

WORKING PAPER · SEPTEMBER 2026

Financial Planning for the Greater Good

An epistemic approach to South Africa's R50.4 trillion protection gap, and the advisory capacity required to reduce it

ABSTRACT

South Africa's R50.4 trillion protection gap—the actuarial shortfall between a household's assessed financial need following the death or disability of an earner and their actual existing resources—compounds at 12.5 per cent annually. Traditional financial education fundamentally misdiagnoses this crisis because consumers do not lack general financial theory; they lack legible, local knowledge of their personal risk and a tractable method to evaluate advisory expertise.

To correct this, this paper introduces SikaSync, a free, non-advisory population screening instrument that calculates an earner's exact rand shortfall. Functioning as an epistemic measurement tool rather than a recommendation engine, SikaSync engineers "salience without shame" by converting an invisible risk into an actionable, portable artefact.

This equips the consumer to hold an advisory relationship to account and routes that demand to a commercially neutral network of vetted practices. Modelled over five years, the architecture would place R0.98 trillion of new cover in force and generate R12.85 million of gross revenue per practice, which is the order of magnitude at which the large-scale, black-owned advisory firms currently missing from the market could be capitalised without fiscal expenditure. The paper offers this as one workable route rather than the only one; what it argues for without qualification is the diagnosis beneath it.

Prof Prince Sarpong, PhD

University of the Free State | Centre for Financial Psychology and Systems | Epistemica Institut

A stakeholder working paper · September 2026

Contents

Executive summary.....	3
PART I — THE PROBLEM IS EPISTEMIC	4
1. From two chapters to one model.....	4
2. The gap.....	4
3. Why the standard remedies fail.....	5
3.1 The empirical case against education.....	5
3.2 The epistemic case: general knowledge is the wrong kind.....	6
3.3 The cognitive island: why “see an adviser” is not an answer.....	6
3.4 The label is itself the injury.....	6
3.5 The newest remedy, and why it has not closed the gap either.....	7
PART II — THE INSTRUMENT	8
4. Screening, not schooling.....	8
5. What the diagnostic is, epistemically.....	8
6. The design specification: salience without shame.....	9
PART III — THE DISTRIBUTION	10
7. The missing firm.....	10
8. Matching, and why it needs no rule.....	11
9. Precedent: distribution as capital formation.....	12
PART IV — THE ECONOMICS	13
10. The economic model.....	13
11. What each stakeholder gets.....	21
PART V — WHAT WOULD MAKE THIS FAIL	21
12. Risks and limitations.....	21
13. Conclusion.....	22
References.....	23
Appendix A Model parameters and formulas.....	24

Executive summary

South Africa's life and disability protection gap stood at R50.4 trillion at 31 December 2024, having widened by R15 trillion — 12.5 per cent a year — over the preceding three years. While ASISA attributes this widening gap to a confluence of economic pressure, behavioural inertia, and financial illiteracy, the sector has primarily funded and applied a single behavioural remedy: financial literacy education. A hypothesis funded while its target variable deteriorates by trillions is not awaiting more evidence. It has been tested.

This paper argues that the gap persists because the sector has misdiagnosed it. The problem is not that South Africans know too little about finance in general. It is that they have no legible knowledge of their own position in particular, and no tractable way to find a practitioner who will supply it. Those are epistemic failures, and they have specific, well-described causes:

- **Education supplies the wrong kind of knowledge.** Financial education transmits general knowledge; the insurance decision turns on local knowledge — the household's dependants, debts, and shortfall (Hayek, 1945).
- **Choosing an adviser is a cognitive island.** Evaluating financial expertise requires financial expertise. Telling an uninsured earner to “go see an adviser” asks them to solve the problem that made them uninsured (Nguyen, 2020).
- **The literacy label is itself an injury.** Calling consumers financially illiterate is a textbook case of testimonial injustice: it discounts their credibility as knowers, depresses epistemic self-confidence, and drives them toward low-diversity information sources like family and friends rather than practitioners (de Bruin, 2021).

The instrument this paper proposes follows directly. A non-advisory diagnostic, free at the point of use, returns an earner's own shortfall. It supplies local knowledge rather than general knowledge; it asks to be read rather than believed, which keeps it outside the advice perimeter and outside the trust problem that burdens robo-advice; and it produces a portable artefact the client carries into an advisory relationship, converting an unsolvable expert-selection problem into a tractable verification problem.

The behavioural specification comes from the profession's own literature. Clients who engage advisers exhibit moderate-to-high physiological arousal together with low financial anxiety; those who do not engage exhibit low arousal and high anxiety, a state described as self-imposed helplessness (Grable, Heo and Rabbani, 2015). The design target is therefore precise and falsifiable: raise salience while suppressing shame. The literacy framing does the exact opposite, which is why it produces disengagement in the segment most able to afford cover.

On distribution, the structural fact is not contested by critics of the industry but conceded by its incumbents. Speaking at the PSG Financial Services conference on 9 May 2024, Momentum Metropolitan Holdings Group CEO Jeanette Marais observed that while 85 per cent of South African personal income taxpayers are Black, Indian or Coloured, “there is not one large black financial advisor firm or advisor group in SA”. She called it an opportunity. This paper argues it is a structural barrier, and that the reason it has survived seven years of enterprise development programmes is that those programmes supply capital, skills and compliance, never clients.

Modelled conservatively on cover in force, net of assumed lapses, a staged national deployment over five years:

- **R0.98 trillion** of cover placed in force for 315,000 households, lifting the national coverage ratio from 38.8 to 39.9 per cent, and to 52.3 per cent by year thirty.
- **R6.43 billion** of regulated intermediary remuneration into a 500-practice network — R12.9 million per practice, or R2.6 million a year.
- **R1.89 billion** of value of new business to insurers, on R6.3 billion of newly written annualised premium.
- **Roughly 1,250** direct professional roles, in the G20 economy with the strongest employment response to growth.

The model requires no fiscal expenditure. It operates on a multi-sided token system where the diagnostic scan is funded by the party capturing the downstream value. Advisers can pre-purchase tokens to distribute to prospects; institutional providers can purchase corporate tokens that exclusively route the resulting five adviser matches to their own tied networks; and development financiers like the National Empowerment Fund (NEF) can fund tokens that route matches strictly to qualifying black-owned practices, directly capitalising the transformation outcomes Marais identifies. Alternatively, un-sponsored clients can pay a direct retail fee for an unconstrained match of five independent specialists.

What this paper does not claim. It does not claim to close the gap: a shortfall compounding at 12.5 per cent a year outruns any single distribution intervention. And it does not claim that financial education is worthless, only that it has been asked to carry a load it is unable to bear.

1. From two chapters to one model

Portfolio Management for Financial Advisors, Volume II makes two arguments that have not previously been put together. Chapter 2, titled Money Doctors, builds the case that financial advisers are healers. It draws on Gennaioli, Shleifer and Vishny's (2015) account of advisers as money doctors whose value rests on trust rather than superior returns, and on Klontz, Van Zutphen and Fries (2016), who show that financial planning shares the common healing factors that Jerome Frank (1974) identified across all schools of psychotherapy: a helping relationship, a setting, an agreed conceptual scheme, and an agreed ritual. The chapter documents the physiological reality of financial stress, and observes — quoting Cummings (2017) — that because far more people face devastating financial situations in their lifetimes than face devastating illness or legal crises, financial planning holds more promise to deliver personal and social benefit than either medicine or law.

Chapter 9, Financial Planning for the Greater Good, argues that the macroeconomic impact of financial planning is barely discussed in financial planning curricula, and that the profession consequently "fails the broader society by failing to leverage its position to push for positive change and the greater good".

Put the two together and a conclusion follows that neither chapter states. If financial advisers are doctors, then the profession should ask what medicine actually does when it wants to change outcomes across a whole population. It does not run a public lecture series on cardiology. It screens. It takes an invisible, asymptomatic, high-consequence risk and makes it visible in the individual patient, then refers that patient to a practitioner. Screening is the highest-leverage population instrument medicine has, and financial planning has no equivalent.

This paper proposes that instrument, sets out its epistemic and behavioural design, describes the distribution architecture required to convert its output into cover, and models the economics of both. It is offered as the operational answer to the challenge Chapter 9 leaves open.

2. The gap

The Association for Savings and Investment South Africa commissions a life and disability insurance gap study every three years, conducted by True South Advisory. The 2025 study, published on 28 October 2025 and measured as at 31 December 2024, found:

Measure	Value
Total insurance need	R82.2 trillion
Cover actually in force	R31.9 trillion
Protection gap	R50.4 trillion
of which death cover	R21.1 trillion
of which disability cover	R29.2 trillion
Coverage ratio (cover ÷ need)	38.8%
Income earners assessed	16.1 million
Average shortfall per earner	R3.1 million
life cover shortfall (need R2.1m, cover R0.8m)	R1.3 million
disability shortfall (need R3.0m, cover R1.2m)	R1.8 million
Additional premium required to close the gap	8.2% of earnings

Table 1 South Africa's protection gap at 31 December 2024. Source: ASISA / True South Advisory, *The South African insurance gap (2025)*. The 8.2 per cent comprises 5.2 per cent for death cover and 3.0 per cent for disability cover. Exact study figures: need R82,225bn, cover in force R31,872bn (of which R5.2tn is disability-grant cover), gap R50,353bn, death gap R21,138bn; the model's appendix anchors use the rounded headline values.

The average shortfall of R3.1 million is roughly twelve times the average earner's gross annual income. Disability accounts for 58 per cent of the total — the risk consumers most reliably underestimate. And the gap is approximately 6.9 times South Africa's 2024 nominal GDP of R7.3 trillion.

The trajectory matters more than the level. Cover in force grew 5 per cent a year over the three years to December 2024. Assessed need grew 9.3 per cent. The gap therefore grew 12.5 per cent a year, from a restated R35.4 trillion at December 2021 to R50.4 trillion. Cover is not falling, it is being outrun.




