

CAPITAL MARKETS IN THE AGE OF AI:
THE FLYWHEEL COMPOUNDS, FOR THE
PLATFORMS POSITIONED TO SPIN IT



MOTIVE PARTNERS

SPEED OF ADOPTION

WHY CAPITAL MARKETS INFRASTRUCTURE IS DIFFERENT

The broader debate about AI and its impact on financial software is continuing to play out in the public markets. Public financial software and solutions companies have traded off roughly 40% over the last year¹, and terminal-value assumptions are being repriced across the category for cost-of-code collapse, seat-pricing redundancy, stack compression, and margin compression. Those concerns are real. They are not, however, evenly distributed across software.

Capital markets infrastructure is a category where the headline risks are dampened by structural features of the sector itself: **(1) gated financial networks, (2) regulatory standing, (3) deeply embedded workflows, (4) decades of proprietary transactional data, and (5) customer relationships and trust measured in decades.** That is why public exchanges and trading platforms are flat over the last year and have actually traded up 11% over the last 3 months.² The question for an investor in this space is not whether AI is disruptive, it is, but where in the capital markets value chain the disruption shows up as opportunity for incumbents, where it shows up as opportunity for new entrants, and where it shows up as terminal-value risk for both.

Each of those structural features is not just defensive cover, it is an **input AI itself needs to run at institutional scale, and each compounds in value when paired with a credible AI strategy.** Gated networks are where execution actually happens; an agent operating outside them is operating blind. Regulatory standing, audit trails, governance gates, supervision, is precisely the harness foundation models lack today. Deeply embedded workflows are the place where the labor budget actually sits, and where AI converts process knowledge into agentic product.

HIGHLIGHT

The single most important diligence question in capital markets software is no longer “what does this software do”, it is “what proprietary data does the workflow generate, what network does it sit inside, and what models can be trained on it.” Everything else flows from that answer.

¹ Source: FactSet as of 5/22/26, Public financial software & solutions companies: Intuit, Fiserv, FIS, Guidewire, Jack Henry, Xero, Clearwater Analytics, Temenos, ACIW, Q2, Upstart, CCC, Alkami, nCino, Blackline, Pagaya, Riskified, Finatext HD, Blend, Open Lending, Expensify

² Source: FactSet as of 5/22/26, Public exchanges and trading platforms: CME, ICE, HKEX, LSE, Deutsche Boerse, Nasdaq, Cboe, Tradeweb, SGX, Euronext, B3, JPX, TMX Group, ASX, Miami International Holdings, MarketAxess, BMV

Proprietary transactional data is the asset most foundation models cannot replicate, and the asset most likely to train a model nobody else can. Multi-decade customer trust is the gating factor on autonomous execution; in capital markets, traders will not hand a \$25M block to an AI whose pricing they do not believe in, and that trust is built one trade at a time over years.

