The Insurance Valuation of LEMNI

An analytical note on pricing a one-time, recoverable premium — and on the arithmetic that governs it.

Scope

This note addresses one question: what is a rational price for an asset whose function is to abate the risk of catastrophic loss on other holdings, where the price is paid once rather than annually and the asset is retained rather than consumed?

The question is worth isolating because it is routinely answered badly. Applying insurance-style benchmarks without an explicit capitalization rate produces numbers that are absurd on their face; asserting that redundant signature schemes make failure risk negligible substitutes a tractable calculation for the risk that actually matters. Both errors run in the same direction, and both are avoidable.

The structure below is deliberately parametric. Three inputs — the annual premium the protection displaces, the rate at which that premium is capitalized, and the annual hazard that the protection ceases to function — determine the result completely. Each is stated, sourced where possible, and open to substitution.

The subject

LEMNI is a fixed-supply asset (10,500,000 tokens, no issuance) whose holding confers a dual-authority protection on other digital-asset holdings. Transfers of protected assets require two independent post-quantum signatures — the proprietary LAIP scheme alongside NIST-standardized Dilithium5 or Falcon-512 — so that compromise of a wallet private key alone cannot move the protected assets. Coverage is denominated in Bitcoin-value: one LEMNI confers protection equal to one Bitcoin's worth of value, with the requirement for any position given by ⌈ protected value ÷ BTC price ⌉, minimum one.

Two consequences follow from that denomination and are used throughout. Bitcoin holders are one-to-one and never re-rate: one whole Bitcoin always takes exactly one LEMNI, at $30,000 or at $300,000. Every other asset rolls with Bitcoin, since its requirement is its value divided by the Bitcoin price. And the supply imposes a hard ceiling — total coverage can never exceed 10.5 million Bitcoin-equivalents, roughly $630 billion at a $60,000 reference price, against an addressable base on covered chains on the order of $1.8 trillion.

The analysis that follows is the valuation section of the fuller LEMNICOIN valuation and market analysis, presented here as a standalone argument.

PART 6: THE INSURANCE VALUATION — A ONE-TIME, RECOVERABLE PREMIUM

LEMNI is priced the way insurance is priced. The holder is buying protection against catastrophic loss — wallet theft today, quantum exposure tomorrow — and the rational price of that protection is anchored to what the same perils cost to transfer in the open market. Two features separate it from an insurance contract, and both raise the rate a holder will rationally bear:

- **The premium is one-time, not recurring.** Payment is made once. Protection persists for as long as the token is held. There is no annual renewal and no premium bleed.
- **The premium is retained, not spent.** What the holder “pays” is converted into LEMNI itself — a fixed-supply, dual-authority, quantum-resistant asset that remains owned and can be sold. The true economic cost is the opportunity cost of the capital while held, less any change in the token's value — a figure that can be zero or negative, which is never true of a sunk premium.

The correct mental model is therefore a one-time, recoverable premium. This section derives what that premium is worth, using observable market prices for the same perils, an explicit capitalization rate, and an explicit hazard that the protection itself fails. Each of those three inputs is stated rather than assumed, so that a reader who disagrees with any of them can substitute their own and re-derive the result.

6.1 What the market actually charges for these perils

There is now a regulated market price for insuring self-custodied Bitcoin, and it is directly comparable to LEMNI's architecture. AnchorWatch, a Lloyd's of London coverholder, underwrites self-custody Bitcoin policies against theft,

loss, and catastrophic events. Its Trident vault requires two independent keysets to complete signatures before a transaction can be broadcast — structurally the same two-independent-authority requirement that LEMNI enforces cryptographically.

This matters more than any modelled estimate: a regulated underwriter has already priced a dual-signature Bitcoin custody structure, and it did not price it at zero.

Benchmark	Annual rate	Deductible	Notes
AnchorWatch insurance premium	0.40% – 0.80%	10% or 25%	Lloyd's-backed, self-custody Bitcoin; the provider's own calculator uses 0.60% for estimation. Minimum policy $250,000.
AnchorWatch all-in (premium + services fee)	0.55% – 2.00%	10% or 25%	Varies by risk profile; many customers quoted at or near 1%. Premium payable in full at policy inception.
Broader institutional crypto crime / custody cover	1.0% – 2.5%	typically 5% – 20%	Higher for hot or operationally exposed setups; lower for strong multi-signature cold storage.
Physical-attack / duress-specific cover	from ~0.55%	varies	Prices the coercion peril rather than the key-compromise peril.
DeFi smart-contract cover (mutual markets)	0.11% – 0.81%	n/a	Narrow, protocol-specific, and only for audited contracts; a floor, not a comparable.

