Behavioral Valuation of Benefits and Harms in Cardiovascular Prevention: A Trial-Based Cumulative Prospect Theory Sensitivity Analysis

Abstract

Aims: To examine how probability weighting, reference points, relative outcome burdens, and assumptions about co-occurring events affect the valuation of cardiovascular prevention benefits and harms.

Methods: A secondary decision-modeling study used published aggregate counts from six randomized comparisons covering aspirin, rosuvastatin, blood-pressure lowering, and a polypill plus aspirin. Each comparison retained its own population, endpoints, and follow-up. Expected normalized endpoint burden was compared with cumulative prospect theory using complete outcome distributions. Outcomes were not assigned empirical disability-adjusted life-year weights. Severity thresholds, alternative reference points, and feasible dependence scenarios were examined.

Findings: In the aspirin comparisons, crude vascular-event reductions were 10.98, 2.40, and 1.87 per 1000 participants in ASCEND, ASPREE, and ARRIVE, respectively; corresponding bleeding increases were 8.91, 10.26, and 5.11. Bleeding definitions differed. Under independence and additive burdens, changing probability weighting from linear to a Prelec parameter of 0.7 shifted the harm-severity thresholds from 1.232 to 1.029, 0.234 to 0.222, and 0.366 to 0.249. The direction of threshold change was not universal across the other comparisons. Dependence assumptions could reverse selected behavioral valuations. With a healthy-state reference, a common loss-aversion coefficient changed scale but not sign.

Conclusion: The model identifies assumption-sensitive benefit–harm trade-offs within selected randomized comparisons. The illustrative scenarios do not estimate patient preferences, clinical net benefit, or public-health priorities.

Keywords: Cardiovascular Diseases; Primary Prevention; Decision Making; Health Communication

In Press

Introduction

Cardiovascular prevention involves comparing uncertain future benefits with possible treatment-related harms. For health education, communicating a relative reduction alone is insufficient: the starting risk, absolute difference, adverse outcomes, and observation period also matter. Statistical literacy research and risk-communication reviews identify reasons to present these quantities transparently and on comparable denominators [1-3]. Nevertheless, a transparent presentation does not imply that every person assigns the same importance to each outcome. Clinical evidence and the valuation of that evidence therefore require separate treatment.

Reviews of patient and clinician expectations have documented inaccurate assessments of intervention benefits and harms [4,5]. Risk-perception research also distinguishes numerical probability from the meaning and salience attached to an adverse event [6]. These observations motivate behavioral analysis, but do not establish that a particular functional form reproduces preferences in a specific population. A numerical behavioral score should not be interpreted as an observed preference, an adherence probability, or a measure of the effectiveness of health education.

Prospect theory introduced reference-dependent valuation, diminishing sensitivity, loss aversion, and nonlinear probability weighting [7,8]. Cumulative prospect theory extends this approach by constructing decision weights from ordered cumulative probabilities [9]. Prelec's weighting function provides a convenient parametric specification [10]. Health-domain preference studies support investigating such mechanisms, while also showing why parameters elicited for particular health tasks require careful interpretation [11]. Values estimated for life duration cannot automatically be treated as coefficients for a normalized composite cardiovascular endpoint.

Several established decision approaches address related but different questions. Multiple-criteria decision analysis makes competing criteria and their weights explicit [12]. Expected-burden calculations summarize outcomes under linear probability weighting. Cumulative prospect theory instead examines how reference dependence and probability weighting alter the evaluation of an outcome distribution. None of these approaches supplies missing clinical evidence, identifies the correct severity of a composite endpoint, or establishes the budget priority of an intervention without the additional inputs required for that decision.

Aspirin provides a useful test case because major primary-prevention trials report both vascular events and bleeding in different populations [13-15]. Rosuvastatin, blood-pressure lowering, and a polypill with aspirin extend the analysis to other prevention strategies and tolerability outcomes [16-18]. The corrected aspirin meta-analysis further illustrates the need to check endpoint-specific effects and publication updates before modeling [19,20]. Unexpected behavioral rankings could reflect mathematical implementation, outcome definitions, or untested joint-event assumptions, rather than measured psychological bias.

