

In-Class Exercise: Sanctions Screening Business Models

Exercise 1: Structured Debate — “Is Napier a Compliance Vendor or a Risk-Insurance Provider?”

Format: Split into two teams. Each team prepares three arguments for its assigned position, one concession acknowledging the strongest counter-argument, and a closing that addresses the concession. Then the class votes after both sides present.

Team A — “Napier Is a Compliance Vendor”

Anchoring evidence: Napier’s core deliverable is an explainable-AI overlay that sits on top of the bank’s legacy list pipeline, producing audit-ready reasoning artefacts that the bank shows to its supervisor. Its per-lookup overlay fee, model-risk deliverables, and attestation packages are consumed by bank compliance teams who themselves remain accountable for the final screening decisions. The overlay is the product.

Team A: Napier Is a Compliance Vendor

Argument I

Argument II

Argument III

Concession *Strongest argument AGAINST your position:*

Closing *How you address the concession:*

Team B — “Napier Is a Risk-Insurance Provider”

Anchoring evidence: Napier prices its overlay on a per-lookup basis bundled with a suppression budget that effectively absorbs a portion of the bank’s false-positive cost. That pricing move moves the cost of adjudicated-but-true-negative alerts from the bank’s operations desk onto the vendor’s income statement, which is the economic structure of a risk-insurance contract rather than a tooling subscription.

Team B: Napier Is a Risk-Insurance Provider

Argument I

Argument II

Argument III

Concession *Strongest argument AGAINST your position:*


Closing *How you address the concession:*

Debrief Questions

- Q1:** Does the answer — compliance vendor or risk-insurance provider — matter for how regulators should supervise Napier? Why or why not?
- Q2:** Could the answer genuinely be “both”? If so, what does that imply about how to classify hybrid tooling-plus-loss-sharing businesses?
- Q3:** Name another business (in any sector) that blurs the line between “sells a tool” and “absorbs the risk that the tool creates”. What tensions does that blurring create?

Exercise 2: Screening Value Chain Mapping

Scenario: The sanctions-screening value chain can be broken into six links. Modern screening vendors attack specific links, and some reposition the vendor inside a link the incumbent was taking for granted. Your task: for each link, identify a vendor from the reference slate, describe the value the vendor creates, and classify the outcome.



Value Chain Link	Vendor tacking It	At-	Value Created	Repositions or Adds?	Moat or Margin Cap?
List Aggregation					
Matching					
Scoring					
Alert Generation					
Adjudication					
Attestation					

Synthesis Questions

- Q1:** Which link in the screening value chain creates the strongest moat for a vendor, and which link most clearly caps its margin? Defend both answers with reference to label accumulation, regulatory barriers, and data ownership.
- Q2:** For at least two of the five vendors on the reference slate, identify a link where the vendor rents rather than owns the underlying function. What would it take for the vendor to convert that rented link into an owned one, and should it?

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 (Napier Is a Compliance Vendor) — sample arguments

Argument I. Napier's core deliverable to each bank is an explainable reasoning artefact that accompanies a sanctions decision, not a contractual absorption of credit or operational loss. The Business Model Canvas reads as a B2B compliance-tooling vendor: the Key Activities are overlay retraining, explanation generation and attestation packaging; the Key Resources are the overlay model, the labelled-outcome history and the supervisory-aligned evidence templates; the Revenue Streams are per-lookup licence fees bundled with an explainability premium. Those blocks describe a vendor, not an insurer.

Argument II. Napier sells into a buyer who remains accountable for the screening decision — the bank's compliance function. The bank's head of financial crime, not Napier, signs the sanctions report and owns the examiner conversation. That division of responsibility is the hallmark of a vendor relationship: Napier supplies the reasoning; the bank owns the outcome. A risk-insurance provider would own the outcome directly.

Argument III. Napier's competitive advantage resides in its model IP, its explanation templates and its supervisory relationships — all assets typical of a compliance-tooling platform. None of those advantages are balance-sheet advantages. A genuine risk-insurance provider's competitive advantage would include underwriting capital, loss reserving and reinsurance capacity, none of which Napier competes on primarily.