Underwriters price security posture, not just asset value. Documented discounts are available for a multi-year record of responsible key management, for geographic separation of keys, for a low public profile, and for larger vaults. Each is a proxy for the same thing LEMNI enforces in code rather than in underwriting: the difficulty of a single compromise producing a total loss.

It is worth being explicit that the cheapest regulated Bitcoin insurance in the market is cheap in part because it already requires two independent signing authorities. That is an argument about the mechanism, not about the token, and it is the most objective evidence available that the architecture reduces loss cost.

What the benchmark premium actually buys — and does not

A stated premium cannot be lifted directly into this model, because it prices a different product against a different set of perils. Three adjustments are required, and they do not run in the same direction.

- **Deductible (raises the implied value of prevention).** The policies are not first-dollar. AnchorWatch customers select a 10% or 25% deductible, assessed against an approved claim benefit — on a $1,000,000 loss under the 25% option, $250,000 is retained by the insured. Prevention eliminates the whole loss including the retained layer, so per unit of loss actually eliminated the premium-equivalent is $a \div (1 - d)$. Grossing up consistently — lower quotes pair with the higher deductible — turns a stated 0.40%–0.80% into roughly 0.53%–0.89%.
- **Peril scope (lowers it, and by more).** The premium buys a bundle. Covered perils include loss of keys to physical events such as fire, flood and wildfire; theft including coercion and violent events; and collusion. LEMNI's dual-authority design addresses remote compromise of the wallet key, and explicitly does not address seed-phrase compromise or coercion. Only the fraction of the bundle's expected loss cost attributable to the covered channel may be claimed — written below as a scope factor s.
- **Claim contingency and structure (raises it modestly).** Indemnity pays only on an approved claim, subject to conditions and exclusions; published guidance notes that policies may exclude broad risk events including a break in Bitcoin's cryptography, which is precisely the quantum peril. Coverage is also capped at a chosen dollar limit that does not scale as Bitcoin appreciates, and is subject to annual re-underwriting and non-renewal. Prevention has no approval step, no limit, and no renewal. These effects are real but small relative to the first two, and are not separately parameterised below.

The scope factor is the load-bearing judgment in this section and it is not observable. A holder whose principal exposure is remote compromise — hot or warm keys, browser-based signing, frequent transaction activity — sits at the high end. A holder already using distributed cold multi-signature sits at the low end, because most of their residual exposure is physical and operational rather than remote. The range used here is 0.35 to 0.60, and a reader who takes a different view should substitute it directly.

Effective annual premium: $a_e^{ff} = a \times s \div (1 - d)$ — ranging from 0.19% to 0.53% of protected value.

A note on population

One further adjustment is available and is deliberately excluded from every figure in this section. The benchmark rate is charged to holders already inside a distributed multi-signature vault with an institutional co-signer; it prices residual risk after that architecture is in place. The holders LEMNI is designed for generally do not have it, and their ground-up remote-compromise hazard is correspondingly higher — plausibly by a factor of one and a half to two and a half.

Applying that factor would raise every figure below proportionally. It is not applied, for two reasons: the multiple is unmeasured, and it is an adjustment that happens to favour the issuer of this paper. It is shown as a labelled sensitivity in section 6.4 rather than folded into the headline, so that no headline figure depends on it.

6.2 From an annual premium to a one-time price

Because the protection persists as long as the token is held, its value is the capitalized value of the effective annual premium it displaces. Write A for that premium in dollars, r for the discount rate the holder applies, and f for the annual hazard that the protection ceases to function:

$$V = A \div (r + f)$$

Dividing through by the protected value P, and substituting the effective premium rate derived in 6.1, the protected value cancels entirely:

$$\text{One-time premium rate} = [\, a \times s \div (1 - d) \,] \div (r + f)$$

This connects directly to the pricing identity used throughout this paper. Because the gate requires [P ÷ BTC] tokens to protect a value P, the price per token is the one-time premium divided by the token count, and P drops out:

$$\text{LEMNI price} = [\, (a \times s \div (1 - d)) \div (r + f) \,] \times \text{BTC price}$$

The premium rate in the identity “LEMNI price = premium rate × BTC price” is therefore not a free parameter. It is the product of four stated inputs: an observable market premium a, the deductible d embedded in that premium, a scope factor s reflecting how much of the covered bundle the protection actually replaces, and a capitalization rate (r + f) that must be defended. No other parameter enters, and every one of the four is open to substitution by the reader.

