Relative Risk, Absolute Risk, and the Arithmetic of Persuasion

Two interventions, one statistical lesson: atorvastatin (Lipitor) and the COVID-19 vaccines

Contents

Executive summary

The same clinical trial can be summarized as “cuts your risk by 95%” or “reduces your risk by less than one percentage point,” and both can be arithmetically correct. The gap between those two sentences is the gap between relative and absolute risk reduction. It is the single most exploited ambiguity in the public communication of medicine, and it sits at the center of arguments about both cholesterol-lowering drugs and COVID-19 vaccines.

This paper uses atorvastatin (marketed as Lipitor) and the COVID-19 vaccines as parallel worked examples. They are not clinically comparable — different diseases, outcomes, and time structures — but they are statistically comparable, and the comparison is instructive precisely because the same trick of framing has been used to oversell one and to undersell the other. The core results:

- **Relative risk reduction (RRR) is approximately invariant to baseline risk.** Atorvastatin lowers major vascular events by roughly a fifth per unit of cholesterol lowering across almost every population studied; the COVID vaccines lowered symptomatic infection by roughly 90–95% (mRNA) regardless of where the trial ran. RRR is a property of the intervention acting on those at risk.
- **Absolute risk reduction (ARR) is the product of that RRR and the baseline event rate, so it inherits the baseline risk and the time horizon.** The same atorvastatin yields a ~0% absolute benefit in someone with no coronary calcium and a 6–8% absolute benefit in someone with a high calcium score. The same Pfizer vaccine yields a 0.84% ARR over a two-month trial window and a far larger cumulative benefit over a full season of near-universal exposure.
- **Number needed to treat (NNT), being 1/ARR, is therefore not a property of the drug at all — it is a property of the population and the clock.** Quoting a trial-window NNT as though it described lifetime benefit is the central error in the “the vaccines are only 0.84% effective” claim; quoting a high-baseline-risk ARR as though it applied to a low-risk person is the mirror-image error that oversells statins to the healthy.

The honest practice is to report both measures together, always attached to a stated baseline risk and time horizon. Neither number is “the truthful one”; each answers a different question.

> **Scope note.** This is a methodological and educational paper about how risk statistics are constructed and communicated. It is not individualized medical advice. Decisions about either medication should be made with a clinician using a person’s own risk profile.

1. Definitions: four numbers from one 2×2 table

Every two-arm trial of a binary outcome reduces to four counts: events and non-events in the treated arm and in the control arm. From the two event rates — the **control event rate (CER)** and the **experimental event rate (EER)** — four summary statistics follow.

Quantity	Definition	What it measures
Relative risk (RR)	EER / CER	How the treated risk compares to the untreated risk, as a ratio

Quantity	Definition	What it measures
Relative risk reduction (RRR)	1 − RR = (CER − EER) / CER	The *proportion* of the baseline risk removed
Absolute risk reduction (ARR)	CER − EER	The *percentage-point* drop in risk
Number needed to treat (NNT)	1 / ARR	How many people must be treated for one to benefit

The decisive identity is:

$$\mathbf{ARR \approx CER \times RRR}$$

Absolute benefit is relative benefit *scaled by* the baseline event rate. This one line explains essentially everything that follows. Because RRR tends to stay roughly constant across populations while CER varies enormously, the *same* intervention produces wildly different ARRs — and therefore wildly different NNTs — depending entirely on who is treated and for how long. RRR and ARR are not competing estimates of one underlying truth; they are two projections of the same data, and they diverge whenever the baseline risk is far from 50%.

A symmetrical pair of definitions governs harms: **absolute risk increase (ARI)** for an adverse effect, and **number needed to harm (NNH) = 1/ARI**. A complete decision weighs ARR against ARI over the same horizon.

2. Why the distinction is not pedantry

When the baseline risk is small — as it is for most preventive interventions over any short window — a large RRR and a small ARR coexist comfortably. Halving a 2% risk (RRR = 50%) is a 1-percentage-point ARR. Halving a 40% risk (also RRR = 50%) is a 20-point ARR. Same drug, same relative effect, twentyfold difference in absolute benefit.

