DeFi: A Market Mechanism for Cybersecurity Risk Insurance

Decentralized Finance & Crypto Workshop @ Scuola Normale Superiore

Björn Hanneke
PhD Candidate at Goethe University Frankfurt, Chair of Information Systems and Information Management

Cybersecurity risk is of ongoing concern of growing DeFi ecosystems

As DeFi TVL has been increasing to, so have cybersecurity-related losses

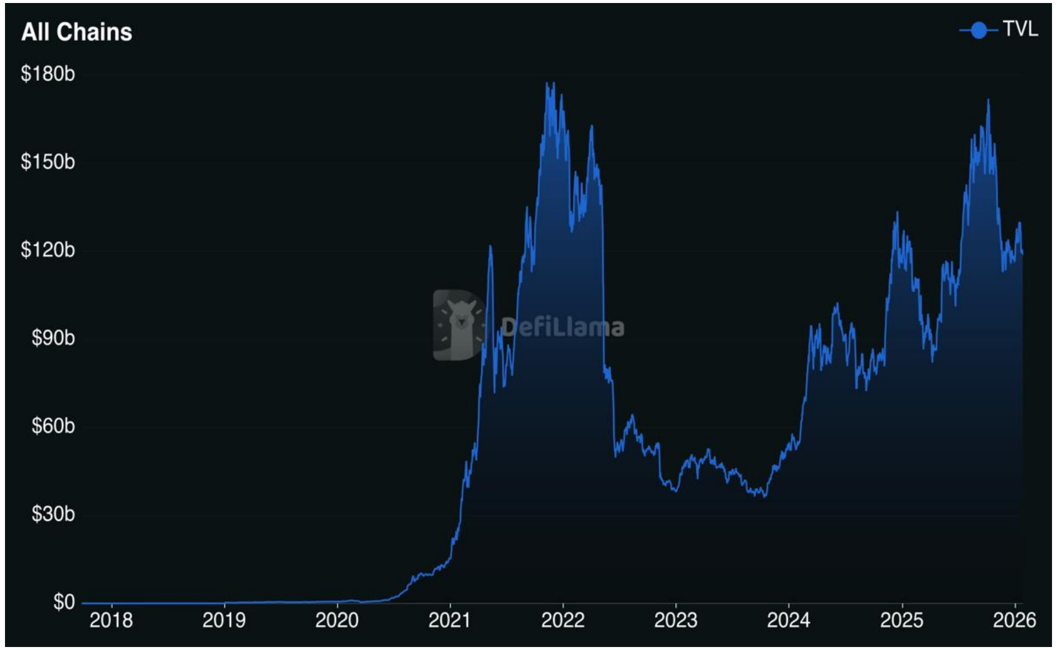

defillama.com/

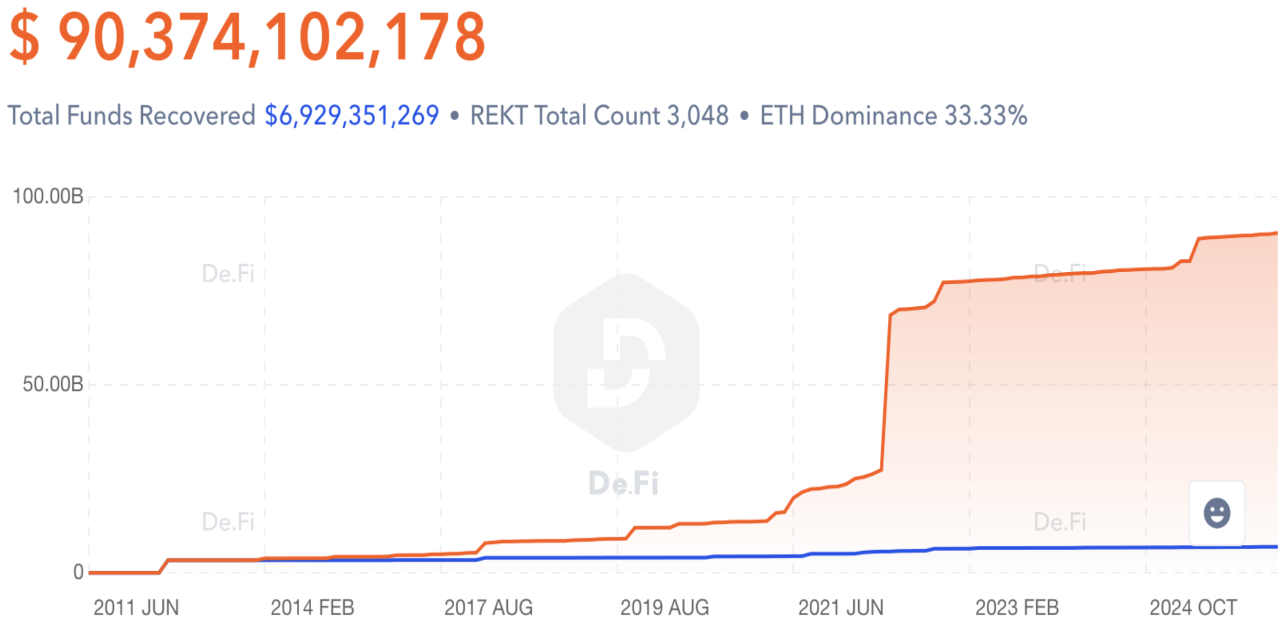

de.fi/rekt-database

Yet, cybersecurity insurance protocols have failed to generate substantial adoption

Name	TVL	Fees 7d	Revenue 7d	Mcap/TVL	Fees 30d
1 Nexus Mutual 1 chain	$136.04m	$122,274	$61,137	0.98	$386,238
2 Easedefi.org 1 chain	$6.12m				
3 3F Mutual 1 chain	$5.55m				
4 Unslashed 1 chain	$4.23m			0.09	
5 Cover Protocol 1 chain	$809,499			0.05	

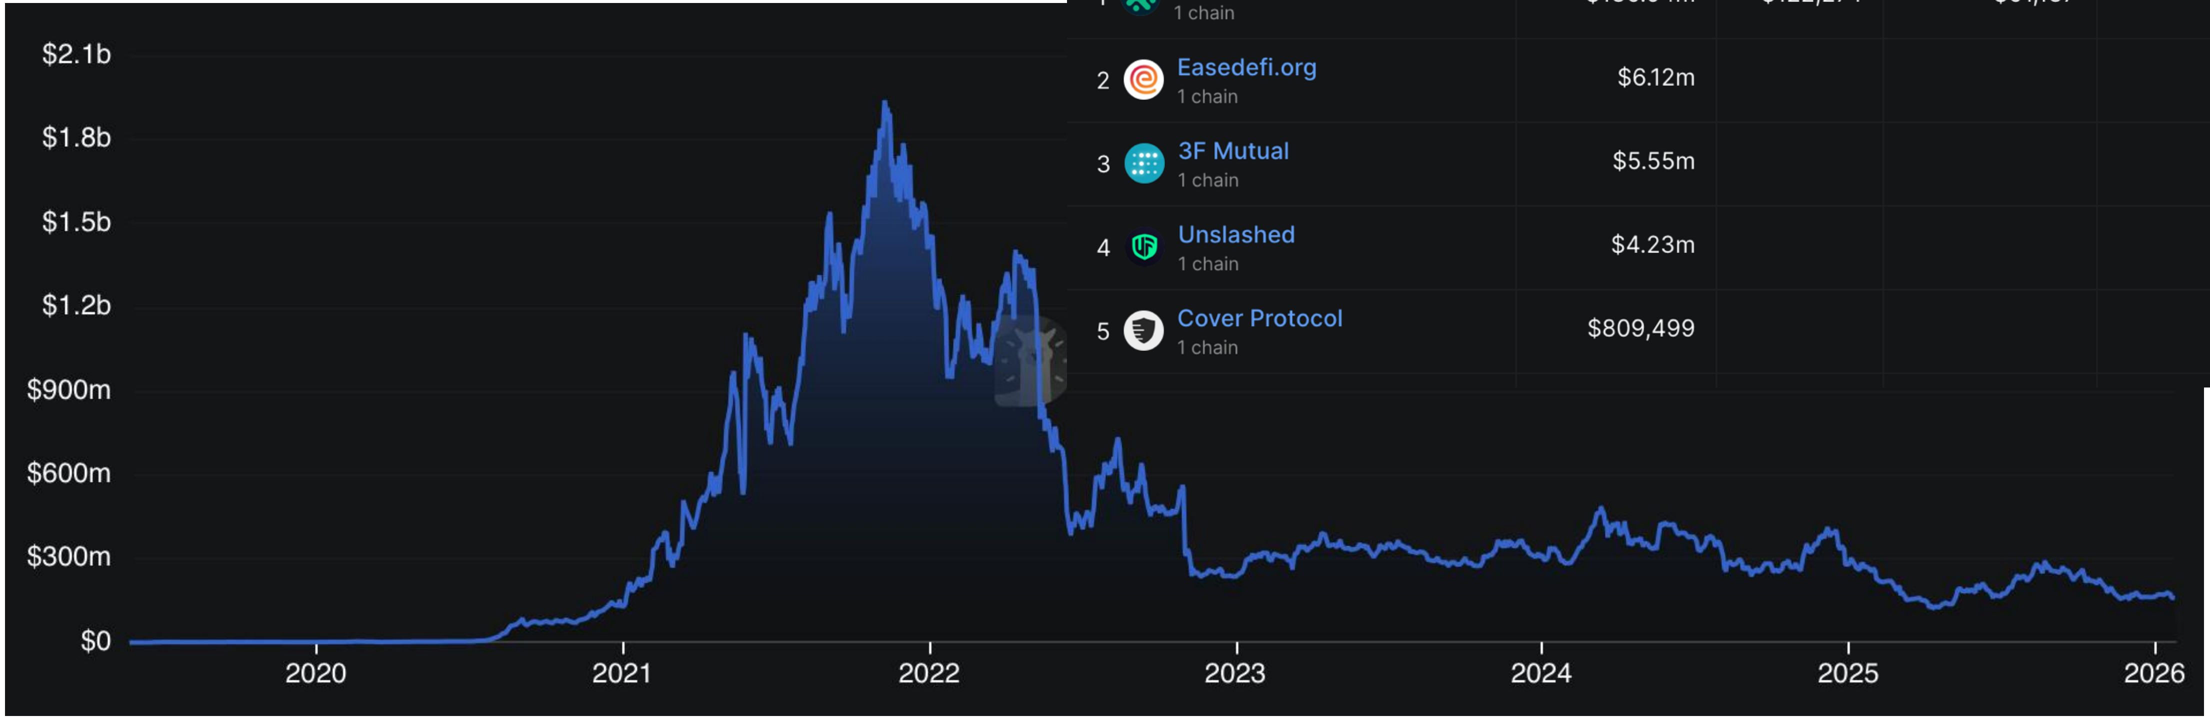

https://defillama.com/

In our mechanism, a trusted Operator verifies losses

1. Loss Determination

- Fully automated → Can't handle ambiguous hacks (partial exploits, MEV)
- Human voting (Nexus Mutual) → Slow, manipulable, hard to scale
- Our choice: Trusted operator verifies → Fast, scales, but requires trust

In our mechanism, compensation is driven by market forces (risk price and utilization)

1. Loss Determination

- Fully automated → Can't handle ambiguous hacks (partial exploits, MEV)
- Human voting (Nexus Mutual) → Slow, manipulable, hard to scale
- Our choice: Trusted operator verifies → Fast, scales, but requires trust