The capitalization rate is where valuations of this kind go wrong

The quantity 1 ÷ (r + f) is the capitalization multiple applied to the annual premium. It is worth being explicit about what different multiples imply, because the result is far more sensitive to this input than to the premium benchmark:

Multiple	Implied r + f	What it asserts
20×	5%	That the protection is as certain and as liquid as a government bond. Not defensible for a novel, illiquid, adoption-dependent asset.
10×	10%	Investment-grade corporate risk. Still assumes near-certainty of continued function and a ready exit.
4× – 6×	17% – 25%	Venture-like discount for an unproven protective technology with thin liquidity. This is the defensible band.
2× – 3×	33% – 50%	Deep-discount case: material doubt about continued function or exit.

A 20× multiple — a 5% discount rate with no hazard term — is the single most common error in valuations of this asset. Applied to an effective 0.36% annual premium it produces a one-time rate of over 7% of protected value, which is self-refuting on its face: no rational holder pays that share of a Bitcoin's value to protect it when reducing the position accomplishes the same risk reduction for free. When a model returns an absurd number, the discount rate is almost always the reason.

Why no growth term is applied

The perpetuity above places (r + f) in the denominator and no growth term. The general form would be V = A ÷ (r + f − g), where g is the expected annual growth rate of the displaced premium in Bitcoin-denominated terms. Setting g to zero is a deliberate choice, and it is the most consequential omission in this section: the denominator shrinks as g rises, the value diverges as g approaches (r + f), and the effect lands hardest on the high end of the range, where the denominator is already smallest.

Three mechanisms could make g genuinely positive:

- **Rising peril salience.** Quantum risk advances toward Q-Day rather than receding, and wallet theft has no end-state on the horizon. If the underwriting market's assessment of either hardens — a major documented loss, a quantum milestone, an insurer withdrawing cover for the channel — the premium a rises in real terms.
- **Widening scope.** If the construction is extended to cover loss channels it does not presently address, the scope factor s rises. That is a product decision rather than a market movement, but its effect on value is the same.
- **Scarcity as the supply subscribes.** Above roughly 35% adoption the 10.5 million slots are contested, and the marginal holder's willingness to pay rises independently of the underlying loss cost.

Each is plausible and none is measured. Together they might support a g of zero to two percent, which would widen the upper bound of the range by roughly a tenth. Three things g may not include, however, and these are the reason it is held at zero:

- **Repricing toward this range.** The largest reason anyone might expect the token to appreciate is that the market may re-rate it from its current implied rate toward the range derived here. That expectation cannot be used as an input to deriving the range without assuming the conclusion. Convergence between price and fundamental is not growth in the fundamental.
- **Bitcoin appreciation.** It is already carried by the 1× / 2× / 3× columns, which state the price conditional on a Bitcoin level rather than embedding a forecast of one. Including it in g as well would count it twice.
- **Speculative demand.** It is a level effect on price, not growth in the worth of the protection, and this paper values the protection.

The figures in section 6.4 therefore contain no assumption that the market ever agrees with them.

What g = 0 does and does not assume about resale

Holding g at zero is not the same as assuming the token is worthless on exit. A perpetuity is algebraically identical to holding the protection for any number of years and then selling at the same real price — the resale is the recovery of the unconsumed remainder of the same stream, not a separate gain to be added on top. What g = 0 asserts is that the resale price equals the purchase price in Bitcoin-denominated terms: no real appreciation, and no real decline. The possibility of losing the position entirely is priced separately and explicitly, in the hazard f.

It is worth noting how little of the modelled value rests on that resale in any case. Discounted at the rates used here, resale proceeds account for about 5% of the low bound's value after ten years and well under 1% after twenty; at the high bound the corresponding figures are about 21% and 4%. The low end of the range is therefore close to a pure insurance valuation in which the token is assumed to be worth nothing at exit — the protection alone accounts for nearly all of it. A reader who wants the strict version of that test can compute it directly: a finite ten-year benefit with zero terminal value gives roughly 51 to 248 basis points, against the 53 to 314 derived below.

