusion criterion.

Summary. PICO → validity → size → applicability; GRADE: start from design, downgrade for bias/inconsistency/indirectness/imprecision, upgrade for large effects; the pyramid is a prior, not a verdict; relative effects travel, absolute benefits (ARR, NNT) are local – recompute before applying.

Topic 13

University Medicine

What this course is, and how to use it

This is the thirteenth course in University Medicine: **clinical foundations** – communication, reasoning and ethics. MED05–MED12 built the science; this course builds the consultation, the reasoning and the judgement that turn science into care.

Scope limit, stated once: a study course, not a medical degree; every consultation and case below is fictional, communication practice is role-play and not assessed clinical skill, ethics statements name their jurisdiction, and nothing here licenses or guides treatment.

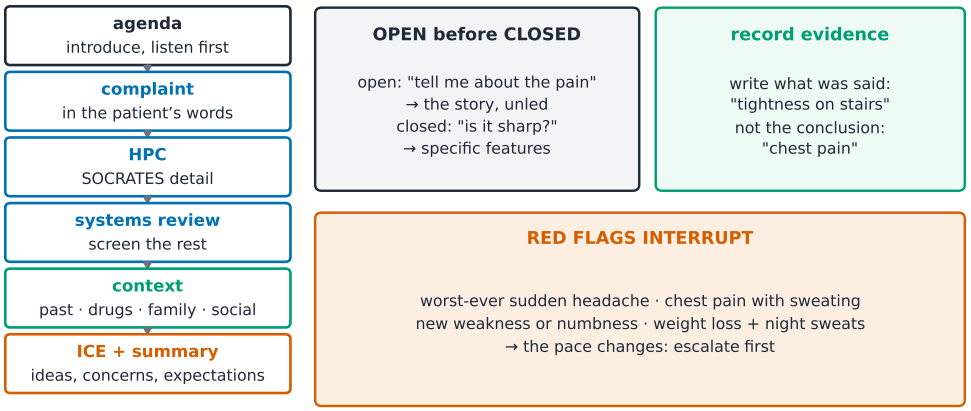
- Units 13.1–13.3 are the encounter: history, examination, reasoning.
- Units 13.4–13.5 are the relationship: communication and ethics.
- Units 13.6–13.7 are the system: safety and public health.

The clinical habit: **structure carries you when knowledge is incomplete** – every framework below is a checklist you can hold under pressure, and every case is something you reason through, not look up.

The consultation and history-taking

Why it matters. The history is the diagnosis most of the time: symptoms, timing and context arrive as words before any test is ordered. A structured history is the clinical interview's grammar.

The idea, with a picture. The standard arc: introducing yourself and setting the agenda → the **presenting complaint** in the patient's words → **history of presenting complaint** (onset, character, radiation, timing, exacerbating/relieving, severity) → **systems review** → past history, medication history (INN names, MED11), family and social history → summarising and ICE: ideas, concerns, expectations ✓.



Open questions open the story; closed questions close the details; red flags interrupt

The derivation (why open before closed). Open questions ("tell me about the pain") elicit the patient's frame without leading; closed questions ("is it sharp?") harvest specific features. Reversing the order primes the patient's answers and loses the story's own sequence – and the sequence IS the diagnostic information (timing patterns from MED08's clocks) ✓.

> **Worked example 1 – the pain frame.** (Fictional consultation, role-play text) "Tell me about the chest discomfort" → patient gives 40 seconds: central, tight, with stairs, two weeks, worse each morning. Six closed follow-ups place character, radiation, timing and associations. The SOCRATES frame is a checklist, not a script ✓.

> **Worked example 2 – clinical link: red flags interrupt.** (Fictional) a headache history unearths "sudden, worst ever, with vomiting" – a **red flag** that changes the encounter's pace whatever the differential suggests ✓. Escalation criteria are taught as patterns (severity, speed, associated deficits), so the history actively screens for them.

