In-Class Exercise: Smart Contract Business Models

Exercise 1: Structured Debate — "Should the Audit-Contest Model Be Treated as Insurance?"

Format: Split into two teams. Each team prepares arguments for its assigned position, then presents. After both sides speak, the class votes — but first, read the debrief questions.

Team A — "It Is Insurance and Should Be Regulated as Such"

Anchoring observation: An audit-contest marketplace pairs a competitive-review fee with a separately staked pool that pays out covered exploits. The pool transfers risk from protocol-team balance sheets onto third-party capital providers in exchange for a premium-like inflow. That is the economic substance of insurance, regardless of how the contract is labelled.

	Team A: Treat It as Insurance
Argument I	
Argument II	
Argument III	
Concession	*Strongest argument AGAINST your position:*
Closing	*How you address the concession:*

Team B — "It Is a Marketplace Service, Not Insurance"

Anchoring observation: The vendor only operates the platform that matches reviewers to protocol teams and intermediates the staking pool. The reviewers carry no balance-sheet risk, the platform itself does not write the cover, and the pool participants choose freely whether to expose their own capital. This is a marketplace plus a peer-to-peer risk-sharing arrangement, not the underwriting of an insurance contract.

	Team B: It Is a Marketplace, Not Insurance
Argument I	
Argument II	
Argument III	
Concession	*Strongest argument AGAINST your position:*
Closing	*How you address the concession:*

Debrief Questions

Q1: Does the answer — insurance or marketplace — matter for how supervisors should oversee the vendor? Why or why not?

Q2: Could the answer genuinely be "both"? If so, what does that imply about the usefulness of inherited regulatory categories for on-chain risk products?

Q3: Name another sector (in finance or beyond) where a platform vendor and a peer-to-peer risk pool blur a similar regulatory boundary. What tensions did that blurring create?

Exercise 2: Audit-Stack Value Chain Mapping

Scenario: The pre-deployment value chain for a smart contract has six links. Different vendors anchor at different links. Your task: for each link, name a vendor type, describe the friction the vendor removes, and predict the long-term outcome.

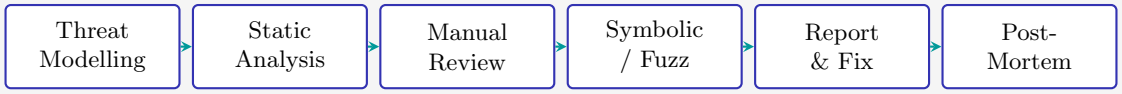

Value Chain Link	Vendor Anchoring It	Friction Removed	Replaces or Improves?	Internal Team Loses or Adapts?
Threat Modelling				
Static Analysis				
Manual Review				
Symbolic / Fuzz				
Report & Fix				
Post-Mortem				

Synthesis Question

Q1: Which link in the audit chain is *most vulnerable* to displacement by an automated tool? Which is *most resistant*? Defend your reasoning with reference to switching costs, reputational rent, and the cost of being wrong.

Facilitator Solutions

Sample answers for instructor reference. These are illustrative; student reasoning may diverge and still be valid.

Exercise 1: Debate Sample Answers

Team A (Treat It as Insurance) — sample arguments

Argument I. The pool transfers loss risk from a protocol-team balance sheet onto third-party capital providers in exchange for a premium-like flow. That is the canonical economic definition of insurance: pre-funded loss-sharing across a pool against a defined event. Calling the same arrangement by another name does not change its supervisory implications.

Argument II. Cover buyers rely on the staked pool's solvency in the same way insurance customers rely on a regulated insurer's reserves. If the pool is undercapitalised relative to the cover written, end users carry a hidden counterparty risk that no consumer-protection regime currently addresses. Insurance regulation exists precisely to detect and constrain that gap.

Argument III. Cross-jurisdictional consistency requires it: an offshore vendor offering a "staking pool" that pays out on a defined trigger to onshore users is offering an insurance product into the onshore market. Treating it as a marketplace would create a regulatory loophole that conventional insurers do not enjoy and would distort competition between on-chain and off-chain cover providers.

Concession. The strongest argument against Team A is that the vendor itself does not underwrite. The cover capital comes from independent participants who choose to stake voluntarily. Forcing the vendor into an insurance regime designed for balance-sheet underwriters may be a category error.

Closing. Regulatory form should follow economic substance. The end user receives loss compensation from a pre-funded pool against a defined event; that is insurance. The vendor's role as platform operator does not exempt the product from supervision – it changes the supervisory toolkit, not the conclusion.