Concession. The strongest argument against Team A is that Napier does bundle a suppression budget with its overlay fee, which contractually absorbs a portion of the bank's false-positive cost and complicates a pure vendor classification.

Closing. The suppression budget is best understood as an incentive-alignment mechanism, not as a balance-sheet insurance product. Its purpose is to give the vendor skin in the game on suppression quality, so that the overlay's accuracy matches what the bank needs. That is the structural feature of an outcome-aligned vendor contract, not of an insurance policy. The vendor classification therefore describes the economic substance.

Team B (Napier Is a Risk-Insurance Provider) — sample arguments

Argument I. Napier carries contractual exposure on the false-positive cost that flows through the overlay, through the suppression-budget construction in its pricing. That exposure is a liability position on a probabilistic outcome, not a software licence. A firm whose profit and loss moves with the noise of its buyer's alert stream is an insurer by economic substance, regardless of how its Canvas boxes are labelled.

Argument II. The revenue mix at Napier is shaped by the suppression budget rather than by pure tooling fees: the stronger the overlay's suppression, the more economic margin the vendor retains; the weaker the suppression, the more the vendor absorbs. That asymmetric reward structure is the defining shape of a risk-taking business, not of a subscription software firm.

Argument III. Regulators have begun to scrutinise bundled-suppression arrangements under outsourcing and operational-resilience frameworks, asking who truly bears the residual risk. The supervisory treatment is calibrated to where the loss actually sits rather than to what the contract is called. If the regulator sees a firm that bears part of the bank's operational loss, the economic substance is an insurer.

Concession. The strongest argument against Team B is that Napier does not hold underwriting capital, is not capitalised as an insurer, and earns a material share of revenue from model-explanation and attestation activities that are squarely in the compliance-tooling space.

Closing. The absence of underwriting capital does not change the fact that Napier earns on a risk outcome and bears part of the buyer's operational cost. A firm that carries a slice of its customer's loss is economically a partial insurer — it simply funds that exposure through model margin rather than through a reinsurance tower. Economic substance over Canvas labels: Napier is a risk-insurance provider that delivers its product through an API.

Debrief Q1 — Regulatory supervision

Whether regulators should supervise Napier as a compliance vendor or a risk-insurance provider depends on the risks the firm actually creates rather than on the preferred self-description. If the firm absorbs operational-loss exposure on bank alerts — even partially, through a suppression budget — outsourcing and operational-resilience rules apply to that exposure directly, and the buyer bank remains accountable for the residual risk. If the firm supplies explanations to banks who themselves carry the decision risk, the relevant oversight is model-risk governance on the vendor side and bank-side accountability for the decisions used. In practice Napier occupies both positions simultaneously, which is why supervisors have increasingly applied third-party-risk-management, operational-resilience and model-risk frameworks in parallel. The classification matters because each framework carries different capital, disclosure and examination obligations.

Debrief Q2 — “Both” as an answer

The answer genuinely can be “both”. Napier is a compliance-tooling vendor with a portion of the false-positive cost carried on its own meter and a portion remaining with the bank. That duality reveals the limits of categorical labels inherited from a world where tooling, underwriting, and loss absorption were bundled inside separate institutions. When a firm unbundles tooling, explanation and loss sharing across a vendor contract, the right regulatory response is functional: supervise each function where it resides, regardless of the firm's overall label. “Both” as an answer implies that classification by function is more robust than classification by firm.

Debrief Q3 — Cross-sector blurring example

A close parallel is a managed-security-services firm that sells a cyber-monitoring tool with a service-level agreement that refunds fees (or pays damages) on undetected breaches. The firm sells tooling but contractually absorbs a portion of the operational risk those tools fail to mitigate. Regulators have struggled to classify it: tooling vendor or cyber-insurer. The tension is directly analogous to Napier's: the firm delivers reasoning as software but carries some of the operational risk that decisions informed by that reasoning create. In both cases, functional regulation — regulating the tooling, explanation or loss-sharing function separately from the vendor's overall label — is better-calibrated than categorical classification.

