

# AI in Wholesale Brokerage: Adoption, Risk, and Regulatory Fit

Mitchell de Vries

Senior Software and AI Engineer at DMA One

mitchelldevries2001@gmail.com

August 20, 2026

## Abstract

This survey examines where AI is actually being used in wholesale brokerage, focused on two domains: the client lifecycle (onboarding, identification, and monitoring) and settlement operations. For each domain I describe the current level of adoption, the risks that come with it, and whether the regulatory frameworks that apply today are fit for purpose in a wholesale context. The picture that comes out is uneven. On the client side, document processing and screening are mature, while the start and end of the onboarding flow are barely AI-driven at all. On the settlement side, AI mostly sits as a thin layer on top of rule-based systems that have been in production for decades. In both domains the honest description is a stack of assistive tools, not autonomous systems, because the cost of a wrong answer in a wholesale setting keeps a human in the loop. The main risks identified are vendor concentration, hallucination, opacity, and the distance between vendor marketing and actual capability. On the regulatory question, the answer is no. MiFID II, ESMA guidance, and the EU AI Act all apply to a wholesale broker in some way, but each was calibrated around retail investor protection. The result is a gap that sits exactly where the operational risk in wholesale brokerage is concentrated. The survey also finds that the academic literature on this segment is thin, which makes the reported gains from adjacent domains hard to verify in a wholesale setting.

## Contents

|          |                                                                  |           |
|----------|------------------------------------------------------------------|-----------|
| <b>1</b> | <b>Introduction</b>                                              | <b>3</b>  |
| <b>2</b> | <b>Brokerage Context</b>                                         | <b>5</b>  |
| 2.1      | Retail Brokerage . . . . .                                       | 5         |
| 2.2      | Wholesale Brokerage . . . . .                                    | 6         |
| 2.3      | The Settlement Process . . . . .                                 | 7         |
| <b>3</b> | <b>AI in Client Lifecycle</b>                                    | <b>11</b> |
| 3.1      | Onboarding . . . . .                                             | 11        |
| 3.2      | Client monitoring . . . . .                                      | 15        |
| 3.3      | Client Support . . . . .                                         | 17        |
| <b>4</b> | <b>AI in Settlement Operations</b>                               | <b>19</b> |
| 4.1      | Trade confirmation and matching . . . . .                        | 19        |
| 4.2      | Reconciliation . . . . .                                         | 22        |
| 4.3      | Risk and collateral management . . . . .                         | 23        |
| <b>5</b> | <b>Risks and Limitations</b>                                     | <b>26</b> |
| 5.1      | Vendor concentration and lock-in . . . . .                       | 26        |
| 5.2      | Hallucination and confident wrong answers . . . . .              | 27        |
| 5.3      | Opacity and accountability gaps . . . . .                        | 28        |
| 5.4      | The gap between vendor marketing and actual capability . . . . . | 29        |
| <b>6</b> | <b>Regulatory Landscape</b>                                      | <b>30</b> |
| 6.1      | MiFID II . . . . .                                               | 30        |
| 6.2      | ESMA Guidance . . . . .                                          | 31        |
| 6.3      | The AFM . . . . .                                                | 31        |
| 6.4      | The EU AI Act . . . . .                                          | 32        |
| 6.5      | DORA . . . . .                                                   | 32        |
| 6.6      | The Wholesale Gap . . . . .                                      | 33        |
| <b>7</b> | <b>Gap analysis</b>                                              | <b>35</b> |
| 7.1      | Adoption gaps . . . . .                                          | 35        |
| 7.2      | Research gaps . . . . .                                          | 35        |
| 7.3      | Regulatory gaps . . . . .                                        | 36        |
| 7.4      | How the gaps interact . . . . .                                  | 37        |
| <b>8</b> | <b>Conclusion</b>                                                | <b>38</b> |
| 8.1      | What the survey found . . . . .                                  | 38        |

|     |                                               |    |
|-----|-----------------------------------------------|----|
| 8.2 | Are the frameworks fit for purpose? . . . . . | 38 |
| 8.3 | What firms and regulators should do . . . . . | 39 |
| 8.4 | Limitations . . . . .                         | 39 |
| 8.5 | Closing . . . . .                             | 40 |

## 1. Introduction

When people talk about AI in finance, they usually talk about the visible side. Robo-advisors, chatbots in a banking app, trading algorithms. Almost all of that attention goes to retail, the side that serves private individuals. However, there is a segment private individuals never see, and AI is arriving there too: wholesale brokerage. A wholesale brokerage serves institutions such as asset managers, hedge funds, and pension funds. The clients are professional, the trade sizes are large, and the operational workload around identifying clients and settling trades is heavy. This is exactly the kind of document-heavy, repetitive work where AI vendors promise the biggest gains.

The problem is that we do not really know how much of that promise is real. The academic literature on AI in finance leans toward retail use cases, and the regulatory frameworks that apply, such as MiFID II and the EU AI Act, were largely calibrated around retail investor protection. So AI is being deployed in a segment that the research barely covers and that the rules were not written for. That is the motivation for this survey.

This paper surveys AI applications in wholesale brokerage in two domains: the client lifecycle (onboarding, identification, monitoring) and settlement operations. This gives three research questions:

1. Where is AI currently being applied in the client lifecycle and settlement operations of wholesale brokerage, and how mature is that adoption?
2. What risks does the use of AI bring with it in a wholesale context?
3. Are the regulatory frameworks that apply today fit for purpose for AI in wholesale brokerage?

It is worth saying up front what the answers turn out to be. The rest of this survey is easier to follow if you already know where it is going.

On the first question, AI in wholesale brokerage today is a stack of assistive tools rather than a set of autonomous systems. Adoption is uneven inside a single process, not just between firms. On the client side, document processing and screening are mature while the first and last steps of the onboarding flow are barely AI-driven at all. On the settlement side the AI mostly sits as a thin layer on top of rule-based engines that have been in production for decades. In both domains a human still signs off on anything material. This is because the cost of a wrong answer scales with the trade size, and in wholesale the trade sizes are large.

On the second question, four risks come out of the survey. Vendor concentration and lock-in, hallucination, opacity and the accountability gap that comes with it, and the

distance between what vendors market and what the tools actually do.

On the third question the answer is no. The frameworks that are specifically about AI are calibrated around harm to individuals. The frameworks that take wholesale operations seriously treat a model like any other system. The two bodies of rules pass each other without meeting, and the space where they do not meet is exactly where the risks above sit. What this means is that a wholesale broker deploying AI today is not breaking any rules, because the rules that would apply to what it is doing have not been written yet.

Out of scope are retail brokerage (except as contrast), AI-driven trading strategies, and portfolio management for end investors. The method is a literature survey. Because peer-reviewed research on wholesale brokerage specifically is thin, I also draw on supervisory publications, industry reports, and vendor material, with the vendor material read critically.

Chapter 2 gives the background on brokerage and settlement. Chapters 3 and 4 cover AI in the client lifecycle and in settlement operations. Chapter 5 discusses the risks, chapter 6 the regulatory frameworks, chapter 7 the gap analysis, and the final chapter concludes.

## 2. Brokerage Context

Most people have bought a stock at some point, or at least know someone who has. You open an app on your phone, search for a company, type in an amount and press buy. The order is filled within seconds and the position shows up in your account. What almost nobody thinks about is what happens after that. The trade doesn't just appear in the account by itself, it passes through a chain of licensed firms, clearing systems, and legal processes before ownership actually changes hands. The brokerage is where that chain starts.

Let's put it simply, a brokerage is a licensed intermediary that executes trades on behalf of clients who cannot access financial exchanges directly. The NYSE, NASDAQ, and equivalents don't let just anyone connect to their systems. Exchange membership requires regulatory approval, significant capital reserves, and ongoing compliance obligations that most investors don't want. A brokerage provides those services for a commission.

That much is common to all brokerages. Where they diverge is in who their clients are, and that divergence turns out to matter enormously for everything from regulatory obligations to operational infrastructure. The industry breaks into two distinct models: retail and wholesale.

### 2.1 Retail Brokerage

Retail brokerages serve private individuals, people investing personal savings through a mobile app or online platform. The rise of these platforms, together with commission-free trading, is what pulled a lot of these investors into the market in the first place [1], [2]. The typical retail client buys shares in individual companies, and increasingly also ETFs. The research from (Milazzo, et al) [3] looks at two decades of Finnish household stock portfolios and finds that more than half of retail investors hold a portfolio made up of only a few stocks. ETFs are the other side of this. A study of over 500,000 Finnish investors [4] shows that ETF adoption is growing among retail investors and that it leads to smaller but better diversified portfolios. Keep in mind that both of these studies use Finnish data, so the exact numbers will not transfer one to one to other markets. The transaction sizes are modest. On the other hand the amount of clients a broker has is enormous.

Because these clients are private individuals with no assumed financial expertise, they get the strongest regulatory protections available. Under MiFID II, brokerages must assess the suitability and appropriateness of investments before providing them, and warn or restrict clients where a product is not appropriate [5]. The basic idea is protective on purpose: small everyday investors are seen as more at risk, in a way that big professional

investors are not. A place where retail investors can start trading is Fidelity<sup>1</sup>, LYNX<sup>2</sup>, and Easybroker<sup>3</sup>.

## 2.2 Wholesale Brokerage

Wholesale brokerage is a different business serving a different world. The clients here are not individuals. They are asset managers running billion-euro portfolios, hedge funds executing complex arbitrage strategies, pension funds managing retirement savings at scale, and banks trading on their own account. Under MiFID II these clients fall into the professional client and eligible counterparty categories, and that classification is what decides how much protection they get [6].

Wholesale clients are permitted to trade across the full range of available instruments: equities, fixed income, derivatives, structured products, and more without the product restrictions that apply in retail. The assumption is that they understand what they're doing and can absorb the associated risks. This does not mean there is no regulation. There is, and it is substantial. What changes is what the regulation is trying to protect. The conduct of business rules that apply to professional clients are a reduced set compared to the retail ones [7], and the duty of care that remains is thinner and, according to [8], not always effective in practice. So the focus shifts away from protecting clients from themselves and toward market integrity, systemic risk, and operational standards instead.

Beyond execution, wholesale brokerages typically offer a broader service stack: proprietary research, direct market access, and hands-on client support tailored to the specific needs of institutional accounts. DMA ONE<sup>4</sup> and Liquidnet<sup>5</sup> are examples of firms operating in this segment. Table 1 lays out the key contrasts between the retail and wholesale brokerage.

Table 1: Key differences between retail and wholesale brokerage

|   | <b>Retail</b>                                 | <b>Wholesale</b>                                |
|---|-----------------------------------------------|-------------------------------------------------|
| 1 | High quantity of transactions                 | Low quantity of transactions                    |
| 2 | Highest level of regulatory client protection | Reduced regulatory client protection            |
| 3 | Low individual transaction value              | High individual transaction value               |
| 4 | Restricted to certain investment products     | Access to the full range of investment products |

The rows in the table pair up. Rows 1 and 3 are the same observation from two sides: retail handles many transactions of low individual value, and wholesale handles few transactions of high individual value. Rows 2 and 4 both follow from who the client is. A private individual gets the strongest protection available and a restricted product set. A professional

<sup>1</sup><https://www.fidelity.com/>

<sup>2</sup><https://www.lynx.nl/>

<sup>3</sup><https://www.easybroker.nl/>

<sup>4</sup><https://www.dma.world/one>

<sup>5</sup><https://www.liquidnet.com/>

counterparty gets neither, because it is assumed to understand what it is buying. What this means is that every difference in the table is a difference in the client relationship. Who the client is decides how much the broker has to protect them, what they are allowed to trade, and how the broker services them day to day.

The back office is another matter. Once a trade is executed, whether it was placed by a private investor or by a pension fund, the same post-trade machinery kicks in. Both types of brokerage move a trade through the same four phases before it is finished. Which raises the obvious question: what does settlement actually involve?