2. Risk Compensation

- Mutual ex-post loss sharing → No upfront price signal
- Explicit yield-based compensation → Requires pricing mechanism
- Our choice: Yield-share function $\gamma(U, P_{risk})$ → Market-driven compensation

In our mechanism, pooled capital for efficiency

1. Loss Determination

- Fully automated → Can't handle ambiguous hacks (partial exploits, MEV)
- Human voting (Nexus Mutual) → Slow, manipulable, hard to scale
- Our choice: Trusted operator verifies → Fast, scales, but requires trust

2. Risk Compensation

- Mutual ex-post loss sharing → No upfront price signal
- Explicit yield-based compensation → Requires pricing mechanism
- Our choice: Yield-share function $\gamma(U, P_{risk})$ → Market-driven compensation

3. Capital Structure

- Pairwise contracts → Bespoke but capital inefficient
- Pooled capital → Efficient but requires standardization
- Our choice: Single pool + parametric coverage → Scalable pooling

We target protocol-level insurance with market-based risk pricing

1. Loss Determination

- Fully automated → Can't handle ambiguous hacks (partial exploits, MEV)
- Human voting (Nexus Mutual) → Slow, manipulable, doesn't scale
- Our choice: Trusted operator verifies → Fast, scales, but requires trust

2. Risk Compensation

- Mutual ex-post loss sharing → No upfront price signal
- Explicit yield-based compensation → Requires pricing mechanism
- Our choice: Yield-share function $\gamma(U, P_{risk})$ → Market-driven compensation

3. Capital Structure

- Pairwise contracts → Bespoke but capital inefficient
- Pooled capital → Efficient but requires standardization
- Our choice: Single pool + parametric coverage → Scalable pooling

Design decision:
Prioritize scalability and explicit compensation

Trade-off:
Operator trust assumption

Separating pricing (speculators) from bearing risk (LPs) reduces information asymmetry without diluting incentives

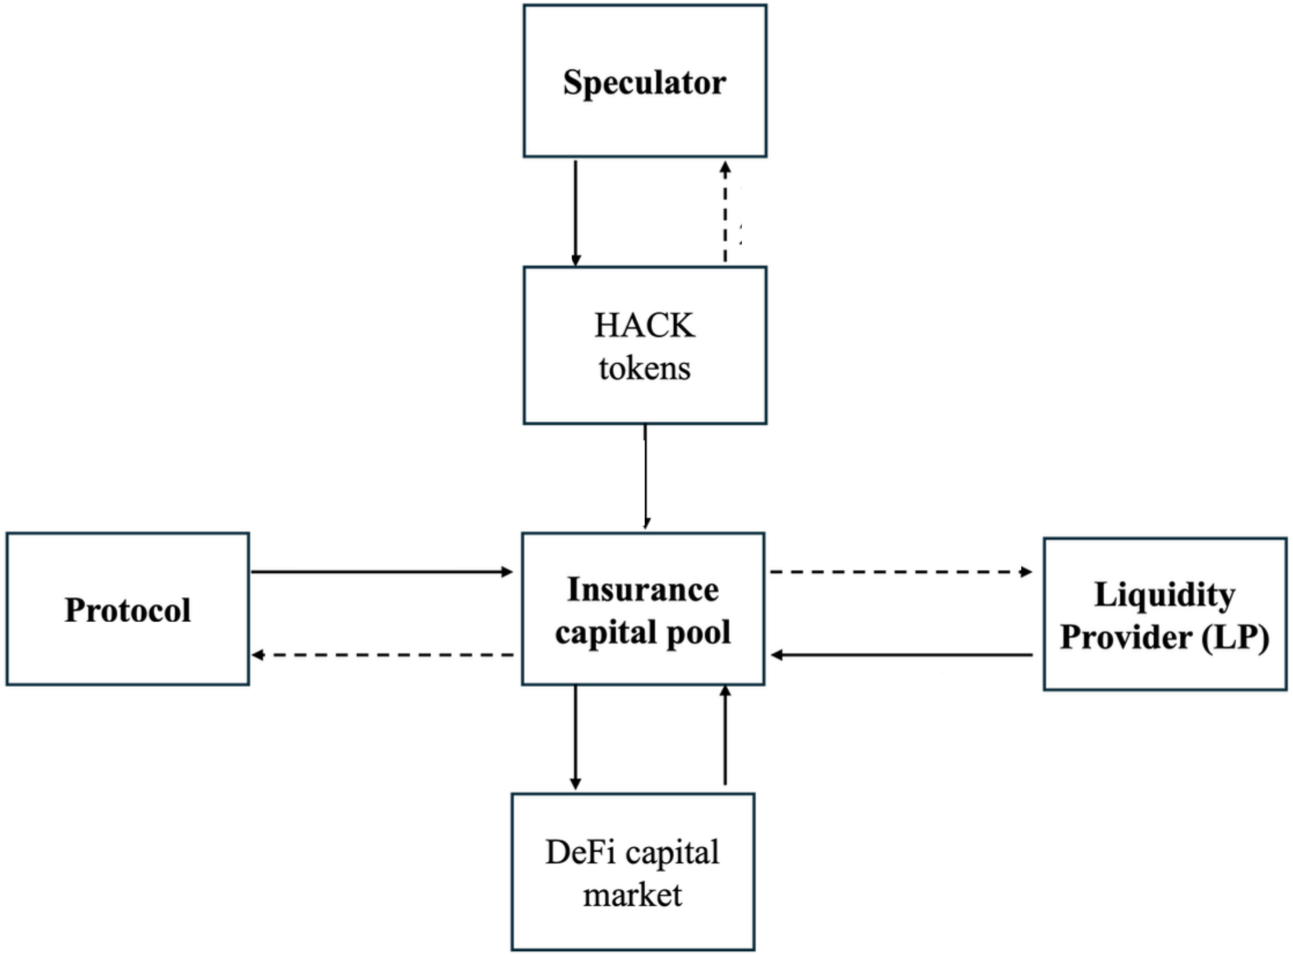

Stakeholders

- (Trusted) "Operator": Sets up and governs the mechanism
- Protocols: Seek insurance
- LPs: Underwrite risk / provide capital
- Speculators: Price risk

Mechanism

- "Shared" Insurance capital pool
- Capital pool covers losses in case of hacks
- Dynamic distribution of pooled capital yield

The Operator’s incentives align with platform growth regarding all stakeholders

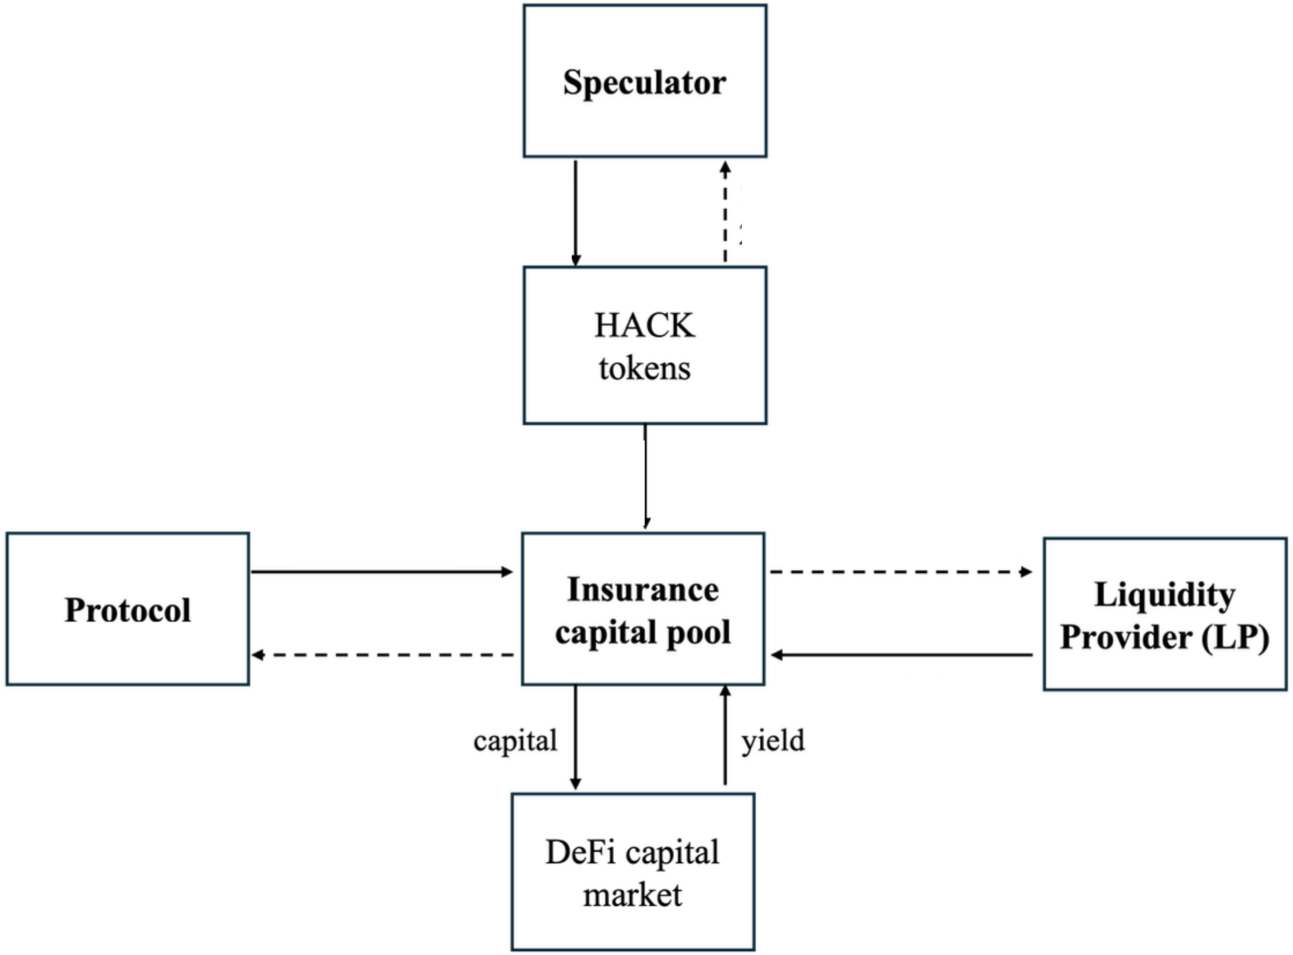

(Trusted) “Operator”

- Operator ensures yield generation by insurance funds, manages the insurance protocol and governs its constraints.
- The operator retains a fixed fee from total pool yield.
- The remaining yield is split between LPs and insured protocols according to the yield-sharing rule.
- The operator is economically disciplined by reputation and repeated interaction.