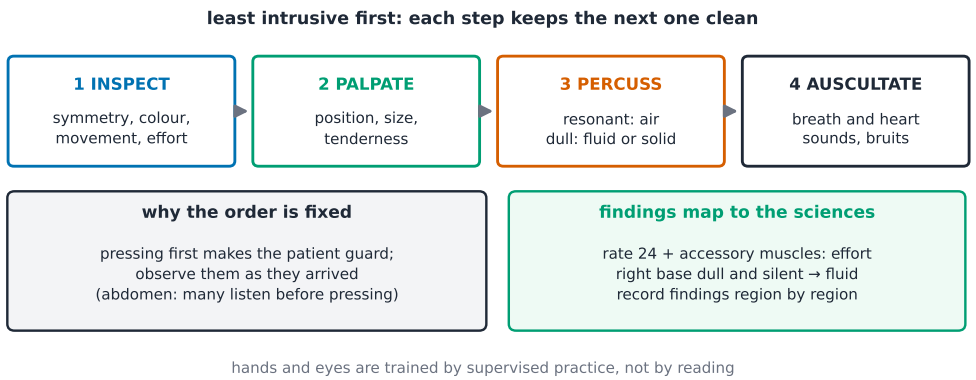
Common mistakes. Writing the history as conclusions ("chest pain") instead of the patient's evidence ("tightness on stairs"). Skipping medication and social history (interactions from MED11; occupation and home drive risk). Never asking what the patient thinks is wrong (ICE) – you cannot align a plan with unvoiced fears.

Summary. Agenda → complaint in patient's words → SOCRATES detail → systems review → past/medication/family/social → ICE and summary; open before closed; red flags interrupt the routine; record in the standard sections.

Principles of physical examination

Why it matters. Examination is physiology made visible (MED05) on anatomy made palpable (MED03): each manoeuvre asks a question a history cannot answer.

The idea, with a picture. The four-step sequence: **inspection** (look before touching: symmetry, colour, movement), **palpation** (feel: position, size, tenderness, crepitus), **percussion** (tap: resonant lung, dull fluid, tympanic gut), **auscultation** (listen: breath sounds, heart sounds, bruits) – ordered least to most intrusive so each step preserves the next ✓.



Least intrusive first: every step keeps the conditions the next step needs

The derivation (why the order is fixed). Palpating a tender abdomen changes the patient's guarding and breath pattern before you inspect; percussing before auscultating the chest disturbs the sounds. The sequence is designed so that each observation is made in the state the patient arrived in ✓ – a checklist against contaminating your own data. (The abdomen is the usual exception: many clinicians auscultate before palpating, because pressing changes bowel sounds – the same principle.)

> **Worked example 1 – reading the respiratory pattern.** (Fictional teaching case) respiratory rate 24 at rest, accessory muscle use on inspection, reduced expansion right side, dull to percussion at the right base, absent breath sounds there: each finding maps to MED05's ventilation arithmetic and MED10's dull-fluid rule – a pleural effusion read from the four steps ✓.

> **Worked example 2 – clinical link: what reading cannot replace.** (As a statement of principle) volumes, consistency and comparative findings (the left side against the right, the patient against the expected) are perceptual skills built by supervised practice; the sequence and the questions are knowledge, the hands and eyes are training ✓.

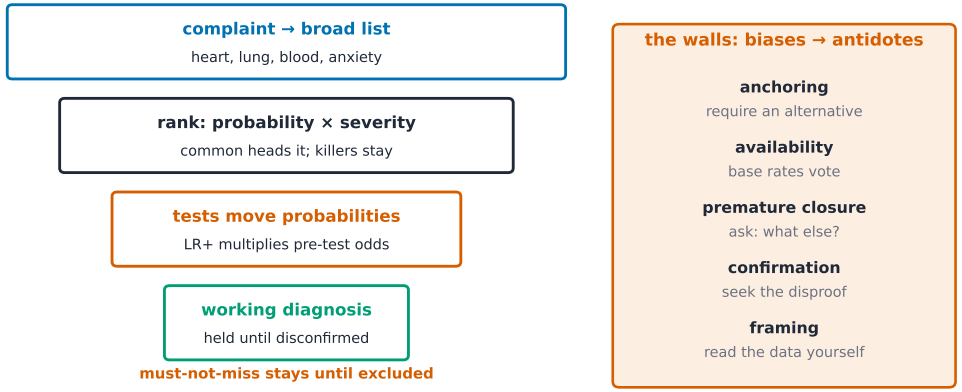
Common mistakes. Skipping inspection (the single highest-yield step). Auscultating through clothing or before percussion. Examining the painful region first (guarding contaminates everything after). Recording "exam normal" instead of what was examined.