6.3 The failure hazard, and why joint-probability multiplication overstates the case

LEMNI requires two independent post-quantum signatures — the proprietary LAIP scheme alongside NIST-standardized Dilithium5 or Falcon-512. It is frequently argued from this that the failure probabilities multiply: if LAIP has some annual probability of catastrophic break and Dilithium has another, the joint probability is their product, and the hazard is therefore negligible.

That argument is directionally right about one channel and wrong about the total. It holds only under statistical independence, and independence does not hold across most of the ways a system of this kind actually fails.

Failure channel	Mitigated by dual signature?	Character of the risk
Cryptanalytic break of one scheme	Yes — fully	A break of LAIP alone, or of Dilithium alone, does not move protected assets. This is the design's genuine and substantial contribution.

Independent simultaneous break of both	Yes — remote	Requires unrelated advances against unrelated mathematics. Reasonably modelled as negligible.
Defect in the enforcement logic	No	The contract or verifier that checks both signatures is a single component. One bug there bypasses both signatures at once.
Shared implementation surface	No	Common signing client, vault, key-derivation path, RNG, or library. A single defect propagates to both signatures.
Seed or credential compromise; coercion	No	Defeats both authorities by construction. Explicitly outside the design's scope.
Attestation / verification service availability	No	Off-chain verification is a single operational dependency.
Supersession by native chain upgrades	n/a	Erodes the premium a rather than raising f. Addressed separately below.

The empirical record supports this decomposition. Essentially no material loss in the history of digital assets has resulted from breaking an underlying signature primitive. Losses have come from implementation defects, contract logic, key management, operational compromise, and social engineering. A model whose only failure mode is cryptanalysis is not modelling the hazard that actually retires capital.

The correct statement is therefore stronger and more credible than the one it replaces. The dual-signature architecture drives the cryptanalytic hazard — which is precisely the quantum peril the design was built for — toward zero, and no single-signature scheme can make that claim. It does not reduce the common-mode hazards, which are set by engineering discipline, audit coverage, and operational practice rather than by the choice of signature scheme. The total hazard f is dominated by those channels.

Reference annual hazard: f ≈ 2% – 8%, dominated by common-mode engineering and operational channels rather than cryptanalysis.

A reader who believes the implementation is unusually well audited should use the low end. A reader weighting the novelty of the construction should use the high end. The point is that f is a parameter of the engineering, not of the mathematics, and it cannot be argued to zero by multiplying two cryptanalytic probabilities together.

6.4 The resulting valuation range

Combining the four inputs gives the following. Each row is a complete, self-contained derivation — premium, deductible, scope, and capitalization rate — and the price columns scale by the Bitcoin reference price only. Lower premium quotes are paired with the higher deductible and the narrower scope, since in practice those travel together.

Case	a	d	s	r + f	Effective a	1× ($60k)	2× ($120k)	3× ($180k)
Conservative	0.40%	25%	0.35	35%	0.19% → 53 bps	$320	$640	$960
Cautious	0.60%	17.5%	0.45	30%	0.33% → 109 bps	$655	$1,309	$1,964
Moderate	0.60%	17.5%	0.50	24%	0.36% → 152 bps	$909	$1,818	$2,727
Constructive	0.80%	10%	0.55	20%	0.49% → 244 bps	$1,467	$2,933	$4,400
Optimistic	0.80%	10%	0.60	17%	0.53% → 314 bps	$1,882	$3,765	$5,647

The modelled range runs from 53 to 314 basis points of protected value. At a $60,000 Bitcoin price those bounds are $320 and $1,882 per LEMNI; at $120,000, $640 and $3,765; at $180,000, $960 and $5,647.

This paper quotes the bounds rather than a point estimate, and no row above is designated a central case. Collapsing the range to a single figure would require a probability distribution over four inputs — the premium, the deductible, the scope factor, and the capitalization rate — and no such distribution is available. The intermediate rows are shown to make the derivation legible, not to imply that the middle of the range is more likely than its ends. A reader who holds a view on any input should read the row that matches it.

These figures are materially lower than a derivation that borrows the stated insurance premium without adjustment. That is the correct result. The stated premium buys indemnity, net of a substantial deductible, against a set of perils considerably wider than the one the dual-authority construction addresses. Claiming the whole premium for one channel of it overstates the case by roughly a factor of two, and the overstatement is the kind a sophisticated reader finds quickly.

