The Cantillon Institute Working Paper Series: Insurance and the Digital Asset Gap Working Paper No. 1 A.K. Menger 2026

The Valuation Problem

Insurance pricing has not caught up to tokenised asset classes. This paper identifies the liability exposure and maps the framework required to close it.

Abstract

The global insurance and reinsurance industry is accumulating tokenised asset exposure it cannot accurately price. Tokenised real-world assets held across institutional blockchain platforms exceeded $15 billion in total value as of December 2024, with Boston Consulting Group projecting market growth to $16 trillion by 2030. Against this trajectory, the actuarial infrastructure required to underwrite these instruments credibly does not exist. No industry-wide loss database covers tokenised asset classes with actuarially meaningful depth. Policy language drafted for traditional property, financial institution, and cargo coverage is being applied without material amendment to instruments with fundamentally different risk architectures. The loss vectors specific to tokenised assets, including smart contract vulnerability, oracle price manipulation, custodial key compromise, and cross-chain bridge failure, have no established loss distribution equivalents in the actuarial literature. This paper identifies three structural dimensions of the valuation problem: the temporal mismatch between the value of an asset at underwriting and its value at time of loss; the legal ambiguity separating a tokenised instrument from its underlying asset; and the absence of a standardised loss taxonomy for digital asset risk events. The aggregate underwriting exposure arising from this gap is not prospective. It is

already present in the policy books of firms that have not recognised it. This paper proposes a minimum-framework approach to actuarial standards that would permit credible initial pricing, and argues that the window for establishing those standards on industry terms is narrowing.

Keywords: tokenised assets, insurance pricing, actuarial modelling, real-world asset tokenisation, smart contract risk, liability exposure, reinsurance

---

1. Introduction

The gap between what the insurance industry is covering and what it is capable of pricing is not new. It widened materially when cyber risk emerged as an insurable class in the late 1990s; the industry spent a decade writing cyber coverage on terms designed for property and casualty, absorbing losses it had no framework to anticipate, before a credible pricing infrastructure began to develop. The tokenised asset problem follows the same structural pattern, at a larger scale and with faster underlying growth.

Tokenised real-world assets, meaning blockchain-registered representations of physical or financial assets including real estate, private credit instruments, infrastructure equity, commodities, and Treasury securities, are not a speculative category. As of December 2024, on-chain tokenised real-world assets excluding stablecoins exceeded $15 billion in total value, a figure that represents organic institutional adoption rather than retail speculation (RWA.xyz, 2025). The growth rate from January 2023 to December 2024 was approximately 640 percent. Boston Consulting Group's projection of $16 trillion in tokenised asset value by 2030 implies a compound annual growth rate of approximately 66 percent from the 2022 base at which that projection was calibrated; given the current institutional pipeline, that trajectory is conservative in the medium term (Amendola & Schneider, 2022).

Against this, the insurance industry's preparation is close to absent. There is no standardised insurance product for tokenised asset risk. There is no industry-agreed taxonomy of digital asset loss events. There is no actuarial loss database with the historical depth required to construct credible loss distributions. What exists is a patchwork of cyber coverage, crime coverage, and professional liability products, each written for different

risk architectures, being applied at discretion by underwriters who do not share a common framework for what they are actually covering.

This paper does not characterise that situation as a scandal. It characterises it as a pricing problem with a calculable cost; and the cost is already accumulating.

---

2. What Actuarial Pricing Actually Requires

Sound actuarial pricing requires three foundational inputs: a defined loss taxonomy, a credible historical loss database with sufficient depth to estimate loss distributions, and a stable relationship between the insured value at underwriting and the recoverable value at loss. Tokenised assets fail, with varying degrees of severity, on all three.

Loss Taxonomy

Traditional property insurance pricing is built on a loss taxonomy developed over centuries. Fire, flood, windstorm, earthquake, theft; each peril has a defined trigger, a measurable frequency, and a severity distribution calibrated from loss experience spanning multiple market cycles. The correlation structure between perils is understood. Reinsurance treaties can be priced accordingly. Probable maximum loss estimates can be validated against historical data.

The loss events specific to tokenised assets do not map onto this taxonomy. Smart contract vulnerabilities are not analogous to any recognised property or casualty peril. A vulnerability in a smart contract governing a tokenised real estate position can result in total loss of the digital position without any change to the underlying physical asset; the legal relationship between those two loss outcomes is untested in most jurisdictions. Oracle manipulation, where an external data feed supplying price information to a smart contract is compromised or falsified, can trigger automated settlement at values that bear no relationship to the underlying asset. Cross-chain bridge failure, the loss vector responsible for the largest single theft events in the digital asset space, constitutes a systemic exposure with no traditional equivalent: a single protocol failure can produce correlated losses across thousands of unrelated positions simultaneously.

