

Research Directions to Cure All Diseases

Medicine has four ways to end a disease, not thousands. Three of them are crowded and the fourth is almost unworked, and the fourth one holds sepsis, ARDS, severe influenza, severe COVID-19, the autoimmune diseases, Ebola and, on the evidence set out below, atherosclerotic cardiovascular disease, which kills more people than anything else does. This paper sets out where the four stand in 2026, which is which, and where the next decade of work should go.

J. Konstapel, Constable Research, Leiden — Leiden,29 October 2026

© 2026 J. Konstapel, Constable Research, Leiden. All rights reserved. Quotation with attribution is permitted; reproduction of the whole requires written permission.

Four operations, not thousands of diseases

Disease catalogues are organised by organ, by pathogen, by mechanism. None of those divisions predicts whether anything can be done. A classification by operation does.

There are exactly four ways a disease state ends. The state resolves on its own. One of the two parties is removed. Something missing is supplied. Or both parties remain, the relation between them does not settle, and what must change is the proportion between them.

Operation	Ends the disease by	Exemplar	Status in 2026
1. Let it close	Doing nothing, and protecting the process	Mild influenza, viral pneumonia, most wounds	Understood; routinely violated
2. Remove one party	Killing or clearing the other participant	Antibiotics, antivirals, surgical excision	Mature; hitting a resistance ceiling
3. Supply what is missing	Replacing an absent component	Insulin, vitamin C, gene editing	Technically solved; economically blocked

Operation	Ends the disease by	Exemplar	Status in 2026
4. Restore proportion	Changing the relation, not the parts	Sepsis, ARDS, severe influenza, autoimmunity, Ebola, atherosclerosis	Almost unworked; largest remaining burden

The asymmetry is historical, not biological. Operation 2 produced the most spectacular results in the history of medicine, and the field reorganised around it: find the agent, kill the agent. Operation 3 followed, and modern genetic medicine has made it technically routine. Operation 4 never acquired a name, a standard trial design, or a regulatory pathway, and it is where the largest remaining burden of death sits.

Three consequences follow immediately, and the rest of this paper develops them.

First, the operation determines the endpoint. A trial testing an operation-2 agent asks whether the agent was cleared. A trial testing operation 4 cannot ask that, because nothing is being cleared, and measuring clearance there produces null results on interventions that work.

Second, the operation determines who to enrol. Operations 1 to 3 enrol by diagnosis. Operation 4 cannot, because the same diagnosis contains patients on both sides of the boundary at which the intervention reverses sign. The evidence for this is not theoretical and is given below.

Third, applying the wrong operation is not merely ineffective. It does harm, and the harm is quantified in several of the trials cited here.

Operation 1: let it close

The first operation is the one no industry is built on. A reference catalogue that files conditions by what can be done with them marks these SELF-LIMITING: the body completes the process itself, and intervention adds nothing. Viral pneumonia and mild influenza are filed that way.

The research direction here is not a treatment. It is knowing where the boundary lies, because intervention on the wrong side of it causes measurable harm.

RECOVERY provides the number. Among hospitalised COVID-19 patients who needed no oxygen, 28-day mortality was 17.0% with dexamethasone against 13.2% with usual care, a rate ratio of 1.22 (95% CI 0.93–1.61). The confidence interval includes 1, so harm is a

direction rather than an established finding, but the point estimate is 3.8 percentage points in the wrong direction, in the group whose process was still completing.

Antibiotics in viral infection are the same error at scale, and the bill for it arrives in the next section.

Direction 1. Build the boundary test, disease by disease. The clinical question is not “does this drug work” but “has this patient’s process stopped completing”. That question has an answer at the bedside today in exactly one condition — COVID-19, where the answer is whether the patient needs oxygen — and it was found by accident, in a subgroup analysis. For every other condition in this category the boundary is unmeasured.

What would count as success: a measurement, available at admission, that separates patients whose process will complete from those whose will not, validated prospectively against outcome without treatment.