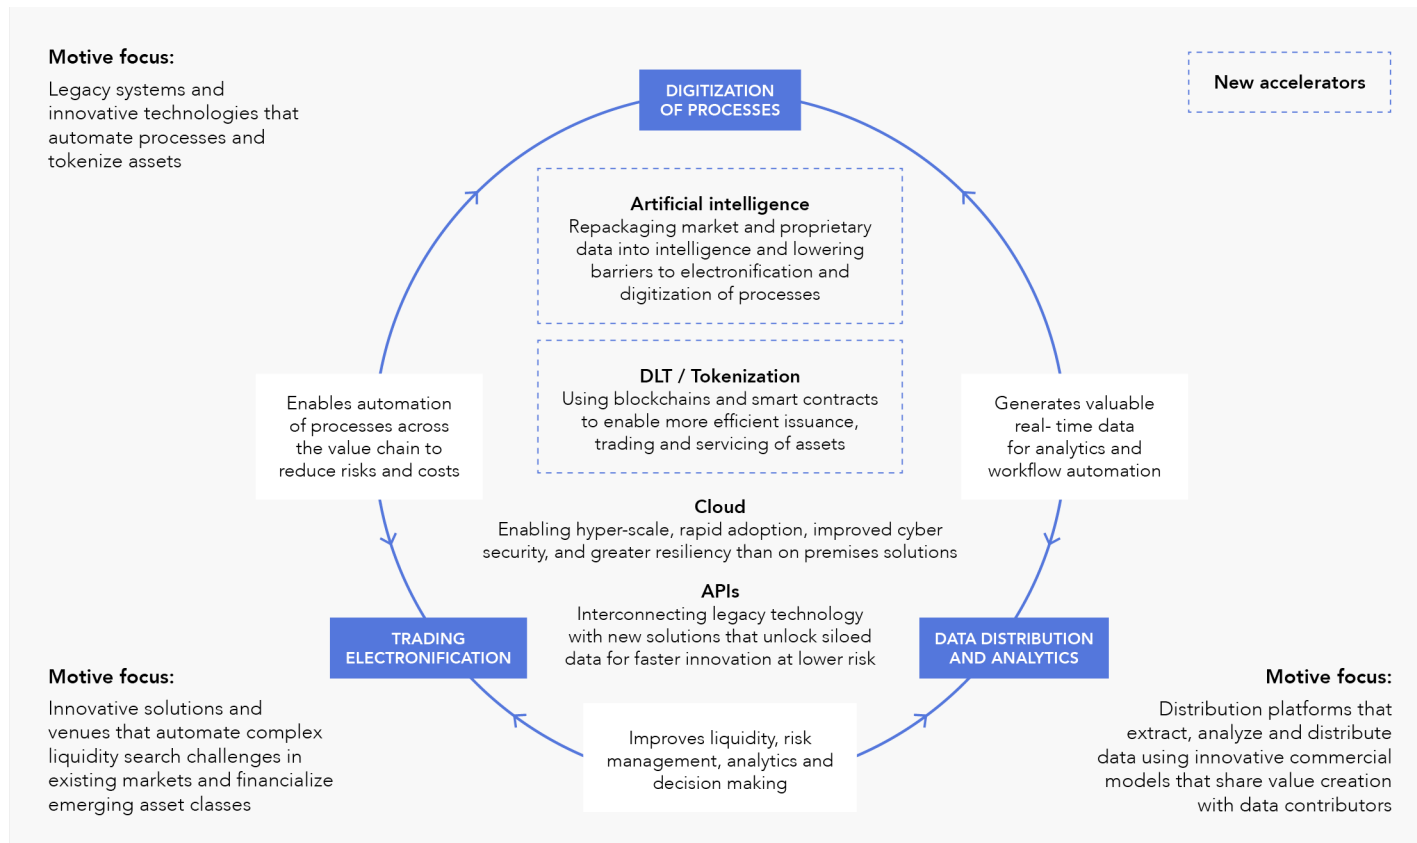
The platforms that own all five compound their advantage as AI matures rather than seeing it erode. The platforms that own none of them are exposed regardless of how many AI features they ship.

THE CAPITAL MARKETS FLYWHEEL

Capital Market efficiency in any given asset class is built on three reinforcing layers of a flywheel: **electronification, digitization, and distribution.** Electronification moves

asset classes from voice and paper onto electronic rails. Digitization turns post-trade, servicing, risk, and compliance workflows from manual to machine-run processes. Distribution monetizes the structured data the first two layers produce. The layers are not independent, each makes the next more valuable. Electronification produces structured workflow data. Digitized workflows generate clean, machine-readable training data. Distributed data creates transparency and attracts more flow back onto the platforms that produced it in the first place. Every turn of the wheel makes the next turn easier.

AI is not a separate, fourth layer. **It is the accelerant on all three**, sharpening the slope of each layer and shortening the time between turns. The platforms operating in more than one layer, with the structural advantages described above, are the platforms that benefit most.



**ELECTRONIFICATION: AI COMPOUNDS WHERE
STRUCTURED FLOW ALREADY EXISTS**

AI in capital markets does not begin with the model, it begins with the venue. You cannot train a fair-value engine on phone calls. The precondition for the next wave of AI products is the fifteen years of patient work that moved trading onto electronic rails, driven by global venues like Tradeweb, MarketAxess, Trumid, ICE, CME, Deutsche Börse, LSEG, and Nasdaq, and by a generation of newer electronic platforms attacking less-liquid corners of the market. US corporate bond electronic trading is now well over 40% of TRACE volume, OTC rates and FX have continued to electronify, and innovation in adjacent asset classes, syndicated loans, private markets secondaries (Forge, EquityZen), tokenized assets, prediction markets, is generating new pools of structured flow that did not exist three years ago.³

