THE INTELLIGENCE LAYER:
WHY CAPITAL MARKETS FINTECH DEMANDS A DIFFERENT KIND OF RESEARCH

A single capital markets fintech vendor might span equity derivatives, regulatory reporting, and post-trade analytics — serving hedge funds in New York under one framework and pension funds in Frankfurt under another. Now try categorizing that in a generic industry database.

Global technology spending across capital markets hit $163 billion in 2023 and is on track to reach $244 billion by 2028, with M&A activity hitting record highs — over 1,350 deals through Q3 2025. New categories are emerging as firms rethink everything from multi-asset trading infrastructure to regulatory technology, and the pace of innovation is compressing product cycles that used to unfold over years into quarters.

This acceleration has created a problem that most industry databases weren't designed to handle. Capital Markets fintech doesn't segment neatly. A single vendor might span multiple asset classes, serve different buyer types across the buy-side and sell-side, and operate under distinct regulatory frameworks depending on geography. The relevant dimensions — asset class, business function, user segment, regulatory construct — intersect in ways that generic industry categories struggle to capture. Academic research has described fintech categorization as a "blurry issue," noting that the lack of standardized classification makes competitive comparison genuinely difficult. The Cambridge Fintech Ecosystem Atlas, one of the more comprehensive efforts, mapped the space across 14 market segments, 63 subsegments, and 118 categories— and that's fintech broadly, before drilling into capital markets specifically.

This complexity would be paralyzing without a recent convergence. Data has become more accessible and granular, and large language models can now process, tag, and structure it at speeds and with consistency that manual approaches couldn't achieve. The next application layer is already here: specialized AI models trained on narrow domains — post-trade settlement workflows, derivatives pricing terminology, regulatory filing structures — are increasingly deployed for nuanced classification tasks that general-purpose models miss, such as distinguishing a vendor focused on post-trade derivatives analytics from one handling equity execution. The infrastructure to automate delivery for specific use cases, whether that's investment screening, competitive mapping, or M&A due diligence, has caught up with the models.

AI capability alone, though, doesn't solve the categorization problem. Research from Glean found that organizations with high-quality, well-structured data achieve factual accuracy above 95%, while those relying on generic data structures achieve 60-70% — a gap that no amount of model sophistication can close on its own. Squirro's 2024 acquisition of Synaptica, a semantic graph specialist, underscored the same point: the acquirer explicitly stated that taxonomies and

knowledge graphs are needed to "ground" generative AI and prevent hallucinations. The AI capability exists. The binding constraint is the structure directing it — the taxonomy that determines whether you get 95% accuracy or 70%.

BUILT BY PRACTITIONERS, POWERED BY AI

The research market today spans everything from financial metrics platforms aggregating earnings data to deep-sector intelligence built around how practitioners actually categorize vendors and workflows. At OPCO, we saw both dynamics from the inside — the industry's structural evolution and the emergence of capabilities that could address it. As practitioners who have spent years evaluating capital markets fintech vendors, mapping competitive sets, and advising on technology strategy, we recognized something that horizontal research platforms cannot prioritize: the multi-dimensional complexity of how this market actually works.

AlphaSense, PitchBook, and CB Insights excel at cross-industry pattern recognition and processing enormous volumes of documents and filings — AlphaSense alone has raised over $1.6 billion and earned a spot in the CNBC Disruptor Top 10. These are serious platforms that deliver real value for broad market intelligence. Their business model, however, requires categories that work across all industries — fintech, healthcare, consumer, industrials — and the segment-level distinctions that capital markets practitioners rely on tend to get compressed in that translation. FinoptiQ, created as a deep, capital markets database, surfaces what horizontal platforms don't: operator workflow structure, segment-level categorization that tracks how a vendor's post-trade analytics capability differs from its equity execution play, and coverage of emerging verticals like crypto infrastructure and DLT settlement, where categories are still being defined.

