



ARKAYA RISK
Building resilience capital

Why evidence changes insurance pricing

The economics of pricing uncertainty, and what happens to the price when governance becomes verifiable. v5 · 1 July 2026 · Strategic paper · shareable.

The question

An insurer does not price risk. It prices what it knows about risk. The two are not the same, and the gap between them is where a large part of every premium is made. A premium covers the losses an insurer expects to pay, but it also covers what the insurer cannot see: the possibility that the true risk is worse than it looks, that the insured knows something the insurer does not, and that the capital held behind the policy must be large because the outcome is uncertain. Reduce what the insurer cannot see, and you reduce the part of the price that pays for not seeing it. This paper sets out how, and why verifiable evidence, rather than more data, is what moves the number.

What a premium is made of

A technical premium has more in it than expected loss. In its simplest honest form it is the expected loss, plus a loading for the uncertainty around that expectation, plus the cost of the capital the insurer must hold against the risk, plus expenses and a margin. Evidence acts on some of these and not others. It does not change the expected loss of a given risk, and it does not touch expenses or the target margin. It acts on two components that are larger than they look: the uncertainty loading, and the cost of capital. Understanding why means separating two kinds of uncertainty that are usually run together.

Two kinds of uncertainty

Some uncertainty is irreducible. A fair coin lands heads half the time, and no amount of information tells you which way the next toss falls. This is process uncertainty, the inherent randomness of outcomes drawn from a known distribution, and evidence cannot shrink it. But a great deal of what an insurer faces is not this. It is uncertainty about the distribution itself: what the true frequency is, what the tail looks like, whether this insured is like the ones in the data or unlike them. This is parameter uncertainty, and it is epistemic, a function of what is known rather than of what is random. Evidence acts only on this second kind. The reason it matters commercially is that the second kind is priced, often heavily and not only as premium: it surfaces as higher deductibles, tighter sub-limits, restrictive terms and, at the limit, exclusions. For new and fast-moving risks it dominates.

The price of not knowing the odds

The distinction is old. Frank Knight drew it in 1921: risk is a situation whose outcome is drawn from a distribution you can describe, uncertainty a situation whose distribution you cannot [1]. Daniel Ellsberg showed in 1961 that people do not treat the two alike: faced with a gamble whose odds are unknown, they demand more to bear it than the mathematics of expected value can justify, an aversion to ambiguity distinct from aversion to risk [2]. Insurers are no exception, and the evidence is from their own desks. Surveys of actuaries, underwriters and reinsurers by Kunreuther, Hogarth and Meszaros found all three groups quoting materially higher prices for ambiguous risks than for well-specified risks with the same expected loss, with the loading rising as the probability or the loss became harder to characterise, and identified insurer ambiguity as a contributor to market failure at the industry level [3]. The ambiguity premium is not a construct from the laboratory. It is observed underwriting behaviour,

and the rational response to an exposure with no credible loss history is exactly what the surveys record: load heavily, cap the cover, or decline it altogether.

This is the mechanism behind the exclusion of poorly understood risks, and it is why the treatment of AI-related governance risk has moved toward exclusion rather than pricing. The exposure is not obviously uninsurable in the actuarial sense; it is uncharacterisable, because the loss history that would let an underwriter set the odds does not yet exist. Uncharacterisable risk is priced at the ambiguity premium or, at the limit, at infinity, which is what an exclusion is. The only thing that shrinks an ambiguity premium is information that lets the underwriter replace a guess about the distribution with something closer to knowledge of it. For a risk with no actuarial past, that information cannot come from history. It has to come from evidence of the present state of the thing being insured.