Figure 1 The gap widened by R15 trillion in three years. Source: ASISA / True South Advisory. The December 2021 figure is as restated in the 2025 study; the 2022 study originally published R34.3 trillion for that date.

ASISA publishes this study with care and candour every three years. The sector reads it, quotes it, and, on the evidence of the numbers themselves, does not act on it at a scale that changes the trend. The purpose of this paper is to move from restating the problem to proposing a mechanism that can be tested against it.

3. Why the standard remedies fail

While the sector concedes that affordability forms a structural baseline to the gap, its response to the variables it can control is dominated by two remedies: educate the consumer, or tell the consumer to consult an adviser. Both fail, and they fail for reasons that are principled rather than incidental. A third factor, which is the language in

which the problem is described, actively makes things worse.




Figure 2 The two standard remedies, why each fails, and what a diagnostic does instead.

3.1 The empirical case against education

The most-cited meta-analysis in this area, covering 201 prior studies, found that financial education interventions explained roughly 0.1 per cent of the variance in downstream financial behaviour, with effects decaying substantially within about twenty months (Fernandes, Lynch and Netemeyer, 2014). Lauren Willis had already named the underlying problem in the American Economic Review three years earlier, in a paper titled simply: *The Financial Education Fallacy* (Willis, 2011).

Intellectual honesty requires the counterpoint. A later meta-analysis using newer randomised trials found considerably larger and more durable effects, and concluded that financial education does affect both knowledge and behaviour (Kaiser, Lusardi, Menkhoff and Urban, 2022). The literature is genuinely contested, and nothing here requires the strong claim that education is useless.

The claim is narrower and harder to escape. Whatever education achieves at the margin, it has not been sufficient here: South Africa funded literacy programmes and adviser credentialing throughout the period in which the gap widened by R15 trillion. On the sector's own measurement, the remedy is not keeping pace with the problem by an order of magnitude. That is a reason to add a different mechanism, not to fund the same one harder.

3.2 The epistemic case: general knowledge is the wrong kind

The deeper problem is that education and insurance decisions traffic in different kinds of knowledge.

Hayek's 1945 distinction is exact. There is scientific or general knowledge — the sort that can be taught in a classroom and written in a textbook — and there is "the knowledge of the particular circumstances of time and place", which never exists in concentrated form but only as dispersed particulars held by individuals. Financial education supplies the first. The decision to buy R3.1 million of cover turns entirely on the second: this earner's income, dependants, debts, existing policies, and the specific risk and consequence for this household if this person dies on Tuesday.

A consumer can complete a financial literacy course, correctly answer every question about how term assurance works, and still have no idea that their family faces a R3.1 million shortfall. While the general knowledge is necessary, it is not sufficient. It is the integration of their own dispersed particulars into a single legible number that changes the decision, and that is the operation education cannot perform, because the educator does not have the particulars.

3.3 The cognitive island: why "see an adviser" is not an answer

The sector's second reflex is to route the consumer to a professional. This is better advice, but it conceals a genuine epistemological obstacle.

Nguyen (2020) describes a cognitive island as a domain in which expertise is required in order to evaluate expertise, where there is no test of competence available to the inexpert, and no adjacent domain that supplies one. Financial advice is close to a paradigm case. To tell a good adviser from a bad one, a consumer must already be able to assess the quality of financial reasoning, which is the capability they lack and the reason they need an adviser. Nguyen's further point is that the resulting error is self-sealing rather than self-correcting: an inexpert person's criteria for choosing an expert are themselves shaped by the expert they chose.

This explains an otherwise puzzling empirical regularity. De Bruin (2021) documents that consumers with low confidence in their financial judgement do not search harder among professionals; they narrow their search to family and friends, and sources whose general reliability they believe they can assess, even where those sources have no relevant competence at all. That is not irrationality. Given a cognitive island, it is a sensible epistemic strategy. It is also how a person ends up uninsured on the advice of a cousin.

Any proposal that begins "the consumer should consult a qualified adviser" therefore assumes away the binding constraint. The constraint is not willingness, it is that the consumer has no way to tell whether the person in front of them is worth consulting.

3.4 The label is itself the injury

The third factor is the language. The sector's habitual description of the problem — financial illiteracy — is not a neutral report of a deficit. It is, in the technical sense, an epistemic injustice.

Fricker (2007) defines *testimonial injustice* as a wrong done to someone specifically in their capacity as a knower, when prejudice causes a hearer to give deflated credibility to their word. De Bruin (2021) applies the framework to financial services directly, and the mechanism he describes operates in both directions. On the supply side, practitioners discount the credibility and competence of clients from disadvantaged groups; his worked example, drawn from Gray (2012), is that lacking a college degree reduces a white applicant's small-business loan approval probability by 0.4 per cent (effectively nil) while reducing a black applicant's by 81 per cent. On the demand side, repeated exposure to the prejudice that one lacks the capacity to understand one's own finances lowers epistemic self-confidence, which in turn suppresses both the intensity and the diversity of information search. De Bruin himself sets the phrase "financial literacy" in quotation marks.

The practical consequence is a closed loop. The sector labels the consumer illiterate; the label depresses the consumer's confidence in their own financial judgement; the depressed confidence narrows their search away from practitioners and toward family; the narrowed search leaves them uninsured; and the sector cites their uninsured state as evidence of the illiteracy it diagnosed. De Bruin, following Swire (1995), calls this self-fulfilling testimonial injustice. It is a mechanism that generates the R50.4 trillion outcome we observe, and no amount of additional education can interrupt it, because education is delivered in the very register that causes the harm.

There is also a plain professional argument. An educated professional executes contracts without a law degree and buys medication without medical training. In those domains, relying on a specialist is understood as an efficient division of labour, not diagnosed as illiteracy. Financial services is unusual in expecting consumers to acquire the practitioner's discipline in

addition to their own, and unique in describing their reliance on professionals as a deficit. If every profession ran literacy programmes on that logic, nobody would have time left for their own work.

Even in law courts, where civil liberty and property rights hang in the balance, individuals retain the right to *pro se* representation without the state diagnosing their literacy and so the state takes no official stance on whether a citizen possesses 'legal literacy'.

3.5 The newest remedy, and why it has not closed the gap either

Since 2023 a fifth remedy has joined the four above: artificial intelligence will close the gap. The claim has not been tested at population scale until now. The Financial Planning Standards Board surveyed 6,206 financial planners across 24 territories, South Africa among them, between November and December 2024 (Financial Planning Standards Board, 2025). Its findings do not support the claim, and the shape of the shortfall is more instructive than the shortfall itself.

Adoption is not the constraint. Nearly two in three planners report that their firm is using AI or will be within twelve months (64 per cent), half hold a positive view of its impact (50 per cent) against 8 per cent negative, and 78 per cent expect it to help them serve clients better. The technology has arrived and the profession is willing. Whatever is failing here, it is not appetite.

The failure is one of placement. Asked which stage of the financial planning process stands to benefit most, planners nominated the analytical core: collecting the client's information, analysing their position, and developing recommendations. Asked where they actually use AI, they answered client communication and engagement (41 per cent), collection of client information (33 per cent) and client risk profiling (30 per cent); on the business-operations side, marketing and promotions (35 per cent), client onboarding (34 per cent) and internal workflow optimisation (33 per cent). The analytical work the profession itself identified as the highest-value target sits far down its own list of deployments: financial plan creation and analysis at 16 per cent, insurance needs analysis at 14 per cent, review of the client's position against the plan at 10 per cent, and monitoring the client's position between reviews at 6 per cent.

The explanation is a division of expertise, not a limitation of the technology. Forty-five per cent of firms deploy third-party vendor solutions against 34 per cent building in-house, and the models most used are general-purpose: large language models at 43 per cent, practice-management suites at 37 per cent. A general-purpose tool can be pointed at a general-purpose problem — drafting a client email, populating an onboarding form, running a campaign — by a builder who understands software and need not understand financial planning. It cannot be pointed at an insurance needs analysis by someone who does not know what one is: which liabilities survive the earner, what an employer benefit already covers, what the statute deems property of the estate, what a coverage ratio is measured against. The deployment map is therefore not a statement about the capability of the technology. It is a map of where domain knowledge was not required in order to build.

The planners corroborate the reading in their own answers. Their most-requested professional development is data analysis and interpretation, at 49 per cent, well ahead of understanding AI fundamentals at 34 per cent; they do not report a technology deficit but an analytical one. Their second-largest concern is the accuracy and reliability of AI outputs, at 42 per cent, which is the predictable behaviour of a general-purpose generative model asked to perform arithmetic nobody specified, and precisely the property a deterministic instrument does not have. Their fourth is the loss of human touch in client relationships, at 24 per cent — a concern about substitution, which is an argument for an instrument that routes to a practitioner rather than one that replaces her.

One finding deserves care, because it is easy to overstate. Asked why the firm adopted AI, respondents cited internal business need most often (61 per cent), ahead of industry demand (52 per cent) and client demand (47 per cent). The question permitted multiple selections and the responses sum to 160 per cent, so this is not evidence that clients are silent: nearly half of adopting firms did report client demand. What it establishes is a ranking rather than an absence. The strongest pull is the firm's own operating requirement, which is consistent with everything above. Work is being automated where the firm feels its own cost, not where the household feels its own risk.

Two limits on what this survey can be asked to prove. It measures adoption by financial planning firms and says nothing directly about consumer-facing tools, so it cannot establish that no instrument of the kind proposed here exists. And it reports a moment — late 2024 — in a fast-moving field. Its value to this argument is narrower and sufficient: the profession has itself identified the analytical core as where AI would matter most, and it has not built there. The gap this paper addresses is not one the market has failed to notice. It has been correctly identified by the practitioners closest to it and left unbuilt, because building it requires domain knowledge that the firms supplying the tools do not have, and the practitioners who hold that knowledge are not resourced to write software.

This is the space the instrument described in Section 5 occupies, and the coincidence is exact rather than rhetorical: insurance needs analysis, the single application this paper concerns itself with, is the fourth-least deployed use of AI on the profession's own list.

