



ARKAYA RISK
Building resilience capital

Why governance becomes a financial signal

*The economics of verifiable evidence as a signal priced across insurance, credit, trade finance and diligence. v6
· 1 July 2026 · Strategic paper · shareable.*

The claim

Insurance is the market where the price of not knowing is easiest to see, because an insurer prices uncertainty explicitly and holds regulatory capital against it. But the mechanism is not peculiar to insurance. Every counterparty that extends capital to another, or accepts exposure to another's conduct, prices what it cannot verify about that party's state, and holds something back against the part it cannot see. A companion paper works the mechanism through in one market, insurance, to the depth an underwriter tests. This paper draws the wider consequence: the same evidence is priced not once but in every market that must form a view of the same risk, and that is what turns governance from an internal cost into a signal the market values.

Governance here means the observable operational state of how an organisation runs itself: the controls, the conduct and the decisions that determine its risk, not the board-level machinery the word sometimes denotes.

One structure, many markets

Strip the markets back and the structure underneath is identical. Wherever one party bears the consequences of another's conduct, it faces two kinds of uncertainty that behave very differently. Some of it is irreducible: the genuine randomness of outcomes drawn from a distribution, which no information can shrink. The rest is epistemic, uncertainty about the counterparty's true state, which information can shrink. Frank Knight drew the line between the two in 1921, and it is the line that matters commercially, because only the second kind responds to evidence and only the second kind need be there at all [1].

On top of this sits a second fact, equally general: the party being financed knows more about itself than the party financing it. This is asymmetric information, and it is the condition that turns uncertainty into mispricing. George Akerlof showed in 1970 that when quality cannot be observed the market prices the average, the good is underpaid and withdraws, and the market erodes [2]. The rational responses to this are the same in every market. Load the price for the uncertainty. Restrict the terms. Withhold capital rather than misprice it. Or demand a signal the counterparty cannot fake and the financier can check, which Michael Spence showed in 1973 is the only thing that separates a good type from a bad one that claims to be good [3]. Insurance, lending, trade finance, acquisition and rating are this one problem wearing different clothes. What follows is the same signal, verifiable governance evidence, doing the same work in each.

The whole argument can be assembled from the results just cited, and the derivation is short enough to state in four steps. One: wherever a party bears the consequences of another's conduct, part of its uncertainty is epistemic, uncertainty about the counterparty's true state rather than the randomness of outcomes [1]. Two: that epistemic component is priced, in the loading, the spread, the discount or the capital withheld, because a financier that cannot verify prices the average or declines to lend at all [2]. Three: a credible signal, costly to the bad type and checkable by the financier, reduces the epistemic component where a bare assertion cannot [3]. Four: therefore, wherever such a signal exists and can be read, the price that contained the uncertainty it removes will tend to move. Each step is an established result; only the conjunction is new, and it is the conjunction the markets below instantiate.

Stated in full: **The Financial Signal Principle:** wherever capital is allocated under asymmetric information, independently verifiable evidence that reduces epistemic uncertainty will tend to be reflected in the price, in the capital held against the exposure, or in the terms of the deal, to the extent that the market can observe, trust and incorporate it. On this reading the insurance premium, the credit spread, the capital charge, the supplier's terms, the acquisition discount and the rating are not six phenomena but one, seen in six markets. It is a synthesis of results already established, not a new theorem, and it earns its authority from them rather than from itself.

Hedged that way, the Principle risks explaining everything and forbidding nothing, so its falsifiable form should be stated with it. The testable claim is a comparative static: within a given market, a counterparty that produces independently verifiable evidence of its governance state obtains measurably better terms than an otherwise identical counterparty that does not. That is a claim the data can refute, market by market, and the argument stakes itself on it. For governance evidence specifically the test has not yet been run, and the argument holds it as an experiment with the result pending; the natural first laboratory is a carrier pilot, where the premium is explicit, the counterfactual is quotable and the result is a number rather than an argument.

The Principle also predicts its own null results, and stating them guards it against explaining everything. The signal is not priced where the counterparty's state is not what is being priced: fully collateralised or asset-based lending, where the security rather than the borrower carries the credit; commoditised personal lines pooled on statutory rating factors, where individual evidence is regulated or averaged away; and guaranteed or sovereign-wrapped exposure, where the guarantor's covenant displaces the obligor's conduct. Where the epistemic component of the price is small, the signal has little to act on, and a market that shows no effect there confirms the mechanism rather than refuting it.