The problem evidence solves is not the insurer's; it is the gap between the parties

There is a second, sharper reason evidence changes the price, and it is the most studied result in the economics of insurance. The insured usually knows more about its own risk than the insurer does. George Akerlof showed in 1970 that when a buyer cannot tell a good item from a bad one, the market prices the average, the good items are underpaid and withdrawn, and quality falls until the market can unravel [4]. Michael Rothschild and Joseph Stiglitz carried the result directly into insurance in 1976: under asymmetric information a single pooled price cannot survive, because the low-risk insured is subsidising the high-risk one and will leave for any insurer that can distinguish them [5]. Akerlof, Spence and Stiglitz shared the 2001 Nobel Memorial Prize for the analysis of markets with asymmetric information; the Rothschild and Stiglitz insurance result is the branch of that work this argument stands on [4][5].

The consequence for pricing is precise. When the insurer cannot separate the well-governed insured from the badly governed one, it prices the pool, and the well-governed insured pays for the sins of the average. The way out, in the theory and in practice, is separation: give the insurer a credible way to tell the risks apart, and each can be priced on its own merits, with the good risk earning the lower price it deserves. What makes this hard is that the insured always says its governance is good. A claim the insured makes about itself carries no information, because a bad risk makes the same claim; in the language of the field it is cheap talk. Separation requires a signal the insured cannot fake and the insurer can check without trusting the insured. Spence's 1973 insight was exactly this: only a costly, credible signal separates types [6]. Verifiable evidence of governance is that signal. Self-reported governance is not, because it is the cheap talk the market already discounts. A signal separates only when it costs the badly governed more to produce than the well-governed, and it is worth sending only when its cost falls below the ambiguity loading it removes. Verifiable evidence drawn from governance the insured is already conducting meets both tests: it is cheap for the firm that has the governance to evidence and dear for the one that does not, which is precisely why a bare assertion, equally cheap to both, cannot.

The cost asymmetry has a second, blunter form, and it is legal rather than economic. A continuous, sealed record of governance state is symmetric evidence: discoverable in litigation, producible to a regulator, and material to the duty of fair presentation the Insurance Act 2015 places on a commercial insured at placement [7]. For the firm whose governance is sound, that record is an asset at every one of those moments, the claim included. For the firm whose governance is not, the same record is a standing exhibit against it, and the rational bad risk declines to hold one. The evidence cuts both ways by design. That is not a defect to be managed but the separation itself: only the party confident in its own conduct will carry a record that can be read against it, which is precisely the costly signal Spence's condition demands and a bare assertion can never be.

Separation answers adverse selection, the concealment of state before the contract is written. Continuous evidence reaches the second information problem as well, the one that begins after it is signed. Bengt Holmström's informativeness principle holds that any observable signal carrying information about the agent's conduct improves the contract that can be written on it [8]. A governance record produced continuously through the policy period is exactly such a signal: the insured that knows its state is legible behaves as if observed, because it is, and the moral hazard loading falls with the asymmetry that fed it. What the economics calls incentive-compatible monitoring, the practitioner will recognise as the oldest rule in the market: the surveyed risk keeps its sprinklers maintained.

This is the crux, and it is why the word verifiable does more work than the word evidence. Governance has never lacked evidence; it has lacked evidence a counterparty can verify without taking the insured's word for it. Evidence that resolves the asymmetry is evidence the insurer can check independently. Everything else reverts to the pooled price.

Capital: the part that compounds

The uncertainty loading is the visible half of the price of not knowing. The cost of capital is the larger, quieter half. An insurer must hold capital against unexpected loss, and under modern solvency regimes the amount is tied directly to the uncertainty of the outcome. Solvency II calibrates an insurer's capital requirement to the value at risk of its own funds over one year at a 99.5 per cent confidence level, the loss it should survive in all but a one-in-two-hundred year [9]. That number rises with the volatility of the loss distribution and with the parameter uncertainty around it. A risk whose distribution is poorly characterised carries not only a higher ambiguity loading in its premium but a higher capital charge behind it, and capital is not free: it must earn a return, and that return is a cost the premium has to recover.

