Institute and Faculty of Actuaries

Behavioural Economics for Insurance Nudges

Quantifying the Actuarial Value of Prevention Interventions

by A. Nanda*

Presented to The Big Pitch, Institute and Faculty of Actuaries, 18 March 2026, London

Title

Behavioural economics for insurance nudges

Author/Authors

A. Nanda*

Abstract

Health insurers systematically underinvest in prevention because costs are immediate while benefits accrue over years—yet actuaries lack a formal mechanism to translate behavioural intervention evidence into pricing-ready claims adjustments. This paper introduces the Behavioural Adjustment Factor (BAF)—a multiplicative actuarial framework that quantifies the claims impact of behavioural interventions by decomposing reach, efficacy, clinical translation, and durability into a single pricing-ready construct. To the best of the author's knowledge, the BAF is the first actuarial framework to decompose behavioural intervention impact into condition-specific claims projections suitable for pricing and reserving.

Drawing on randomised controlled trial evidence, well-designed interventions can triple smoking cessation rates, double medication adherence, and reduce hospital admissions by 29%—but programme architecture fundamentally determines population-level impact. Disease management targeting existing chronic conditions generates robust returns (ROI 3.78:1), while lifestyle management does not reliably produce claims savings (ROI 0.48:1). A worked hypertension example yields a BAF range of 0.82–0.94, with the conservative end recommended as the default pending UK calibration. The framework provides confidence intervals, Monte Carlo integration for Solvency II, milestone-based pilot funding, and a clear path from evidence to practice.

Keywords

Behavioural adjustment factor; Value-based insurance design; Chronic disease management; Reach–efficacy frontier; Actuarial modelling

Correspondence details

*Correspondence to: Ankit Nanda, FSA, MAAA, FCA. E-mail: anknanda@gmail.com

1. Introduction

Health insurers face a fundamental temporal mismatch: the costs of preventive interventions are incurred immediately, while the benefits—reduced morbidity, lower claims, reduced mortality—accrue over years or decades. This mismatch creates systematic underinvestment in prevention, even when the long-term actuarial case is compelling (Baicker, Cutler and Song, 2010). Policyholders exhibit the same temporal bias, discounting future health benefits heavily relative to immediate costs.

Consider the economics: a single hospitalisation for diabetic ketoacidosis costs £10,000–£20,000 (NHS England, 2023), yet medication adherence rates for statins fall to approximately 50% within one year of initiation (Lemstra et al., 2012). Cardiovascular disease generates over £9 billion in annual NHS spending. The actuarial question is not whether prevention has value—the clinical evidence is overwhelming—but rather how to quantify and price the behavioural interventions that make prevention actionable.

Behavioural economics identifies predictable cognitive biases—present bias, loss aversion, status quo bias—that cause individuals to make choices inconsistent with their long-term interests. These biases can be leveraged to design interventions, often called nudges (Thaler

and Sunstein, 2008), that guide individuals toward healthier behaviours without restricting choice.

Despite this evidence, actuaries lack a formal mechanism to translate behavioural effects into pricing-ready claims adjustments. Existing approaches—health economics programme evaluations that stop at effect sizes, wellness vendor ROI claims that conflate correlation with causation—do not produce outputs compatible with actuarial pricing, reserving, or capital standards. This paper proposes the Behavioural Adjustment Factor (BAF) to fill that gap: a multiplicative construct that decomposes reach, efficacy, clinical translation, and durability into a single adjustment factor suitable for condition-specific claims projections, with uncertainty quantified throughout for Solvency II compatibility.

No prior actuarial framework, to the author's knowledge, decomposes behavioural intervention impact into this form. Health economics programme evaluations assess cost-effectiveness but do not produce pricing factors. Vendor wellness ROI models typically conflate disease management and lifestyle effects, lack component-level decomposition, and cannot map to Solvency II capital requirements or satisfy APS X2 documentation requirements. The BAF addresses each of these limitations.

The remainder of this paper develops and applies the BAF framework. Section 2 outlines the research methodology. Section 3 reviews the evidence base. Section 4 develops the BAF framework with a worked example. Section 5 examines condition-specific applications. Section 6 addresses ethical considerations. Section 7 proposes pilot studies. Section 8 discusses implications for actuarial practice. Section 9 addresses limitations. Section 10 concludes. Pricing actuaries will find the BAF formula, worked example, and net pricing impact (Section 4) most immediately applicable; reserving actuaries should attend to the durability modelling and 3-Year Rule; and product designers and regulatory stakeholders will find the ethical framework (Section 6) and pilot roadmap (Section 7) most relevant.

2. Research Methodology

The research employs a structured synthesis-to-framework approach: a systematic review of empirical evidence followed by the construction of an applied actuarial model grounded in that evidence. The evidence base is drawn from peer-reviewed RCTs published in leading journals, including the New England Journal of Medicine, Health Affairs, and the Journal of the American College of Cardiology. Studies were selected for methodological rigour, sample size (≥200), and relevance to insurer-actionable interventions. All monetary values have been converted to GBP (at USD 1.25 per GBP).

A formal meta-analysis was considered but judged impracticable given the heterogeneity of endpoints across studies. The parameter ranges presented are derived from structured qualitative synthesis with explicit inclusion criteria rather than from formal pooling. The multiplicative structure was selected over aggregate approaches because it isolates the contribution of each stage, enabling actuaries to update individual parameters as new evidence emerges without re-estimating the entire model. Alternative approaches considered but discounted include aggregate wellness ROI modelling (rejected for conflating disease and lifestyle management effects) and difference-in-differences designs—a quasi-experimental

method that compares the change in outcomes over time between a treatment group and a control group, netting out pre-existing trends to isolate the causal effect of an intervention (rejected as requiring insurer-held claims data not publicly available for initial calibration). Parameter distributions (Beta, Log-normal, Gamma) are assigned to each component to enable Monte Carlo simulation for uncertainty quantification compatible with Solvency II capital requirements.

