Eliquis and Its Dose: A quantitative review of the benefits, the bleeding risks, and the evidence behind the 2.5 mg reduction

Summary

Apixaban, sold as Eliquis, is an oral anticoagulant taken twice a day to prevent strokes in people with atrial fibrillation and to treat and prevent venous blood clots. It is the most widely used drug of its class in the United States, taken by several million Medicare enrollees alone (CMS, 2023). The comparison a patient cares about is taking the drug versus not taking it, and that trial has never been run; the evidence has to be assembled from a chain of comparisons. Assembled, it is strong. Against aspirin, which is nearly the untreated state, apixaban cut strokes by more than half, about 21 events per 1,000 patient-years, with no measurable increase in major bleeding (Connolly et al., 2011). Against warfarin, the older standard it displaced, the margins are thinner: about 3 fewer strokes or systemic embolisms, 10 fewer major bleeds, and 4 fewer deaths per 1,000 patient-years (Granger et al., 2011).

The dosage story is less tidy. The approved rule cuts the dose from 5 mg twice daily to 2.5 mg twice daily only when a patient meets at least two of three criteria: age 80 or older, weight 60 kg or less, or serum creatinine 1.5 mg/dL or higher. That rule rests on pharmacokinetic modeling rather than on a randomized comparison of the two doses, and fewer than 5% of trial patients received the reduced dose (Zeitouni et al., 2020). In practice the reduced dose is used far more often than the rule allows. In a cohort of about 46,000 US apixaban patients, 61% of those on the reduced dose did not meet the criteria (Gouda et al., 2026). Whether that common underdosing costs strokes or lives is unsettled, and the observational studies that seem to show it does carry a familiar confounding signature. This paper walks through the benefit ledger, the bleeding ledger, and what is actually known about the dose, and closes with a small Bayesian network that prices the open question. Its verdict, with the uncertainty shown rather than hidden: the half dose most likely keeps the greater part of the protection, a central estimate of 53% stroke-risk reduction against the full dose’s 61%, but the evidence leaves room for much less, and only a trial that compares the doses directly can close the gap.

The drug and its two doses

Apixaban is a direct inhibitor of factor Xa, one of the enzymes in the clotting cascade. Compared with warfarin, the drug it largely displaced, it needs no INR

blood monitoring, has few food interactions, and is given at a fixed dose. The convenience has a flip side. Its half-life is about 12 hours, which is why it is taken twice daily and why missed doses matter: much of the anticoagulant protection is gone within a day of stopping (Bristol Myers Squibb and Pfizer, 2021).

For atrial fibrillation the standard dose is 5 mg twice daily. The label reduces this to 2.5 mg twice daily only when at least two of the three criteria above are met. One point of frequent confusion is worth settling early: 2.5 mg twice daily is also the correct labeled dose for two entirely different situations, extended treatment after a venous clot and clot prevention after hip or knee replacement (Bristol Myers Squibb and Pfizer, 2021). A 2.5 mg prescription is therefore not automatically an underdose. This paper is about the atrial fibrillation indication, where the two-of-three rule governs.

The scale involved makes the dosage question consequential. Eliquis was the largest single drug expense in Medicare Part D when CMS selected it for price negotiation, with roughly 3.7 million enrollees using it in the year ending May 2023 (CMS, 2023). The negotiated price, $231 for a 30-day supply, took effect on January 1, 2026 (Pfizer, 2025).

The benefit ledger

A patient deciding whether to take Eliquis is not choosing between Eliquis and warfarin. The choice is between taking an anticoagulant and not taking one. The trial that would answer that question directly, apixaban against placebo, has never been run and never will be. By the time apixaban was developed, anticoagulation was already the standard of care for atrial fibrillation at elevated stroke risk, and randomizing such patients to nothing was considered unethical. So the number the patient wants has to be assembled from a chain of older and newer comparisons. This section builds the chain explicitly.

The oldest link is warfarin against placebo, tested in six trials before anticoagulation became standard: adjusted-dose warfarin cut stroke by 64% (95% CI, 49 to 74) (Hart et al., 2007). Aspirin, tested against placebo in eight trials, managed about 22% (95% CI, 6 to 35) (Hart et al., 2007). Aspirin, in other words, sits much closer to nothing than to real anticoagulation.

The newest link is direct and randomized. AVERROES assigned 5,599 patients considered unsuitable for warfarin to apixaban or aspirin. Strokes and systemic embolisms fell from 3.7 to 1.6 per 100 patient-years, a hazard ratio of 0.45, while major bleeding was statistically indistinguishable, 1.4 versus 1.2 per 100 patient-years (Connolly et al., 2011). Two caveats apply. The trial was stopped early for benefit, and trials stopped early tend to overstate effects (Bassler et al., 2010). And the bleeding confidence interval, 0.74 to 1.75, is wide enough to hide a substantial increase.