None of these loss types appear in standard insurance loss taxonomies. Without taxonomy, there is no pricing. Without pricing, the policy is a contingent liability without an adequate reserve.

Loss History

The Chartered Insurance Institute's guidance on minimum data requirements for credible loss distribution modelling specifies at least ten years of loss history at adequate exposure volume to permit stable frequency and severity estimates (CII, 2019). The broader reinsurance convention, reflected in Swiss Re's catastrophe modelling standards, requires minimum return-period validation against at least two major loss events within the model's calibration period (Swiss Re Institute, 2023).

Tokenised real-world assets as an asset class have less than five years of meaningful market history at institutional scale. The digital asset space more broadly, which provides the only available proxy data, has produced significant loss events; Chainalysis documented $3.8 billion in cryptocurrency hack losses in 2022 alone, including the $625 million Ronin Network bridge compromise in March 2022 and the $320 million Wormhole bridge exploit in February 2022 (Chainalysis, 2023). These events occurred in a speculative, largely unregulated market with custody and governance standards materially below what institutional tokenised asset platforms are beginning to implement.

Loss Event	Date	Documented Loss
Ronin Network bridge compromise	March 2022	$625 million
Wormhole bridge exploit	February 2022	$320 million
Nomad bridge exploit	August 2022	$190 million
Total documented 2022 hack losses	Full year	$3.8 billion

Source: Chainalysis, 2023. Note: these events occurred in markets with governance and custody standards below current institutional-grade tokenisation platforms. Direct actuarial application to

institutional RWA books requires material adjustment; the proxy data is directionally informative, not calibration-ready.

The loss history is insufficient for actuarial calibration on two counts: it is too short, and the proxy data it draws from does not match the institutional risk environment of the market now being insured. Using 2022 DeFi hack data to price a 2026 tokenised Treasury fund is an error of category, not merely of calibration.

Valuation Stability

The third requirement, a stable and recoverable insured value, is where the tokenised asset problem is most structurally distinct from the cyber analogy. Property and casualty insurance operates on the principle that the insured value at underwriting approximates the replacement cost at time of loss, with agreed valuation methodologies applied at settlement. This approximation holds for physical assets because underlying replacement cost, while variable, moves at comprehensible rates.

Tokenised assets can reprice continuously, in real time, based on factors that include but are not limited to the value of the underlying asset. A tokenised position in a real estate fund may carry an insured value at underwriting that diverges from both the underlying property value and the secondary market price of the token at time of loss. In volatile market conditions, those three values can diverge by 40 percent or more within a single trading session. No standard policy language addresses this three-way divergence.

A policy written on the tokenised market value will produce a settlement that bears no reliable relationship to the actual economic loss sustained by the insured. A policy written on the underlying asset value will leave the digital position uninsured in the event of a loss that damages the token without affecting the underlying asset. Most policies in the market today are neither; they are crime or cyber policies drafted for different asset classes and applied by analogy.

---

3. The Layered Valuation Problem

The valuation problem has three distinct layers, and each layer creates a separate failure mode in the insurance relationship. Addressing one layer without the others does not produce a workable policy framework.

Layer One: Temporal Valuation Mismatch

Insurance policies are typically written on declared values at inception, updated at renewal, and settled against values at time of loss. For most insured asset classes, the relationship between these three values is reasonably stable. For tokenised assets, it is not.

Tokenised real estate positions can be priced at a discount or premium to net asset value based on liquidity conditions in the secondary token market. A fund with $100 million in underlying real estate exposure may have a token market capitalisation of $73 million in a period of digital asset market stress; the same fund may trade at $118 million under conditions of high demand. If a loss event occurs during the discount period, the insured's economic loss is the token value, not the NAV. If the policy was written on NAV and the settlement is on NAV, the insurer has overpaid by a material amount. If the policy was written on token market value at inception and the loss occurs during a period of premium pricing, the settlement undercompensates the insured.

None of the standard adjustable-value policy mechanisms designed for equity or property exposures resolve this cleanly. The temporal mismatch is not a technical detail; it is a gap that will produce contested settlements at scale.

Layer Two: Legal Identity of the Token

The relationship between a tokenised instrument and its underlying asset is not legally settled in any major jurisdiction. In the event of a loss affecting the digital layer, whether through smart contract exploit, custody failure, or blockchain fork, does the holder of the tokenised instrument retain full legal claim to the underlying asset? In the event of a loss affecting the underlying asset, whether through property damage, corporate insolvency, or sovereign default, does the tokenised instrument holder hold a claim equivalent to that of a traditional instrument holder?

