Apixaban and Its Dose

A Bayesian network for the dosing decision in atrial fibrillation, written for clinical readers.
Murray Cantor, Healthy Skepticism papers, August 2026.

Summary

The apixaban label halves the dose for patients meeting two of three criteria (age 80 or above, weight 60 kg or below, creatinine 1.5 mg/dL or above), yet no randomized comparison of 5 mg with 2.5 mg twice daily exists in that population, and in practice most dose reduction happens off label. This paper presents a Bayesian network that treats the dose as a formal intervention with three values, 5 mg, 2.5 mg, and 0 mg twice daily, and returns annual absolute probabilities of stroke or systemic embolism, major bleeding split into intracranial and extracranial, and a fatal-or-disabling composite, for a patient described by sex, age, weight, creatinine, AF burden, and prior stroke or TIA. Rates are anchored to ARISTOTLE, AVERROES, and their subgroup analyses; dose-response for stroke and for bleeding is learned as Weibull exposure curves by Bayesian synthesis of published anchors, with cross-trial anchors down-weighted and disagreements between sources carried forward as posterior width rather than resolved by assumption. On-treatment predictions reproduce the observed ARISTOTLE renal-band rates exactly; prescribing-pattern checks the model was not fitted to come out within a point or two of observation. The central finding: half dose likely retains most of the stroke protection (53% risk reduction versus no treatment, 5th to 95th percentile 29 to 74, against 61% for full dose), the drug’s bleeding excess is almost entirely extracranial, and no anticoagulant is by far the most dangerous of the three interventions for every patient profile examined. An anchor sensitivity analysis shows the half-dose conclusion leans most on the PCORnet null; the paper quantifies exactly how much rides on how one prices its residual confounding. The model is distributed in two audited, mutually validating forms: dependency-free Python, and an Excel workbook whose formulas contain the entire network, so the model can be run, interrogated, and edited without writing any code.

The clinical problem

ARISTOTLE established apixaban 5 mg twice daily against warfarin (stroke or systemic embolism 1.27 versus 1.60 per 100 patient-years; major bleeding 2.13 versus 3.09; Granger et al., 2011). A 2.5 mg twice-daily dose was assigned, inside the trial, to patients meeting at least two of three criteria: age 80 or above, weight 60 kg or below, serum creatinine 1.5 mg/dL or above. About 4 to 5% of the trial met the rule. The rule’s logic is exposure matching: each criterion proxies reduced clearance, and halving the dose in such patients was intended to reproduce the standard exposure rather than to test a lower one. As a consequence, the trial

cannot say what 2.5 mg does in a patient who does not meet the criteria, and it randomized nobody between the two doses.

Practice has not respected that boundary. In a 2026 PCORnet analysis of over 400,000 apixaban users, 15.4% were on the reduced dose and 61% of those did not meet the label criteria (Gouda et al., 2026). The same analysis found no detectable stroke excess in the underdosed group (adjusted HR 1.05, 95% CI 0.90 to 1.23), an observation that any account of the dose-response has to accommodate. Meanwhile the patient's own question is not the trial's question. Warfarin is no longer the alternative that patients weigh; the live alternatives are the full dose, the half dose, and no anticoagulant at all. Every comparison in this paper is therefore against no treatment, with warfarin appearing only inside evidence chains.

Dose as an intervention, not a covariate

Observational contrasts between doses are confounded by indication: the patients who receive 2.5 mg are selected precisely for markers of frailty and clearance. The PCORnet cohort illustrates the trap. Off-label reduced-dose patients had higher mortality (adjusted HR 1.77) with no matching excess of stroke or bleeding, the signature of sicker patients dying of other causes rather than of a dosing effect. Conditioning on dose received answers "how do patients on 2.5 mg fare," which is dominated by who they are. The clinical question is interventional: what would this patient's risk be if we set the dose.

The model makes that distinction mechanical (Pearl, 2009). Prescribing is represented explicitly as arrows from age, weight, and creatinine into the dose node, either the label's deterministic 2-of-3 rule or a real-world policy calibrated to PCORnet. A do() query severs those arrows and fixes the dose, so "who gets 2.5 mg" and "what 2.5 mg does" cannot be conflated. Both kinds of query are available, and the gap between them is itself informative.