Reduce the epistemic uncertainty and both move, though not on the same clock. The premium's uncertainty loading falls wherever the insurer prices the risk on its own analysis. The capital effect is real but conditional: it flows through in full for insurers whose capital models can reflect the reduction, through an internal model or undertaking-specific parameters, and only weakly for those on the prescribed standard formula, whose calibration is fixed regardless of any single insured's evidence. The condition points at where the effect lands first. At Lloyd's, syndicate capital is set on full internal models reviewed by the Corporation and held to Solvency UK tests and standards, so the market that writes governance-adjacent risk most acutely is also the market whose capital machinery can reflect the evidence in full [10]. Where it does flow through, the capital held against the risk falls, and with it the cost-of-capital component of the price. This is why the effect of evidence is larger than a line-by-line reading of a single premium suggests. Lower uncertainty on one risk lowers the capital against one policy; lower uncertainty across a book lowers the capital against the whole portfolio and the reinsurance that sits above it. The saving compounds where the capital is held, not only where the premium is quoted. The deeper proposition is therefore not merely a better price but less capital standing behind the same risk, which is a different and larger claim.

The claim, and the dispute

Evidence acts at one more point, the moment of loss. A policy responds, or does not, and the argument about whether it should is expensive: loss-adjustment expense, subrogation, litigation, and the retrospective reconstruction of what the insured actually did. A contemporaneous, verifiable record of the insured's governance state at the relevant time materially narrows much of that argument, because the facts of what was in place, and when, are not themselves in dispute even where causation still is. It also changes the timing of the whole exercise. Today the governance signal is priced retrospectively, after a failure, at full cost. A record that is legible before the event lets the same signal be priced prospectively, at renewal, refinancing or exit, where the insured still has the time and the agency to act on it. Moving the signal from after the loss to before it is worth more than the adjustment saving alone.

What evidence does not do

The honest core of the argument is a limit. Evidence does not change the true risk of a well-run or a badly run insured. It changes what is known about it, and it moves the price toward the truth. For the genuinely well-governed insured, that is good news: the price falls, because the premium it was paying for the insurer's ignorance and for the pool's bad risks falls away. For the badly governed insured, the same mechanism works against it: the evidence reveals the risk, and the price rises or the cover is declined. Evidence is not a discount. It is accuracy, and accuracy is what the well-governed party is buying. Nor does evidence repeal the underwriting cycle. It acts on the risk-specific component of the price, the part an insurer sets on the merits of the individual risk; in a hard market, where capacity is short and capital dear, the gain may show as an advantage over an unevidenced peer rather than as an absolute fall. Evidence changes the insurer's view of the risk, not the market's price of bearing it. A market in which governance is priced on evidence rewards the organisations that

can prove their governance and penalises those that only assert it, which is the outcome a well-functioning insurance market is supposed to produce and, under asymmetric information, currently cannot.

Why continuous, and why now

Two conditions turn evidence from an idea into a working input. It must be current, because a governance record taken once a year is, at the moment a price is set, already an assertion about a state that has since moved; stale evidence reverts to inference. And it must be independently verifiable, because evidence the insurer cannot check reverts to cheap talk. Only evidence that is both continuous and independently verifiable resolves both the staleness and the asymmetry at once.

The reason this has become pressing rather than merely interesting is that the cadence of the underlying risk has changed. Governance was always produced periodically, on the audit and attestation cycle, and that cadence was tolerable when the environment moved slowly. It no longer does. Where a corporation's exposure changes continuously, a periodic record is stale on arrival, and the gap between what the insured can assert and what the insurer can verify widens rather than closes. The instrument the market is missing is not more data. It is a way to make the present state of governance continuously and verifiably legible, so that the price can be set on what is known rather than on what is feared.

The precedent

The market has priced verified state before, and the practice is older than the theory. Engineering-surveyed property underwriting has done it since the factory mutuals of the nineteenth century: Zachariah Allen fireproofed his Rhode Island mill in 1835, was refused a premium reduction, and founded the mutual that priced surveyed state instead, and the highly protected risk category his successors built still works the same way, an independent surveyor verifying construction, sprinklers, separation and process controls, the premium conditioned on the surveyed state rather than the insured's assertion, HPR status lost and repriced when the verified controls lapse.