This produces two opposite rhetorical strategies, and both appear in the wild:

- **To make an intervention sound impressive, report the RRR alone.** This is the standard way trial results reach the press: “95% effective,” “cuts heart attacks by a third.” It is not false, but without the baseline risk it is uninterpretable.
- **To make an intervention sound useless, report the trial-window ARR alone and call it the “real” effectiveness.** This is the move behind “the vaccines are less than 1% effective.” It is also not false as arithmetic, but it silently substitutes a short trial’s attack rate for a person’s actual exposure risk.

The remedy in both cases is identical: never report one without the other, and always state the baseline risk and time horizon they are computed over.

3. Worked example A — Atorvastatin (Lipitor)

3.1 The relative effect is remarkably stable

The Cholesterol Treatment Trialists' (CTT) Collaboration, pooling individual data from dozens of statin trials, found that major vascular events fall by approximately **22% per 1.0 mmol/L (~39 mg/dL) reduction in LDL cholesterol**, and that this proportional reduction is close to *constant* across baseline-risk strata. [1] This is the empirical fact that makes RRR portable: the drug removes about the same *fraction* of risk almost regardless of how much risk you started with.

3.2 The absolute effect tracks baseline risk

Hold the RRR roughly fixed and vary the population, and the ARR moves with the baseline event rate exactly as ARR ≈ CER × RRR predicts:

Trial / population	Setting	RRR	ARR	Horizon
ASCOT-LLA (atorvastatin 10 mg, hypertensives)	Primary prevention	36% (fatal CHD + non-fatal MI)	1.1%	3.3 yr [2]
ASCOT-LLA, broader composite	Primary prevention	20% (total CV events + revasc.)	1.9%	3.3 yr [2]
CARDS (atorvastatin 10 mg, type 2 diabetes)	Primary prevention (elevated risk)	37% (major CV events)	3.2%	3.9 yr [3]
TNT (intensive vs. moderate atorvastatin), age ≥65	Secondary prevention	22%	2.3%	~4.9 yr [4]

The most vivid illustration comes from stratifying *the same drug* by coronary artery calcium (CAC), a direct marker of baseline atherosclerotic risk. In analyses of statin therapy by CAC score, patients with a **CAC of 0 derived essentially no benefit** (hazard ratio ≈ 0.99), whereas patients with **CAC > 400 saw a 42–44% RRR translating into a 6.3–8.5% ARR (NNT ≈ 12–13).** [5] Identical pharmacology; the patient's baseline risk alone moves the absolute benefit from negligible to substantial. This is why modern guidelines do not prescribe statins off a cholesterol number in isolation — they prescribe off *estimated absolute risk*, because that is what determines whether the ARR is worth the harms and costs.

3.3 The harm side of the ledger, in absolute terms

A complete accounting weighs ARR against absolute harms over the same horizon:

- **Statin-associated muscle symptoms (SAMS)** are the most commonly *reported* adverse effect, but blinded randomized and re-challenge data indicate that the large majority are not pharmacologically caused — the symptoms are largely a nocebo effect, even though they are real and should not be dismissed. [6]
- **Rhabdomyolysis**, the feared severe myopathy, is genuinely rare at standard atorvastatin doses — on the order of a few cases per 100,000 patient-years — and

concentrated among patients with interacting drugs, renal impairment, or advanced age. [6]

- **New-onset diabetes** is the best-quantified metabolic harm. The Sattar 2010 meta-analysis found roughly a **9% relative increase**, about **one extra case per 255 patients treated over four years**. [7] High-intensity regimens carry a larger signal — a ~36% relative increase, an absolute annual excess near 1.3%, NNH on the order of ~21 over several years. [8] A useful framing from that literature: roughly **one additional diabetes diagnosis for every three major cardiovascular events prevented** — and the two outcomes are not of equal severity. [9]

The structure to notice: the *benefit* (ARR) scales up with baseline cardiovascular risk, while several *harms* are closer to fixed per-patient hazards. High-risk patients therefore have a strongly favorable benefit-harm ratio; very-low-risk patients can have a benefit so small in absolute terms that the harms and burdens dominate. The RRR never told you which patient you were.

4. Worked example B — The COVID-19 vaccines

4.1 The headline relative effect

The registration trials reported relative risk reductions against symptomatic COVID-19 of approximately **95% (Pfizer–BioNTech), 94% (Moderna), 91% (Gamaleya), and 67% (Janssen and Oxford–AstraZeneca).** [10] These are the “95% effective” figures that dominated coverage.