The model

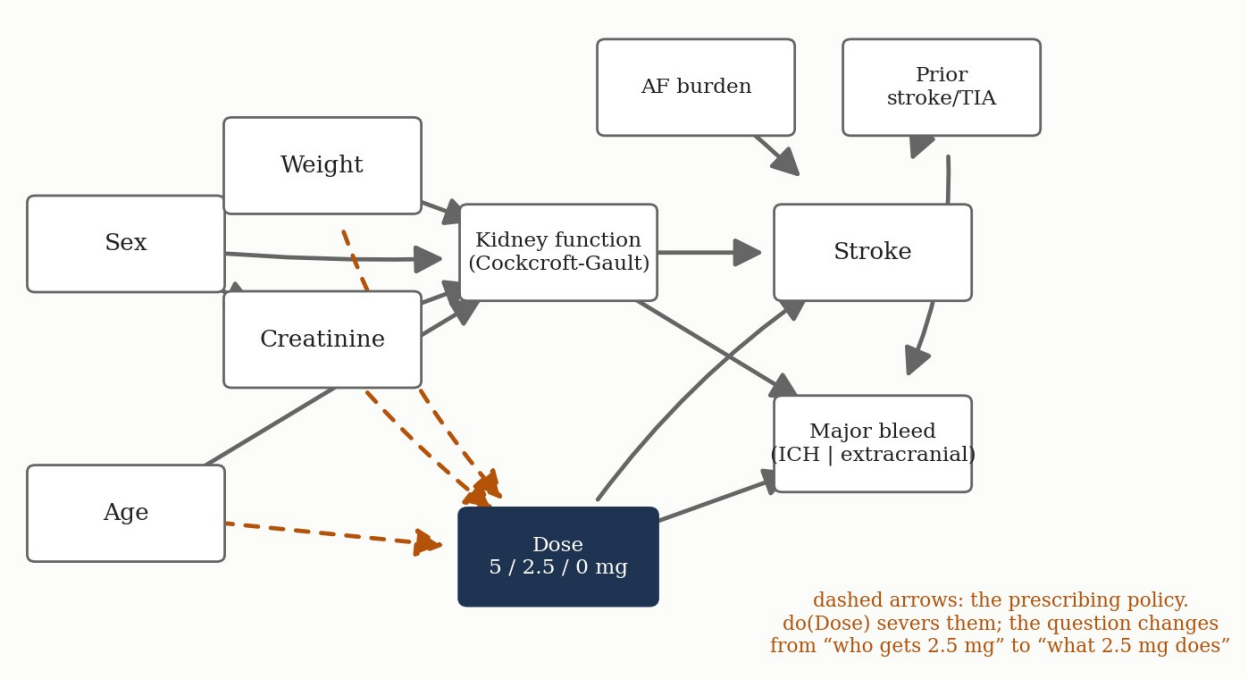

Figure 1. The network. Sex, age, weight, and creatinine determine creatinine clearance (Cockcroft-Gault); AF burden and prior stroke/TIA are observation nodes; the dose node takes 5, 2.5, or 0 mg twice daily. Dashed arrows are the prescribing policy, severed under do().

Observed on-treatment event rates anchor the model where the evidence is strongest. Stroke or systemic embolism and major bleeding on 5 mg are taken directly from the ARISTOTLE renal subgroups (stroke 1.05, 1.46, 2.39 per 100 patient-years for CrCl above 80, 50 to 80, and 25 to 50; Hohnloser et al., 2012), with the band below 25 mL/min an explicit extrapolation. Everything else is expressed as a multiplier on those rates, and each multiplier is learned, not assumed.

Exposure. Relative exposure scales the dose by a renal multiplier,

with m rising from 1.00 above 80 mL/min to 1.37 below 25. The cap follows from pharmacokinetics: apixaban is about 27% renally cleared, so even complete renal failure cannot raise exposure beyond roughly 1/(1 − 0.27). This term produces a clinically important self-correction: 2.5 mg in a patient with CrCl 25 to 50 sits at e = 0.625, not 0.5, so the reduced dose gives up less in exactly the patients the rule targets.

Dose-response for stroke. The fraction of the full log relative effect retained at exposure e is modeled as a Weibull curve,

so that the relative risk at exposure e, against no treatment, is

The family spans the two positions clinicians actually hold, a threshold (effect flat above a knee, k large) and proportionality (F(e) ≈ e), and lets the evidence place the curve between them. The full-exposure effect against no treatment is itself uncertain: it is chained from AVERROES (HR 0.45 versus aspirin; Connolly et al., 2011) and the aspirin-versus-placebo meta-analysis (RRR 22%; Hart et al., 2007), giving a central 65% risk reduction with a standard error of 0.26 on the log scale. That uncertainty enters the posterior as its own dimension rather than being frozen at the point estimate.

Dose-response for bleeding. Major bleeding gets the same machinery with its own curve,

a second Weibull posterior. Nothing forces bleeding and stroke to trade off along the same curve; the data decide each separately.

AF burden and prior stroke. Burden b (fraction of time in AF) scales the untreated stroke hazard through a learned curve g(b) with a floor at low burden, reflecting the temporal dissociation between episodes and strokes. Prior stroke or TIA multiplies the stroke hazard by 2.44 and bleeding by 1.43 (Easton et al., 2012), renormalized so the 19%/81% trial mix reproduces the band rates. In both cases the drug's relative effect is taken as constant across strata, licensed by the reported interaction nulls ($p = 0.71$ and 0.69 for prior stroke; consistency across AF type in Al-Khatib et al., 2013, and in ARTESiA).

Calibration is per posterior draw. Only ratios of the dose-response curve enter the stroke and bleeding calculations, so for every draw from every posterior, do(5 mg) reproduces the observed ARISTOTLE band rates exactly, and the untreated baseline is backed out per draw. Uncertainty therefore concentrates where the evidence is weakest, in the do(0 mg) and do(2.5 mg) rows, which is where it belongs. Inference is exact enumeration over the discrete joint; there is no simulation error.

Evidence synthesis

The dose-response, bleeding, and burden curves are fitted as grid posteriors from published summary statistics, each anchor a Gaussian likelihood on the log rate-ratio scale, in the tradition of the confidence profile method (Eddy et al., 1992). Two judgments recur and deserve scrutiny. First, cross-drug anchors (edoxaban, dabigatran, low-intensity warfarin) are used for curve shape under an exchangeability assumption, with standard errors widened to pay for it; they inform, but cannot dominate, the apixaban-specific evidence. Second, when

anchors disagree, the posterior is allowed to stay wide. The PCORnet null wants a flat stroke curve at half dose; ENGAGE's 30 versus 60 mg arms want a steep one; the bands in Figures 2 and 3 are the honest resolution. Table 1 lists every anchor with its source and role.