This study examined the conditions under which trial-based benefit-harm valuations change sign under linear and cumulative prospect theory specifications. Its contribution is an explicit separation of observed marginal event proportions, assumed relative burdens, and assumed behavioral mechanisms. The purpose was to identify sensitive trade-offs for risk-communication and preference research, not to construct a national ranking or estimate an economic return.

Materials & Methods

This secondary decision-modeling analysis used published aggregate trial counts. Model construction and uncertainty reporting were informed by guidance on conceptual clarity and the distinction between parameter and structural uncertainty [21,22]. Six comparisons from five trials were selected purposively to provide documented benefit-harm examples across four drug-based prevention strategies. This was not a systematic review, a network meta-analysis, or an estimate of the effectiveness of all available prevention interventions. The comparisons were retained as separate cases throughout.

Evidence selection and outcome extraction

ASCEND enrolled people with diabetes without evident cardiovascular disease; ASPREE enrolled initially healthy older adults; and ARRIVE enrolled people with moderate estimated cardiovascular risk while excluding diabetes and high bleeding risk [13-15]. HOPE-3 contributed separate factorial comparisons of rosuvastatin and candesartan plus hydrochlorothiazide [16,17]. TIPS-3 contributed its polypill-plus-aspirin cell against double placebo [18]. The HOPE-3 comparisons overlap within one trial and were not treated as independent replications. Table 1 specifies the populations, denominators, and observation periods.

In Press

Table 1. Trial populations and randomized comparisons used in the scenario analysis

Trial / comparison	Population	Treatment / control n	Follow-up, years	Selected safety endpoint
ASCEND [13]	Diabetes; no evident CVD	7740 / 7740	7.4 mean	Major bleeding
ASPREE [14]	Initially healthy older adults	9525 / 9589	4.7 median	Major hemorrhage
ARRIVE [15]	Moderate estimated risk; no diabetes	6270 / 6276	5.0 median	GI bleeding, mostly mild
HOPE-3 statin [16]	Intermediate risk; no CVD	6361 / 6344	5.6 median	Muscle symptoms; cataract surgery
HOPE-3 BP [17]	Intermediate risk; no CVD	6356 / 6349	5.6 median	Symptom-related permanent discontinuation
TIPS-3 combination [18]	Intermediate risk; no established CVD	1429 / 1421	4.6 mean	Discontinuation for dizziness or hypotension

CVD: cardiovascular disease; GI: gastrointestinal; BP: blood pressure. HOPE-3 comparisons are factorial marginal comparisons from one trial, not independent trials. TIPS-3 uses the polypill-plus-aspirin cell versus double placebo. Endpoints and follow-up are not harmonized across studies. Source: original reports [13-18].

The six comparisons were selected as worked examples with randomized arm denominators and published counts for both a vascular composite and at least one selected safety outcome. The three aspirin trials contrast populations and bleeding definitions; the other three comparisons extend the examples to a statin, a blood-pressure regimen and a combination regimen, including one case with two modeled harms. This scope illustrates these structures without implying representative coverage of primary-prevention evidence. Counts, denominators, endpoint definitions and follow-up summaries were transcribed from the primary reports [13-18] into a comparison-level input table. The input records document extraction on 29 October 2026; this is the extraction date, not a database search cutoff. Table 2 makes the extracted counts auditable. The corrected aspirin meta-analysis and its correction [19,20] provide context only; neither supplied modeled probabilities. The selection was purposive and was not based on a systematic database search. No formal database search window is reported.

Safety outcomes were retained as reported: major bleeding or hemorrhage, gastrointestinal bleeding of any severity, muscle pain or weakness, cataract surgery, or treatment discontinuation due to specified symptoms. These are distinct outcomes; omitted events were outside the model. The cataract surgery count was included as a trial-specific safety signal.