Hartford Steam Boiler, chartered in 1866 out of Hartford's Polytechnic Club after a decade in which American boilers exploded roughly once every four days, inverted the relationship entirely: the inspection was the product and the policy was its guarantee. The founders' stated premise was that inspection would make boilers safe, and that insurance would function as the incentive to inspect and the warrant of the inspection's quality, safety first and indemnity second, and the model proved durable enough that HSB's engineers still perform the jurisdictional inspections statute requires today [11]. Governability as the precondition of insurability is not a thesis awaiting proof. It has been a charter premise since 1866.

And trade credit insurance runs the continuous version at market scale now. The credit insurer holds live financial evidence on the buyers its insured sells to, sets a limit against each, and moves the limit when the evidence moves, cover expanding and contracting through the policy period rather than waiting for renewal; where limits are written non-cancellable, the same monitored evidence sets the terms at each renewal instead. Either way the structure is the one this paper argues for: a counterparty's state, continuously observed, continuously priced. What the whole precedent line lacked was portability. The factory mutual's survey, HSB's inspection report and the credit insurer's buyer file are each proprietary to the insurer that produced them, priced once and read by no one else, so the same state is re-underwritten, at cost, by every counterparty that needs a view of it. The governance signal keeps the underwriting behaviour and removes the enclosure: one verified record, read by every counterparty entitled to it, at once.

The recent precedents carry the rest. Motor insurance repriced itself around telematics: a continuously observed driving state, verified by the insurer rather than asserted by the insured, moved individual premiums off the pooled proxies of age and postcode and onto evidenced conduct, and it moved them in both directions, which is the separation result working in the open. Cyber underwriting already consumes externally scanned security ratings between renewals, and its experience carries the design lesson rather than the design. Those ratings compress a scanned surface into an opaque score whose method the insured cannot interrogate and whose validity the underwriter must take on trust, which is the defect the credit-rating precedent already taught. The governance signal keeps the

first property and refuses the second: continuous verified state, no score, so that the evidence is the product and the pricing judgment stays where it belongs, with the underwriter.

The consequence

Put together, the argument is simple. A premium prices uncertainty as much as loss; a large part of that uncertainty is epistemic and therefore reducible; verifiable evidence reduces it; and the reduction flows to the insured through a lower uncertainty loading, to the insurer through less capital held against the risk, and to both through a cheaper and faster claim. None of it changes the true risk. It changes what is known about the risk, and in insurance, what is known about the risk is what is priced. Insurance has always been the price of uncertainty. Evidence is what lowers it, for the party that can prove its case. One honesty remains owed. The precedents establish that this market pays for verified state; what it pays for verified governance state specifically has not yet been demonstrated, and the claim is held here as a hypothesis with its experiment named. The natural first demonstration is a carrier pilot, where the premium is explicit, the counterfactual is quotable and the answer arrives as a number.