What it constrains	Anchor and source
On-treatment stroke/SE and major bleeding by CrCl band	ARISTOTLE renal subgroups (Hohnloser et al., 2012); apixaban-arm bleeding via the 2.13/2.61 arm ratio (Granger et al., 2011). Reproduced exactly at do(5 mg).
Full-dose effect versus no treatment	AVERROES HR 0.45 versus aspirin (Connolly et al., 2011) chained with aspirin RRR 22% versus placebo (Hart et al., 2007): RRR ≈65%, ln-scale SE 0.26, carried as a posterior dimension.
Stroke dose-response shape F(e)	Zeitouni et al. (2020), on-label 2.5 mg consistency; Gouda et al. (2026) PCORnet null, SE inflated for residual confounding; ENGAGE 30 versus 60 mg and the low-intensity-warfarin era as class anchors with widened SEs.
Bleeding dose-response shape and magnitude	RE-LY 110 versus 150 mg, 2.71/3.11 per 100 pt-yr (Connolly et al., 2009) and ENGAGE 30 versus 60 mg, 1.61/2.75 (Giugliano et al., 2013) as class shape anchors; AVERROES magnitude 1.4 versus 1.2 (HR 1.13, CI 0.74 to 1.75), SE inflated because the baseline is aspirin.
ICH versus extracranial split	ICH is 15.5% of major bleeds on 5 mg (ARISTOTLE, 0.33/2.13 per 100 pt-yr); held dose-flat because AVERROES found apixaban ICH equal to aspirin ICH (0.4 versus 0.4, HR 0.85, CI 0.38 to 1.90).
AF burden curve g(b)	ARISTOTLE paroxysmal versus persistent, 0.98 versus 1.52 per 100 pt-yr (Al-Khatib et al., 2013); KP-RHYTHM top burden tertile aHR ≈3.1 (Go et al., 2018); ARTESiA subclinical baseline (Healey et al., 2024).
Prior stroke / TIA	ARISTOTLE subgroup, 19% prevalence (Easton et al., 2012): stroke ×2.44 (2.46 versus 1.01), bleeding ×1.43 (2.84 versus 1.98); drug effect consistent across strata.
Severity weights (death or severe disability)	ICH 0.76 and extracranial 0.03 (Fang et al., 2007); stroke from 30-day mortality of 19.6% on and 27.7% off anticoagulation (Fang et al., 2012) plus an assumed 35% severe-disability share among survivors, the one unverified weight.

Table 1. Evidence anchors. Every entry is an editable line in the model's parameter block.

Results

The stroke dose-response posterior

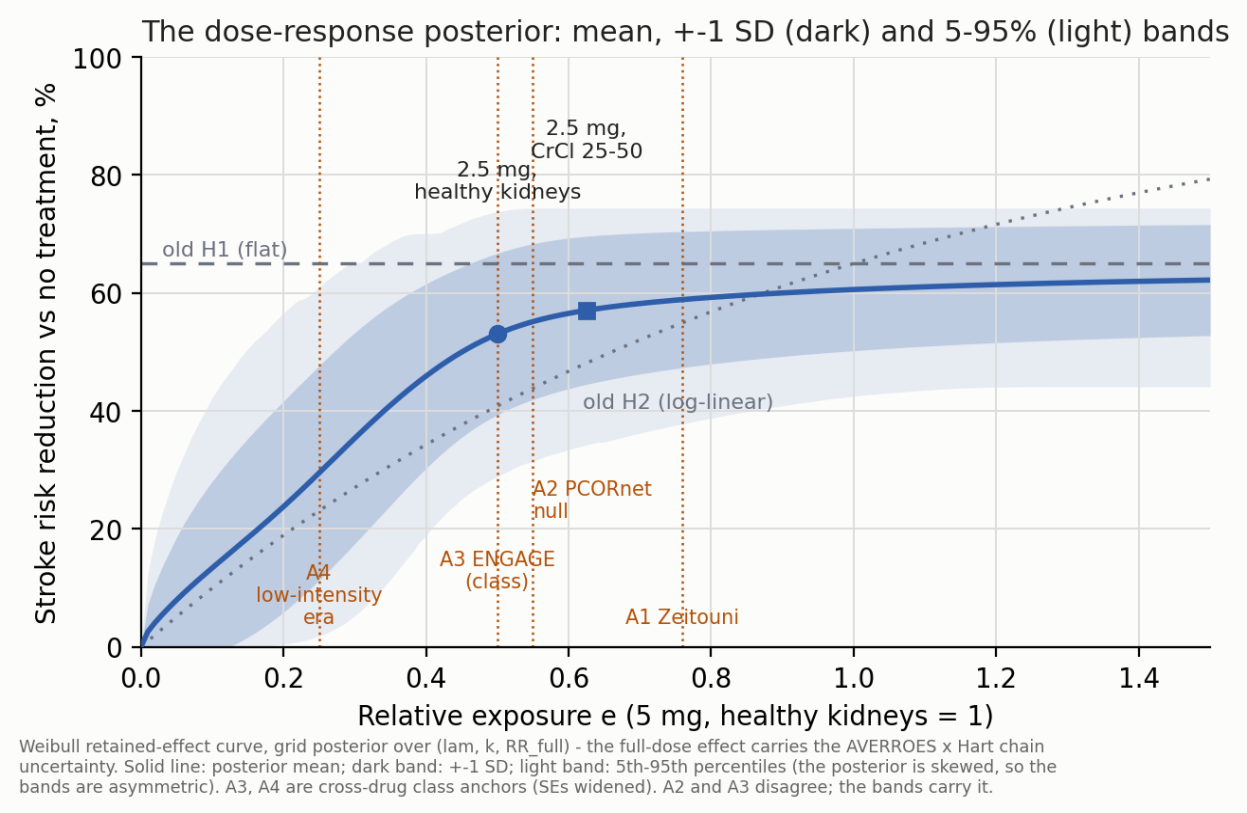

Figure 2. Posterior stroke risk reduction versus no treatment as a function of relative exposure: median, 5th to 95th percentile band, anchors, and the two legacy hypotheses as reference lines.