### 2.3 The Settlement Process

A trade is not finished when it's executed. Execution just means two parties have agreed on a price. The securities still need to change hands, and so does the money. Until both of those things happen, neither side has fully delivered on their side of the deal. The process of making that happen is called settlement.

More precisely, settlement is the exchange of securities for payment, formally known as *Delivery versus Payment* (DvP). The two legs are linked so that neither party is left exposed, a principle set out by the central-bank community in the early 1990s [9]. Weiss [10] describes the whole settlement circuit as the “plumbing” of financial markets. That comparison is accurate. Settlement is what actually moves the cash and the securities between parties after a trade is agreed, and like real plumbing it sits in the background where nobody looks at it.

Getting from trade agreement to completed settlement involves several layers of infrastructure. The most important piece to understand is the *Central Counterparty Clearing House*, or CCP. When a trade is executed, the CCP inserts itself in the middle, it becomes the buyer to every seller, and the seller to every buyer. If one counterparty defaults, the CCP covers the other side. The other party gets made whole regardless. CCPs also *net* positions across participants: instead of settling each individual trade, they calculate what each firm owes or is owed across all of its activity in a given instrument and settle a single net figure. For retail brokers routing thousands of small client orders this is efficient, although the size of the netting benefit depends on how concentrated clearing is across CCPs [11].

While wholesale trades utilize this same CCP infrastructure, the post-trade pipeline for institutional market participants is far more fragmented. Institutional trades pass through a more complex chain of custodians holding client securities, prime brokers providing financing and clearing, correspondent banks managing cross-border cash, central securities depositories (CSDs) maintaining the official ownership records. Each of these intermediaries runs its own systems. Getting them all to agree on what needs to move, when, and

to whom requires coordination that is partly automated and partly manual. [12] covers this infrastructure in detail. The consequence of all this complexity is straightforward: more things can go wrong, and when they do, a human usually has to fix it.

The settlement lifecycle itself breaks into four phases, regardless of brokerage type:

1. **Trade Capture:** Immediately after execution, the details: instrument, price, quantity and counterparty are recorded.
2. **Confirmation and Matching:** Both sides of the trade independently submit their record of what was agreed. If the two records match, the trade can proceed. If they don't, there's a discrepancy that has to be resolved before anything else can happen.
3. **Clearing:** Net obligations are calculated across all matched trades, typically with a CCP stepping in to guarantee delivery and reduce the bilateral exposure between counterparties.
4. **Settlement:** The actual transfer of securities to the buyer, cash to the seller, legal ownership updated. The trade is done.

Figure 1 shows how these phases connect for a typical wholesale trade. The diagram runs top to bottom and groups the flow into three bands: execution, clearing and settlement. Keep in mind that the four phases listed above and the three bands in the figure are the same process, since trade capture and confirmation happen within minutes of each other and both sit inside the execution band. It is worth reading the figure for one thing in particular. In the clearing and settlement bands the boxes come in horizontal pairs. The buy side back office sits opposite the sell side back office, the buyer's custodian sits opposite the seller's custodian, and a CCP or a central depository sits in the middle of each pair. Those pairs are the point. There is no single system that knows the whole trade. Each side keeps its own books and acts on its own instructions, and the party in the middle can only act once the two sides line up. So every horizontal pair in the figure is a place where the two sides can disagree, and every disagreement is a potential fail. Two other things the figure flags are worth carrying forward. At the end of the execution band no money and no shares have moved yet, all that exists is an agreement. And in the clearing band both parties monitor margin and collateral themselves, which is not how it works in retail where the broker does it on the client's behalf. That is the same split we saw in Table 1, now visible in the operational flow.



In wholesale brokerage, the gap between confirmation and final settlement is where most of the day-to-day operational difficulty sits. Trades that fail to match, or that don't settle on their intended date, are called *fails* or *late settlements*. Resolving them involves manual investigation, communication between multiple intermediaries, and often time pressure, and chains of unresolved fails can spread to other participants and affect overall market liquidity [13], [14]. Settlement windows in Europe are typically T+2 [15] [16], meaning two business days after execution. That window is about to get shorter, and we come back to what that means for the back office in Chapter 4.

Missing the settlement window is not free. Chapter 3 of the Central Securities Depositories Regulation (CSDR) sets out a settlement discipline regime that applies cash penalties for fails and provides for mandatory buy-ins, and it covers any firm settling a security in one of the European CSDs [17]. A fail also leaves the non-defaulting party exposed to the counterparty until the trade is closed out [18]. For trades worth hundreds of millions, this is not a small problem. And the fail rates themselves are not marginal. According to [19] the EU settlement efficiency rate lags behind other regions, equities settle worse than bonds, and ETFs are the worst of all, which points at a structural problem rather than at incidental mistakes.

The manual workload is high, the stakes are significant, and the processes are structured enough that automation is plausible. The following sections examine where AI is being applied, starting with client-facing functions and moving into the operational back office.

### 3. AI in Client Lifecycle

Before a wholesale brokerage settles or sends trade details to the executing brokerage a significant amount of work has already been done. The client is identified, verified, risk-assessed and on-boarded. This process can become quite complex for institutional clients. They have for example complex legal structures or cross-border operations. Furthermore once the client has been on boarded the brokerage must continuously monitor the client. Updating risk profiles and ensure ongoing compliance. Apart from all this the brokerage is also providing a service to the client. The client expects to receive client support, research about different securities and monitoring of the trades. In this section we will examine the different AI tools that are already being used by wholesale brokerage to improve the client life cycle.

We divide the client life cycle in three parts. On boarding, client monitoring and client support. Identification and verification are not treated as a separate part here. This is because in practice they sit inside the on boarding flow as two of its steps, and we discuss them there.

#### 3.1 Onboarding

Client onboarding is an important entry point for financial institutions. It is there to establish a relationship with a new client, and after that the firm has to maintain it. These tasks range from operational to regulatory, and in the past they relied heavily on manual work spread across different teams. Much of what a client has to provide, and which laws a brokerage has to adhere to, comes from the regulators, and it depends on the jurisdiction. A brokerage in the Netherlands works under different rules than one in South Africa. Even inside the EU the rules are not applied in the same way: member states use different definitions of a politically exposed person and run their beneficial ownership registers in different formats, so a firm active in several countries has to adjust its process per country [20].

Before the paperwork starts there is a conversation about scope. Which markets the client wants to access, which instruments they will trade, and which services they are actually buying. There are four common ones. Direct market access (DMA), where the client sends orders straight to the exchange through the broker's membership instead of asking the broker's desk to execute for them. Research, meaning the broker's own analyst coverage on securities and markets. Settlement, where the broker handles the post-trade leg and delivers the securities and cash instead of the client's own custodian doing it. And financing, which covers margin lending and securities lending so the client can take a position larger than the cash it has posted. Everything downstream scales with what is agreed here. A pension fund trading equities looks very different from a hedge fund

running cross-border derivatives.

From there the flow runs through seven steps. Table 2 sets out what each one asks of the client and who does the work.

Table 2: The manual onboarding flow for an institutional client

|   | <b>Step</b>                      | <b>What the client provides</b>                                                                                                                                                                                | <b>Who does it</b>                                                                    |
|---|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| 1 | Pre-onboarding scoping           | Nothing formal yet. Markets, instruments, and which of the four services are wanted                                                                                                                            | The sales desk and the client, in conversation                                        |
| 2 | Document collection              | Articles of incorporation, authorized signatory lists, UBO disclosures, tax forms such as FATCA W-8BEN-E or CRS self-certifications, proof of regulated status, investment mandates, source of funds or wealth | The onboarding team requests, the client supplies                                     |
| 3 | Identity and entity verification | Passport and proof of address for each authorized signatory and each beneficial owner                                                                                                                          | The onboarding team, against company registers, LEI databases and regulator registers |
| 4 | Screening                        | Nothing further. The names already supplied are the input                                                                                                                                                      | Compliance, against sanctions lists (OFAC, EU, UN), PEP lists and adverse media       |
| 5 | Risk classification              | Expected transaction volumes and business activity                                                                                                                                                             | Compliance, using jurisdiction, entity structure, activity and the screening output   |
| 6 | Compliance review and sign-off   | Whatever is missing or inconsistent, usually over several rounds                                                                                                                                               | A compliance officer, who approves, requests more information, or rejects             |
| 7 | Account setup and activation     | Settlement instructions and user details                                                                                                                                                                       | Operations and IT, in the broker's internal systems                                   |

Three things in that table are what make the process heavy rather than merely long. The first is that the document pack in step 2 multiplies. A single asset manager may run dozens of sub-funds, and each one is a separate legal entity with its own pack. The second is step 3. The usual threshold for a beneficial owner is a holding of more than 25%, but who counts as a beneficial owner is applied differently across member states, which is one of the reasons the registers are hard to rely on [21], [22]. For a client with a layered ownership tree that turns step 4 into dozens of separate screening checks for one onboarding, and the work that does deal with layered ownership comes from UBO register analysis rather than a brokerage setting [23]. The third is step 6. For institutional clients with complex structures the review can take weeks and it loops back to the client several times. The delay is not only an internal cost. It also holds up the client's access to the service, and for firms operating across several EU countries the absence of one single process is a large part of why the review drags [20], [24]. High-risk clients are routed into Enhanced Due Diligence at step 5, which adds another round of investigation and senior compliance sign-off on top of that.

This is the manual version, and it is also where AI has the clearest entry points: document collection, identity verification, screening, and risk classification are repetitive and rules-

driven enough to be partly or fully automated. But the adoption is uneven, and that unevenness is itself part of the picture. Figure 2 maps the same seven steps against how far AI has actually reached into each one. The pattern is easiest to describe in three bands: a mature core in the middle of the flow, an emerging layer around the risk decision, and two ends that are barely AI-driven at all.

**The mature core: collection, verification, screening.** The middle of the flow is where AI is genuinely established. Document collection is one of the more mature use cases. OCR combined with NLP, often packaged together as Intelligent Document Processing (IDP), reads articles of incorporation, tax forms, authorized signatory lists, and beneficial ownership disclosures, and pulls the structured fields out of them [25], [26]. Ling et al. [25] show that combining OCR with NLP feature extraction and a classification model, with a human still in the loop for validation, improved processing efficiency by over 200% while keeping the error rate under 0.5%. Form logic engines also adjust the required document pack based on entity type and jurisdiction, so a low-risk client in a clean jurisdiction is not asked for the same depth as a high-risk one [24], [27]. Vendors here include Quantiphi (Dociphi) and Synapse Analytics (Doxter).

Identity and entity verification follow the same pattern. AI-driven document analysis matches the photo ID submitted by an authorized signatory to the underlying document, and biometric facial recognition matches the live person in front of the camera to that ID [28], [29]. For entities, AI agents pull from LEI registers, national company houses, and regulator databases automatically rather than relying on a compliance officer to do the lookup [27]. Sumsu and Trulioo are the most commonly named vendors for KYC and KYB. The output of this step is essentially a yes/no on whether the entity and the people behind it are who they claim to be.

Screening is the third mature step, and here the value of AI is specifically in cutting noise. Sanctions and PEP screening done with traditional rule-based string matching produces a very high false positive rate, because legitimate clients often share name fragments with sanctioned individuals, especially across non-Latin scripts [30], [31]. ML-based name




Figure 2: The onboarding flow with AI maturity per step.

matching reduces those false positives, and NLP-based adverse media monitoring scans news in multiple languages and flags negative coverage connected to the parties [32]. ComplyAdvantage is the most prominent vendor and reports cost reductions of up to 50%, though that figure comes from its own marketing material and should be read with care. What can be said with more support is that the published research reports measurable gains for ML-based name matching over rule-based matching [30], [31]. Keep in mind that the academic work on watch-list filtering is thin, and Alkhalili et al. [31] point at compliance risk as the reason there is so little of it. The work that does deal with layered ownership comes from analysis of UBO registers rather than a brokerage setting [23], so whether these gains hold up in a wholesale onboarding is not established.