Chaining the links gives the drug-versus-nothing estimate, at the price of an inference; what follows is arithmetic assembled from the trials above, not a trial result. If aspirin trims about 22% from the untreated rate, the 3.7 per 100 patient-years observed in the aspirin arm implies an untreated rate near 4.7. The apixaban rate of 1.6 against that baseline is a reduction of roughly two-thirds, about 30 strokes prevented per 1,000 patient-years, with no measurable addition to the 12 major bleeds per 1,000 patient-years the aspirin arm already had. The absolute benefit scales with baseline risk, which varies severalfold with age and comorbidity, so the same two-thirds reduction is worth far more to an 80-year-old with a prior stroke than to a 65-year-old with neither.

Warfarin enters the story for two reasons. First, ARISTOTLE, the pivotal trial, randomized 18,201 patients between apixaban and warfarin for a median of 1.8 years and showed that apixaban beats the drug that beat placebo: fewer strokes, less major bleeding, and lower mortality (Granger et al., 2011). Apixaban therefore inherits the placebo-controlled warfarin evidence with a margin to spare. Second, ARISTOTLE is where nearly all the evidence about the dose lives, which matters for the rest of this paper. Table 1 gives its results in absolute terms, as rates per 100 patient-years; a rate of 1.27 means that among 1,000 patients treated for a year, about 13 had the event.

Outcome	Apixaban (%/yr)	Warfarin (%/yr)	Hazard ratio (95% CI)
Stroke or systemic embolism	1.27	1.60	0.79 (0.66 to 0.95)
Ischemic or uncertain stroke	0.97	1.05	0.92 (0.74 to 1.13)
Hemorrhagic stroke	0.24	0.47	0.51 (0.35 to 0.75)
Major bleeding	2.13	3.09	0.69 (0.60 to 0.80)
Intracranial hemorrhage	0.33	0.80	0.42 (0.30 to 0.58)
Death from any cause	3.52	3.94	0.89 (0.80 to 0.99)

Table 1. ARISTOTLE outcomes, rates per 100 patient-years (Granger et al., 2011). Hazard ratios below 1 favor apixaban.

The marketing language for these results was a 21% reduction in stroke, a 31% reduction in major bleeding, and an 11% reduction in death. All three relative claims are accurate. The absolute differences are 0.33, 0.96, and 0.42 per 100 patient-years. Put another way, treating 1,000 patients with apixaban instead of warfarin for one year prevents about 3 strokes or embolisms, about 10 major bleeds, and about 4 deaths. Both framings describe the same trial; only one of them tells a patient how much the choice matters. The thin absolute margins are a fact about the comparator, warfarin, not about the drug; quoted without that context they can make the drug look either miraculous or marginal at will. The

general habit of quoting the relative number alone is the subject of an earlier paper in this series, *The Arithmetic of Persuasion*.

Table 1 contains a subtlety that the headline conceals. The stroke benefit over warfarin came almost entirely from hemorrhagic strokes, the kind caused by bleeding into the brain, which anticoagulants themselves provoke. Ischemic strokes, the kind the drug exists to prevent, were not significantly different between the arms. Against warfarin, then, apixaban mainly wins by causing less harm, not by preventing more clots. That is a real and valuable win. It is just a different claim than most patients hear.

The bleeding ledger

Anticoagulants prevent clots by making blood slower to clot everywhere, so the cost side of the ledger is bleeding. On apixaban in ARISTOTLE, major bleeding ran at 2.13 per 100 patient-years. That is lower than warfarin, not low. Among 1,000 patients on the drug for a year, about 21 had a major bleed.

The composition of the bleeding matters as much as the total. Apixaban roughly halved intracranial hemorrhage relative to warfarin, 0.33 versus 0.80 per 100 patient-years, and intracranial bleeds are the ones that kill and disable (Granger et al., 2011). Gastrointestinal bleeding, the most common site of major bleeding, showed no significant difference from warfarin (Hylek et al., 2014). So the drug shifted the mix of bleeds away from the worst kind while leaving the most common kind about the same.

Two practical notes complete the risk picture. First, reversal. Warfarin has a cheap, familiar antidote in vitamin K. Apixaban has a dedicated reversal agent, andexanet alfa, approved in 2018 under the FDA’s accelerated pathway, along with off-label use of prothrombin complex concentrates; the agent is expensive and not stocked everywhere (FDA, 2018). Second, the same short half-life that punishes missed doses helps in an emergency, since most of the anticoagulant effect clears within a day or two of stopping (Bristol Myers Squibb and Pfizer, 2021).