Protocols provide “costly” collateral but receive a yield-share if no hack occurs, effectively reducing their insurance costs

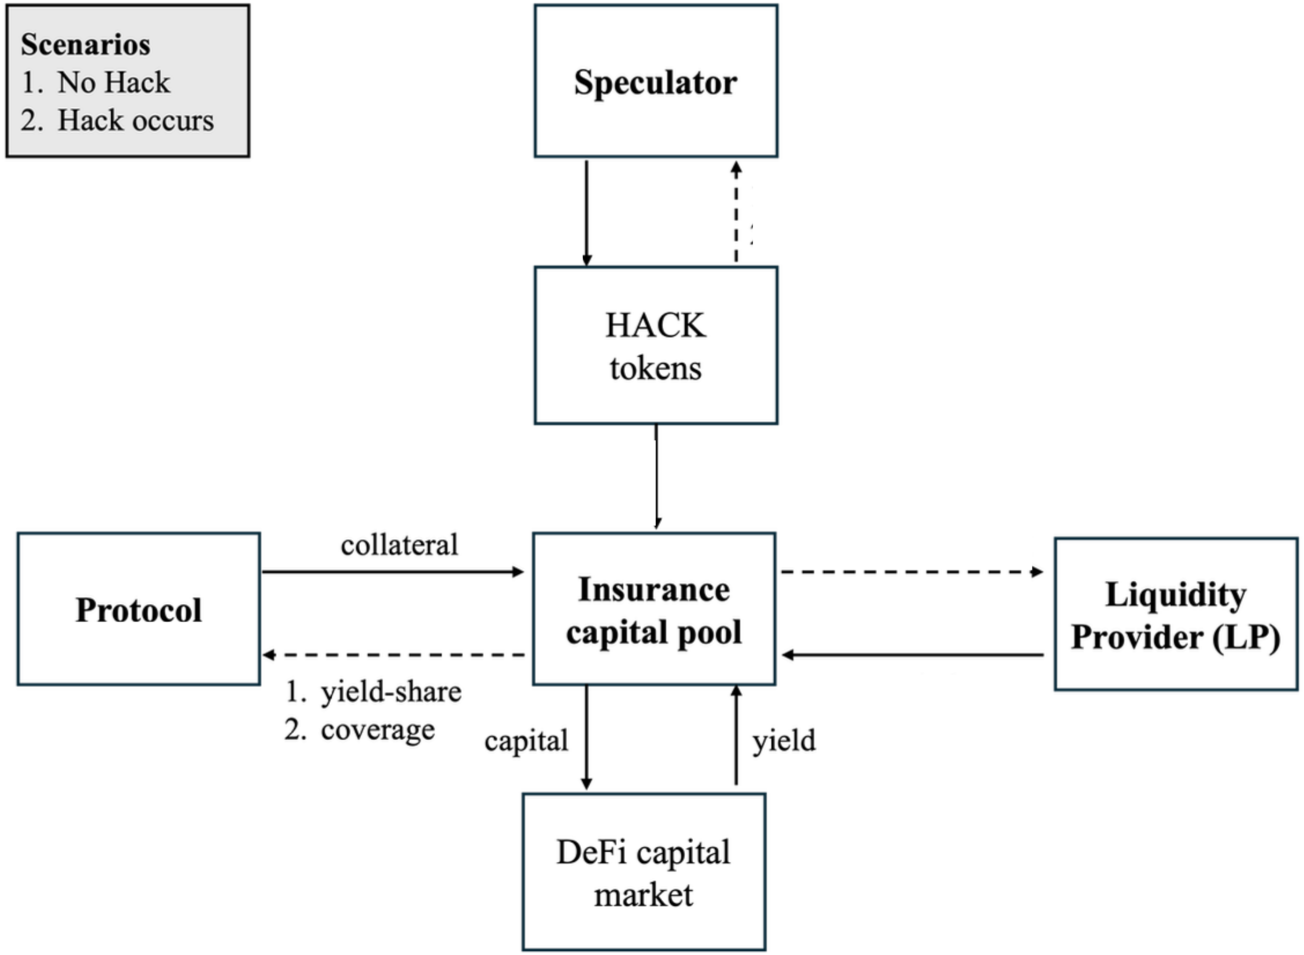

Protocols

- Seek insurance for irreducible cybersecurity risk
- Have to provide collateral for coverage

$$\boxed{\text{coverage} = \mu \cdot C_{\mathrm{C}}^{\theta} \cdot (1+\xi)}$$

- The more collateral (C_C), the higher the coverage (where $\mu>0$ and $\theta \in (0,1)$ calibrate scale and concavity); scaling (ζ) for positive security alignment (audits, etc.)
- In case of hack, collateral is forfeited but does not reduce insurance payouts
- In case of no hack, participate in generated yield

LPs are compensated for bearing irreducible cybersecurity risk

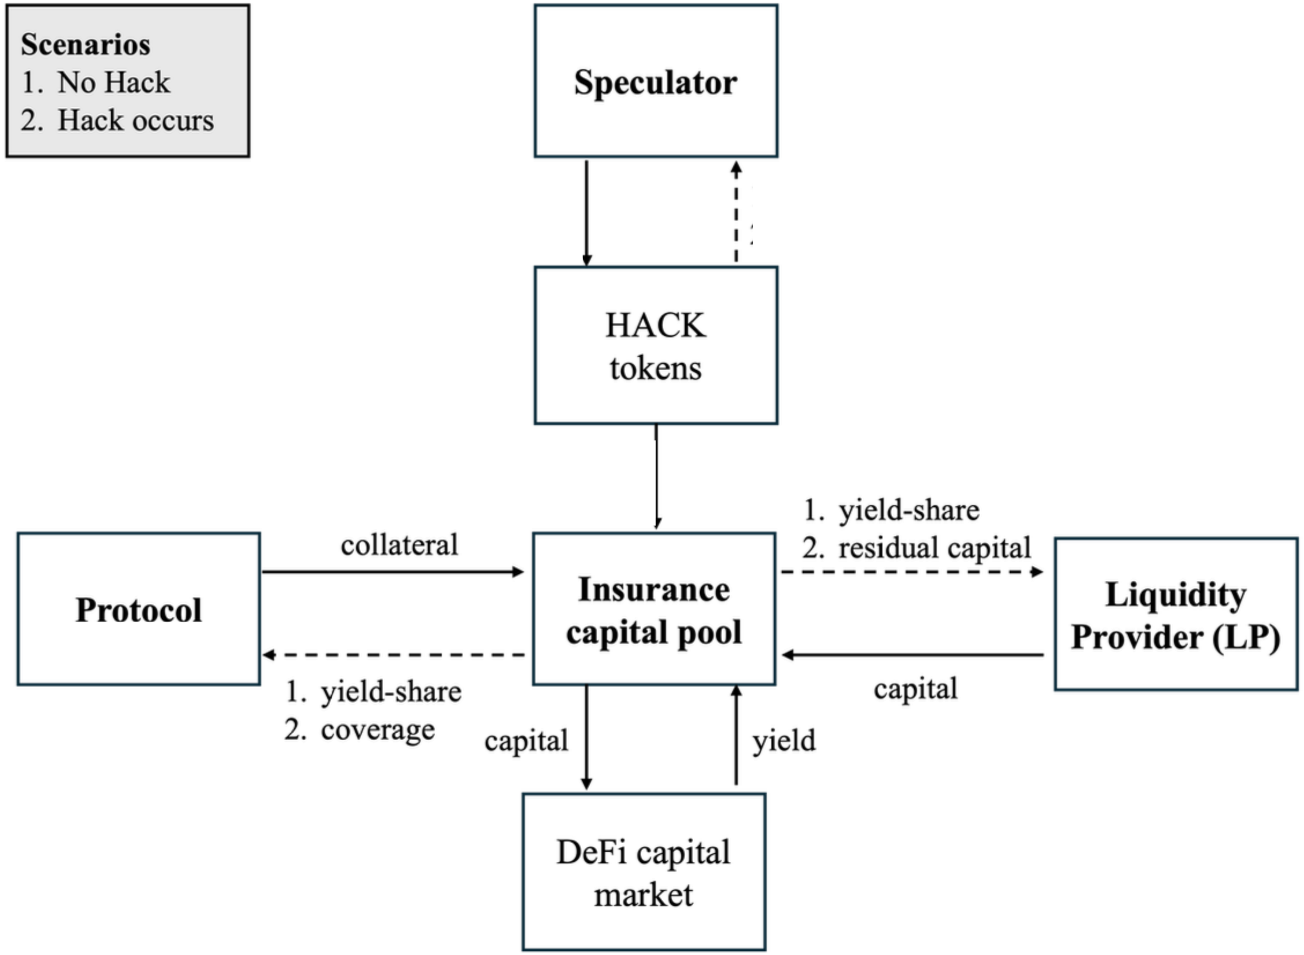

Liquidity Provider

- LPs supply capital to earn a share of the insurance yield
- Capital is exposed to insured losses in hack states
- Even if the pool earns only the market return, LP capital returns can be higher, because LPs gain exposure to:
  - yield generated by protocol collateral,
  - fees from risk pricing (HACK tokens)

→ LPs are compensated for bearing irreducible cybersecurity risks

Speculators provide forward-looking risk signals

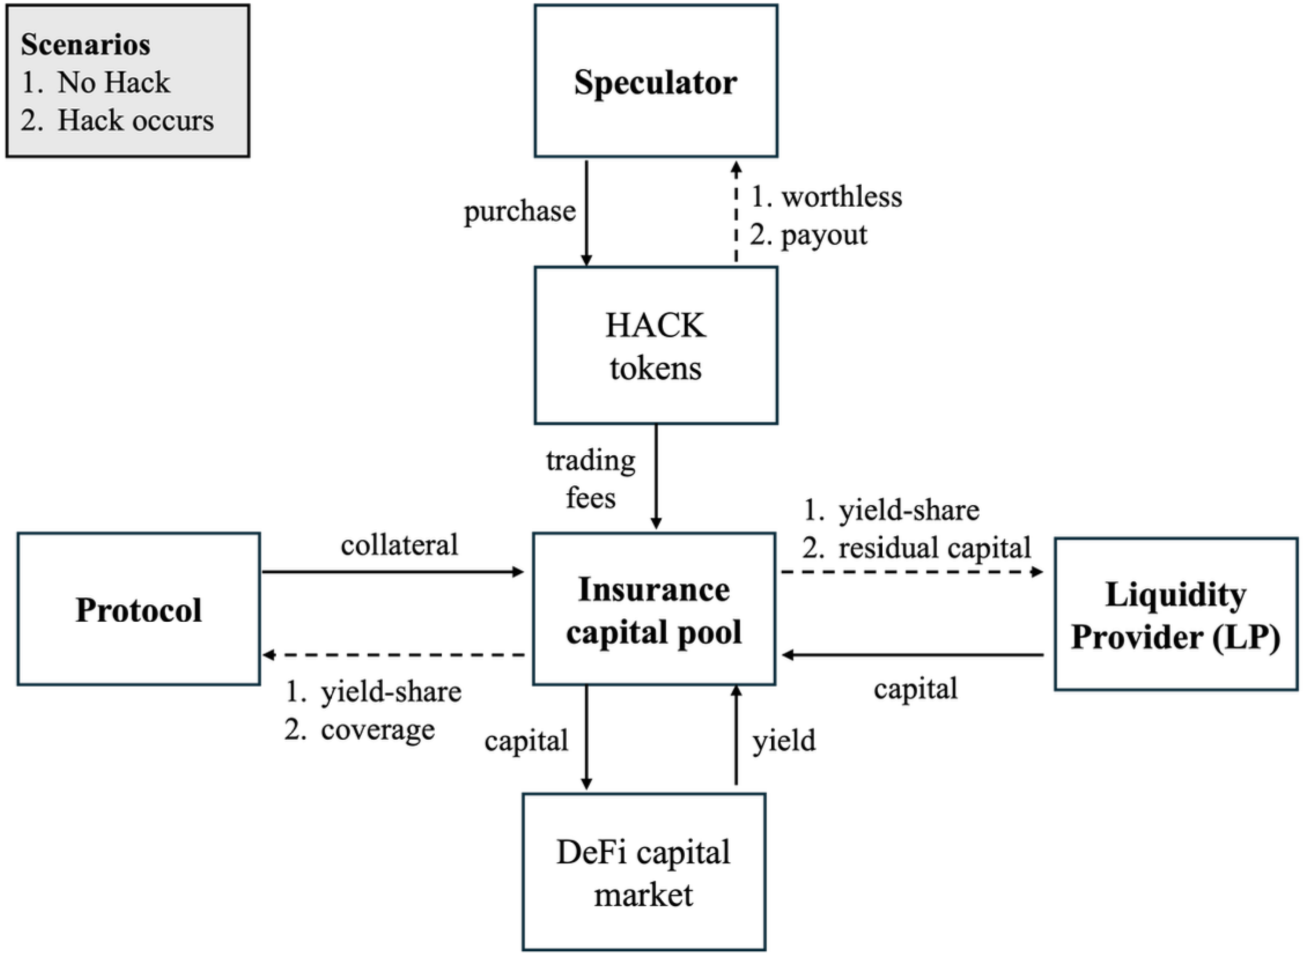

Speculators

- Speculators trade **binary HACK tokens** at **multiple maturities (e.g., quarterly expiries)** (i.e., Polymarket-style prediction markets)
- Prices aggregate forward-looking beliefs about hack likelihoods
- the fair value being the discounted expected payout:

$$\boxed{V_{\text{HACK}}(T) = p_{\text{hack}}(T) \cdot DF(T).}$$

- Prices do *not* determine payouts or transfers directly
- Prices only affect policy bounds (e.g., yield-sharing --> *more details to follow*)