Summary. Inspect → palpate → percuss → auscultate, least intrusive first; compare sides and systems; every finding maps to MED03 anatomy and MED05 physiology; findings recorded per region, never as "normal".

Clinical reasoning and differential diagnosis

Why it matters. Reasoning is the bridge from findings to action – and knowing its failure modes (the biases) is patient safety for the mind.

The idea, with a picture. The **differential diagnosis** funnel: anatomy of the complaint → broad candidate list → rank by probability × severity ("what is common, what can kill") → tests move probabilities (MED12's likelihood ratios!) → the working diagnosis → safe disposition of what remains ✓.



Broad to ranked to tested to working – and the funnel's walls are the biases to watch

The derivation (why probability and severity both rank). Ranking by probability alone misses rare killers; by severity alone over-tests the common. The standard compromise: must-not-miss diagnoses stay on the list until excluded, however unlikely; probable benign diagnoses head the working list ✓. Each test updates with its LR (MED12 §4): a positive with LR+ 9 multiplies the pre-test odds by 9 – reasoning that computes.

> **Worked example 1 – the ranked list.** (Fictional case) 58-year-old with acute breathlessness: candidates from anatomy (heart, lung, blood, anxiety) → ranked: heart failure, infection, clot (the must-not-miss), anaemia. Each finding and test reorders: orthopnoea raises failure; a sharp drop in oxygen saturation with a clear chest X-ray raises the clot ✓.

Worked example 2 – clinical link: the bias roll-call. (Named patterns) anchoring (first impression fixed), availability (yesterday's memorable case), **premature closure** (too early a winner), confirmation bias (seeking agreement), framing (the referral letter's spin) ✓. Each bias has a structural antidote: "what else?" for closure, a required alternative for anchoring, blind first impressions for framing.

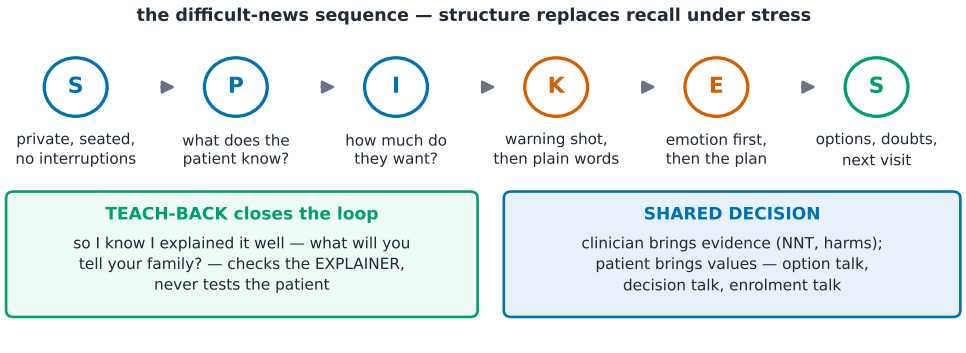
Common mistakes. Treating the working diagnosis as final (it is the current best, held until disconfirmed). Chasing rarity (the base rate always votes). Letting the first test decide (tests update odds, they rarely rewrite them alone – MED12's PPV arithmetic explains why).

Summary. Complaint → broad list → rank by probability × severity → tests as LR updates → working diagnosis + disposition; must-not-miss stays until excluded; biases: anchoring, availability, premature closure, confirmation, framing – each with its structural antidote.

Communicating with patients

Why it matters. Care travels through conversation: understanding, decisions, and bad news all have structures that keep the patient a partner rather than a target.