The result is a structural gap. Bessemer Venture Partners put it directly in their 2025 State of AI report: "The future of AI is vertical." Tidemark's 2025 Vertical SaaS Benchmark, drawing from over 200 companies across 20 sectors, found that fintech-focused vertical platforms achieve 96% gross retention compared to roughly 90% for horizontal counterparts. Vertical SaaS is growing 2-3x faster than horizontal. Nearly 90% of executives view vertical platforms as the future of the sector. The economics validate what practitioners already experience: depth creates value that breadth cannot replicate.

FinoptiQ started with what we knew from operating in the market — which companies compete in which of the capital markets' 13 distinct technology verticals, where categorization breaks down in existing databases, and how buying decisions map to vendor capabilities. The result is a three-level taxonomy (Major, Minor, and Sub-sector tags) with an asset-class overlay, indexing roughly 3,800 capital markets fintech companies across the US, UK, and EU. Most fintechs straddle at least two functional categories — a company might serve both pre-trade analytics and post-trade processing — so the system uses dual-tagging to map vendors at the neighborhood level, not just at the borough level. A proprietary LLM harvests over 40 standardized attributes per company, from funding history and asset-class focus to client mix and regulatory posture, and then structures that data, augmenting it with qualitative color from OPCO's network of over 10,000 industry contacts — pricing corridors, product gaps, and competitive positioning that no public filing captures.

The research supports this hybrid approach. JPMorgan's work on human-in-the-loop machine learning demonstrated that interactive, expert-guided AI systems augment and significantly outperform fully automated alternatives. Broader research on document processing in financial services found that combining AI with human expert validation achieves

accuracy rates up to 99.9% for critical financial documents — a threshold that pure automation consistently falls short of. The combination of practitioner judgment and AI processing isn't a workaround — it's the only architecture that delivers the reliability these decisions require.

WHY STRUCTURE DETERMINES ACCURACY

Capital markets professionals — investors, M&A advisors, strategic planners, technology buyers — need competitive intelligence that reflects how the market actually segments. When the underlying taxonomy is wrong or incomplete, the downstream consequences are measurable. SRS Acquiom's 2025 M&A Due Diligence Study found that 44% of leaders cite a lack of quality third-party data as the greatest barrier to effective due diligence. Among boutique investment banks, 40% identified incomplete information on a target company as a top hurdle. And the CFA Institute notes that insufficient due diligence contributes to M&A failure rates between 70% and 90%.

FinoptiQ addresses this by delivering research that starts with the right structure. Search for vendors in a specific segment, and the results reflect practitioner-designed categories, not a generic taxonomy that requires manual re-sorting. A scoring engine weighs traction, functional fit, and valuation sanity to surface the most relevant targets per theme — so a PE firm originating a deal thesis, a bank assembling a buyer universe, or an asset manager benchmarking vendors can move from question to output in hours rather than weeks. Reports auto-render as Excel, JSON, or board-ready fact sheets, and a real-time diff-engine flags changes in ownership, C-suite moves, and funding rounds daily, keeping competitive sets current without manual monitoring.

Horizontal databases do what they're built to do: broad coverage across industries, fast access to filings and financials, and reliable search at scale. Where FinoptiQ goes further is practitioner-designed categorization — the kind that lets an investor see which vendors compete on post-trade derivatives analytics versus which ones handle equity execution, without manually re-sorting results. And because the structure is right from the start, the output travels: the same competitive set that informs a screening call renders as the Excel workbook or board-ready fact sheet that closes the loop.

The capital markets fintech landscape will continue to evolve. New vendors will emerge, categories will shift, and the market's multi-dimensional complexity will only increase. If you're spending time re-sorting results, maintaining shadow taxonomies in spreadsheets, or manually correcting competitive sets, your data infrastructure doesn't match the market you're operating in. That's the gap FinoptiQ was built to close—and we'd welcome the conversation about what it would look like for your specific use case.

For questions or inquiries, please reach out to finoptiq@opcollc.com