MYLIFE.IE · WORKING PAPER SERIES · MWP-2026-05

Serious Illness Cover in Ireland

Claims, Definitions and Fair Value for Consumers

Donal Milmo-Penny, QFA FLIA · SMP Financial (mylife.ie)
Research input: Norah McNulty, QFA CB

Version 1.2.2 · July 2026
Companion workbook: MWP-2026-05_audit_workbook_v1.2.2.xlsx

Serious Illness Cover in Ireland: Claims, Definitions and Fair Value for Consumers

Mylife Working Paper Series — MWP-2026-05
Donal Milmo-Penny, QFA FLIA — SMP Financial (mylife.ie)
Research input: Norah McNulty, QFA CB.
v1.2.2 — July 2026

Cite this paper: Milmo-Penny, D. (2026). *Serious Illness Cover in Ireland: Claims, Definitions and Fair Value for Consumers.* Working Paper MWP-2026-05, v1.2.2, July 2026. mylife.ie.

Abstract

Specified serious-illness insurance pays a lump sum on diagnosis of one of a defined list of conditions, and it is bought by a large share of Irish protection customers. Because the benefit is triggered by a contractual definition and not by a diagnosis alone, two policies naming the same conditions can differ materially in what they are expected to pay — a gap reflected in Irish serious-illness claims being admitted at lower rates than life-cover claims. This paper measures that difference across the five domestic life offices. It maps each office's policy conditions to a common disease classification, values them against independently sourced Irish incidence for a portfolio of ten representative lives (entry ages 30 to 50, both sexes), and expresses the result as the expected present value of the serious-illness benefit relative to the premium — a fair-value index. The whole calculation is held in an accompanying audit workbook, and incidence is drawn only from epidemiological sources, kept separate from claims experience.

At every age and for both sexes, the index varies across offices by several times more than premium does: value is driven far more by how conditions are defined than by price, and no office leads across the whole age range. Because incidence is sex-specific while premiums are, by regulation, unisex, an identical contract at an identical premium delivers less expected benefit to a woman than to a man, and this gap is wider at older ages; its size varies with the office's disease mix, being widest for the cardiovascular-strong office and narrowest for the cancer-strong one. Read from the insurer's side, the same figures identify, without imputing intent, which lives return the greatest gross margin under unisex pricing. Sensitivity testing and an equal-weighted market check confirm the findings do not depend on the discount rate, the wording-to-coefficient mapping or the population weighting. The paper's contribution is a sourced and reproducible method for valuing serious-illness wording — computed from disclosed inputs, with its key assumptions sensitivity-tested and shown not to affect the findings — and evidence that wording, not price or the length of the conditions list, should be central to how the cover is compared and advised.

Contents

1. Introduction

1.1 Purpose and question

Specified serious-illness cover pays a lump sum when the insured is diagnosed with one of a defined set of conditions. It is among the most widely held forms of protection in Ireland, and yet it is among the hardest for a consumer to compare, because the benefit turns not on whether a person falls ill but on whether their illness meets a contractual definition. This paper asks a single, practical question: across the domestic market, how much does the choice of provider change the value of the cover a consumer actually receives, and does that difference come from price or from the wording of the contract?

The paper answers it by measurement rather than assertion. It builds a standardised framework in which each provider's policy conditions are mapped to a common disease classification and valued against independently sourced Irish incidence, producing — for each of a portfolio of representative lives spanning the working ages — the expected present value of the serious-illness benefit and its ratio to the premium. That ratio, the fair-value index, is the paper's central instrument. Every figure is computed in an accompanying audit workbook from disclosed inputs, so that the results can be reproduced and challenged.

1.2 Why the definition matters

The distinctive feature of serious-illness cover is that a medically real and serious event may still fail to trigger a claim. Death, the event insured by life cover, is binary and readily proven; a serious-illness benefit depends on a threshold — a cancer's stage, a heart attack's biomarker level, a stroke's persistent deficit. The public claims evidence shows the consequence directly: Irish serious-illness claims are admitted at appreciably lower rates than life-cover claims, and much of that difference reflects definitions being met or missed rather than events occurring or not. If wording determines whether real events become payable, then wording is not a technical footnote but a central determinant of value — and one that is largely invisible where products are compared on premium and on the number of conditions listed.

1.3 Contribution

The paper makes two contributions. The first is methodological: a transparent, auditable way to translate heterogeneous policy wording into comparable value, holding the medical event fixed while letting contractual treatment vary. The second is empirical: applied to the five domestic offices, the framework shows that value differs across providers by far more than price does, that no office leads across the whole age range, and that — because incidence is sex-specific while premiums are, by regulation, unisex — the value of an identical contract depends on the profile of the insured. Together these support a research-first approach to serious-illness selection, in which the contract is matched to the individual rather than a single product being named as best.

1.4 Structure of the paper

Section 2 sets out the data sources and the evidential standards that govern their use, including the separation of incidence from claims experience. Sections 3 to 7 develop the framework: the disease taxonomy and contract normalisation (Section 3), the claims evidence base (Section 4), the wording-assessment framework (Section 5), the mathematical model (Section 6), and the base case and portfolio of representative lives (Section 7). Section 8 reports the results and Section 9 tests their sensitivity; Sections 10 and 11 discuss the implications and conclude. Formal propositions and proofs, the reference list, the data-provenance record, data notes and a declaration of interests are given in the appendices.

2. Data sources and evidential standards

2.1 The source hierarchy

The paper draws on a closed hierarchy of sources, in order of evidential weight. First are primary market and contract sources: the whole-of-market claims report (mylife.ie Claims 2025), the insurers' own claims disclosures and policy conditions (insurers' 2025 claims disclosures and policy conditions), the published clause-level policy-conditions ranking (All-Illness Conditions Ranking v1.2, mylife.ie) from which the coverage coefficients are derived, and the quotation system (mylife.ie whole-of-market quotation system, 24 Jul 2026) from which premiums are taken. Second are professional-actuarial, supervisory and reputable registry and epidemiological sources — the World Health Organization's disease classification, the national cancer registry, national hospital-audit and discharge data, the European renal registry, and the relevant supervisory reference data. Third are peer-reviewed academic sources, used where national data are incomplete. Fourth is internal working material. General journalism and competitor marketing are excluded as evidential authorities. The full, dated reference list is given in Appendix B.

2.2 The incidence–claims firewall

One rule governs the use of sources above all others. The incidence against which policy wording is valued is built exclusively from epidemiological and population data. Serious-illness claims statistics — paid rates, claim counts, cause shares — are never used to construct incidence, nor to weight illness families within the valuation; they serve only to establish which families are material in practice and, at the end, as an ex-post check that the incidence-weighted family mix is not implausible against observed claims. The reason is methodological necessity. Claims data reflect the combined influence of incidence, wording, underwriting, portfolio mix and distribution, and using them to calibrate incidence would confound the very effects the paper seeks to separate, admitting a circularity in which claims are explained by weights drawn from those same claims. The dominance of malignant disease in the results is therefore an output of independently sourced incidence; its broad agreement with the claims cause-mix is a validation of the model, not an input to it.

2.3 Excluded sources and international evidence

General journalism and non-technical commentary may surface a relevant fact but do not provide a robust basis for quasi-actuarial comparison, and are not relied upon; competitor marketing is likewise excluded. International evidence is used only where it is official, primary or methodologically strong and where Irish data are unavailable or incomplete, and the reason for its use is stated where it arises.

2.4 The disease base and the commercial-sensitivity boundary

The disease taxonomy is anchored to the World Health Organization's International Classification of Diseases, Tenth Revision (ICD-10) — the revision in which the Irish incidence sources the wording is valued against (the National Cancer Registry and HIPE) are themselves coded, so that taxonomy and incidence share one coding system rather than being bridged across revisions (Section 3.1). It provides a neutral medical reference against which contractual definitions are normalised. The paper distinguishes between what must be published for the analysis to be auditable and what is withheld as commercially sensitive: the formulas, assumptions, illness-family structure and derived coefficients needed to reproduce the results are all disclosed in the audit workbook, while the full internal condition library and operational implementation that would let a third party reconstruct the firm's advisory capability as a commercial product are withheld on commercial grounds alone. That omission does not affect the reproducibility of any published result. The author's commercial interest in the offices analysed is declared in Appendix E; the safeguards described in this section are the paper's answer to it.