Sensitivity to the population adjustment

Section 6.1 identified a further factor k — the ratio of remote-compromise hazard in LEMNI's target population to that in the already-vaulted population the benchmark rate is charged to. It is excluded from the figures above. Its effect on both bounds of the range is proportional:

Population factor k	Implied rate range	1× ($60k)	3× ($180k)
1.0 — as used throughout	53 – 314 bps	$320 – $1,882	$960 – $5,647
1.5	80 – 471 bps	$480 – $2,824	$1,440 – $8,471
2.0	107 – 627 bps	$640 – $3,765	$1,920 – $11,294
2.5	133 – 784 bps	$800 – $4,706	$2,400 – $14,118

This table is shown for completeness, not as a claim. No figure elsewhere in this paper uses k above 1.0.

What the range does not assume

- **It does not assume quantum.** The premium benchmarks priced above are for present-day theft and loss perils, and the scope factor is applied against that same present-day bundle. The quantum leg is additive to this range, not a component of it.
- **It does not assume the quantum leg survives.** Native chain post-quantum upgrades would erode the quantum premium over time. They would not address the separation of ownership authority from the ledger key, which is what the theft benchmark prices.
- **It does not assume full utilization.** The figures are per-token prices conditional on the supply being used for protection. Section 6.5 addresses the quantity side.

6.5 Quantity: constant token demand across cycles

The number of tokens demanded is invariant to the market cycle. The gate requires ⌈P ÷ BTC⌉ tokens to protect a value P. If the covered market scales with Bitcoin, the numerator and denominator move together and the token count is unchanged:

	1× (BTC $60k)	2× (BTC $120k)	3× (BTC $180k)
Potential covered market	~$1.8T	~$3.6T	~$5.4T
LEMNI demanded at full adoption	30.0M	30.0M	30.0M
Against 10.5M supply	2.9×	2.9×	2.9×

At every cycle, full adoption would demand roughly 30 million tokens against a fixed 10.5 million supply. What changes across cycles is not the token demand but the dollar premium standing behind each token. Adoption below roughly 35% leaves slots idle and the rate demand-limited; above it the supply is oversubscribed and the constraint shifts to scarcity.

Adoption of covered market	LEMNI demanded	vs. 10.5M supply
10%	3.0M	0.29× — demand-limited
35%	10.5M	1.0× — fully subscribed
50%	15.0M	1.4× oversubscribed
100%	30.0M	2.9× oversubscribed

6.6 Conditions, limitations, and what would falsify this

The range in section 6.4 is conditional. It holds only if the following are true, and each is a testable proposition rather than an assumption of good faith:

- **The gate is accepted.** Holders must treat one LEMNI as conferring protection on one Bitcoin of value. If the market does not accept the denomination, the premium framing does not apply and the token is priced on other grounds entirely.
- **The protection functions.** The dual-authority construction must hold under adversarial conditions in production, not only in specification. Independent audit of the enforcement logic — not of the signature schemes — is the load-bearing evidence, because that is where the hazard concentrates.
- **The supply is utilized.** Per-token prices assume tokens are held for protection. Tokens held purely for speculation do not create the premium the model capitalizes, though they do compete for the same fixed supply.
- **The scope estimate holds.** The scope factor asserts that a specific share of self-custody loss cost runs through remote wallet-key compromise rather than through seed compromise, coercion, physical loss, or user error. It is the least evidenced input in the model. Published loss-channel data, or an underwriter willing to quote the channel separately, would materially tighten it in either direction.
- **An exit exists.** The premium is recoverable only to the extent a holder can sell. The current float is thin; a marginal quoted price on a concentrated liquidity pool is not a price at which size can transact, and fully-diluted valuation computed from it is not realizable value.

The model would be falsified by: a demonstrated bypass of the enforcement logic; loss-channel evidence placing remote key compromise well below the assumed scope range; native post-quantum chain upgrades arriving alongside a credible native solution to key-compromise theft; sustained failure of adoption to approach the subscription threshold; or the emergence of a competing dual-authority construction without the patent constraint.

Basis of presentation

This section is a valuation model, not a forecast, an offer, or investment advice. It is prepared by the issuer of the asset it describes and should be read with that interest in view. The figures are conditional outputs of stated assumptions, and every input — the premium benchmark, the discount rate, the hazard, the adoption rate — is uncertain and may be substituted by the reader. No return, timeline, or price outcome is projected, promised, or implied. Any prospective purchaser should conduct independent verification of the technical claims and independent assessment of the assumptions before relying on any figure here.