AI accelerates this layer in two ways. First, it **lowers the barrier to electronification of harder asset classes** by automating workflow steps that historically required a human such as chat parsing, RFQ capture, dealer selection, and pre-trade pricing on instruments with no continuous quote. LLM-based trade-ticket capture from unstructured chat (now in production at multiple credit venues), ML-driven dealer selection on rate swaps (e.g., AiSNAP)⁴, and AI-enhanced exchanges like IEX applying ML to price discovery all compress the cost of moving an asset class from voice to electronic. Second, it **increases the value of electronification for the venue itself**, every additional electronic trade produces structured data that trains the next generation of pricing, liquidity, and execution models, deepening the venue's moat against AI-native challengers and against the foundation labs increasingly building vertical capability of their own.

The beneficiaries are venues with **scale, network depth, and proprietary execution data**, incumbents whose moats AI compounds rather than threatens. Beyond the global venues named above, well-positioned beneficiaries include credit-focused platforms like Trumid, loan-trading platforms like Octaura, private markets secondaries such as Forge, and the new wave of execution-intelligence platforms being built on top of existing rails (e.g. smart order routers for fixed income, multi-protocol routing engines, AI-driven dealer selection). The losers are venues with subscale flow or weak network effects, which can buy AI features, but cannot generate the proprietary data needed to make them differentiated.

3 Source: The DESK, 'E-trading flat as US credit volumes balloon', January 2026 (fi-desk.com); Coalition Greenwich, 'U.S. Corporate Bond Trading in 2024 by the Numbers', 2025. Investment-grade electronic trading approached 50% of TRACE volume; high-yield remained at approximately 32%. (<https://www.fi-desk.com/e-trading-flat-as-us-credit-volumes-balloon/>)

4 How Tradeweb is Unlocking Data to Enhance Liquidity Discovery (<https://www.tradeweb.com/newsroom/media-center/insights/blog/how-tradeweb-is-unlocking-data-to-enhance-liquidity-discovery/>)

DIGITIZATION: AI IS RE-PRICING SERVICE REVENUE AT SOFTWARE MARGINS

The second layer is where the dollars are. For every \$1 of software spend inside a major financial institution, roughly \$10 to \$26 sits in adjacent labor and professional services performing the work that software supports.⁵ AI changes what addressable means. When software stops supporting a workflow and starts executing it, the relevant TAM is no longer the software budget, it is the labor budget sitting next to it. Even at 10% labor addressability the software TAM doubles; at the higher end of that range modest penetration translates into multi-x expansion. The Ramp AI Index now shows 68% of US finance businesses hold a paid enterprise AI subscription⁶, second only to technology itself, the demand-side conditions for capturing that labor budget are already in place.

AI is accelerating this layer through two reinforcing patterns. The first is **incumbents converting proprietary process knowledge into agentic product**. Broadridge has deployed agentic AI across post-trade BPO for 40+ institutional clients, automating unstructured email classification, data retrieval, response drafting, and SLA-based prioritization.⁷ SS&C deployed 3,300+ digital workers internally first, automating credit agreement processing 95% faster and NAV reconciliation, then commercialized the same agents to clients through its AI Gateway, generating over \$200M in annual cost savings and adding a new high-margin revenue line on top of fund admin.⁸ Each of these is service revenue, historically priced at services margins, being **re-priced at software margins**. This is the SaaSification of BPO economics and one of the cleanest examples in any vertical software category of a moat-driven TAM expansion.

The second pattern is **software platforms compounding their own productivity**. Across vertical fintech, the majority of new code is now written with AI assistance. Engineering headcount has stayed flat or grown slower than output across multiple public-company disclosures. The structural implication: every platform serving capital markets is, in effect, getting cheaper to run at the same time its addressable market is expanding, a combination that has historically been very rare in software.

The beneficiaries are platforms with **deep workflow ownership and proprietary process knowledge**, which is precisely where the digitization layer is concentrated: post-trade utilities and BPO platforms with the data exhaust to build defensible agents (SS&C, BetaNXT, FNZ); buy-side investment accounting, reporting, and reconciliation platforms with cross-client data and embedded workflows (Clearwater Analytics, SimCorp); collateral, margin, and reg-eligibility platforms automating high-volume rule-driven work (OSTTRA, LSEG PTS, Transcend). The most exposed are middle-of-the-market legacy software vendors with workflow surface area, but thin data moats and BPO-heavy outsourcers whose underlying labor is the most directly AI-addressable.