4.2 The absolute effect over the trial window

Computed over the same trials — each only a few months long, during which only a small fraction of participants were exposed and infected — the **absolute** risk reductions were far smaller: roughly **1.3% (AstraZeneca), 1.2% (Moderna and Janssen), 0.93% (Gamaleya), and 0.84% (Pfizer).** [10] Independent appraisals of the published data reproduced these figures closely (Pfizer ARR ≈ 0.7%; Moderna ARR ≈ 1.1%). [11] A GRADE synthesis across five trials likewise found RRRs spanning 45–96% but ARRs of only about 6–17 events per 1,000 over the trial windows. [12]

The 0.84% figure is arithmetically correct. It became the basis for the widely circulated claim that the vaccines were “really” less than 1% effective — a claim that even the lead author of the paper that highlighted the ARRs publicly rejected as a misreading. [13]

4.3 Why the small trial-window ARR does *not* mean small benefit

This is the crux, and it is the same identity ARR ≈ CER × RRR read in the other direction. The trial ARR is small because the **trial-window CER was small** — only a few percent of the placebo group caught symptomatic COVID during the brief follow-up. But the trial window is not the decision horizon. Over a full season or the course of a pandemic, with eventual exposure approaching universal, the relevant *cumulative baseline risk* is vastly higher than a two-month attack rate, so the *cumulative* ARR is correspondingly larger. The methodological literature makes the point directly: **NNT and ARR are properties of the population’s baseline risk and the follow-up duration, not intrinsic properties of the intervention; for a fixed RRR, ARR is larger when event rates are higher** — earlier in an outbreak, during high transmission, or in

high-risk individuals — and shrinks as risk falls. [14][15] Substituting a short trial's attack rate for a person's real exposure risk is exactly the error the "<1% effective" framing commits.

It is the same error, run in reverse, as quoting a high-risk patient's 8% statin ARR to a low-risk person. In one case a low baseline risk is used to make a real benefit look trivial; in the other a high baseline risk is used to make a marginal benefit look compelling. The arithmetic is honest in both; the *transport* of the number to the wrong baseline is the dishonesty.

A further subtlety specific to vaccines: the policy-relevant outcome is usually not symptomatic infection but **severe disease, hospitalization, and death**, whose baseline rates vary across orders of magnitude by age and comorbidity. The absolute benefit of vaccination against death in an 80-year-old and in a healthy 20-year-old differ by far more than their relative benefits do — again because ARR carries the baseline risk that RRR discards.

4.4 The harm side of the ledger, in absolute terms

The principal serious adverse signal for the mRNA vaccines is **myocarditis/pericarditis**, and it is sharply concentrated by age, sex, and dose:

- Across the broad population the excess is small — on the order of ~2.7 cases per 100,000 vaccinated in early Israeli data. [16]
- The risk concentrates in **young males after the second dose**: CDC estimates around **40 cases per million second doses in males aged 12–29**, with later commercial-claims analyses around 27 per million in males 12–24 and surveillance signals for some booster/age combinations running higher still (with wide confidence intervals). [16][17][18]
- Crucially, the comparator is not zero: **SARS-CoV-2 infection itself causes myocarditis** at rates that match or exceed the vaccine-associated rate in most demographic groups — the notable exception being adolescent and young-adult males, where the vaccine-associated rate can exceed the infection-associated rate. [16]

As with statins, the benefit (ARR against severe COVID) rises steeply with baseline risk — age, comorbidity, local transmission — while the most serious harm is concentrated in a specific low-baseline-risk subgroup. That opposing structure is precisely why recommendations became risk-stratified rather than uniform: the favorable benefit-harm balance in an older adult is not the same calculation as in a healthy adolescent male, even though the *relative* efficacy is similar in both.