3. Literature Review

3.1 Core behavioural principles

Traditional economic theory assumes rational utility maximisation. The empirical reality is starkly different: medication non-adherence affects 40–60% of patients with chronic conditions (Osterberg and Blaschke, 2005). Three cognitive biases are most relevant, drawing on work recognised by the Nobel Memorial Prize in Economic Sciences awarded to Richard Thaler in 2017 (Royal Swedish Academy of Sciences, 2017). Present bias: individuals disproportionately weight immediate costs relative to future benefits. A patient may understand that taking medication reduces heart attack risk in ten years, but the immediate inconvenience looms larger; interventions providing immediate rewards can counteract this. Loss aversion: losses loom larger than equivalent gains, typically by 2:1 (Kahneman and Tversky, 1979). Deposit contracts, where participants risk their own money, prove more effective than equivalent rewards among those who accept them. Status quo bias: individuals stick with default options; changing from opt-in to opt-out enrolment can double participation rates without altering the underlying choice set (Madrian and Shea, 2001).

3.2 Evidence from randomised controlled trials

Financial incentives for smoking cessation. Volpp et al. (2009) randomised 878 employees to financial incentives totalling £600 versus information alone. Quit rates at 9–12 months were 14.7% versus 5.0% (adjusted OR 3.16; $p < 0.001$), implying ROI of approximately 7:1. Halpern et al. (2015), in a five-arm RCT with 2,538 participants, reveal the reach–efficacy frontier (Table 1).

Intervention Architecture	Reach	Efficacy	Population Effect	Source
Reward-based (£640)	90%	15.7%	14.1%	Halpern et al. (2015)
Deposit-based (£120 + £520)	14%	52.3%	7.3%	Halpern et al. (2015)
VBID copay reduction	100% (automatic)	7–14% ↓ nonadherence	7–14%	Chernew et al. (2008)
Disease management	~60%	£109 PMPM savings	~£65 PMPM	Caloyeras et al. (2014)

Vouchers (pregnancy, UK)	~85%	27% vs 12% quit	~23%	Tappin et al. (2022)

Table 1. The reach–efficacy frontier. (Sources: Halpern et al., 2015; Chernew et al., 2008; Caloyeras et al., 2014; Tappin et al., 2022)

The reach–efficacy frontier presents insurers with a strategic optimisation problem: reward-based programmes maximise population impact; deposit-based programmes maximise per-participant ROI; VBID copay reductions maximise automatic coverage. Mitchell et al. (2023), in a network meta-analysis of financial incentives for physical activity, corroborate the finding that incentive architecture matters more than incentive magnitude. The BAF framework makes this strategic choice explicit rather than resolves it.

UK evidence from Tappin et al. (2022; CPIT III), randomising over 1,000 pregnant smokers across multiple NHS sites to standard Stop Smoking Services versus services plus up to £400 in vouchers contingent on biochemically verified abstinence, provides the strongest direct UK calibration anchor. Quit rates at end of pregnancy were 26.8% versus 12.3% (adjusted OR 2.78, 95% CI 1.94–3.97), with effects persisting at 6–12 months post-partum. NICE-aligned evaluations show the incentives are highly cost-effective even allowing for substantial post-partum relapse. Critically for equity (Section 6), results were broadly consistent across Index of Multiple Deprivation quintiles.

Medication adherence and value-based insurance design. Chernew et al. (2008) studied VBID copay reductions of approximately 50% for five chronic medication classes. Adherence increased 7–14% versus controls with identical disease management—demonstrating that copay reductions and disease management are complements, not substitutes. The BETTER-BP trial (Dodson et al., 2025) randomised 400 hypertension patients to mHealth incentive lotteries, achieving twice the adequate adherence rate (71% versus 34%; RR 2.04; 95% CI 1.58–2.63). Notably, this adherence gain did not produce a statistically significant systolic blood pressure reduction (−6.7 vs −5.8 mmHg, $p \approx 0.62$), underscoring that E and C must be treated as distinct parameters rather than assumed to co-move. Separately, meta-analyses of digital adherence interventions in high-income settings report pooled adherence improvements of approximately 10–23 percentage points (Thakkar et al., 2016), providing a digital baseline against which lottery-based interventions sit at the upper end of the plausible efficacy range.

The disease versus lifestyle distinction. Caloyeras et al. (2014) reported seven years of PepsiCo data with 22,204 matched pairs. Disease management achieved £109 PMPM savings, 29% hospital admission reduction, and ROI of 3.78:1. Lifestyle management showed no significant savings and ROI of 0.48:1. The ‘wellness saves money’ narrative must be disaggregated: disease management works; lifestyle management does not reliably translate to claims savings.

In summary, the evidence supports five intervention archetypes with distinct actuarial profiles: (i) disease management for existing chronic conditions (ROI 3.78:1; high confidence); (ii) financial incentives for smoking cessation (ROI 5–7:1 per quitter; moderate UK transferability with direct CPIT III evidence); (iii) VBID copay reductions for medication adherence (7–14% nonadherence reduction; moderate transferability); (iv) preventive

screening nudges via defaults and reminders (10–20pp uptake increase; strong UK transferability); and (v) general lifestyle wellness (ROI 0.48:1; no reliable claims savings). This hierarchy structures the BAF confidence intervals and tiered ROI expectations.

3.3 UK calibration and generalisability

The evidence base is drawn predominantly from US-based studies. This section clarifies which BAF components can be informed by UK evidence today. Reach: UK RCTs consistently demonstrate material uptake increases—a cluster RCT of behaviourally informed NHS Health Check invitations increased uptake from 18.2% to 30.0% (Sallis et al., 2016), and Hafner et al. (2018), analysing over 400,000 Vitality participants, found loss-framed incentives produced a 34% sustained activity increase over 24 months. Efficacy: UK primary care datasets (CPRD, QResearch, OpenSAFELY) provide evidence; Boonmanunt et al. (2023), synthesising 35 RCTs, found deposit contracts achieved RR 1.63 with post-intervention persistence. Clinical translation: Strong UK evidence via NICE guidance and QOF-linked outcomes; NHS App pilots and opt-out enrolment designs show directionally consistent effects across a range of preventive services (NHS Digital, 2023). Claims impact: Limited public evidence, but Vitality (2024) reports that highly engaged members claim approximately 28% less than less-engaged members. Durability: Limited UK-specific evidence beyond 24 months; this is the primary gap requiring insurer-led pilots.