Observed differences and normalized burdens

Within each arm, a crude event proportion was calculated as the number of participants experiencing an endpoint divided by the number randomized to that arm. Each comparison retained its original follow-up period, without conversion to a common ten-year horizon. These proportions do not account for censoring and are not substitutes for the published time-to-event estimates. Descriptive intervals for between-arm risk differences were calculated using the Newcombe method based on Wilson intervals [23]. No joint confidence interval for net burden was calculated because within-person covariance between endpoints was unavailable.

Let p_T and p_C denote vascular-event proportions in the treatment and control arms, respectively, and let q_T and q_C denote the corresponding proportions for a selected harm. Define the absolute vascular-event reduction as

$$A = p_C - p_T,$$

and the absolute increase in the selected harm as

$$B = q_T - q_C.$$

A vascular composite event was assigned one normalized burden unit, whereas the selected harm was assigned $h \geq 0$ units. The additive expected-burden difference, defined as expected control burden minus expected treatment burden, was

$$G(h) = (p_C + hq_C) - (p_T + hq_T) = A - hB.$$

When $A > 0$ and $B > 0$, the zero threshold was

$$h^* = \frac{A}{B}.$$

Thus, $G(h) > 0$ when $h < h^*$, and $G(h) < 0$ when $h > h^*$. This threshold expresses a conditional trade-off: h is neither an empirical disability-adjusted life-year weight nor an elicited patient preference.

For the rosuvastatin comparison with two modeled harms, the additive expected-burden difference was

$$G(h_m, h_c) = A - h_m B_m - h_c B_c,$$

where m denotes muscle symptoms and c denotes cataract surgery. Here, h_m and h_c are their respective normalized burdens, and

$$B_m = q_{T,m} - q_{C,m}, \quad B_c = q_{T,c} - q_{C,c}$$

are their respective absolute increases under treatment. The zero boundary was therefore

$$h_m B_m + h_c B_c = A,$$

which defines a line rather than a single threshold. Normalizing heterogeneous composite endpoints to one unit does not make those units equivalent across trials. Consequently, thresholds and scores were interpreted within each comparison, without combining them into an intervention league table or a monetary valuation.

Complete outcome distributions and cumulative valuation

For one vascular endpoint and one harm, let p and q denote their respective within-arm probabilities, and let c denote the probability that both occur. The four mutually exclusive states—neither endpoint, the vascular endpoint only, the harm only, and both endpoints—had probabilities

$$1 - p - q + c, \quad p - c, \quad q - c, \quad c,$$

respectively. A valid joint distribution required

$$\max(0, p + q - 1) \leq c \leq \min(p, q).$$

Under additive burdens, the corresponding state burdens were

$$0, \quad 1, \quad h, \quad 1 + h,$$

where one vascular event was assigned one normalized burden unit and the selected harm was assigned $h \geq 0$ units. An alternative maximum-burden rule assigned

$$\max(1, h)$$

to the joint state, while leaving the other state burdens unchanged. This alternative examined potential overlap in clinical consequences; neither rule was presumed to reconstruct actual lifetime disease burden.

With two harms, the model used the eight combinations of three binary endpoints. Independence supplied a transparent reference distribution, not an empirical claim. For a normalized burden b, the health outcome was defined as

$$y = -b,$$

and its deviation from a common reference point r was

$$x = y - r = -b - r.$$

The primary reference was

$$r = 0,$$

representing the absence of modeled endpoint burden. An alternative reference was the negative expected control burden under the reference distribution:

$$r = -\mathbb{E}_{C,ref}[b] = -\sum_k p_{Ck}^{ref} b_k,$$

where p_{Ck}^{ref} denotes the probability of state k in the control arm's reference distribution and b_k denotes its burden. The reference point was common to the treatment and control arms and was held fixed when testing alternative dependence structures within a specified scenario.