Team B (It Is a Marketplace, Not Insurance) — sample arguments

Argument I. The vendor is a platform operator that matches reviewers to protocol teams and supports a separate, optional, peer-to-peer staking pool. The pool participants are sophisticated, voluntarily expose their own capital, and earn a fee for doing so. This is structurally a marketplace plus a peer-to-peer risk-sharing arrangement, not the underwriting of an insurance policy.

Argument II. The reviewers earn discovery rewards through a competitive contest, not commissions on policy sales. The cover stakers earn a share of the contest fee, not premium income. Both compensation logics are platform-economics primitives – listing fees and revenue sharing – not insurance underwriting.

Argument III. Forcing the vendor into a traditional insurance regime would impose capital and licensing burdens that match a balance-sheet underwriter, not a software platform that intermediates voluntary participants. That would either kill the model or push it into jurisdictions with lighter regimes – weakening, not strengthening, end-user protection.

Concession. The strongest argument against Team B is that end users perceive the product as protection against loss, regardless of how its mechanics are described. If users rely on a cover payout the way they would rely on insurance, supervisory attention is hard to justify withholding.

Closing. A new product category needs a new supervisory framework, not a forced fit into the inherited insurance regime. Disclosure obligations, capital adequacy of the pool, and conduct rules for the platform are all reasonable – but applying twentieth-century insurance regulation

wholesale to a peer-to-peer risk pool is more likely to drive the activity offshore than to protect the user.

Debrief Q1 — Supervisory implications

The classification matters because the supervisory toolkit is calibrated to the risks the regulator expects. Insurance supervision focuses on reserve adequacy, claim-paying ability, conduct-of-sale rules, and resolution regimes. Marketplace supervision focuses on conflicts of interest, disclosure, market abuse, and anti-money-laundering compliance. If the vendor sells loss compensation that users rely on for protection, the gap between the two regimes leaves real consumer-protection holes – pool insolvency risk, claim-handling disputes, and cross-border enforcement difficulties. The answer matters because it determines which holes get filled.

Debrief Q2 — "Both" as an answer

The answer can genuinely be "both". The vendor operates a software marketplace and intermediates an arrangement that is functionally an insurance product. That duality reveals the limits of regulatory categories that assume a single firm controls both the platform and the balance sheet. If "both" is correct, supervisors need a hybrid framework: marketplace conduct rules for the platform, plus prudential-style oversight of the pool and consumer-protection rules for the cover product. Few jurisdictions have built such a framework yet – which is itself the regulatory-arbitrage opportunity that some vendors exploit.

Debrief Q3 — Cross-sector example

Online platform-based lending illustrates the same blurring. Early peer-to-peer lenders described themselves as marketplaces matching individual lenders to individual borrowers; supervisors eventually concluded that, for retail participants, the substance was credit intermediation requiring banking-style oversight. The parallel to audit-and-cover marketplaces is direct: a platform-economics framing collides with a substantive financial-services activity, and the question is whether to grow a new framework or to extend an old one. In peer-to-peer lending the answer mostly became "extend the old one," which reshaped the industry. The same path is plausible for audit-and-cover platforms.

Exercise 2: Audit-Stack Value-Chain Mapping Sample Answers

Value Chain Link	Vendor Anchoring It	Friction Removed	Replaces or Improves?	Internal Team Loses or Adapts?
Threat Modelling	Specialist consultancy	Lack of attacker-mindset experience inside the protocol team	Improves	Adapts
Static Analysis	Tool vendor (open-source plus paid)	Manual scanning for known vulnerability patterns	Replaces	Adapts
Manual Review	Trail of Bits-style firm	Reputational and depth deficit on a senior-engineer review	Replaces (for the named report)	Loses
Symbolic / Fuzz	Certora-style verification engine	Coverage gap that ad-hoc test suites cannot close	Improves	Adapts
Report & Fix	Audit firm plus client engineering	Time and credibility cost of writing a public report	Improves	Adapts
Post-Mortem	Specialist incident-response provider	Lack of a structured lessons-learned playbook	Improves	Adapts

Synthesis Question Sample Answer

The most vulnerable link is Static Analysis. Open-source scanners and continuous-integration plug-ins compound rapidly: every new vulnerability pattern becomes a community-contributed rule, and the marginal cost of running one more scan is near zero. A vendor that sells nothing more than a paid scanner has very thin defensibility, because the same checks become free shortly after they are introduced. The most resistant link is Manual Review by a top-tier firm. The deliverable is not a list of bugs but a signed report that downstream parties – token-listing committees, integration partners, institutional users – treat as credible. That credibility is reputational rent that compounds across engagements, is difficult to copy, and is intentionally rationed because it depends on a small senior-engineer roster. Switching costs are high because the signature itself is the product, and the cost of being wrong on a major engagement is severe. Internal teams can adopt automated tools but cannot manufacture reputation in the same way.