2.5 Unisex pricing and the Gender Directive

Since 21 December 2012, life offices operating in Ireland have been prohibited from using the sex of the applicant as a rating factor in pricing protection cover. The prohibition derives from Council Directive 2004/113/EC (the "Gender Directive") (Council Directive 2004/113/EC), which implements the principle of equal treatment between men and women in the access to and supply of goods and services. Article 5(1) provides that the use of sex in calculating premiums and benefits must not give rise to differences in individuals' premiums and benefits; Article 5(2) had, however, permitted proportionate sex-based differentiation where supported by relevant actuarial and statistical data, and life offices across the European Union relied on that derogation. In Case C-236/09 (Association Belge des Consommateurs Test-Achats ASBL and Others v Conseil des ministres) (CJEU C-236/09, *Test-Achats*), the Court of Justice of the European Union held, on 1 March 2011, that the derogation was incompatible with Articles 21 (non-discrimination) and 23 (equality between women and men) of the Charter of Fundamental Rights, and declared Article 5(2) invalid with effect from 21 December 2012. From that date, all new insurance contracts — including the level term assurance and accelerated serious-illness cover examined here — must be priced on a unisex basis (European Commission, 2012/C 11/01): a male and a female applicant of the same age, smoking status and health are charged the same premium for the same cover, notwithstanding that their underlying morbidity differs.

This is central to interpreting the results. Because the premium is identical across the sexes while serious-illness incidence is not, the premium denominator of the fair-value index does not vary by sex; any difference in fair value between a male and a female life is therefore a pure benefit-value effect, arising from the sex-specific pattern of disease and not from price. The widening of that gap with age is a direct consequence of pricing a sex-differentiated risk on a sex-neutral basis.

---

3. Disease taxonomy and contract normalisation

3.1 A neutral disease base

The analysis begins from a disease classification that is independent of any insurer's drafting. The World Health Organization's International Classification of Diseases, Tenth Revision (ICD-10) (WHO ICD-10) is used as that reference base — the revision in operational use by the Irish data sources drawn on here (the National Cancer Registry and the Hospital In-Patient Enquiry), so that the incidence codes and the taxonomy are the same system. (ICD-11 has been the mandated WHO revision since 1 January 2022, and WHO ceased maintaining ICD-10 in 2018; ICD-10 is retained here because the Irish registries this paper draws on — the NCRI and HIPE — still code in ICD-10 themselves, so a forward-mapping to ICD-11 would add nothing to the analysis until the source data migrate. That mapping is recorded in Appendix D as not-yet-performed rather than merely optional.) Its role here is classificatory, not actuarial: it provides a medically neutral vocabulary within which disease entities can be identified and grouped consistently across providers, so that contractual variation can be studied against a fixed medical backdrop rather than against each insurer's own terminology.

The distinction this enables is central to the whole paper. Serious-illness insurance does not insure diseases in the abstract; it insures contractually defined trigger states that sit within, and often at the margins of, medically recognised disease families. Holding the medical identity of an illness constant while allowing the contractual definition to vary is what makes a like-for-like comparison possible. Where the paper refers to specific ICD-10 groupings, those assignments are recorded in the audit workbook rather than asserted in the text, so that the mapping can be inspected and challenged directly.

3.2 The four analytical layers

Each insurable condition in each provider's wording is decomposed into four layers:

1. **Disease identity** — the ICD-10 family or grouping to which the condition belongs.
2. **Contractual trigger** — the state the wording actually requires for the benefit to become payable (severity, stage, biomarker, persistence, procedural route).
3. **Benefit structure** — whether payment is full, partial, ancillary, child-specific or subject to a booster.
4. **Exclusion architecture** — carve-outs, staging exclusions and limiting clauses.

This layered representation is the object on which the rest of the model operates. Its purpose is to separate the probability that a disease occurs (layer 1, a medical question) from the probability that the wording recognises the occurrence as payable (layers 2–4, a contractual question) and from the amount then payable (layer 3). Those are different things, and the model keeps them distinct throughout.

3.3 Aggregation to illness families

Because no two offices draft on a common basis, condition-level comparison is aggregated to illness families that are clinically and actuarially coherent — malignant disease (with an explicit separation of invasive from early-stage forms), cardiac events, cerebrovascular events, demyelinating and neurodegenerative disease, major organ failure, and so on. The aggregation rule is deliberately conservative: where a condition can be assigned unambiguously to a family it is included; where it straddles families or is materially ancillary it is treated separately or set aside and discussed rather than forced into the main valuation. This avoids the false precision of raw condition-counting, under which a long list of narrowly defined benefits can appear equivalent to a shorter list of broad ones.

---

4. Claims evidence base

4.1 The whole-of-market claims anchor

The principal empirical description of the market is the mylife.ie Life Insurance Claims in Ireland — 2025 whole-of-market report (mylife.ie Claims 2025), which consolidates the publicly disclosed 2025 claims statistics of the five domestic offices in scope. It records

approximately €919.2 million of protection claims paid across more than 18,200 individual claims, and it establishes three facts that shape the analysis: that the market pays claims at scale; that life-cover paid rates cluster at very high levels (broadly 97–99% where disclosed; the top of that range reflects at least one office's combined-benefit disclosure — life, serious-illness and income protection together — rather than a strictly life-only figure, consistent with the disclosure-scope caveat in Section 4.4) while serious-illness paid rates sit materially lower (broadly 87–90%); and that malignant disease dominates the serious-illness book, accounting for roughly 60–68% of serious-illness claims across the offices that publish a cause breakdown, with heart attack, stroke and multiple sclerosis following.

4.2 The definition problem

The gap between life-cover and serious-illness paid rates is the single most important signal in the public data. Death is largely an evidential event: it is binary and readily proven. A serious-illness claim is structurally different — a claimant may have a medically real and serious condition and still fail to meet the contractual threshold, because a myocardial infarction did not reach a biomarker cut-off, a stroke did not leave the required persistent deficit, or a malignancy was in-situ or low-grade. The lower serious-illness paid rate is therefore, in large part, a definition phenomenon. That observation is the empirical motivation for the coverage-intensity coefficient introduced in §5: if definitions determine whether medically real events become payable, then wording is not a peripheral matter but a central determinant of value.

4.3 How the claims data are — and are not — used

The claims evidence is used for two purposes only: to establish which illness families are material in practice, and, at the end of the analysis, as an ex-post reasonableness check that the incidence-weighted family mix emerging from the model is not wildly at odds with observed claims. It is **not** used to set incidence, and it is **not** used to weight the families inside the expected-present-value calculation.

This firewall is deliberate and is enforced in the workbook. Incidence is built exclusively from epidemiological and population sources (national cancer registry data, hospital-discharge and audit data, disease registries). Claims statistics — paid rates, claim counts, cause shares — never enter the incidence inputs or the family weighting. The reason is that claims data reflect the joint effect of incidence, wording, underwriting, portfolio mix and distribution; using them to calibrate incidence would confound the very things the model is trying to separate, and would risk the circularity of "explaining" claims with weights derived from those same claims. The cancer dominance visible in the results is therefore an output of independently sourced incidence, not an assumption imported from the claims report — and the fact that the two broadly agree is a check on the model, not an input to it. That agreement can be quantified. For the reference life, aggregated across the five offices, the model attributes roughly 68% of expected serious-illness value to malignant disease against the 60–68% cancer share of the serious-illness claims book. That 68% is the equal-weighted male/female aggregate for the reference life; the male and female components (about 58% and about 79% respectively, the difference reflecting the sex-specific disease mix) bracket the claims range, so a market-realistic sex mix places the aggregate at its upper end. The two measures are constructed differently (expected value of benefit versus share of paid claims), so exact equality is neither expected nor required; the point is that an incidence model built with no reference to the claims data lands inside the claims-observed range, which is precisely the validation the firewall is designed to permit.