**The emerging layer: risk classification and review.** One step further and the picture gets less settled. Supervised ML models score clients on jurisdiction, entity structure (opacity, use of SPVs, ownership depth), expected behaviour, and the output of the screening step [33], [34], and the score then drives which track the client goes onto: low-risk fast-track, standard, or EDD. Reite et al. [34] show that client risk classification models can predict which clients will later produce suspicious transactions, and that adding accounting and credit-score data improves this a lot. Synapse Analytics' Konan is one of the named risk-scoring tools. But this is also the step where the concerns about explainability and bias are sharpest. A risk score that cannot be explained to a client or a regulator is a problem [35], [36], and we return to it in the risks chapter.

Compliance review is the newest use case of all, and the interesting development here is agentic AI. AI agents identify missing or inconsistent documents, communicate directly with the client to request what is needed, and only escalate to a human compliance officer when the case actually needs judgment. Early academic work on agentic systems in this domain exists, but it is aimed at the screening side rather than the full review [37]. According to a 2025 case study from Nagent AI, an agent deployed for a wholesale bank's onboarding workflow saved relationship managers three to four hours per day on document chasing and follow-up. This is a single vendor case study rather than an independent evaluation, so it says more about what is being claimed than about what is proven. The sign-off decision itself is still made by a human. What the AI changes is how much of the file is already clean by the time it reaches that human.

**The manual ends: scoping and activation.** The two ends of the flow are, tellingly, the least AI-touched. Pre-onboarding scoping is still mostly a human conversation about scope, services and instruments. Some brokerages use predictive analytics on prospect data to flag which prospects are a good commercial fit, and AI-guided conversational assistants sometimes help clients through the initial digital intake [38], but the bulk of the value sits downstream. At the other end, account setup and activation is closer to

Robotic Process Automation (RPA) than to learned AI. Bots create CRM records, set up settlement instructions, provision market access, and configure user accounts across the brokerage's internal systems. RPA is rules-based and not really "AI" in the modern sense, but it slots into this step well and is often deployed alongside the AI tools in the earlier stages [25], [27]; one 2023 study [27] shows an RPA bot with ML capabilities cutting KYC document verification from around 30 minutes to under a minute. Vendors here include HighRadius for B2B financial onboarding and WonderBotz for high-volume account creation.

Taken together, AI adoption in onboarding is not a single tool but a stack of them, and the stack is deepest exactly in the middle of the flow. The mature core does real work, the risk decision and the review are moving but still carry governance questions, and the two ends stay manual or rules-based. The end-to-end process therefore improves less than any single step would suggest, a point the gap analysis returns to. The next subsection looks at what happens once the client is live, where the AI focus shifts from verification to monitoring.

### 3.2 Client monitoring

Client monitoring is a time consuming task that can be overlooked quite easily. This is because nothing visibly breaks when it is done badly. The on boarding either finishes or it does not, but monitoring just runs in the background for as long as the relationship exists, and the failure only becomes visible when a regulator asks about it or a counterparty defaults.

The biggest reason a brokerage firm needs to do client monitoring is the regulatory one. According to [39] the relevant frameworks are MiFID II, MiFIR, MAR (Market Abuse Regulation), and the AML directives. Currently also AMLD5 and AMLD6 and moving to the new EU AML Regulation. In practice these translate into four monitoring tasks that a wholesale brokerage has to keep running.

- **Transaction monitoring for money laundering.** The brokerage watches client activity for patterns that look like laundering or sanctions evasion and files a report when something is suspicious [33], [40]. The traditional setup is a rule engine with thresholds, for example a payment above a certain amount to a high-risk jurisdiction.
- **Market abuse surveillance.** A firm that arranges or executes transactions has to have systems in place to detect and report suspicious orders and transactions. This covers insider dealing and manipulation patterns such as layering, spoofing and wash trades [41], [42]. This is a wholesale-specific concern. In retail the regulator is mostly worried about the client harming themselves, in wholesale it is worried about the client harming the market.

- **Counterparty credit and exposure monitoring.** Between trade date and settlement date the brokerage carries the risk that the client cannot pay or deliver. So exposure, margin and collateral have to be tracked continuously, in some cases intraday [43]. The collateral side of this is covered in Chapter 5, because it sits closer to settlement operations than to the client relationship.
- **Ongoing KYC review.** The risk rating assigned at on boarding does not stay valid forever. Ownership changes, a director becomes a PEP, a jurisdiction gets added to a high-risk list. The classic approach is a periodic review on a calendar cycle, where the interval depends on the risk band of the client [20].

So where does AI actually come in. The pattern across all four tasks is the same. The detection layer is traditionally rule-based, rules produce a lot of noise, and AI is being introduced either to rank that noise or to replace the rules entirely.

- **Alert triage and false positive reduction.** Rule-based transaction monitoring is widely reported to produce very high false positive rates [40], [44]. What this means in practice is that an analyst spends most of their day closing alerts that were never going to be anything. This is the alert fatigue problem [45], and it is the most researched AI application in this whole subsection. Supervised ML models are trained on historically closed alerts and used to score or rank new ones, so the analyst sees the likely-true alerts first [46], [47]. The research from (Halford, et al) [46] is useful here because it reports a model that was actually put into production, where 19% of alerts were automatically put on hold and 18% were escalated straight to the second line, and the time to file a report dropped by 61%. Keep in mind that this does not remove the rules, it sits on top of them.
- **Behavioural anomaly detection.** Instead of a fixed threshold, an unsupervised model learns what normal activity looks like for a specific client and flags deviations from it. Common approaches are clustering, isolation forests and autoencoders [36], [48]. The appeal is that it can catch patterns nobody wrote a rule for. The problem is data. Labelled suspicious cases are scarce in this domain to begin with [33], [36], and a wholesale brokerage has a few hundred clients rather than a few million, so there is even less history per client to learn a baseline from. I would say this is one of the reasons the wholesale side lags behind retail here, although I have not found a study that measures it directly.
- **Network and graph analytics.** Institutional clients sit inside ownership structures, and the interesting signal is often in the links rather than in a single transaction. Graph-based methods are used to detect shared beneficial owners, circular flows between related entities, and counterparties that are connected through an in-

intermediate layer [23], [47], [49]. In the research of (Naser Eddin, et al) [47] a triage model using graph features on real bank data reduced false positives by 80% while still catching over 90% of the true positives. This maps onto the wholesale problem better than transaction-level anomaly detection does, because the complexity in wholesale is structural.

- **Perpetual KYC.** Instead of reviewing a client every one, three or five years, NLP-based adverse media monitoring and continuous list screening run in the background and trigger a review when something actually changes [32], [37]. This shifts the review from calendar-driven to event-driven. Fenengo, Napier AI and ComplyAdvantage are the vendors most often named here.
- **Communications surveillance.** For market abuse, NLP models are applied to trader chat, email and voice recordings to flag language that suggests collusion or information sharing [42]. Behavox, SteelEye and Nasdaq Trade Surveillance are examples of tools in this space.

The maturity here is lower than in on boarding, and the reason is interesting. In retail the business case for AI monitoring is volume. There are millions of clients and no realistic way to look at them by hand. In wholesale the client count is low, so the volume argument is much weaker. What I would argue is that the case for AI in wholesale monitoring rests on complexity instead of volume, and complexity is harder to sell to a budget holder than a headcount saving is. This is my own reading of the situation and not a finding from the literature.

There is also a difference in what a mistake costs. A retail false negative is one small account. A wholesale false negative can be a counterparty with a large open position, which makes firms conservative about letting a model close an alert on its own. On top of that, uncertainty about how a supervisor will assess the model is itself a barrier to deploying one [50]. What we see in practice is that models are used to rank and prioritize, and the decision to close a case stays with a human [46], [51].

### 3.3 Client Support

Wholesale clients don't ask the same kinds of questions as retail clients. A retail client might ask why their app is not loading or how to reset a password. An institutional client is asking things like whether a US fund can buy Chinese A-shares through the broker, what documentation a Polish regulator needs for a derivative trade, or what the cut-off time is for the Japanese market today. These questions are operational and jurisdictional, and getting them wrong has consequences. If a desk gives the wrong settlement cycle for Brazil, the trade fails and somebody pays a penalty.

To handle this, wholesale brokerages run dedicated information service teams. These are specialists who know specific markets, settlement procedures, country-specific regulations, and the tax forms required for each jurisdiction. A large broker might have separate desks for the Americas, Asia, and EMEA (Europe, the Middle East and Africa), with overlapping coverage so that someone is available during every market's trading hours. This is expensive. The work requires experienced people, and the knowledge mostly sits in their heads, in old emails, and in scattered internal documentation.

This is where AI is starting to be useful. A common application is a chatbot or AI assistant that answers routine market access and settlement questions by retrieving from internal knowledge bases, a technique known as retrieval-augmented generation [52]. If a client asks what the settlement cycle is for South Korean equities, or what tax form a German pension fund needs before trading US treasuries, the system can answer in seconds. What this means is that the information service team is no longer the first line of contact for every question. They become the second line, dealing with the harder cases where a human still needs to think. This is the same pattern that (Andrade, et al) [53] describe in a commercial bank, where the virtual assistant shortened the queues at the call centre and let the human agent take on the more complex cases. On top of that there is a real cost benefit. Maintaining global information service teams across multiple time zones is one of the larger operational costs in wholesale brokerage, and routing the routine queries to an assistant reduces the need for round-the-clock human coverage in every region [54], [55].

However, this is also one of the less mature use cases. The problem is the cost of being wrong. Large language models can hallucinate, producing fluent and confident answers that are factually incorrect or unsupported [56]. If a chatbot tells a client that a market has T+2 settlement when it is T+1, the client may act on that and the trade fails. In a wholesale context the trade size is large, so the penalty is large. This is why most firms currently deploy these tools internally first, as an aid to the information service team rather than directly to clients. A human still sees the answer before it goes out. Techniques such as retrieval grounding, confidence thresholds and forced citation of the source document could reduce how often this happens, but the research does not show that it can be removed entirely [56]. We come back to this in Chapter 4.

## 4. AI in Settlement Operations

Client facing AI applications attract a lot of attention. Because they give you a visual representation into future technologies. But the backbone of wholesale brokerage has an equally important domain for AI adoption which is, compared to the client life cycle underexplored in the literature. This backbone is trade settlement, reconciliation and fail management.

### 4.1 Trade confirmation and matching

The first step after a trade execution is the confirmation and matching. As a client you are able to buy or sell a security. The counter party will then send an email or different kind of message with a confirmation about the trade, this could be a SWIFT message. A SWIFT message is a standardized electronic message used by financial institutions to communicate trade and settlement instructions [12]. In chapter 3 we covered the client lifecycle. In this chapter we are going to cover the operational side of the back office. How and where the securities actually switch places.

Before going into the tooling it is worth being precise about what matching actually is, because the word gets used for two different things. Confirmation is the counterparty telling you what it thinks the trade was. Matching is the act of comparing that against what you think the trade was, field by field, and deciding whether the two records describe the same trade. Nothing settles until they do. What this means is that matching is a gate. A trade that is executed but not matched is not a trade that is on its way to settlement, it is a trade that is stuck.

The comparison is done on a fixed set of economic and settlement fields. The ones that have to agree are:

- **Instrument:** which security, usually identified by ISIN or CUSIP.
- **Direction:** who is buying and who is selling.
- **Quantity:** how many shares or what nominal amount.
- **Price:** the executed price, and from that the gross consideration.
- **Trade date and settlement date:** when it was executed and when it is due to settle.
- **Fees and net amount:** commission, taxes, and the final cash figure that will actually move.
- **Settlement instructions:** which account at which custodian the securities and

the cash should go to. These are the SSIs from Figure 1 in chapter 2.