The idea, with a picture. The SPIKES sequence for difficult news: Setting (private, seated, uninterrupted) → Perception (what the patient already knows) → Invitation (how much they want to know) → Knowledge (a warning shot, then plain language, no jargon) → Empathy (respond to emotion) → Strategy and Summary ✓. **Teach-back** closes the loop: "so I know I explained it well – what will you tell your family?" – checking the doctor's explanation, never testing the patient.



Setting, Perception, Invitation, Knowledge, Empathy, Strategy – and teach-back to verify the loop

The derivation (why structure helps hardest conversations). Under stress both parties lose working memory; a sequence externalises the plan so no step depends on recall. Warning shots ("I'm afraid the news is not what we hoped") let the patient brace before the content; perception-first corrects the starting point; empathy-before-plan keeps the relationship while delivering facts ✓ – each step is cognitive first-aid.

Worked example 1 – the role-play script. (Fictional consultation) Setting: door closed, alarm off, sit level with the patient. Perception: "what is your understanding of why the tests were done?" Invitation: "do you want all the detail or the outline?" Knowledge: warning shot, two plain sentences, pause. Empathy: silence, then "this is not what you were hoping to hear." Strategy: options, uncertainty, next appointment ✓ – the script as role-play practice, labelled as such.

Worked example 2 – clinical link: shared decision-making. (Fictional) two options with different trade-offs, settled by **shared decision-making**: the clinician brings the evidence (MED12's NNT, harms), the patient brings values; the decision is built together – option talk, decision talk, enrolment talk ✓. Numbers are framed as natural frequencies (an NNT of 13 from MED12 becomes "about 8 of 100 people like you benefit") because patients reason in counts.

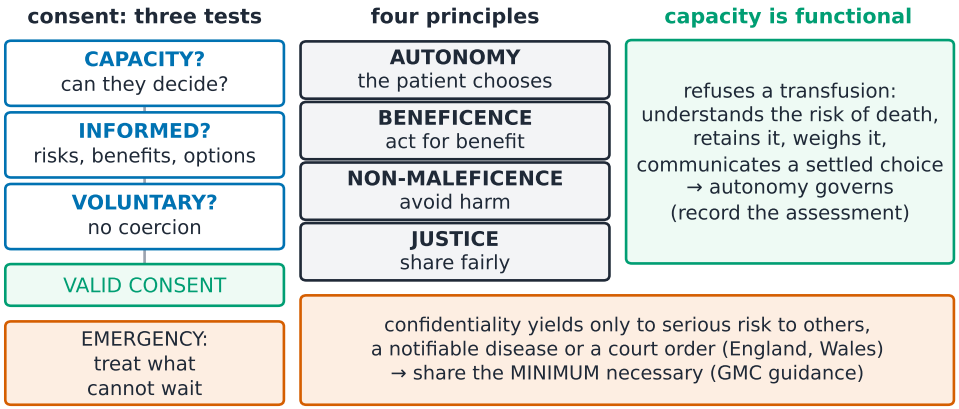
Common mistakes. Delivering facts before perception (correcting the wrong picture or none). Using teach-back as a quiz ("do you understand?" measures politeness, not understanding). Offering options without evidence, or evidence without options.

Summary. SPIKES: Setting, Perception, Invitation, Knowledge, Empathy, Strategy; teach-back checks the explainer; shared decisions = evidence + values, numbers as natural frequencies; communication practice is role-play, not assessed skill.

Medical ethics: autonomy, consent and confidentiality

Why it matters. Ethics is the framework that keeps clinical decisions the patient's own: four principles, a capacity test, and the narrow, law-governed exceptions to **confidentiality**.

The idea, with a picture. The four principles: **autonomy** (the patient's informed choice governs), **beneficence** (act for benefit), non-maleficence (avoid harm), justice (fair distribution) ✓. Consent is valid when the patient has capacity, is informed, and volunteers – and **capacity** is decision-specific: understand, retain, weigh, communicate ✓.