Transportability approach. To translate international RCT estimates into UK PMI priors, trial effect estimates for Reach and Efficacy are reweighted by observable demographics (age band × comorbidity count × postcode-level deprivation tercile) to a reference UK PMI distribution. Where microdata are unavailable, two bounding scenarios are presented—a 'younger/healthier' reweight and an 'older/more comorbid' reweight—to show plausible directional shifts. Insurers can replace the reference weights with their own member distribution.

4. The BAF Framework

This section develops the BAF framework: defining components, justifying the multiplicative structure, formalising durability, presenting confidence intervals, specifying stochastic distributions, and demonstrating application through a worked hypertension example.

Table 2 defines the notation used throughout.

Symbol	Definition	Range
R	Reach: proportion accepting the intervention	0–1.0
E	Efficacy: behavioural change conditional on participation	0–1.0
C	Clinical Translation: claims reduction per unit behavioural change	0–1.0
D	Durability: persistence of effect at time t	0–1.0
λ	Monthly exponential decay rate	0.03–0.08

h	Habit formation proportion	0.15–0.50

Table 2. BAF framework notation.

Figure 1. The BAF framework pipeline.

Reach (R) → Efficacy (E) → Clinical Translation (C) → Durability (D) → BAF = 1 − (R × E × C × D)

BAF = 1 − (R × E × C × D) ... (1)

Adjusted Claims = Baseline Claims × BAF ... (2)

where each component maps to a distinct evidence domain. The framework is appropriate for condition-specific projections and wellness-embedded pricing; it is inappropriate for aggregate wellness loading or short-term (<3-year) lifestyle projections.

The Clinical Translation component (C) warrants specific comment, as it represents the most heterogeneous stage. C maps behavioural change to claims reduction through a condition-specific dose–response function. Because the marginal health benefit of additional adherence typically diminishes once a patient approaches clinical targets, the recommended base case adopts a non-linear diminishing-returns specification: $C(\Delta) = \alpha_0 \times (1 - \exp(-\gamma \times \Delta))$, where Δ is the adherence gain in percentage points, α_0 is the maximum attainable claims reduction at full adherence, and γ controls the curvature. This form is biologically plausible—for hypertension, the relationship between blood pressure reduction and cardiovascular events follows a log-linear gradient (Lewington et al., 2002). Chen et al. (2022), in a meta-analysis of 402,201 patients, report that a 20pp increase in cardiovascular medication adherence is associated with 8% lower CV events and 30% lower MI—providing an empirical anchor for α_0. The linear approximation $C \approx \alpha \times \Delta$ should be retained as an upper-bound sensitivity case. For hypertension (37pp adherence gain), the non-linear form yields $C \approx 0.45–0.55$ versus 0.60 under linearity. The functional form of C should vary by condition: for cardiovascular adherence, the diminishing-returns specification is recommended as the base case; for general wellness and lifestyle interventions, the linear form should be retained, since effect sizes are typically small and operate within the approximately linear portion of the dose–response curve; for diabetes (HbA1c-driven), a threshold-saturation model is biologically plausible but calibration requires individual-level claims-linked HbA1c data not yet available in the UK PMI context.

A design choice warrants explanation. The framework maintains four components, even though current evidence often supports calibration of E×C as a single composite. Maintaining the four-component structure ensures that, as claims-linkage data accumulate from insurer-led pilots, actuaries can update E and C independently—and can identify where in the causal chain a programme underperforms.

4.1 What actuaries uniquely contribute

The BAF delivers distinctive actuarial value: it translates to pricing and reserving standards (not just programme evaluation); its confidence intervals support Solvency II capital calculations; it integrates with existing IBNR, trend, and large claims processes; it addresses

adverse selection from voluntary participation; and it provides the assumption documentation structure required by APS X2 (IFoA, 2016).

4.2 Multiplicative structure

The BAF represents a sequential causal pipeline: an intervention must be accepted, must change behaviour, must translate clinically, and must persist. Each component is defined conditionally on the previous stage, so the product represents a valid factorisation. If any stage fails, the overall effect is zero. This mirrors standard actuarial decompositions of exposure, frequency, and severity.

Interactions can arise in practice—durability may depend on initial efficacy intensity; high-reach programmes may exhibit different decay profiles. To address this, Monte Carlo simulation propagates correlated uncertainty via a Gaussian copula, confidence intervals are calibrated from end-to-end evaluations, and the framework targets condition-specific application rather than portfolio-wide aggregation.

4.3 The 3-Year Rule

Claims savings from behavioural interventions in chronic disease typically require 36+ months to materialise. This paper formalises this lag as 'The 3-Year Rule,' synthesising findings from Caloyeras et al. (2014), who found cost differences became statistically significant only after Year 3. An important exception: acute post-event (secondary prevention) adherence interventions may show savings within 12–18 months. Pilots of less than 24 months should use intermediate endpoints. Pricing should not assume immediate impact.

4.4 Durability and effect decay

Effect persistence is modelled with a habit formation component structurally analogous to ultimate improvement factors in mortality projection:

Effect(t) = Effect_initial × [(1 − h) × e^(−λt) + h] ... (3)

where λ is the monthly decay rate and h is the habit formation proportion. The Durability component D in the BAF formula is $D(t) = (1 - h) \times \exp(-\lambda t) + h$.

The practical impact is substantial. Under base assumptions (λ=0.05, h=0.30), effect retention at 24 months is 51%. Under pessimistic assumptions (λ=0.08, h=0.15), retention falls to 27%. In plain terms: a nudge that initially reduces cardiovascular claims by £100 PMPM would, under base assumptions, still save £51 PMPM at the two-year mark; under pessimistic assumptions, only £27 PMPM. The entire viability of a behavioural intervention programme can therefore turn on durability parameters for which UK-specific evidence is thinnest.

Critically, habit formation varies systematically by intervention mechanism. Deposit/commitment contracts show the highest persistence ($h \approx 0.35–0.50$), supported by Boonmanunt et al. (2023) finding that only deposit contracts maintain statistically significant post-intervention effects. Loss-framed incentives show $h \approx 0.30–0.45$, consistent with Hafner et al. (2018) finding 34% sustained activity increase at 24 months. Reward-based incentives

show lower persistence (h ≈ 0.15–0.30), with Halpern et al. (2015) reward arms showing faster decay. Default/opt-out nudges show the highest persistence (h ≈ 0.40–0.55) due to status quo bias. VBID copay reductions (h ≈ 0.20–0.35) are contingent on benefit design persistence. Evidence from the Vitality Habit Index (Vitality and LSE, 2024) suggests that durability may also vary by age cohort, with older (65+) chronic-disease cohorts exhibiting higher effective habit formation than younger adults under similar incentive designs—a finding with direct pricing relevance for PMI. These ranges are inferred from study-level persistence patterns rather than directly estimated; calibrating h for specific mechanisms is a priority for the pilot studies proposed in Section 7.