Not every field is treated the same. A difference in the instrument or the quantity is a real economic break and the trade cannot proceed. A difference of a few cents in the net amount, caused by a rounding difference or an FX rate applied at a slightly different time, is usually inside a tolerance and the match is allowed through. So matching is not a plain equality check, it is a comparison with a tolerance per field, and setting those tolerances is itself an operational decision.

There are also two ways of doing the comparison. In *local matching* both sides send their version to each other and each one compares against its own book, which means the same work is done twice and any disagreement has to be resolved over email or phone. In *central matching* both sides submit their version to a shared utility, and that utility does the comparison and tells both sides the result at the same time. Central matching is faster and it removes the argument about whose version is correct, which is why it is the direction the industry has moved in. The most widely used central matching utility for institutional trades is DTCC's CTM. Keep in mind that CTM is a trade processing service and not a CSD. The CSD comes later in the chain and handles the settlement itself, which is a point we return to below.

Markets are moving from T+2 to T+1. The US did this on 28 May 2024 [57], [58]. The UK and the EU are planning their own transitions for late 2027 [58]. Under T+1, the window between trade execution and settlement is just one working day. This changes the timing of matching. Under T+2, trades could be matched overnight in batches, since there was a spare day before settlement. Under T+1, that slack disappears: matching has to happen the same day, intraday or close to real-time. For example, a trade executed at 16:00 on the New York Stock Exchange (NYSE) could previously be matched at the end of the day. Now it has to be affirmed on trade date, with the industry deadline set in the evening of the same day.

It is worth noting that shortening the cycle is not purely an improvement. The research from (Gandy, et al) [59], based on 44 interviews with firms in settlement and operations roles, found a perceived trade-off between the efficiency gains of a shorter cycle and an increased risk of settlement failure. Less time to match means less time to fix a break before the settlement date arrives. This is exactly why the automation question became urgent rather than optional [57].

The traditional matching process is a mix of channels. But the most common are: SWIFT messages, FIX, Bloomberg messages or email attachments. The first two are structured, so a machine can read them directly. The last two are not, and for those the back office has to read the confirmation and key the comparison in by hand [57], [60]. Different

counterparties use different formats. This means a broker may communicate differently with different clients. Here we can already see the first use case of NLP-based extraction. AI models read the unstructured confirmations (email bodies, attached PDFs, Bloomberg chat) and pull out the structured fields listed above so they can be submitted to the matching utility. A company that is already doing this is Xceptor. ClearDox does something similar but more centered at commodity confirmations.

Once the data is structured it is submitted to the central matching utility. After the trade is matched, the settlement instructions flow on to the CSD, which is the infrastructure that holds the official ownership record and performs the transfer of securities against cash. Not every trade goes through central matching. Trades on alternative venues that self-clear, and very illiquid names, are often still matched bilaterally. As seen in Table 3 there are different CSDs for different markets.

Table 3: Major Central Securities Depositories by market

| <b>Market</b>                  | <b>CSD</b>                                                        |
|--------------------------------|-------------------------------------------------------------------|
| United States                  | Depository Trust Company (DTC, part of DTCC)                      |
| Belgium / France / Netherlands | Euroclear                                                         |
| Germany / Luxembourg           | Clearstream                                                       |
| United Kingdom                 | Euroclear UK & International (CREST)                              |
| Japan                          | Japan Securities Depository Center (JASDEC)                       |
| Hong Kong                      | Hong Kong Securities Clearing Company (HKSCC)                     |
| China                          | China Securities Depository and Clearing Corporation (ChinaClear) |

The matching utilities and the CSDs both run on rules-based software, and their value comes from operational engineering and network effects rather than from machine learning. A matching engine that compares seven fields against a tolerance does not need a model. But the layer in front of it, the bit that turns a PDF into a structured submission, is where AI actually plays a role.

The second use case is error handling. When two records do not match, somebody has to figure out why. ML-based classifiers learn from historical error patterns and route exceptions to the right team without manual triage [61]. Some platforms also use AI to highlight clauses in OTC confirmations that deviate from a standard template, so a reviewer is pointed at the lines that need attention rather than reading the whole document. The technique behind this is clause-level classification against a reference template, which has been demonstrated for contract review in other domains [62], [63]. This is closer to assisted review than full automation.

As we can see some aspects are more developed than others. Goldman Sachs reported a reduction in same-day affirmation exceptions and settlement fails in the quarter after

rolling out CTM's Match-to-Instruct workflow, and buy-side adoption of the same workflow grew 91% in 2023. The AI layer sitting on top of those utilities is newer. NLP-based confirmation extraction works well for structured asset classes like cash equities. It still struggles with bespoke OTC documents where the language is non-standard and the cost of mis-parsing is large. Most wholesale firms today run these tools in a hybrid setup where the AI proposes the extracted record and a human still signs off on higher-value or non-standard trades.

## 4.2 Reconciliation

Reconciliation is the process of making sure that the firm's own record of its trades, positions, and cash agrees with the records held by every other party in the chain. A wholesale broker needs its internal book to match what the custodian holds, what the prime broker reports, what the exchange or CCP confirms, and what the counterparty thinks happened. When all of these agree, the firm has a reliable picture of where it stands. When they do not, the difference is called a break and somebody has to figure out why.

There are several types of reconciliation that run in parallel. Cash reconciliation checks that the cash balances held at banks match the firm's internal cash ledger. Position reconciliation checks that the firm's record of how many shares of a given security it holds matches what the custodian reports. Trade reconciliation checks that executed trades have been booked correctly and reported to the right downstream systems. Each of these runs daily, sometimes multiple times a day, and each can produce breaks that the operations team has to investigate. Broadridge estimates that reconciliation activities account for 30 to 40% of total back-office labor costs at financial firms, although this figure is from their own white paper and should be read as an industry estimate rather than peer-reviewed data.

The reason reconciliation is so labor-intensive is the variety of data sources. The firm's internal booking system speaks one format. The custodian sends statements in another. The prime broker uses yet a different one. SWIFT messages, CSV exports, PDF statements, and proprietary file formats all arrive on different schedules. Before any matching can happen the data has to be normalised, which means parsing it, mapping the fields to a common schema, and handling missing or inconsistent values. Once the data is normalised, traditional reconciliation engines apply rule-based matching: this trade ID equals that trade ID, this amount within a tolerance of that amount, this date matches that date. Everything that does not match becomes a break, and from there a human investigates.

Reconciliation is one of the areas where AI vendors are most active and where the mar-

keting language is most generous. A closer look at what the tools actually do is worth the effort, because a meaningful share of "AI reconciliation" is a better rule engine with a machine learning layer on top.

The honest assessment is that most reconciliation work in a wholesale broker can be handled with proper IDs and proper rules. If trades carry consistent reference numbers across systems, or if the systems agree on at least a stable suffix, a rule can match them. If the activity in an account is single-purpose, a daily sum compared against the bank statement reconciles it. If the bank fees follow a known fee schedule, the expected total can be calculated and compared directly. None of this needs machine learning. It needs operational discipline and well-designed reference data.

AI starts to genuinely add value in a smaller set of cases. The first is tolerance bands that are not fixed. The research from [64] shows that learned anomaly thresholds outperform fixed rules on general ledger data, and the same principle applies to FX-driven amount differences in reconciliation. A learned tolerance per currency pair or per counterparty catches more real matches than a single band applied across the board. The second is one-to-many groupings where there is no shared reference, for example one invoice covering many trades with no per-trade breakdown. A model trained on historical groupings can propose the grouping where a rule cannot. The third is onboarding new data sources, where the firm does not yet have rules for a particular custodian or product and needs to start matching immediately. The fourth is anomaly detection on already-matched records, where the model flags matches that the rule accepted but that do not fit the firm's historical pattern.

Some area are more mature than other in the client onboarding and settlement flow. This is also seen in the matching subsection. Rule-based reconciliation engines such as Broadridge, Duco, NeoXam, and SmartStream have been in production for decades and are widely used. The AI layer sitting on top of them is newer and the genuine value of that layer depends heavily on the state of the firm's underlying data. The academic literature on AI applied to securities reconciliation specifically is also thin. Most peer-reviewed work sits in adjacent domains: general ledger anomaly detection, payment reconciliation, and transaction fraud detection. What work does exist on trade exception handling in capital markets [61], [65] sits in low-impact venues and has not been widely cited, which is itself part of the picture. The techniques transfer, but the absence of established brokerage-specific research is notable and is a point this survey returns to in chapter 7.

### **4.3 Risk and collateral management**

This subsection covers the operational side of risk and collateral handling in a wholesale broker: tracking exposures, processing margin calls, and managing intraday cash and

securities flows. The related topic of counterparty credit risk monitoring is part of the client monitoring discussion in chapter 3 and is not repeated here. The split is deliberate. Credit risk monitoring sits next to KYC and client surveillance because it is about understanding the counterparty. Operational risk and collateral handling sit in the back office because they are about moving cash and securities on time.

The operational problem can be described in plain terms. Some trades close almost immediately. If a client buys shares today, the trade settles a day or two later and the relationship is over. Other trades stay open for months or years. A swap, a repo, or a securities loan keeps running, and during that time the value of the trade keeps changing as the market moves.

This is where collateral comes in. While the trade is open, the side with the loss has to post collateral, which is cash or high-quality securities held by the other side as protection in case of a default. The amount of collateral is recalculated every day based on the current value of the trade. If the market has moved overnight, more collateral has to be posted. This is called a margin call, and the deadline for delivering the assets is short, usually the same business day.

A wholesale broker operating in an agency model does not hold these positions itself. The client holds the position. The broker handles the operational side of the process on the client's behalf: calculating what is owed, instructing the movement of assets, reconciling the positions afterwards, and making sure cash and collateral are in the right account at the right time. The broker is not exposed to the market move, but it is exposed to operational risk. A missed margin call, an incorrect calculation, or a delayed settlement instruction can still cost the firm money in penalties, reputational damage, or client claims.

There are two operational areas where AI is starting to appear in this work.

The first is intraday liquidity and cash flow forecasting. A wholesale broker has cash moving in and out of its accounts throughout the day: client settlements, broker fees, custody charges, FX conversions, and so on. The treasury team needs to know whether cash will be available when an outgoing payment is due. The research from (Singh, et al) [66] applies machine learning to liquidity forecasting at corporate and subsidiary level and outperforms the classical in-house approach in six out of nine forecasts, which is a real improvement but not a clean sweep. The same idea is applied at the intraday level, where the goal is to estimate the cash position at each point in the day so that payments can be released or held back accordingly. (Liermann, et al) [67] do this by detecting patterns in intraday flows per customer and aggregating the forecast by currency. More recent work looks at whether a general purpose language model can handle the decision layer

on top of that forecast, balancing settlement speed against liquidity usage in a wholesale payment system [68]. This is a time-series prediction problem with a clear ground truth, which is the realised cash flow, and it is the kind of problem where ML tends to do well.

The second is margin call processing. A typical wholesale broker receives margin call notifications from multiple counterparties and custodians, each in their own format. Bloomberg messages, SWIFT messages, PDF statements, and emails all arrive in the same morning window. The broker has to read each one, identify which client account it applies to, verify the figure against its own calculation, and instruct the settlement. NLP-based extraction and classification tools, the same family of tools described in the trade confirmation subsection, are being applied here too [61]. The broker still validates and signs off, but the first pass of reading and routing the notice is handled by software.

Both of these are useful but neither is dramatic. Liquidity forecasting is an improvement over manual or spreadsheet-based forecasting. Margin call processing is an improvement over manual reading of incoming notices. In both cases the AI sits on top of an existing process rather than replacing it.

The academic literature in this area is also limited for the agency case. Most of the published work on machine learning in margin and collateral is on the principal side and assumes the firm runs its own derivatives book. The papers on initial margin approximation [69], [70] are about banks calculating SIMM for their own portfolios, not about brokers processing margin calls on behalf of clients. The result is that the AI value proposition for an agency wholesale broker, while real, is narrower than the headline literature suggests. It mostly comes down to forecasting cash and reading documents faster.