Where the 2.5 mg rule came from

The two-of-three rule was written into the ARISTOTLE protocol before the trial began, on pharmacokinetic grounds. Early-phase studies showed that drug exposure rises in the elderly, in the light, and in the renally impaired: roughly 32% higher exposure in older subjects, about 20 to 27% higher in subjects under 50 kg, and about 38% higher in severe renal impairment (Vu et al., 2021). Patients likely to combine these features were assigned half the dose to bring their expected blood levels back toward the middle of the range.

The clinical evidence behind the rule is thinner than its regulatory status suggests. Only 4.2% of properly dosed ARISTOTLE patients, around 760 people, met the criteria and received 2.5 mg (Zeitouni et al., 2020). Within that small group the drug performed well: relative to warfarin, the reductions in stroke, bleeding, and death were consistent with those in the standard-dose cohort, with no significant interaction, even though measured drug exposure was about a quarter lower. The investigators read this as evidence of a relatively wide therapeutic window.

A skeptical reader should notice three things about this rule. First, the criteria are sharp cutoffs laid over continuous variables. To take a deliberately constructed illustration, an 81-year-old weighing 59 kg gets half the dose of a 79-year-old weighing 61 kg, though nothing in their physiology changes at those boundaries. This series treats sharp cutoffs on smooth risk curves as a recurring fault in medical evidence, and the dose rule inherits it. Second, the rule itself has never been randomized. No trial has taken patients who meet the criteria and assigned them to 5 mg versus 2.5 mg, so the claim that the reduced dose is right for them rests on modeled blood levels plus a subgroup of a few hundred, compared against warfarin rather than against the other dose. Third, the exposure differences the rule corrects for, 20 to 40%, are about the size of the exposure difference the trial itself showed the drug tolerating without any detectable change in relative effect.

Figure 1 shows why the cutoff style matters more here than it might elsewhere. Kidney function sits on both sides of the ledger at once: as eGFR falls, the stroke rate the drug exists to prevent roughly doubles, and the major bleeding rate nearly triples (Hohnloser et al., 2012). Both rise smoothly across the whole range; there is no break anywhere for a threshold to latch onto. The label collapses that continuous, two-sided tradeoff into a single creatinine value, and serum creatinine is itself a rough proxy: the same 1.5 mg/dL corresponds to nearly normal kidney function in a heavyset 55-year-old man and to substantially impaired function in a small 80-year-old woman. In the trial data, apixaban held its advantage over warfarin in every renal band, and its bleeding advantage was largest where the kidneys were worst (Hohnloser et al., 2012).