DeFi primitives define utilization and risk-based yield-sharing

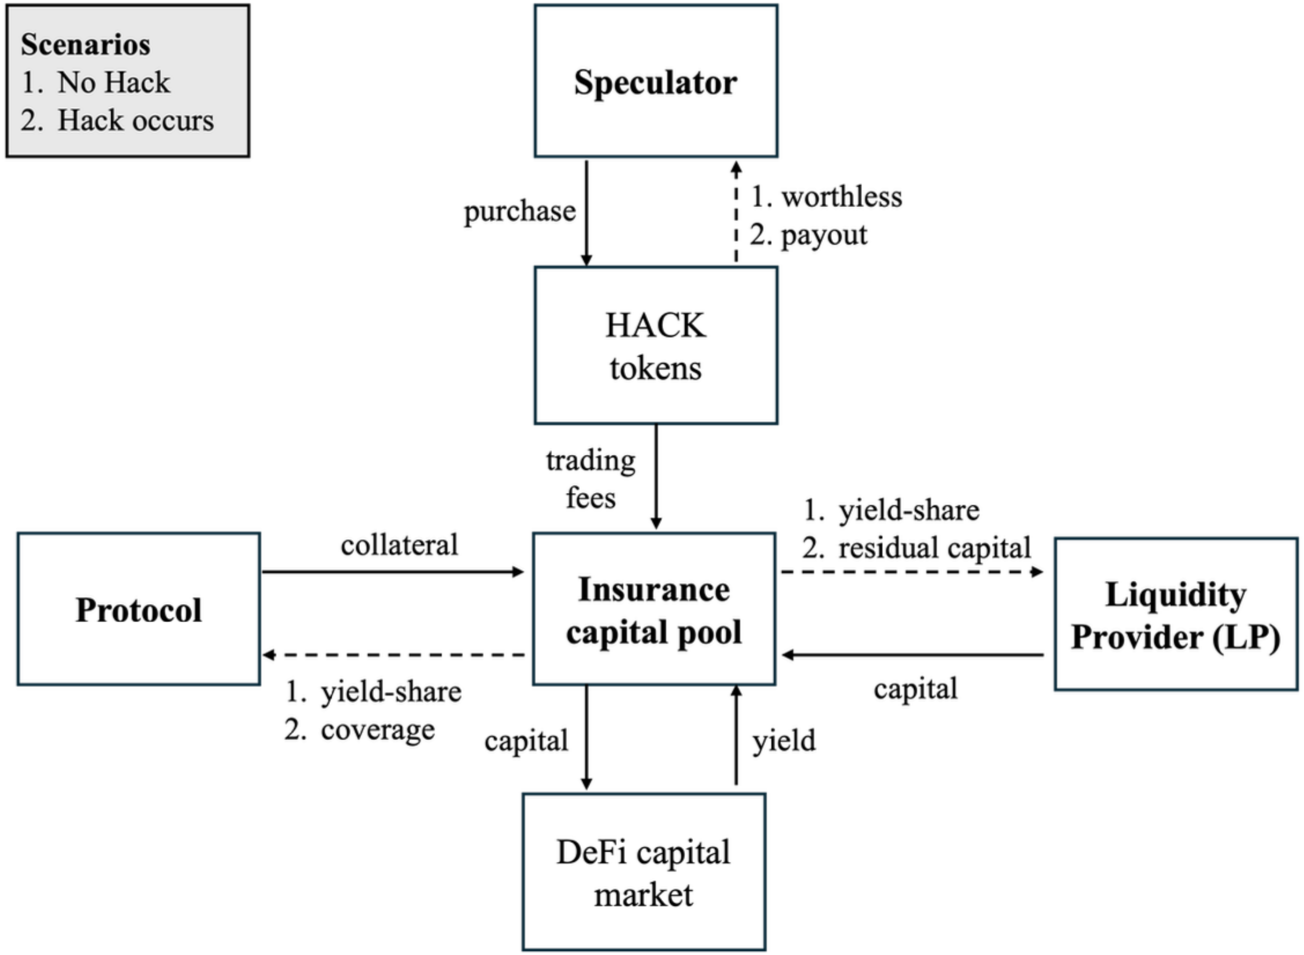

Core mechanisms

- **Utilization** measures how much LP capital is committed relative to coverage obligations:

$$U = \frac{\text{coverage}}{C_{LP}}$$

- **Risk signal** from insurance price index:

$$P_{\text{risk}}(t) = \sum_i \omega_i \, P_{\text{HACK}}(T_i, t), \qquad \omega_i \geq 0, \ \sum_i \omega_i = 1,$$

- **Risk-based yield-sharing:** LP compensation increases when capital is scarce or perceived risk rises:

$$\gamma(t) = \alpha \left(\frac{U}{U_{\text{target}}} \right)^{\beta} + (1 - \alpha) \left(\frac{P_{\text{risk}}(t)}{P_{\text{anchor}}(t)} \right)^{\delta}$$

Stakeholder objectives ensure incentive alignment

Protocol Objectives

- **Chooses collateral C_C** to trade off opportunity cost against tail-risk protection

$$\begin{aligned}\pi_{\text{protocol}} = p_{\text{hack}} \cdot \mathbb{E}\big[\min\big(\text{coverage}, \text{Loss}\big)\big] \\ + (1-\gamma)\,(1-\varphi)\,Y_{\text{total}} - C_{\text{C}}\,\mathbb{E}[r_{\text{market}}] \\ + \rho_{\text{P}}\,p_{\text{hack}}\,\mathbb{E}\big[\min\big(\text{coverage}, \text{Loss}\big)\big]\,.\end{aligned}$$

LP Objectives

- Supplies LP capital to maximize expected profit under insured loss exposure
- Participation governed by an expected-return constraint

$$\pi_{\text{LP}} = \gamma\,(1-\varphi)\,Y_{\text{total}} - p_{\text{hack}}\,\mathbb{E}\big[\min\big(\text{coverage}, \text{Loss}\big)\big] - C_{\text{LP}}\,\mathbb{E}[r_{\text{market}}]$$

We derive analytical properties of the mechanism and its stakeholders

- **Theorem 1: Existence of Three-Party Equilibrium**
  - Protocols, LPs, and Speculators have mutually consistent best responses under continuity/compactness and monotone policy-feedback assumptions (fixed-point existence).
- **Proposition 1: Truthful Risk Assessment by Speculators**
  - In equilibrium, prices are incentive-compatible signals of hack likelihoods and aggregate dispersed information.
- **Proposition 2: LP Dynamics and Participation Bounds**
  - LP capital adjusts endogenously until expected returns satisfy a participation constraint; resulting in bounded exposure and stable pool dynamics.
- **Proposition 3: Sustainable Undercapitalization Bounds**
  - Utilization caps and yield-sharing bound insured exposure relative to capital, preventing runaway undercapitalization and supporting solvency under adverse realizations.

Anti-cyclical, risk-driven compensation and leverage adjustments support insurance solvency

1. **Protocols choose collateral optimally**
   - Collateral balances opportunity cost against tail-risk reduction.
2. **LP capital adjusts endogenously**
   - Capital enters or exits until participation constraints are satisfied.
3. **Risk signals provide input for policy feedback and controls (but do not affect payouts directly)**
   - Market-implied hack probabilities tighten utilization caps and adjust yield-sharing.
4. **Negative feedback ensures stability**
   - Higher risk/utilization $\Rightarrow$ tighter leverage and higher compensation $\Rightarrow$ restoring LP incentives.
   - Lower risk/utilization $\Rightarrow$ looser leverage and lower compensation $\Rightarrow$ weakening LP incentives.
5. **Equilibrium is self-enforcing**
   - No stakeholder can improve payoff unilaterally given prices, policy bounds, and pool responses.

The mechanism is robust across a wide parameter range (stress test, not calibration) and demonstrates self-sustaining properties