Operation 2: remove one party

This is the operation that built modern medicine, and it still works beautifully where it works at all. Tuberculosis is curable on a four-drug regimen. Malaria is curable on artemisinin combination therapy. Bacterial pneumonia is curable on amoxicillin. Hepatitis C is cleared in eight weeks by glecaprevir/pibrentasvir.

Two facts about this operation should govern how much further investment it receives.

It is running into a wall. Deaths directly attributable to bacterial antimicrobial resistance were 1.14 million in 2021 and are forecast at 1.91 million in 2050, a rise of 67.5%. Associated deaths rise from 4.71 million to 8.22 million, a rise of 74.5%. Cumulatively, 39 million deaths are forecast between 2025 and 2050. Deaths in adults over 70 rose by more than 80% between 1990 and 2021.

The wall is structural. The operation works by removing a self-replicating party, and a self-replicating party selects for escape. Every success narrows the space of future successes. This is the only one of the four operations with that property.

Where it works, the binding constraint has stopped being the drug. Hepatitis C is the cleanest case. About 47 million people live with chronic infection; roughly 68 million were ever infected between 2015 and 2024. Of those, about 24 million knew their diagnosis, which is 35%, and about 13 million had received direct-acting antivirals by the end of 2024 — 54% of the diagnosed and 19% of the ever-infected. Around 44 million people were never diagnosed at all. The WHO targets for 2030 are a 90% reduction in new infections and a 65% reduction in mortality against 2015.

A curative eight-week regimen exists and four out of five people who need it have not had it. No further pharmacology closes that gap.

Direction 2. Treat delivery as the research problem, not the molecule. For the diseases where operation 2 already succeeds, the remaining mortality is a diagnostic and logistics problem, and it is larger than anything a new compound would add.

Direction 3. Stop paying for operation 2 where operation 1 applies. Antibiotic use in self-limiting viral illness is the main driver of the resistance curve above. The boundary test of Direction 1 is also the main available lever on AMR, and it is cheaper than any new antibiotic class.

Direction 4. For resistant organisms, change operation rather than compound. Phage, anti-virulence and host-directed approaches all amount to abandoning removal. The strategic case for them is not that they are promising but that removal has a mathematical ceiling and the others do not.

The same ceiling, inside one patient

The resistance curve above is a population measured over decades. Oncology shows the identical structure inside one person, over months.

Osimertinib is the best first-line removal agent in its class for EGFR-mutated non-small-cell lung cancer. In FLAURA, median progression-free survival was 18.9 months against 10.2 months on the earlier agents, a hazard ratio of 0.46, which is close to a doubling. Median overall survival was 38.6 months (95% CI 34.5 to 41.8) against 31.8 months (26.6 to 36.0), a hazard ratio of 0.80 at $P = 0.046$, which is 6.8 months. At three years, 28% of patients were still on the regimen against 9% on the comparator, which is to say that 72% were not.

The distance between those two figures is the whole argument. The time before return nearly doubles; survival gains a fifth. What is removed is removed, and what remains is what the removal did not reach, and what remains is what continues. This is the antimicrobial resistance curve with the population replaced by one body and the decades replaced by eighteen months.

Cancer is therefore not a second instance of operation 2 waiting for a better compound. Nothing foreign is present and nothing is missing: what does not settle is division itself. That places it structurally in the fourth row, and it is the one fourth-row condition for which no measured boundary marker of the kind described below has yet been established. Naming that marker is the open problem, not finding the next inhibitor.

Direction 4 therefore reads across from infection to oncology unchanged. The strategic case for abandoning removal does not rest on the alternatives being promising. It rests on

removal having a ceiling that is structural rather than technical, in a population and in a body alike.

Operation 3: supply what is missing

Where a component is absent, supplying it ends the disease. Insulin and vitamin C are the classical cases. The modern form is genetic, and it now works.