5. The two cases side by side

Dimension	Atorvastatin (Lipitor)	COVID-19 vaccines (mRNA)
Typical reported RRR	~22% per 1 mmol/L LDL; 20–37% for events in trials [1][2][3]	~94–95% vs. symptomatic infection (trials) [10]
Trial-window ARR	~1–3% over 3–4 yr, rising with baseline risk [2][3]	~0.84–1.3% over ~2–4 mo trial window [10][11]

Dimension	Atorvastatin (Lipitor)	COVID-19 vaccines (mRNA)
What inflates the gap	Low absolute event rate in lower-risk arms	Low attack rate over a short trial window
How ARR scales	With cardiovascular baseline risk (e.g., CAC 0 → ~0%; CAC >400 → 6–8%) [5]	With exposure/transmission and with age–comorbidity for severe outcomes [14][15]
Principal serious harm	Rhabdomyolysis (rare); new-onset diabetes (NNH ~21–255 by intensity) [6][7][8]	Myocarditis, concentrated in young males post-dose-2 (~tens per million) [16][17]
Direction of common spin	RRR used to *oversell* to low-risk people	Trial-window ARR used to *undersell* across the board

The interventions are not clinically equivalent and should not be traded off against each other. What they share is the statistical anatomy: a stable RRR, an ARR that lives or dies on baseline risk and time horizon, and an NNT that is a fact about a population rather than a fact about a drug.

6. Methodological traps to flag in any risk claim

1. **A bare RRR with no baseline risk.** “Cuts risk by X%” is uninterpretable until you know X% of *what*. Ask for the CER.
2. **A bare ARR or NNT with no time horizon.** Both are monotone in follow-up duration; a per-trial ARR is not a lifetime ARR. Ask “over how long?”
3. **Transporting an NNT between populations.** Because NNT = 1/(CER × RRR), an NNT measured in a high-risk trial does not apply to a low-risk clinic population, and vice versa. NNT travels only if baseline risk travels with it.
4. **Attack rate masquerading as exposure risk.** A vaccine trial’s control-arm infection rate over two months is a lower bound on real-world cumulative exposure, not an estimate of it.
5. **Choosing the metric to fit the conclusion.** RRR for the things you favor, ARR for the things you don’t, is the signature of motivated framing. The discipline is to report both, every time.
6. **Composite endpoints.** Both literatures lean on composites (e.g., “major vascular events,” “symptomatic COVID of any severity”); a large RRR on a soft composite can coexist with a small effect on hard outcomes like mortality. Check what is actually in the endpoint.

7. A decision-theoretic reading

For a quantitative reader the cleanest framing is expected net benefit per person over a defined horizon T:

E[net benefit] = (baseline risk over T) × RRR × (value of the event avoided) − P(harm) × (disvalue of harm) − costs

Two structural facts fall out immediately:

- The **first term is linear in baseline risk**, because RRR is roughly constant and the baseline risk multiplies through. The **harm terms are closer to baseline-risk-independent**. So expected net benefit is an increasing function of baseline risk that crosses zero at some threshold — below which the intervention is net-negative for that individual, above which it is net-positive. Risk-stratified guidelines (ASCVD-risk thresholds for statins; age-and-comorbidity prioritization for COVID vaccines) are exactly the operationalization of this threshold.
- **RRR is the right input for the benefit term; ARR is the right output for the decision.** You estimate the stable relative effect from trials, multiply by the *individual's* baseline risk over the *relevant* horizon to get that person's ARR, and weigh it against that person's absolute harms. Reporting RRR alone hides the multiplication; reporting trial ARR alone hard-codes the wrong baseline. Both numbers are necessary inputs to the same calculation.

The popular debate treats RRR and ARR as rival claims about whether a drug "works." They are not. They are the relative effect and the absolute effect, and the whole content of the disagreement is usually a smuggled assumption about whose baseline risk and what time horizon to use.

8. Conclusion

Relative risk reduction tells you how good an intervention is at removing risk from those who carry it. Absolute risk reduction tells you how much a particular population, over a particular stretch of time, actually stands to gain. The first is a property of the intervention; the second is a property of the intervention *and the person and the clock together*. Atorvastatin and the COVID-19 vaccines have been pushed and attacked using the same lever — quoting whichever of the two numbers serves the argument and detaching it from the baseline risk that gives it meaning. Quote both; anchor both to a stated baseline and horizon; and the spin, in either direction, has nowhere left to hide.

Appendix A — Worked example for teaching: one relative effect, opposite absolute risks

The lesson of this paper reduces to a single arithmetic move. Fix the relative risk reduction; vary only the baseline event rate; recompute the absolute benefit.

> **ARR = CER × RRR and NNT = 1 / ARR = 1 / (CER × RRR)**

Because NNT is inversely proportional to baseline risk, a perfectly constant relative effect produces absolute benefits — and "number needed to treat / vaccinate" — that differ by an order of magnitude or more across populations. The intervention has not changed. The patient has.