The posterior places the knee of the curve near e of 0.35 to 0.45 and rises steeply through the half-dose region. At full exposure the risk reduction is 61% (43 to 74). At e = 0.5, the off-label half dose in a patient with normal renal function, the posterior mean retained protection is 53% (29 to 74): most of the effect, probably, but with a 5th percentile low enough to keep the question live. At e = 0.625, the on-label half dose in a patient with CrCl 25 to 50, retention rises to 57% (34 to 74). The label's arithmetic partially protects its own patients; the exposed flank is the off-label reduction in patients with good clearance, and even there the PCORnet null keeps the posterior from endorsing the steep curve ENGAGE alone would imply.

Bleeding, and what kind of bleeding

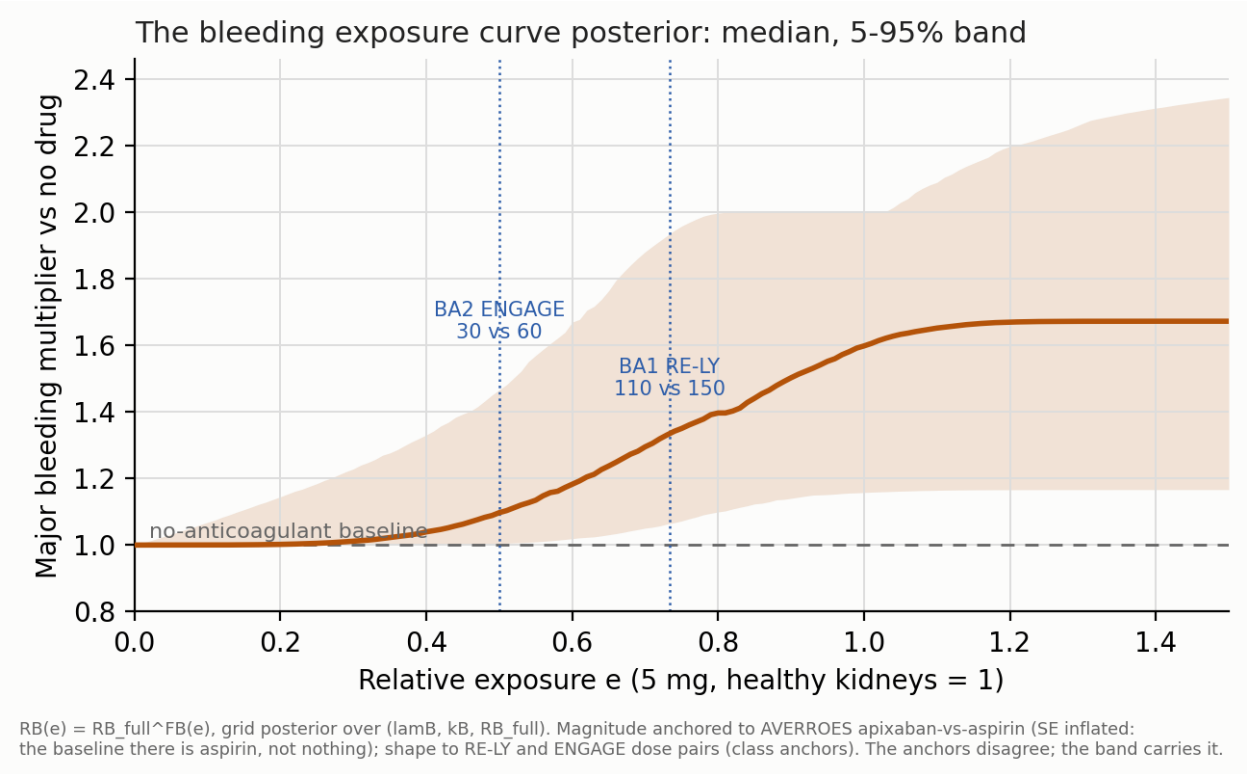

Figure 3. Posterior major-bleeding multiplier versus no anticoagulant as a function of exposure, with the RE-LY and ENGAGE dose-pair anchors.

The bleeding posterior implies a full-dose multiplier of roughly 1.5 against no anticoagulant (the posterior compromises between ENGAGE's steep dose gradient and AVERROES's near-null against aspirin; the band is wide). The clinically decisive fact is not the total but the split. Apixaban's intracranial hemorrhage rate in AVERROES equaled aspirin's, so the model holds ICH flat across all three interventions, and the entire dose-responsive excess is extracranial, a class of bleed with 3% death-or-severe-disability against 76% for ICH (Fang et al., 2007). On the outcomes that resemble stroke in their finality, the drug is close to free; its cost is paid almost entirely in survivable bleeding.