⁵ Motive Analysis, Derived from 10-K data of four major US banks in North America. Ratios range from 1:10 to 1:26 for derived software to labor and professional services spend. This analysis is illustrative based on a sample of four major US banks; the ratio may vary significantly across institution types, geographies and sub-sectors.

⁶ Source: Ramp AI Index, <https://ramp.com/data/ai-index> (accessed May 2026).

⁷ Source: Broadridge Financial Solutions press release, May 2026 (<https://www.broadridge.com/press-release/2026/broadridge-deploys-agentic-ai>).

⁸ Source: SS&C Technologies press release via BusinessWire, June 2025 (<https://www.businesswire.com/news/home/20250630282734/en/SSC-Blue-Prism-Recognized-as-a-Gartner-Magic-Quadrant-RPA-Leader-for-the-Seventh-Consecutive-Year>); SS&C WorkHQ launch press release, April 2026 (<https://www.businesswire.com/news/home/20260429512507/en/SSC-Unveils-WorkHQ-to-Power-Enterprise-Agentic-Automation>).

DISTRIBUTION: WHERE PROPRIETARY DATA BECOMES AN INTELLIGENCE LAYER

The third layer is the most underappreciated. Every serious financial institution has two decades of transactional, pricing, and behavioral data, but very few have it in a state where an agent can actually use it. The bottleneck is not the model, it is the plumbing underneath.

AI is accelerating this layer by **collapsing the cost of repackaging proprietary data as an intelligence layer** rather than as a feed. The firms moving fastest are the ones that already own scarce data and have figured out how to commercialize it at a much higher level of the stack. Bloomberg trained a 50B-parameter finance-domain LLM on its proprietary content, something no general-purpose model can replicate.⁹ S&P's Kensho launched an LLM-ready API to make its structured datasets plug directly into customer GenAI models.¹⁰ AlphaSense aggregates 500M+ documents to feed AI-driven research workflows.

The pattern matters: each of these firms already owned the data, AI just gave them a way to monetize it at a higher level of the stack, and at outcome-based rather than license-based economics. The beneficiaries are trading venues and exchanges with **critical mass of proprietary trading and pricing data** and platforms with **scarce, non-public, workflow-generated data**: exchange and venue data, derivatives histories, fund-level alternatives data, private markets reference data (Preqin, With Intelligence, Pitchbook), post-trade exception and lifecycle data, and credit-specific datasets that feed AI-native research (9fin, AlphaSense). The exposed are aggregators of public or near-public data and commoditized news/reference services, which is exactly the workflow an AI-native player can replicate at lower cost, and the part of the data sub-sector where multiples have contracted most sharply over the last twelve months.

HIGHLIGHT

Electronify the trade, digitize the workflow, distribute the data, each turn of the flywheel makes the next one easier, and AI sharpens the slope of every layer. The platforms operating in more than one layer compound rather than just grow.

⁹ Source: Bloomberg press release, 'Introducing BloombergGPT, Bloomberg's 50-billion parameter large language model, purpose-built from scratch for finance', March 30, 2023 (<https://www.bloomberg.com/company/press/bloomberggpt-50-billion-parameter-llm-tuned-finance/>)

¹⁰ Source: S&P Global / Kensho product announcement. The Kensho LLM-Ready API is a publicly documented S&P Global product

WHERE THE FLYWHEEL RUNS FASTEST, AND WHERE IT STALLS

Not every sub-sector of capital markets infrastructure is equally well-positioned to spin the flywheel.

Capital markets software is the strongest hunting ground. Front-to-back OMS/EMS, post-trade platforms, and integrated trading systems own the workflow end-to-end, sit inside regulated networks, and are adjacent to the largest single labor budget in financial services. The moats are sufficient, not maximal, and the labor budget being addressed is the largest. Risk concentrations are in middle-office point solutions that don't own workflow end-to-end, and in legacy under-invested platforms that have left the door open for AI-native challengers.

