Sanctions Screening Business Models

Selling a verdict per transaction — and externalising the false-alarm bill onto the buyer

Digital Finance

vs.

The Vendor Stall

Pay per check.
No refunds.

The Back Office

My desk pays
for every false match.

"The vendor sells the lookup. The bank pays for every mistaken one."

Why Does Every Extra Lookup Shrink the Vendor's Risk but Grow the Bank's Cost?

The Externalised-Cost Wedge
A sanctions-screening vendor sells its buyer — a bank, a payment processor, a correspondent house — a single verdict per lookup: hit or clear. The vendor charges a per-lookup fee regardless of the verdict quality. The cost of each false hit is absorbed downstream, inside the bank's alert-triage operations.

- Every additional watch-list entry raises lookup lift but enlarges the surface the bank must adjudicate; every fuzzy-match relaxation tightens coverage and floods the queue.
- The vendor is paid the same whether the match is sharp or noisy; the bank absorbs the adjudication bill.
- The deeper the list, the more the vendor earns; the louder the queue, the more the bank pays.

This is the founding tension of the business: the product is the verdict, but the cost of acting on a verdict is borne by somebody else. Whoever reshapes the fee so the verdict and its downstream cost live on the same meter changes the industry.

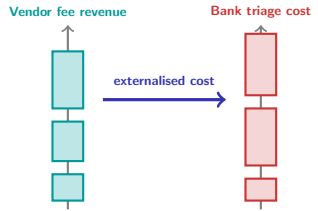

Both columns grow per lookup — but on different meters.

Osterwalder BMC anchor — Value Proposition (verdict per lookup) separates from Cost Structure. The bill for noise sits on the buyer's meter, not the vendor's.

Which Canvas Blocks Separate a List Vendor from an AI-Overlay Vendor?

The Business Model Canvas Applied
Osterwalder's Business Model Canvas separates any venture into nine blocks. Mapping a classical sanctions-list publisher and a modern AI-overlay vendor onto the same canvas shows that most blocks overlap — both sell into the same compliance buyer, both anchor on a verdict per lookup — but three blocks diverge sharply, and those three blocks encode who absorbs the false-alarm bill.

- **Key Resources:** a list vendor owns a curated roster of sanctioned names and associates; an overlay vendor owns a scoring engine, a labelled alert-outcome history, and the adjudication playbooks that shape how banks interpret results.
- **Key Activities:** the list vendor curates, versions and republishes the roster; the overlay vendor retrains the matcher, tunes thresholds, and produces explanation artefacts the bank can show to its supervisor.
- **Revenue Streams:** the list vendor earns per-record subscription or per-lookup fees on the roster; the overlay vendor earns a per-lookup premium bundled with a suppression budget — a portion of the false-positive cost is written back into the fee instead of landing on the bank's desk.

What overlaps: Customer Segments (banks and payment firms), Channels (API into the payment-message pipeline), Customer Relationships (B2B SaaS support), Value Proposition (a verdict per lookup), Key Partners (data brokers, regulator advisories), Cost Structure surface (compliance, engineering, legal). A sanctions vendor is an alert-economics company wearing list-vendor clothes.

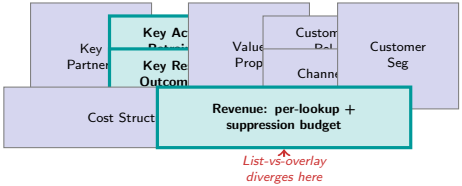

Osterwalder BMC anchor — three canvas blocks shift, six overlap. Outcome labels replace roster entries; suppression budget replaces raw per-lookup.

How Does Refinitiv World-Check Turn a Name List into a Two-Sided Data Platform?

Two-sided platform = a service that sells access to one user group by first attracting another; value rises on both sides as each grows.

The World-Check Platform (Refinitiv World-Check, United Kingdom)

Refinitiv World-Check, headquartered in London under London Stock Exchange Group ownership, did not begin as a platform. It was a research desk that compiled sanctioned-person and politically-exposed-person profiles from open sources. Over time, it evolved into what it really is now: a platform that connects contributing researchers on one side with consuming compliance teams on the other. The central roster sits in the middle and is priced per lookup, regardless of which side originated a given record.

- **Contributor side:** open-source analysts, regulatory filings, court records, and in-house researchers who enrich the roster with new entities, aliases, and linkages.
- **Consumer side:** banks, payment firms, corporate compliance teams, and correspondent houses that query the roster on every transaction or onboarding.
- **Cross-side network effect:** more consumers surface more edge-case aliases and negative matches, which feed back into research; richer research draws more consumers; the roster compounds rather than depreciates.

The chicken-and-egg wedge: Refinitiv World-Check solved it by carrying the research labour itself at launch, producing a roster that had immediate enterprise value before any consumer feedback loop was asked of a bank. Once the roster existed, the two-sided platform could be stood up.

The platform-economics insight: World-Check is not selling a file of names. It is selling the proximity of ongoing research to the consuming compliance decision, and the per-lookup cadence that lets both sides grow together.