## 5. Risks and Limitations

Up to now this survey has described where AI is being applied in wholesale brokerage. In this chapter we will discuss where the technology falls short, and where its deployment creates risks that did not exist before.

A theme runs through several of the earlier chapters: AI tooling is real but often less mature than vendor marketing suggests. The reconciliation subsection in chapter 4 made this point explicitly. Four risks are discussed, each specific to the wholesale context. The order moves from the industry wide picture down to the level of an individual procurement decision.

### 5.1 Vendor concentration and lock-in

Most wholesale firms access AI through third-party vendors rather than building it themselves. Financial organisations integrate AI in one of three ways, in-house, outsourced, or through an ecosystem of partners, and which of those is realistic depends heavily on the scale of the organisation [71]. Large banks with significant engineering capability build internally for some use cases, but for the majority of wholesale firms, particularly mid-sized agency brokers, the AI capability described in chapters 2 and 3 reaches them through vendor products. Securities market participants are becoming more reliant on third-party providers for tasks they did not previously outsource, and AI is a large part of why [72]. This is true across the range of use cases: KYC tooling, document parsing, reconciliation, margin processing, and intraday liquidity forecasting.

The list of vendors is short. For KYC and onboarding, names like ComplyAdvantage, NICE Actimize, Refinitiv, and LexisNexis dominate. For reconciliation, Broadridge, Duco, NeoXam, and SmartStream cover most of the market. For trade confirmation, Xceptor and DTCC's ITP services hold most of the wholesale flow. The underlying LLMs that power the newer AI features come from an even smaller group: OpenAI, Anthropic, Google, with some open-source models from Meta in the mix. A wholesale firm running modern AI tooling is therefore three layers deep in dependencies, and each layer is concentrated in its own right.

This creates several risks. The first is operational. If a vendor has an outage or a data breach, the firm cannot easily switch to a competitor because the integration cost is high and the data formats are not interchangeable. The research from (Gozman, et al) [73], based on regulatory data and interviews with technology risk specialists at the FCA, points at exactly this: the absence of interoperability standards between providers is one of the main things that turns a supplier relationship into a concentration risk. The second is commercial. A vendor with limited competition can raise prices, change licensing terms,

or deprecate features the firm relies on. The third is systemic. If a single vendor has an outage across its client base, the wholesale market loses a function at the same time across many firms [74]. This is a different shape of risk from the vendor concentration that already exists in market data (Bloomberg, Refinitiv) because the AI vendor is making decisions, not just providing information.

Keep in mind that concentration risk is not new and it is not unique to AI. [74] makes the point that the same argument was made about cloud adoption, and that it is not obvious the risk would be smaller if firms ran everything on their own infrastructure. The question is not whether to depend on third parties but how that dependency is managed.

Regulators have started to notice this. The EU's Digital Operational Resilience Act (DORA), which applies from 17 January 2025, establishes an oversight framework for critical third-party ICT providers in financial services [75]. Supervisory attention is also widening from outsourcing in the narrow sense to operational resilience across all third-party relationships, which is where AI vendors sit [72]. Whether DORA and similar frameworks are enough to address AI-specific concentration is something this survey returns to in chapter 6.

## **5.2 Hallucination and confident wrong answers**

The most distinctive failure mode of modern AI, particularly LLM-based systems, is the production of outputs that are confidently stated but factually wrong. The model has no reliable internal signal that it is uncertain, so a wrong answer arrives in the same confident tone as a correct one [56]. The term for this is hallucination, and it has become a research field in its own right [76].

In a wholesale context the cost of being wrong is large because trade sizes are large. Chapter 2 raised this point about client-facing chatbots. The same problem appears across most of the AI use cases in chapter 4. An NLP model extracting trade details from a confirmation email might misread a settlement date by a day. An LLM generating a break narrative might attribute the difference to the wrong cause. A document parsing tool might pull a notional figure from the wrong section of a PDF. In each case the output looks correct on the surface, and a human reviewer working through hundreds of items in a morning may not catch the error.

This is not only a problem with free-form text generation. Hallucination also shows up in structured information extraction, which is the task most of the tools in chapter 4 are actually doing. A study on extraction from medical texts found that models handled some fields consistently and were noticeably less reliable on others [77], so the error rate is not uniform across the fields you are pulling out. On top of that, extraction accuracy degrades when the input document does not look like what the model expects. The

research from (Biswas, et al) [78] tested three current multi-modal models on documents that were simply rotated slightly, the kind of skew you get from a scanned page, and found that extraction accuracy dropped for all of them and that hallucinations increased. Wholesale documents are not standardised. Each counterparty has its own template and each custodian its own statement layout. This is exactly the condition under which hallucination is most likely.

The mitigation in practice is to keep a human in the loop. Most firms run AI tools as assistants rather than autonomous agents. The AI proposes a value, a human signs off, the work gets done. There is a large body of work on technical mitigation as well, from retrieval grounding to uncertainty estimation [79], but none of it removes the problem entirely. So the human review stays, and it limits the efficiency gain. The AI cannot fully replace the operations analyst because the analyst is still required to verify each output. The headline "AI processes confirmations automatically" needs the footnote "with human review of every record above a certain threshold". The same pattern appeared in the reconciliation and margin call discussions in chapter 4.

### **5.3 Opacity and accountability gaps**

A question in any regulated process is who is responsible when something goes wrong. The traditional answer involves a chain of human accountable owners: an analyst makes a decision, a manager approves it, a compliance officer reviews it, and the record of who decided what is preserved for the regulator. AI complicates this chain.

The opacity issue is twofold. The first is technical. Modern ML models, especially deep learning models, are not easily interpretable. A model trained to detect suspicious transactions can flag a counterparty as high-risk, but explaining exactly why, in terms a compliance officer can defend to a regulator, is harder than for a rule-based system, since deep-learning models are widely characterised as opaque "black boxes" [80]. The second is procedural. Even when an explanation can be produced, the question of who is accountable for the decision becomes blurred. The data scientist who trained the model, the vendor who built it, the operations analyst who acted on its output, and the firm's compliance function are all involved. None of them has full visibility into the others. This is the accountability gap that [81] describes for algorithmic systems in general: the decision is distributed across parties in a way that the existing accountability mechanisms were not designed for.

This matters because regulators expect a named human owner for material decisions. Under the UK Financial Conduct Authority's Senior Managers Regime and similar frameworks in the EU, individuals are personally accountable for the decisions taken under their function. (Zetsche, et al) [82] argue that these post-crisis manager responsibility systems

are in fact the most promising external framework for governing AI in finance, precisely because they force a human back into the loop by making a named person answerable for what the system does. If an AI declines a counterparty incorrectly and the firm is sued, the accountable manager has to explain the decision. "The model said so" is not a defence. The same applies in reverse. If an AI fails to flag a counterparty that should have been flagged, the manager is still accountable for the gap.

#### **5.4 The gap between vendor marketing and actual capability**

The final risk is the gap between what is marketed and what is actually delivered. There is a name for this in the literature. It is called AI washing, and it means exaggerating, misrepresenting, or superficially signalling AI capability in order to gain a commercial or reputational advantage [83]. The term is borrowed from greenwashing and the mechanism is the same.

A meaningful share of what is sold as "AI" in wholesale brokerage is not AI in the modern sense. It is configurable rule engines, fuzzy string matching, optical character recognition, and statistical models that have existed for decades, packaged with a UI and marketing language that suggests machine learning. (Elsayed, et al) [83] call this technical capability inflation and treat it as one of four distinct types of AI washing, alongside marketing, strategic signalling, and governance-based washing. This was discussed in the reconciliation subsection in chapter 4 and the pattern generalises. A vendor product labeled "AI-powered" may contain a small ML layer sitting on top of an otherwise traditional system.

The incentives explain why. The label is commercially valuable. Firms want to be seen to be adopting AI. Vendors want to sell at AI prices. What the research shows is that this works in the short run and not in the long run. The research from (Song, et al) [84] looked at US firms between 2018 and 2023 and found that unsubstantiated AI claims first attract investors and then produce negative market reactions and sustained underperformance. In marketing research the same cycle is described as AI washing followed by AI booing, where the backlash comes from expectations that were never going to be met [85].

The risk for a wholesale firm is not that these tools do not work. Many of them work fine, they are just not doing what the label says. The risk is that the firm makes decisions on the basis of what it was told the tool does. A process gets redesigned around an expected improvement, headcount is planned against it, or a manual control is removed because the model is supposed to catch the same thing. The gap only becomes visible later, either when something fails or when the expected improvement does not show up in the numbers. By that point the old process is gone.

## 6. Regulatory Landscape

The regulatory response to AI in financial services has picked up speed over the last few years. This is driven by concerns over market stability, consumer protection, and systemic risk [86]. However, most of the regulatory guidance that exists today was written with retail investor protection as its main concern. What this means is that wholesale brokerage sits in a kind of regulatory grey zone. The rules that govern AI in finance were mostly not built for a world of professional counterparties trading with each other, and that mismatch is the argument of this section.

### 6.1 MiFID II

The starting point for any discussion of trading regulation in Europe is MiFID II, the Markets in Financial Instruments Directive, which was adopted in 2014 and has applied since January 2018 [5]. MiFID II sets out requirements around algorithmic trading, best execution, and transaction reporting. Firms that use algorithms to trade have to have systems and controls in place to make sure those algorithms behave, including testing before deployment, kill switches, and limits on order-to-trade ratios [87]. On best execution, firms have to take all sufficient steps to get the best possible result for their clients, and they have to be able to show it. On top of that, transaction reporting requires firms to report the details of their trades to regulators so the market can be monitored for abuse. Together with the Market Abuse Regulation, this was the European response to the technological risks of automated trading [88].

The problem is timing. MiFID II predates the current wave of machine learning. It was designed for the rule-based, deterministic algorithms that dominated automated trading at the time, not for models that learn from data and whose behaviour is not fully predictable in advance [86]. The word “algorithm” in MiFID II does not map cleanly onto a neural network, and the directive gives little guidance on how a firm is supposed to test or explain a model whose logic it cannot fully trace. The research from (Raschner, et al) [89] looks specifically at the MiFID II algorithm testing requirements and finds that although the regime is prescriptive in tone, the parts that actually matter for controlling an algorithm are left too vague, and that this gets worse for systems based on AI or ML.

There is a deeper reason for this than just age. MiFID II is built on the principle of technological neutrality, which means it regulates the activity rather than the technology used to perform it. (Azzutti, et al) [90] argue that this neutrality makes the framework functionally enabling but normatively silent: it does not stop a firm from deploying AI, but it also says nothing specific about the risks that AI brings. Because the obligations attach to activities rather than to systems, the framework struggles to deal with AI applications whose risk profile varies enormously depending on where in the firm they sit. That is a

useful way to think about the whole of this chapter. The rules are not absent, they are just aimed somewhere else.

## 6.2 ESMA Guidance

The European Securities and Markets Authority (ESMA) is the body that coordinates supervision across the EU, and it has started to look at AI more directly. ESMA has published work on the use of AI in investment services and has flagged the risks around transparency, governance, and the reliance on third-party model providers [91]. It has also issued guidance specifically on algorithmic trading under MiFID II [87]. However, the focus of most of this work is again on the retail side. The main worry ESMA keeps returning to is the protection of retail investors, for example the risk that AI-driven advice or chatbots give unsuitable recommendations to individuals who cannot judge the quality of that advice for themselves [91]. For a wholesale broker dealing with asset managers and hedge funds, that concern does not really apply. The clients are professionals and are assumed to understand what they are doing, which is the client classification logic we saw in chapter 2. So while ESMA's guidance is useful, a large part of it is aimed at a risk that wholesale brokerage does not carry.