4.4 Limitations of the claims base

The claims report is a normalised public view, not a statutory common-template return. Disclosures differ in scope, granularity, rounding and the treatment of group and ancillary benefits. The report is accordingly used to establish the broad empirical shape of the market — scale, product mix, the dominance of malignant disease, the centrality of definition satisfaction — and not to support office-level actuarial benchmarking, which the public data cannot bear.

---

5. Wording assessment framework

5.1 Purpose: separating reach from amount

The wording framework converts heterogeneous policy definitions into the model's contractual variables. Its organising principle is the separation of **reach** — how much of a medically relevant illness family a contract actually responds to — from **amount** — how much of the sum assured is then payable. These are captured by two distinct objects, the coverage-intensity coefficient and the benefit fraction, and they are not conflated: a contract may have broad reach at a reduced amount, or narrow reach at full amount, and only by valuing both can the two be compared.

5.2 Principal wording dimensions

Each condition is assessed across the dimensions that recur in serious-illness wording and that materially affect whether a benefit is triggered: diagnostic threshold (what evidence is required); severity threshold (the stage or grade the disease must reach); temporal persistence (how long a deficit or symptom must last); procedural route (whether payment is tied to a particular intervention); benefit tiering (full versus partial); and exclusion architecture (explicit carve-outs). These dimensions are the mechanisms through which a common medical event is turned into a provider-specific payable state.

5.3 The coverage-intensity coefficient

For illness family i and provider j, the coverage-intensity coefficient w(i,j) measures the share of the medically relevant illness space that the contract effectively recognises. Conceptually, if $\Omega(i)$ denotes the medically relevant event space for family i and $T(i,j) \subseteq \Omega(i)$ the subset that qualifies as a payable trigger under provider j's wording, then

$$w_{i,j} = \frac{\mu(T_{i,j})}{\mu(\Omega_i)},$$

where μ is a normalised measure of relevant disease space. This is the interpretation of the coefficient — the fraction of the disease family that is contractually reachable.

The operative w is a transparent proxy for that conceptual quantity, not a literal integral over disease space: it is derived from a published clause-level ranking of every provider's serious-illness policy conditions (the "All-Illness" serious-illness policy-conditions ranking, version 1.2) (All-Illness Conditions Ranking v1.2, mylife.ie), in which each condition is assigned an ordinal wording rank and a coverage tier. The rank is converted to a triggerability score on a linear scale and gated by the tier, so that conditions which are not covered, or are covered only for children, contribute zero. The exact mapping from rank and tier to w is set out in the audit workbook and is sensitivity-tested. The coefficient is an ordinal wording proxy for the conceptual quantity above, not a literal measure, and it is attributed to that public ranking research.

One limitation of this construction must be stated plainly, because it is the point at which the declared conflict of interest in Appendix E most directly touches the numbers. The ordinal placement of providers within the ranking comes from a single, internally produced source — mylife.ie's own "All-Illness" conditions ranking — and the firm that produces it holds commercial appointments with every office ranked. The sensitivity in Section 9.4 tests robustness to the *mapping* from rank to coefficient, holding placement fixed; it does not, and cannot, test the placement itself against an independent assessor. Every downstream quantity — expected value, fair value, the sex gap and the market ranking — is therefore conditional on that single ranking being correct in its *relative* ordering of providers. The safeguards in Section 2 constrain how the ranking is used, but they do not substitute for an independent second opinion on the ranking itself. A blinded re-rank of a subset of headings by an independent assessor, reported as a rank correlation (Spearman or Kendall) against the ranking used here, is the natural triangulation and would strengthen the ranking's external validity. The ordinal placement should accordingly be read as the paper's most conflict-exposed input, and the results as conditional upon it.

5.4 The benefit fraction

The benefit fraction b(i,j) is the share of the serious-illness sum assured payable when a trigger is met, and it is keyed to the coverage tier: full-payment conditions take b = 1; partial-payment conditions take their contractual fraction (0.50 in the base case, the market-standard fraction, sensitivity-tested); ancillary benefits take a small fraction; and conditions gated out by tier take zero. Booster structures are held at the base sum in the central case and treated as a separate sensitivity, so that a feature not comparable across the market does not distort the base comparison.

5.5 The timing factor

The timing factor c(i,j,t) captures features that affect when, or under what persistence conditions, a trigger can operate. It defaults to unity, and is departed from only where a definition's persistence, survival or timing requirement is a genuine differentiator; such features are placed here rather than absorbed silently into w, so that reach and timing remain conceptually distinct.

5.6 From condition to family

The assessment is performed at condition (heading) level, and the valuation is carried out at that level too: expected present value is summed heading-by-heading, with each heading weighted by its own sourced incidence. This is important, because it means the illness-family emphasis in the model is determined by incidence, not by any discretionary weighting scheme. Where family-level summaries are reported, they are aggregations of heading-level results. No family-weight vector is imposed, and in particular no weighting derived from claims cause-shares is applied — the weighting of the model is the incidence function of §6.

5.7 Commercial sensitivity and the boundary with the advisory engine

The formulas, tiers, coefficients and assumptions that generate the published results are all visible in the workbook. What is not published is the full internal condition library and operational implementation that would allow a third party to reconstruct mylife.ie's advisory engine as a commercial proposition; that material is commercially sensitive and is withheld on that basis alone.

A distinction of substance, not only of confidentiality, should be drawn here. mylife.ie operates a proprietary pre-underwriting and matching capability ("Health Gate") that also rests on clause-level wording analysis. But the wording score it produces is customer-specific: it is designed to move with an individual's health and circumstances, and the same provider can score differently for two different customers. The coefficient w in this paper is a different object — a population-level, profile-independent property of a provider's wording, identical for every insured. The paper's coefficients are therefore not the advisory engine's outputs; they are derived from the public v1.2 ranking. The Health Gate capability and the firm's published statement of method are relevant only as context — they explain why the research base for this paper exists — and are not the provenance of any figure reported.

---

6. Mathematical model and notation

6.1 Objective

The model measures, on a standardised comparative basis, the expected present value of serious-illness benefits under differing policy definitions. It is not a reserving tool, a claims forecast for any office, or a substitute for underwriting. Its purpose is to quantify how differences in contractual trigger structure and benefit design alter expected value for a given life under stated assumptions.

6.2 Notation

Let i index illness headings (aggregated to families for reporting) and j index providers; $t = 1, \ldots, n$ index policy years; x the entry age; S

the serious-illness sum assured; and $v = (1+r)^{-1}$ the annual discount factor. Then:

- $q_i(x,t)$ — the baseline probability that a life aged x first experiences an event in family i during policy year t (a medical incidence object, sourced from epidemiology);
- $w_{i,j}$ — the coverage-intensity coefficient (a contractual object, §5.3);
- $c_{i,j}(t)$ — the timing factor (§5.5);
- $b_{i,j}(t)$ — the benefit fraction (§5.4); in the base case it is independent of t, the time argument being retained so that future booster or deferred-benefit structures can be represented without altering the model's form.

The separation of the medical object q from the contractual objects w, c, b is the structural heart of the model.

6.3 Effective insured incidence

The effective insured incidence — the rate at which medically observable events become contractually payable triggers — is

$$q^*_{i,j}(x,t) = q_i(x,t)\, w_{i,j}\, c_{i,j}(t).$$

A coefficient below unity means the contract recognises only part of the medically relevant family; the product qw is precisely the phenomenon the claims evidence points to when serious-illness paid rates fall below life-cover paid rates despite substantial claim volumes.

6.4 Expected present value

The expected present value of the serious-illness benefit for provider j, entry age x, term n is

$$EPV_j(x,n) = \sum_{t=1}^{n} \sum_{i} q_i(x,t)\, w_{i,j}\, c_{i,j}(t)\, b_{i,j}(t)\, S\, v^t.$$

The model is intentionally modular: age, sex, duration and staged-benefit refinements enter by refining a component function, not by altering the structure.