The value function was defined as

$$v(x) = \begin{cases} x^{\alpha}, & x \geq 0 \\ -\lambda(-x)^{\beta}, & x < 0 \end{cases}$$

where x denotes the deviation of a health outcome from the common reference point, α and β govern curvature in the gain and loss domains, respectively, and $\lambda > 0$ is the loss-aversion coefficient. Probability weighting followed the Prelec function,

$$w(p) = exp[-(-ln\, p)^{\gamma}], \quad 0 < p < 1, \quad \gamma > 0,$$

with $w(0) = 0$ and $w(1) = 1$ [9,10]. The notation w^- and w^+ distinguishes weighting in the loss and gain domains; both used the same γ in the reported scenarios. Outcomes were ordered from worst

In Press

to best. Loss decision weights were calculated as differences between weighted cumulative probabilities accumulated from the worst outcome, whereas gain decision weights were calculated using cumulative probabilities accumulated from the best outcome. Each arm's value was

$$V = \sum_{i=1}^{n} \pi_i v(x_i),$$

and the comparator-adjusted valuation was

$$S = V_T - V_C,$$

where V_T and V_C denote treatment and control values, respectively. Both arms were evaluated relative to the same reference point. Absolute risk reductions were not themselves treated as complete lotteries; valuation used each arm's complete modeled outcome distribution.

For reproducibility, let p_i denote the probability of outcome x_i, with $p_i \geq 0$ and $\sum_{i=1}^{n} p_i = 1$. For ordered outcomes satisfying

$$x_1 \leq \cdots \leq x_m < 0 \leq x_{m+1} \leq \cdots \leq x_n,$$

the decision weight for a loss was

$$\pi_i^- = w^-\left(\sum_{j=1}^{i} p_j\right) - w^-\left(\sum_{j=1}^{i-1} p_j\right), \quad i = 1, \ldots, m,$$

and the decision weight for a gain or zero outcome was

$$\pi_i^+ = w^+\left(\sum_{j=\mathrm{i}}^{n} p_j\right) - w^+\left(\sum_{j=\mathrm{i}+1}^{n} p_j\right), \quad i = m+1, \ldots, n.$$

Accordingly, each arm's value could also be written as

$$V = -\lambda \sum_{i=1}^{m} \pi_i^- (-x_i)^{\beta} + \sum_{i=m+1}^{n} \pi_i^+ x_i^{\alpha}.$$

Empty sums were defined as zero, and equal-valued outcomes were aggregated by summing their probabilities before weighting. Outcomes at the reference point contributed zero because $v(0) = 0$. This procedure preserves the distinction between cumulative decision weights and separately transformed event risks, while making outcome ordering and the treatment of ties and zero outcomes explicit and reproducible.

The reference analysis compared linear weighting (γ=1) with illustrative nonlinear weighting (γ=0.7), using independence, additive burdens, a healthy-state reference, and linear outcome values (α=β=1). These settings isolate the effect of probability weighting. Under this reference, a common positive λ rescales both arm values without changing the sign of their difference. Alternative reference points, curvature, joint burden rules and dependence were exploratory. No preference data were collected or used to calibrate the parameters.

Sensitivity analysis and verification

The aspirin sensitivity analyses evaluated five relative harm burdens, five probability-weighting parameters, three loss-aversion coefficients, two reference points, linear or curved outcome-value functions, and additive or maximum-burden rules for joint events. The parameter grids were

$$h \in \{0.1;\ 0.25;\ 0.5;\ 1;\ 2\},$$

$$\gamma \in \{0.5;\ 0.7;\ 0.85;\ 1;\ 1.15\},$$

and

$$\lambda \in \{1;\ 1.6;\ 2.1\}.$$

These specifications produced 600 combinations per aspirin comparison and 1800 combinations across the three comparisons. They were scenario evaluations, not independent observations. Analyses of the other interventions added 1260 combinations, including a two-harm burden grid for rosuvastatin. The extension used relative burdens drawn from

$$h \in \{0;\ 0.01;\ 0.05;\ 0.1;\ 0.25;\ 0.5;\ 1\},$$

with λ=1.6. The extension retained linear outcome values (α=β=1); curved values (α=0.95 and β=0.8) were examined only in the aspirin analyses. Of the 1260 extension scenarios, 980 involved rosuvastatin and 140 each involved blood-pressure lowering and the polypill combination. Separate threshold calculations examined each selected harm individually; these diagnostic calculations omitted other harms and were labeled accordingly. Threshold curves varied γ from 0.5 to 1.2 at 36 equally spaced points:

In Press

$$\gamma_k = 0.5 + \frac{k-1}{35}(1.2 - 0.5), \quad k = 1, \ldots, 36.$$

The γ grid spans nonlinear weighting on both sides of the linear case; the λ grid includes no loss aversion (1) and larger multipliers (1.6 and 2.1) to illustrate sensitivity when both gains and losses can occur. The curved α and β specification contrasts diminishing sensitivity with linear outcome values. These choices are motivated by the structure of CPT and the task dependence of health-domain estimates [9-11], not by estimated parameter distributions. Harm grids include burdens below, equal to, and (for aspirin) above the normalized vascular event burden; they test possible trade-offs rather than assert clinical severity equivalence. The alternative control reference tests whether the conclusion changes when the benchmark is the expected burden of the comparator rather than absence of modeled events.

For aspirin, the joint probability was evaluated at five equally spaced points spanning its feasible interval and at independence. For rosuvastatin, the vascular endpoint, muscle symptoms and cataract surgery generate eight binary states in each arm. Under independence, a state's probability is the product of the corresponding event or non-event probabilities. In the common-uniform alternatives, let U be uniform on [0,1]. An endpoint with marginal probability p occurs when U<p for a co-directed indicator or U>1−p for an oppositely directed indicator. The vascular indicator is co-directed; each harm can take either direction. Splitting [0,1] at every indicator cutoff assigns each interval to its unique state, with interval length as probability. The same construction applies to the other extension cases with one harm. These constructions preserve the published margins and deliberately stress-test overlap; they are not clinical estimates of co-occurrence.

Treatment and control dependence patterns were allowed to vary separately to examine the consequences of unidentified within-arm overlap. For the tested sets of treatment and control distributions, the sensitivity envelope was calculated as follows.

$$S_{min} = \min_{d_T \in \mathcal{D}_T} V_T(d_T) - \max_{d_C \in \mathcal{D}_C} V_C(d_C),$$

and

$$S_{max} = \max_{d_T \in \mathcal{D}_T} V_T(d_T) - \min_{d_C \in \mathcal{D}_C} V_C(d_C).$$

The reference point and other scenario parameters were held fixed during these comparisons. The tested families did not exhaust all possible joint distributions; consequently, their envelopes were neither sharp identification bounds nor statistical confidence intervals.

Implementation checks verified nonnegative state probabilities summing to one,

$$p_i \geq 0, \quad \sum_{i=1}^{n} p_i = 1,$$

preservation of marginal endpoint probabilities, and equivalence between cumulative valuation and the expected-burden difference calculated using the same state-burden rule. Under fully linear specifications,

$$\alpha = \beta = \gamma = \lambda = 1.$$

the implementation was required to reproduce the difference in expected control and treatment burdens:

$$S = E_C[b] - E_T[b].$$

Under the additive rule, this reduced to $G(h) = A - hB$, or $G(h_m, h_c) = A - h_m B_m - h_c B_c$ for two harms. Under the maximum-burden rule, expected burdens were calculated from the corresponding joint-state burdens rather than from the additive formula.

Additional checks compared analytical and numerical thresholds, verified that additive expected burdens were invariant to dependence when marginal probabilities were preserved, and confirmed loss-aversion scaling under the healthy-state reference:

$$S(\lambda) = \lambda S(1).$$

Numerical tests of first-order stochastic dominance and invariance to splitting an outcome into equal-valued states were also retained. Reporting distinguishes verification of mathematical implementation from clinical validation, consistent with modeling transparency guidance [24]. Supplementary Tables S1-S2 and the accompanying reproducibility archive provide selected scenario results, sourced event counts, code, figure coordinates, and model assumptions.

Findings

The principal finding was that nonlinear probability weighting and assumed co-occurrence could change the modeled benefit–harm trade-off within a trial. The direction of the change varied by

comparison and selected harm. Table 2 reports the underlying crude event differences and their descriptive intervals.