Financial markets infrastructure, exchanges, venues, clearing, post-trade utilities, have the deepest moats of any sub-sector, but the most limited TAM expansion. Workflows are already substantially automated; AI primarily delivers margin improvement and incremental tokenization-driven new product. The investable angle is cost takeout and selective innovation, not transformation.

Capital markets data is bifurcated. Core data with proprietary differentiation, exchange feeds, derivatives histories, fund-level data on alternatives, private markets reference data, sits on the right side of the line and benefits cleanly from AI through higher monetization of existing assets. Aggregated public data and enriched public data sit on the wrong side; these are exactly the workflows an AI-native player can replicate.

Trading platforms and tech-driven broker-dealers are mixed. Retail platforms with strong brand and network, Robinhood, Trade Republic, are well-positioned to capture retail interest in new asset classes including crypto, prediction markets, private markets, and alternatives. Pure-UX retail apps with no data or network are exposed.

**AI INNOVATION IN ELECTRONIC CREDIT TRADING:
A TRUMID CASE STUDY¹¹**

The clearest way to see how the flywheel turns in practice is to look at a single asset class over time. Credit, and our portfolio company Trumid specifically, is a useful case study because the sequence: market maturation generating quality data, quality data plus time invested in generating trusted analytics, trusted analytics unlocking AI-native execution, is unusually clear. The same pattern is now being run in adjacent asset classes, and the same pre-conditions apply.

Market maturation made the innovation possible. A decade ago, US corporate bond trading happened almost entirely over the phone and via chat, dealer-to-client, with virtually no electronic data exhaust. Today, electronic trading is well over 40% of TRACE volume, and the participant base has expanded dramatically, not just dealers and traditional asset managers, but ETF issuers, systematic credit funds, non-bank liquidity providers, and increasingly retail-oriented allocators. That maturation, driven by venues like Trumid and others over more than a decade, is the precondition for everything now happening in credit AI. Without the volume of structured electronic data those venues have generated, none of the AI-native innovation now hitting the market would be possible. You cannot train a fair-value model on phone calls.

Client receptivity has shifted in parallel. Capital markets clients have historically been, rightly, conservative about adopting new trading workflows, both because of regulatory considerations and because desks have decades of best practices that genuinely work. AI has changed that posture meaningfully in the last 18 months. Buy-side and sell-side traders are actively asking for new tooling rather than resisting it. That receptivity is what is letting the next wave of features move from pilot to production at a pace that was not possible even three years ago.

¹¹ Source: Trumid company materials

HIGHLIGHT

Trust takes years to build and seconds to lose, and you have to earn it every single day. In capital markets, the gating factor on AI adoption is not model quality, it is trust.

Trust is the precondition for AI in trading. In capital markets, the gating factor on AI adoption is not model quality, it is trust. A trader will not let an AI system route a \$25 million block trade based on a fair-value estimate they do not believe in. The way platforms earn that trust is by building underlying data and analytics products first, and letting clients see them work, repeatedly, over years, before layering automation on top.

Trumid's **Fair Value Model Price (FVMP)** is a clean illustration of the pattern. It generates ML-driven prices on roughly 26,000 USD-denominated corporate bonds spanning all major sectors in IG, HY, and EM, refreshing every 30 seconds, with confidence-scored mid and bid-offer levels and size-adjusted pricing across six liquidity buckets. Clients now widely treat it as a reliable, objective measure of fair value because they have seen it perform against actual prints across thousands of trades over multiple years. **Trumid PT Pricer** extended that fair-value foundation to portfolio trading, estimating where a list of names should transact as a portfolio versus single-name execution, both pre-trade and in-session, with liquidity scores and confidence scores layered on top.

That foundation, trusted pricing, liquidity context, and confidence scoring, is what unlocks everything coming next. Without it, no client would let an AI execute on their behalf. With it, clients are now actively trusting the platform, within defined trader criteria, to evaluate and execute opportunities. The same dynamic applies to any AI-native execution product in capital markets: years of trusted analytics are the competitive advantage, and that advantage compounds as more clients use it.