Exagamglogene autotemcel edits BCL11A to restore fetal haemoglobin. Across more than 100 treated patients — 46 with sickle cell disease and 56 with transfusion-dependent beta-thalassaemia — 92.3% of evaluable sickle cell patients had no vaso-occlusive crisis for at least 12 consecutive months, with a longest crisis-free period of 54.8 months and a mean of 27.9 months. In thalassaemia, 94.2% were transfusion-independent for at least 12 consecutive months, longest 59.4 months, mean 31 months. Edited allele levels stayed stable, indicating editing in long-term haematopoietic stem cells.

That is a single administration converting a lifelong transfusion-dependent disease into its absence, sustained for approaching five years.

The constraint on this operation is not biology. It is that each therapy is built once, for one mutation, at a cost that only large patient populations can amortise, which excludes most of the conditions the operation would solve. The number of single-gene disorders is in the thousands; the number with an approved genetic therapy is in the tens.

Direction 5. Make the platform the product, not the therapy. The research target is a delivery and editing platform where changing the target sequence does not restart development, manufacturing or regulatory review. Every proposal in this field should be judged by what it does to the cost of the *second* indication, not the first.

Two adjacent developments bear on this. In vivo generation of engineered cells, using targeted lipid nanoparticles rather than cells manufactured outside the body, removes the per-patient manufacturing step that dominates current cost. And computational protein design has reached the point where binders and antibodies are generated against specified targets with atomic accuracy, which shortens the front end of the same pipeline.

Direction 6. Audit the boundary between operation 3 and operation 4. Not every deficiency is a deficiency. Where supplying the missing component fails to end the disease, the classification was wrong and the condition belongs in the fourth row. That audit is cheap, it uses trial results that already exist, and it would redirect programmes that are currently failing for a reason nobody has named.

Operation 4: restore proportion

Here both parties remain. Nothing is cleared and nothing is missing. What is wrong is the relation between the parts, and the catalogue marks these RELIEF ONLY: nothing removes it. Severe influenza is filed that way. So, structurally, are sepsis, ARDS, severe COVID-19, the autoimmune diseases and Ebola virus disease, for which there is no licensed vaccine or therapeutic against the Bundibugyo species behind the 2026 outbreak in the Democratic Republic of the Congo.

The first thing to correct is an assumption carried over from operation 2: that more response is better. It is not.

In a Sierra Leone cohort of 112 Ebola patients, 220 of 246 differentially abundant transcripts were *higher* in those who died, with stronger interferon signalling and stronger acute-phase responses. Fibrinogen alpha, beta and gamma were approximately 54-fold, 28-fold and 19-fold higher in fatal cases. Meanwhile viral load did not separate outcome where it mattered: cycle threshold predicted at the extremes, with Ct 12 suggesting death and Ct 30 survival, but around the cohort mean of 21.4 outcome was close to even, while host-gene classifiers reached 79–92% accuracy in that same range.

In this row, the size of the response is not a measure of defence. It is a measure of what is not settling.

Why trials in this row keep returning null

Because the intervention reverses sign inside the enrolled population, and a single pooled estimate averages the two halves.

RECOVERY shows the reversal inside one randomisation:

Respiratory support at entry	Dexamethasone	Usual care	Difference	Rate ratio
None	17.0%	13.2%	+3.8 pp	1.22 (0.93–1.61)
Oxygen only	21.5%	25.0%	–3.5 pp	0.80 (0.70–0.92)
Invasive ventilation	29.0%	40.7%	–11.7 pp	0.65 (0.51–0.82)

Mortality under usual care rises monotonically across the strata — 13.2, 25.0, 40.7 — and the rate ratio falls monotonically — 1.22, 0.80, 0.65. The same drug at the same dose, in one trial, changes sign.

HARP-2 shows what happens when a trial does not stratify. It randomised 540 ARDS patients to simvastatin or placebo and found no overall survival difference. A secondary analysis applied latent class analysis to baseline data alone, without using outcomes: a two-class model fitted better than one ($P < 0.0001$), splitting 353 patients (65%) hypoinflammatory and 186 (35%) hyperinflammatory. Baseline 28-day mortality was 39% against 17% between the classes. Within the hyperinflammatory class, simvastatin significantly improved 28-day survival ($P = 0.008$), with survival differing across the four treatment-by-class groups at $P < 0.0001$.