Table 2. Arm-level event counts, crude risks, and differences by outcome type

Panel	Comparison	Vascular endpoint[1]	Treatment n/N (%)	Control n/N (%)	Absolute reduction per 1000 (95% CI)[2]
A. Vascular outcomes	ASCEND	Vascular death, MI, stroke or TIA; excludes confirmed intracranial hemorrhage	658/7740 (8.50%)	743/7740 (9.60%)	10.98 [1.94, 20.03]
	ASPREE	Fatal coronary disease, nonfatal MI, fatal/nonfatal stroke, or hospitalization for heart failure	448/9525 (4.70%)	474/9589 (4.94%)	2.40 [-3.69, 8.48]
	ARRIVE	CV death, MI, unstable angina, stroke, or TIA	269/6270 (4.29%)	281/6276 (4.48%)	1.87 [-5.31, 9.06]
	HOPE-3 statin	CV death, nonfatal MI or nonfatal stroke	235/6361 (3.69%)	304/6344 (4.79%)	10.98 [3.97, 18.02]
	HOPE-3 BP	CV death, nonfatal MI or nonfatal stroke	260/6356 (4.09%)	279/6349 (4.39%)	3.04 [-3.99, 10.07]
	TIPS-3 combination	CV death, MI, stroke, resuscitated cardiac arrest, heart failure or revascularization	59/1429 (4.13%)	83/1421 (5.84%)	17.12 [1.10, 33.37]
B. Selected safety outcomes	ASCEND	Major bleeding	314/7740 (4.06%)	245/7740 (3.17%)	8.91 [3.04, 14.82]
	ASPREE	Major hemorrhage	361/9525 (3.79%)	265/9589 (2.76%)	10.26 [5.22, 15.34]
	ARRIVE	Any GI bleeding (mostly mild)	61/6270 (0.97%)	29/6276 (0.46%)	5.11 [2.17, 8.19]
	HOPE-3 statin	Muscle pain or weakness	367/6361 (5.77%)	296/6344 (4.67%)	11.04 [3.30, 18.80]
	HOPE-3 statin	Cataract surgery	241/6361 (3.79%)	194/6344 (3.06%)	7.31 [0.98, 13.67]
	HOPE-3 BP	Permanent discontinuation for symptomatic hypotension, dizziness or light-headedness	217/6356 (3.41%)	130/6349 (2.05%)	13.67 [8.02, 19.40]
	TIPS-3 combination	Discontinuation for dizziness or hypotension	45/1429 (3.15%)	22/1421 (1.55%)	16.01 [4.89, 27.64]

Notes: n = participants experiencing the endpoint; N = participants randomized to the corresponding arm; CI = confidence interval; CV = cardiovascular; MI = myocardial infarction; TIA = transient ischemic attack.

[1] Vascular endpoint definitions: ASCEND: vascular death, MI, stroke, or TIA, excluding confirmed intracranial hemorrhage. ASPREE: fatal coronary disease, nonfatal MI, fatal or nonfatal stroke, or hospitalization for heart failure. ARRIVE: CV death, MI, unstable angina, stroke, or TIA. Both HOPE-3 comparisons: CV death, nonfatal MI, or nonfatal stroke. TIPS-3: CV death, MI, stroke, resuscitated cardiac arrest, heart failure, or revascularization.

[2] Absolute reductions in Panel A were calculated as control risk minus treatment risk; absolute increases in Panel B were calculated as treatment risk minus control risk. Estimates and descriptive 95% CIs were calculated from crude proportions using the Newcombe–Wilson method [23], without adjustment for censoring. Each comparison retained its original follow-up: ASCEND, mean 7.4 years; ASPREE, median 4.7 years; ARRIVE, median 5.0 years; HOPE-3, median 5.6 years; and TIPS-3, mean 4.6 years.

Under independence, additive burdens, a healthy-state reference, and linear outcome values, changing γ from 1 to 0.7 moved the severity thresholds from 1.232 to 1.029 in ASCEND, 0.234 to 0.222 in ASPREE, and 0.366 to 0.249 in ARRIVE. Figure 1 shows the continuous scenario relationships. Below a threshold the treatment had a positive selected-outcome valuation; above it the valuation was negative. These are point-estimate thresholds conditional on hypothetical relative burdens, not clinical treatment cutoffs.