Example A1 — A statin at a *fixed* 25% relative effect

Holding the relative risk reduction constant at a realistic ~25% for major vascular events, the absolute benefit is entirely a function of who is treated:

Population (illustrative)	5-yr baseline risk, CER	RRR (fixed)	ARR = CER × RRR	NNT = 1/ARR
Low-risk primary prevention (e.g., zero coronary calcium)	2.5%	25%	0.6%	≈ 160
Moderate-risk primary prevention	10%	25%	2.5%	40
Secondary prevention / high calcium score	30%	25%	7.5%	≈ 13

Same pill, same relative effect. You must treat roughly **160** low-risk people, but only about **13** high-risk people, to prevent one event. A bare "cuts risk by 25%" is identical across all three rows and tells a clinician nothing about which patient is worth treating. (Baseline rates here are round illustrative figures; the constant-RRR assumption is the real, well-replicated finding of the CTT meta-analysis [1].)

Example A2 — A COVID-19 vaccine at a *fixed* high relative effect

The same exercise for a vaccine exposes a second trap on top of the baseline-risk dependence — the *time window* over which the baseline risk is measured:

Scenario (illustrative unless noted)	Baseline risk, CER	RRR (fixed)	ARR	Number needed to vaccinate
Trial window (~2 months, low attack rate), symptomatic infection	~0.9% *(real anchor)*	95%	~0.85% *(real anchor)*	≈ 118
One high-transmission season, symptomatic infection	30%	95%	28.5%	≈ 4

Scenario (illustrative unless noted)	Baseline risk, CER	RRR (fixed)	ARR	Number needed to vaccinate
Severe COVID, healthy young adult, one season	0.1%	90%	0.09%	≈ 1,100
Severe COVID, elderly / comorbid, one season	5%	90%	4.5%	≈ 22

Same vaccine. The number needed to vaccinate ranges from about **1,100** (preventing severe disease in a low-risk young adult) to about **4** (preventing symptomatic infection at peak transmission). The widely quoted trial-window figure — ARR ≈ 0.84%, NNV ≈ 118 [10] — is **one unrepresentative point** on this curve, not "the effectiveness of the vaccine." (Only the first row's figures are the real reported trial values; the remaining rows are illustrative scenario numbers to show the mechanism.)

The contrast in one line, for a slide

> The same vaccine: treat about **4** people to prevent one symptomatic case at peak transmission, or about **1,100** to prevent one severe case in a healthy young adult. The drug did not change — the baseline risk and the time window did.

Classroom note

Two distinct things move the absolute number, and students should be able to name which is operating:

1. **Baseline risk of the population** (compare rows *within* either table). This is legitimate and is exactly what risk-stratified guidelines respond to.
2. **The time horizon over which baseline risk is accumulated** (compare the trial-window row to the seasonal rows in A2). Quoting an absolute number computed over the wrong horizon is the more serious error, because it is not a valid estimate of the decision-relevant ARR at all — not merely an incomplete one.

Suggested exercise: give students a fixed RRR and three baseline risks and have them compute ARR and NNT; then give them one ARR and ask them to recover the implied baseline risk. The point lands when they realize NNT alone does not identify an intervention — it identifies an intervention *applied to a population over a period*.