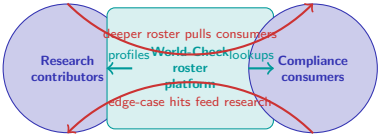

Research labour seeded the first roster.

Platform-economics anchor — World-Check sells the proximity of research to adjudication. Both sides reinforce each other through surfaced aliases and enriched profiles.

How Did Dow Jones Risk Move Its Product Stack from the List Era to the Overlay Era?

Unbundling = pulling one service out of a historical bundle and offering it alone; *rebundling* = stacking adjacent services back onto that foothold once trust is established. Clayton Christensen (Harvard Business School) argued disruptors start narrow and cheap, earn trust, then expand upward.

Dow Jones Risk's Stack Shift (Dow Jones Risk & Compliance, United States)
Dow Jones Risk & Compliance, headquartered in New York, entered the market as a classical list publisher: a curated adverse-media and watchlist feed that a bank could run against its payment messages. From that narrow wedge it has rebundled outward, with each adjacent product chosen so that it reuses the same compliance buyer, the same governance stack, and the same per-lookup billing shape. The product stack has slid upward into AI-overlay territory without breaking the API contract that the bank's integration team committed to years ago.

- **Core wedge (list era):** the curated adverse-media and watchlist feed delivered as a batch or streaming pipe.
- **Matcher wedge:** a fuzzy-match scoring engine that reads the same feed but delivers a calibrated verdict, not a raw hit.
- **Adjudication wedge:** alert-triage tooling that turns a hit into a disposition, with a suppression budget written back into the price.
- **Risk-intelligence wedge:** enrichment content — entity graphs, ultimate-beneficial-owner lineage, adverse-media excerpts — that the analyst reads next to the hit.
- **Attestation wedge:** package the adjudicated outcomes as regulator-ready attestations and model-card artefacts.

The rebundling logic is strict: each new product reuses the same buyer, governance rails, and billing cadence. The wedge becomes the rail for every subsequent product.

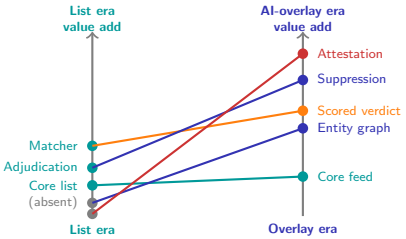

Christensen's unbundling-rebundling cycle — the wedge that cleared the first bank integration becomes the rail for every adjacent product. The slopegraph shows the shift.

Where in the Screening Value Chain Does LexisNexis Bridger Insert Itself?

The Screening Value Chain (LexisNexis Bridger, United States)
Evans and Wurster argued that information-rich value chains deconstruct — each link can be contested independently by a specialist. LexisNexis Bridger, part of RELX's LexisNexis Risk Solutions and operated from North Carolina, is the unusual case of a vendor that attacks the **Matching** link rather than the list underneath it, turning the matcher itself into the product the bank integrates with.

- **List Aggregation** — source lists from regulators, open research, secondary publishers.
- **Matching** — fuzzy-string, phonetic, transliteration, and identity-resolution scoring.
- **Scoring / Thresholding** — calibrated decision producing a verdict per lookup.
- **Alert Generation** — queue population and routing across business lines.
- **Adjudication** — analyst disposition, escalation, and evidence capture.
- **Attestation / Reporting** — regulator-ready audit trail and periodic filings.

The strategic move: Bridger attacks **Matching**, not the list. The vendor treats the list as a commodity input and converts its edge into the matcher's precision-recall curve. The downstream bank retains full ownership of adjudication and attestation; the vendor owns the part where most of the false-positive cost is created. That single repositioning converts a commodity list subscription into a defensible engine subscription.

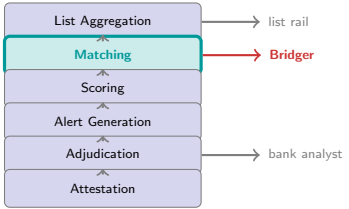

List stays commodity. Matcher becomes the product.

Evans-Wurster anchor — deconstruction can reposition the product within the chain. Bridger attacks Matching, not the list, and that repositioning is the business.

Is Napier's AI-Overlay Arbitrage a Durable Moat or a Regulator-Dependent Window?

Regulatory arbitrage = a firm earns profit specifically because it faces a lighter rulebook than its competitors, not because it is better at the underlying business; the advantage lasts only as long as the rulebook gap does. A *moat* = a competitive advantage that rivals cannot easily copy.

The Explainability-Arbitrage Wedge (Napier, United Kingdom)

Napier, headquartered in London, built its business in the gap between two supervisory truths: banks must now show auditable, explainable reasoning for every sanctions decision, and incumbent list vendors were built for a decade when raw hit rates were the metric and auditability was an afterthought. Napier sells the explainable-AI overlay that turns a hit into a defensible decision, together with the workflow that packages that decision for the examiner.