A trial filed as negative contained a benefit. The stratifier was computable from measurements taken before treatment.

SAVE-MORE shows the converse. Anakinra was given to hospitalised COVID-19 patients selected by plasma suPAR of 6 ng/mL or more, excluding those with $\text{PaO}_2:\text{FiO}_2$ below 150 and those already ventilated. The primary ordinal outcome improved with an odds ratio of 0.36 (95% CI 0.26–0.49; $P < 0.001$) and 28-day mortality was 3.2% against 6.9%.

That trial treated patients who were not yet severely hypoxic, and it worked, because entry was decided by a measured marker rather than by clinical stage. The boundary is therefore not lateness and not crude severity. It is a state that can be measured, and a patient can be past it while still looking mild.

Direction 7. Stratify at randomisation on a measured marker, and report the strata separately. Every trial in this row should be designed this way. A single pooled estimate is not a conservative choice here; it is a design that destroys the signal.

Operation 4 outside the intensive care unit

Everything above comes from critical care, where the fourth row is easiest to see because death arrives within days. If the row were confined there it would be a small category. It is not.

CANTOS gave canakinumab, an antibody to interleukin-1-beta, to 10,061 people who had had a myocardial infarction within thirty days and whose baseline hsCRP was 2 mg/L or above, median 4.20. Entry was therefore already decided by a measured marker, as in SAVE-MORE, and not by diagnosis alone. Three results matter here.

The lipid route was not touched. Median LDL cholesterol at 48 months moved by +1.7% on the 150 mg dose ($P = 0.004$) against -1.4% on placebo. LDL went slightly up and events

went down. Whatever changed, it was not the quantity at which fifty years of operation-2 and operation-3 thinking has been aimed.

Pooled, the effect is real and small. Primary events, meaning myocardial infarction, stroke or cardiovascular death, ran at 3.86 per 100 person-years on 150 mg against 4.50 on placebo: hazard ratio 0.85 (0.74 to 0.98). Recomputed crudely from those rates the ratio is 0.858, a reduction of 14.2%.

The statistical space was divided across the wrong thing. The trial split its alpha across three doses of one drug. The 150 mg dose passed its prespecified threshold of 0.02115 with $P = 0.02075$, a margin of 0.0004. The 300 mg dose reached a hazard ratio of 0.86 (0.75 to 0.99) and was declared non-significant, because its threshold was 0.01058 and its P value 0.0314, nearly three times that. Two doses of one antibody produced the same effect; one counted and one did not.

And then the division that matters: on-treatment hsCRP at three months, below or above 2 mg/L.

Endpoint	hsCRP below 2 mg/L	hsCRP 2 mg/L or above
MI, stroke or cardiovascular death	0.75 (0.66 to 0.85)	0.90 (0.79 to 1.02)
Cardiovascular death	0.69 (0.56 to 0.85)	0.99 (0.82 to 1.21)
Death from any cause	0.69 (0.58 to 0.81)	1.05 (0.90 to 1.22)
Number needed to treat over 5 years	16	57

Death from any cause falls by nearly a third on one side of the line and crosses 1 on the other. The number needed to treat differs 3.6-fold. The share of patients reaching the lower state rose with dose: 44%, 55% and 65%.

This division is weaker evidence than HARP-2 or SAVE-MORE, and the difference should be stated plainly rather than left for a reader to find. HARP-2 classified patients on baseline data alone, and SAVE-MORE decided entry on a marker before randomisation. CANTOS divided on a measurement taken three months after treatment began, so the two groups are defined partly by how they responded, and a group that responds may simply be a group with more room to improve. What CANTOS establishes is therefore not that this marker identifies a treatable stratum, but that a pooled estimate can conceal a division large enough to reverse the sign of mortality. The trial that would settle it is the one the structure

points at: enter on the baseline marker, dose, measure at three months, and randomise again.