Exercise 2: Screening Value Chain Mapping Sample Answers

Value Link	Chain	Vendor	Value Created	Repositions or Adds?	Moat or Margin Cap?
List Aggregation		Refinitiv World-Check (curated roster with two-sided research feedback)	Turns an aggregation step into a platform whose roster compounds with each lookup edge case	Repositions (aggregation becomes platform)	Moat (research labour compounds across buyers)
Matching		LexisNexis Bridger (fuzzy, phonetic, transliteration and identity-resolution scoring)	Converts a commodity list into a defensible matcher subscription by attacking the matching link directly	Repositions (list stays commodity, matcher becomes product)	Moat (precision-recall improvements hard to replicate)
Scoring		Dow Jones Risk (calibrated verdict as part of the stacked re-bundle)	Produces a scored verdict rather than a raw hit, inside the same buyer and billing cadence as the underlying feed	Adds (new product layer on existing integration)	Moat if suppression quality holds; margin cap if commoditised
Alert Generation		Dow Jones Risk (adjudication toolkit re-bundle)	Routes the scored verdict into a case-management queue already integrated with the bank's compliance workflow	Adds (new product layer on existing integration)	Moat (workflow lock-in inside buyer's perimeter)
Adjudication		SymphonyAI Sensa (alert-ranking overlay)	Re-orders the alert queue so likely-true hits surface first, conditioned on the buyer's own labelled-disposition history	Repositions (ranking becomes a distinct product inside the adjudication step)	Moat at label-rich buyers; margin cap at label-sparse buyers
Attestation		Napier (explainability overlay and regulator-ready packaging)	Wraps every decision in audit-ready reasoning aligned to supervisory expectations	Adds (attestation emerges as its own purchasable artefact)	Moat if supervisory cadence keeps rising; margin cap if regulator ships standard toolkit

Synthesis Question 1 Sample Answer

The link that creates the strongest moat for a screening vendor is Matching. Once a vendor has accumulated labelled-outcome history on fuzzy, phonetic and transliteration variants across many bank integrations, the precision-recall curve of its matcher compounds in a way a new entrant cannot replicate without years of lookup flow. Each additional bank integration adds new aliasing patterns and edge cases that feed back into the matcher, sharpening it further. Regulatory barriers reinforce the moat: the matcher's explanation artefacts, suppression budgets and governance documentation become part of the vendor's asset base, all reusable across the label history. The link most likely to cap margin is List Aggregation when viewed in pure commodity terms. Upstream primary sources — regulators, sanctions-issuing authorities, secondary publishers — retain pricing power, and the vendor is a buyer. As the vendor's coverage scales, its source-licensing bill scales with it, which caps gross margin unless the vendor can compound the list into a platform (as Refinitiv World-Check does through two-sided research dynamics).

Synthesis Question 2 Sample Answer

LexisNexis Bridger rents much of its list input: it consumes regulatory and secondary-publisher feeds as raw material and converts them through its matcher. Converting those rented inputs into owned ones would require Bridger to build its own research desk — effectively replicating what Refinitiv World-Check maintains — so that list coverage could be produced under a publishing relationship Bridger directly owns. The case for converting is gross-margin expansion and strategic independence from secondary publishers; the case against is the operational burden of running a sustained research operation, and the head-start advantage that incumbent publishers already hold in primary-source access. SymphonyAI Sensa rents its training fuel: its ranker is only as strong as the labelled-alert history the buyer brings, and that history is owned by the bank, not by Sensa. Converting that rented fuel into an owned asset would mean negotiating data-sharing agreements that pool adjudicated alerts across many banks into a common training set, which is slower and politically complicated but yields a ranker that would be much harder to displace once the pool achieves critical size. Whether either vendor should convert is a strategic choice between speed-to-scale (renting) and margin-durability (owning); the right answer depends on how long the vendor expects the wedge to remain open before a standard toolkit or an incumbent bundle appears.