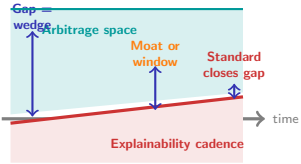

- **The gap:** supervisory expectation on explainability keeps rising; legacy list pipelines produce hits that do not carry the evidence trail the supervisor expects.
- **The wedge:** run the overlay for the bank, deliver a packaged reasoning artefact, keep the model version history audit-ready.
- **The moat thesis:** as explainability cadence rises, the per-lookup overlay premium becomes a structural add-on that the bank cannot walk back from once it is embedded.
- **The window thesis:** if regulators ship standard explainability toolkits or incumbent list vendors ship bundled overlays, the per-lookup premium collapses.

The dual reading is deliberate. Napier can become a required second line of defence — a moat — or a transitional overlay layer that disappears when the standard matures. The business model bets that procurement inertia at banks outpaces standardisation at regulators.

Arbitrage-to-moat anchor — Napier converts rising explainability cadence into a per-lookup overlay premium. Durability depends on who standardises first.

Why Does SymphonyAI Sensa's Alert-Ranking Model Thrive at Large Banks but Stall at Small Ones?

The Sensa Context Test (SymphonyAI Sensa, United States)
SymphonyAI Sensa, headquartered in Palo Alto, is not a list vendor. It is an overlay whose business model is stitched to the plumbing of a buyer who already generates a thick stream of alerts. Sensa provides the ranking engine that re-orders a queue the bank already has and promotes the likely-true hits to the top. The ranker is the gate; the alert stream is the fuel.

- **Where it thrives:** large banks with deep alert volume, standardised case-management systems, and enough labelled historical dispositions to train a ranker that the supervisor will accept.
- **Where it stalls:** small banks whose alert stream is thin, whose disposition history is sparse, and whose case management is still spreadsheet-driven — the ranker has nothing to re-order and no labels to learn from.
- **The dependency:** the vendor's value proposition is not only the ranker but also the labelled-alert history the buyer brings. Without that history the ranker has nowhere to point its confidence.

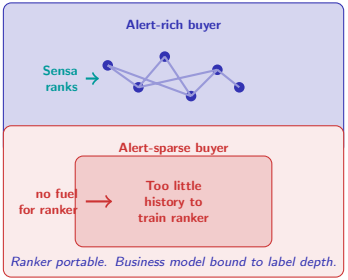

The lesson: a sanctions-screening overlay is not a universal product. Its viable buyers are the ones whose internal alert plumbing can feed the model the labels it needs. The ranker travels in principle; the business model travels only as far as the labelled-alert history does.

Context-dependency anchor — Sensa's business model is bound to labelled-alert depth at the buyer. The ranker travels; the training fuel does not.

Which Five Tests Separate a Durable Screening Vendor from a Fragile One?

The Five-Test Synthesis

1. **Friction test:** does the vendor close a gap the bank is already paying to cover? Refinitiv World-Check applied: the bank already pays for ongoing research labour; the roster replaces that labour at scale.
2. **Platform test:** does each new participant improve the product for the next? World-Check applied: consumer queries surface edge-case aliases that sharpen the researcher-side roster.
3. **Rebundling test:** can the vendor add adjacent products on the same billing cadence and the same buyer? Dow Jones Risk applied: matcher, adjudication, entity graph and attestation all reuse the compliance buyer and per-lookup rail.
4. **Infrastructure test:** is the vendor repositioning inside an existing link or building a new one? LexisNexis Bridger applied: attacks Matching rather than the list, turning a commodity list into a matcher subscription.
5. **Arbitrage test:** is the wedge the gap between supervisory expectation and bank tooling — and is that gap widening? Napier applied: explainability cadence outpaces legacy-list evidence trails.

A sixth meta-test applies across all five: the context test (Sensa). Does the buyer's internal alert plumbing generate enough labelled history to power the vendor's overlay? If not, even a perfect ranker has nowhere to apply its confidence. Durable vendors pass at least three of the five tests above *within* a buyer that passes the context test.

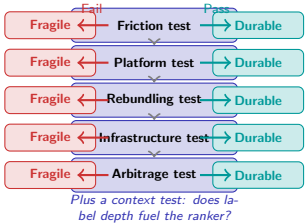

Synthesis anchor — durable screening vendors pass at least three of five tests within a buyer whose alert history can fuel the overlay.

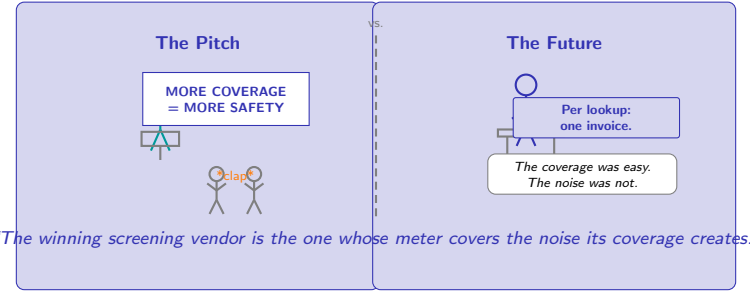

"The winning screening vendor is the one whose meter covers the noise its coverage creates."