PART II — THE INSTRUMENT

4. Screening, not schooling

Medicine faced an analogous problem and solved it. Cervical cancer, hypertension, diabetes and HIV are all conditions that are asymptomatic in their early stages, catastrophic if untreated, and treatable if caught. In none of these cases did public health conclude that the answer was to teach the population enough pathology to self-diagnose. It built screening programmes: a cheap, standardised test that converts an invisible individual risk into a visible individual result, followed by referral into clinical care.

South Africa knows this model better than most countries. Although several educational endeavours were deployed, the national HIV response did not resort to immunology literacy. It asked them to know their status — and then built the referral pathway from a positive result into treatment. Testing is not education. It is the production of one specific, personal, actionable fact.

Underinsurance has the identical structure. It is asymptomatic since nothing about being underinsured feels like anything until the day it matters. It is catastrophic on realisation. It is entirely treatable in advance, at a cost the ASISA study puts at 8.2 per cent of earnings. And it is invisible to the person carrying it. Every property that makes screening the right instrument in medicine holds here.

Chapter 2 of *Portfolio Management for Financial Advisors* argues at length that the financial adviser occupies the position of a healer of financial stress. This paper simply takes that claim to its operational conclusion. A profession of money doctors without a screening programme is a profession that waits in its consulting rooms for the small minority who already suspect they are ill. In South Africa that minority is 9 per cent of households (Momentum and the Bureau of Market Research, 2025). Screening is how the other 91 per cent find out they should come in.

The analogy also imports medicine's cautions, and the paper adopts them explicitly. Screening programmes carry documented harms: overdiagnosis, false positives, and anxiety induced in people who would never have become symptomatic. A financial diagnostic can fail in analogous ways — it can overstate need, it can be tuned to manufacture urgency, and it can be captured by whoever pays for it. Section 12 treats these as live risks rather than rhetorical concessions.

5. What the diagnostic is, epistemically

Calling the instrument a diagnostic rather than a recommendation engine is not branding. It is a claim about what the user is being asked to do with the output, and it has three consequences that matter.

5.1 It asks to be read, not believed

Alvarado (2023) argues that AI systems are first and foremost epistemic technologies: they are deployed in contexts of inquiry, they operate on epistemic content, and they perform epistemic operations. The appropriate question about such a system is therefore not whether it is useful but whether it is trustworthy as a source of belief, and here the literature draws a distinction the financial sector has largely ignored.

A robo-adviser issues a recommendation. A recommendation is an assertion, and accepting it requires the user to extend testimonial trust to a system that cannot bear responsibility for what it asserts, which is why philosophers of testimony have struggled to classify machine outputs as testimony at all. Freiman (2024) proposes a third category, technology-based belief, distinct from both testimony-based belief with a human source and instrument-based belief of the sort produced by a thermometer or a blood test.

A diagnostic sits close to the instrument end of that spectrum. It does not tell the user what to do; it reports a measurement of their own position, computed from inputs they supplied, by a method that can be shown. The user is not asked to defer to it. They are asked to read it, the way one reads a blood pressure result. This is a genuine epistemological difference rather than a rhetorical dodge, and it maps cleanly onto the regulatory distinction between advice and factual information, which is what allows the instrument to be deployed by employers, funds and public bodies that could not distribute advice.

5.2 It preserves epistemic autonomy

Ahlstrom-Vij (2013) defines epistemic paternalism as interfering with an agent's inquiry, for that agent's epistemic good, without needing their consent. A recommendation engine is paternalistic in this sense: it substitutes its own conclusion for the user's deliberation. A diagnostic does not. It changes the user's evidential position and then hands the deliberation back.

This paper does not claim epistemic paternalism is always wrong; Ahlstrom-Vij's defence of it is serious. The claim is narrower: where a paternalistic design and an autonomy-preserving design can achieve the same behavioural outcome, the autonomy-preserving one is to be preferred — and in this case it is also the design that avoids algorithm aversion, stays outside the advice perimeter, and does not require the user to solve the trust problem before they can benefit. Three independent arguments converge on the same design choice.

5.3 It builds a bridge off the cognitive island

The most important property is the one that addresses Section 3.3. The diagnostic does not ask the consumer to evaluate an adviser. It gives them a portable, legible statement of their own position: a number, a shortfall, a specific claim about their household, which they can carry into any advisory conversation.

That transforms the problem. The consumer still cannot assess whether an adviser is broadly competent, which remains a cognitive island. But they can now check something much narrower and entirely tractable: does this adviser's recommendation address the shortfall the diagnostic identified, and if not, why not? An unanswerable question about general expertise becomes an answerable question about a specific gap. The client acquires the standing to hold the relationship to account, which is what epistemic autonomy actually consists of, on the field's own account: not self-reliance, but the capacity to govern one's dependences.

It also, incidentally, protects the adviser. A client who arrives with a documented shortfall and leaves with cover addressing it has a record of why the advice was given, which is a materially better compliance position than a reconstructed needs analysis.

6. The design specification: salience without shame

The preceding sections establish what the instrument must be. This section states what it must do, and the specification comes from the profession's own literature rather than from first principles.

Chapter 2 of *Portfolio Management for Financial Advisors* reports a finding from Grable, Heo and Rabbani (2015). In an exploratory clinical study measuring physiological arousal alongside self-reported financial anxiety, the people most likely to engage a financial adviser exhibited moderate to high physiological arousal combined with low financial anxiety. The people least likely to engage exhibited the opposite: low physiological arousal and high financial anxiety, a combination the literature associates with a kind of self-imposed helplessness. Related work finds that physiological stress is negatively associated with the desire to change (Britt, Lawson and Haselwood, 2016), and that financial shame functions as what Klontz and Klontz (2009) call an emotional glue trap, keeping people stuck in patterns they will not discuss and will not seek help for.

This yields a falsifiable engineering target. An effective financial diagnostic must raise arousal, the risk must become salient, personal and present, while simultaneously holding anxiety and shame down. Salience without shame. Raise both and the user disengages into helplessness; raise neither and nothing happens at all.

Read against this specification, the sector's dominant approach is close to a worst case. Literacy programmes raise anxiety, by telling the consumer the problem is their own deficiency, while raising salience hardly at all, because they never state the consumer's personal number. High anxiety, low arousal, which is the profile Grable and colleagues associate with non-engagement. The sector has been optimising, unintentionally, for the quadrant that does not act.

The design implications are concrete and testable:

- **State the number, not the norm.** "Your household faces a R2.1 million shortfall" raises salience. "South Africans are underinsured" does not.
- **Attribute the gap to circumstance, not character.** The output must read as a measurement, not a verdict. Nothing in it should imply the user was ignorant or irresponsible.
- **Never use the vocabulary of literacy or education.** It is the shame trigger, and Section 3.4 explains why it is also unjust.
- **Pair every finding with an available action.** Arousal without a route to resolution converts into anxiety. The referral is not a commercial afterthought; it is what keeps the intervention on the correct side of the specification.

- **Normalise.** Chapter 2 identifies normalising the client's experience as one of the ten ways advisers act as healers. A diagnostic can do this by placing the user's position in context without moralising about it.

This specification is the paper's most immediately testable contribution. It predicts that two diagnostics delivering identical information in different registers will produce materially different conversion rates, and that the difference will track measured shame rather than measured comprehension. That is a hypothesis a pilot can falsify.

The instrument specified here is not a proposal to build software. It exists. [SikaSync Financial MRI](#), an Epistemica instrument, is a working implementation of Parts II and III: a non-advisory diagnostic that returns a household's own position in rand, engineered against this section's specification, with the choice architecture of Section 8.1 and the client-elective matching layer of Section 8.3 already in production. Its non-advisory character is enforced in software rather than promised in policy: every generated passage is checked against a prescriptive-language doctrine before it reaches the user, so the instrument reports positions and cannot drift into recommendations. The doctrine implements the position developed in [Applied AI Epistemics: Rethinking Cognition, Behavior and Assessment in the Age of Generative Intelligence](#) (Sarpong, forthcoming 2027): the system is a mirror rather than an oracle; it measures, and the human deliberates. The platform is deployed in South Africa and four other regulatory jurisdictions, its calculation engine is versioned and held by a locked test suite, and the independent methodology review that Section 12 calls for is the governance step that scaled deployment would add. What remains unproven is the conversion parameter, and that is what the pilot is designed to measure.

PART III — THE DISTRIBUTION

7. The missing firm

A diagnostic that produces demand and has nowhere to send it is a public information campaign with extra steps. The instrument only closes gaps if there is a distribution layer capable of absorbing what it activates. In South Africa that layer has a specific and well-documented deficiency, and the most useful statement of it comes from inside the incumbent industry rather than from its critics.

Speaking on the second day of the PSG Financial Services annual conference on 9 May 2024, Jeanette Marais, the Group CEO of Momentum Metropolitan Holdings, one of South Africa's largest life insurers, set out the demographics of the addressable market and then named what is missing from the supply side.

"When analysing the consumer market from a demographic point of view, only 15% of personal income taxpayers are white. So at least 85% of official income earners are black, Indian and coloured, and they can afford some form of insurance, savings or investment products."

"Yet, while 85% of individuals who [could] potentially contribute to an annuity or pension fund are black, there is not one large black financial advisor firm or advisor group in SA — another opportunity for us as an industry."

Jeanette Marais, Group CEO, Momentum Metropolitan Holdings, PSG Financial Services conference, 9 May 2024, reported in [News24 Business](#), 10 May 2024. The bracketed insertion is the publication's.

In the same address she set out the scale of the advice shortage: approximately 3,000 practising CFP professionals and 21,000 licensed Category I advisers serve the country, against an estimated need for 4,000 more planners; roughly 750,000 households earning above R30,000 a month make no retirement provision; approximately one million households in that income band hold no life insurance; only about a quarter of the labour force provides for retirement at all; and some 90 per cent of her group's payouts to clients are facilitated by a financial adviser. Two of those figures have since moved: the Institute now counts 5,001 CFP professionals, and the regulator licenses providers as entities rather than advisers, recording 11,336 authorised Category I and IV providers (Financial Planning Institute, 2025; Financial Sector Conduct Authority, 2025).