Credit: the spread and the capital behind it

A lender prices a loan on the chance the borrower does not repay, and that chance is a function of how uncertain the borrower's condition is. Robert Merton's 1974 model made the link exact: the spread on a firm's debt rises with the volatility of its asset value and with its leverage, because both widen the range of outcomes in which the debt is not made whole [4]. The more uncertain the borrower's state, the wider the spread, for the same reason an ambiguous risk carries a heavier insurance loading.

Darrell Duffie and David Lando sharpened the point in 2001 in the way that matters most here. When the lender cannot observe the borrower's condition directly and receives only periodic, imperfect accounting information, the credit spread does not fall to zero even at short horizons; the incompleteness of the information is itself priced, and better information compresses the spread [5]. This is the credit market stating, in its own formal language, that the quality of what can be verified about a borrower is a price, not a courtesy. At the limit the lender stops pricing the uncertainty and simply withholds: Joseph Stiglitz and Andrew Weiss showed in 1981 that under imperfect information credit rationing is an equilibrium, the lender lending less than the borrower will pay for at any rate, because raising the rate worsens the pool it attracts [6]. Rationing is the lending analogue of the insurance exclusion, and it has the same cause.

The effect reaches the balance sheet as well as the price. Under the internal-ratings approach to bank capital, a lender holds capital against its own estimates of default and loss, and must add a margin of conservatism that grows precisely as the data and the method weaken, the supervisory expression of the same idea, that unresolved estimation and model risk must itself be capitalised [7]. The same epistemic uncertainty that widens the spread thickens the capital held behind the loan. This is the insurance result reproduced in another market with the capital effect intact: reduce what the lender cannot verify, and both the spread and the capital move.

Capital-raising shows the signal directly. Hayne Leland and David Pyle showed in 1977 that an entrepreneur who knows more about a venture than the market can credibly signal its quality only by doing something a weaker type would not, such as retaining a large stake, a costly signal in Spence's exact sense [8]. Verifiable governance evidence is a cheaper and more general instance of the same move: a credible signal of the state of the thing being financed.

Trade and supplier finance

The pattern repeats wherever performance is extended on credit. The rate at which a receivable is discounted, or a supplier is granted terms, embeds the financier's uncertainty about whether the counterparty will keep performing and keep conducting itself soundly. In a long supply chain the problem compounds: a buyer's own exposure runs through suppliers it cannot see, whose failures of governance become its disruption and its liability. Evidence of a counterparty's governance state that the financier or the buyer can verify without taking the counterparty's word for it lowers the discount and widens the terms on offer, for the same reason it lowers a premium and compresses a spread.

The deal: diligence, price and the holdback

Acquisition is diligence under a deadline. A buyer prices a target on an estimate of what diligence cannot fully close out, and protects the unresolved remainder with structure: escrow, holdback, indemnities, representations-and-warranties cover. Each of these is a price paid for residual uncertainty about the target's true state, its liabilities, its compliance, its conduct before the buyer arrived. It is the Duffie and Lando result in a private-market setting: better information, a smaller spread, here expressed as a smaller discount and a lighter retention. Contemporaneous, verifiable governance evidence narrows the residual that the escrow exists to cover, at least the part of it that turns on the target's present and continuing conduct, as distinct from undiscovered historical liabilities, which no current-state evidence can reach. A narrower residual is worth real money at the point of sale, where it is priced prospectively rather than litigated later.

The intermediary: ratings and the limits of periodic assertion

Rating agencies and auditors exist to do exactly what separation requires, resolve asymmetry at scale and produce a signal that many counterparties can rely on at once. Their difficulty is structural. The signal is periodic, refreshed on a cycle, and in the dominant model it is paid for by the party being rated, which sits awkwardly against Spence's condition that a credible signal be costly to the bad type rather than merely issued to it. Continuous, independently verifiable evidence does not displace the intermediary. It changes what the intermediary certifies, from an assertion about a state at the last review to a state that can be checked between reviews, and it narrows the window in which a rating can be current on paper and stale in fact. It does not follow that the periodic, issuer-paid model is safe: a signal the market can check for itself, continuously and independently, puts competitive pressure on any certificate that is refreshed slowly and paid for by its subject, even where it stops short of replacing it.

The penalty on silence

The Principle claims the price falls for the party that proves its state. The harder half of the claim is that the price rises for the party that does not, and the mechanism deserves to be stated rather than assumed, because the signal is voluntary and no badly governed firm will volunteer its own exposure. The mechanism is inference from silence. Sanford Grossman and Paul Milgrom showed, in the same year, that where a party can credibly disclose its quality at low cost, disclosure unravels from the top: the best discloses to escape the average, the next-best discloses to escape the residual pool, and silence is ultimately read as the worst case, because only the worst has nothing to gain from proof [9] [10]. Once verifiable governance evidence exists and is cheap for the sound firm to produce, not producing it becomes informative, and the badly governed firm comes to be priced on the inference its silence invites.