Durability stress scenario. Under h = 0.10 and λ = 0.10 (rapid decay), D(24 months) = 0.18, and the hypertension BAF rises to 0.94—confirming the programme is not viable under rapid decay. This is not presented as likely, but establishes the durability floor below which the business case breaks down.

4.5 BAF confidence intervals and stochastic modelling

Point estimates overstate certainty. Table 3 presents BAF ranges alongside stochastic distribution specifications for Solvency II capital modelling.

Intervention	BAF (Point)	95% CI	Component Distributions
Disease management (chronic)	0.86	0.79–0.93	R ~ Beta(α,β); E×C ~ Log-normal; λ ~ Gamma
Smoking cessation incentives	0.91	0.85–0.96	Monte Carlo with Gaussian copula for correlated draws
VBID medication adherence	0.94	0.89–0.98	VaR/TVaR at 95th percentile for capital relief
Preventive screening nudges	0.95	0.91–0.98	
Lifestyle wellness	1.00	0.95–1.05	CI spans 1.00: positive ROI cannot be assumed

Table 3. BAF estimates, confidence intervals, and stochastic modelling approach.

Note: BAF < 1.00 = claims reduction. Ranges derived from end-to-end programme evaluations. Where end-to-end evidence is unavailable, ranges reflect the author's actuarial judgement informed by component-level evidence—making the synthesis evidence-informed rather than evidence-based in the strict sense. (See Appendix B for evidence strength ratings by parameter.)

For external communication (board presentations, regulatory submissions, policyholder disclosures), actuaries may wish to present a deliberately wider 'public conservative' interval that incorporates transferability uncertainty not fully reflected in the evidence-based CI. Table 3a presents both intervals side by side.

Intervention	BAF (Point)	Evidence CI	Public CI	Widening Rationale
Disease management	0.86	0.79–0.93	0.75–0.95	Case-mix + claims translation

Intervention	BAF (Point)	Evidence CI	Public CI	Widening Rationale
Smoking cessation	0.91	0.85–0.96	0.82–0.96	Relapse + UK prevalence
VBID adherence	0.94	0.89–0.98	0.86–0.98	Formulary + copay differences
Lifestyle wellness	1.00	0.95–1.05	0.90–1.10	No assumed savings

Table 3a. Evidence and public conservative confidence intervals by intervention (illustrative).

Note: Public conservative CI widens the evidence CI by approximately 3–5 percentage points to accommodate cross-jurisdictional transferability, claims-translation uncertainty, and population mix variation. The lifestyle public CI explicitly spans 1.10, signalling that net costs (not savings) remain plausible. Insurers with access to UK claims data should calibrate their own widening factors.

For the hypertension worked example (Section 4.6), illustrative parameterisations are: $R \sim \text{Beta}(81, 9)$ centering on 0.90; $E \times C \sim \text{Log-normal}(\mu=-0.92, \sigma=0.15)$ centering on 0.40; $\lambda \sim \text{Gamma}(\text{shape}=6.25, \text{scale}=0.008)$ centering on 0.05. A minimum of 10,000 Monte Carlo iterations is recommended. Sensitivity runs should introduce correlation structures—recommended scenarios include $\rho(E,D) \in \{-0.3, 0, +0.3, +0.6\}$ and $\rho(R,E) \in \{-0.3, 0, +0.3\}$—imposed via a Gaussian copula. Illustrative simulation results confirm robustness: under moderate negative $\rho(R,E) = -0.3$, the 95th-percentile BAF rises by approximately 0.02–0.03 from the independence case.

Illustrative Solvency II mapping. Suppose a UK PMI insurer writes £100m in health premium with cardiovascular claims comprising 15% (£15m). The hypertension BAF point estimate of 0.82 applied to the targeted subpopulation (20% of CV members enrolled) yields an expected claims saving of £0.54m. The BAF VaR_{95} of 0.89 gives a conservative reduction of £0.33m. The difference (£0.21m) represents the additional capital buffer required if the intervention underperforms—a statement directly compatible with Solvency II capital requirements.

Net pricing impact. To embed the BAF in a pricing formula:

BAF_Net = BAF − (Programme Cost ÷ Target Claims Base) ... (4)

If BAF_Net > 1.0, the programme costs more than the savings it produces—a simple go/no-go check.

In practice, programme costs are not deterministic: per-member costs vary with uptake rates, incentive redemption, and platform scaling, while claims savings are uncertain and lagged. For full integration with the Monte Carlo framework, actuaries should model programme costs as a distribution (e.g., Gamma or Normal, parameterised from vendor contracts and historical programme data) and draw costs jointly with BAF components in each simulation iteration, producing a BAF_Net distribution rather than a point estimate. This is particularly important for go/no-go decisions at the pessimistic tail, where cost overruns and BAF underperformance may coincide.

4.6 Worked example: Hypertension adherence nudge

This example demonstrates how a seemingly large behavioural effect becomes a modest claims adjustment once population reach, clinical translation, and effect decay are applied.

Consider a reward-based mHealth lottery (per BETTER-BP; Dodson et al., 2025) for hypertensive PMI members:

Reach: 90% (reward-based acceptance per Halpern et al., 2015).

Efficacy: 37 percentage point adherence increase (71% versus 34% in BETTER-BP).