That last figure requires some emphasis. Nine in ten benefits reach households through an adviser. Advice is not an optional enhancement to the protection system in South Africa; it is the delivery mechanism. Any strategy for closing the gap that does not expand advice capacity is, on the industry's own data, a strategy for not closing it.

7.1 An opportunity, or a structural barrier?

Marais characterised the absence of a large black-owned advisory firm as "another opportunity for us as an industry". Her framing is generous and, from where she sits, reasonable. This paper takes a different view, and the difference is the argument.

An opportunity is something the market has not yet got around to exploiting. A structural barrier is something the market's own mechanics prevent. The evidence points to the second, and the strongest evidence comes from the industry's own transformation programmes.

Enterprise development in advisory distribution	Result
Momentum Masibambane Broker Initiative, with the ASISA Enterprise Development Fund	8 black-owned brokerages graduated (2017)
Momentum intermediary coaching cohort, reported 2021	60 black intermediaries, 22 women
ASISA ESD Fund black-owned SMMEs supported, F2025	45
Santam SIGDP emerging black brokers	752, writing R1.4 billion in premium
ESD targets achieved, 2021 (B-BBEE Commission)	61%
Large black-owned advisory firms or groups in South Africa, 2024	Zero

Table 2 Seven years of enterprise development in advisory distribution, and the outcome its own beneficiary industry reports. Sources: MMI Holdings media release, 10 August 2017; FAnews, 2021; Momentum Group Sustainability Report 2025; Santam via the FIA, 15 May 2026; B-BBEE Commission; Marais, News24, 10 May 2024.

The Masibambane programme launched in 2017 and graduated eight brokerages. Seven years later the group's own chief executive reports that there is still not one large black-owned advisory firm in the country. That is not a programme that has yet to run long enough. It is a programme calibrated to produce many small practices, being asked to explain the absence of large ones.

The diagnosis this paper offers is specific. Enterprise development supplies capital, skills and compliance infrastructure. It does not supply clients. A practice does not fail because its principal lacks a qualification — the country has 5,000 certified CFP professionals. It fails because client acquisition in a market where only 9 per cent of households use an adviser is prohibitively expensive, and because the assets-under-management model that sustains established practices is structurally unavailable to advisers serving clients who have not yet accumulated assets.

Protection advice breaks that constraint, because it decouples practice revenue from client wealth and anchors it to client income instead. A client with no investable assets and a R25,000 monthly salary is not a viable assets-under-management prospect, but is a fully viable risk-advice client, and their household needs the cover considerably more than a wealthy client needs another portfolio review. Demand, not capital, is the missing input. That is what the diagnostic supplies, and it is the one thing seven years of enterprise development has never been able to manufacture.

8. Matching, and why it needs no rule

8.1 A commercially neutral front end

The matching layer is governed by how the diagnostic scan is funded, ensuring the entity carrying the cost of the instrument captures the corresponding commercial or developmental return. Following a scan, the tool suggests five vetted advisers who best match the client's profile. The pool from which these five are drawn is determined by the token used to initiate the scan:

Corporate Tokens: Institutions purchase scan tokens for public deployment; the five suggested matches are drawn exclusively from that provider's tied advisory network.

Development Tokens: Financiers fund tokens to engineer structural transformation; matches are routed strictly to qualifying black-owned practices to capitalise the capacity the industry acknowledges is missing.

Practice Tokens: Individual advisers acquire and distribute tokens; the resulting scan defaults the originating adviser into the match set.

Client-Funded: Users who pay directly for their own scan receive an unconstrained match of five independent specialists, based purely on service fit and elective demographic homophily.

This design removes the friction of a political or commercial choice at the point of conversion. The routing logic is embedded in the token distribution, preserving the diagnostic's epistemic integrity as an objective measurement while efficiently pricing its output as a qualified pipeline.

8.2 Homophily does the work

The distributional outcome does not require a routing rule, because of a regularity the author has already documented in Chapter 2 of *Portfolio Management for Financial Advisors*. Discussing Frank's third common healing factor — the agreed-upon conceptual scheme — that chapter cites Stolper and Walter's study *Birds of a Feather*, which finds that homophily, the affinity of individuals for others like them (McPherson, Smith-Lovin and Cook, 2001), has a powerful and significant influence on whose recommendations people follow, and that client-adviser matching increases the likelihood of following financial advice (Stolper and Walter, 2019).

The South African evidence is unusually direct. Sanlam's market research has found that 26 per cent of black respondents used a broker or adviser against 64 per cent of white respondents, and Sanlam has reported that close to 95 per cent of black financial advisers' clients are black African (Becker, 2023). The preference is already operating. It is operating inside a distribution layer too small to absorb the demand.

Combine that with Marais's demographic arithmetic (85 per cent of taxpayers, zero large firms) and the conclusion follows without anyone designing it. Open a marketplace of vetted advisers to a population that is overwhelmingly black, and demand flows along lines of existing preference into practices that are currently sub-scale. The transformation outcome is a consequence of arithmetic and documented behaviour, and so this requires no political defence and creates no consumer friction.

8.3 How the matching engine protects the client

The matching layer is client-elective and privacy-preserving by construction, and this should be understood as part of the compliance case rather than an implementation detail.

- **The client selects the dimensions.** Optional dimensions — language, faith, gender, age, marital status, parental status and race — are applied only where the client chooses them. The default is no demographic matching at all.
- **Matching is computed server-side.** Advisers' demographic attributes are never transmitted to the client device. The client receives a neutral list of matches and a fit score; the underlying identity values do not leave the server.
- **Lived experience is a first-class dimension.** Advisers may describe experience such as navigating black tax, first-generation professional status, or single parenthood. These are frequently more predictive of genuine advisory fit than any demographic label, and they are chosen rather than ascribed.
- **Only verified, licensed advisers can receive a client.** Regulatory authorisations are verified by a human reviewer against the relevant regulator's own records, are set by the platform operator, and cannot be self-edited; matches are filtered to the client's regulatory jurisdiction, and no financial information is shared without explicit consent.

A client who selects no dimensions receives advisers ranked on service fit alone. The system encodes a preference the client actually expressed, a stronger legal and ethical position than either ignoring the preference or imposing one on their behalf.

8.4 The advisory network

Demand routed to practices that cannot convert it is wasted. The network layer is therefore a capability programme, not a lead list. Participating practices receive structured development in financial therapy and behavioural coaching — a field with a peer-reviewed literature since 2010 (Grable, McGill and Britt, 2010), and treated at length in Chapter 2 of the companion volume.

The rationale follows directly from Section 6. A client arriving from a diagnostic has just been shown evidence that their family is exposed. That is a state of elevated arousal, and whether it converts into action or into helplessness depends on how the adviser handles the first ten minutes. An adviser trained to work with that reaction rather than past it converts better and, more importantly for the model's economics, retains the client.

9. Precedent: distribution as capital formation

The proposition that an advice margin can capitalise a class of enterprises is not speculative in South Africa. It has been executed here before, at scale, and the institutions it built still dominate the market.

In December 1917 a group meeting in Cape Town resolved to create an Afrikaans-controlled insurer. Santam was registered on 28 March 1918 with Willie Hofmeyr as chairman, and Sanlam followed on 8 June 1918 as its life assurance offshoot. Hofmeyr chaired Sanlam's board until his death in 1953. Sanlam's early motto was "*Uit die volk gebore om die volk te dien*"

— born from the people to serve the people. Verhoef (2008) analyses the company as elite social capital deliberately deployed for the economic upliftment of a poor community.

The mechanism was distribution. Afrikaans-speaking agents sold policies to Afrikaans-speaking households, and the premiums, commissions and institutional capital stayed inside the community rather than flowing to established competitors. What began as an insurance company became a capital formation engine.

The sequence deserves accuracy. The first Ekonomiese Volkskongres met in Bloemfontein in October 1939, where Sanlam's actuary M. S. Louw proposed an Afrikaner finance house. Federale Volksbeleggings followed in 1940 — capitalised not by Sanlam's premium reserves but by a public share issue that Sanlam prepared and promoted. The direct conversion of premium income into industrial equity came later, through Bonuskor in 1946, which channelled policyholders' bonuses into shareholdings in Afrikaner-led enterprises. By the 1980s these holdings were consolidated in Sankorp. Sanlam today is one of Africa's largest financial services groups (Verhoef, 2018).

This history must be cited honestly. The Afrikaner economic movement was interwoven with the Broederbond and the National Party, was directed at white upliftment specifically, and helped finance the political project that produced apartheid (O'Meara, 1983). Nothing in that record is a model for anything. What is separable, and what scholars treat as a case of ethnic economic mobilisation (Giliomee, 2008), is the financial mechanism: insurance distribution margin, retained within a community, is one of the most effective engines of group capital formation in South African economic history. The mechanism can be operated inside a constitutional, race-neutral, consumer-led framework. The argument here is that it should be.

There is a final turn that makes the point better than any argument. In 1993 Sankorp sold 10 per cent of Metropolitan Life to a black consortium through Methold, and later a further 20 per cent of its remaining interest; Methold was renamed New Africa Investments Limited before listing in 1994. It is widely recorded as South Africa's first major black economic empowerment transaction, and it was consciously modelled on the Afrikaner empowerment experience. The institution built by community distribution went on to seed the template for its successor. This paper proposes to complete that sequence by supplying the input the 1993 transaction could not: the demand pipeline that makes an advisory business viable without concessional capital.

A note on social engineering. The design proposed here deliberately shapes an environment in which particular outcomes become more likely, and it is better to say so than to pretend otherwise. Popper's (1945) distinction is the relevant one: utopian social engineering seeks to remake a society toward a blueprint and cannot be corrected by evidence, whereas piecemeal social engineering makes specific, bounded interventions whose effects can be measured and reversed. This model is piecemeal in that sense. It changes one variable: whether a person knows their own shortfall and can reach a practitioner they can trust, and every claim made for it is stated as a testable quantity. A pilot that returns a null result falsifies it. That is the difference between engineering and ideology.