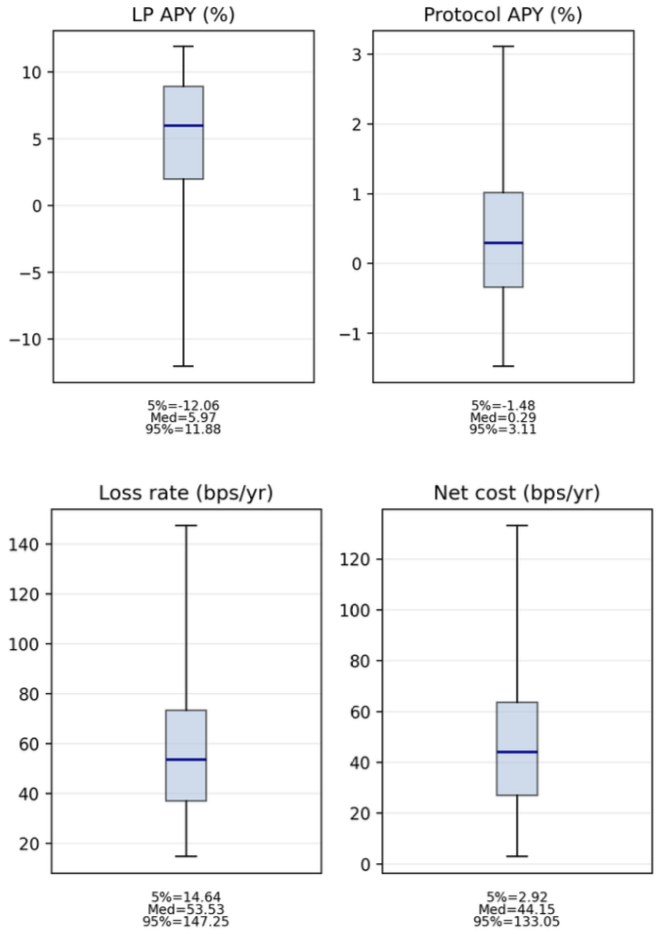

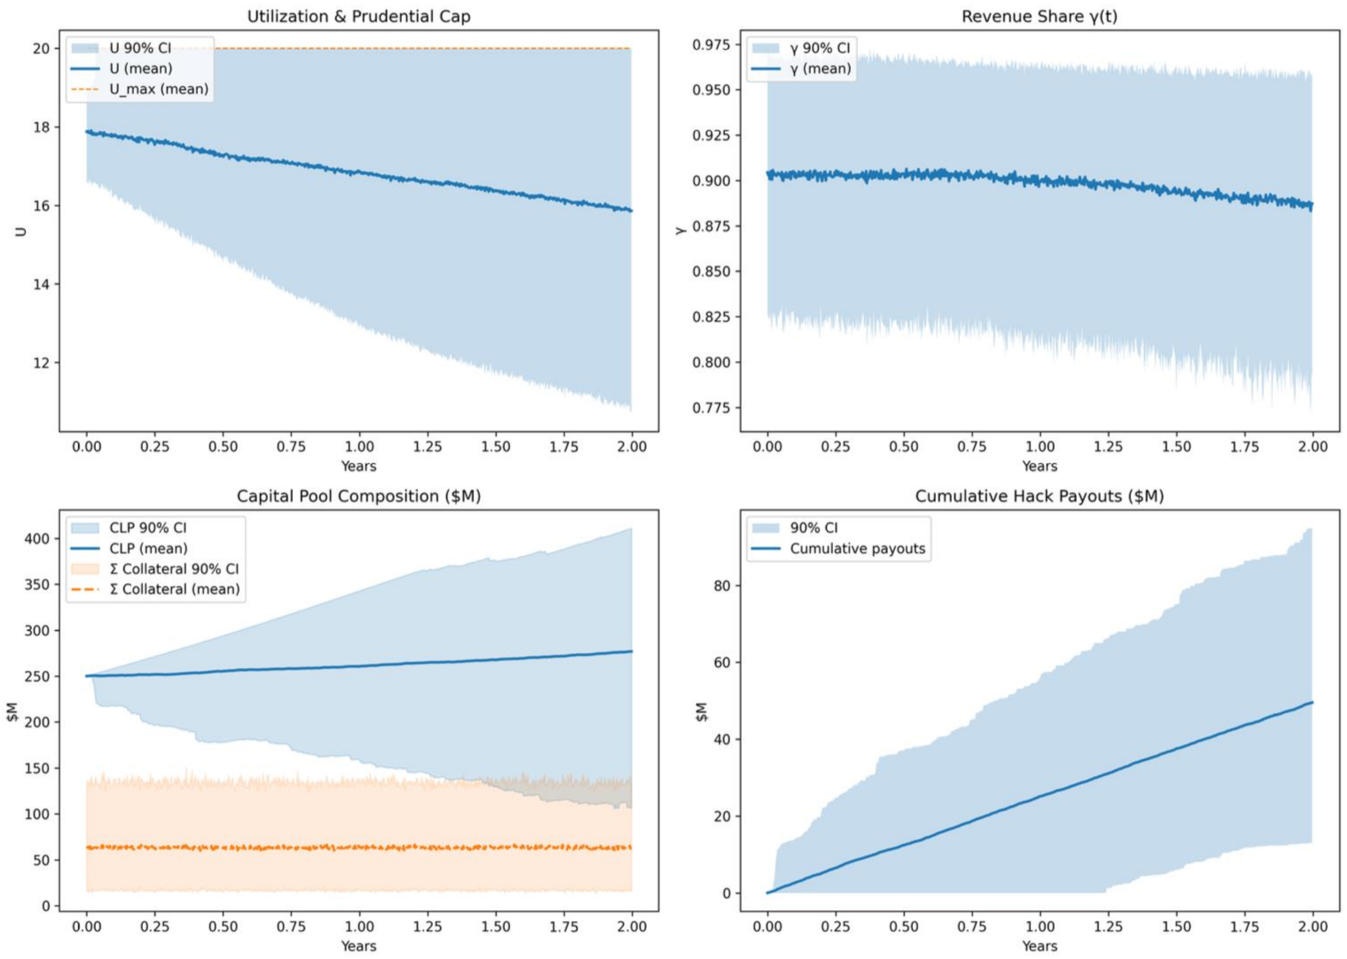

Real world implementation hinges on a few assumptions

1. **Trusted operator as coordinating entity**
   - Centralized loss verification, parameter updates, and enforcement trade full trust minimization for deployability and capital efficiency.
2. **Governance and oversight constraints**
   - Operator discipline via reputation, DAO oversight, and transparent on-chain accounting.
3. **Mechanism operations and liquidity requirements**
   - Risk-pricing accuracy depends on sufficient participation and market depth
   - Collateral/liquidity provision and withdrawal must be defined to prevent manipulation
4. **Parameter calibration and robustness**
   - Utilization caps, yield-sharing functions, and collateral multipliers require calibration but admit wide stability regions.
5. **Failure modes and stress scenarios**
   - Extreme correlated hacks, oracle outages, or operator failure shift the mechanism into conservative regimes (coverage reduction, capital withdrawal).

Designing the market mechanism involves a trilemma

1. **DeFi amplifies irreducible cybersecurity risk** through larger attack surface and asset exposure.
2. **Protocol-level, market-based insurance** aligns incentives and reduces moral hazard.
3. **Separating risk pricing from risk bearing** enables forward-looking, incentive-compatible pricing.
4. **Explicit trade-off**: scalable and capital-efficient pooling requires trusted coordination.
5. **Positioning**: complementary to fully decentralized designs, optimized for deployability and scalability.

Ecosystem-level cybersecurity risk insurance can improve capital efficiency, competitiveness, and security robustness

1. **Ecosystem foundations as insurers:**
   Foundations can deploy insurance as shared infrastructure to protect core protocols and user funds.
2. **Ecosystem treasuries as anchor investors:**
   Treasury capital can seed insurance pools, crowding in external LPs and stabilizing early participation.
3. **Insurance as an internal rescue mechanism:**
   Insurance pools can function as pre-funded, rule-based recovery funds after major incidents.
4. **Insurance as a competitive advantage:**
   Ecosystems offering credible, scalable insurance may attract protocols, developers, and long-term capital.
5. **Toward modular risk infrastructure:**
   The mechanism can be adapted across ecosystems, asset classes, and governance structures.
6. **Alternative pricing and signal designs are possible:**
   The mechanism is compatible with different market structures and risk-signal sources.

Thank you for having me today!

- **Please reach out any time if you have questions or want to build the mechanism**

- **Contact**
  - Email: hanneke@wiwi.uni-frankfurt.de
  - Whatsapp: +49 151 188 49 403
  - https://www.wiim.uni-frankfurt.de/en/team/bjoern-hanneke
  - https://www.linkedin.com/in/bhanneke/

Backup

Table 1. Major Model Parameters.

Symbol	Definition	Domain / Units
C_{C}	Protocol's posted collateral	$\mathbb{R}_{>0}$
C_{LP}	Liquidity providers' (LP) risk capital	$\mathbb{R}_{>0}$
C_{S}	Trading-fee inflows from insurance pricing layer	$\mathbb{R}_{\geq 0}$
C_{total}	Total capital in insurance pool	$\mathbb{R}_{>0}$
Y_{total}	Absolute yield generated by pool capital	$\mathbb{R}_{\geq 0}$
r_{pool}	Return rate on insurance pool capital	$\mathbb{R}_{\geq 0}$
TVL	Total value locked (protocol assets)	$\mathbb{R}_{\geq 0}$
φ	Operator fee from pool yield	$[0, 1]$
μ	Coverage amplification factor	$\mathbb{R}_{>0}$
θ	Coverage concavity parameter	$(0, 1)$
ξ	Coverage security multiplier (audits, bounties)	$(0, 1]$
α	Weight on utilization vs. risk in γ	$[0, 1]$
β	Convexity parameter for utilization in γ	> 0
δ	Convexity parameter for risk price in γ	> 0
ω_i	Weight of expiry T_i in P_{risk}	> 0
coverage	Maximum insurable amount (cap)	$\mathbb{R}_{\geq 0}$
$\gamma(U, P)$	Yield-share function	$[0, 1]$
γ_{min}	Minimum yield-share required by LPs	$[0, 1]$
r_{LP}	Realized LP return	$\mathbb{R}_{\geq 0}$
U	Utilization: coverage/C_{LP}	$[0, \infty)$
$\mathrm{U}_{\mathrm{max}}(t)$	Dynamic leverage ceiling (Eq. (5))	policy bounds (e.g., $[1, 3]$)
$\mathrm{U}_{\mathrm{target}}$	Target utilization level	$[0, \infty)$
$P_{\mathrm{HACK}}(T_i, t)$	Market price of HACK token for expiry T_i	$[0, 1]$
$P_{\mathrm{risk}}(t)$	Weighted average hack probability across expiries	$[0, 1]$
$P_{\mathrm{anchor}}(t)$	Annualized hack probability (anchor)	$[0, 1]$
$p_{\mathrm{hack}}(T)$	Probability of hack before expiry T	$[0, 1]$
$\widehat{p}_{1\mathrm{Y}}(t)$	Market-implied annualized hack probability	$[0, 1]$
$\widehat{\lambda}(t)$	Estimated hazard rate (Eq. 4)	$\mathbb{R}_{\geq 0}$
$DF(T)$	Discount factor for maturity T	$(0, 1]$
ρ_{P}	Protocol risk-aversion coefficient	$(0, \infty)$
ρ_{LP}	LP risk-premium coefficient	$(0, \infty)$
r_{market}	External market return rate	$\mathbb{R}_{\geq 0}$
π_{protocol}, π_{LP}	Profit functions of protocol and LPs	$\mathbb{R}$
κ_U	Sensitivity of prudential cap $U_{\mathrm{max}}(t)$ to $\widehat{p}_{1\mathrm{Y}}(t)$ (Eq. (5))	$\mathbb{R}_{\geq 0}$
κ_{LP}	LP capital adjustment speed in dynamics (Thm. 2(ii))	$\mathbb{R}_{>0}$

Analytical Propositions

Theorem 1: Existence of Three-Party Equilibrium

For any continuous yield-share function $\gamma : [0, 1] \times \mathbb{R}_+ \to [0, 1]$ that is increasing in both utilization and risk price P_{risk}, there exists a Nash equilibrium $(C_{\text{C}}^*,\ C_{\text{LP}}^*,\ S^*)$, where C_{C}^* is the optimal protocol collateral, C_{LP}^* is the optimal LP capital supply, and S^* is the optimal speculator demand.

Proof. We construct the equilibrium through best-response functions. Let the strategy spaces be compact intervals: $C_{\text{C}} \in [0, W_{\text{C}}]$, $C_{\text{LP}} \in [0, W_{\text{LP}}]$, and $S \in [0, W_{\text{S}}]$, where the upper bounds represent wealth constraints.