The answer in both cases is: it depends. The dependence is on jurisdiction, on the structure of the tokenisation platform, on the terms of the custodial arrangement, and on the specific nature of the loss event. That indeterminacy is an underwriting risk, because it means the insurer cannot establish in advance what it is actually covering.

The Lloyd's Market Association published guidance in 2022 acknowledging the unresolved legal status of tokenised assets as a material underwriting consideration (LMA, 2022). The guidance documented the problem. It did not resolve it. Documentation alone does not create a pricing framework.

Layer Three: Smart Contract Risk as a Distinct Peril

In a tokenised asset structure, the smart contract is not merely a record-keeping mechanism; it is the instrument of ownership, transfer, and settlement. A flaw in the contract can create total loss of the digital position irrespective of what occurs at the underlying asset level. The contract can be exploited by external actors. It can contain logic errors that produce unintended outcomes under specific market conditions. It can be rendered inoperable by changes to the underlying blockchain protocol.

These are not low-probability events. Certik documented over 1,700 significant on-chain security incidents in 2023, with total identified losses across all categories exceeding $1.8 billion for the calendar year (Certik, 2024). The loss distribution is highly skewed: the twenty largest incidents in any given year account for the majority of total losses. That distribution pattern is familiar to catastrophe reinsurers; the return-period analysis required to price it does not yet exist for smart contract risk.

Standard cyber insurance does not cover smart contract failure as a defined peril. Most cyber policy forms were drafted when smart contracts did not constitute a material coverage category. Firms that have attempted to extend coverage have done so through manuscript endorsements with language that remains contested in the absence of case law. These endorsements carry sub limits that do not reflect the actual size of tokenised positions being covered, and they contain carve-outs for losses arising from "coding errors" or "protocol vulnerabilities" that would exclude the majority of documented smart contract loss events.

The three layers interact. A tokenised asset with a compromised smart contract, in a jurisdiction that has not settled the legal relationship between the token and the underlying asset, at a time when the token market value diverges materially from NAV, will produce a coverage dispute with sufficient ambiguity to sustain parallel litigation in multiple jurisdictions. This is not a projection; it is a description of a scenario the industry will encounter as the market grows.

---

4. What Current Policies Actually Cover

The coverage available for tokenised asset risk today falls into four categories: cyber insurance, crime insurance, directors and officers liability, and professional indemnity. Each covers a different subset of the risk. None covers the full exposure. The overlaps and the gaps are both consequential.

Cyber Coverage

Cyber insurance, a market generating approximately $15 billion in gross written premium globally as of 2024, was designed to cover losses arising from the compromise of information systems (Munich Re, 2025). It covers data breach costs, business interruption arising from system failure, and, in more comprehensive forms, funds transfer fraud. It was not designed to cover the loss of a tokenised asset position following a smart contract exploit, because smart contract-mediated asset ownership did not exist as a category when standard policy forms were written.

The practical effect is that cyber underwriters are applying policy language drafted for traditional IT environments to tokenised asset loss events through claims-time interpretation. The interpretation varies by insurer, by underwriter, and by the specific facts of the loss event. There is no industry consensus on whether a smart contract exploit constitutes a "computer system compromise" under standard cyber policy language. The absence of consensus means claims are resolved through negotiation rather than contract.

Crime Coverage

Financial institution crime coverage, including the Bankers Blanket Bond and its successors, covers employee dishonesty, forgery, theft of money and securities, and

computer fraud. It does not cover loss arising from protocol-level failures in external blockchain infrastructure. The crime policy assumes that the loss event involves a human actor committing a defined criminal act; a smart contract vulnerability that is algorithmically exploited by an automated agent does not fit that structure.

Several London Market underwriters have extended crime coverage to include "virtual asset theft" through specific endorsements. These endorsements carry sub limits that do not reflect the actual exposure size of the positions being covered, and they contain carve-outs for losses arising from "coding errors" or "protocol vulnerabilities" that would exclude the majority of documented smart contract loss events.

D&O and Professional Indemnity

Directors and officers liability and professional indemnity coverage address liability arising from decisions made in the governance of the tokenised asset platform, not losses to the underlying asset position itself. These forms are relevant to the operators of tokenisation platforms; they do not address the insured interests of investors holding tokenised positions.

The aggregate effect is coverage fragmentation: multiple policy types apply partial coverage to different aspects of the same loss event, with contested overlaps and definitive gaps at the boundaries. In a complex loss scenario, the coverage litigation will be as expensive as the loss itself.