The patient-level ledger

Intervention	Stroke/SE [5-95%]	Major bleed [5-95%]	ICH	Extracranial	Fatal or disabling
Patient A: age 80+, weight above 60 kg, creatinine 1.5–2.0 (meets the 2-of-3 rule)					
do(5 mg)	2.10 [1.89–2.27]	3.94	0.61	3.33	1.57
do(2.5 mg)	2.34 [1.95–2.89]	3.06 [2.36–3.85]	0.61	2.45	1.66
do(0 mg)	5.83 [3.64–8.60]	2.47 [1.77–3.37]	0.61	1.86	3.61
Patient B: age 65–74, weight above 60 kg, creatinine below 1.0 (meets no criterion)					

do(5 mg)	1.12 [1.01–1.21]	1.87	0.29	1.58	0.81
do(2.5 mg)	1.30 [1.05–1.63]	1.39 [1.07–1.79]	0.29	1.10	0.88
do(0 mg)	3.07 [1.90–4.57]	1.20 [0.91–1.61]	0.29	0.91	1.87

Table 2. Annual probabilities (%) under each intervention, posterior means with 5th to 95th percentile bands. On-treatment 5 mg rows are calibrated to observed rates, so their bleeding percentiles collapse; ICH is dose-flat by assumption.

Two readings of the table matter. Within each patient, the gap between either active dose and do(0 mg) dwarfs the gap between the doses: for Patient A, stopping anticoagulation raises the fatal-or-disabling rate from about 1.6 to 3.6 per 100 patient-years, while the full-versus-half question moves it by less than 0.1 with overlapping bands. The dosing controversy is real but second-order; the first-order error is the patient who, offered a confusing choice, takes nothing. Second, for Patient B the half dose is off label, and the model prices it honestly: a probable but uncertain loss of protection (stroke 1.30 versus 1.12) bought with a real reduction in extracranial bleeding (1.10 versus 1.58). Reasonable clinicians can land on either side of that trade for a bleeding-prone patient; the model's contribution is to put numbers and bands on what is being traded.

AF burden

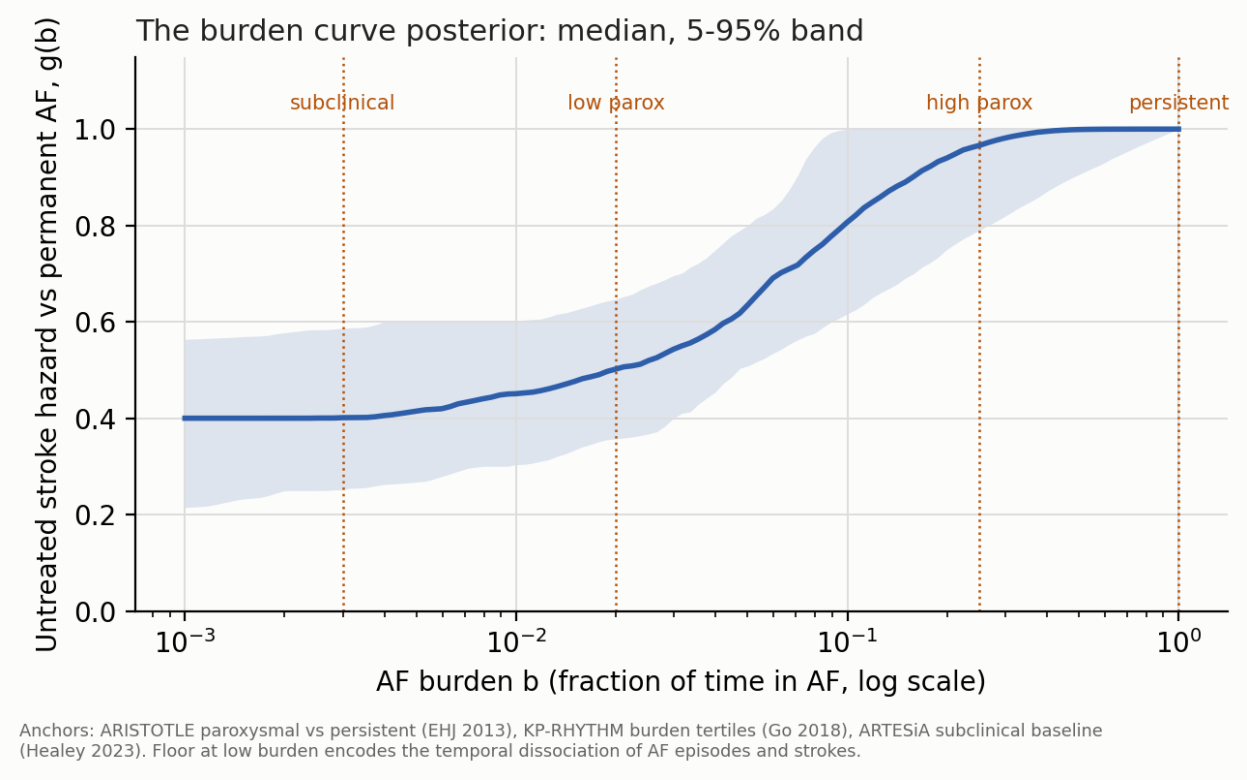

Figure 4. Posterior untreated-hazard multiplier g(b) against AF burden, normalized to permanent AF, with the four burden states marked.