## 6.3 The AFM

In 2019 De Nederlandsche Bank (DNB) published a discussion paper setting out general principles for the responsible use of AI in the financial sector, built around soundness, accountability, fairness, ethics, skills and transparency [92]. These six are usually referred to as the SAFEST principles. The AFM joined that work later. In 2024 the two supervisors published a joint report on the impact of AI on the financial sector and on supervision, which sets out what they expect from firms that deploy AI and how they intend to supervise it [93].

More recently the AFM has been vocal in the wider conversation about AI regulation in the Netherlands and about how the EU AI Act will land on financial firms. The AFM's position is that existing sectoral rules, MiFID II and the rest, already cover a lot of what AI does in finance, and that the AI Act sits on top of these rather than replacing them [93]. What this means in practice is that a Dutch wholesale broker deploying AI has to read several frameworks at once and work out how they overlap. This is not a Dutch problem specifically. The interaction between the horizontal AI Act and the existing sector-specific rules is the main compliance difficulty identified for European banks more broadly [94]. The AFM's guidance is helpful here, but even the AFM's material is written mostly with the protection of the end investor in mind. The specific operational risks of wholesale brokerage, for example an unsupervised model making settlement decisions on trades worth hundreds of millions, are not really the thing the guidance was written to

address [95].

## 6.4 The EU AI Act

The newest piece of the picture is the EU AI Act, which entered into force in 2024 and is being phased in over the following years [96]. The AI Act takes a risk-based approach. It sorts AI systems into categories, unacceptable risk (banned outright), high risk (heavily regulated), limited risk (transparency obligations), and minimal risk (largely unregulated) [96]. For financial services the important category is high risk. Systems used for creditworthiness assessment and credit scoring of natural persons are classed as high risk and face strict requirements around data quality, documentation, human oversight, and transparency [96], [97].

However, once again the high-risk trigger is tied to consumer-facing use. Credit scoring of a natural person is high risk. A model that groups unmatched settlement records for an institutional client is not obviously caught at all. This is not only my reading of the AI act but also that of Martins. (Martins Pereira, et al) [98] argue that the notion of high risk at the centre of the AI Act leaves out financial systemic risk altogether, and that this exclusion cannot be justified either by technological neutrality or by proportionality, because the systemic risk created by financial AI is not covered by the existing macroprudential toolkit either. In other words the gap is real and it was noticed by legal scholars before it was noticed by supervisors. More broadly, [99] argues that several important provisions of the AI Act do not actually follow the risk-based logic the Act claims for itself. So the AI Act, for all its detail, still leaves most wholesale brokerage AI in the lightly regulated tiers, judged by a risk framework that was calibrated around harm to individuals.

## 6.5 DORA

One framework in this list was not written with the retail investor in mind at all. The Digital Operational Resilience Act (DORA), which applies from January 2025, is about operational resilience: ICT risk management, incident reporting, resilience testing, and an oversight regime for critical third-party providers [75]. Chapter 5 discussed it in the context of vendor concentration. It matters here because it is the one piece of the picture whose concern is not whether a client was treated fairly but whether the firm can keep operating.

That makes DORA the most natural home for the wholesale AI risks described in this survey. A model embedded in a settlement workflow is an ICT dependency, and a vendor supplying that model is a third-party provider. Legal analysis of the AI Act has pointed at exactly this, arguing that the ICT risk governance obligations under DORA are an effective route for addressing the gaps the AI Act leaves open in finance [100]. Keep in

mind though that DORA is not an AI framework. It says nothing about model validation, explainability, bias, or whether an output can be defended to a regulator. It treats a machine learning model the same way it treats a database. So DORA covers the wholesale context but not the AI-specific risk, which is close to the mirror image of the problem with the AI Act.

## 6.6 The Wholesale Gap

Putting these frameworks side by side, a pattern shows up, although it is a more specific pattern than it first appears.

The obvious version of the argument is that everything discussed here has retail investor protection near its centre. That is broadly true of MiFID II's conduct rules, of ESMA's AI work, of the AFM and DNB material, and of the AI Act's high-risk triggers. But it would be too strong to say that no European regulation addresses the wholesale context. Chapter 3 covered the Market Abuse Regulation, which is about market integrity rather than client protection. It also covered CSDR, whose settlement discipline regime is aimed squarely at institutional post-trade operations. And DORA, above, is about the firm's own resilience. None of these is a retail framework.

So the gap is narrower and sharper than “the rules are all about retail”. It is best seen along two axes at once, which is what Figure 3 sets out. One axis asks whether a framework is specifically about AI or silent on it. The other asks whether it is aimed at harm to individuals or at wholesale operations. Sorting the frameworks this way, the frameworks that are specifically about AI, the AI Act, ESMA's guidance, the AFM and DNB work, all sit on the retail side, calibrated around harm to individuals. The frameworks that take wholesale operations seriously, MAR, CSDR and DORA, all sit on the silent-on-AI side, treating a model no differently from any other system. The two bodies of rules pass each other without meeting.

The empty top right of the figure is the whole point. Lets look at two examples. A model that scores the creditworthiness of a private individual counts as high risk under the AI Act. A model that decides which settlement needs to get human attention does not. However, that second model sits inside the settlement process, and that is a process CSDR takes seriously enough to fine firms over. So there is no rule that is both AI-specific and aimed at wholesale operations. And that empty space is exactly where the risks from the earlier chapters live.

Those risks sit in that empty quadrant: unsupervised algorithmic trading, models making or informing settlement decisions on very large trades, and weak governance around models the firm did not build itself [86], [101]. This is not a criticism of the regulators. The harm to a private individual who is mis-sold a product is exactly the kind of thing




Figure 3: AI-specific frameworks are calibrated around individual harm; wholesale-focused frameworks say nothing about AI. The two bodies of rules pass each other without meeting, which leaves the top-right quadrant, AI-specific rules for wholesale operations, empty.

regulation should protect against, and it makes sense that the rules grew up around it. It is also not clear that adding more rules is the answer, since a comparison of AI oversight across nine jurisdictions found that most of them already have the vocabulary of trustworthy AI without the enforcement mechanisms to back it [102].

What it does mean is that a wholesale broker deploying AI today is not really breaking any rules, because the rules that would apply to what it is doing have not been written yet. It is left to read across frameworks built for a different kind of client, or for a different kind of risk, and work out for itself where the lines are, filling in by hand the quadrant the regulation has so far left blank.

## 7. Gap analysis

The previous chapters described where AI is being applied in wholesale brokerage, where it falls short, and how the regulatory frameworks line up against it. This chapter names what is missing, split into three kinds: gaps in what firms have adopted, gaps in what the research covers, and gaps in what the rules address.

### 7.1 Adoption gaps

The first thing that stands out across chapters 2 and 3 is that AI adoption in wholesale brokerage is uneven. Adoption is mature in some steps and barely present in others, sometimes inside the same process. On the client side, document processing [25], [26] and sanctions screening [30], [31] are mature and widely used, risk classification and agentic review are newer and still raise governance questions [35], and the first and last steps of the flow, pre-onboarding scoping and account setup, are barely AI-driven at all. On the settlement side the same pattern shows up. NLP-based extraction works well for structured asset classes but struggles with bespoke OTC documents where the language is non-standard [78]. Reconciliation is a thin AI layer on rule-based engines that have been in production for decades, and the value of that layer depends heavily on the firm's underlying data **kumar2023payment**, [64]. Margin call processing and liquidity forecasting are useful but narrow [66], [67]. In every case the AI sits on top of an existing process rather than replacing it.

What this means is that the honest picture of adoption is a stack of assistive tools, not autonomous systems. The mature steps pull ahead while the manual steps stay where they are, so the end-to-end process improves less than the individual gains suggest. And because the cost of a wrong answer is large in a wholesale setting, most firms keep a human signing off every material output [56], [79], the same pattern the monitoring literature shows [51]. The analyst is still required, which puts a ceiling on how much of the work can actually be taken away.

### 7.2 Research gaps

The academic research on AI applied to wholesale brokerage specifically is thin. The work exists in adjacent domains, but it does not sit squarely on the wholesale back office. For reconciliation, most peer-reviewed work is on general ledger anomaly detection [64], payment reconciliation **kumar2023payment**, and transaction fraud detection, and what does exist on trade exception handling sits in low-impact venues [61], [65]. For risk and collateral, the published work assumes a principal firm running its own derivatives book **mostafa2021im**, [69], not an agency broker processing margin calls on behalf of clients. The monitoring literature is built on retail-scale data, where a bank has millions of clients

to learn a baseline from [33], [36], while a wholesale brokerage has a few hundred. And the broader AI-in-finance literature leans toward robo-advice, retail chatbots [53], [55], and consumer credit scoring [43], [71]. In each case the research that exists answers a different question than the one a wholesale broker is asking.

This is a real gap and not just an academic complaint, because numbers from a different setting do not transfer cleanly. A study showing large gains from machine learning in payment reconciliation **kumar2023payment** was run in an environment where reference data was poor, and a wholesale operation with clean IDs and maintained rules will not see the same gain. Firms that read the headline number and expect it to hold in their own setup are being misled by a research gap, not by the vendors alone.

### 7.3 Regulatory gaps

The obvious version of this gap, that every framework governing AI in finance has retail protection at its centre, is too strong. It is broadly true of MiFID II's conduct rules, ESMA's AI work [91], the AFM and DNB material [92], [95], and the AI Act's high-risk triggers [96], [97]. But MAR is about market integrity, CSDR's settlement discipline regime is aimed at institutional post-trade operations, and DORA is about the firm's own resilience. So the gap is narrower and sharper: the frameworks that are specifically about AI are calibrated around harm to individuals, and the frameworks that are specifically about wholesale operations say nothing about AI. The two bodies of rules pass each other without meeting.

The mismatch is not subtle once you look for it. MiFID II predates the current wave of machine learning, was written for deterministic algorithms, and regulates the activity rather than the technology, which leaves it enabling but silent on AI-specific risk [86], [89], [90]. ESMA's guidance keeps returning to retail investors, a concern that does not really apply to a broker dealing with asset managers and hedge funds [91]. The AI Act ties its high-risk trigger to consumer-facing use, so a model scoring an individual's creditworthiness is high risk while a model grouping unmatched settlement records for an institutional client is not obviously caught at all [96], [98]. The result is that the risks this survey actually found, an unsupervised model informing settlement decisions on very large trades, weak governance around bought models, and vendor concentration, sit outside the main focus of every framework discussed [98], [101]. DORA is the one framework that moves in the right direction, because its critical third-party provisions speak directly to vendor concentration [72], [75] and legal analysis has argued its ICT risk governance is the most workable route for closing the gaps the AI Act leaves open in finance [100]. But it addresses resilience, not opacity or accountability. It treats a machine learning model much the same way it treats a database.

## 7.4 How the gaps interact

These three gaps are not independent. Because there is little independent evidence on what works in a wholesale setting, firms lean on vendor claims when they decide what to buy, which is the AI washing problem from chapter 6 [83], [84]. Because the tools are then deployed with a human signing off, there is no clean measurement of what the model would have done on its own, so the evidence base stays thin. The regulatory gap sits underneath both: as long as wholesale AI stays in the lightly regulated tiers, there is no external pressure to keep the kind of audit trail that would produce evidence in the first place, a pattern also visible across jurisdictions that have adopted the vocabulary of trustworthy AI without the supervisory mechanisms to enforce it [102]. There is also a second-order effect: uncertainty about how a supervisor would assess a model is itself a reason firms hold back from deploying one [50], so a vague regime not only fails to produce evidence, it slows the adoption that would generate it.

Closing these gaps therefore needs more than better tools. It needs measurement of AI performance in wholesale operations that is specific enough to compare across firms, and a framework that asks for that measurement. The conclusion returns to what that would look like.

## 8. Conclusion

This survey set out to answer three things. Where AI is actually being used in wholesale brokerage across the client lifecycle and settlement operations, what risks that use brings with it, and whether the regulatory frameworks that apply today are fit for purpose in a wholesale context.