The harms are in the same record. Fatal infection or sepsis ran at 0.31 per 100 person-years against 0.18 ($P = 0.02$), an excess of 0.13. Against that, fatal cancer ran at 0.45 against 0.64 ($P = 0.02$), a reduction of 0.19, and the primary endpoint fell by 0.64. That arithmetic is not a defence of the drug, which was never licensed for this indication. It is a description of the fourth row: nothing is removed, and the cost appears exactly where the capacity to settle an infection is reduced.

The consequence for the classification is the one that matters most in this paper. Atherosclerotic cardiovascular disease is the largest single cause of death in the world. If it belongs in the fourth row, and the LDL result above is the strongest available argument that it does, then the fourth row is not a residual category holding rare and desperate conditions. It is the largest of the four, and it is the one with no standard trial design, no accepted biomarker class and no regulatory precedent.

Direction 10. Divide the statistical space across strata, not across doses. CANTOS spent its alpha on three doses of one antibody and had none left for the division that reversed the sign of mortality. This is the cheapest correction in the paper and it is purely procedural: where a fourth-row trial must choose, prespecify fewer doses and more strata. The dose can be found afterwards; the stratum cannot be recovered once the result has been filed as negative.

The first clean demonstration that operation 4 is real

Autoimmune disease is the purest case in the fourth row. Nothing is missing. Nothing foreign is present. The relation is wrong, and treatment has consisted of holding it down indefinitely.

Eight patients with treatment-resistant systemic lupus erythematosus received a single dose of CD19-directed CAR T cells between March 2021 and June 2023. All eight entered remission with SLEDAI scores of zero, and all came off immunosuppressants including glucocorticoids. Autoantibodies disappeared, with one antibody in one patient as the exception, and stayed negative through last follow-up, 12 to 24 months after treatment. Naive B cells returned within months without the disease returning, and patients retained vaccine responses to measles, mumps, rubella, varicella zoster, Epstein-Barr, tetanus and pneumococcus.

The B cell compartment was not removed permanently. It was depleted and it came back, and what came back did not reproduce the disease. The relation reset.

This is the operation the fourth row needs, demonstrated once, in eight people. The caveats are real: a small and ethnically narrow cohort, cytokine release syndrome and neurotoxicity as risks, no approval for this indication in the United States, and a cost of roughly \$375,000 to \$425,000 per infusion before associated costs.

But the result is categorical rather than incremental. The patients are not better managed; they are off treatment with no disease activity. Nothing in the first three operations predicts that outcome for lupus.

Direction 8. Pursue reset rather than suppression across the fourth row. The question to carry from lupus into sepsis, ARDS and severe viral disease is not which mediator to block, but which compartment can be made to re-form in a settled state. Depletion followed by reconstitution is one mechanism of that; it is unlikely to be the only one.

The cost structure is also the point at which the directions converge. In vivo generation of engineered cells, by targeted nanoparticle delivery rather than external manufacture, is the step that would move a reset from an eight-patient demonstration to a usable intervention. That is the same platform problem as Direction 5, and solving it once serves both rows.

What it costs to pick the wrong operation

Alzheimer's disease has absorbed more investment than any other target in the fourth row, under an operation-2 programme: identify the agent, remove the agent. Amyloid was the agent. The removal works.

The clinical result does not follow. In the lecanemab phase 3 trial over 18 months, CDR-SB declined 1.21 points on treatment against 1.66 on placebo from a baseline of 3.2 — a difference of 0.45 points, reported as 27% slowing. For donanemab in the low-to-medium tau group over 76 weeks, iADRS declined 10.2 against 13.1 from a baseline of 104, a difference of 3.4 points, reported as 33%.

Set those against the minimal clinically important differences: 1.0 on CDR-SB in mild cognitive impairment and 1.6 in mild dementia; 5 and 9 respectively on iADRS. The lecanemab effect is 0.45 of the smaller threshold. One analysis concludes that changes of about 0.5 on CDR-SB and about 3 on iADRS are imperceptible to patients, families and physicians.