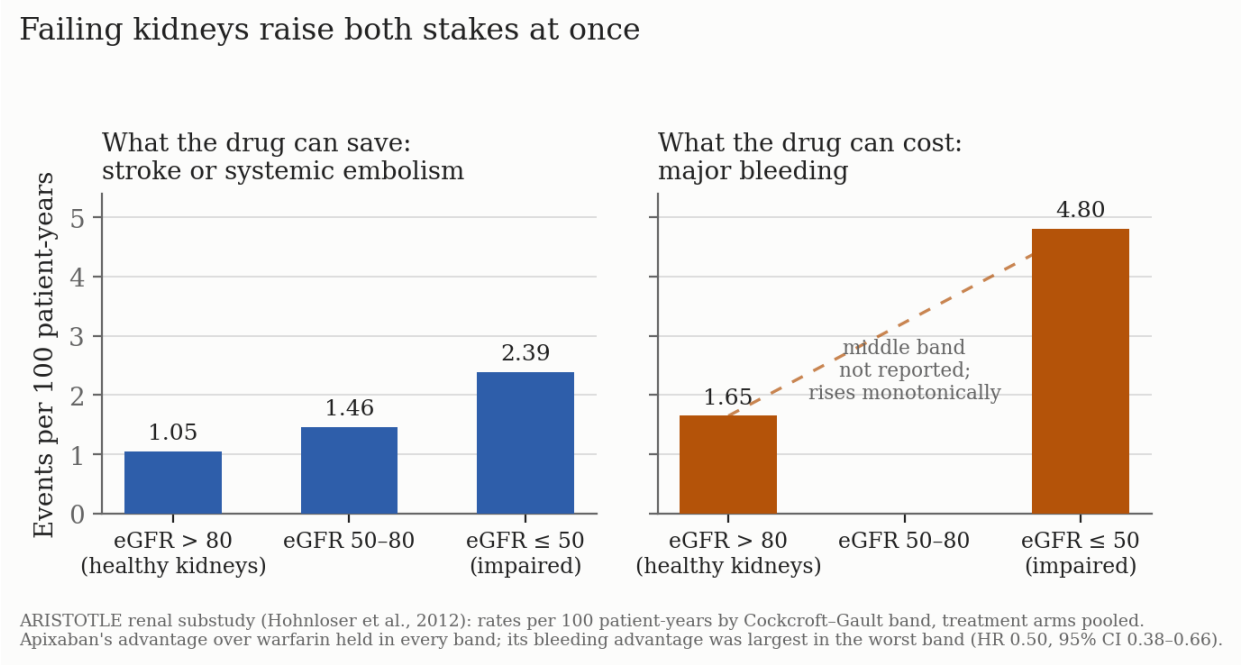

Figure 1. Both sides of the ledger rise as kidney function falls (Hohnloser et al., 2012). Rates per 100 patient-years, ARISTOTLE, treatment arms pooled; the middle bleeding band is not reported in the accessible text.

None of this means the rule is wrong. It means the rule is a reasonable engineering guess that has held up in limited data, which is a weaker warrant than most prescribing labels are assumed to carry.

Underdosing in practice

In the trial, about one patient in twenty qualified for the reduced dose. In practice, the reduced dose is used in 20 to 60% of apixaban-treated patients, depending on the setting (Zeitouni et al., 2020). The most direct recent accounting comes from a cohort of 45,947 apixaban users with atrial fibrillation drawn from US health systems: 15.4% were on the reduced dose, and 61% of those did not meet the label criteria (Gouda et al., 2026). A three-center study of adults 65 and older found 15% of the whole cohort underdosed (Campbell et al., 2024). The pattern is easy to understand. Bleeding is visible and feels attributable to the prescriber; the stroke that a full dose would have prevented is invisible. Faced with a frail patient near the cutoffs, many clinicians split the difference.

What does underdosing cost? Here the literature turns treacherous, and in an instructive way. In the US cohort, patients inappropriately given the low dose died at 1.77 times the adjusted rate of appropriately full-dosed patients. Taken causally, that is alarming. But the same analysis found no significant difference in ischemic stroke, with an adjusted hazard ratio of 1.05, and no significant difference in major bleeding or intracranial hemorrhage (Gouda et al., 2026). An anticoagulant dose

can only affect death through clotting and bleeding. A 77% mortality excess with neither more strokes nor less bleeding is not biologically coherent as a drug effect. It is exactly what confounding by indication looks like: clinicians reserve the low dose for sicker, frailer patients, and sicker patients die more of everything. Adjustment removes only the confounders that were measured. The older-adult study shows the same signature, no difference in stroke or bleeding but mortality of 10.9% versus 1.4% (Campbell et al., 2024). These studies establish that underdosing is common. They do not establish what it costs, in either direction, and the mortality figures should not be quoted as if they did.

The most interesting patients are those who meet exactly one criterion, since the label sends them to 5 mg while many prescribers hesitate. The ASPIRE study followed 1,944 such patients in Korea, about half on each dose, and found no significant differences at one year: stroke or embolism 0.7 versus 0.9 per 100 patient-years favoring the low dose, major bleeding 1.0 versus 0.5 disfavoring it, with confidence intervals running from 0.28 to 1.59 and 0.44 to 4.35 respectively (Lee et al., 2025). With 15 strokes and 15 major bleeds in total, this study mostly measured its own size. Finding no significant difference in a study this small is not evidence the doses are equivalent; the earlier paper in this series on the Bernoulli fallacy, *How Sure Is Sure Enough*, covers why. What would settle the question is a randomized trial of 5 mg versus 2.5 mg in these borderline patients. None has reported.

Study	Population	Finding on dosing
ARISTOTLE dose analysis (Zeitouni et al., 2020)	18,073 trial patients; 4.2% on 2.5 mg by criteria	Relative benefit vs warfarin consistent at both doses; exposure about 24% lower at 2.5 mg
US health systems cohort (Gouda et al., 2026)	45,947 apixaban users, 2015 to 2019	61% of reduced-dose patients did not meet criteria; mortality signal (HR 1.77) without stroke or bleeding differences, consistent with confounding
Three academic centers (Campbell et al., 2024)	286 matched patients, age 65+	No stroke or bleeding difference; mortality 10.9% vs 1.4%, same confounding signature
ASPIRE (Lee et al., 2025)	1,944 patients with exactly one criterion, Korea	No significant differences at 1 year; intervals too wide to show equivalence

Table 2. The evidence on reduced-dose apixaban outside the label criteria. Only the first row comes from randomized treatment assignment, and even there the dose itself was not randomized.