Burden shifts the whole ledger without changing the relative drug effect. For Patient A with persistent AF the untreated stroke rate is 7.0 per 100 patient-years; with brief device-detected episodes it is 2.9, and on 5 mg it falls to about 1.0, now below the unchanged bleeding cost. That inversion is the ARTESiA landscape: in subclinical AF the absolute benefit shrinks toward the bleeding excess, and the decision genuinely tightens. As a cross-check the model was not tuned to, its population-level subclinical untreated rate is 1.93 (0.99 to 3.23) per 100 patient-years, bracketing the roughly 1.6 implied by ARTESiA's aspirin arm (Healey et al., 2024).

Prior stroke or TIA

Secondary prevention concentrates everything. For Patient A with a prior stroke or TIA, the untreated rate is 11.2 (7.0 to 16.5) per 100 patient-years, 5 mg brings it to 4.0, and the fatal-or-disabling composite falls from 6.6 untreated to 2.7 on full dose. Bleeding rises too (5.2 on 5 mg), but the asymmetry in severity means the composite still favors treatment by a factor of about 2.5. The model applies the Easton multipliers within every stratum, so the same query works for the patient who is both post-stroke and renally impaired, which is where the dosing rule and the secondary-prevention imperative collide.

Prescribing policies at the population level

Set every patient in the modeled 65-plus AF population to 5 mg and the population stroke rate is 1.42 (1.28 to 1.54) per 100 patient-years; the label's 2-of-3 rule gives 1.44; real-world prescribing, with its 61% off-label reductions, gives 1.45; universal half dose would give 1.62. Two conclusions follow. The label rule as written costs almost nothing at population level, because it reduces few patients and partially exposure-matches the ones it does. And even widespread off-label reduction produces a population signal (about 0.01 per 100 patient-years here) far too small for observational data to detect, which is exactly why the PCORnet null, reassuring as it is, cannot settle the individual-level question that the do() queries answer.

Where the evidence disagrees, and how much it matters

The half-dose question turns on two observations that cannot both be taken at face value. In ENGAGE AF-TIMI 48, the low-dose edoxaban arm suffered roughly 40% more strokes than the standard arm, randomized evidence that halving the dose of a factor Xa inhibitor costs real protection, but evidence from a different drug with a different renal clearance fraction, admitted here only as a class anchor with a widened standard error (Giugliano et al., 2013). In PCORnet, more than 60% of reduced-dose apixaban patients did not meet the label criteria, and their adjusted stroke hazard was 1.05 (0.90 to 1.23) against full-dose patients, apixaban-specific evidence at enormous scale, but observational: prescribers reduce the dose on exactly the frailty and clearance signals that also predict stroke, and residual

confounding could mask a real excess (Gouda et al., 2026). The model inflates the PCORnet standard error from 0.08 to 0.15 to price that risk, a stated judgment rather than a derived quantity.

Bayesian synthesis does not choose between the two; each anchor pulls on the curve with weight inverse to its assigned variance, and the posterior band carries what remains unresolved. But a reader deserves to know which anchor is load-bearing, so Table 3 refits the entire model under variations of the anchor set and reports the quantity the controversy is about, the protection retained at half exposure, alongside the central case's half-dose stroke rate.

Anchor-set variant	RRR at e = 0.5	RRR at e = 0.625	Patient A stroke, 2.5 mg
Baseline (all four anchors)	53 [29–74]	57 [34–74]	2.34 [1.95–2.89]
Drop the ENGAGE class anchor	55 [30–74]	58 [35–74]	2.30 [1.94–2.84]
Drop the PCORnet null	44 [13–70]	51 [23–73]	2.71 [2.00–3.93]
PCORnet at face value (SE 0.08)	57 [33–74]	60 [37–74]	2.22 [1.93–2.63]
ENGAGE at trial precision (SE 0.15)	50 [27–70]	55 [33–74]	2.43 [1.98–2.99]

Table 3. Anchor sensitivity. Each row refits the full posterior. RRR entries are percent stroke risk reduction versus no treatment, posterior mean [5th to 95th percentile]; the last column is annual percent for the Table 2 Patient A.

The table says something sharper than "the model is uncertain." The synthesis is nearly indifferent to the ENGAGE anchor: drop it and every number moves by a point or two, because the tighter PCORnet likelihood already dominates that region of the curve. The load-bearing anchor is the PCORnet null itself. Remove it and retained protection at half exposure falls from 53% to 44%, the 5th percentile collapses from 29% to 13%, and Patient A's half-dose stroke estimate grows a long right tail. A reader's posture toward this model should therefore track their posture toward one question: how much residual confounding lives in an adjusted hazard ratio of 1.05 from prescription data? Anyone who answers "plenty" should read the drop-PCORnet row as their version of the model, and it counsels noticeably more caution about the half dose. The workbook makes rerunning any of these variants a matter of editing one standard-error cell.

What would settle the question is equally concrete. A randomized comparison of the doses in criteria-negative patients is the clean answer and is unlikely to be run: at event rates near 1 per 100 patient-years, excluding a hazard ratio of 1.3 requires tens of thousands of patient-years per arm. Short of that, two designs would tighten the posterior materially. Exposure-bridging studies, anti-Xa activity or plasma level measured directly in reduced-dose patients, would collapse uncertainty on the e axis, where Zeitouni et al. (2020) is currently the only apixaban-specific point. And cohort analyses restricted to single-criterion patients,

where confounding by indication is structurally weaker, are beginning to appear (Lee et al., 2025) and slot directly into the anchor list as they mature.