The harms are not imperceptible. Nearly one in four lecanemab patients developed brain swelling or bleeding, and more than one in three on donanemab. Treatment-related adverse events ran at 45% and 89%. Severe bleeding occurred in 5 patients against 1 on placebo for lecanemab, and 7 against 2 for donanemab. The number needed to harm is about 3 for both; the number needed to treat could not be computed from the published data.

This is what applying operation 2 to a fourth-row condition produces: the agent is successfully removed, the relation is unchanged, the benefit sits below the threshold of perception, and the harm does not.

The correct reading is not that amyloid is irrelevant. It is that removal was the wrong operation, for the same reason that killing the virus is the wrong operation in severe influenza. The research question becomes which relation in the brain has stopped settling, and what would make it settle — a fourth-row question, which has never been the one the field was funded to ask.

Direction 9. Re-classify the major chronic diseases by operation before funding another removal programme. Neurodegeneration, metabolic disease, chronic kidney disease, heart failure and the major psychiatric disorders have all been pursued predominantly as operation-2 or operation-3 problems. Each should be re-examined against the four rows, and the trial results that already exist are sufficient to do it.

The directions, collected

#	Direction	What is built	Judged by
1	Build the boundary test, disease by disease	An admission measurement separating processes that will settle from those that will not	Prospective separation of untreated outcomes
2	Treat delivery as the research problem where a cure exists	Diagnosis and linkage at population scale	Share of the ever-infected treated, not SVR rate
3	Stop applying removal where nothing needs removing	Point-of-care discrimination of self-limiting illness	Reduction in inappropriate antimicrobial courses
4	For resistant organisms and returning tumours, change operation, not compound	Phage, anti-virulence, host-directed; in oncology, a settling target rather than a further inhibitor	Durability under selection, not initial potency
5	Make the editing platform the product	Delivery and editing where retargeting does not restart development	Cost of the second indication
6	Audit the operation-3 / operation-4 boundary	Re-classification of failed replacement programmes	Programmes correctly redirected
7	Stratify at randomisation on a measured marker	Trial designs reporting strata separately	Sign reversal detected rather than averaged away

#	Direction	What is built	Judged by
8	Pursue reset rather than suppression	Interventions that let a compartment re-form settled	Drug-free remission, not disease control
9	Re-classify the major chronic diseases by operation	A published operation assignment per disease	Removal programmes stopped where the row is wrong
10	Divide the statistical space across strata, not across doses	Protocols prespecifying fewer doses and more strata	Effects retained instead of filed as negative

Three of these — 2, 3 and 6 — require no new science at all. They are re-allocations, and they act on larger numbers than anything in the pipeline.

Two — 5 and 8 — converge on one engineering problem: generating engineered cells inside the body rather than manufacturing them outside it. That single capability would move both the supply row and the proportion row.

And two, 7 and 10, are free. Both are changes in how a trial is written down. They cost nothing beyond the discipline of fixing strata before allocation and of spending the statistical space on patients rather than on doses. HARP-2 shows that the first recovers effects currently being discarded, and CANTOS shows what the second costs when it is not done.

What follows for each reader

For clinicians. Before reaching for an intervention, name the operation. Where the process is still settling, suppression is not neutral: in RECOVERY's no-oxygen group the point estimate ran 3.8 percentage points against the patient. Where an operation-2 cure exists, the question that decides mortality is whether the patient was ever diagnosed, not which regimen to choose.

For trialists. A null result in the fourth row is not evidence of no effect; it is evidence of an unstratified population, until a stratified analysis says otherwise. Fix the strata before allocation, choose the stratifier as a measured marker rather than a clinical stage, and publish the strata separately. The archive of null immunomodulation trials is a research resource, not a closed file.

For funders. Three of the ten directions require no new science and act on the largest numbers: closing the hepatitis C diagnostic gap, stopping removal where nothing needs removing, and re-classifying failed replacement programmes. Two more, 7 and 10, cost nothing but a change in how protocols are written. These compete poorly for funding because they produce no molecule, and they are where the available mortality reduction is concentrated.