Clinical translation: Two bounding scenarios illustrate why the E/C distinction matters—indeed, the BETTER-BP trial itself demonstrates the distinction in operation. *Optimistic* (E×C = 0.40): if the 37pp adherence gain translates to clinical outcomes in line with epidemiological dose-response gradients—approximately 0.5 mmHg per 1pp adherence, 18.5 mmHg SBP reduction, and weighted 40% CV claims reduction (Lewington et al., 2002; Chen et al., 2022)—the resulting BAF is 0.82 (18% claims reduction). This scenario represents the upper bound, appropriate only if adherence translates to sustained BP control at population scale. *Conservative* (E×C = 0.15–0.20): consistent with the BETTER-BP finding that adherence gains did not produce significant blood pressure reduction in-trial, and the well-documented gap between surrogate endpoint improvements and downstream event reduction in short-horizon trials, yielding BAF = 0.92–0.94 (6–8% claims reduction, assuming weaker durability under this scenario). This conservative scenario—the Prudent Default—is recommended for initial pricing and reserving until UK claims-linkage data are available. A UK-adjusted mid-case (E×C = 0.30, applying a 25% haircut for prescribing patterns and GP gatekeeping) yields BAF ≈ 0.86 (14% reduction, net savings ~£20 PMPM). The plausible actuarial range is therefore BAF = 0.82–0.94.

Durability: h=0.3, λ=0.05 → 51% effect retention at 24 months.

BAF calculation: Optimistic: BAF = 1 − (0.90 × 0.40 × 0.51) = 0.82; 95% CI 0.71–0.89. Conservative: BAF = 1 − (0.90 × 0.17 × 0.51) = 0.92. For initial pricing and reserving, the conservative scenario is the recommended default.

Interpretation: Baseline cardiovascular claims of £500 PMPM would adjust to £410–£470 PMPM for the intervention cohort, a reduction of £30–£90 PMPM against programme costs of approximately £50 PMPM. Table 4 illustrates durability sensitivity.

Scenario	h	λ	D (24m)	BAF	Net Savings
Pessimistic	0.15	0.08	0.27	0.90	−£1 PMPM
Base case	0.30	0.05	0.51	0.82	£42 PMPM
Optimistic	0.50	0.03	0.74	0.73	£84 PMPM
Rapid decay	0.10	0.10	0.18	0.94	−£18 PMPM

Table 4. Sensitivity of hypertension BAF to durability assumptions (R=0.90, E×C=0.40). Net savings = (£500 − Adj PMPM) − £50 programme costs. The rapid decay scenario establishes a quantitative floor: the BAF signals 'do not proceed' unless h ≥ 0.15.

5. Applications

Preventive screening. Increasing colorectal screening from 45% to 57% shifts detection toward earlier stages: Stage I treatment costs £15,000–£20,000 versus £100,000+ for Stage IV (NICE, 2020). BAF: 0.92–0.95 for colorectal claims over 5 years. Annual diabetic retinal exams detect retinopathy before vision loss (NHS Diabetic Eye Screening Programme, 2023) (early treatment £1,000–£2,000 versus £20,000+ advanced).

Chronic disease management. Per Caloyeras et al. (2014), disease management generates ROI of 3.78:1. The UK Prospective Diabetes Study Group (1998) demonstrated that intensive blood-glucose control reduced diabetes-related complications by 12% and microvascular endpoints by 25%, establishing the causal pathway from adherence to claims reduction. The 29% hospital admission reduction is the single most powerful claims impact figure in the evidence base—affecting both frequency and severity development in reserving.

Mental health (emerging). Treatment adherence rates for depression and anxiety are notably low (40–50%). Recent meta-analyses suggest mHealth nudges may yield 5–10pp adherence improvements, although the evidence base is less mature (Goldberg et al., 2022).

Directional UK PMI categories. Cardiovascular claims are the highest-confidence BAF application (high per-event severity, strong adherence-to-event evidence). Diabetes-related claims span routine management and complications, with the 29% admission reduction providing the clearest frequency signal. Musculoskeletal claims present a more complex application with wider uncertainty bands pending condition-specific data.

6. Ethical Considerations

Behavioural interventions raise ethical questions that actuaries must address in product governance documentation, framed through the IFoA's public interest duty and the FCA's Consumer Duty (PS22/9, effective July 2023).

Autonomy and transparency. Nudges preserve freedom of choice, but three principles are proposed: transparency (members know interventions are deployed), genuine optionality (easy opt-out), and alignment of interests.

Actuarial fairness and selection risk. If a pricing model already uses 'disease management participation' as a rating factor, applying a BAF would double-count savings. Actuaries must document which effects are captured in underwriting versus BAF adjustments. For separating causal impact from self-selection, three approaches are recommended: intention-to-treat estimates from randomised encouragement trials (highest rigour, consistent with the evidence hierarchy in NICE, 2013); propensity-score stratification matching on pre-intervention claims trajectories; or applying the BAF at its 95th-percentile worst-case value as a conservative default.

Equity and differential response. Tudor Hart's (1971) inverse care law—that the availability of good medical care tends to vary inversely with the need for it—has an uncomfortable commercial analogue. Behavioural nudge programmes are most easily adopted by members who least need them: higher-income, higher-literacy, lower-comorbidity individuals with the disposable income to commit deposits and the cognitive bandwidth to

respond to incentive framing. The members with the highest chronic disease burden and the greatest potential for claims reduction are, on average, those least likely to participate.

Halpern et al. (2015) found deposit-based nudges—the highest-efficacy mechanism—exhibited approximately 20% lower uptake in lower socioeconomic groups. Michie et al. (2009) found financial incentive programmes disproportionately attracted participants with greater health literacy. Hart's original concern was supply-side distribution of healthcare services; the parallel here operates on the demand side, through differential uptake of voluntary programmes. The underlying equity concern is the same—those with the greatest need benefit least—even though the mechanism differs. Default-based mechanisms tell a different story: Madrian and Shea (2001) demonstrated that automatic enrolment nearly eliminated participation gaps across income quintiles. Without deliberate design safeguards, BAF-based pricing produces a cross-subsidy flowing in the wrong direction: non-participants bear a proportionately larger share of costs, and non-participation correlates with deprivation. The FCA's Consumer Duty (PS22/9) requires firms to deliver good outcomes for all retail customers; a programme that systematically excludes lower-income members raises questions across the Duty's price-and-value and products-and-services outcomes. Table 5 proposes design principles for equitable implementation.

Principle	Implementation
Default-first architecture	Use universal defaults as primary pathway; offer deposit/loss-framed options as supplements
Tiered incentive calibration	Graduate reward values by deprivation (analogous to Pupil Premium)
Stratified BAF reporting	Report BAF by deprivation tiers; where data unavailable, reduce R by 15–25% for most deprived quintile
Multichannel delivery	Offer telephone, in-person, and culturally adapted modalities alongside digital

Table 5. Design principles for equitable BAF implementation.