A Bayesian network for the dose

The preceding sections leave the dosing question in an unsatisfying state: the criteria are proxies, the observational literature is confounded, and the one unmeasured quantity, what 2.5 mg does in patients outside the criteria, is exactly

what a prescriber needs to know. A small Bayesian network organizes what is known and prices what is not. Its purpose is not prediction for an individual but bookkeeping: it makes every assumption explicit, computes what the published numbers jointly imply, and shows precisely which assumption the controversy turns on.

Structure

Figure 2 shows the network. Sex influences creatinine and weight; age, sex, weight, and creatinine determine kidney function through the Cockcroft–Gault formula; kidney function drives both outcomes; and the dose influences both outcomes as well. Two further facts about the patient enter as observations: how much of the time the heart is actually in fibrillation, which scales the untreated stroke hazard, and whether there has been a prior stroke or mini-stroke, which raises both hazards (Easton et al., 2012). Bleeding is split into the two kinds that differ most in consequence, bleeding into the brain and bleeding everywhere else. The dose node takes three values, 5 mg, 2.5 mg, or 0 mg, the last standing for no anticoagulation at all. In ordinary operation the dose follows a prescribing policy, the dashed arrows: either the label rule or prescribing as observed in practice. The queries that matter sever those arrows with the do-operator (Pearl, 2009). Setting the dose by intervention separates two questions the observational studies of the previous section conflate: who gets 2.5 mg, which is a fact about prescribing, and what 2.5 mg does, which is a fact about the drug.

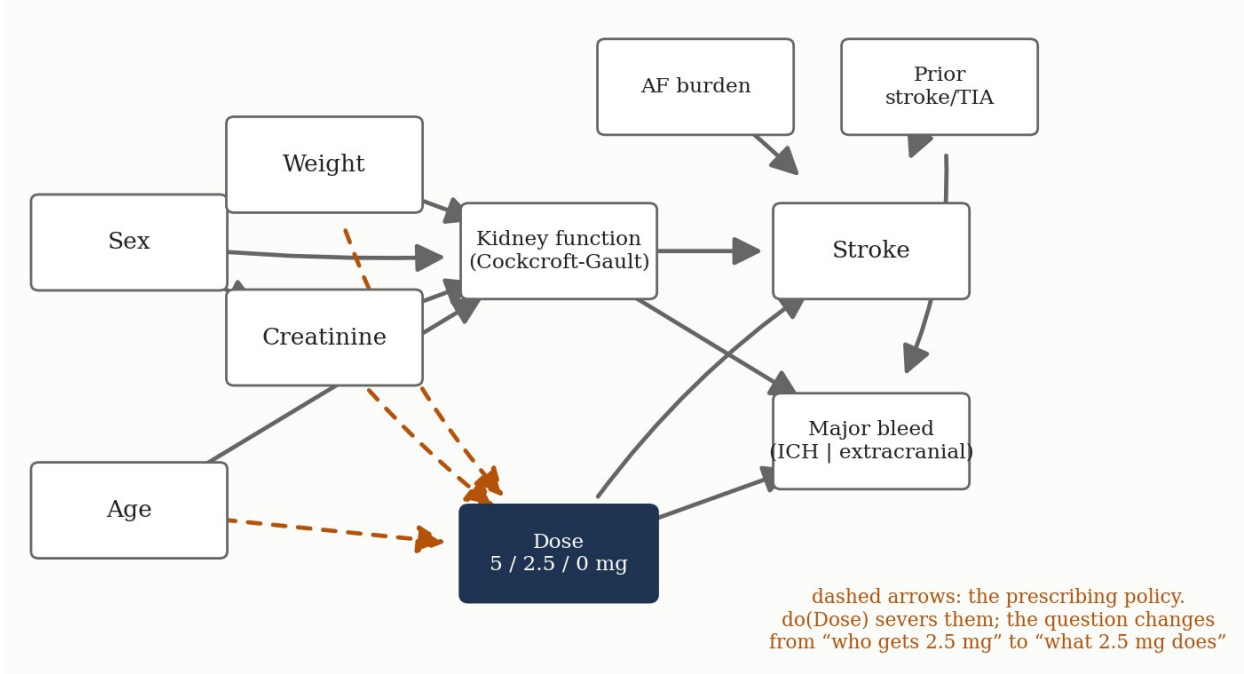

Figure 2. The network. The dose node is the intervention point; dashed arrows carry the prescribing policy and are severed under do(dose).