Capacity → information → voluntariness → valid consent – and the emergency exception when the test cannot be run

The derivation (why autonomy leads). The principles can conflict: a beneficent plan the patient refuses. Modern medical ethics resolves toward autonomy for competent adults – the patient's values are the arbiter of benefit – with beneficence shaping what is offered ✓. Capacity is not diagnosis-specific (a diagnosis does not remove it); it is functional and task-specific, tested per decision.

> **Worked example 1 – the capacity test in action.** (Fictional) a patient refuses a transfusion. Do they understand the information (the transfusion, the serious risk of death without it), retain it long enough to decide, weigh it against their own values, and communicate a settled choice? Yes to all four → capacity → autonomy governs, and the team documents the assessment ✓.

> **Worked example 2 – clinical link: when confidentiality yields.** (Declared jurisdiction: England and Wales) confidentiality may be breached without consent when there is a serious risk to others (e.g. a notifiable disease – Public Health (Control of Disease) Act 1984) or a legal requirement (a court order); the response must be the minimum necessary, and the GMC's guidance frames the judgment ✓. Where jurisdictions differ, the course names its frame rather than pretending one answer.

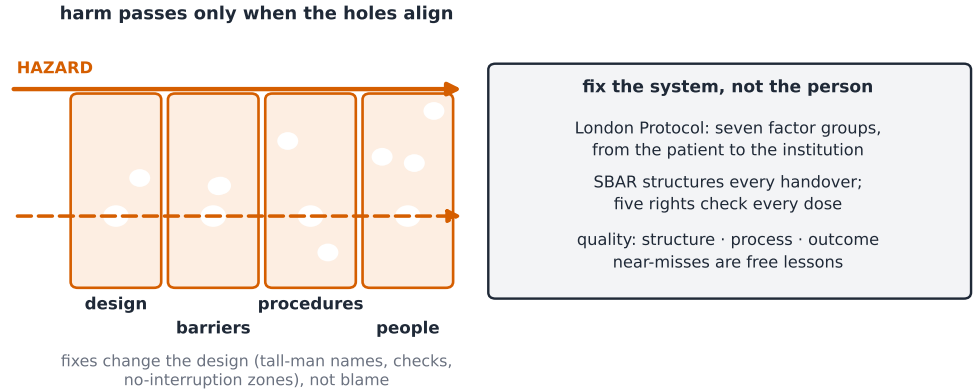
Common mistakes. Treating capacity as global (it is decision-by-decision; a patient may hold capacity for one decision and not another). Calling a consent form consent (the form records, the conversation consents). Breaching confidentiality on convenience – the threshold is serious harm and the route is the minimum necessary.

Summary. Autonomy, beneficence, non-maleficence, justice – autonomy leads for competent adults; consent = capacity + information + voluntariness; capacity: understand, retain, weigh, communicate, per decision; breach only for serious risk to others or law, minimum necessary (jurisdiction named; GMC framework cited).

Patient safety and quality

Why it matters. Errors are made by systems of people and tools; treating them as personal failings guarantees repetition. The systems view is the science of the **near-miss** .

The idea, with a picture. Reason's **Swiss cheese model** : hazards meet layers of defence (design, barriers, procedures, people), each with holes; harm passes when the holes align ✓. The response to an incident is systemic: what allowed the hole, not who stood in it.



Harm passes only when holes align – so incident analysis hunts the alignment, not the person

The derivation (why analysis must be structured). After an incident, memory reconstructs a tidy story; structured methods (the London Protocol's seven factor groups: patient, task and technology, individual, team, work environment, organisation and management, institutional context) force the search across every layer ✓. The same logic powers the tools: the **SBAR handover** SBAR (Situation, Background, Assessment, Recommendation) standardises what working memory drops under pressure – a checklist for conversations (MED11 §7's five rights for the bedside).

> **Worked example 1 – the tenfold near-miss.** (Fictional, from MED11 §7) a tenfold dose reached the bedside and was caught: the London Protocol reads it as prescribing (similar names), dispensing (no independent check), administration (interruption at the bedside) – three holes that aligned ✓; the fixes are design changes (tall-man lettering, separate stocks, do-not-interrupt zones), not blame.