The theory also states the limit honestly, and the limit is the adoption path. Unravelling completes only where disclosure is near-costless and the market knows the firm could have disclosed. Where producing the evidence still carries real cost, or a counterparty cannot yet tell inability from unwillingness, silence stays ambiguous and the penalty is blunted. Which is why the pricing of silence arrives with adoption rather than before it: in the early market the evidenced firm earns a premium for proof, and only as production becomes standard does the unevidenced firm begin to pay a discount for its absence. The reward leads; the penalty compounds behind it.

Why the signal is not already there

If the price of governance is latent in all these markets, the obvious question is why they do not price it already. The answer is that until recently they could not, because the evidence did not exist in a form a market can use. Governance has been recorded as attestation rather than proof, produced to a periodic cycle rather than continuously, bespoke to each firm rather than comparable across firms, and locked inside the organisation rather than portable to a counterparty who could check it. Evidence with those properties is the cheap talk each market already discounts, so each market did the rational thing and priced the uncertainty instead. What was missing was never the governance. It was a way to make the governance state standardised enough to compare, verifiable enough to trust, and current enough to rely on. That is a demanding set of conditions, and meeting it is the real work, not a detail to be waved through. But the barrier has been the form of the evidence, not the economics, and the economics is what has been in place since the results cited here were written.

Why it compounds: the non-rival signal

Here is the property that makes the wider claim more than a list of markets. Governance produced once and made verifiable is consumed by many counterparties at the same time. The insurer setting a premium, the bank setting a spread, the acquirer setting a price, the agency setting a rating and the large customer running supplier diligence are all forming a view of the same underlying risk, and a single verifiable signal answers all of them. The cost of producing it is borne once; the benefit is priced in every market that must form that view. A signal with this property is not a feature of one transaction. It is infrastructure, in the precise sense that its value rises with the number of parties who can draw on it without using it up. The insurance precedents show the inefficiency this removes: the factory mutual's survey, Hartford Steam Boiler's inspection and the trade credit insurer's buyer file each priced verified state, but each was proprietary to the insurer that produced it, so the same state was re-underwritten, at cost, by every counterparty that needed a view of it. The signal is those precedents with the enclosure removed: produced once, verified independently, and read by every entitled counterparty at once; the companion insurance paper carries the precedent line in full. This is what converts governance from an unpriced internal process, a cost that earns nothing outside the firm, into a produced asset with a price in several markets simultaneously.

The markets do not take the signal up at the same pace, and the value is realised over a sequence rather than all at once. Insurance is the natural first consumer, because it prices uncertainty explicitly and is required to hold capital against it, so it has both the machinery and the regulatory reason to reflect the signal soonest. Credit follows, where the pricing of information is already formalised in the spread and in the capital rules. The less standardised markets, trade finance among them, reprice last, as the signal proves portable and the means to read it matures. The non-rival value is real, but it accrues along that path, not on the first day, and the penalty on silence accrues along the same path behind it.

What this does not do

The honest core carries across from the narrower argument. Verifiable evidence does not change the counterparty's true creditworthiness, solvency or conduct. It changes what others can verify about them, and it moves every price toward the truth, which is good news for the well-governed and bad news for the badly governed in credit exactly as in insurance. It does not repeal the credit cycle or the liquidity cycle any more than it repeals the underwriting cycle; it acts on the counterparty-specific component of the price, not the market's price of bearing risk at large. Nor is it a measure of the tail. It reduces what a counterparty can conceal, not the irreducible uncertainty of extreme outcomes, and a signal of this kind becomes dangerous the moment it is trusted as if it did more, because a risk measure over-trusted is a cause of fragility rather than a guard against it, as the pre-2008 faith in credit ratings showed. It lowers the price of the knowable; it does not abolish the unknowable. It works only where the evidence is standardised enough to be compared and verifiable enough to be trusted, failing which it reverts to the cheap talk each of these markets already discounts. And markets price it only to the extent they can observe and act on it: how far it moves any given price depends on that market's maturity, its regulation and its capacity to take the signal in, and in some it is barely priced today. And in every market the signal earns its price only where it is cheaper for the well-governed to produce than for the badly governed, the Spence and Leland-Pyle condition, which evidence drawn from governance a sound firm is already conducting happens to satisfy and a bare assertion never

can. The asymmetry is partly legal: a sealed, contemporaneous record is discoverable in litigation and producible to a regulator, an asset to the firm whose conduct it evidences and a standing exhibit against the firm whose conduct it would betray, which is a cost no bare assertion carries and no badly governed firm will bear.