6.5 Full- and partial-payment decomposition

Partitioning the headings into full-payment ($\mathcal{I}_F$) and partial-payment ($\mathcal{I}_P$) sets gives

$$EPV_j = EPV^F_j + EPV^P_j,$$

with each term the corresponding restricted sum. This decomposition is reported in the results because it isolates how much of a provider's value comes from earlier or lower-severity triggers as opposed to classic full-severity events — one of the paper's substantive questions.

6.6 Derived measures

Three comparative measures are derived from the EPV.

- **Fair-value index.** For a given life with serious-illness premium P_j, the fair-value index is

$$FV_j = \frac{EPV_j}{PV(P_j)},$$

the expected present value of benefit per unit of the present value of premium. The premium denominator is integral to the paper, not an optional overlay: a valuation of benefit that never confronts price cannot speak to value. It is computed on an un-decremented, pre-expense basis (no lapse, mortality, expense or profit loading on either side), so that its comparative content is clean; its absolute level is not a literal money's-worth.

- **Partial-payment share.** $PPS_j = EPV^P_j / EPV_j$, the share of value attributable to partial or early-stage benefits.
- **Effective coverage index.** `ECI_{i,j} = (\sum_t q^{*}_{i,j})\big\,/\,(\sum_t q_i)`, the fraction of medical incidence that becomes effective insured incidence for family i under provider j. This is a derived output — the realised coverage after wording and timing — and is distinct from the input coefficient w.

No scenario re-weighting is folded into the incidence or the effective-incidence terms; where consumer-profile scenarios are examined, they are handled in the sensitivity analysis by varying which families are summed, never by a multiplier inside q.

6.7 Propositions

Four elementary propositions follow and are proved in the appendix.

- **P1 (monotonicity).** For fixed q, c, b, S, v, EPV_j is weakly increasing in each $w_{i,j}$.
- **P2 (dominance of an added payable state).** If two contracts are identical except that one adds a payable state with strictly positive incidence and benefit fraction, that contract has weakly higher EPV.
- **P3 (effect of narrowing).** A stricter threshold that reduces $w_{i,j}$ weakly reduces the expected value of that family, all else equal.

- **P4 (widening gap).** For fixed q_i, the gap between medical and effective insured incidence, $q_i - q^*_{i,j}$, is increasing in $(1 - w_{i,j}\, c_{i,j})$.

These are simple, but they provide the formal basis for the empirical claims: that wording breadth raises value, that early-stage payable states contribute value in proportion to incidence and amount, and that contractual narrowing has a measurable cost.

6.8 The workbook as audit trail

Every quantity above is computed in the accompanying audit workbook (MWP-2026-05 audit workbook), which holds the operative q, w, c, b, discount and premium inputs and the resulting EPV, decomposition and fair-value outputs. The equations in this section are the analytical summary of that workbook; the workbook is the numerical record.

7. Base case and the portfolio of representative lives

7.1 From a single reference life to a portfolio

Earlier papers in this series valued a single reference life, a device that isolates contractual differences at one point in the age range. Because serious-illness incidence varies steeply with age and differs by sex, while premiums do not (Section 2.5), a single life cannot show how value is distributed across the people who actually buy the cover. This paper therefore evaluates a portfolio of representative lives spanning the working-age range, and then weights them to a market view. The single 35-year-old of the earlier work is retained as one member of that portfolio, so the results remain continuous with, and reconcilable to, the prior analysis.

7.2 The representative lives

The portfolio comprises entry ages 30, 35, 40, 45 and 50, evaluated for both a male and a female life — ten representative lives in all. Each is a non-smoker in standard health, taking level term cover to age 65 (so the term shortens as entry age rises, from 35 years at age 30 to 15 years at age 50). Holding the product and the sums assured constant across the ten lives means the comparison isolates the effect of age, sex and provider, and nothing else.

7.3 Sums assured

Every life takes life cover of €322,000 with accelerated serious-illness cover of €161,000 (50% of the life sum), the convention inherited from the series (mylife.ie Working Paper Series). Sums assured are held constant across ages so that the fair-value ratio, which normalises by premium, reflects contractual value rather than the size of cover purchased. In the workbook the serious-illness sum is carried as a formula (0.5 × 322,000) rather than a constant.

7.4 Unisex premiums

Premiums are the monthly cost of the accelerated serious-illness benefit, taken as the difference between a life-plus-illness quotation and a life-only quotation for the same cover, sourced from the whole-of-market quotation system (mylife.ie whole-of-market quotation system, 24 Jul 2026) on a single date (24 July 2026) on a guaranteed (non-reviewable) basis for all five offices. Because pricing is unisex by law (Section 2.5), premiums are quoted by age only; the same premium applies to a male and a female life of the same age. This is the hinge of the analysis: with an identical denominator across the sexes, every difference in fair value between a male and a female life is a difference in expected benefit, not in price.

7.5 Age-specific incidence

For each life the expected present value steps through attained age. Malignant disease is projected on the National Cancer Registry's five-year age bands (NCRI incidence by age, 2017–2019), so incidence rises with the life as it ages from entry to 65; early-stage breast disease is treated the same way on its own age bands. The cardiovascular and neurological families are projected on the broad hospital-audit age bands (HIPE, 2017–2019) (under-45 and 45–64), weighted by the discounted time the life spends in each. Every incidence input is the same sourced quantity used elsewhere in the workbook; only the age range traversed, and the discounting, differ between lives.

7.6 Discounting

Benefits and premiums are discounted at a central real rate of 3.0% (EIOPA EUR risk-free rate) (tested at 2% and 4% in Section 9). Because both sides of the ratio are discounted identically, the discount rate moves the level of the index but not the comparison between providers or between the sexes.

7.7 The market view

To express the ten lives as a single market picture, the fair values are weighted by the age and sex structure of the population using the CSO Census 2022 five-year age groups (30–54) (CSO PxStat PEA11 / Census 2022). This general working-age structure is a transparent proxy for the buying population; because the true buying population skews somewhat younger, and younger lives receive higher fair value, the weighting is, if anything, conservative with respect to the level of value reported. Critically, the market-level figures are reported alongside an equal-weighted alternative, and the two are almost identical, so the market conclusions do not depend on the weighting. The conservative reading just given concerns the level of value; the male–female gap could in principle move differently under a reweighting toward a younger or more male-skewed buyer population, but the equal-weighted check bounds the gap as well as the level — the two weightings agree to within 0.01 on both — so neither conclusion rests on the population structure.

7.8 What is computed, and how it is validated

For each representative life the workbook computes, by formula from the sourced inputs, the expected present value of the benefit by illness block, the total, the premium present value, and the fair-value index; and it aggregates these to the market-level male, female and

average-buyer figures. The engine is validated inside the workbook: the age-35 members of the portfolio reproduce the single-life result of the base engine to the third decimal, a reconciliation carried as a live formula that returns zero. No figure in the chain is entered as a computed constant.

7.9 Limitations of the base case

The base case abstracts from customer-specific underwriting and the pre-underwriting matching that operates in practice, from smoker status (all lives are non-smokers), and from lower-frequency families and early-stage sites outside the base-case scope. The cardiovascular age bands are broad. These abstractions are intentional and disclosed: they isolate the contractual value of the mainstream proposition on a stable, auditable base, to which customer-specific and lower-frequency refinements can later be added.

---

8. Results

Provider-level results are shown anonymised as Provider A–E; the figures are outputs of the accompanying audit workbook.

8.1 What is reported

For each of the ten representative lives (entry ages 30, 35, 40, 45 and 50, male and female) the model produces the expected present value of the serious-illness benefit and, against the unisex premium for that age, the fair-value index (Section 7). The index is a comparative measure of expected benefit per unit of premium on a common, un-decremented, pre-expense basis; its level is not a literal money's-worth, and it is the comparison across providers, ages and sexes that carries the content. The results are then weighted to a market view using the CSO age-sex structure. Values are reported to two decimal places.