Worked example 2 – clinical link: measuring quality. (Structure– process–outcome) a service can be judged by its structure (staff, equipment), its process (were the right steps taken – checklists, handover compliance) and its outcome (results, risk-adjusted) ✓. Process measures respond fastest to improvement; outcome measures convince the sceptic – quality work reads both.

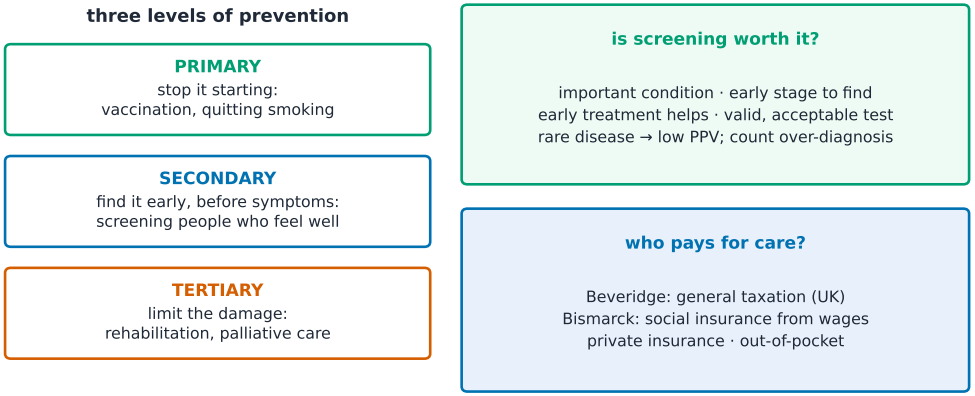
Common mistakes. Stopping at "human error" (the beginning of the analysis, not the end). Testing defences only when they fail (drills and near-miss reporting keep holes visible). Equating a checklist with safety (it is one layer; the culture that speaks up is another).

Summary. Swiss cheese: harm = aligned holes across layers; analyse with the London Protocol's seven factor groups; SBAR standardises handovers; five rights standardise the bedside; quality = structure + process + outcome, and near-misses are free lessons.

Public health and health systems

Why it matters. Medicine works one patient at a time; health works populations at a time – prevention levels, screening arithmetic and financing models decide what care exists to be delivered.

The idea, with a picture. Prevention levels: primary (prevent disease: vaccination – MED07), secondary (detect early: **screening**), tertiary (limit damage: rehabilitation) ✓. Screening is judged by explicit criteria: important condition, treatable early stage, valid test, acceptable, and the programme's arithmetic (MED12's PPV – the pre-clinical pool governs the yield).



Primary stops the disease, secondary catches it early, tertiary limits its damage – screening arithmetic gates the middle level

The derivation (why screening arithmetic decides). A screen tests the well: its yield is PPV, and PPV lives on prevalence (MED12 §4) – so the pre-clinical detectable pool (a screening interval short enough that the pool is real) is the design variable ✓.

Over-diagnosis (finding disease that would never have mattered) is the hidden cost the criteria guard against.

Worked example 1 – the levels in one disease. (Fictional teaching example) a smoking-related cancer: primary – cessation programmes; secondary – targeted screening of the high-risk (pool arithmetic); tertiary – rehabilitation and palliation ✓. Each level has its own denominator and its own success measure.

Worked example 2 – clinical link: financing shapes access. (Named models) general taxation (the Beveridge model – a national health service), social insurance (the Bismarck model – funds from wages), private insurance, out-of-pocket ✓. The models set who is covered, who pays at the point of care, and what gatekeeping exists – the macro-decisions that every consultation inherits.

Common mistakes. Calling any early test screening (screening tests the WELL; testing the symptomatic is diagnosis). Judging screening by cases found (count over-diagnosed harm too: the criteria demand both). Comparing systems by spending alone (coverage and outcomes differ at equal spend).

Summary. Primary/secondary/tertiary; screening criteria: important, treatable-early, valid, acceptable + pool arithmetic (PPV) with over-diagnosis counted; financing: Beveridge (tax), Bismarck (social insurance), private, out-of-pocket – each shaping coverage and access.