Communicating the ledger to patients

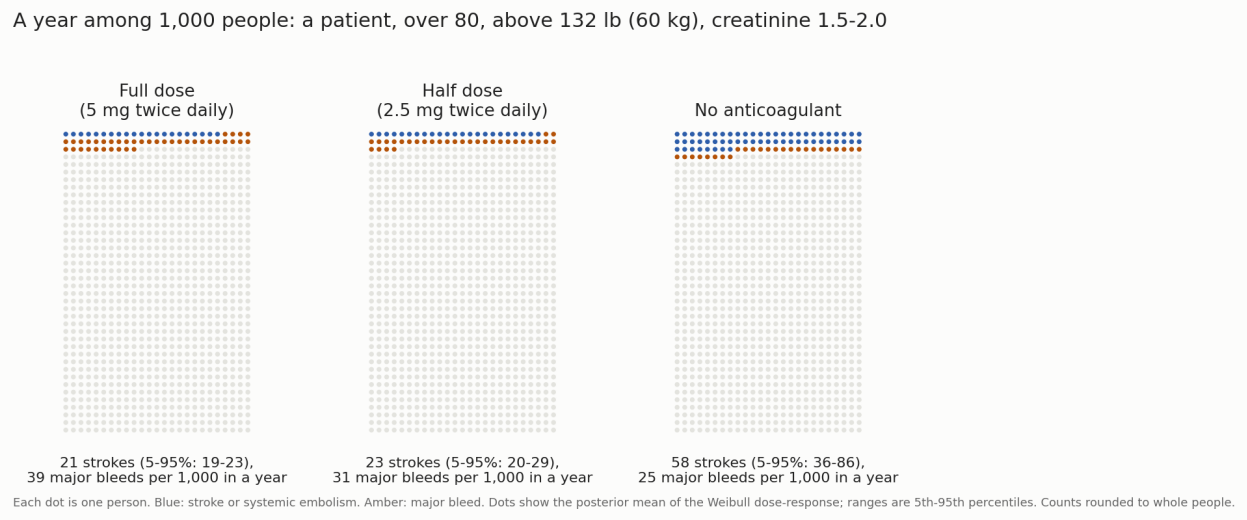

Figure 5. The same ledger as an icon array for Patient A: each dot one person in a year, blue for stroke or systemic embolism, amber for major bleed.

For shared decision-making the model renders any scenario as an icon array of 1,000 patients, the format with the best evidence base in risk communication. The arrays make the central asymmetry visible without arithmetic: moving from the left panel to the right, the blue dots nearly triple while the amber dots thin only modestly, and none of the added amber carries the finality of the added blue.

Validation

By construction, do(5 mg) reproduces the observed ARISTOTLE band rates exactly for every posterior draw. Checks the model was not fitted to: the share of the modeled population assigned 2.5 mg is 5.4% under the label rule, against 4 to 5% observed in ARISTOTLE (Zeitouni et al., 2020), and 13.5% under the real-world policy, against 15.4% in PCORnet (Gouda et al., 2026); the subclinical-AF untreated rate brackets ARTESiA as above. The companion spreadsheet implements the identical model in formulas alone and agrees with the reference implementation to the fourth decimal on every scenario tested, so the numbers can be audited without running code.

Limitations

The model is assembled from published aggregates, not fitted to patient-level data, and it is honest about what that buys and what it costs. Cross-drug shape anchors assume exchangeability of dose-response across factor Xa inhibitors and

dabigatran; the SEs are widened but the assumption is not testable within the model. Separability is assumed twice: the drug's relative effect is taken as constant across burden and prior-stroke strata (supported by interaction nulls, which are weak instruments), and the burden and prior-stroke multipliers are treated as independent. ICH dose-flatness rests on a single wide confidence interval. Age acts only through Cockcroft-Gault, so two patients in the same CrCl band get the same rates regardless of age; the renal bands absorb much of the age gradient but not the part that runs through fall risk and frailty. Probabilities are annual, with no time dimension, no competing mortality, and no recurrence; over a decade the ledger compounds in ways the model does not represent. CrCl below 25 is extrapolation, and the trials' exclusion of such patients is itself information. The stroke disability weight (35% of survivors severely disabled) is the one severity parameter without a verified source. None of these limitations is buried in code: each corresponds to a named, commented, editable line.

The model as a spreadsheet

Most clinical readers will meet this model not as code but as an Excel workbook, and the workbook is not a front end to a hidden engine: it IS the model. Its Engine sheet holds the exact enumeration of the joint distribution, 1,152 patient states, as ordinary visible formulas; the conditional probability tables, the Cockcroft-Gault computation, the 2-of-3 rule, and the do() severing of the prescribing arrows are all cell formulas a reader can click on and trace. The only pieces Excel does not compute are the three learned posterior curves, whose grid fits require the companion script; their output enters the workbook as clearly marked pasted tables, and everything downstream of those tables is live arithmetic.

The working surface is the Scenario sheet. A reader describes a patient with dropdowns: sex, age band, weight (in pounds, six bands, with 132 lb marking the label's 60 kg criterion), serum creatinine, AF burden, and prior stroke or TIA, any of which can be left at "Any" to integrate over the population prior. One table then shows all five interventions at once, do(5 mg), do(2.5 mg), do(0 mg), the label policy, and real-world prescribing, against all five outputs: stroke or systemic embolism, total major bleeding, ICH, extracranial bleeding, and the fatal-or-disabling composite. A single statistic selector switches the whole table between the posterior mean and the 5th or 95th percentile, and a chart beside the table redraws the three do() interventions as absolute annual risks with 5th-to-95th whiskers the moment any observation changes. Every constant the model uses, the trial rates, the severity weights, the prior-stroke multipliers, the prescribing probabilities, sits on the Parameters sheet as an editable cell annotated with its source, and the sheet's built-in checks (the 2.5 mg shares against ARISTOTLE and PCORnet, the weight and multiplier normalizations) recompute alongside any edit.