Using the profit functions (9)–(11), we define the best responses as follows:

$$C_{\text{C}}^*(C_{\text{LP}}, S) = \arg \max_{C_{\text{C}} \in [0, W_{\text{C}}]} \pi_{\text{protocol}}(C_{\text{C}}, C_{\text{LP}}, S) \tag{13}$$

$$C_{\text{LP}}^*(C_{\text{C}}, S) = \arg \max_{C_{\text{LP}} \in [0, W_{\text{LP}}]} \pi_{\text{LP}}(C_{\text{C}}, C_{\text{LP}}, S) \tag{14}$$

$$S^*(C_{\text{C}}, C_{\text{LP}}) = \text{equilibrium zero expected profits} \tag{15}$$

The utility (profit) functions are continuous because coverage$(\cdot)$ is concave (Eq. 2), $\gamma(\cdot, \cdot)$ is bounded and continuous (Eq. 8), and expectations are taken over compact support. For any fixed choices of the other stakeholders, each stakeholder's optimization problem has a unique solution that varies continuously with those fixed choices; hence the best-response mappings in (13)–(15) are single-valued and continuous. The joint best-response operator is therefore a continuous self-map on a convex, compact strategy set. By Brouwer's fixed point theorem, the joint best response mapping

$$(C_{\text{C}}, C_{\text{LP}}, S) \mapsto (C_{\text{C}}^*(C_{\text{LP}}, S),\ C_{\text{LP}}^*(C_{\text{C}}, S),\ S^*(C_{\text{C}}, C_{\text{LP}}))$$

has a fixed point $(C_{\text{C}}^*,\ C_{\text{LP}}^*,\ S^*)$, which constitutes the Nash equilibrium [29].

In equilibrium, the protocol chooses collateral such that the marginal utility of additional coverage equals its opportunity cost; LPs supply capital until expected returns match the market rate plus risk premium; and speculators arbitrage HACK token prices until they equal risk-neutral hack probabilities. No stakeholder can unilaterally deviate to improve its payoff.

Proposition 1: Truthful Risk Assessment

Under competitive markets with risk-neutral speculators, equilibrium HACK token prices reveal true risk-neutral probabilities of a material exploit, ensuring incentive compatibility in the insurance pricing layer.

Proof. Consider a speculator with subjective belief $\tilde{p}$ about the probability of a hack before expiry T. Suppose the market price of a HACK token implies a risk-neutral probability p_{market}.

- If $\tilde{p} > p_{\text{market}}$, the speculator buys HACK tokens, increasing demand and pushing the price upward.
- If $\tilde{p} < p_{\text{market}}$, the speculator sells or abstains, reducing demand and driving the price downward.

In equilibrium, with free entry and sufficient participation, marginal traders become indifferent between buying and selling, so $\tilde{p} = p_{\text{market}}$. Hence, p_{market} converges to the true risk-neutral probability of a hack event. Equilibrium HACK prices therefore truthfully aggregate information and provide incentive-compatible insurance pricing signals to the yield-share mechanism.

Analytical Propositions

Proposition 2: LP Dynamics and Participation Bounds

(i) Participation bound. Liquidity providers participate only if the yield-share satisfies

$$\gamma(U, P_{\text{risk}}) \geq \gamma_{\min}, \tag{16}$$

where $\rho_{\text{LP}} > 0$ denotes the minimum risk premium LPs require for bearing irreducible cybersecurity exposure:

$$\gamma_{\min} = \frac{C_{\text{LP}} \cdot (r_{\text{market}} + \rho_{\text{LP}}) + p_{\text{hack}} \cdot \mathbb{E}[\min(\text{coverage}, \text{Loss})]}{(1 - \varphi)\, r_{\text{pool}}\, C_{\text{total}}}. \tag{17}$$

(ii) Self-stabilization. If LP capital adjusts according to

$$\frac{\partial C_{\text{LP}}}{\partial t} = \kappa\big[r_{\text{LP}}(U) - (r_{\text{market}} + \rho_{\text{LP}})\big], \qquad \kappa_{LP} > 0, \tag{18}$$

then utilization U converges to a stable equilibrium U^* such that

$$r_{\text{LP}}(U^*) = r_{\text{market}} + \rho_{\text{LP}}. \tag{19}$$

Proof ((i) Participation Bound). Define the LP's realized return as

$$r_{\text{LP}} = \frac{\gamma(U, P_{\text{risk}})\,(1 - \varphi)\, Y_{\text{total}}}{C_{\text{LP}}} - \frac{p_{\text{hack}}\, \mathbb{E}[\min(\text{coverage}, \text{Loss})]}{C_{\text{LP}}}. \tag{20}$$

Substituting $Y_{\text{total}} = r_{\text{pool}} C_{\text{total}}$, imposing $r_{\text{LP}} \geq r_{\text{market}} + \rho_{\text{LP}}$, and solving for $\gamma(U, P_{\text{risk}})$ yields Eq. (17).

Corollary (minimum pool yield). Rearranging Eq. (17) gives

$$r_{\text{pool}} \geq \frac{C_{\text{LP}} \cdot (r_{\text{market}} + \rho_{\text{LP}}) + p_{\text{hack}}\, \mathbb{E}[\min(\text{coverage}, \text{Loss})]}{\gamma(U, P_{\text{risk}})\,(1 - \varphi)\, C_{\text{total}}}. \tag{21}$$

Hence, the pool must outperform the market yield whenever γ is low or expected hack losses are high, while yield parity suffices when utilization and γ remain moderate.

Proposition 3: Sustainable Undercapitalization Bounds

The system remains solvent with probability at least $1 - \varepsilon$ if utilization satisfies the prudential bound:

$$U \leq U_{\max}(t), \tag{22}$$

where $U_{\max}(t)$ is defined in Eq. (5) as a decreasing function of the market-implied annualized hack probability $\widehat{p}_{1Y}(t)$.

Proof. Solvency requires that available LP capital covers realized losses in a hack event with probability at least $1 - \varepsilon$:

$$\mathbb{P}(C_{\text{LP}} \geq \text{Loss} \cdot \mathbb{1}_{\text{hack}}) \geq 1 - \varepsilon. \tag{23}$$

Since $\text{Loss} = \min(\text{coverage}, L_{\text{true}} \cdot \text{TVL})$, where L_{true} is the realized fractional loss, we require

$$\mathbb{P}\big(C_{\text{LP}} \geq \min\big(\text{coverage}, L_{\text{true}} \cdot \text{TVL}\big)\big) \geq 1 - \varepsilon. \tag{24}$$

When realized losses exceed the coverage limit ($L_{\text{true}} \cdot \text{TVL} > \text{coverage}$), the constraint simplifies to

$$\mathbb{P}(U \cdot L_{\text{true}} \leq 1) \geq 1 - \varepsilon. \tag{25}$$

For conservative regimes ($U \leq 1$), solvency holds trivially; only leveraged states $U > 1$ require explicit bounds. By definition of the prudential cap $U_{\max}(t)$, which is calibrated as a monotone decreasing function of the market-implied annualized probability of a hack $\widehat{p}_{1Y}(t)$, solvency is guaranteed whenever $U \leq U_{\max}(t)$, completing the proof.

Corollary (arbitrage-free equilibrium). Together with Proposition 1 (Truthful Risk Assessment) and Proposition 2(i) (LP participation), the mechanism admits an arbitrage-free, incentive-compatible equilibrium across all three stakeholders.

The mechanism is robust across a wide parameter range (stress test, not calibration)

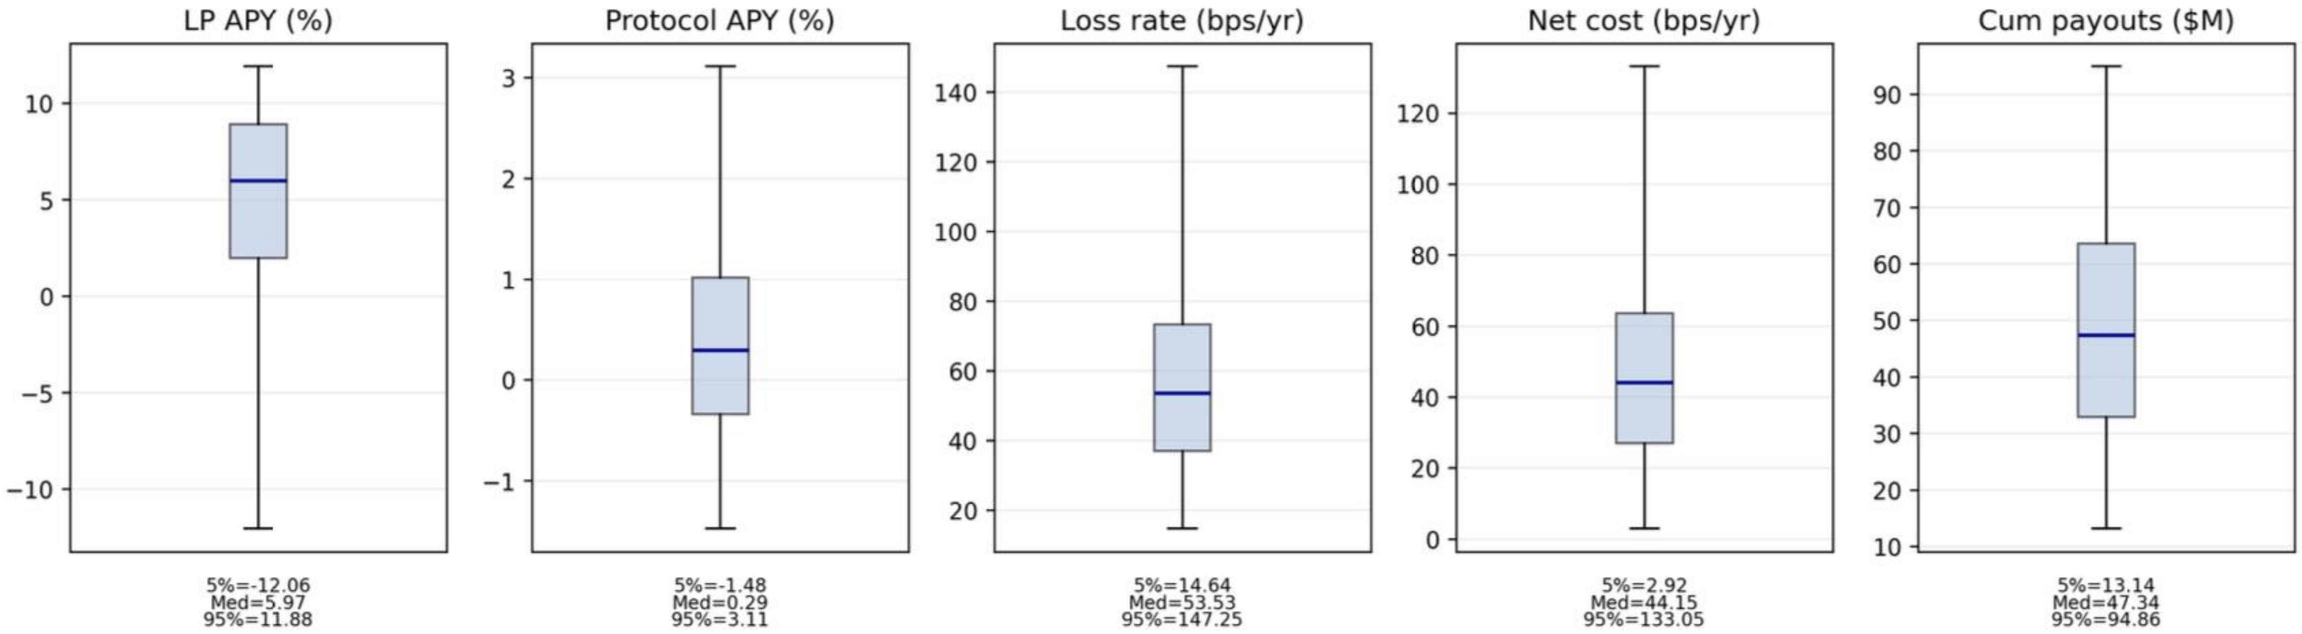

Fig. 2. Run-level distributions (5–95% intervals) for LP APY, protocol APY, empirical loss rates, net protocol cost per $M·year, and cumulative payouts.