Appendix A: Three methodological notes

A.1 Basis-point conversion

A capitalized value expressed in dollars must be divided by the protected value before being described in basis points, and the conversion is a frequent source of order-of-magnitude error. A $20,000 capitalized value against a $60,000 Bitcoin is 33.3% of protected value — 3,333 basis points, not 333. The discipline that prevents this is to work in the rate throughout: compute a ÷ (r + f) directly as a percentage of protected value, and convert to a token price only at the last step by multiplying by the Bitcoin price. No dollar figure need enter the derivation at all.

The error matters beyond arithmetic hygiene, because the correct figure is usually a signal that the model is wrong. A one-time premium of 33% of protected value would never be paid by a rational holder, since reducing the position achieves comparable risk reduction at no cost. A result in that region indicates the discount rate is too low, not that the protection is extraordinarily valuable.

A.2 Discount-rate selection

Capitalizing a perpetual benefit at a risk-free or near-risk-free rate asserts that the benefit stream is as certain as a sovereign bond. For a novel protective technology with thin liquidity, adoption dependency, and unaudited components, that assertion is not available. The capitalization multiple is 1 ÷ (r + f), and it should be defended as a discount rate rather than chosen as a multiple, because the multiple form obscures what is being claimed. The band used in this note — 17% to 35% — corresponds to multiples of roughly 3× to 6× and reflects venture-like uncertainty with an explicit hazard term.

A.3 Independence in joint-failure estimates

Multiplying the failure probabilities of two redundant components requires those components to fail independently. In a signing architecture they generally do not: the enforcement logic that verifies both signatures, the client that

produces them, the key-derivation path, the randomness source, and the holder's own credentials are shared. A defect in any of these is a common-mode failure that redundancy does not address.

The practical consequence is that the redundancy argument should be made narrowly and made well. Dual independent post-quantum signatures are a genuine and substantial defence against cryptanalytic compromise, which is exactly the quantum peril, and that claim is strong enough to stand without exaggeration. Extending it to a claim of negligible total failure risk both overstates the case and invites a sophisticated reader to discount the whole analysis. A model that names its common-mode hazards and prices them is more persuasive to an institutional audience than one that argues them away.

A.4 Adjusting a market premium before using it

A quoted insurance premium is a price for a specific product, and three properties of that product have to be stripped out before the number can be used as an input to anything else. It is not first-dollar — a deductible, often a percentage rather than a fixed sum, sits between the loss and the payment. It covers a bundle of perils rather than the single channel a prevention mechanism addresses. And it pays contingently, on an approved claim, subject to exclusions that may include the very tail event of interest.

The first and third adjustments raise the implied value of prevention; the second lowers it, generally by more. Applying the premium without adjustment, as though it were a first-dollar price for exactly the covered channel, roughly doubles the result. The failure is not arithmetic — it is treating a market price as more comparable than it is, which is the same class of error as A.2 and A.3.

Appendix B: Summary of parameters

Parameter	Value used	Basis
a — stated annual premium	0.40% – 0.80%	Regulated Lloyd's-backed self-custody Bitcoin underwriting; provider's own estimation midpoint is 0.60%.
d — deductible embedded in a	10% or 25%	Selected by the insured at policy inception; assessed against an approved claim benefit.
s — peril-scope factor	0.35 – 0.60	Share of the covered bundle's expected loss cost attributable to remote wallet-key compromise. Judgment, not observable.
Effective premium a × s ÷ (1 − d)	0.19% – 0.53%	What the protection actually displaces.
r — discount rate	15% – 25%	Venture-appropriate for an illiquid, adoption-dependent claim.
f — annual failure hazard	2% – 8%	Dominated by enforcement-logic, shared-implementation and operational channels; cryptanalytic component negligible.
r + f — capitalization rate	17% – 35%	Multiple of roughly 3× to 6×.
g — real growth in the premium	0% (held at zero)	Excludes repricing toward this range, Bitcoin appreciation, and speculative demand. See 6.2.
k — population factor	1.0 (excluded)	Shown as sensitivity only; not used in any headline figure.
Resulting one-time rate	53 – 314 bps (bounds; no point estimate)	[a × s ÷ (1 − d)] ÷ (r + f).
BTC reference price	$60,000 (1×)	Consistent with the supply-ceiling and coverage figures used throughout.
Fixed supply	10,500,000	No issuance.