Actuaries should track participation and outcomes by deprivation quintile annually. Where the bottom quintile shows participation below 60% of the top—a threshold at which illustrative pricing simulations indicate the cross-subsidy from non-participants exceeds approximately 5% of baseline premium, triggering material fairness concerns under Consumer Duty—the pricing actuary should flag this in the Actuarial Function Report and recommend design modifications. An amber threshold at 0.60–0.75 should trigger targeted outreach and channel diversification; a red threshold below 0.60 should trigger suspension of differential pricing pending design review. To operationalise this monitoring beyond the participation ratio, actuaries should also track: early attrition rates by deprivation quintile as a leading indicator, since differential attrition can erode equity even where initial uptake is balanced; and interim behavioural endpoints (adherence rates, screening completion) by quintile, which are observable within 6–12 months and do not require waiting for the 3-Year Rule claims horizon to detect inequitable programme effects. This is consistent with the IFoA's Ethical Charter commitment to consider the public interest implications of actuarial work. A deeper tension remains unresolved: if the most effective interventions (deposit contracts, loss-framed incentives) systematically exclude lower-SES members, and if the most equitable interventions (defaults, VBID) have lower per-participant efficacy, the actuary

faces a genuine trade-off between maximising population health impact and maximising ROI. Whether the actuary's duty under the FCA's Consumer Duty and the IFoA's public interest obligation requires prioritising equity over marginal efficiency gains is a question that this paper surfaces rather than resolves; it is an appropriate subject for future IFoA professional guidance on behavioural intervention pricing.

7. Pilot Study Design

The target population comprises members with chronic conditions and proportion of days covered (PDC) below 80%. A minimum of 2,000 per arm is required to detect a 5pp improvement with 80% power (two-proportion z-test, baseline $p_0 = 0.50$, $\alpha = 0.05$). Duration is 18 months active plus 6-month follow-up; per the 3-Year Rule, full claims impact requires longer observation. The control group receives identical disease management without behavioural enhancement.

The recommended primary design is a randomised encouragement trial with intention-to-treat (ITT) analysis: all eligible members are randomised to receive or not receive an invitation to participate, and outcomes are analysed by randomisation assignment regardless of actual uptake. This yields unbiased estimates of the average effect of being offered the programme (the policy-relevant parameter for pricing). As robustness checks, actuaries should apply difference-in-differences (DiD) with propensity-score-matched controls using pre-intervention claims trajectories, and instrumental variable (IV) estimation using randomised encouragement as the instrument to estimate the local average treatment effect among compliers.

The UK health system has already delivered large-scale, multi-site behavioural trials closely analogous to the pilots proposed here: CPIT III randomised over 1,000 pregnant smokers across multiple NHS trusts; London's cervical screening programme sent GP-endorsed SMS to hundreds of thousands of women (NHS England, 2023b); Vitality's partnership with RAND enrolled over 400,000 members across three countries. These precedents confirm the logistical requirements are challenging but within demonstrated capability.

Key operational considerations. Effective BAF calibration requires linking behavioural data with claims outcomes—requiring insurer-integrated platforms or data sharing agreements compliant with UK GDPR. Deposit-based programmes require members to commit personal funds, which may face regulatory and reputational challenges; reward-based and VBID programmes face lower consent barriers. mHealth interventions require digital platforms capable of real-time data capture and integration with insurer administrative systems. If BAF-adjusted pricing leads to differential premiums based on programme participation, insurers should ensure compliance with equalities legislation and early regulatory engagement is recommended.

Phase	Cost	Duration	Go/No-Go
1: Feasibility	£40–60k	6 months	≥1 insurer committed
2: Single-insurer pilot	£120–180k	24 months	≥10pp PDC improvement

3: Multi-insurer consortium	£350–500k	36 months	Effect replication across ≥2 insurers
4: Long-term follow-up	£60–100k	+24 months	Significant claims reduction
Total	**£570–840k**	**5+ years**	

Table 6. Phased research programme with success criteria. Each phase delivers specific calibration outputs: Phase 2 calibrates R and E; Phase 3 provides sample sizes for C; Phase 4 generates longitudinal evidence for D. The primary objective of the £570k–£840k investment is to replace international proxies with UK-calibrated α parameters—the principal gap identified in Section 9.

8. Implications for Actuarial Practice

Adoption of the BAF would shift actuarial practice in several respects. Pricing actuaries would move from static trend assumptions to intervention-conditional projections. Reserving actuaries would model behavioural decay in IBNR, with the 29% admission reduction affecting frequency development. Programme evaluators would assess ROI at condition level, with the disease/lifestyle distinction becoming standard. Product designers would embed nudge architecture into benefit design, with VBID structures and incentive timing becoming actuarial parameters. Stop-loss actuaries would recognise that the BAF affects large claim probability disproportionately, since the interventions target precisely the high-cost chronic conditions that drive excess loss development. Equally important, the BAF provides a rigorous basis for declining vendor proposals: a programme producing a BAF indistinguishable from 1.00 can be rejected on transparent, documented grounds rather than subjective scepticism.

Tiered ROI expectations should guide investment: disease management at 3.5–4.0:1 (high confidence), targeted behavioural incentives at 5.0–7.5:1 (requiring proper design), and general lifestyle wellness at 0.5–1.5:1 (should not be assumed positive).

Where data permit, actuaries should estimate BAF parameters stratified by age band, comorbidity count, and socioeconomic segment rather than applying a single portfolio-wide adjustment. As BAF-adjusted projections accumulate experience, insurers should implement ex post validation by comparing predicted BAF-adjusted claims to realised claims on a quarterly or annual basis, with persistent deviations triggering parameter re-estimation through Bayesian updating of the Monte Carlo priors.

The BAF supports dynamic reserving under Solvency II by incorporating behavioural uncertainty into capital requirements. For UK PMI, where NHS interactions influence private claims (e.g., delayed NHS care driving PMI utilisation), the BAF can model nudge effects on waiting list avoidance, reducing claims volatility. UK calibration is primarily constrained by claims translation rather than behavioural response—reinforcing the case for insurer-led pilots rather than undermining the framework itself.