How the conditional tables were built

Creatinine given sex comes from the NHANES III population survey: means of 1.16 mg/dL for men and 0.96 for women, with the spread fitted to the reported fractions above 1.5 mg/dL, 9.7% of men and 1.8% of women (Jones et al., 1998). Weight is carried in six bands, in pounds, with the label's low-weight criterion falling at 132 lb and the priors set to match the trial prevalence of that criterion. Kidney function is deterministic Cockcroft–Gault. Stroke and bleeding rates by kidney function come from the ARISTOTLE renal bands (Hohnloser et al., 2012), and the effect of treatment comes from the chained estimate of the benefit ledger, about a 65% reduction, carried as a distribution rather than a constant because the chain has uncertainty of its own. The effect of 2.5 mg outside the criteria is the quantity no trial has measured, and an earlier draft of this model carried the two extreme readings as rival hypotheses: the reduced dose keeps the full effect, or effect simply follows exposure. Neither extreme deserves belief, so the model instead does what a Bayesian does with an unmeasured quantity: it learns a curve. Protection as a function of drug exposure is given a flexible shape and fitted to every published observation that bears on it, the trial subgroup that took 2.5 mg under the rule (Zeitouni et al., 2020), the large cohort in which off-label half-dosing showed no detectable stroke excess (Gouda et al., 2026), and, more loosely, the dose comparisons inside trials of the sister drugs dabigatran and edoxaban (Connolly et al., 2009; Giugliano et al., 2013). Those sources disagree, and the fitted curve carries their disagreement as a visible range instead of resolving it by assumption. Bleeding gets the same treatment with its own learned curve, and its excess is weighed by consequence: in the trial that compared apixaban with aspirin, bleeding into the brain was no more common on the drug, so the drug's added bleeding is almost entirely of the recoverable kind, and the model scores severity with published outcome rates (Fang et al., 2007).

Before asking it anything, the network passes two checks it was not fitted to. Under the label policy it assigns the reduced dose to 5.4% of the population; the trial observed 4 to 5% (Zeitouni et al., 2020). Under the real-world policy, derived from the finding that 61% of reduced doses are off-label (Gouda et al., 2026), it puts 13.5% of patients on 2.5 mg; the cohort observed 15.4%. The input distributions reproduce both numbers without being tuned to them.

What the model says

Figure 3 shows what the evidence, taken together, says about the curve. Protection climbs steeply with exposure and then levels off. At half exposure, which is where the off-label half dose puts a patient with healthy kidneys, the central estimate is a 53% reduction in stroke risk against no treatment, most of the full dose's 61%, but the band runs from 29% to 74%: probably most of the protection, possibly far less. The label's own arithmetic partly protects the patients the rule was written for, because in reduced kidney function the same 2.5 mg

produces more exposure and the retained protection rises toward 57%. The two old hypotheses appear on the figure as reference lines; what the evidence supports lies between them, nearer the optimistic one without reaching it.