---

5. The Liability Exposure Already on the Books

The exposure described in the preceding sections is not prospective. It is already underwritten, because the growth of tokenised assets has not waited for the insurance industry to develop appropriate coverage.

As of December 2024, an estimated $15 billion in tokenised real-world assets sat on major institutional blockchain platforms, with coverage arrangements that range from partial to essentially absent (RWA.xyz, 2025). The majority of large institutional tokenised asset platforms are required by their investors, and in some cases by regulation, to carry

insurance coverage. That coverage exists; it consists of the imperfect patchwork described in the preceding section.

Platform / Instrument	Approximate AUM, December 2024
Tokenised U.S. Treasury instruments (all platforms)	$3.5 billion
BlackRock BUIDL Fund	$530 million
Franklin Templeton BENJI	$360 million
Ondo Finance USDY/OUSG (combined)	$590 million
Private credit tokenisation (Securitize, Hamilton Lane, KKR combined)	$1.1 billion

Source: RWA.xyz, December 2024. Platform concentrations represent single-event loss accumulation risk that has not been modelled as a catastrophe exposure class.

The question of whether that coverage will respond on the terms the insured intends has not been tested in litigation at scale. The testing will occur during the first major loss event that triggers claims across multiple policy types simultaneously. At $15 billion in current exposure growing at rates that imply $100 billion within three to four years, the probability of such an event occurring before the industry has established coherent coverage frameworks is not small.

The reinsurance dimension compounds the problem. Primary insurers carrying tokenised asset exposure have ceded portions of that exposure to reinsurers who, in the majority of cases, are pricing the reinsurance based on the same inadequate actuarial foundations as the primary insurer. The cedant's uncertainty is propagated through the treaty structure. A major tokenised asset loss event would produce simultaneous claims disputes at the primary and reinsurance levels, in an environment where the policy language has not been validated and the loss taxonomy has not been established.

The concentration risk warrants specific attention. The tokenised asset market, while growing, is currently dominated by a small number of platforms. A platform-level loss

event at any of the dominant concentrations would produce correlated exposure across the policy books of multiple insurers simultaneously. Catastrophe reinsurers will recognise the profile: it is a risk structure consistent with a single event generating industry-wide loss accumulation, and the industry has not modelled it as such.

Pricing this exposure requires quantitative humility. The data required for a precise aggregate loss estimate does not exist. What can be stated is that the ratio of insured value to actuarially defensible premium across the tokenised asset class is currently indeterminate; underwriters are writing a number whose denominator they cannot calculate. In any other class of business, that is the definition of inadequate pricing.

---

6. A Framework for Closure

Establishing a credible pricing framework for tokenised asset risk requires action at four levels: taxonomy, data infrastructure, policy language, and regulatory classification. Each is a precondition for the next.

Taxonomy

The first requirement is an industry-agreed taxonomy of tokenised asset loss events. This does not require consensus on pricing; it requires only consensus on which categories of event constitute a covered loss and what the trigger structure for each category is. The taxonomy should distinguish at minimum between: custody layer failure, covering key compromise and custodian insolvency; smart contract layer failure, covering exploit, logic error, and protocol upgrade failure; oracle layer failure, covering price manipulation and feed disruption; bridge and interoperability failure, covering cross-chain protocol compromise; and governance layer failure, covering unauthorised smart contract modification through governance mechanism exploitation.

Each is a distinct peril with a distinct trigger structure, a distinct probable severity distribution, and a distinct set of risk mitigation measures. Treating them as a single undifferentiated "digital asset risk" category is an error of the same order as treating fire, flood, and earthquake as a single peril.

The International Association of Insurance Supervisors published a consultation in 2024 on digital asset risk classification that did not produce a binding taxonomy (IAIS, 2024). The Lloyd's Market Association's 2022 guidance represents the most advanced public documentation of the problem. It has not been followed by a standard-form solution. Both represent analytical effort that has not been converted into operational infrastructure.

Data Infrastructure

The second requirement is a shared loss database at the industry level. Cyber insurance developed CyberAcuView as a mechanism for sharing anonymised loss data to build actuarial credibility in an emerging class. A functionally equivalent mechanism for tokenised asset losses is absent. The data required to construct such a database exists; on-chain transactions are public by design, and loss events are documented in near real time by blockchain analytics firms including Chainalysis, Elliptic, and TRM Labs.

The barrier is not data availability. It is the absence of an industry body willing to coordinate collection and actuarial analysis of that data in a form usable for pricing. Swiss Re and Munich Re, the two largest reinsurers globally by premium volume, have the analytical infrastructure to lead this effort. Neither has moved beyond preliminary research publication.