9. Limitations and Future Research

The BAF is presented as a structuring tool for actuarial judgement, not a fully calibrated model. Efficacy calibration risk: RCT parameters reflect controlled conditions; real-world

deployment may yield lower effect sizes due to implementation variability and differences in participant motivation. As a practical rule, actuaries should haircut RCT effect sizes by 20–40% when using them as pricing priors—20% for automated mHealth platforms with high fidelity, 40% for complex multi-component programmes. Vendor ROI overstatement: the wellness industry has a documented tendency to overstate ROI through historical controls, regression to the mean, and conflation of correlation with causation (Mattke et al., 2013). The BAF counters this by requiring RCT-grade evidence and component-level decomposition. International transferability: the majority of the evidence originates from US employer-sponsored insurance, where benefit design and cost structures differ materially from UK PMI. The framework is structurally valid but parametrically provisional until UK data are incorporated. RCT representativeness: participants in behavioural RCTs are typically younger, more educated, and more health-literate than the general insured population; subgroup analyses reinforce the need for stratified BAF application. Tail-risk and model risk: the multiplicative structure may understate tail risk where multiple components deteriorate simultaneously; if the interaction between efficacy and durability is sub-multiplicative, the framework would overstate claims-reduction potential. This model risk should be disclosed under APS X2 requirements, with explicit statement of the conditions under which the multiplicative assumption is expected to hold.

Table 7 summarises the US-to-UK parameter adjustment direction.

Component	US Estimate	UK Direction	UK Plausible Range
R (Reach)	0.85–0.95	Similar or lower	0.70–0.90
E (Efficacy)	0.30–0.50	Lower (10–30%)	0.20–0.45
C (α)	0.40–0.60	Uncertain	0.25–0.60 (principal gap)
D (h, λ)	h=0.15–0.50	Similar	h=0.15–0.50; λ=0.03–0.08

Table 7. US-to-UK parameter adjustment (indicative — to be replaced by UK pilot estimates).

Research priorities include: UK claims-translation calibration through insurer-led pilots (the principal gap is the α parameter mapping behavioural change to PMI claims); multi-insurer consortia modelled on the US Health Care Cost Institute or ABI data-sharing frameworks for adequately powered claims-level studies; BAF extensions for mental health, musculoskeletal conditions, and high-cost specialty therapies (GLP-1 receptor agonists, PCSK9 inhibitors); integration with value-based care contracts as UK PMI moves toward outcomes-linked reimbursement; and development of IFoA professional guidance establishing documentation standards for behavioural adjustment factors in pricing and reserving.

10. Conclusions

This paper has introduced the Behavioural Adjustment Factor—to the best of current knowledge, the first actuarial framework to decompose behavioural intervention impact into evidence-grounded, Solvency II-compatible components. The BAF is structurally sound and ready for piloting with conservative UK priors; it should be understood as a structuring device for disciplined evidence generation and actuarial judgement rather than a fully calibrated pricing tool.

The evidence is compelling but nuanced. Disease management generates robust returns (ROI 3.78:1, 29% hospital admission reduction). Lifestyle management does not reliably produce savings—and the BAF's value lies equally in providing a disciplined mechanism for declining investment in programmes that do not meet the evidentiary threshold. The reach–efficacy frontier creates strategic choices. The 3-Year Rule warns against premature evaluation. The worked example demonstrates that a large behavioural effect (37pp adherence gain) produces a modest claims adjustment (BAF 0.82–0.94) once reach, clinical translation, and decay are applied.

The IFoA, insurers, and the actuarial profession are called upon to: adopt the BAF framework for evaluating behavioural programmes, replacing aggregate wellness ROI claims with condition-specific, evidence-based adjustments; fund the milestone-based pilot beginning with a £40–60k feasibility study, with durability parameters as the highest-priority calibration target; and develop professional guidance on integrating behavioural adjustment factors into pricing and reserving standards.

These conclusions carry important caveats. The evidence base is predominantly US-sourced, the durability parameters carry material uncertainty, and selection effects in voluntary programmes remain difficult to isolate from causal impact. The framework should be treated as structurally sound but parametrically provisional until UK insurer-led pilots provide calibration data.

The broader vision is of actuaries evolving from passive pricers of risk to active shapers of health outcomes—designing interventions, quantifying impact, and embedding behavioural architecture into products. The potential—for insurers, for members, and for population health—is substantial. The question is whether the profession will invest in generating that evidence and embrace the opportunity to shape risk, not merely measure it.

Acknowledgements

The author is grateful to the Institute and Faculty of Actuaries for the opportunity to develop this work through The Big Pitch programme. Any errors remain the author's own.

Appendix A: Statistical Power Calculations

Endpoint	Detectable Effect	n per Arm	Feasibility
Adherence	Δ = 5pp (0.50→0.55)	1,565	Feasible
Adherence	Δ = 10pp (0.50→0.60)	388	Feasible
Claims (PMPM)	5% (≈£20)	2,510	Marginal
Claims (PMPM)	7% (≈£28)	1,280	Feasible

Table A1. Minimum sample size per arm (α=0.05 two-sided, power=0.80). Claims detection at Δ<7% requires multi-insurer consortia.

Appendix B: Evidence Strength Ratings

Parameter	Range	Key Sources	Strength
R: Reward-based	0.85–0.95	Halpern et al. (2015); Hafner et al. (2018)	★★★
R: Deposit-based	0.10–0.20	Halpern et al. (2015); Boonmanunt et al. (2023)	★★★
E: Smoking cessation	14.7–52.3%	Volpp et al. (2009); Tappin et al. (2022)	★★★
E: Medication adherence	7–37pp	Chernew et al. (2008); Dodson et al. (2025)	★★
C (α): CV events	0.25–0.60	Chen et al. (2022); Lewington et al. (2002)	★★★
D (h): Habit formation	0.15–0.50	Halpern et al. (2015); Lally et al. (2010)	★
D (λ): Decay rate	0.03–0.08	Hafner et al. (2018)	★

Table B1. Evidence strength: ★★★ = Strong (multiple RCTs, consistent); ★★ = Moderate; ★ = Limited. Insurers should prioritise generating data for ★-rated parameters.