Stylized simulations confirm stable coverage and capital dynamics

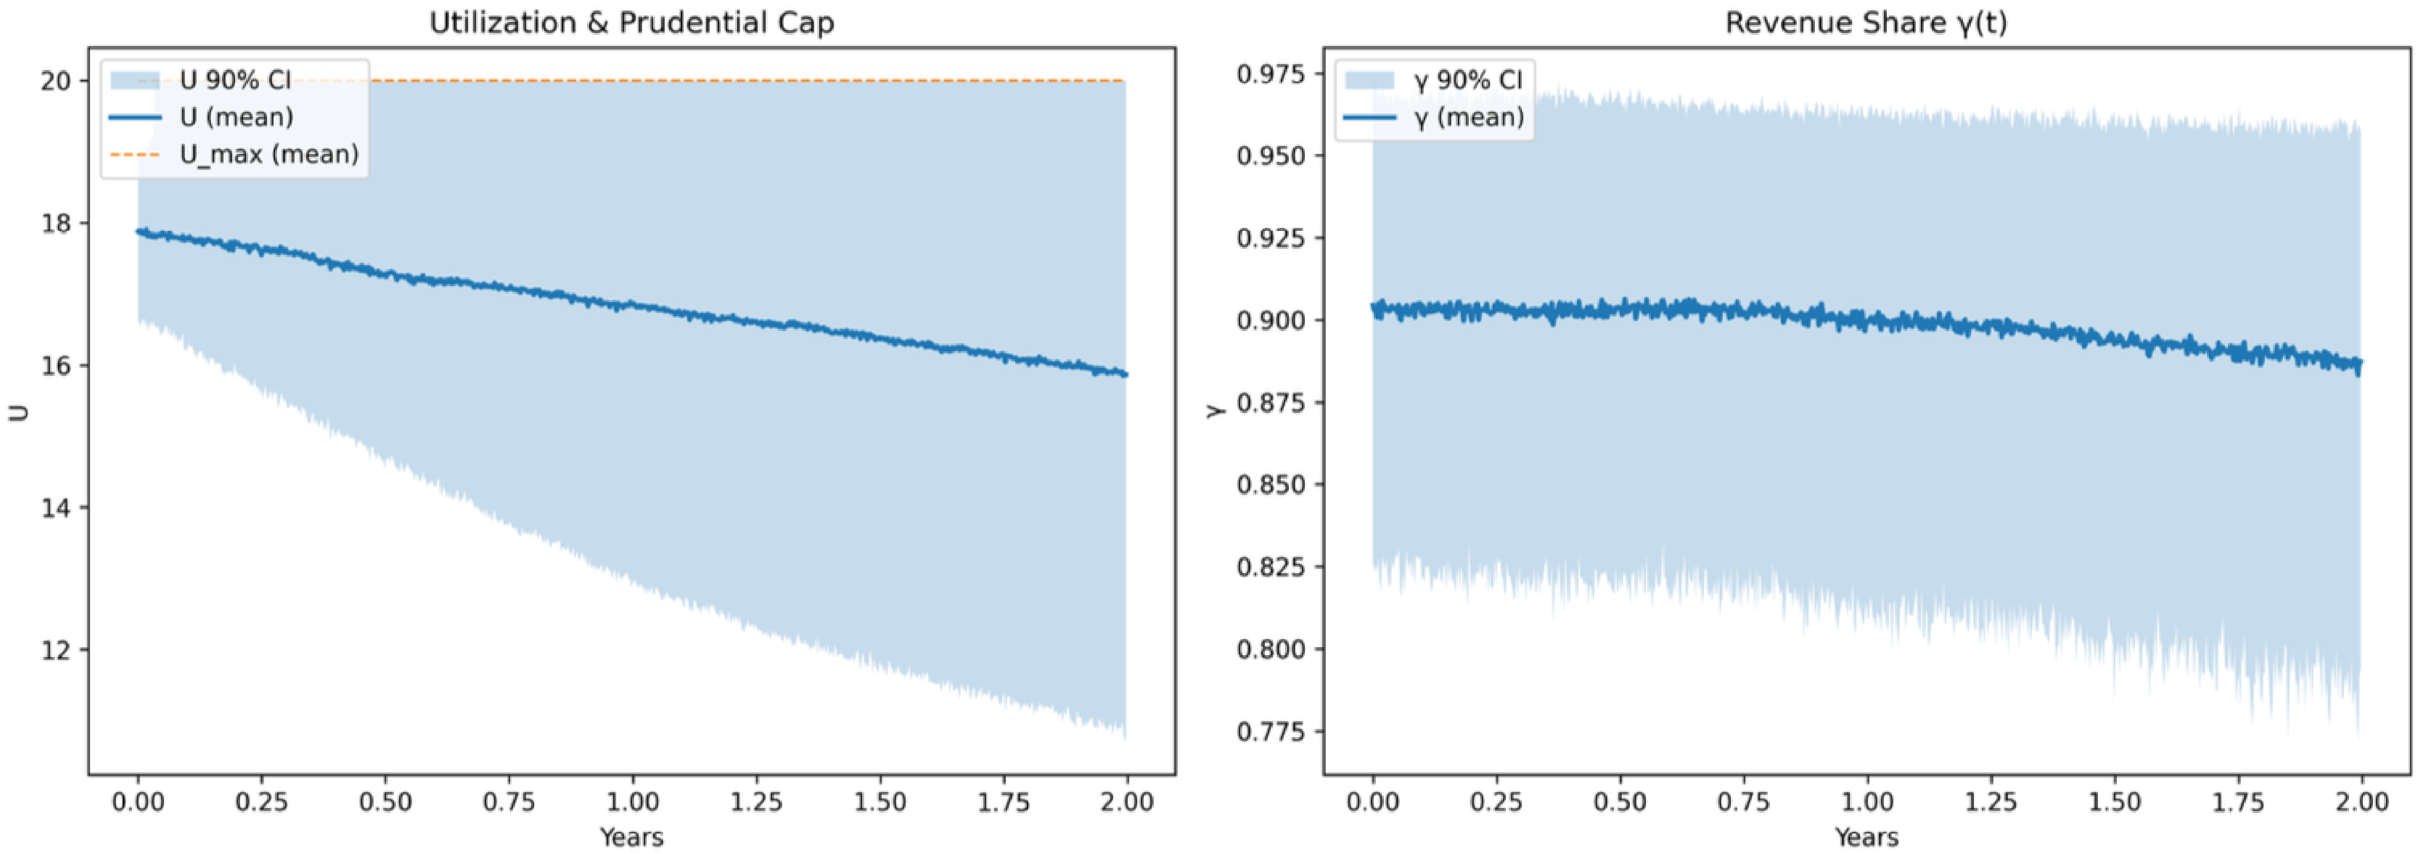

Fig. 3. Simulation dynamics (mean with 90% bands) across 1000 runs. Top-left: utilization vs. prudential cap. Top-right: simulated yield-share γ_t. Bottom-left: LP capital and aggregate posted collateral. Bottom-right: cumulative hack payouts.