Policy Language

Standard-form policy language for tokenised asset coverage should be treated as an immediate project. The Lloyd's Lab and the International Underwriting Association have demonstrated the capacity to develop standard-form coverage for emerging risks on timescales of twelve to eighteen months. The absence of standard form for tokenised asset risk in 2026 reflects a prioritisation decision, not a capacity constraint.

The minimum required policy language development covers three elements: a defined insured event trigger that explicitly distinguishes between the digital layer and the underlying asset layer; a defined valuation mechanism that specifies which of the three value references (token market price, NAV, or underlying asset replacement cost) applies at settlement; and a defined exclusion structure that is explicit about what remains uninsured rather than leaving exclusions to claims-time interpretation.

Regulatory Classification

The regulatory dimension is outside the insurance industry's unilateral control but not outside its capacity for coordinated advocacy. The Prudential Regulation Authority in the United Kingdom, the European Insurance and Occupational Pensions Authority, and the Federal Insurance Office in the United States have not yet classified tokenised asset exposure as a distinct risk category requiring specific capital treatment under solvency frameworks.

Under Solvency II, tokenised asset exposures are currently aggregated with the broader asset classes they represent, whether real estate, equities, or fixed income, without adjustment for the additional risk layers described in this paper. That classification produces capital charges that understate the actual risk and reserves that are inadequate relative to the true liability. Regulators who receive the loss event before they receive the taxonomy will set the taxonomy on enforcement terms; that is a structurally worse outcome for the industry than establishing it on analytical terms in advance.

Conclusion

The insurance industry has a documented history of under-preparing for emerging asset class risk and absorbing the consequence in loss ratios that damage the affected lines for a decade following initial mispricing. Asbestos liability cost the global industry an estimated $70 billion or more in total incurred losses and required fifteen years of adverse development before the scope was fully understood (Swiss Re, 2002). Cyber insurance losses in the 2017 to 2020 period produced combined ratios above 100 in multiple consecutive underwriting years, driving market correction that significantly reduced coverage availability before pricing discipline was restored. The pattern is consistent enough that it does not require a theory of institutional failure to explain it; it requires only an observation about how insurance markets systematically under-price novel risk until the novel risk produces a loss large enough to reset the market.

The tokenised asset market is at an early stage of that cycle. The exposure is growing at 640 percent over two years. The actuarial infrastructure is absent. The policy language has not been tested. The regulatory classification has not been established. The window in

which the industry can establish a credible pricing framework on its own terms, before a loss event forces the lesson at greater cost, is measurable in years rather than decades.

What distinguishes the current situation from the cyber and asbestos precedents is the transparency of the problem. The data is public. The loss events have already occurred, at smaller scale, in adjacent markets. The loss vectors are identifiable in advance. The gap between what is being insured and what can be priced is visible to anyone who examines the policy forms against the risk architecture.

The cost of closing that gap on the industry's own initiative is the cost of taxonomy development, data infrastructure investment, and policy language drafting. That cost is non-trivial; it is not comparable to the cost of absorbing a systematic loss event across an inadequately priced book at the scale this market will reach. The probability of the loss event occurring before the framework is in place is already higher than any actuary pricing a known, visible, and growing exposure concentration would find professionally defensible. The subsequent papers in this series address each of the four closure requirements in turn; the present paper has established what is at stake if they are not met.

References

Amendola, S., & Schneider, M. (2022). *Relevance of on-chain asset tokenisation in 'crypto winter': Building the foundation of an asset management value chain*. Boston Consulting Group and ADDX.

Certik. (2024). *Web3 security report 2023*. Certik Security.

Chainalysis. (2023). *The Chainalysis 2023 crypto crime report*. Chainalysis Inc.

Chartered Insurance Institute. (2019). *Actuarial data requirements and standards for general insurance*. CII Technical Publication.

International Association of Insurance Supervisors. (2024). *Application paper on the supervision of insurer activities related to crypto-assets*. IAIS.

Lloyd's Market Association. (2022). *Digital assets: Considerations for the Lloyd's market*. LMA Guidance Note.

Munich Re. (2025). *Cyber insurance: Risks and trends 2024/2025*. Munich Re.

RWA.xyz. (2025). *Real world asset tokenisation market data, December 2024*. RWA.xyz Analytics.

Swiss Re. (2002). *Asbestos: The long and expensive tail of liability*. Swiss Re Sigma No. 6/2002.

Swiss Re Institute. (2023). *Natural catastrophe modelling: Standards and calibration methodology*. Swiss Re Institute Technical Paper.