The workbook is validated rather than trusted. Every results cell has been compared against the reference Python implementation across scenarios, interventions, and statistics, agreeing to four decimal places, and the network itself has been rebuilt independently in pgmpy, a standard Bayesian network library, matching the reference to machine precision, including the do() semantics. The script can also read its parameters from the workbook, making the spreadsheet the single source of truth for a reader who works there. A reader who disputes any anchor, the PCORnet standard error, the exchangeability discount, the ICH flatness, the disability weight, changes one cell or one line and watches every number in this paper regenerate. That is the intended mode of disagreement: the model's claim is not that its numbers are right, but that its numbers follow from stated evidence by stated arithmetic, and that competing numbers should be made to do the same. The workbook is hosted with this paper at healthy-skepticism.com/eliquis.html in a locked form whose only editable cells are the Scenario inputs, so a clinician can run any patient in a minute with no risk of silent modification; the open workbook, the script, and the validation checks are distributed with this paper through the Healthy Skepticism papers library for readers who want to change more than the inputs.

References

Al-Khatib, S. M., Thomas, L., Wallentin, L., et al. (2013). Outcomes of apixaban vs. warfarin by type and duration of atrial fibrillation: results from the ARISTOTLE trial. European Heart Journal, 34(31), 2464–2471.

Connolly, S. J., Ezekowitz, M. D., Yusuf, S., et al. (2009). Dabigatran versus warfarin in patients with atrial fibrillation. New England Journal of Medicine, 361(12), 1139–1151.

Connolly, S. J., Eikelboom, J., Joyner, C., et al. (2011). Apixaban in patients with atrial fibrillation. New England Journal of Medicine, 364(9), 806–817.

Easton, J. D., Lopes, R. D., Bahit, M. C., et al. (2012). Apixaban compared with warfarin in patients with atrial fibrillation and previous stroke or transient ischaemic attack: a subgroup analysis of the ARISTOTLE trial. Lancet Neurology, 11(6), 503–511.

Eddy, D. M., Hasselblad, V., and Shachter, R. D. (1992). Meta-Analysis by the Confidence Profile Method: The Statistical Synthesis of Evidence. Academic Press.

Fang, M. C., Go, A. S., Chang, Y., et al. (2007). Death and disability from warfarin-associated intracranial and extracranial hemorrhages. American Journal of Medicine, 120(8), 700–705.

Fang, M. C., Go, A. S., Chang, Y., et al. (2012). Thirty-day mortality after ischemic stroke and intracranial hemorrhage in patients with atrial fibrillation on and off anticoagulants. Stroke, 43(7), 1795–1799.

Giugliano, R. P., Ruff, C. T., Braunwald, E., et al. (2013). Edoxaban versus warfarin in patients with atrial fibrillation. New England Journal of Medicine, 369(22), 2093–2104.

Go, A. S., Reynolds, K., Yang, J., et al. (2018). Association of burden of atrial fibrillation with risk of ischemic stroke in adults with paroxysmal atrial fibrillation: the KP-RHYTHM study. JAMA Cardiology, 3(7), 601–608.

Gouda, P., Harrington, J., Arps, K., et al. (2026). Clinical patterns and appropriateness of apixaban dosing in patients with atrial fibrillation. JACC: Advances, 5(8), 102954.

Granger, C. B., Alexander, J. H., McMurray, J. J. V., et al. (2011). Apixaban versus warfarin in patients with atrial fibrillation. New England Journal of Medicine, 365(11), 981–992.

Hart, R. G., Pearce, L. A., and Aguilar, M. I. (2007). Meta-analysis: antithrombotic therapy to prevent stroke in patients who have nonvalvular atrial fibrillation. Annals of Internal Medicine, 146(12), 857–867.

Healey, J. S., Lopes, R. D., Granger, C. B., et al. (2024). Apixaban for stroke prevention in subclinical atrial fibrillation. New England Journal of Medicine, 390(2), 107–117.

Hohnloser, S. H., Hijazi, Z., Thomas, L., et al. (2012). Efficacy of apixaban when compared with warfarin in relation to renal function in patients with atrial fibrillation: insights from the ARISTOTLE trial. European Heart Journal, 33(22), 2821–2830.

Jones, C. A., McQuillan, G. M., Kusek, J. W., et al. (1998). Serum creatinine levels in the US population: Third National Health and Nutrition Examination Survey. American Journal of Kidney Diseases, 32(6), 992–999.

Lee, S.-R., Choi, J., Kwon, S., et al. (2025). Apixaban outcomes in atrial fibrillation patients with a single-dose reduction criterion: ASPIRE 1-year results. European Heart Journal – Cardiovascular Pharmacotherapy, 11(5), 403–411.

Pearl, J. (2009). Causality: Models, Reasoning, and Inference (2nd ed.). Cambridge University Press.

Zeitouni, M., Giczewska, A., Lopes, R. D., et al. (2020). Clinical and pharmacological effects of apixaban dose adjustment in the ARISTOTLE trial. Journal of the American College of Cardiology, 75(10), 1145–1155.