The consequence

Insurance is where the mechanism is clearest, because insurers price uncertainty openly and are made to hold capital against it. But pricing what cannot be verified, and holding capital against it, is what every financier does. A lender does it in the spread and the margin of conservatism, an acquirer in the discount and the escrow, a rating agency in the notch, a buyer in the terms it grants a supplier. Verifiable governance evidence reduces the epistemic component of all of them, because it is the same uncertainty being priced in each. The party that can prove its state does not merely buy cheaper insurance. It borrows, sells, raises and is rated on better terms, on the strength of one signal produced once. Governance, made verifiable, stops being a cost the firm carries and becomes a credential the market prices.

References

1. **Frank H. Knight, Risk, Uncertainty and Profit, 1921.** Separates risk, an outcome drawn from a knowable distribution, from uncertainty, a situation whose distribution cannot be known. The spine of the argument: only the epistemic, reducible part of a financier's uncertainty responds to evidence, and it recurs in every market that prices a counterparty. Verified.
2. **George A. Akerlof, "The Market for 'Lemons'," Quarterly Journal of Economics, 84(3), 488-500, 1970.** When quality is unobservable the market prices the average, good quality withdraws, and the market can unravel. The general form of the problem that insurance, credit, trade finance and diligence each face in their own terms. Awarded, with Spence and Stiglitz, the 2001 Nobel Memorial Prize. Verified.
3. **A. Michael Spence, "Job Market Signaling," Quarterly Journal of Economics, 87(3), 355-374, 1973.** Only a signal that is costly to the bad type, and cheaper for the good type, separates types under asymmetric information. The reason a self-issued or issuer-paid assertion resolves nothing and independently verifiable evidence does, in every market examined here. Verified.
4. **Robert C. Merton, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," Journal of Finance, 29(2), 449-470, 1974.** The prototype structural model of default: a firm's credit spread rises with the volatility of its asset value and its leverage. Ties the price of credit directly to the uncertainty of the borrower's state, the same quantity evidence reduces. Verified.
5. **Darrell Duffie and David Lando, "Term Structures of Credit Spreads with Incomplete Accounting Information," Econometrica, 69(3), 633-664, 2001.** When investors observe only periodic, imperfect accounting information, the credit spread does not vanish even at short horizons; the incompleteness of information is itself priced, and better information compresses the spread. The formal statement that information quality is a price in credit markets. Verified.
6. **Joseph E. Stiglitz and Andrew Weiss, "Credit Rationing in Markets with Imperfect Information," American Economic Review, 71(3), 393-410, 1981.** Shows credit rationing is an equilibrium under imperfect information: lenders withhold rather than raise rates, because a higher rate worsens the applicant pool. The lending analogue of the insurance exclusion, with the same informational cause. Verified.
7. **Basel internal-ratings-based approach, margin of conservatism (BCBS, 2006; EBA Guidelines on PD and LGD estimation, effective 2021).** Requires a bank to add a margin of conservatism to its risk-parameter estimates that grows as the data and method weaken, so that unresolved estimation and model risk is itself capitalised. The banking twin of the Solvency II capital effect: estimation uncertainty is capitalised, so reducing it releases capital, not only price. Verified against BCBS and EBA texts.
8. **Hayne E. Leland and David H. Pyle, "Informational Asymmetries, Financial Structure, and Financial Intermediation," Journal of Finance, 32(2), 371-387, 1977.** Carries costly signalling into capital markets: an informed entrepreneur signals quality by taking a stake a weaker type would not.

Establishes that the signalling logic behind the governance argument is already how capital markets resolve asymmetry. Verified.

9. **Sanford J. Grossman, “The Informational Role of Warranties and Private Disclosure about Product Quality,” *Journal of Law and Economics*, 24(3), 461-483, 1981.** Where quality can be credibly disclosed at low cost, sellers disclose and silence is read as the worst case. One half of the unravelling result: the mechanism by which the unevidenced party comes to be priced on the inference its silence invites. Verified.

10. **Paul R. Milgrom, “Good News and Bad News: Representation Theorems and Applications,” *Bell Journal of Economics*, 12(2), 380-391, 1981.** The companion unravelling result: a sender who can prove favourable states reveals them, and concealment is penalised recursively until all states stand revealed. Dye (*Journal of Accounting Research*, 1985) supplies the limit relied on here: where the market cannot tell inability to disclose from unwillingness, unravelling is partial, which is why the penalty on silence arrives with adoption rather than before it. Verified.

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: v6. Approved for library publication by the founder, 3 July 2026.

AKR AIPS · v6