For industry. The fourth row is the largest unserved category in medicine, and on the evidence above it contains the largest cause of death in the world. It is unserved because it has no standard trial design, no accepted biomarker class and no regulatory precedent for a stratified claim. The company that establishes that pathway, meaning a marker, a stratified design and a label that names a stratum rather than a diagnosis, defines the category. Two of the three crowded rows are closing: operation 2 is a mature market with a resistance ceiling that holds in a population and inside a single patient alike, and operation 3 is technically solved and economically blocked, with its unlock being platform cost rather than new biology. The commercial argument for the fourth row is not that it is virtuous. It is that it is where the volume is and where no one has a design.

One caution for all four. This paper’s framing does not predict that any particular intervention will work. It predicts where to look, which populations to enrol, and which negative results deserve re-examination. Every claim it makes is refutable by the trials it proposes.

Verification and limitations

Every figure was recomputed from the reported counts or percentages before being quoted.

Check	Recomputed	As reported
AMR attributable deaths, 2021 to 2050	1.14 to 1.91 m, +67.5%	+67.5%
AMR associated deaths, 2021 to 2050	4.71 to 8.22 m, +74.5%	+74.5%
Hepatitis C diagnosed share	24 / 68 = 35%	36%
Hepatitis C treated, of ever infected	13 / 68 = 19%	not stated
Casgevy patients	46 + 56 = 102	more than 100

Check	Recomputed	As reported
HARP-2 class sizes	$353 + 186 = 539$	539 analysed
Lecanemab CDR-SB difference	$1.66 \text{ minus } 1.21 = 0.45;$ $0.45 / 1.66 = 27\%$	0.45; 27%
Lecanemab effect against MCID	$0.45 / 1.0 = 0.45$	MCID 1.0 at MCI
SAVE-MORE relative mortality reduction	$(6.9 \text{ minus } 3.2) / 6.9 = 54\%$	55%
RECOVERY control mortality across strata	$13.2 < 25.0 < 40.7$	strictly increasing
RECOVERY rate ratios across strata	$1.22 > 0.80 > 0.65$	strictly decreasing
CANTOS 150 mg crude rate ratio	$3.86 / 4.50 = 0.858$, a 14.2% reduction	HR 0.85 (0.74 to 0.98)
CANTOS 50 mg and 300 mg crude ratios	$4.11 / 4.50 = 0.913;$ $3.90 / 4.50 = 0.867$	HR 0.93; HR 0.86
CANTOS 150 mg significance margin	$0.02115 \text{ minus } 0.02075 = 0.00040$	threshold met
CANTOS 300 mg against its threshold	$0.0314 / 0.01058 = 2.97\text{-fold over}$	not significant
CANTOS NNT ratio across the marker	$57 / 16 = 3.6\text{-fold}$	NNT 16 and 57
CANTOS net events per 100 person-years	$0.64 + 0.19 \text{ minus } 0.13 = 0.70$	not stated
FLAURA progression-free survival	$18.9 / 10.2 = 1.85\text{-fold,} +8.7 \text{ months}$	HR 0.46 (0.37 to 0.57)
FLAURA overall survival	$38.6 / 31.8 = 1.21\text{-fold,} +6.8 \text{ months}$	HR 0.80, $P = 0.046$
FLAURA patients off the regimen by 3 years	$100 \text{ minus } 28 = 72\%$	28% still on regimen

Two discrepancies are left standing rather than smoothed. On the stated hepatitis C numbers, 13 million treated of 24 million diagnosed computes to 54%, while the source describes about 20% of diagnosed people as having received direct-acting antivirals; the two statements cannot be reconciled from the published figures, so both are recorded here and the paper quotes only the shares of the ever-infected. And the diagnosed share computes to 35% against a stated 36%, which is rounding.

Crude ratios computed from the RECOVERY percentages sit further from 1 than the published rate ratios, because the trial reports age-adjusted values. The published values are the ones quoted in the text.