Stylized simulations confirm stable coverage and capital dynamics

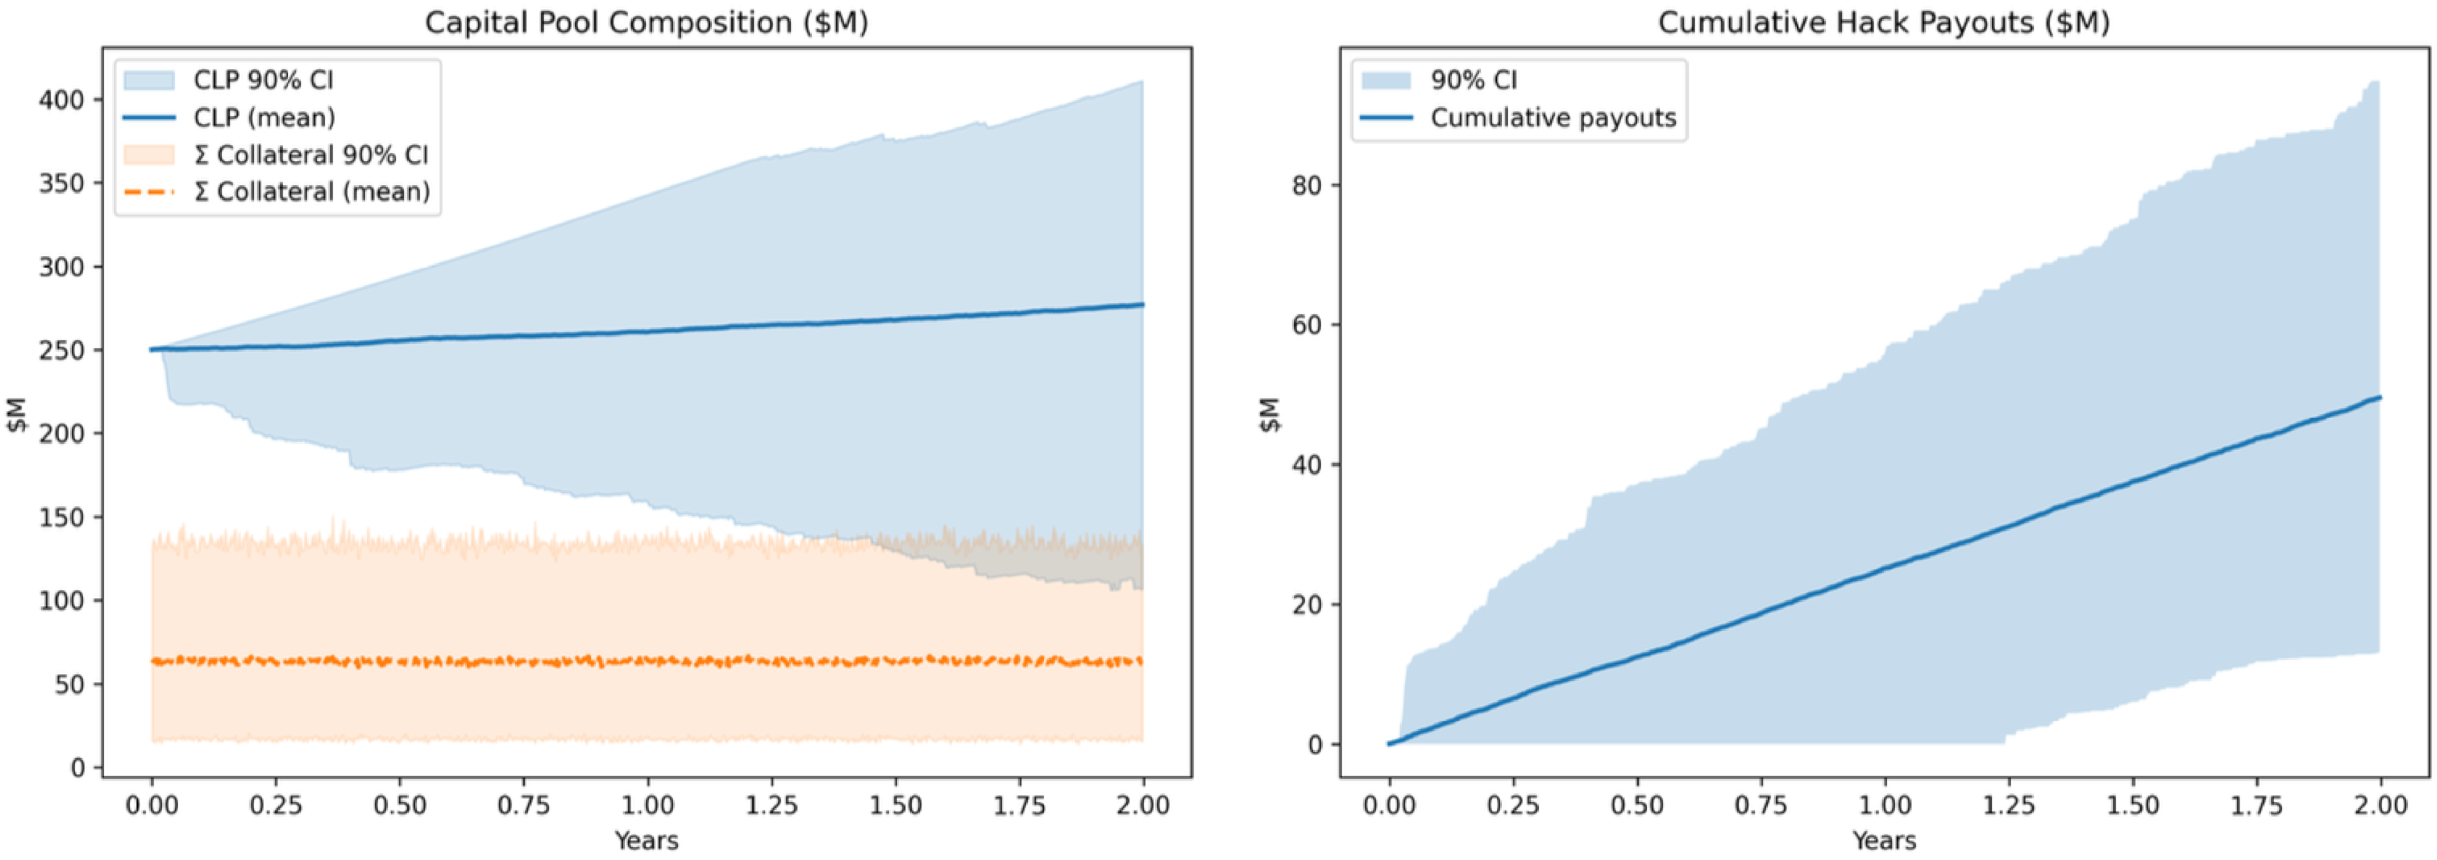

Fig. 3. Simulation dynamics (mean with 90% bands) across 1000 runs. Top-left: utilization vs. prudential cap. Top-right: simulated yield-share γ_t. Bottom-left: LP capital and aggregate posted collateral. Bottom-right: cumulative hack payouts.

Table 2. Baseline simulation parameters used in Section 6.

Parameter	Description	Value
n_{days}	Horizon (daily steps)	2×365
$n_{\text{protocols}}$	Number of protocols	500
$n_{\text{mc\_runs}}$	Monte–Carlo runs	1000
C^0_{LP}	Initial LP capital	\$250M
μ	Coverage scale (Eq. (2))	5.0
θ	Coverage concavity (Eq. (2))	0.10
φ	Operator fee on pool yield	0.01
$U_{\min}$	Lower bound for utilization cap	1.0
$U_{\max}$	Fixed prudential cap in simulation	20.0
κ_U	Prudential cap sensitivity	0.0
α	Weight on utilization in γ	0.6
β	Utilization convexity in γ	1.0
δ	Risk-price convexity in γ	0.7
U_{target}	Target utilization	10.0
r_{market}	External benchmark return	5% p.a.
r_{pool}	Pool return	10% p.a.
ρ_{LP}	LP risk premium	0.5% p.a.
κ_{LP}	LP capital adjustment speed	2.0
$\text{fee}^{\text{base}}_{\text{annual}}$	Base speculator fee (annualized)	3% p.a.
$\text{fee}^{\text{hack}}_{\text{jump}}$	Additional fee on hack day (annualized)	10% p.a.
ω	Term-structure weights $(T_1, \dots, T_4)$	$(0.40, 0.30, 0.20, 0.10)$

Protocol Population Figure 4 shows the distribution of TVL, risk aversion parameters, and security multipliers for the 500 protocols used in the simulation. The TVL distribution is heavy-tailed with median \$8.77M and maximum \$939.86M, closely matching empirical DeFi concentration patterns (top 10% hold 52% of TVL).

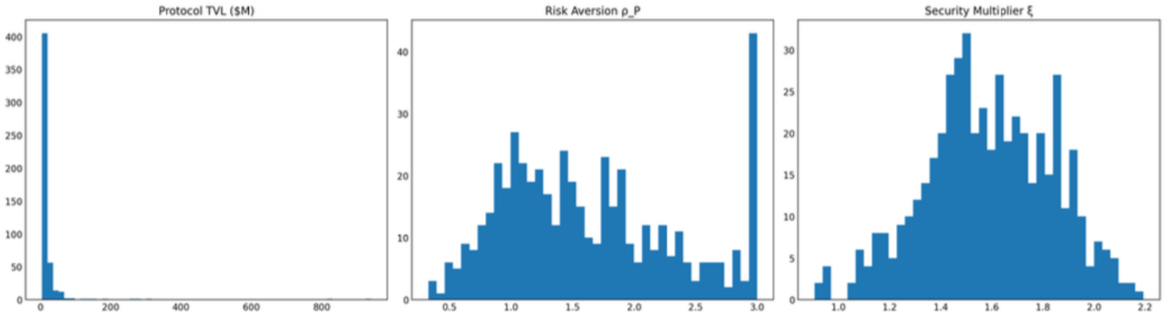

Fig. 4. Distribution of protocol characteristics used in the simulation (TVL, risk aversion, security multipliers).

Policy Bound: Dynamic leverage ceiling

The mechanism enforces a dynamic leverage ceiling linked to the *market-implied annualized hack probability* from the insurance pricing layer. Let quarterly expiries be $T_i \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\}$ (in years) with observed HACK prices $P_{\text{HACK}}(T_i, t)$. Estimate an annualized constant-hazard rate $\widehat{\lambda}(t)$ by least squares:

$$\widehat{\lambda}(t) = \arg\min_{\lambda \geq 0} \sum_i \left(1 - e^{-\lambda T_i} - P_{\text{HACK}}(T_i, t)\right)^2, \tag{5}$$

and define the annualized implied probability $\widehat{p}_{1\text{Y}}(t) = 1 - e^{-\widehat{\lambda}(t)}$. The prudential cap then tightens as implied risk increases:

$$U_{\max}(t) = U_{\min} + \frac{U_{\max}^{\text{hi}} - U_{\min}}{1 + \kappa\, \widehat{p}_{1\text{Y}}(t)}, \qquad U_{\min} < U_{\max}^{\text{hi}},\ \kappa_U > 0. \tag{6}$$

Pool-Level Solvency vs. Protocol-Level Pricing

- **Key distinction**
  - **Solvency is global:** Liquidity providers face a single pool balance sheet.
  - **Pricing is local:** Each protocol is priced based on its own risk.
- **Utilization decomposition**
  - Each protocol occupies a slice of pool capacity:
  - Protocol utilization = coverage provided ÷ total LP capital
  - Total pool utilization is the sum of all protocol slices.
- **Prudential constraint**
  - Aggregate utilization must remain below a dynamic maximum.
  - The maximum utilization tightens when the pool-level risk signal increases.
- **Interpretation**
  - Riskier protocols face higher yield-sharing (i.e., more expensive insurance).
  - If aggregate exposure becomes too large, leverage tightens for the entire pool.
  - This separates: **who pays more** (pricing) from **how much risk the pool can bear** (solvency).

Designing the market mechanism involves a trilemma

Explicit Compensation for Irreducible Risk

Nadler et al. (2022). DeFi Risk Transfer: Towards A Fully Decentralized Insurance Protocol

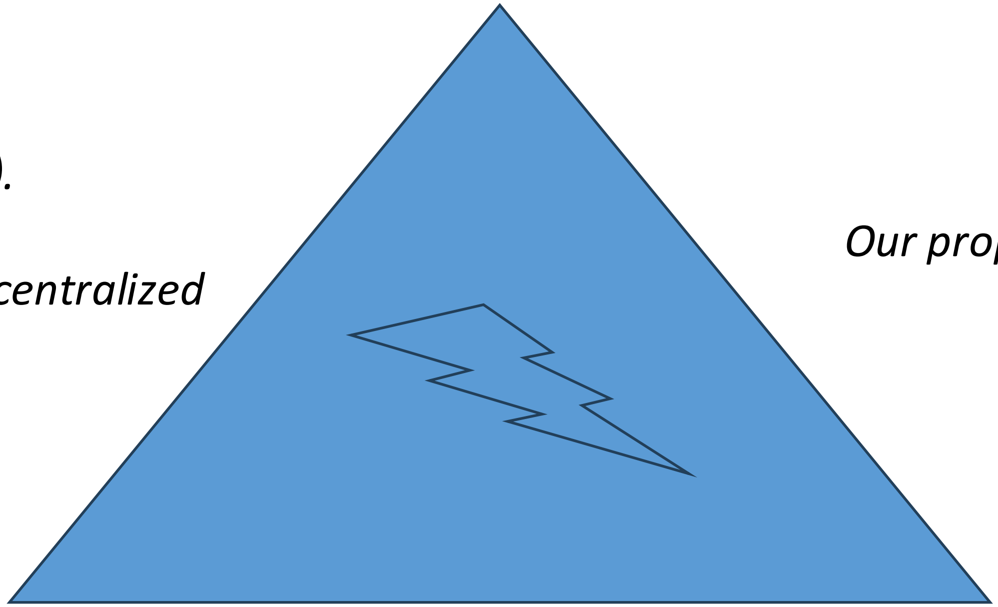

Our proposed mechanism

Trust-Minimized Loss Determination

Scalable, Capital-Efficient Risk Pooling

DeFi cybersecurity insurance design is constrained by irreducible trade-offs

- **Trust-minimized loss determination** vs. discretionary management

→ *Example:* Nexus Mutual requires humans to vote on claims → slow, manipulable.
Full automation: cannot handle complex hacks like partial exploits or MEV attacks.

- **Explicit compensation for irreducible risk** vs. mutual / ex-post pricing

→ *Paying insurers explicitly upfront is incompatible with purely mutual, after-the-fact loss sharing*

- **Scalable, capital-efficient pooling** vs. pairwise designs

→ *Using one capital pool to insure many protocols does not work with bespoke, pairwise contracts*

→ ***Any viable mechanism must choose where to sit in this design space***

We target protocol-level insurance with market-based risk pricing

- **Shift the insured entity**: insure entities (protocols) rather than individual users
- **Internalize moral hazard**: insured protocols post collateral
- **Externalize information aggregation**: markets provide forward-looking risk prices
- **Irreducible cyber risk is pooled:** Liquidity providers (LP) supply capital when underwriting risk offers competitive expected returns.

→ At a high level: protocols post collateral, LPs provide capital, and speculators price cybersecurity risk.