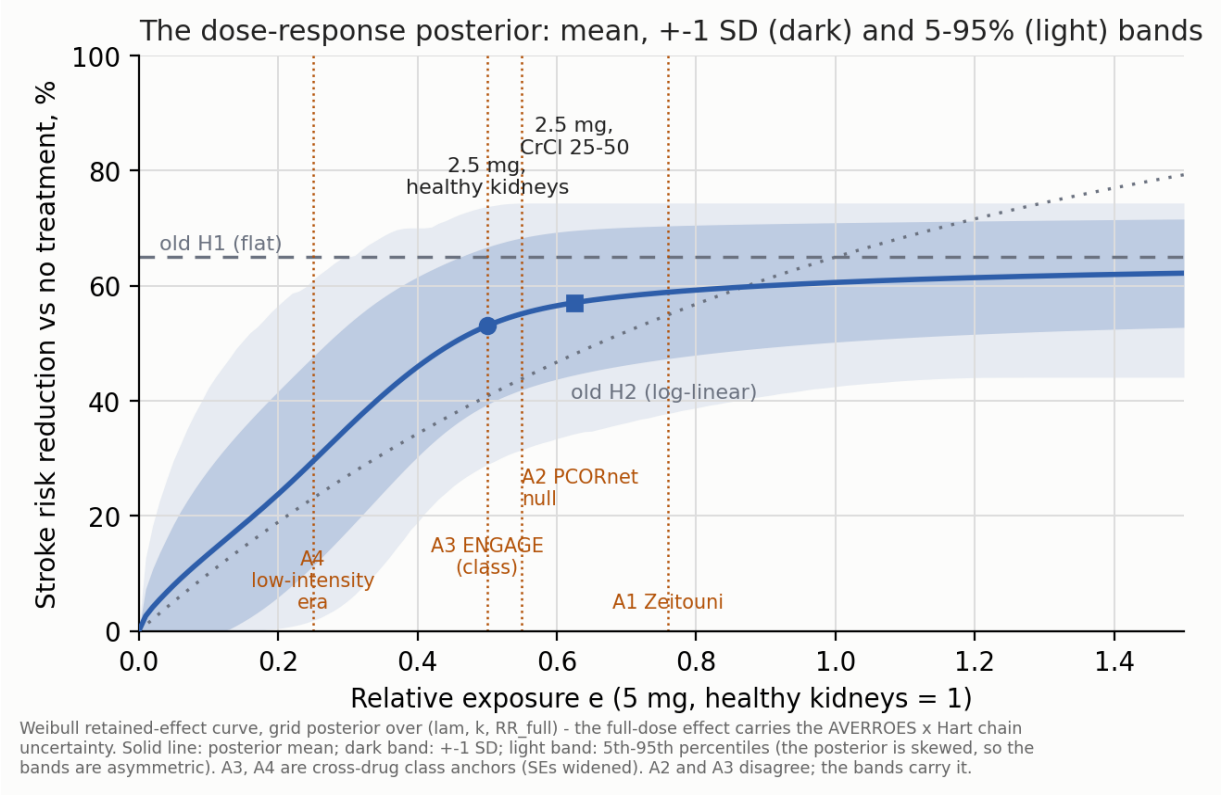

Figure 3. The learned dose-response curve: stroke-risk reduction against exposure, with the posterior mean, a one-standard-deviation band (dark), and the 5th to 95th percentile band (light). The two dose markers show 2.5 mg with healthy and with reduced kidneys; the old rival hypotheses are the flat and sloping reference lines.

Figure 4 turns the curve into people, for the paper's running example of a patient over 80 with reduced kidney function. Out of 1,000 such patients in a year, about 21 have a stroke on the full dose, about 23 on the half dose, and about 58 on no anticoagulant; major bleeds run 39, 31, and 25. Weigh the events by consequence and the picture sharpens. Bleeding into the brain sits near 6 per 1,000 at every dose including none, the drug's added bleeding is the kind most people recover from, and the strokes it prevents kill or permanently disable about half their victims, so counting only events that kill or disable, the full dose gives about 16 per 1,000 against 36 untreated. The dominant fact survives everything the model has learned: the expensive choice is not the wrong dose but no dose.

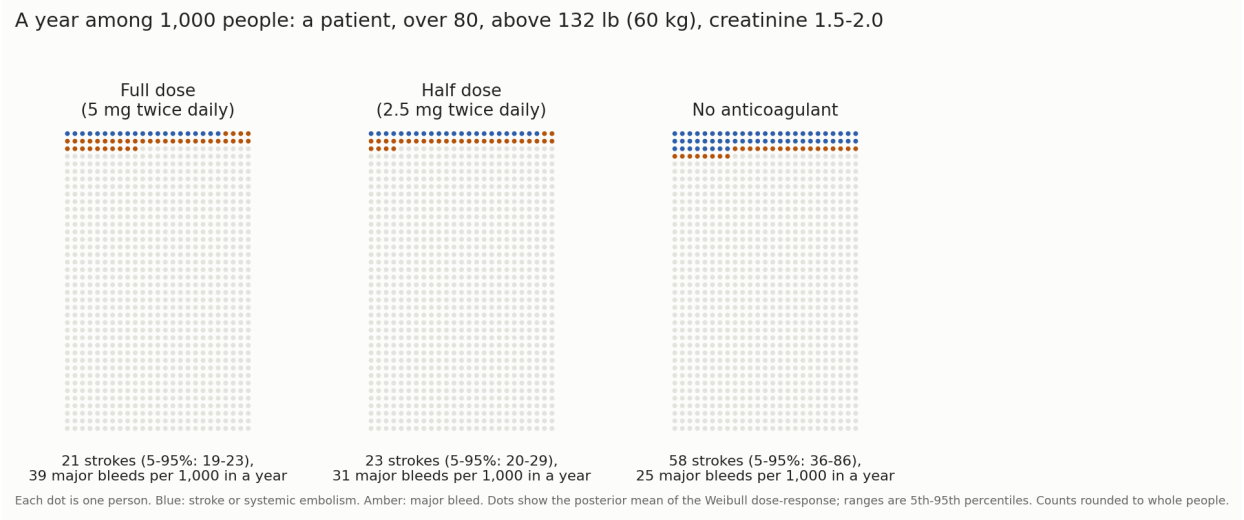

Figure 4. A year among 1,000 patients like the running example, under each choice. Each dot is a person: blue, a stroke; amber, a major bleed. Ranges in the captions are the 5th to 95th percentiles.

At the population scale the network prices the prescribing policies directly. Setting every patient in the modeled population to 5 mg gives an annual stroke rate of 1.42%, with a band of 1.28 to 1.54 from the chain's uncertainty. The label rule gives 1.44%, and prescribing as observed in practice, with its 61% off-label reductions, gives 1.45%: two or three added strokes per 10,000 patient-years, real but far too small for any observational study to detect, which is exactly why the cohort null could not settle the question. Universal half-dosing, which nobody proposes, would cost about 2 per 1,000. The observational mortality claims of the previous section are not needed to see the stakes; the network prices them from the trial numbers, and it prices the uncertainty along with them.