### 8.1 What the survey found

On adoption, the pattern across both domains is the same. In the client lifecycle, document processing and screening are mature and widely deployed, identity verification and risk classification are moving that way with open governance questions, and the steps at either end of the flow are barely AI-driven at all. In settlement operations the tools are newer but the shape is identical: NLP extraction is useful for structured asset classes and weaker on bespoke documents, reconciliation is a thin AI layer on decades-old rule engines, and margin and liquidity forecasting are genuine but modest improvements. AI in wholesale brokerage today is a stack of assistive tools sitting on top of existing processes, not a replacement for them. The move to T+1 is likely to increase the pressure here, because the window for manual intervention gets shorter.

On risks, four came out of the survey.

- Vendor concentration leaves most firms three layers deep in dependencies they cannot easily swap out.
- Hallucination puts a ceiling on automation, because the cost of a wrong answer scales with trade size and trade sizes here are large.
- Opacity and accountability gaps sit awkwardly against regimes that expect a named human owner for material decisions.
- Distance between vendor marketing and actual capability means firms redesign processes around gains that do not arrive.

### 8.2 Are the frameworks fit for purpose?

On the evidence in this survey, no, not for wholesale specifically. The problem is not that MiFID II, ESMA's guidance, the AFM and DNB's work, or the AI Act are absent. Each applies to a wholesale broker in some way. The problem is that each was calibrated around harm to individuals, and the harms this survey found are not of that shape. MiFID II's treatment of algorithms was written for deterministic systems and does not map onto models that learn. The AI Act's high-risk trigger follows consumer-facing use, so scoring a person's creditworthiness is caught while a model informing settlement on a very large

institutional trade is not obviously caught at all. DORA is the closest fit for the vendor concentration problem, but it addresses operational resilience rather than opacity and accountability.

So a wholesale broker deploying AI today is not breaking rules. It is operating in a space where the rules that would govern what it is actually doing have not been written. That is a different situation from being under-regulated in the ordinary sense, and I think it is the more interesting finding of this survey. The gap is not a matter of enforcement. It is a matter of fit.

### **8.3 What firms and regulators should do**

For firms, four things follow from this survey.

- Record what the model proposed and what the human decided. Then investigate where the two diverged.
- Have clean data before buying or making an AI model. Most of the times clean data and maintained rules perform better on reconciliation than an AI model.
- Vendor concentration is a real risk and threat it that way.
- Extend risk governance for AI models to the operational and compliance tools. That is where deployments is currently being outrun by governance

For regulators, the main change is to stop using consumer harm as the trigger for scrutiny. In wholesale the relevant question is not whether a natural person is exposed but whether a model informs decisions where an error is large, hard to reverse, or propagates to other firms. Beyond the trigger, three things would help: clarity on what documentation and testing is expected for a model a firm bought rather than built, clarity on where accountability sits when a vendor model contributes to a decision, and attention to concentration among AI vendors specifically, which DORA opens the door to but does not fully address. The AFM is well placed for this work. It has been more active than most European supervisors on algorithm supervision, and its position that sectoral rules and the AI Act sit on top of each other is the right starting point. What is missing is a reading of that stack from a wholesale perspective rather than a retail one.

### **8.4 Limitations**

Several things limit what can be concluded here. The survey relies on published sources, vendor material, and my own exposure to wholesale operations. It does not include systematic interviews or a survey of firms, so claims about what is deployed in practice are weaker than claims about what is available, and vendor material is promotional even

when read critically. Where the peer-reviewed work on wholesale settlement AI is thin, I have leaned on adjacent domains and argued that the techniques transfer, which is a reasonable argument but an argument, not evidence. Finally, the regulatory picture is moving. The AI Act is being phased in, the EU's move to T+1 is set for 2027, and some specific claims in chapter 5 will date quickly, even if the underlying mismatch between retail-calibrated rules and wholesale risk is likely to outlast them.

## **8.5 Closing**

The thing I would want a reader to take away is that the interesting story in wholesale brokerage is not that AI is arriving. It is that AI is arriving into a part of the financial system that nobody has written rules for and that almost nobody has studied directly. The tools themselves are less impressive than the marketing suggests and more useful than the sceptics allow. What is genuinely missing is the evidence to tell those two apart, and a framework that would require someone to produce it.

## Reference

- [1] F. Simonn, “Past, present, and future research trajectories on retail investor behaviour: A composite bibliometric analysis and literature review,” *International Journal of Financial Studies*, 2025.
- [2] P. S. Nair et al., “Determinants of mobile apps adoption by retail investors for online trading in emerging financial markets,” *Benchmarking: An International Journal*, 2022.
- [3] M. Milazzo et al., “Heterogeneity of household stock portfolios in a national market,” *Physica A: Statistical Mechanics and its Applications*, 2025.
- [4] Z. Wu et al., “Diversification or distortion? the role of ETFs in retail investor portfolios and performance,” *Journal of Financial Stability*, 2026.
- [5] European Parliament and Council of the European Union, *Directive 2014/65/eu of 15 may 2014 on markets in financial instruments and amending directive 2002/92/ec and directive 2011/61/eu (recast)*, Official Journal of the European Union, L 173, 12.6.2014, pp. 349–496, MiFID II, 2014. [Online]. Available: <https://eur-lex.europa.eu/eli/dir/2014/65/oj/eng>.
- [6] V. Colaert, “Welke bescherming voor welke belegger? cliëntenclassificatie pre en post MiFID,” 2007.
- [7] M. Kruithof et al., “A differentiated approach to client protection: The example of MiFID,” in *The Failure of Contracting: Financial Crisis, Financial Services and General European Contract Law*, 2010.
- [8] T. Loonen et al., “The effectiveness of MiFID provisions for professional clients: A critical review,” *Journal of Financial Regulation and Compliance*, 2019.
- [9] Committee on Payment and Settlement Systems, “Delivery versus payment in securities settlement systems,” Bank for International Settlements, Basle, Tech. Rep., Sep. 1992. [Online]. Available: <https://www.bis.org/cpmi/publ/d06.htm>.
- [10] D. M. Weiss, *After the Trade Is Made: Processing Securities Transactions*, 3rd ed. New York: Portfolio (Penguin Group), 2006, ISBN: 9781591841272.
- [11] D. Duffie and H. Zhu, “Does a central clearing counterparty reduce counterparty risk?” *The Review of Asset Pricing Studies*, vol. 1, no. 1, pp. 74–95, 2011. DOI: [10.1093/rapstu/rar001](https://doi.org/10.1093/rapstu/rar001).
- [12] D. Loader, *Clearing, Settlement and Custody*, 3rd ed. Butterworth-Heinemann (Elsevier), 2019, ISBN: 9780128186909.
- [13] J. Devriese and J. Mitchell, “Liquidity risk in securities settlement,” *Journal of Banking & Finance*, vol. 30, no. 6, pp. 1807–1834, 2006.
- [14] P. Gurrola Perez et al., “Securities settlement fails network and buy-in strategies,” 2019.

- [15] Euronext, *T+1 programme*, Euronext, 2025. [Online]. Available: <https://www.euronext.com/en/regulation/t1-programme>.
- [16] European Securities and Markets Authority, “Final report: The shortening of the settlement cycle in the european union,” ESMA, Tech. Rep., Nov. 2024, Recommends 11 October 2027 as the transition date to T+1. [Online]. Available: <https://www.esma.europa.eu/press-news/esma-news/esma-proposes-move-t1-october-2027>.
- [17] B. Meenaghan, “CSDR settlement discipline costs and how to reduce or avoid them,” *Journal of Securities Operations & Custody*, 2020.
- [18] M. Fleming et al., “When the back office moved to the front burner: Settlement fails in the treasury market after 9/11,” *Economic and Policy Review*, 2004.
- [19] J. Benito, “Navigating settlement efficiency in a world of CSDR and T+1,” *Journal of Securities Operations & Custody*, 2025.
- [20] H.-J. Von Haenisch et al., “Know your customer: Unravelling the challenges,” *Journal of Payments Strategy & Systems*, 2024.
- [21] P. Gilmour, “Lifting the veil on beneficial ownership,” *Journal of Money Laundering Control*, 2020.
- [22] T. M. Matras, “Harmonisation of the beneficial ownership registers in European Union,” *Journal of Money Laundering Control*, 2025.
- [23] S. Jafarnejad et al., “A risk-based AML framework: Finding associates through ultimate beneficial owners,” in *2024 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics (CIFEr)*, IEEE, 2024.
- [24] T. Aamer et al., “Improving digital onboarding processes for financial services: A multivocal literature review,” *Baltic Journal of Modern Computing*, 2023.
- [25] X. Ling et al., “Intelligent document processing based on RPA and machine learning,” in *2020 Chinese Automation Congress (CAC)*, IEEE, 2020.
- [26] M. Pandey et al., “AI-based integrated approach for the development of intelligent document management system (IDMS),” *Procedia Computer Science*, 2023.
- [27] M. M. Y. G. Somda et al., “Implementation of robotic process automation to decrease the time required for KYC onboarding process,” in *2023 6th International Conference on Artificial Intelligence and Big Data (ICAIBD)*, IEEE, 2023.
- [28] M. Abdul-Al et al., “The evolution of biometric authentication: A deep dive into multi-modal facial recognition: A review case study,” *IEEE Access*, 2024.
- [29] P. Ruiu et al., “Enhancing eID card mobile-based authentication through 3D facial reconstruction,” *Journal of Information Security and Applications*, 2023.
- [30] Z. Chen, L. D. Van Khoa, E. N. Teoh, A. Nazir, E. K. Karuppiah, and K. S. Lam, “Machine learning techniques for anti-money laundering (AML) solutions in suspicious transaction detection: A review,” *Knowledge and Information Systems*, vol. 57, no. 2, pp. 245–285, 2018. DOI: [10.1007/s10115-017-1144-z](https://doi.org/10.1007/s10115-017-1144-z).

- [31] M. Alkhalili et al., “Investigation of applying machine learning for watch-list filtering in anti-money laundering,” *IEEE Access*, 2021.
- [32] R. Khandpur et al., *Adverse media mining for KYC and ESG compliance*, arXiv preprint, PREPRINT, not peer reviewed, 2021.
- [33] R. Jensen et al., “Fighting money laundering with statistics and machine learning,” *IEEE Access*, 2022.
- [34] E. J. Reite et al., “Improving client risk classification with machine learning to increase anti-money laundering detection efficiency,” *Journal of Money Laundering Control*, 2024.
- [35] P. T. Mazumder, “Explainable machine learning pipelines for customer risk scoring in anti-money laundering: A management and governance perspective,” *Journal of Data Analysis and Critical Management*, 2025.
- [36] D. V. Kute et al., “Deep learning and explainable artificial intelligence techniques applied for detecting money laundering: A critical review,” *IEEE Access*, 2021.
- [37] P. Chernakov et al., *An agentic LLM framework for adverse media screening in AML compliance*, arXiv preprint, PREPRINT, not peer reviewed, 2025.
- [38] V. Thokal et al., “Review and analysis of the literature: Artificial intelligence-based digital transformation of automated customer onboarding,” in *Proceedings of the International Conference on Recent Advances in Science and Engineering Technology (RAiSE-2023)*, 2024.
- [39] G. Schaecken Willemaers, “Client protection on european financial markets — from inform your client to know your product and beyond: An assessment of the PRIIPs regulation, MiFID II/MiFIR and IMD 2,” *Revue Trimestrielle de Droit Financier*, 2014, Autumn 2014 issue. Also available at SSRN 2494842. [Online]. Available: [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2494842](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2494842).
- [40] B. Oztas et al., “Transaction monitoring in anti-money laundering: A qualitative analysis and points of view from industry,” *Future Generation Computer Systems*, 2024.
- [41] P. Mazzarisi et al., “A machine learning approach to support decision in insider trading detection,” *EPJ Data Science*, 2022.
- [42] P. Priyadarshi et al., “A comprehensive review on insider trading detection using artificial intelligence,” *Journal of Computational Social Science*, 2024.
- [43] R. Tse et al., “Machine learning-driven credit risk: A systemic review,” *Neural Computing and Applications*, 2022.
- [44] A. N. Bakry et al., “Automatic suppression of false positive alerts in anti-money laundering systems using machine learning,” *The Journal of Supercomputing*, 2023.
- [45] S. Tariq et al., “Alert fatigue in security operations centres: Research challenges and opportunities,” *ACM Computing Surveys*, 2025.