PART IV — THE ECONOMICS

10. The economic model

The preceding sections argue that the instrument should work. This section asks whether it pays for itself, and for whom. The model is deterministic and fully specified; every figure below is computed from the parameters in Appendix A.

Three conventions govern every table that follows. Cover is reported in force rather than ever sold, so each gap-closure figure carries the persistency haircut of Table 3 and the model never takes credit for a policy that lapsed. All amounts are in December 2024 rands, neither inflated nor discounted. And practice revenue is gross regulated remuneration before costs, tax and clawback, which is the top line of a business rather than its profit.

The coverage ratio carries one qualification, and it cuts against the model rather than for it. The ratio is computed against a fixed denominator — the assessed need of R82.2 trillion at December 2024 — so it moves only by what this model adds. The underlying trend is held still, and on the sector's own measurement that trend is downward, since need has been growing at 9.3 per cent a year against 5 per cent for cover. A ratio rising to 52.3 per cent by year thirty is what the model contributes, not where the country would stand.

One constraint inside the commission arithmetic should be stated rather than left to inference. Part 3A caps primary commission at the lesser of two limbs, and which limb binds depends on the premium-paying term. At R1,500 a month the 85 per cent limb yields R15,300 while the 3.25 per cent limb yields R585 for each year of term; the two are equal at a term of

about twenty-six years, and below that the 3.25 per cent limb binds and pays less. The per-policy figures therefore assume a term of roughly twenty-six years or longer — a thirty-five-year-old insuring to age sixty-five is the case in mind. Because both limbs scale with the premium, that threshold is identical in every scenario of Table 8.




Figure 3 The diagnostic-led distribution model. Demand is activated by legibility rather than prospecting, routed through a commercially neutral choice architecture, and converted by practices equipped with behavioural and financial-therapy capability.

10.1 Anchors and assumptions, separated

The model distinguishes strictly between externally verified anchors and modelled assumptions. Anchors are sourced in Appendix A. Assumptions are labelled, defended, and tested in Section 10.6.

Modelled assumption	Value	Basis
Scan-to-policy conversion	10%	Against 1.5–3% typical cold-lead conversion; assumes diagnostic priming and behavioural capability
Cover purchased per client	R3.1m	The full average per-earner shortfall from the ASISA study
Premium per client	R1,500 / month	13% below the R1,715 implied by ASISA's own 8.2%-of-earnings estimate
Commission realisation	90% of cap	Allows for clawback, product mix and sub-cap arrangements
Policies in force at year 5 / 30	90% / 72%	Persistency haircut applied to all gap-closure figures
PVNBP per policy	R180,000	Annualised premium × 10 discounted premium-years
Insurer VNB margin	3% of PVNBP	SA reported range 0.6–2.8% blended; higher for pure protection books
Network size	500 practices	Phase-one target

Table 3 Modelled assumptions. Every figure in Sections 10.2 to 10.6 is generated from these inputs together with the verified anchors in Table 1.


One assumption carries most of the weight: the 10 per cent conversion rate. It is the difference between a transformative model and an ordinary one, and it is the parameter most exposed to challenge. Section 10.6 therefore models a 5 per cent case, and Section 12 treats it as the primary risk. Section 6 explains why a diagnostic engineered for salience without shame should outperform a cold lead by this margin, but explanation is not evidence, and a pre-registered pilot should measure it before capital is committed.

10.2 Deployment path and coverage

The table reports cover in force, not cover ever sold, and the coverage ratio rather than the rand gap. The rand gap grows with nominal earnings; the coverage ratio does not, which makes it the honest measure of whether the country is gaining ground.

Horizon	Policies sold	In force	Cover added	% of gap	Coverage	Premium in force
Baseline	—	—	—	—	38.8%	—
Year 5	350,000	315,000	R0.98tn	1.94%	39.9%	R5.67bn
Year 10	1,000,000	850,000	R2.63tn	5.23%	42.0%	R15.30bn
Year 15	1,800,000	1,458,000	R4.52tn	8.97%	44.3%	R26.24bn
Year 20	2,500,000	1,950,000	R6.04tn	11.99%	46.1%	R35.10bn
Year 30	5,000,000	3,600,000	R11.16tn	22.14%	52.3%	R64.80bn

Table 4 Deployment path on an in-force basis. Coverage ratio is measured against the assessed need of R82.2 trillion at December 2024. Cumulative policies of 5 million over thirty years represent 31 per cent of the current earner base.



Coverage ratio measured against the assessed need of R82.2 trillion (December 2024). Modelled on cover in force, net of assumed lapses.

Figure 4 National protection coverage ratio under the model.

10.3 What one policy produces


The unit economics are set by regulation, and the regulation is more restrictive than is often assumed. Under Part 3A of the Regulations to the Long-term Insurance Act 52 of 1998, primary commission on a recurring-premium risk policy is capped at the lesser of 3.25 per cent of total premiums over the premium-paying term and 85 per cent of the first year's annualised premium. Secondary commission may not exceed one third of primary, and is payable from the start of the second premium period. Maximum total intermediary remuneration is therefore 113.33 per cent of one year's premium, paid across roughly thirteen months. Commission is not clawed back at a cliff but recalculated month by month against a published scale, and is fully earned only at 24 months; death, disability and health-event terminations are excluded.

A correction to how this model is often presented. Secondary commission is a single further tranche capped at one third of primary, it is not a perpetual annual renewal stream on the book. Advisory revenue in risk business is therefore flow-

based: it is earned as new business is written, not as an annuity on business already written. Any model that projects a permanent renewal percentage on an accumulating book materially overstates practice revenue. The figures in this paper do not.

Per policy at R1,500 per month (R18,000 annualised)	Amount	Timing
Primary commission cap (85.00%)	R15,300	At inception
Secondary commission cap (28.33%)	R5,100	From month 13
Regulatory maximum (113.33%)	R20,400	Over ~13 months
Modelled realisation at 90% of cap	R18,360	Used throughout
Insurer present value of new business premiums	R180,000	Discounted
Insurer value of new business at 3%	R5,400	At inception
Cover placed in force for the household	R3,100,000	Until lapse or claim

Table 5 Per-policy economics. The protective benefit to the household is approximately 170 times the advice margin that mobilises it.



A single R1,500-per-month policy. The logarithmic scale is necessary because the protective benefit is roughly 170 times the advice margin that mobilises it — that asymmetry, not the size of the commission, is the economic case.

Figure 5 What one policy produces for each party, on a logarithmic scale.

10.4 Phase one: the enterprise development engine

The first five years are where the enterprise development case is made or lost. The question is not whether the aggregate flows are large (they are) but whether the flow per practice is enough to convert a micro-practice into the kind of institution Marais describes as absent.

R2.57 million a year of gross commission revenue is the threshold figure. It sustains a principal, two to three support professionals, a compliance function and the systems a licensed practice requires and this is the difference between a sole proprietor buying their own leads and a business with a balance sheet, a succession plan and a saleable client book. Aggregate that across a network and it is also, for the first time, the raw material from which a firm of scale could consolidate.

The qualification: 140 new fully underwritten clients a year per practice is demanding. It is achievable for a practice receiving pre-qualified introductions rather than prospecting for them, but it assumes the routing engine performs. Practice capacity, not consumer demand, is the more likely binding constraint, and network onboarding should be paced accordingly.

Phase one, years 1–5	Value
Diagnostic scans deployed	3,500,000
Policies written at 10% conversion	350,000
New annualised premium written	R6.30 billion
Total intermediary remuneration	R6.43 billion
Insurer value of new business	R1.89 billion
Cover placed in force	R0.98 trillion
New advised clients per practice	700 (140 per year)
Gross commission revenue per practice	R12.85 million (R2.57 million per year)
Direct professional roles across the network	≈ 1,250

Table 6 Phase-one outcomes across a 500-practice network. Revenue is gross regulated remuneration before practice costs.

Whether the programme funds itself. The diagnostic carries a list price of R100 a scan, and the model prices sponsorship at list rather than at cost: an instrument subsidised below its own economics is one that ends when the subsidy does. Phase one's 3.5 million scans therefore represent R350 million over five years, against the R6.43 billion of regulated remuneration those same scans mobilise — 5.4 per cent. The ratio matters more than either figure. Diagnostic spend is not a cost that must be raised before the mechanism can begin; on these assumptions it is an input recovered roughly eighteen times over out of the flow it creates, inside the same five-year window. The ratio holds at the level of a single policy: ten scans at R100 produce one policy carrying R18,360 of intermediary remuneration and R3.1 million of cover, which is R1,000 of diagnostic spend for every R3.1 million of protection placed in force. The significance is structural rather than promotional. A remedy whose running costs are met from the activity it generates does not depend on the persistence of any institution's appetite for funding it, and that independence is the property most conspicuously absent from the education programmes examined in Section 3. Who carries that cost is a separate question from whether it can be carried, and the model does not assume a philanthropic funder. The parties who pay are the parties who capture the return. An adviser practice funds scans as client-acquisition expenditure against the regulated remuneration of Section 10.3; an employer or retirement fund funds them against a workforce whose protection gap is its own contingent liability and whose financial distress it already absorbs in absenteeism, two-pot withdrawals and staff turnover; a product house funds them against the new business the mechanism originates. Each is buying something it would otherwise buy less efficiently, and none is subsidising the others. This is what distinguishes the proposal from a financial-education programme, which must be funded by a party that receives nothing measurable in return and is therefore always the first line cut.

What concessionary capital would add. It follows that concessionary capital is not a precondition for the mechanism described here. Where it were deployed — the instruments of Section 10.4, and the informally employed majority identified in Section 10.5.1 — it would buy reach into segments the commission economics genuinely cannot serve, rather than underwriting the model's viability. The distinction is worth drawing because the two are often conflated in discussions of financial inclusion: a remedy that requires funding in order to exist and a remedy that can absorb funding in order to extend are different objects, and only the second is being described.