8.2 Fair value across the age range

Table 8.1 — Fair-value index, male lives, by entry age

Provider	30	35	40	45	50
A	0.90	0.94	0.96	0.95	0.87
B	1.02	1.03	0.92	0.98	0.92
C	0.89	0.96	0.98	0.95	0.80
D	0.74	0.72	0.75	0.72	0.63
E	0.22	0.25	0.30	0.27	0.24

Table 8.2 — Fair-value index, female lives, by entry age

Provider	30	35	40	45	50
A	0.81	0.83	0.82	0.76	0.66
B	1.00	0.98	0.85	0.84	0.74
C	0.65	0.69	0.68	0.62	0.51
D	0.56	0.52	0.53	0.48	0.42
E	0.19	0.21	0.25	0.21	0.18

Three features of these tables organise the discussion that follows: the dispersion between providers at any given age, the way the ordering changes with age, and the difference between the two tables.

8.3 Dispersion between providers persists at every age

At every entry age and for both sexes, the distance between the strongest and weakest provider is large. For a male life it runs from roughly 0.22–0.30 (Provider E) to roughly 0.92–1.03 (Provider B) across the ages; for a female life, from roughly 0.18–0.25 to 0.74–1.00. That spread is many times the spread in premium, which differs across providers by only a few euro a month at each age. The central finding — that contractual value is driven far more by wording and benefit design than by price — is therefore not an artefact of any single age; it holds across the whole working-age range.

8.4 No provider leads across the whole male age range

The ordering is not fixed. For male lives, Provider B leads at ages 30, 35, 45 and 50, but at age 40 it falls to third: Provider C moves ahead (0.98) and Provider A second (0.96), with B at 0.92. Provider C — strongest on the cardiovascular families — is at or near the front through the middle of the age range (leading at 40, level with A at 45) before falling back at 50, when the short remaining term limits how much cardiovascular incidence the projection accumulates. For female lives the ordering is more stable — Provider B ahead, then A, then C at every age — because value for women is more heavily cancer-weighted, where B and A are strongest. The practical implication is that a single market-wide "best value" claim is not supportable: the best contract depends on the age and sex of the life.

8.5 The value gap between the sexes

The most consequential result is the difference between Tables 8.1 and 8.2. At every age and for every provider, the female fair value is lower than the male — the same contract, at the same premium, delivers less expected benefit to a woman — and the gap is wider at older entry ages than at younger ones. For Provider B it rises from about 0.03 at age 30 to about 0.17 at age 50; for Provider A, from about 0.09 to 0.21. For the cardiovascular-strong providers the gap is larger throughout and broadly widens into the mid-forties before easing slightly at 50: Provider C runs 0.23, 0.27, 0.30, 0.33, 0.28 across the five ages, and Provider D 0.19, 0.19, 0.22, 0.24, 0.21. (These gaps are computed from the unrounded fair values in the audit workbook and may therefore differ by up to 0.01 from the difference of the rounded entries in Tables 8.1–8.2 — for example, Provider C at age 30 is 0.8875 − 0.6546 = 0.23, not 0.89 − 0.65 = 0.24.) The easing at age 50 reflects the short remaining term, which compresses both sexes' expected values; it does not reverse the pattern, since the gap at 50 remains well above the gap at 30 for every provider. The mechanism is set out in Section 2.5 and is not a matter of

interpretation: premiums are unisex by law, while serious-illness incidence is sex-specific and its male component — cardiovascular disease and, at older ages, prostate cancer — is larger, relative to the female cancer advantage, at older ages than at younger. Pricing a sex-differentiated risk on a sex-neutral basis produces a value difference that grows as the sex-specific risks diverge.

8.6 The market view

Weighting the ten lives by the CSO age-sex structure gives the market-level picture in Table 8.3. The average-buyer figures are reported alongside an equal-weighted alternative; the two agree to within a hundredth for every provider, so the market conclusions do not depend on the weighting.

Table 8.3 — Market-level fair value (CSO-weighted)

Provider	Male	Female	Average buyer	Male–female gap
A	0.93	0.78	0.85	0.15
B	0.97	0.88	0.93	0.09
C	0.92	0.64	0.78	0.28
D	0.71	0.50	0.61	0.21
E	0.26	0.21	0.24	0.05

Two things stand out. First, the average-buyer ordering — B ahead of A, then C, then D, with E far behind — confirms that the value dispersion survives aggregation. Second, the size of the male–female gap varies systematically with the provider's disease mix: it is widest for Provider C (0.28), whose contractual strengths lie in the male-skewed cardiovascular families, and narrowest among the leaders for Provider B (0.09), whose strength is in cancer, where female incidence is high at working ages. Under unisex pricing, a contract built around male-skewed conditions transfers the most value away from female buyers; a cancer-weighted contract transfers the least. Provider E's small gap (0.05) is a different phenomenon — it reflects uniformly low value to both sexes, not balance.

8.7 Where price, not wording, moves the result

The exception is instructive. Provider B's male fair value dips at age 40 (to 0.92, behind Providers C and A) and its female value falls between ages 35 and 40. This is not a wording effect: Provider B's wording is unchanged across ages. It is a pricing effect — Provider B's serious-illness premium at age 40 is materially higher than its peers' (about €102 a month against roughly €80–€94 for the others), and the higher denominator pulls the ratio down. It is the one place in the analysis where price, rather than definition, drives the outcome, and it is reported precisely because it is the exception that locates the rule: elsewhere, differences in value track differences in wording, not price.

8.8 What the results do not show

The fair-value index measures modelled contractual value per unit premium under stated assumptions. It is not a claims-payment record, an underwriting assessment, a service measure or a solvency ranking, and it does not capture the customer-specific underwriting and matching that determine real-world outcomes for a given applicant. Different offices may rationally position for different segments. The results identify where structural differences in value arise, how large they are, and how they fall across age and sex; they do not establish that any one office is preferable to another for every consumer.

---

9. Sensitivity and scenario analysis

This section tests how far the results of §8 depend on the assumptions that are least firmly fixed. Each is varied around the base case, and throughout a distinction is drawn between the *absolute level* of the fair-value index — which several assumptions move — and the *cross-provider ordering, the sex gap and the market-level conclusions*, which are the paper's actual findings and which prove robust across the whole portfolio.

9.1 Discount rate

The central real rate of 3.0% is varied to 2.0% and 4.0%. Because incidence, and hence the benefit cash flow, is more back-loaded than the level premium, the fair-value level rises as the rate falls and falls as it rises. Illustrated at the age-35 anchor for a male life (other ages and the female lives behave identically in direction):

Provider	FV at 2.0%	FV at 3.0%	FV at 4.0%
A	1.00	0.94	0.88
B	1.10	1.03	0.96
C	1.02	0.96	0.90
D	0.76	0.72	0.68
E	0.27	0.25	0.24

The level moves by roughly ±6–7% across the range for the four higher-value providers (A–D); Provider E, whose base value is very low (about 0.25), responds more and asymmetrically — about +8% at a 2% rate and −4% at 4% — as percentage changes on a small base naturally do. But at every rate, and at every entry age, the ordering is unchanged and the distances between providers — and between the sexes — are preserved almost exactly. Because both sides of the ratio are discounted identically, the discount rate affects how the index is read in absolute terms, not the comparison.

9.2 Stroke calibration

The stroke input is the I60–I69 hospital grouping (HIPE, 2017–2019) scaled to the national acute-stroke total (NOCA — Irish National Audit of Stroke) by a factor of 0.79 (§7.5). Removing that adjustment (factor 1.0, the raw grouping) is the relevant upper-bound test and isolates the largest single calibration adjustment, again shown at the age-35 male anchor:

Provider	FV, factor 1.0 (raw)	FV, factor 0.79 (adopted)
A	0.96	0.94
B	1.06	1.03
C	1.01	0.96
D	0.76	0.72
E	0.26	0.25

The adjustment moves the cardiovascular-strong providers most — Provider C by about 0.05, Provider D by about 0.04 — and the others little. It therefore affects the size, but not the direction, of the male–female gap (which is driven by the cardiovascular families), and it leaves the provider ordering unchanged under either treatment at every age. The calibrated figure is adopted; this test shows the unadjusted alternative would change no comparative conclusion.