The direction of change was not universal. In the single-muscle-outcome rosuvastatin calculation, the threshold increased from 0.994 to 1.032, whereas the cataract-only threshold decreased from 1.502 to 1.405. The blood-pressure symptom-discontinuation threshold decreased from 0.222 to 0.210, and the polypill-combination threshold from 1.070 to 0.928. Thus, nonlinear weighting did not invariably reduce the tolerance for a selected harm. Each single-harm threshold omitted other endpoints, including the other rosuvastatin safety outcome.

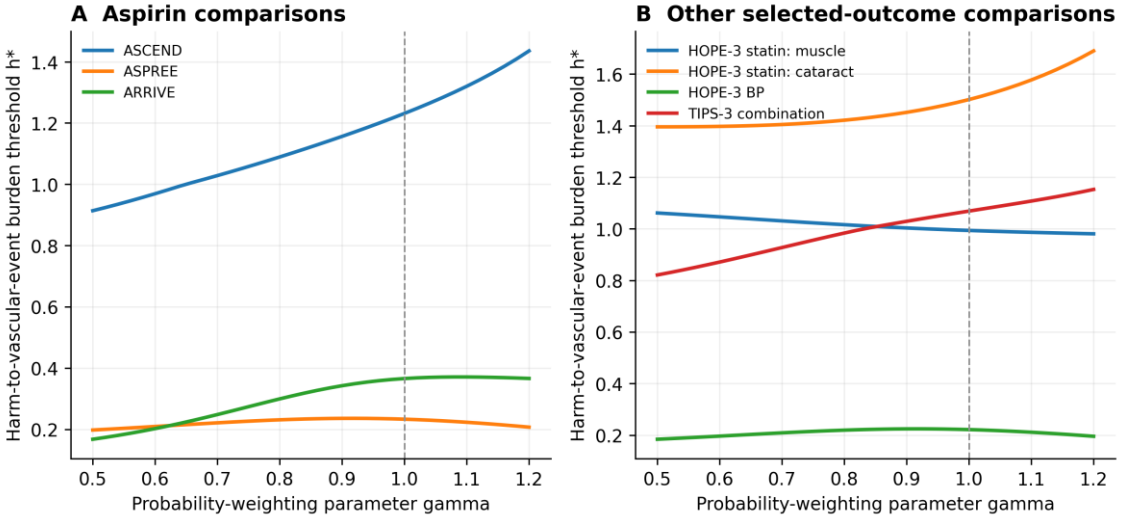


Figure 1. Severity thresholds across probability-weighting assumptions.

Figure 1 compares each treatment with its own control at its original trial follow-up (Table 1), using one vascular endpoint and one selected harm. γ is the Prelec weighting parameter; γ=1 is linear weighting. The vertical axis h* is the harm burden per one normalized vascular-event unit at which the modeled difference is zero. Curves identify trial and selected harm. Assumptions: independence, additive burdens, healthy-state reference, and α=β=1. No annualization or cross-trial severity calibration was applied. Source: arm counts in Table 2 [13-18].

The two-harm rosuvastatin boundary was as follows.

$$G = 0.0109754 - 0.0110371\, h_m - 0.0073070\, h_c.$$

Figure 2 compares its linear and nonlinear zero boundaries with the dependence scenarios. At $h_m = h_c = 0.1$, the additive linear difference was 9.14 normalized units per 1000 participants. The independent CPT difference was positive, but its sign could reverse under tested dependencies. Similar reversals occurred in selected aspirin and blood-pressure scenarios. For the polypill

combination at $h = 0.1$, the valuation remained positive across the tested common-uniform family, a result limited to that burden and endpoint definition.