10.5 The public-sector channel

The Government Employees Pension Fund had approximately 1.27 million active members at 31 March 2025, holding assets of some R2.7 trillion (Government Employees Pension Fund, 2025). It is the largest single concentration of formally employed, middle-income earners in the country. It is also the clearest case in the economy of a household whose retirement provision is secure and whose exposure to a death or disability shock is nonetheless real — not because the Fund's provisions are deficient, but for three reasons that sit outside the Fund's rules entirely.

The first is a timing mismatch written into the benefit formula. Where a member has ten or more years of pensionable service, the death gratuity is the sum of two components — 6.72 per cent of final salary multiplied by pensionable service, and one fifty-fifth of final salary multiplied by pensionable service and by five — each computed on a service base enhanced by the lesser of one third of pensionable service or the unexpired period, capped at five years (Government Employees Pension Fund, 2024). The surviving spouse's annuity is computed on a different base again, adding the full unexpired period without that cap. The benefit therefore rises with service, while a household's obligations peak early. A member with ten years' service

on a R400,000 salary leaves a gratuity of roughly R843,000 at precisely the point in life when the bond has just been registered, the vehicle is financed and the children are young; a colleague on the same salary with thirty years' service leaves roughly R2.21 million against a bond that is nearly settled. The cover is smallest where the claim on it would be largest, and at ten years it stands against an average per-earner shortfall the ASISA study puts at R3.1 million. This is not a defect in the Fund, since a service-based formula is the correct way to fund a pension. It is a defect in the assumption that a pension formula also constitutes life cover.

The second reason is behavioural, and it is produced by the Fund's own comprehensiveness. Because death, disability and ill-health benefits arrive automatically with employment, they are never chosen and never priced by the member, and what is never priced is rarely reviewed. Middle-income public servants consequently under-purchase supplemental life and salary-protection cover relative to the commercial debt they carry. The automatic benefit crowds out the discretionary one. A member who has never seen the gratuity quantified against the outstanding bond has no reason to imagine a gap, and the Fund has no mandate to point one out. This is the legibility problem of Section 1 in its purest institutional form: the information exists, the entitlement exists, and the household still cannot see the shortfall between them.

The third pressure is demographic and recent. The gratuity and the spouse's annuity were designed around a dependency profile in which children left the household on completing school or a first degree. That profile no longer describes South Africa. In the first quarter of 2026 the unemployment rate for those aged 25 to 34 was 40.6 per cent, and 10.6 million people aged 15 to 34 were outside the labour force altogether (Statistics South Africa, 2026a). A public servant's adult children are now statistically likely to remain financially dependent well past graduation, and some extend that dependency further through postgraduate study. The household's real dependency period has lengthened by years while the benefit that would fund it has not changed. A gratuity adequate for a widow and two schoolchildren in 1996 is a materially different instrument from the one required by a widow supporting two unemployed graduates in 2026.

Offering the diagnostic to this membership is administratively simple since the instrument carries no product recommendation. As Section 5.1 argues, a diagnostic asks to be read rather than believed; no employer or fund distributes financial advice by making one available.

Conversion	Policies	Cover placed	New premium	Adviser remuneration
5%	63,377	R0.20tn	R1.14bn	R1.16bn
10%	126,754	R0.39tn	R2.28bn	R2.33bn
15%	190,131	R0.59tn	R3.42bn	R3.49bn

Table 7 Illustrative outcomes for the GEPF active membership alone.

10.5.1 The uncovered private market

The public-sector channel is the most concentrated market, not the largest one. It is worth stating plainly what the private sector already has, because the gap is not where it is usually assumed to be.

Retirement coverage in the formal sector is usually described as high, but the figure normally quoted is old. The two-thirds to four-fifths range that circulates in industry commentary traces to National Treasury's 2004 discussion paper, which put coverage of formal-sector employees at 66 to 84 per cent (National Treasury, 2004). Current measurement is materially lower: of 14.3 million employees in the fourth quarter of 2025, 6.3 million contribute to a pension or retirement fund, or 44.4 per cent (Statistics South Africa, 2026b). Since public-sector employees are close to universally covered, the implied formal private-sector rate is lower still. On retirement provision that segment is under-provided rather than uncovered, which is a different problem requiring a different argument; but it is less well covered than the standard citation suggests.

The genuinely uncovered market is defined by the absence of employment protections rather than by sector. Statistics South Africa's informal-employment measure — workers who lack access to social insurance including pensions and paid leave, and by extension to employer-arranged risk benefits — stood at 5.7 million people, or 33.5 per cent of the 17.1 million employed, in the fourth quarter of 2025 (Statistics South Africa, 2026b). On the separate classification by sector of the employing enterprise, 12.3 million worked in the formal sector, 3.7 million in the informal sector and 1.1 million in the household sector. Nor is informal employment confined to the poorly educated: 18 per cent of employed university graduates are within it. Measured against the Fund's 1.27 million active members, the population with no employer-arranged retirement or risk provision whatsoever is more than four times larger. The comparison should be read with one methodological caution: Statistics South Africa refined its definitions of informality from the third quarter of 2025, so these estimates are not comparable with series published before that date.

Within this population sits a segment that matters disproportionately to the economics of Section 10.3: own-account workers, sole proprietors and the owners of small registered businesses. They have no occupational scheme for the structural reason that they are the employer, yet many earn middle incomes, service commercial debt and support the same dependants as their salaried peers. They are simultaneously among the least covered and the most able to pay for cover of any group in the economy, and no distribution channel is presently organised around them.

This asymmetry has a consequence the model should state rather than elide. The practice economics set out in Section 10.3 depend on a premium the household can actually sustain, and they do not transfer to the informal 5.7 million on any honest reading: a monthly premium of R1,500 is not available to a street trader. The public-sector and own-account channels are therefore where the self-funding mechanism works on its own terms, while the informal majority is where the enterprise-development and concessionary instruments of Section 10.4 are the appropriate route. Treating the two as a single addressable market would overstate the reach of the commission model and understate the case for development capital. The correct reading is one diagnostic serving two funding rails.

10.5.2 The assumption of adequacy

The two preceding channels were described by who is covered. This one is described by what being covered is assumed to mean. A household inside an employer arrangement treats cover as a state rather than a quantity: the member knows that he is covered, and does not know for how much, against what, or whether the amount bears any relation to what his dependants would need. Having cover and being covered are treated as the same fact, and the distance between them is where the exposure sits.

That distance is measurable. Against the R50.4 trillion shortfall with which this paper opens, cover in force meets 39 per cent of assessed need — 37 per cent for death and 40 per cent for disability — down from a restated 44 per cent three years earlier (True South Advisory, 2025). Per earner, the position is more legible than the aggregate: life cover of R2,090,229 is required against R774,408 held, and disability cover of R3,028,179 against R1,209,610 held, the latter including the state disability grant. The gap is widening while employment and scheme membership are not falling, which locates the problem somewhere other than access.

The same study locates it precisely, and in terms that bear directly on this section: middle- and upper-income earners, it records, remain heavily under-insured despite access to group schemes. Group arrangements are not a marginal contributor to national cover. They account for approximately 40 per cent of life cover in force and, once state grants are set aside, roughly 63 per cent of disability cover. The population that believes an employer has resolved the question is therefore large, and on the study's own figures its belief is not well founded.

Three structural features explain why, and none is a defect in any particular scheme. The first is that an employer benefit is denominated in the wrong unit. Group cover is set as a multiple of pensionable salary because salary is the variable an employer administers and can price. A household's requirement is set by its liabilities and its dependency period — the outstanding bond, the years until the youngest child is self-supporting, the income a surviving spouse would not earn. The two quantities are unrelated, and a salary multiple will approximate a household's need only by coincidence. It cannot track the events that change that need, since neither a new bond nor a third child alters the multiple. The benefit is stable precisely where the need is volatile.

The second is that the multiple is neither generous nor static, and that what schemes are designed to provide is not what members are found to hold. The Sanlam Benchmark survey, which asks employers about scheme design, records average approved death benefits in umbrella funds of 2.8 times salary in 2025, down from 3.1 times in 2019, with lump-sum disability multiples falling from 2.8 to 2.4 times over the same period (Sanlam Corporate, 2025). Alexforbes, measuring actual in-fund benefits across its administered book rather than surveying employers, reported an average insured death benefit of 1.94 times annual salary and recorded that a majority of members in the funds analysed were entitled to no insured death benefit at all (Alexander Forbes, 2020). The two figures are not in conflict; they answer different questions, and the unit of analysis must be stated whenever either is quoted. The distance between them is itself part of the phenomenon, and no member can determine from anything in his possession which of the two describes him.

The third is that the cover which does exist is not the member's to direct. Where a benefit is approved, and so payable by the fund, section 37C of the Pension Funds Act 24 of 1956 removes it from the deceased's estate and places its distribution with the trustees, who must identify dependants and apportion the benefit as they consider equitable. A beneficiary nomination is a factor the trustees must weigh and is not binding on them, a position settled in *Mashazi* and reaffirmed by the Constitutional Court in *Mutsila*, which quotes *Mashazi* for exactly this proposition (Republic of South Africa, 1956; *Mashazi v African Products Retirement Benefit Provident Fund*, 2003; *Mutsila v Municipal Gratuity Fund and Others*, 2025). A member who has completed a nomination form in favour of his spouse has expressed a preference rather than given an instruction,

and will ordinarily believe otherwise. One exception runs the other way: where the fund traces no dependant within twelve months and the member has designated a nominee, section 37C(1)(b) requires payment to that nominee, and the board may not deviate from it.