9.3 Contract-side parameters

Two contract-side choices are immaterial to the comparison. Raising the serious-illness sum assured from 50% to 100% of life cover scales expected benefit linearly while raising the premium slightly more than proportionally, so the index falls by about 2–3% uniformly across providers (a 30% fraction moves it the other way by a similar margin); the ordering is unchanged. Varying the partial-payment fraction from 0.50 to 0.25 moves the female index of the partial-paying providers by well under two per cent — the only early-stage family with incidence is breast in-situ, and one provider pays it in full regardless — and does not affect the ranking. This low sensitivity is a property of the current family scope, not of the model in general: were further early-stage or partial-payment families to be added (a future step flagged in Section 7.9), the fraction could become materially more influential and the present robustness would not automatically carry over. Neither parameter, on the current scope, bears on the structural findings.

9.4 The coverage-intensity mapping

The most consequential modelling choice is the conversion of the ordinal wording ranks into the cardinal coefficient w (§5.3), which the base case does linearly. Its influence is specific and bounded. Because any strictly increasing mapping preserves rank order, the ordering of providers within each family — and therefore the broad ordering of the results — is invariant to the choice of mapping. What the mapping affects is the *cardinal spread*: a more convex mapping widens the gap between the strongest and weakest wordings, a flatter one narrows it. The finding that value dispersion is several times premium dispersion is therefore qualitatively robust, but the precise multiple is a property of the linear mapping and should be read as illustrative rather than exact.

9.5 Market weighting

The market-level figures (§8.6) weight the ten lives by the CSO age-sex structure (CSO PxStat PEA11 / Census 2022). To confirm that they do not depend on that weighting, they are computed a second way with equal weights across the entry ages. The two agree to within 0.01 for every provider on the average-buyer index and on the male–female gap. The market conclusions — the ordering of offices, and the finding that the female penalty is widest for the cardiovascular-strong office and narrowest for the cancer-strong one — are therefore properties of the per-age results, not artefacts of the population weights. Because the true buying population skews younger than the general working-age structure, and younger lives receive higher value, the CSO weighting is if anything conservative as to the level of value.

9.6 Summary

Across every test the substantive findings are preserved: contractual value varies far more than price at every age; no office leads across the age range; the female fair value is lower at every age and widens with age; and the size of that gap tracks the office's disease mix. What the sensitivities move is the absolute level of the index (discount rate), the size but not the direction of the cardiovascular contribution (stroke calibration), and the cardinal width of the spread (the coverage mapping). None disturbs the comparative or structural conclusions, and the market-level results are robust to the population weighting. The scenario in which an individual's own risk profile is emphasised is not a separate test but the portfolio itself, read life by life: because the model is additive across families, a cancer-weighted profile reproduces the cancer-led ordering and a cardiovascular-weighted profile the cardiovascular-led ordering, which is the formal basis for the paper's claim that the best-value contract is conditional on the individual.

---

10. Discussion and implications

10.1 What the findings mean

Across the working-age range and both sexes, the same pattern recurs: the value a serious-illness contract delivers is determined far more by how its conditions are defined than by what it costs. The dispersion between providers is several times the dispersion in premium at every entry age (Section 8.3), and it survives aggregation to the market level (Section 8.6). The corollary is that value is profile-conditional. No provider leads across the whole age range for men, and the ordering below the leader reshuffles with age as the disease mix shifts from cancer-dominated at younger ages to increasingly cardiovascular in middle age (Section 8.4). Two distinct mechanisms should be kept apart here. The male age-instability is thin: it rests on a single reordering at age 40, and that reordering is driven by a *pricing* anomaly — Provider B's unusually high premium at that age (Section 8.7) — not by wording. The sex-gap finding, by contrast, is the structural, wording-driven result, present at every age and for every provider. The paper's central thesis that wording rather than price drives value is carried by the sex gap and by the dispersion result, not by the male age-instability, which is the one place price visibly bites. A recommendation that names a single "best value" contract for the market as a whole is not supported by the evidence; the defensible unit of analysis is the individual life.

10.2 Unisex pricing and the value gap between the sexes

The clearest structural result is that, at every age and for every provider, an identical contract at an identical premium delivers less expected benefit to a woman than to a man, and that the gap is wider at older entry ages (Section 8.5). This follows directly and unavoidably from the legal framework set out in Section 2.5: premiums are unisex, while serious-illness incidence is sex-specific. It is not evidence of mispricing or of any defect in a particular contract; it is the arithmetic consequence of pricing a sex-differentiated risk on a sex-neutral basis. What the analysis adds is that the size of the effect is not uniform across the market. Because the gap is driven by the balance between male-skewed (cardiovascular) and female-skewed (cancer) families, it is widest for the provider whose contractual strengths are cardiovascular (Provider C, a market gap of about 0.28) and narrowest among the leaders for the cancer-strong provider (Provider B, about 0.09). The design of the wording, interacting with unisex pricing, determines how much value is transferred between the sexes.

10.3 The same result read from the insurer's side

Because the fair-value index is the ratio of expected benefit to premium, its reciprocal is proportional to the gross expected margin the insurer retains per premium euro, before expenses, lapse and the effects of underwriting. Read in that direction, and without imputing any intent, the same figures describe which lives are most and least profitable to insure at the unisex price: value falls with age and is lower for women at every age, so the largest gross margins arise on older lives and, at each age, on female lives, while the thinnest arise on younger male lives. This is a factual consequence of the numbers, not a claim about how contracts are designed or marketed; real profitability further depends on expenses, persistency and selection, which lie outside this model. It is noted because it is the mirror image of the consumer-value finding and follows from the same arithmetic.

10.4 Implications for consumers and advice

The results support a research-first, circumstance-aware approach to serious-illness selection. Because value turns on wording, the number of conditions a policy lists is a poor guide to what it is worth; a contract may name many conditions yet define the material ones narrowly. Because value is profile-conditional, the contract that serves a 30-year-old man best is not necessarily the one that serves a 50-year-old woman best. And because unisex pricing means women and older buyers already receive less expected benefit per premium euro, the return to shopping on definition — rather than on price or on the length of the conditions list — is greatest precisely for those buyers. Matching contract structure to the individual's age, sex and family history is therefore not a refinement but the core of suitable advice.

10.5 Implications for disclosure and the market

The analysis illustrates why premium and condition-count, the two dimensions on which serious-illness cover is most often compared, are together insufficient to convey value. A comparison that is auditable — sourced incidence, an explicit wording measure, and computation exposed for inspection — recovers a dimension of value that those two miss. The lower-frequency families examined here (multiple sclerosis (O'Connell et al., 2017), end-stage kidney disease (ERA Registry; Irish National Renal Office)) reinforce the point from the other direction: they are genuine covered conditions but contribute little to expected value, so their presence or absence on a conditions list moves the headline count without materially moving value. Value and coverage breadth are not the same thing, and the difference is measurable.

10.6 Limitations

The findings rest on a base case with disclosed abstractions. The wording measure is an ordinal proxy converted to a cardinal coefficient, and while the provider ordering is invariant to any monotone conversion, the exact width of the value spread depends on the mapping (Section 9). Incidence is modelled on population data rather than on the subset of events that present as claims; the expected present value is un-decremented and pre-expense, so it is a comparative index and not a money's-worth in the strict actuarial sense. Cardiovascular incidence is available only on broad age bands, and several families and early-stage sites are not yet included. The stroke input is calibrated to the national acute-stroke total in the absence of a narrower extract. None of these affects the direction of the structural findings, but each bounds the precision of the levels, and each is documented in the provenance record.