References

1. **Frank H. Knight, Risk, Uncertainty and Profit, 1921.** Draws the foundational distinction between risk, an outcome drawn from a knowable probability distribution, and uncertainty, a situation whose distribution cannot be known. The basis for treating the reducible, epistemic part of an insurer's uncertainty separately from the irreducible, process part. Verified.
2. **Daniel Ellsberg, "Risk, Ambiguity, and the Savage Axioms," Quarterly Journal of Economics, 75(4), 643-669, 1961.** Demonstrates that decision-makers are averse to ambiguity, uncertainty without a defined probability measure, over and above their aversion to measurable risk. The mechanism behind the ambiguity premium an underwriter loads onto a poorly characterised exposure. Verified.
3. **Howard Kunreuther, Robin Hogarth and Jacqueline Meszaros, "Insurer Ambiguity and Market Failure," Journal of Risk and Uncertainty, 7(1), 71-87, 1993.** Survey evidence from actuaries, underwriters and reinsurers: all three groups price ambiguous risks materially above equivalent well-specified ones, and insurer ambiguity aversion is identified as a contributor to market failure at the industry level. The desk-level, empirical confirmation of the Ellsberg mechanism, and the measured counterpart to this paper's central claim. Verified.
4. **George A. Akerlof, "The Market for 'Lemons'," Quarterly Journal of Economics, 84(3), 488-500, 1970.** Shows that when quality cannot be observed, the market prices the average, good quality withdraws, and the market can unravel. With Spence and Stiglitz, awarded the 2001 Nobel Memorial Prize for the analysis of markets with asymmetric information. Verified.
5. **Michael Rothschild and Joseph Stiglitz, "Equilibrium in Competitive Insurance Markets," Quarterly Journal of Economics, 90(4), 629-649, 1976.** Carries asymmetric information directly into insurance: a single pooled price cannot survive, and insurers must let insureds sort themselves by offering terms that separate the risks. The formal basis for the separation argument. Verified.
6. **A. Michael Spence, "Job Market Signaling," Quarterly Journal of Economics, 87(3), 355-374, 1973.** Establishes that only a costly, credible signal, one a low type cannot cheaply imitate, separates types under asymmetric information. The reason self-reported governance does not resolve the asymmetry and independently verifiable evidence does. Verified.
7. **Insurance Act 2015, Part 2, duty of fair presentation.** Requires a commercial insured to disclose every material circumstance it knows or ought to know, or failing that to give the insurer sufficient information to put a prudent insurer on notice of the need to make further enquiries. The statutory frame under which a contemporaneous governance record is an asset to the sound insured and an exposure to the unsound one: the legal half of the separating cost. Verified against the Act as enacted.
8. **Bengt Holmström, "Moral Hazard and Observability," Bell Journal of Economics, 10(1), 74-91, 1979.** The informativeness principle: any observable signal that carries information about the agent's action improves the contract that can be written on the outcome alone. Extends the argument from adverse selection to moral hazard: a continuous governance record is contractible information

about conduct during the policy period, so the observed insured behaves and the moral-hazard loading falls with the asymmetry that fed it. Verified.

9. Solvency II Directive, Solvency Capital Requirement. Sets an insurer's capital requirement as the one-year value at risk of its own funds at a 99.5 per cent confidence level, the one-in-two-hundred-year loss. Capital scales with the volatility and the parameter uncertainty of the loss distribution, which is why reducing epistemic uncertainty reduces capital held, not only premium charged. Verified against EIOPA's definition.

10. Lloyd's Capital Guidance, Society of Lloyd's, current edition. Requires managing agents to calculate each syndicate's one-year and ultimate SCR using a full internal model, reviewed by Lloyd's and compliant with Solvency UK tests and standards; new entrants transition via the Lloyd's Standard Model while their model is built. The basis for the claim that the capital channel of the evidence effect flows in full at Lloyd's. Verified against Lloyd's published guidance.

11. The Hartford Steam Boiler Inspection and Insurance Company, chartered 1866; the Polytechnic Club of Hartford, 1857. Founded on the stated premise that inspection would improve boiler safety and that insurance would function as the incentive to inspect and the guarantee of a quality inspection, in a decade when American boilers exploded roughly once every four days. HSB inspectors hold National Board commissions and perform jurisdictionally mandated inspections today. The 1866 charter form of the proposition that governability is the precondition of insurability. Verified against HSB and Munich Re corporate histories and the US Congressional Record, 2016.

Production Note

Standard: Arkaya Production Standard v1.

Drafting tools: Claude (Anthropic). Research synthesis assisted; all assertions verified against primary sources.

Verification: Primary sources named in the numbered references; each reference carries substantive content. Citation audit complete pre-delivery.

Human approver: David J McKibbin, Founder, Arkaya Risk.

Version: v5. Approved for library publication by the founder, 3 July 2026.

AKR AIPS · v5