

## Post-Class Summary: Sanctions Screening Business Models

### Key Frameworks

#### Business Model Canvas Applied to Screening Vendors

The Business Model Canvas decomposes any venture into nine interlocking elements. Mapping a classical sanctions-list publisher and a modern AI-overlay vendor onto the same canvas shows that most blocks overlap — both sell into banks, both deliver via an API, both anchor on a per-lookup verdict. Three blocks diverge sharply and determine the gross-margin profile: Key Resources (a matcher engine, a labelled-outcome history and supervisory-aligned explanation templates, rather than a roster of names), Key Activities (continuous overlay retraining, threshold tuning and explanation generation, rather than roster curation), and Revenue Streams (per-lookup licence fees bundled with a suppression budget, rather than per-record subscription or access fees). Those three blocks encode the economic difference between a passive data publisher and an active verdict vendor.

#### Platform Economics for Screening

Many modern sanctions-screening vendors operate as multi-sided platforms. Refinitiv World-Check connects research contributors on one side with compliance consumers on the other. Cross-side network effects are the engine: more lookups surface more edge-case aliases, which feed back into research, which deepens the roster, which attracts more compliance consumers, which in turn surface more edge cases. The chicken-and-egg problem is real: a new platform has no roster depth until researchers produce one, but no researchers sign up until consumers are present to make the labour economical. The dominant solution is to carry the research labour internally at launch — staffing a desk that produces immediate coverage — so the platform has enterprise value before any consumer is asked to subscribe.

#### Unbundling–Rebundling Applied to Vendors

Christensen's disruption cycle applies to B2B data vendors as readily as to consumer FinTech. Dow Jones Risk illustrates the pattern most cleanly: it entered with a single narrow wedge (a curated adverse-media and watchlist feed), earned integration trust, and rebundled outward into a matcher, an adjudication toolkit, an entity-graph enrichment product and a regulator-ready attestation package. The rebundling rule is strict: each new product must reuse the same compliance buyer, the same governance stack, and the same per-lookup billing cadence. The cycle is completed when the vendor's product footprint inside the bank's perimeter resembles the full screening stack the bank would otherwise have built in-house.

#### Value Chain Deconstruction of Screening

Evans and Wurster argued that information-rich value chains deconstruct. Sanctions screening is a textbook case: List Aggregation, Matching, Scoring, Alert Generation, Adjudication and Attestation are each contestable independently. The unusual move is that deconstruction can reposition the vendor inside the chain rather than only slicing it into smaller pieces. LexisNexis Bridger attacks Matching specifically, treating the underlying list as a commodity input and converting its edge into the matcher's precision-recall curve. Deconstruction is therefore not only about slicing existing links; it can reshape which link the vendor claims as its defensible product while leaving the neighbouring links with their incumbent owners.

#### Regulatory Arbitrage in the Explainability Gap

Sanctions-screening overlay vendors operate in the gap between supervisory expectation (audit-ready reasoning, model-card documentation, suppression justification, ongoing monitoring) and legacy-list capability (raw hit generation with minimal evidence trail). Napier monetises that gap by selling the explainable-AI overlay together with a suppression budget that writes part of the false-positive cost

back into the fee. Whether the business is a durable compliance moat or a transitional overlay depends on which side of the arms race moves faster. If supervisory cadence continues to rise faster than legacy list pipelines can produce their own evidence trails, the per-lookup premium compounds. If regulators ship standard explainability toolkits or if incumbent list vendors bundle overlays into their existing subscriptions, the premium collapses. The arbitrage test asks directly which trajectory dominates.

## Company Cases Summary

Company	Value Creation Mechanism	Key Framework	What Makes It Different
Refinitiv World-Check	Two-sided data platform connecting research contributors with compliance consumers; per-lookup pricing with research-driven roster compounding	Platform Economics	Research labour carried at launch, producing roster depth before inviting consumers
LexisNexis Bridger	Matcher-centric product that attacks the Matching link directly and leaves list aggregation as a commodity input	Value Chain Deconstruction (repositioning)	Converts a commodity list subscription into a defensible matcher subscription
Dow Jones Risk	Rebundling from a single adverse-media feed into matcher, adjudication, entity graph and attestation on the same buyer and billing rails	Unbundling-Rebundling	Each new wedge must reuse the compliance buyer and per-lookup cadence
Napier	Explainable-AI overlay with a suppression budget that absorbs part of the bank's false-positive cost	Regulatory Arbitrage → Compliance Moat	Monetises the widening gap between supervisory explainability and legacy-list evidence trails
SymphonyAI Sensa	Alert-ranking overlay that re-orders the bank's existing queue using the buyer's labelled-disposition history	Context-Dependent Business Model	Business model is tied to the buyer's label depth, not to raw ranker portability

## The Five-Test Framework

Use these five tests to evaluate any sanctions-screening vendor's strategic position:

- 1. Friction test.** Identify the single friction the vendor removes from the bank's existing process.  
*Application to Refinitiv World-Check:* the friction removed is the bank's inability to keep an in-house roster of sanctioned persons current at the cadence the supervisor now expects, without building a research desk.
- 2. Platform test.** Determine whether the vendor benefits from cross-side network effects between two or more market sides.  
*Application to Refinitiv World-Check:* research contributors and compliance consumers reinforce each other through surfaced aliases, enriched profiles and deepening roster coverage.
- 3. Rebundling test.** Assess whether the vendor is adding adjacent products on the same buyer and the same billing cadence, or stalling on a single wedge.  
*Application to Dow Jones Risk:* the feed, matcher, adjudication, entity graph and attestation layers all reuse the compliance buyer and per-lookup billing rail.

**4. Infrastructure test.** Ask whether the vendor repositions inside an existing link or competes in an already-occupied one.

*Application to LexisNexis Bridger:* attacks Matching rather than List Aggregation, turning the matcher itself into the defensible product and leaving the list as a commodity input.

**5. Arbitrage test.** Evaluate whether the wedge is the gap between supervisory expectation and legacy-list evidence capability, and whether that gap is widening.

*Application to Napier:* the wedge is rising explainability cadence outpacing legacy pipelines; whether it is a moat or a window depends on which side standardises first.

A sixth meta-test applies across all five: the context test (SymphonyAI Sensa). Does the buyer's internal alert plumbing generate enough labelled-disposition history to fuel the overlay? If the history is thin, the business model stalls even when the ranker works.

## Connections to Other Topics

The frameworks above connect directly to several other course themes. The cost-centre-to-margin transformation discussed in the RegTech BM package explores the same compliance-as-moat dynamic from a broader vantage, while the privacy-compliance BM package explores the mirror problem: pricing the absence of data leakage rather than the absence of a false positive. The alternative-credit-scoring BM material examines a closely related two-sided platform pattern — lender partners and origination channels reinforcing each other through labels — where the labelled-outcome dataset plays the role that the labelled-disposition history plays in screening. Finally, the cybersecurity BM material in the risk-management lesson shows the closest structural analogue to Napier's suppression-budget construction: a monitoring vendor that prices its product partly on the losses it prevents, bringing the vendor's incentives closer to the buyer's operational reality.