10.7 Interpretive posture

The offices compared operate in a regulated market and pay the majority of the claims they receive; the paper measures structural differences in contractual value without promoting or disparaging any of them, and treats divergence in wording as a fact to be quantified rather than as evidence of misconduct, since it can reflect legitimate differences in target market, underwriting philosophy and benefit-cost design. The author's commercial relationship with the offices analysed, and the methodological safeguards adopted in response, are declared in Appendix E.

---

11. Conclusion

This paper set out to measure how much the choice of provider changes the value of specified serious-illness cover, and where that value comes from. Evaluated across a portfolio of representative lives rather than a single reference life, the answer is consistent and quantified: the value delivered by otherwise-similar contracts differs across the five domestic offices by several times more than their premiums do, at every working age and for both sexes. Value is a property of the wording, not principally of the price, and the number of conditions a policy lists is a poor guide to it.

Two structural findings follow from evaluating the portfolio. First, value is profile-conditional: no office leads across the whole age range, and the ranking below the leader shifts with age as the disease mix moves from cancer-dominated toward cardiovascular. Second, and most consequentially, because premiums are unisex by law while serious-illness incidence is sex-specific, an identical contract at an identical premium delivers less expected benefit to a woman than to a man — a gap that is wider at older ages and is largest where the wording is strongest on the male-skewed cardiovascular families. The same arithmetic, read from the insurer's side, describes which

lives return the greatest margin under unisex pricing. None of this is an allegation of mispricing; it is the measurable consequence of pricing a sex-differentiated risk on a sex-neutral basis, and it is exactly the kind of structure that a premium-and-condition-count comparison cannot see.

The method is the contribution as much as the numbers. The analysis is built in a published workbook from sourced, separately-held incidence and reported with its limitations disclosed — reproducible and open to challenge, its key assumptions sensitivity-tested and shown not to alter the structural findings. That is the standard to which comparative claims about protection value should be held. The immediate practical implication is that suitable advice matches contract structure to the individual's age, sex and family history, and that the return to comparing on definition, rather than on price or the length of the conditions list, is greatest for the buyers — women and older applicants — who under unisex pricing already receive the least.

Appendices

Appendix A — Propositions and proofs

This appendix states, and proves, four properties of the valuation model of Section 6. They are properties of the model's arithmetic rather than empirical claims; their purpose is to make explicit what the fair-value construction implies.

The quantity being analysed. For a life aged x over a term of n years, the expected present value of provider j's serious-illness benefit is the sum — across illness families and policy years — of the benefit expected in each year, discounted to the present:

$$EPV_j \;=\; \sum_{t=1}^{n} \sum_{i} q_i(x,t)\; w_{i,j}\; c_{i,j}(t)\; b_{i,j}(t)\; S\; v^t.$$

The factors are those defined in Section 6: q is the age-specific incidence of the family, a non-negative rate drawn from epidemiology; w is the coverage intensity, a number between 0 and 1 measuring how fully the wording responds to that family; c is the timing factor and b the benefit fraction paid on a trigger, each between 0 and 1; S is the sum assured, which is positive; and v is the annual discount factor, between 0 and 1.

It helps to separate the one factor the wording controls — the coverage intensity — from everything else. Call the rest a family's **full-coverage value**: the value it would contribute if the wording responded to it completely. For a given family and provider this is

$$V_{i,j} \;=\; \sum_{t=1}^{n} q_i(x,t)\; c_{i,j}(t)\; b_{i,j}(t)\; S\; v^t,$$

which is non-negative, because every factor in it is. Each family then contributes its coverage intensity multiplied by its full-coverage value, and the expected value of the whole contract is the sum of those contributions across families. The four propositions read directly off this form.

Proposition 1 — broader wording raises value. With incidence, timing, benefit, sum assured and discount held fixed, the expected value rises, or is unchanged, whenever any coverage intensity rises. The expected value depends on each coverage intensity in a straight line whose slope is that family's full-coverage value, and that value is never negative; a quantity changing at a non-negative rate cannot fall as its input rises.

Proposition 2 — an added payable state can only add value. Take two contracts identical except that the second pays on one further state — a condition, or a stage of a condition, on which the first pays nothing. The second's expected value is at least as high, and strictly higher when that state carries real incidence and a positive benefit. The difference between the two is exactly the contribution of the extra state — its coverage intensity times its full-coverage value — which is non-negative, and strictly positive when the state's incidence, timing, benefit and coverage are all positive.

Proposition 3 — narrowing has a determinate cost. If a stricter definition lowers the coverage intensity of a family, all else equal, the value of that family falls, or is unchanged. The change in the expected value is the reduction in coverage intensity multiplied by the family's full-coverage value; since that value is non-negative and the coverage intensity has fallen, the product is a loss, never a gain.

Proposition 4 — the gap between medical and insured risk widens as wording narrows. For each family, the shortfall between the medically real incidence and the incidence the contract actually pays on is the medical incidence multiplied by the fraction of it the contract does not reach:

$$\mathrm{shortfall}_{i,j}(t) \;=\; q_i(x,t)\,\bigl[\,1 - w_{i,j}\,c_{i,j}(t)\,\bigr].$$

Because coverage intensity and timing each lie between 0 and 1, the bracketed fraction also lies between 0 and 1, so the shortfall is never negative; and since it is the medical incidence times that fraction, it grows in direct proportion as the covered fraction falls. Narrower wording therefore opens a wider gap between the risk a person carries medically and the risk their policy insures.

In sum. Broader wording raises expected value (Proposition 1); an additional payable state with real incidence and benefit adds value (Proposition 2); narrowing carries a determinate cost (Proposition 3); and the divergence between medically real and contractually payable events grows precisely as coverage narrows (Proposition 4). Each is a statement about the model's construction, not an empirical finding.

Appendix B — References

The references are grouped by the evidential hierarchy used throughout the paper: Tier 1 — the whole-of-market claims report and insurers' own primary disclosures and policy conditions; Tier 2 — professional-actuarial, supervisory and reputable epidemiological/registry sources, together with the governing legal instruments; Tier 3 — peer-reviewed academic sources; Tier 4 — internal working material. General journalism and competitor marketing are excluded as evidential authorities.

Tier 1 — market and contract primary sources

- mylife.ie. *Life Insurance Claims in Ireland — 2025* (whole-of-market report). June 2026. [Operative claims source; superseded the 2024 edition.]
- Aviva, Irish Life, New Ireland (Bank of Ireland Life), Royal London Ireland and Zurich Life. Policy conditions and published annual claims disclosures for the five domestic offices, as consolidated in the whole-of-market claims report above.
- mylife.ie. *All-Illness Serious Illness Policy Conditions Ranking*, version 1.2, 27 April 2026 (135 canonical headings; 675 provider-condition entries). Source of the coverage-intensity coefficients.
- mylife.ie whole-of-market quotation system. Whole-of-market quotes for the ten representative lives, prepared 24 July 2026.