None of this is concealed. All of it is disclosed — in fund rules, benefit statements, conditions of service and the Act itself. It is nonetheless not legible in the sense this paper uses, because the entitlement is recorded in one place, the liability in another, and the arithmetic that would relate them is performed nowhere and owed by no one. The observation is not confined to academic commentary. Writing from inside the employee-benefits profession, Peens (2026) argues that institutional and retail advice have been organised as separate industries around a member who has only one financial life, with the consequence that a retail adviser may know a client's retirement annuity, tax-free savings account and estate plan without ever asking for the pension fund benefit statement that is frequently the client's largest asset. His conclusion states the structural point exactly: every party can perform its mandate correctly, every service level can be met and every governance report can be green, and the member can still reach a poor financial outcome. The profession has defined with precision where each party's responsibility ends, and has not settled where responsibility for the member's whole position begins.


This is why the instrument proposed here is a diagnostic rather than an educational intervention. A member who has been taught what a salary multiple is still cannot say whether his own is sufficient, because sufficiency is a calculation about his particular liabilities that no party is presently obliged to perform. What is missing is not comprehension but arithmetic, and arithmetic is the one thing that can be supplied at scale and without advice.

10.6 Sensitivity

The base case is not the case for the model. The conservative case is. At half the assumed conversion, two-thirds the assumed cover and two-thirds the assumed premium, the model still places R0.32 trillion of cover in force in five years and delivers R4.28 million of revenue per practice — enough to sustain a viable licensed business.

Scenario	Conv.	Premium	Cover	Policies	In force	Adviser revenue	Per practice
Conservative	5%	R1,000	R2.0m	175,000	R0.32tn	R2.14bn	R4.28m
Base	10%	R1,500	R3.1m	350,000	R0.98tn	R6.43bn	R12.85m
ASISA-anchored	10%	R1,700	R3.1m	350,000	R0.98tn	R7.28bn	R14.57m
Ambitious	15%	R1,800	R3.1m	525,000	R1.46tn	R11.57bn	R23.13m

Table 8 Five-year sensitivity on 3.5 million scans. The ASISA-anchored row uses the R1,715 monthly premium implied by the study's own 8.2%-of-earnings estimate for closing an average earner's full gap.



Assumes a 500-practice network sharing the modelled policy flow. Revenue is gross regulated remuneration before practice costs, and is earned as new business is written rather than as an annuity on the existing book.

Figure 6 Practice-level economics across scenarios.

10.7 Transmission to the wider economy

Claims about GDP effects are easy to make and hard to defend, so this section states mechanisms and bounds rather than a single growth number. Four channels operate.

- **Financial sector value added.** Finance, real estate and business services is the largest industry in the South African economy. New intermediated premium adds directly to measured sector output, though the increment is modest relative to a R7.6 trillion (2025) economy.
- **Consumption stabilisation.** Household final consumption is roughly two-thirds of GDP. The ASISA study estimates that a household losing an earner must cut expenditure by about a third, or find several thousand rand a month in replacement income. Cover converts those contractions into insured events. This defends existing output rather than creating new output — a distinction aggressive versions of this argument tend to blur.
- **Capital formation.** Life insurers ended 2025 with R5.2 trillion in assets and are among the largest domestic allocators of long-term capital. Premium growth deepens that pool and supports gross fixed capital formation.
- **Employment.** IMF estimates place South Africa's short-run employment response to growth at the top of the G20 (An, Gonzalez Prieto, Loungani and Mishra, 2016). Output growth in South Africa converts into jobs more readily than in almost any comparable economy, which raises the return on any intervention that expands a labour-intensive professional service.

The direct employment effect is the one that can be stated with confidence: approximately 1,250 professional roles across a 500-practice network in phase one, plus principals, plus insurer-side underwriting, administration and claims capacity. Against official unemployment above 30 per cent, that is a small number nationally and a significant one for the specific labour market it serves.

11. What each stakeholder gets

Stakeholder	What the model delivers
Households	Cover in force against the events that most reliably destroy black intergenerational wealth. R3.1 million of protection for roughly R1,500 a month, reached without being sold to or characterised as illiterate.
Advisory practices	R12.85 million of gross revenue per practice over five years, from pre-qualified introductions rather than paid prospecting. Revenue anchored to client income rather than client assets.
Insurers	R6.3 billion of newly written annualised premium and R1.89 billion of value of new business in phase one, from a segment currently outside the advised market — and distribution transformation delivered commercially rather than through procurement spend.
NEF and development financiers	Enterprise development with a revenue model attached. Capital deployed into practices that have a demand pipeline is repayable capital; capital deployed into practices that do not is a grant by another name.
The profession	A route to the large black-owned advisory firm the industry says is missing, built by commercial unit economics rather than quota, and an answer to the question of what financial planning is for.

Table 9 Stakeholder value.

PART V — WHAT WOULD MAKE THIS FAIL

12. Risks and limitations

A model presented without its failure modes is a brochure. These are the conditions under which this one does not work.

- **The model does not close the gap.** At 12.5 per cent annual growth, the rand shortfall outruns any single distribution intervention. Over thirty years the model places R11.16 trillion of cover in force and lifts the national coverage ratio from 38.8 to 52.3 per cent. That is a material change in national protection. It is not closure, and this paper does not claim otherwise.
- **Conversion is the load-bearing assumption.** Ten per cent is defensible against a 1.5 to 3 per cent cold-lead baseline given diagnostic priming, but it is unproven at scale in this market. It should be measured in a pre-registered pilot before capital is committed, with a commitment to publish a null result as readily as a positive one.

- **A diagnostic can be captured.** This is the screening analogy's sharpest warning. An instrument tuned to overstate need becomes a lead-generation engine wearing clinical clothing, and the epistemic argument in Part II collapses the moment the output stops being an honest measurement. Governance of the calculation, independent review of the needs methodology, published assumptions, and separation between whoever funds deployment and whoever sets the parameters, is not optional hygiene. It is what the entire case rests on.
- **Screening has documented harms.** Medicine records overdiagnosis, false positives and induced anxiety. The financial analogue is a household persuaded to over-insure at the expense of more urgent needs, or one left alarmed and unserved because no adviser followed up. Section 6's specification mitigates the second; only honest calibration mitigates the first.
- **Persistency, not conversion, determines whether practices survive.** Commission is clawed back if premiums cease within 24 months. A practice that converts well and retains badly will fail, and will damage the network's standing with underwriters.
- **Advisory revenue is flow-based, not annuity-based.** Risk commission is earned as new business is written. Durable practice value requires converting risk clients into ongoing advice relationships over time. The risk policy is the entry point; the advice relationship is the asset.
- **Practice capacity is the likely bottleneck.** 140 new underwritten clients per practice per year is demanding. Network growth must be paced to onboarding capability, and the model should throttle demand rather than degrade service quality.
- **Affordability constrains the addressable market.** Closing an average earner's full gap costs 8.2 per cent of earnings. For lower-income earners that is not available, which is why the model targets the middle-income segment and does not claim to solve underinsurance among the working poor.
- **Regulatory change is a live risk.** Commission on risk business remains governed by Part 3A of the Long-term Insurance Act Regulations, but Retail Distribution Review proposals have been carried into the Conduct of Financial Institutions framework rather than implemented as drafted. The Conduct of Financial Institutions Bill was introduced in the National Assembly on 17 April 2026 and is not yet in force. A future reform of risk-commission structures would change the model's unit economics.
- **Epistemic dependency is a real objection.** A population that outsources self-knowledge to an instrument may reason less well about its own finances over time. The answer this paper offers is structural rather than reassuring: because the diagnostic terminates in a human advisory conversation rather than an automated transaction, it restores deliberation rather than replacing it. That answer should be tested, not assumed.
- **Compliance is continuous, not initial.** The diagnostic must remain non-advisory in substance and not merely in labelling; demographic data collection must satisfy POPIA on lawful basis, minimisation and consent; and referral arrangements must not become undisclosed inducements.

13. Conclusion

A protection gap that widens by R15 trillion in three years, while the primary educational remedy prescribed for it is continuously funded and applied, exposes a policy framework whose underlying diagnosis has not earned its position. This paper has argued that the diagnosis is wrong in a specific and correctable way. The problem is not that South Africans know too little about finance in general. It is that they have no legible knowledge of their own position, no tractable route to a practitioner who would supply it, and a sector that describes their predicament in a vocabulary which measurably makes it worse.

The alternative is unglamorous and mechanical. Make the exposure legible rather than teaching the theory. Engineer the disclosure for salience without shame, because the profession's own research says that is the combination that produces action. Hand the client a portable statement of their own position so that they can hold an adviser to account rather than having to evaluate one blind. Let them choose who they see, on dimensions they select. And allow the regulated economics of advice to do what they have done before in this country: build institutions out of distribution.

The distributional consequence follows from demographic weight and documented preference, not from a rule anyone has to defend. That is the model's most important property. It does not ask South Africans to choose sides at the moment they are deciding whether to protect their families. It asks them only to look at their own number, and then makes sure somebody qualified, and somebody they trust, is waiting on the other side of that decision.

Jeanette Marais called the absence of a large black-owned advisory firm an opportunity for the industry. It is more than that. It is the specific structural fact that keeps 85 per cent of South Africa's income earners underserved by the profession that exists to serve them, and it has survived every intervention aimed at it because those interventions supplied everything except clients. This model supplies clients. Whether it supplies enough of them is an empirical question, and this paper has been written so that the question can be settled rather than argued.

Chapter 2 of *Portfolio Management for Financial Advisors* argued that financial advisers are money doctors. Chapter 9 argued that the profession should take responsibility for its macroeconomic footprint. This is what those two claims look like when they are made to meet: a screening programme for financial mortality risk, a referral pathway into a profession rebuilt to receive it, and an arithmetic that pays for itself. Financial planning for the greater good need not remain a curricular aspiration. The mechanism is built and deployed, the economics are regulated and knowable, and the first cohort can be measured within eighteen months.