Limitations. The CAR-T lupus result rests on eight patients with 12 to 24 months of follow-up in a single centre, and is not an approved indication. The HARP-2 subphenotype analysis is secondary and its published abstract gives no per-arm percentages within each class, so none are quoted. The RECOVERY no-oxygen stratum has a confidence interval including 1, so harm there is a direction rather than an established finding. The CANTOS division is on a measurement taken three months after treatment began, which is a weaker design than HARP-2 or SAVE-MORE, and the two groups it produces are defined partly by how they responded; the section says so where the figures appear. CANTOS also carried an excess of fatal infection and sepsis, 0.31 against 0.18 per 100 person-years, and canakinumab was never licensed for this indication. The FLAURA comparison is against first-generation inhibitors rather than against no removal at all, and patients in the comparator arm could cross over, so the overall survival difference understates the agent and says nothing about whether removal was the right operation. The classification of conditions by operation was checked in detail against one chapter of one reference catalogue, not exhaustively, and the placement of cancer in the fourth row is an argument from the shape of its failures, not a demonstration: no boundary marker for it is proposed here. And the four-operation scheme itself is a way of organising existing results. Being exhaustive by construction, it cannot be refuted as a scheme; it derives nothing and predicts no effect size. What is refutable is every row assignment in it, and that is where it should be attacked.

Sources

Stratified and biomarker-guided trials

- [RECOVERY — dexamethasone, subgroup outcomes \(The Bottom Line\) and the same trial summarised independently \(First10EM\).](#)
- [Calfee et al. — ARDS subphenotypes and differential response to simvastatin, HARP-2 secondary analysis.](#)
- [SAVE-MORE — early anakinra guided by suPAR, phase 3.](#)

- [RECOVERY — tocilizumab \(REBEL EM\)](#) and [the European Respiratory Society summary](#).
- [CAPE COD, NEJM 2023](#), [Critical Care Reviews summary](#) and [CAPE COD against ESCAPe](#).

Ebola virus disease

- [Liu et al., Genome Biology 2017](#) — transcriptomes, fibrinogen fold-changes, cycle-threshold limits, classifier accuracy.
- [ECDC — Ebola disease: multi-country outbreak 2026](#) and [LSHTM on the absence of a licensed vaccine for Bundibugyo virus](#).

Cell therapy and autoimmunity

- [American College of Rheumatology — CD19 CAR-T in refractory lupus: long-term remission with retained vaccine responses](#).

Gene editing

- [Vertex long-term exa-cel data in sickle cell disease and beta-thalassaemia](#).

Antimicrobial resistance

- [GRAM Project forecasts, University of Oxford](#).

Hepatitis C

- [Working towards hepatitis C elimination \(Contagion Live\)](#) — burden, cascade of care and WHO 2030 targets.

Alzheimer's disease

- [Lecanemab and donanemab: an illustrated perspective on the data, eNeuro 2024](#).

Design and delivery platforms

- [De novo design of protein structure and function with RFDiffusion](#) and [atomically accurate de novo design of antibodies](#).
- [In vivo CAR T cell generation to treat cancer and autoimmune disease](#).

Atherosclerotic cardiovascular disease

- [CANTOS, trial summary with entry criteria, per-dose hazard ratios, prespecified thresholds, LDL change and harms \(WikiJournalClub\)](#) and [the same trial summarised independently \(American College of Cardiology\)](#).

- [The CANTOS on-treatment hsCRP analysis, with hazard ratios by stratum \(Residual Risk Reduction Initiative\)](#) and [the numbers needed to treat by stratum \(TCTMD\)](#).

Oncology

- [FLAURA, first-line osimertinib: progression-free survival and hazard ratios \(Cancer Management and Research\)](#) and [the overall survival analysis with confidence intervals and three-year regimen retention \(Journal of Oncology Navigation and Survivorship\)](#).

Framework

- J. Konstapel, *The Vacuum.Net Theory*, Ninth Edition, Constable Research, Leiden, 27 October 2026 — *Foundations* 9.2 and *Physics* 9.0.
- J. Konstapel, *Where Illness Comes From, Completed*, Constable Research, 20 October 2026 — the cure statuses used to classify conditions.
- J. Konstapel, *Depth of Discharge*, Constable Research, 29 October 2026 — the stratification argument set out in full.