Appendix C: UK Evidence at a Glance

Domain	UK Evidence	Key Numbers	BAF Component(s)
Smoking cessation incentives (pregnancy)	CPIT III RCT (n≈1,000)	Quit 26.8% vs 12.3%; aOR 2.78; IMD-stratified	E, D, C, Equity

Domain	UK Evidence	Key Numbers	BAF Component(s)
Screening reminders	NHS Health Checks cluster RCT; London cervical screening SMS	Health Checks: +11.8pp; Cervical: +4.8–5.9pp	R
Physical activity incentives	RAND/Vitality Apple Watch study (n=400,000+); Vitality Habit Index with LSE (2024)	34% sustained activity increase over 24 months; 38–58% mortality risk reduction	R, E, D, C
Insurer claims impact	Vitality UK Health Claims Insights 2024	Platinum members ≈28% lower claims; up to 5 extra years life expectancy	C (directional)
Digital adherence	Meta-analyses of SMS, apps, electronic monitoring	+10–23pp adherence vs control	E
Behaviour-change guidance	NICE PH49	Recommends goal-setting, feedback, monitoring; ≥1-year evaluation	All (design, monitoring)

Table C1. UK evidence at a glance.

References

Baicker, K., Cutler, D. and Song, Z. (2010) 'Workplace wellness programs can generate savings', Health Affairs, 29(2), pp. 304–311.

Boonmanunt, S., et al. (2023) 'Evaluation of the effectiveness of behavioral economic incentive programs for goal achievement', Annals of Behavioral Medicine, 57(4), pp. 277–287.

Caloyeras, J.P., et al. (2014) 'Managing manifest diseases, but not health risks, saved PepsiCo money over seven years', Health Affairs, 33(1), pp. 124–131.

Chen, C., et al. (2022) 'Adherence with cardiovascular medications and the outcomes in patients with coronary arterial disease', Clinical Cardiology, 45(12), pp. 1220–1228.

Chernew, M.E., et al. (2008) 'Impact of decreasing copayments on medication adherence within a disease management environment', Health Affairs, 27(1), pp. 103–112.

Dodson, J.A., et al. (2025) 'Effect of mHealth incentive lottery on blood pressure medication adherence: The BETTER-BP randomized clinical trial', Journal of the American College of Cardiology.

Financial Conduct Authority (2022) PS22/9: A new Consumer Duty. London: FCA.

Goldberg, S.B., et al. (2022) 'Mobile phone-based interventions for mental health', PLoS Digital Health, 1(1), e0000002.

Hafner, M., Pollard, J. and Van Stolk, C. (2018) Incentives and physical activity. RAND Corporation, RR-2870.

Halpern, S.D., et al. (2015) 'Randomized trial of four financial-incentive programs for smoking cessation', New England Journal of Medicine, 372(22), pp. 2108–2117.

Institute and Faculty of Actuaries (2016) Actuarial Professional Standard X2: Review of Actuarial Work. London: IFoA.

Kahneman, D. and Tversky, A. (1979) 'Prospect theory', Econometrica, 47(2), pp. 263–291.

Lally, P., et al. (2010) 'How are habits formed', European Journal of Social Psychology, 40(6), pp. 998–1009.

Lemstra, M., et al. (2012) 'Proportion and risk indicators of nonadherence to statin therapy', Canadian Journal of Cardiology, 28(5), pp. 574–580.

Lewington, S., et al. (2002) 'Age-specific relevance of usual blood pressure to vascular mortality', The Lancet, 360(9349), pp. 1903–1913.

Madrian, B.C. and Shea, D.F. (2001) 'The power of suggestion', Quarterly Journal of Economics, 116(4), pp. 1149–1187.

Mattke, S., et al. (2013) Workplace wellness programs study: Final report. RAND Corporation, RR-254-DOL.

Michie, S., et al. (2009) 'Low-income groups and behaviour change interventions', Journal of Epidemiology and Community Health, 63(8), pp. 610–622.

Mitchell, M.S., et al. (2023) 'Financial incentives for physical activity and weight loss', British Journal of Sports Medicine, 57(12), pp. 754–761.

NHS Diabetic Eye Screening Programme (2023) Annual Report 2022–23. London: NHS England.

NHS Digital (2023) 'NHS App pilot evaluations'. Available at: https://digital.nhs.uk/services/nhs-app.

NHS England (2023) National Schedule of NHS Costs 2022–23. London: NHS England.

NHS England (2023b) ‘Screening text reminder programme’. Available at: https://transform.england.nhs.uk/.

NICE (2013) Guide to the Methods of Technology Appraisal 2013 (PMG9). National Institute for Health and Care Excellence.

NICE (2014) Behaviour change: Individual approaches (PH49).

NICE (2020) Colorectal cancer (NG151).

Osterberg, L. and Blaschke, T. (2005) ‘Adherence to medication’, New England Journal of Medicine, 353(5), pp. 487–497.

Royal Swedish Academy of Sciences (2017) The Prize in Economic Sciences 2017: Richard H. Thaler. Press release, 9 October 2017. Stockholm.

Sallis, A., et al. (2016) ‘Pre-notification and reminder SMS text messages with behaviourally informed invitation letters’, BMC Public Health, 16, p. 1168.

Tappin, D.M., et al. (2022) ‘Financial incentives for smoking cessation in pregnancy: randomised controlled trial’, BMJ, 379, e071522.

Thakkar, J., et al. (2016) ‘Mobile telephone text messaging for medication adherence in chronic disease’, JAMA Internal Medicine, 176(3), pp. 340–349.

Thaler, R.H. and Sunstein, C.R. (2008) Nudge. New Haven: Yale University Press.

Tudor Hart, J. (1971) ‘The inverse care law’, The Lancet, 297(7696), pp. 405–412.

UK Prospective Diabetes Study Group (1998) ‘Intensive blood-glucose control with sulphonylureas or insulin’, The Lancet, 352(9131), pp. 837–853.

Vitality (2024) Health Claims and Insights Report 2024. London: VitalityHealth.

Vitality and London School of Economics (2024) The Vitality Habit Index. London: LSE Consulting.

Volpp, K.G., et al. (2009) ‘A randomized, controlled trial of financial incentives for smoking cessation’, New England Journal of Medicine, 360(7), pp. 699–709.