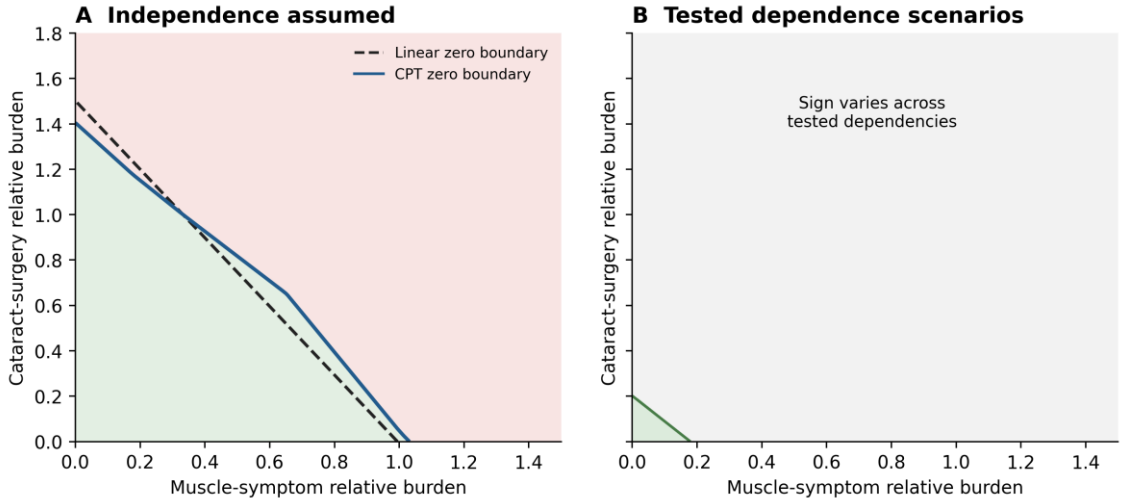


Figure 2. Two-harm rosuvastatin sensitivity surface.
Panel A compares the zero boundaries for linear weighting and illustrative CPT (γ=0.7) under independence. Both axes express selected harm burden per one normalized vascular-event unit. Green and pink indicate positive and negative CPT differences, respectively. Panel B retains the same CPT settings but varies dependence; green indicates positive throughout, pink negative throughout, and gray a change of sign across the tested configurations. Both panels use additive burdens, linear outcome values, and the healthy-state reference; the arm denominators and median 5.6-year follow-up are in Tables 1–2. Regions represent structural scenarios, not confidence or identification bounds. Source: HOPE-3 [16].

All declared implementation checks passed. In particular, a common lambda changed scale but not sign under the healthy-state reference. Results obtained by altering dependence were structural sensitivity results, not evidence that such dependence patterns occur in practice. No national number of events prevented, empirical DALYs averted, cross-intervention ranking, cost-effectiveness ratio, or behavioral uptake probability was generated.

Supplementary Table S1 reports selected scenarios varying the reference point, outcome-value curvature, and joint-burden rule. Curvature comparisons are confined to ASCEND; the rosuvastatin examples retain linear outcome values. The supplement also supplies the underlying randomized denominators and endpoint counts (Table S2).

Discussion

The threshold is specific to the chosen endpoint pair, follow-up and trial population. Composite vascular outcomes combine events with different consequences; normalizing each to one unit does not create a common severity scale. Crude proportions ignore censoring and do not put trials with different follow-up on a shared time horizon. The safety outcomes range from bleeding to symptoms, surgery and discontinuation. They are selected outcomes, and a threshold can change if other relevant outcomes are included.

For health education, the simulations can help identify trade-offs worth explaining and testing in a specific population. Presenting absolute risks with consistent denominators and explicit time horizons is supported by separate risk-communication and decision-aid evidence [1-3,25]. These simulations do not test communication effectiveness, adherence or clinical decisions.

Symptom-related discontinuation was assigned a hypothetical preference burden. Because discontinuation may already influence intention-to-treat efficacy estimates, no additional cardiovascular penalty for nonadherence was deducted.

Published marginal counts do not identify whether the same individuals experienced vascular and safety outcomes. This does not affect an additive linear expected burden, but it can change cumulative weighting or the maximum-burden valuation. The tested dependence patterns and illustrative CPT parameters therefore describe structural sensitivity. Missing within-person covariance precludes a