What that trust is now enabling. The product velocity in credit over the last six months tells the story. (1) **Trumid Smart Voice**, launched publicly in late April, uses LLMs to convert unstructured trader chat into pre-populated trade tickets in seconds. Initially launched in December for dealer sales and expanded to sell-side traders and buy-side clients in April, it has already processed over \$56 billion in traded volume and eliminated an estimated 77,000 manual clicks. Adoption has been viral. (2) **Trumid PT-to-RFQ list optimization**, going live in coming weeks, is a proprietary analytics and workflow tool designed to

evaluate a list of bonds and help traders determine whether it should be executed via Trumid PT, RFQ, or a mix of both. (3) **Trumid Full Self Trading (FST)**, built into the platform since late 2025 and now broadly available with dozens of clients integrating the automation tool into their daily trading workflow and trusting it to execute across multiple Trumid protocols. FST lets a trader input a block order with parameters (size, urgency, price tolerance, completion deadline) and the platform handles execution end-to-end.

Two things are worth flagging beyond the product list. First, none of these features exists in isolation, each is a layer on top of the trust accumulated by the layer below, and none of them would work without the years of FVMP and PT Pricer data that came before. Second, all of this is being delivered while the underlying engineering organization compounds its own productivity, the majority of Trumid's code has been written with AI assistance for roughly 18 months, headcount in the tech team has stayed flat over the past two years while output has roughly doubled.

Why this matters as a case study. Trumid is not the only firm doing this kind of work, their peers in credit, and players across rates, FX, and equities, are running similar playbooks. But the sequence here is extremely clear: market maturation generates quality data, quality data plus the time invested generates trusted pricing, trusted pricing unlocks AI-native execution features, and AI-native delivery infrastructure makes the whole thing economically viable. You cannot skip any of those steps. AI did not suddenly arrive in credit trading this year, it is the payoff for fifteen years of patient work building electronic infrastructure, accumulating data, and earning client trust one trade at a time. The firms that put in that work first are the ones now reaping the AI dividend, and the same template will play out across every other asset class still climbing the electronification curve.

POSITIONING

The next decade of returns in capital markets will not come from picking the single biggest AI winner. It will come from **identifying scaled, network-rich, data-rich category leaders that still have their workflow ownership and customer relationships intact, and re-founding them on a modern intelligence layer**, and from backing the early-stage builders shaping the categories where there is no incumbent yet.

That standard maps directly to where we are concentrating capital: scaled electronic trading venues with proprietary execution data; post-trade and buy-side infrastructure platforms where labor budgets are large and workflows are embedded; intelligence-layer plays on top of scarce, proprietary capital markets data; and the integrated build-out of private markets infrastructure. The last fifteen years in capital markets were about electronification. The next ten are about intelligence, applied on top of those rails, and we are positioning the portfolio around that transition across the full risk curve.

HIGHLIGHT

AI did not suddenly arrive in credit trading this year, it is the payoff for fifteen years of patient work building electronic infrastructure, accumulating data, and earning client trust one trade at a time. The firms that put in that work first are the ones now reaping the AI dividend.

**CAPITAL MARKETS IN THE AGE OF AI:
THE FLYWHEEL COMPOUNDS, FOR THE
PLATFORMS POSITIONED TO SPIN IT**

P 10/11

ABOUT MOTIVE PARTNERS

Motive Partners is a private investment firm exclusively focused on financial technology and technology-enabled business services companies, investing from early-stage ventures to growth equity and buyout in North America and Europe. The firm invests across five subsectors: banking & payments, capital markets, data & analytics, insurance, and wealth and asset management. Motive Partners applies its proven investor, operator, innovator (IOI) model across its

portfolio, combining deep financial technology expertise and proven operational rigor to accelerate growth and value creation. With offices in New York, London, and Berlin, the firm provides differentiated insight, connectivity, and capabilities to create long-term value in financial technology companies.

More information on Motive Partners can be found at www.motivepartners.com.

MOTIVE PARTNERS

This paper is published by Motive for informational and educational purposes only. It does not constitute investment advice, a recommendation, or a solicitation to invest in any fund or security. The views and opinions expressed, including those attributed to Motive employees, reflect their views as of the date of publication and are subject to change without notice. Quantitative estimates and market sizing analyses, including the 10x TAM expansion estimate, represent Motive's own analytical framework based on the sources indicated and are illustrative only; actual market dynamics may differ materially. Case study data is sourced from publicly available company disclosures and industry reports; Motive has not independently verified these figures. Forward-looking statements reflect Motive's current views and involve known and unknown risks and uncertainties. Actual developments may differ materially. Nothing herein should be relied upon as the basis for any investment decision.