References

- Ahlstrom-Vij, K. 2013. Epistemic paternalism: a defence. Basingstoke: Palgrave Macmillan.
- Alexander Forbes. 2020. Risk benefits and expenses deductions. Member Watch™ 2019–2020. Sandton: Alexander Forbes.
- Alvarado, R. 2023. AI as an epistemic technology. *Science and Engineering Ethics*, 29, article 32.
- An, Z., Gonzalez Prieto, N., Loungani, P. and Mishra, S. 2016. Does growth create jobs? Evidence for advanced and developing economies. *IMF Research Bulletin*, 17(3): 1, 5–8.
- Association for Savings and Investment South Africa. 2025. South Africa's life and disability insurance shortfall widens to R50.4 trillion. Media release, 28 October 2025.
- Becker, E. 2023. Tackling the "audacious goal" of boosting black advisers. *Citywire South Africa*, 20 June 2023.
- Britt, S.L., Lawson, D.R. and Haselwood, C.A. 2016. A descriptive analysis of physiological stress and readiness to change. *Journal of Financial Planning*, 29(12), pp. 45–51.
- Cummings, B.F. 2017. Financial planners and advisors. In: Baker, H.K., Filbeck, G. and Ricciardi, V. (eds.) *Financial behavior: players, services, products, and markets*. New York: Oxford University Press, pp. 97–117.
- de Bruin, B. 2021. Epistemic injustice in finance. *Topoi*, 40(4), pp. 755–763.
- Fernandes, D., Lynch, J.G. Jr. and Netemeyer, R.G. 2014. Financial literacy, financial education, and downstream financial behaviors. *Management Science*, 60(8), pp. 1861–1883.
- Financial Planning Institute of Southern Africa. 2025. FPI surpasses 5 000 CFP professionals. Media release, 30 January 2025.
- Financial Planning Standards Board. 2025. *Impact of AI on financial planning: global research report 2025*. Denver, CO: FPSB Ltd.
- Financial Sector Conduct Authority. 2025. Integrated report 2024/25. Pretoria: FSCA.
- Frank, J.D. 1974. *Persuasion and healing: a comparative study of psychotherapy*. Revised edition. New York: Schocken Books.
- Freiman, O. 2024. Analysis of beliefs acquired from a conversational AI: instruments-based beliefs, testimony-based beliefs, and technology-based beliefs. *Episteme*, 21(3), pp. 1031–1047.
- Fricker, M. 2007. *Epistemic injustice: power and the ethics of knowing*. Oxford: Oxford University Press.
- Gennaioli, N., Shleifer, A. and Vishny, R. 2015. Money doctors. *Journal of Finance*, 70(1), pp. 91–114.
- Giliomee, H. 2008. Ethnic business and economic empowerment: the Afrikaner case, 1915–1970. *South African Journal of Economics*, 76(4), pp. 765–788.
- Government Employees Pension Fund. 2024. Rules of the Government Employees Pension Fund. Pretoria: GEPF.
- Government Employees Pension Fund. 2025. Annual report 2024/25. Pretoria: GEPF.
- Grable, J., Heo, W. and Rabbani, A. 2015. Financial anxiety, physiological arousal, and planning intention. *Journal of Financial Therapy*, 5(2), Article 2.
- Grable, J., McGill, S. and Britt, S. 2010. The Financial Therapy Association: a brief history. *Journal of Financial Therapy*, 1(1), Article 1.
- Gray, J.W. 2012. Bi-polar: college education and loans to small businesses headed by black females. *Review of Black Political Economy*, 39(3), pp. 361–371.
- Hayek, F.A. 1945. The use of knowledge in society. *American Economic Review*, 35(4), pp. 519–530.
- Kaiser, T., Lusardi, A., Menkhoff, L. and Urban, C. 2022. Financial education affects financial knowledge and downstream behaviors. *Journal of Financial Economics*, 145(2), pp. 255–272.
- Klontz, B.T. and Klontz, T. 2009. *Mind over money: overcoming the money disorders that threaten our financial health*. New York: Broadway Business.
- Klontz, B.T., Van Zutphen, N. and Fries, K. 2016. Financial planner as healer: maximizing the role of financial health physician. *Journal of Financial Planning*, 29(12), pp. 52–59.

Marais, J. 2024. Address to the PSG Financial Services annual conference, 9 May 2024. Reported in G. Theunissen, "SA needs 4 000 more financial planners and a large black-owned advisory firm, says MMH", News24 Business, 10 May 2024.

Mashazi v African Products Retirement Benefit Provident Fund 2003 (1) SA 629 (W).

McPherson, M., Smith-Lovin, L. and Cook, J.M. 2001. Birds of a feather: homophily in social networks. Annual Review of Sociology, 27, pp. 415–444.

MMI Holdings. 2017. Eight brokerages graduate from Momentum's Masibambane programme. Media release, 10 August 2017.

Momentum Group and Bureau of Market Research. 2025. Momentum Financial Advice Research Report 2024.

Mutsila v Municipal Gratuity Fund and Others (CCT 228/23) [2025] ZACC 17; 2026 (1) SA 1 (CC).

National Treasury. 2004. Retirement fund reform: a discussion paper. Pretoria: National Treasury.

Nguyen, C.T. 2020. Cognitive islands and runaway echo chambers: problems for epistemic dependence on experts. Synthese, 197, pp. 2803–2821.

O'Meara, D. 1983. Volkskapitalisme: class, capital and ideology in the development of Afrikaner nationalism 1934–1948. Cambridge: Cambridge University Press.

Peens, J. 2026. Your client has one financial life. Why advise them in two? Citywire South Africa, 20 August 2026.

Popper, K.R. 1945. The open society and its enemies. London: Routledge.

Republic of South Africa. 1956. Pension Funds Act 24 of 1956 (as amended), s 37C. Pretoria: Government Printer. Section 37C was inserted by s 24 of the Financial Institutions Amendment Act 101 of 1976 and is not part of the Act as originally enacted.

Republic of South Africa. 1998. Long-term Insurance Act 52 of 1998, Regulations, Part 3A, as amended.

Sanlam Corporate. 2025. The Sanlam Benchmark 2025 insights report. 44th edition. Bellville: Sanlam.

Sarpong, P. 2024. Portfolio management for financial advisors, Volume II. Centre for Financial Psychology and Systems. ISBN 9798877002630.

Sarpong, P. Forthcoming 2027. Applied AI epistemics: rethinking cognition, behavior and assessment in the age of generative intelligence. Cham: Springer.

Statistics South Africa. 2026a. Quarterly Labour Force Survey, Quarter 1: 2026. Statistical release P0211. Pretoria: Statistics South Africa.

Statistics South Africa. 2026b. Quarterly Labour Force Survey, Quarter 4: 2025. Statistical release P0211. Pretoria: Statistics South Africa.

Stolper, O. and Walter, A. 2019. Birds of a feather: the impact of homophily on the propensity to follow financial advice. Review of Financial Studies, 32(2), pp. 524–563.

Swire, P.P. 1995. Equality of opportunity and investment in creditworthiness. University of Pennsylvania Law Review, 143(5), pp. 1533–1559.

True South Actuaries and Consultants. 2025. The South African insurance gap (2025). Commissioned by ASISA. Report dated 3 October 2025.

Verhoef, G. 2008. Nationalism, social capital and economic empowerment: SANLAM and the economic upliftment of the Afrikaner people, 1918–1960. Business History, 50(6), pp. 695–713.

Verhoef, G. 2018. The power of your life: the Sanlam century of insurance empowerment, 1918–2018. Oxford: Oxford University Press.

Willis, L.E. 2011. The financial education fallacy. American Economic Review, 101(3), pp. 429–434.

Appendix A Model parameters and formulas

The model is deterministic and fully specified by the parameters below. All figures in Sections 10.2 to 10.6 are computed from these inputs; no figure is entered by hand.

A.1 Verified anchors

Parameter	Value	Source
Protection gap	R50.4tn	ASISA / True South, Dec 2024
Total assessed need	R82.2tn	ASISA / True South, Dec 2024
Cover in force	R31.8tn	ASISA / True South, Dec 2024
Gap at Dec 2021 (restated)	R35.4tn	ASISA 2025 release
Gap growth rate	12.5% p.a.	ASISA 2025 release
Income earners	16.1m	ASISA definition, Dec 2024
Shortfall per earner	R3.1m	R1.3m life + R1.8m disability
Cost to close per earner	8.2% of earnings	5.2% death + 3.0% disability
Primary commission cap	85%	LTIA Regulations, Part 3A
Secondary commission cap	⅓ of primary	LTIA Regulations, Part 3A
Life industry assets	≈R5.2tn	ASISA
GEPF active members	≈1.27m	GEPF Annual Report 2024/25

Parameter	Value	Source
CFP professionals	5,001	FPI, cited by Marais, May 2024
Authorised Category I and IV FSPs	11,336	FPI, cited by Marais, May 2024
Households using an adviser	9%	Momentum / BMR, 2025

A.2 Formulas

Cover in force at horizon t = policies sold(t) $\times$ persistency(t) $\times$ average cover

Coverage ratio(t) = $(R31,872bn + \text{cover in force}(t)) \div R82,225bn$

Adviser remuneration per policy = annualised premium $\times$ 113.33% $\times$ 90% realisation

Insurer value of new business per policy = annualised premium $\times$ 10 $\times$ 3%

Practice revenue = (policies sold in period $\times$ remuneration per policy) $\div$ 500 practices

A.3 Cross-checks performed

- R3.1m per earner $\times$ 16.1m earners = R49.9tn against the published R50.4tn — a 1% difference from rounding of the per-earner averages.
- The study's exact figures (need R82,225bn less cover of R31,872bn) reproduce the published gap of R50,353bn; the rounded anchors preserve the identity $R82.2tn - R31.8tn = R50.4tn$. Coverage ratios are computed on the exact figures throughout, which is why the baseline reads 38.8 per cent rather than the 38.7 per cent the rounded anchors would imply.
- The modelled premium of R1,500 per month is 13 per cent below the R1,715 implied by the study's own 8.2 per cent of average monthly earnings.
- Five million policies over thirty years equals 31 per cent of the current earner base.

Centre for Financial Psychology and Systems · Epistémica Institut