References

1. Cholesterol Treatment Trialists' (CTT) Collaboration. Efficacy and safety of LDL-lowering therapy; ~22% RRR in major vascular events per 1.0 mmol/L LDL-C reduction, approximately constant across baseline risk. Summarized in: Circulation (2022), "Beyond 10-Year Risk: A Cost-Effectiveness Analysis of Statins for the Primary Prevention of Cardiovascular Disease." doi:10.1161/CIRCULATIONAHA.121.057631
2. ASCOT-LLA. Atorvastatin 10 mg vs. placebo in hypertensive patients; fatal CHD + non-fatal MI: 36% RRR, 1.1% ARR; total CV events + revascularization: 20% RRR, 1.9% ARR, over a median 3.3 years. (Sever PS, et al., Lancet 2003; figures as tabulated in trial documentation.)
3. CARDS — Collaborative Atorvastatin Diabetes Study. Atorvastatin 10 mg in type 2 diabetes; major CV events 37% RRR, 3.2% ARR over a median 3.9 years. (Colhoun HM, et al., Lancet 2004; figures as tabulated in trial documentation.)
4. TNT — Treating to New Targets. Intensive vs. moderate atorvastatin in stable CHD; ~22% RRR overall, ~2.3% ARR in patients ≥65. (LaRosa JC, et al., NEJM 2005; elderly subgroup per National Lipid Association editorial, lipid.org.)
5. Statin benefit stratified by coronary artery calcium (CAC): negligible benefit at CAC 0 (HR ≈ 0.99); 42–44% RRR and 6.3–8.5% ARR (NNT ≈ 12–13) at CAC > 400. JACC: Cardiovascular Imaging (2021), editorial on plaque and relative vs. absolute statin benefit. doi:10.1016/j.jcmg.2021.07.005
6. Statin Safety and Associated Adverse Events: A Scientific Statement From the American Heart Association. Arterioscler Thromb Vasc Biol. (2019). doi:10.1161/atv.0000000000000073 (SAMS largely nocebo; rhabdomyolysis rare at standard doses.)
7. Sattar N, Preiss D, Murray HM, et al. Statins and risk of incident diabetes: a collaborative meta-analysis of randomised statin trials. Lancet 375(9716):735–742 (2010). (~9% relative increase; ~1 extra case per 255 over ~4 years.)
8. Effects of statin therapy on new-onset diabetes and worsening glycaemia: individual participant data meta-analysis. Lancet Diabetes & Endocrinology (2024). doi:10.1016/S2213-8587(24)00040-8 (High-intensity statins: 36% relative increase; absolute annual excess ~1.27%; NNH ~21.) See also the linked commentary, doi:10.1016/S2213-8587(24)00217-1.
9. Preiss D, Sattar N. Statins and the risk of new-onset diabetes: a review of recent evidence. (Framing of ~1 additional diabetes case per 3 major CV events prevented.) PMID:21897230.
10. Olliaro P, Torreele E, Vaillant M. COVID-19 vaccine efficacy and effectiveness — the elephant (not) in the room. Lancet Microbe 2:e279–e280 (2021). doi:10.1016/S2666-5247(21)00069-0 (RRRs ~67–95%; ARRs ~0.84–1.3% over trial windows.)
11. Brown RB. Outcome reporting bias in COVID-19 mRNA vaccine clinical trials. Medicina (Kaunas) (2021). PMC7996517. (Independent appraisal: Pfizer ARR ≈ 0.7%; Moderna ARR ≈ 1.1%.)
12. Comparative benefits and risks associated with currently authorized COVID-19 vaccines (GRADE synthesis): RRR 45–96%, ARR ~6–17 per 1,000 across trials. PMC9788423.

13. PolitiFact fact-check (May 28, 2021) and AAP/AFP fact-checks, quoting Piero Olliaro rejecting the “vaccines <1% effective” interpretation of the ARR figures.
14. Correia LCL, Matias D. COVID-19 vaccines: effectiveness and number needed to treat. Lancet Microbe 2:e281 (2021). doi:10.1016/S2666-5247(21)00119-1 (NNT/ARR is a property of population baseline risk, not of the treatment.)
15. Olliaro P, et al. Authors’ reply. Lancet Microbe (2021). doi:10.1016/S2666-5247(21)00120-8 (For a fixed RRR, ARR rises with event rate — high-risk groups, earlier in a programme — and falls as risk declines.)
16. Vaccine-Associated Myocarditis Risk in Context. American College of Cardiology (Feb 2022). (Population excess ~2.7 per 100,000; infection-associated myocarditis matches or exceeds vaccine-associated rates except in young males.)
17. CDC / Bozkurt B, et al. Myocarditis with COVID-19 mRNA vaccines. Circulation (2021). doi:10.1161/CIRCULATIONAHA.121.056135; CDC ACIP estimates ~40 cases per million second doses in males 12–29.
18. FDA. Updated warning in labeling of mRNA COVID-19 vaccines regarding myocarditis/pericarditis (2025): ~8 per million doses ages 6 mo–64 yr; ~27 per million in males 12–24 (2023–2024 formulation, days 1–7).

Prepared as a methodological white paper on relative vs. absolute risk. Figures are drawn from the cited primary trials, meta-analyses, and surveillance sources; illustrative thresholds (e.g., the expected-net-benefit crossover) are conceptual rather than empirically fixed and are labeled as such. Not individualized medical advice.