- [46] E. Halford et al., “Developing a scoring model for managing money laundering transactions using machine learning,” *Journal of Money Laundering Control*, 2025.
- [47] A. Naser Eddin et al., *Anti-money laundering alert optimization using machine learning with graphs*, arXiv preprint, PREPRINT, not peer reviewed, 2021.
- [48] D. Labanca et al., “Amaretto: An active learning framework for money laundering detection,” *IEEE Access*, 2022.
- [49] F. Johannessen et al., *Finding money launderers using heterogeneous graph neural networks*, arXiv preprint, PREPRINT, not peer reviewed, 2023.
- [50] A. Alonso Robisco et al., “Measuring the model risk-adjusted performance of machine learning algorithms in credit default prediction,” *Financial Innovation*, 2022.
- [51] F. Jalalvand et al., “Alert prioritisation in security operations centres: A systematic survey on criteria and methods,” *ACM Computing Surveys*, 2024.
- [52] P. Lewis et al., “Retrieval-augmented generation for knowledge-intensive NLP tasks,” in *Advances in Neural Information Processing Systems*, vol. 33, Curran Associates, Inc., 2020, pp. 9459–9474.
- [53] I. Andrade et al., “Increasing customer service efficiency through artificial intelligence chatbot,” *Revista de Gestão*, 2022.
- [54] S. Hwang et al., “Toward a chatbot for financial sustainability,” *Sustainability*, 2021.
- [55] R. Eustaquio-Jiménez et al., “Chatbots for customer service in financial entities: A comprehensive systematic review,” *Journal of Infrastructure, Policy and Development*, 2024.
- [56] S. Farquhar, J. Kossen, L. Kuhn, and Y. Gal, “Detecting hallucinations in large language models using semantic entropy,” *Nature*, vol. 630, no. 8017, pp. 625–630, 2024. DOI: [10.1038/s41586-024-07421-0](https://doi.org/10.1038/s41586-024-07421-0).
- [57] L. Ryan et al., “Trading operations: Intelligent automation and the T+1 mandate,” *Journal of Securities Operations & Custody*, 2023.
- [58] D. Petiteville, “Lessons from T+1 settlement: Risk mitigation and future implications,” *Journal of Securities Operations & Custody*, 2025.
- [59] A. Gandy et al., “Industry preparedness for accelerated settlement,” 2023.
- [60] R. Walley, “Accelerating the settlement cycle: An opportunity to modernise across the industry,” *Journal of Securities Operations & Custody*, 2023.
- [61] A. Mohammed et al., “Leveraging natural language processing for trade exception classification and resolution in capital markets: A comprehensive study,” *Journal of Cognitive Computing and Cybernetic Innovations*, 2025.
- [62] S. Graham et al., “Natural language processing for legal document review: Categorising deontic modalities in contracts,” *Artificial Intelligence and Law*, 2023.

- [63] H. Zhou et al., “Intelligent detection on construction project contract missing clauses based on deep learning and NLP,” *Engineering, Construction and Architectural Management*, 2023.
- [64] M. Schreyer, T. Sattarov, D. Borth, A. Dengel, and B. Reimer, *Detection of anomalies in large scale accounting data using deep autoencoder networks*, arXiv preprint arXiv:1709.05254, PREPRINT, not peer reviewed, 2017. [Online]. Available: <https://arxiv.org/abs/1709.05254>.
- [65] K. Janamolla et al., “Integrating blockchain and AI for efficient trade exception handling: A case study in cross-border settlements,” *Journal of Cognitive Computing and Cybernetic Innovations*, 2025.
- [66] V. Singh, B. Choubey, and S. Sauer, “Liquidity forecasting at corporate and subsidiary levels using machine learning,” *Intelligent Systems in Accounting, Finance and Management*, vol. 31, no. 3, e1565, 2024. DOI: [10.1002/isaf.1565](https://doi.org/10.1002/isaf.1565).
- [67] V. Liermann et al., “Intraday liquidity: Forecast using pattern recognition,” in *The Impact of Digital Transformation and FinTech on the Finance Professional*, Palgrave Macmillan, 2019.
- [68] I. Aldasoro et al., “AI agents for cash management in payment systems,” 2025.
- [69] X. Ma, S. Spinner, A. Venditti, Z. Li, and S. Tang, *Initial margin simulation with deep learning*, SSRN working paper 3357626, WORKING PAPER, not peer reviewed, Mar. 2019. [Online]. Available: [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3357626](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3357626).
- [70] J. Pérez Villarino and Á. Leitaó, *On deep learning for computing the dynamic initial margin and margin value adjustment*, arXiv preprint arXiv:2407.16435, 2024. [Online]. Available: <https://arxiv.org/abs/2407.16435>.
- [71] A. Ashta and H. Herrmann, “Artificial intelligence and fintech: An overview of opportunities and risks for banking, investments, and microfinance,” *Strategic Change*, 2021.
- [72] M. Pierides et al., “Key themes of resiliency, outsourcing and third-party risk management regimes,” *Journal of Securities Operations & Custody*, 2025.
- [73] D. Gozman et al., “Cloud sourcing and mitigating concentration risk in financial services,” 2020.
- [74] J. W. Gulliver et al., “Cloud adoption in the financial sector and concentration risk,” 2023.
- [75] European Parliament and Council of the European Union, *Regulation (eu) 2022/2554 of 14 december 2022 on digital operational resilience for the financial sector*, Official Journal of the European Union, L 333, 27.12.2022, pp. 1–79, DORA. Applies from 17 January 2025, 2022. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2022/2554/oj>.

- [76] L. Huang et al., “A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions,” *ACM Transactions on Information Systems*, 2023.
- [77] A. B. Das et al., “Hallucinations and key information extraction in medical texts: A comprehensive assessment of open-source large language models,” in *2025 IEEE EMBS International Conference on Biomedical and Health Informatics (BHI)*, IEEE, 2025.
- [78] A. Biswas and W. Talukdar, *Robustness of structured data extraction from in-plane rotated documents using multi-modal large language models*, arXiv preprint, PREPRINT, not peer reviewed, 2024.
- [79] S. M. T. I. Tonmoy et al., *A comprehensive survey of hallucination mitigation techniques in large language models*, arXiv preprint, PREPRINT, not peer reviewed, 2024.
- [80] W. J. Yeo, W. van der Heever, R. Mao, E. Cambria, R. Satapathy, and G. Mengaldo, “A comprehensive review on financial explainable AI,” *Artificial Intelligence Review*, 2025. DOI: [10.1007/s10462-024-11077-7](https://doi.org/10.1007/s10462-024-11077-7).
- [81] M. Busuioc, “Accountable artificial intelligence: Holding algorithms to account,” *Public Administration Review*, 2020.
- [82] D. Zetzsche, D. Arner, R. Buckley, and B. Tang, “Artificial intelligence in finance: Putting the human in the loop,” 2020.
- [83] N. Elsayed, *AI washing and the erosion of digital legitimacy: A socio-technical perspective on responsible artificial intelligence in business*, arXiv preprint, PREPRINT, not peer reviewed, 2026.
- [84] X. Song et al., “AI washing: Strategic disclosure and backlash,” *Finance Research Letters*, 2026.
- [85] S. Ozturkcan et al., “Responsible AI in marketing: AI booing and AI washing cycle of AI mistrust,” *International Journal of Market Research*, 2025.
- [86] OECD, “Artificial intelligence, machine learning and big data in finance: Opportunities, challenges and implications for policy makers,” OECD Publishing, Paris, Tech. Rep., 2021. DOI: [10.1787/98e761e7-en](https://doi.org/10.1787/98e761e7-en).
- [87] European Securities and Markets Authority, “Supervisory briefing on algorithmic trading in the EU,” European Securities and Markets Authority, Paris, Tech. Rep. ESMA74-1505669079-10311, Feb. 2026. Accessed: Aug. 11, 2026. [Online]. Available: [https://www.esma.europa.eu/sites/default/files/2026-02/ESMA74-1505669079-10311\\_Supervisory\\_Briefing\\_on\\_Algorithmic\\_Trading\\_in\\_the\\_EU.pdf](https://www.esma.europa.eu/sites/default/files/2026-02/ESMA74-1505669079-10311_Supervisory_Briefing_on_Algorithmic_Trading_in_the_EU.pdf).
- [88] T. Čuk and A. Van Waeyenberge, “European legal framework for algorithmic and high frequency trading (MiFID II and MAR),” *European Journal of Risk Regulation*, 2018.

- [89] P. Raschner, “Algorithms put to test: Control of algorithms in securities trading through mandatory market simulations?,” 2021.
- [90] A. Azzutti, “AI governance after MiFID II: Beyond (mere) technological neutrality?” *ERA Forum*, 2026.
- [91] European Securities and Markets Authority, “Artificial intelligence in EU securities markets,” ESMA, Paris, TRV Risk Analysis ESMA50-164-6247, Feb. 2023. DOI: 10.2856/851487. [Online]. Available: [https://www.esma.europa.eu/sites/default/files/library/ESMA50-164-6247-AI\\_in\\_securities\\_markets.pdf](https://www.esma.europa.eu/sites/default/files/library/ESMA50-164-6247-AI_in_securities_markets.pdf).
- [92] De Nederlandsche Bank, “General principles for the use of artificial intelligence in the financial sector,” De Nederlandsche Bank, Amsterdam, Discussion paper, Jul. 2019, The SAFEST principles.
- [93] Autoriteit Financiële Markten and De Nederlandsche Bank, “The impact of AI on the financial sector and supervision,” AFM and DNB, Amsterdam, Tech. Rep., 2024. [Online]. Available: <https://www.afm.nl/~profmedia/files/rapporten/2024/ai-rapport-en.pdf>.
- [94] J. Wittlin et al., “Artificial intelligence act compliance assurance challenges for european banks,” *Information & Communications Technology Law*, 2026.
- [95] Autoriteit Financiële Markten and De Nederlandsche Bank, “The impact of AI on the financial sector and supervision,” AFM and DNB, Amsterdam, Tech. Rep., Jun. 2024. [Online]. Available: <https://www.afm.nl/~profmedia/files/rapporten/2024/ai-rapport-en.pdf>.
- [96] European Parliament and Council of the European Union, *Regulation (eu) 2024/1689 of 13 june 2024 laying down harmonised rules on artificial intelligence (artificial intelligence act)*, Official Journal of the European Union, L, 2024/1689, 12.7.2024, 2024. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>.
- [97] G. Mazzini and F. Bagni, “Considerations on the regulation of AI systems in the financial sector by the AI Act,” *Frontiers in Artificial Intelligence*, 2023.
- [98] C. Martins Pereira et al., “The european union’s approach to artificial intelligence and the challenge of financial systemic risk,” 2025.
- [99] M. Ebers, “Truly risk-based regulation of artificial intelligence: How to implement the EU’s AI Act,” *European Journal of Risk Regulation*, 2024.
- [100] Anonymous, “The european AI Act’s impact on financial markets: From governance to co-regulation,” 2023.
- [101] J. Lee and F. L’heureux, “Regulation of algorithmic trading: Frameworks for human supervision and direct market interventions,” *European Business Law Review*, 2022.

- [102] C. García-Llorente et al., “Algorithmic governance in banking: A comparative analysis of risk-based and accountability-oriented oversight,” *Journal of Banking Regulation*, 2026.