What the model is not

The network is assembled, not fitted. Every table traces to a published number, and the joints between them, above all the chained relative effect, carry the uncertainty of the chain. It knows nothing about time: probabilities are annual, and a patient who has a stroke enters no second year. It stops at creatinine clearance 25, where the trial itself stopped, and its lowest band leans on extrapolation. It contains no confounding because effects are imposed rather than estimated, which is both its virtue and its limit: it computes what the published evidence implies, and its central curve remains a synthesis rather than a measurement, which only a randomized comparison of the doses could replace. It also leaves out what a medical chart contains: interacting medicines, falls, bleeding history. What it contributes is the shape of the decision, with honest widths: a towering cost for taking nothing, a bleeding price that is real but mostly recoverable, and a half-dose question that is genuinely open by exactly the amount the bands in Figure 3 show.

Reading the ledger

The evidence supports a few plain conclusions. Measured against the alternative a patient actually faces, taking nothing, apixaban at its labeled doses is one of the better-supported preventive drugs in medicine: a reduction in stroke of roughly two-thirds, worth on the order of 30 events per 1,000 patient-years at trial-typical baseline risk, with the drug's own bleeding rate of about 21 major bleeds per 1,000 patient-years as the cost. The comparison with warfarin is a prescriber's question, and its thin margins are not the measure of the drug. Most of the live uncertainty is not about the molecule but about the prescribing, and it concentrates in one place: the 2.5 mg dose given to people who do not meet the label criteria, which describes most of the people receiving it.

For a patient or a family member, that suggests a short list of questions worth asking. Which dose is being prescribed, and if it is 2.5 mg, which two of the three criteria were met? If the answer is one criterion or none, what is the specific reasoning, and has the alternative of the full dose been weighed against evidence rather than against unease? When was kidney function last checked, since creatinine is one of the criteria and can drift? And do any other medicines interact: strong inhibitors of CYP3A4 and P-glycoprotein, such as ketoconazole or ritonavir, raise apixaban levels enough that the label halves the dose or advises avoidance (Bristol Myers Squibb and Pfizer, 2021).

The model itself travels with this paper. A locked copy of the spreadsheet that computes these numbers for any patient is on the paper's page at healthy-skepticism.com, built for a patient and cardiologist to run together in a minute during a visit, alongside a short guide for patients and a technical companion paper for clinicians.

One economic point belongs on the list because it interacts with the pharmacology. The drug's protection fades within about a day of a missed dose, so stretching a supply by skipping doses converts an expensive effective drug into an expensive ineffective one. For Medicare patients the negotiated $231 monthly price that took effect in January 2026 should ease that pressure (Pfizer, 2025); anyone still rationing doses should raise cost with the prescriber rather than solve it at the pillbox.

This paper is information, not medical advice. Dosing decisions belong with the prescriber, who knows the particulars. The purpose here is the same as elsewhere in this series: to show where the evidence is strong, where it is thin, and how to tell the difference when the numbers are quoted at you.

References

Bassler, D., Briel, M., Montori, V. M., et al. (2010). Stopping randomized trials early for benefit and estimation of treatment effects: systematic review and meta-regression analysis. JAMA, 303(12), 1180–1187.

Bristol Myers Squibb and Pfizer (2021). Eliquis (apixaban) prescribing information. U.S. Food and Drug Administration label.

Campbell, A. M., Pae, E., Lee, E., et al. (2024). Off-label reduced dose apixaban in older adults with atrial fibrillation and associated outcomes. Annals of Pharmacotherapy, 58(6). doi:10.1177/10600280231199137.

CMS (2023). Medicare drug price negotiation: selected drug list for initial price applicability year 2026. Centers for Medicare & Medicaid Services.

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.

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.

FDA (2018). FDA approves Andexxa (andexanet alfa), the first reversal agent for factor Xa inhibitors. U.S. Food and Drug Administration.

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.

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.

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.

Hylek, E. M., Held, C., Alexander, J. H., et al. (2014). Major bleeding in patients with atrial fibrillation receiving apixaban or warfarin: the ARISTOTLE trial. Journal of the American College of Cardiology, 63(20), 2141–2147.

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.

Pfizer (2025). Pfizer statement on publication of Eliquis maximum fair price. Pfizer Inc.

Vu, A., Qu, T. T., Ryu, R., et al. (2021). Critical analysis of apixaban dose adjustment criteria. Clinical and Applied Thrombosis/Hemostasis, 27. doi:10.1177/10760296211021158.

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.