Tier 2 — professional, supervisory, registry, epidemiological and legal sources

- World Health Organization. *International Classification of Diseases, Tenth Revision (ICD-10)*. Disease taxonomy, matching the coding of the Irish registry and hospital-audit sources.
- National Cancer Registry Ireland. Incidence-by-age statistics, invasive and in-situ, 2017–2019 (crude age/sex-specific rates); *Cancer in Ireland* annual statistical report (national anchors, 2019–2021).
- International Agency for Research on Cancer. *GLOBOCAN 2022* (Ireland). Cross-check anchor only.
- Healthcare Pricing Office. *Hospital In-Patient Enquiry (HIPE) Annual Report*, Table 3.11: acute myocardial infarction (ICD-10 I21–I22) and cerebrovascular disease (ICD-10 I60–I69) discharges by age and sex, 2017–2019.
- National Office of Clinical Audit. *Irish Heart Attack Audit* (STEMI, cross-check) and *Irish National Audit of Stroke* (national stroke totals and profile).
- ERA Registry (European Renal Association). *Annual Report* (kidney replacement therapy incidence, European average 152 pmp); Irish National Renal Office (Ireland-specific ESKD incidence, ~90–110 pmp).
- Central Statistics Office. PxStat table PEA01 (Estimated Population by Age Group, Sex and Year). Denominators for non-cancer rate derivation.
- Central Statistics Office. PxStat table PEA11 (Estimated Population by Single Year of Age, Sex and Year) and Census of Population 2022 table FY006B (Population by Single Year of Age, Sex and County/City). Source of the market age-sex weighting.
- European Insurance and Occupational Pensions Authority (EIOPA). Monthly risk-free interest-rate term structure (euro). Basis for the discount rate.
- Council Directive 2004/113/EC of 13 December 2004 implementing the principle of equal treatment between men and women in the access to and supply of goods and services. OJ L 373, 21.12.2004, p. 37.
- Court of Justice of the European Union. Case C-236/09, *Association Belge des Consommateurs Test-Achats ASBL and Others v Conseil des ministres*, judgment of 1 March 2011. OJ C 130, 30.4.2011. [Article 5(2) invalid with effect from 21 December 2012; unisex premiums and benefits for new contracts thereafter.]
- European Commission. *Guidelines on the application of Council Directive 2004/113/EC to insurance in the light of Test-Achats*. 2012/C 11/01, OJ C 11, 13.1.2012.

Tier 3 — academic

- O'Connell K, Tubridy N, Hutchinson M, McGuigan C. Incidence of multiple sclerosis in the Republic of Ireland: a prospective population-based study. *Multiple Sclerosis and Related Disorders* 2017;13:75–80. doi:10.1016/j.msard.2017.02.010. PMID 28427707.

Tier 4 — internal / series

- Mylife Working Paper Series: *The Bank Premium* (MWP-2026-02) and *The Decreasing-Term Anachronism* (MWP-2026-03), source of the inherited reference life; *Longevity Insurance* (MWP-2026-04), source of the EIOPA discount convention.
- MWP-2026-05 audit workbook (this paper's computational record).

Appendix C — Data provenance and status of model inputs

Input	Source	Vintage	Status
Coverage intensity w	v1.2 policy-conditions ranking	Apr 2026	Sourced (public ranking); ordinal proxy, sensitivity-tested
Benefit fraction b	v1.2 coverage tiers	Apr 2026	Sourced for tier; partial fraction 0.50 (market-standard); sensitivity-tested (§9.3)
Timing factor c	model assumption	—	Assumption (=1 unless persistence material)
Invasive cancer incidence (sex-specific)	NCRI incidence-by-age	2017–2019	Sourced; reconciles to NCRI 2019–21 anchors and GLOBOCAN 2022
Breast in-situ (DCIS)	NCRI in-situ incidence	2017–2019	Sourced
Cervix / skin / melanoma in-situ	NCRI in-situ	2017–2019	Excluded, with reasons documented
Heart attack (AMI)	HIPE Table 3.11, principal dx I21–I22	2017–2019	Sourced; **broad age bands**; discharge (not person) counts
Stroke	HIPE I60–I69, calibrated to the INAS acute-stroke total (factor ≈0.79, computed by formula)	2017–2019	Acute-adjusted; the I60–I69 overcount is removed by calibration to the national acute-stroke total (§9.2)
Multiple sclerosis	O'Connell et al. 2017	2014–	Modelled working-age from the national rate (O'Connell 2017)

		2015	
Kidney failure (ESKD)	Irish National Renal Office (~90–110 pmp)	2017–2023	Modelled working-age (M/F 8/5 per 100k) on the Irish incidence basis; immaterial to ranking
Premiums (by age 30-50)	mylife whole-of-market quotation system	24 Jul 2026	Sourced; guaranteed (non-reviewable) basis confirmed for all five offices; life-only and life+SI captured, SI cost derived by formula
Discount rate r	EIOPA euro RFR neighbourhood	—	3.0% central, EIOPA euro basis; tested at 2%/4% (§9.1)
Population denominators	CSO PxStat PEA01	matched years	Sourced
Market weighting (age-sex)	CSO PxStat table PEA11 (Population by Single Year of Age, Sex and Year); Census 2022 table FY006B	latest	CSO population by age and sex; market results robust to the weighting (equal-weighted check, §8.6)
Portfolio engine	audit workbook (Portfolio, Prem_portfolio)	—	Fair value by age and sex; reconciles to the single-life engine (diff = 0)
Reference life / sums assured	MWP-2026-02/03; 50% rule	—	Sourced (series) / stated assumption

Appendix D — Data notes and modelling assumptions

Appendix C summarises the model's inputs and their treatment. This note records the principal modelling assumptions and the checks that bound their influence; each is examined in the sensitivity analysis of Section 9.

- **Wording coefficient.** The coverage-intensity coefficient is an ordinal proxy derived from a single published conditions ranking; the ordinal placement of providers is the paper's most externally-dependent input (Section 5.3), and the results are conditional on it. An independent re-ranking would strengthen its external validity.
- **Partial-payment fraction.** Partial-payment benefits are valued at 0.50 of the sum assured, the market-standard fraction; on the current family scope the results are insensitive to this choice (Section 9.3).
- **Cerebrovascular incidence.** Stroke is taken from the I60–I69 hospital grouping calibrated to the national acute-stroke total (Section 9.2); this is the single largest calibration adjustment.
- **Discount rate.** Benefits and premiums are discounted at a 3.0% real central rate on the EIOPA euro basis, tested at 2% and 4% (Section 9.1); because both sides of the ratio are discounted identically, the rate moves the level of the index but not the comparisons.
- **Market weighting.** The ten representative lives are weighted to a market view using CSO population by age and sex; the market results are robust to the weighting, agreeing with an equal-weighted alternative to within 0.01 (Sections 8.6 and 9.5).
- **Disease classification.** ICD-11 has been the mandated WHO revision since 1 January 2022; ICD-10 is used throughout because the Irish incidence sources (the NCRI and HIPE) are coded in it, so that taxonomy and incidence share one coding system (§3.1).

None of these assumptions affects the structural findings; Section 9 sets out the tests in full.

Appendix E — Declaration of interests

SMP Financial, trading as mylife.ie, holds commercial agency appointments with each of the five domestic life offices analysed in this paper and earns commission on the protection business it places with them. This is an inherent conflict of interest, disclosed here so that the analysis may be judged on its method rather than on the author's independence.

The paper's safeguards against that conflict are methodological, and are set out in the body: a single framework applied uniformly to every office (Sections 5 to 7); a closed, stated evidential hierarchy and a strict separation of epidemiological incidence from claims experience (Section 2); computation exposed in full in the audit workbook; and anonymised presentation of provider-level results (Provider A–E), with the identity key held in the workbook. The offices analysed operate in a regulated market and pay the majority of the claims they receive; the paper's aim is to measure structural differences in contractual value without promoting or disparaging any of them.

About the author

Donal Milmo-Penny, QFA FLIA, is Research Lead at mylife.ie. He has more than twenty years of experience in Irish financial services, protection and client advisory work. He is a Qualified Financial Adviser (QFA) and a Fellow of the Life Insurance Association (FLIA), and is a former Chairman of PIBA and Director of Brokers Ireland.

About the mylife.ie Working Paper Series

The mylife.ie Working Paper Series publishes original analysis of the Irish life assurance, mortgage-protection and serious-illness markets, to inform consumers, policymakers and the wider research community. It is published by SMP Financial Ltd, trading as mylife.ie, regulated by the Central Bank of Ireland (C42382). Papers are made available without restriction and do not constitute regulated financial advice.

Use of AI

This paper was prepared with the assistance of an AI system, used under the author's direction for drafting, structuring and computational support. It was not used to generate data: the incidence, premium and claims figures derive entirely from the sources listed in Appendix B, and every sourced figure has been checked against its primary source. All quantitative results are computed in the accompanying audit workbook and are reproducible from the disclosed inputs; the workbook, not the prose, is the computational record. The author retains full responsibility for the method, the analysis, the findings and any errors.

Copyright and disclaimer