



EUROPEAN CENTRAL BANK
EUROSYSTEM

Discussion Paper Series

Álvaro Cartea, Jean-Edouard Colliard,
Thierry Foucault, Peter Hoffmann,
Rob Graumans, Jean-David Sigaux

Artificial intelligence and financial markets

No 32 / September 2026



Disclaimer: This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Discussion papers

Discussion papers are research-based papers on policy relevant topics, offering a broader and more balanced perspective. While being partly based on original research, they place the analysis in the wider context of the literature on the topic. They also consider explicitly the policy perspective, with a view to develop a number of key policy messages. Their format offers the advantage that alternative analyses and perspectives can be combined, including theoretical and empirical work. The selection and distribution of discussion papers are subject to the approval of the Director General of the Directorate General Research.

Abstract

Artificial intelligence is transforming financial markets by enabling investors and intermediaries to extract more information from increasingly abundant data through advances in algorithms and computing power. We discuss how AI changes information production and decision-making, labor demand, and financial intermediation, with a particular emphasis on securities markets. We then survey the implications of this transformation for market efficiency, competition, informational frictions, and financial stability.

Keywords: artificial intelligence, market efficiency, agency conflicts, competition, financial stability.

JEL Codes: G1, G14, G23, G28, E6.

Non-technical summary

Artificial intelligence (AI) is rapidly changing the finance industry. Financial intermediaries are investing heavily in data, computing infrastructure, algorithms, and specialized staff. Survey evidence shows that AI adoption is widespread and accelerating, especially among large firms and those active in securities markets. Moreover, finance jobs are among those most exposed to AI-driven automation. This is not surprising because information is at the heart of almost every financial activity. Financial intermediaries collect and process information to assess investments, manage risk, and allocate capital. By combining large and diverse datasets with increasingly powerful algorithms and computing capabilities, AI enables them to make more accurate predictions and automate decision-making.

This development raises important questions for economists and policymakers. Will AI make markets more efficient and help direct capital to its most productive uses? Or will it concentrate power in the hands of a small circle of tech-savvy firms? Will it make financial information more accessible to everyone, or instead increase the advantage of sophisticated institutions that can afford proprietary data, advanced algorithms, and large computing resources? And how might these changes affect financial stability?

The paper is centered around the idea that financial intermediaries use data and human expertise to produce information and make decisions, such as evaluating investments and constructing portfolios. AI reshapes this process by increasing the productivity of data: better algorithms and greater computing power enable firms to extract more information from larger datasets. We discuss how this technological shift affects finance jobs, and then consider the wider implications for market efficiency, competition, conflicts of interest, and financial stability. Part of our analysis focuses on firms that can be considered “early

adopters” of AI-powered trading such as proprietary trading firms and quantitative asset managers.

AI is a potential source for efficiency gains as more accurate forecasts can bring prices closer to their “fundamental value” and thereby improve valuations, the screening of investment projects, and ultimately the allocation of capital in the economy. However, it is not guaranteed that these benefits will materialize. Instead, investors may use AI to generate information that is privately profitable, but socially unproductive—for example, fast trading on public information releases or the identification of others’ trading patterns in order to anticipate their behaviour. In the extreme, such strategies can even weaken the incentives to produce socially useful information, and thus undermine market efficiency.

AI is also likely to reshape competition in financial markets. The large investments required in data, computing, and model development create economies of scale, with the potential to increase market concentration and limit competition. However, AI could also lower barriers to entry and spur the entry of new firms. Consistent with these opposing economic forces, a descriptive analysis of Dutch supervisory paints an overall mixed picture for the case of proprietary trading firms. Finally, we discuss how autonomous pricing algorithms can help sustain noncompetitive outcomes, and argue that their unpredictable behaviour complicates policy intervention.

The paper also discusses how AI may shape key financial market frictions such as adverse selection and moral hazard. We highlight the channels through which AI may either amplify existing information asymmetries or level the playing field. Further, we discuss how AI can reshape existing agency conflicts such as those between employers and employees, and give rise to new ones that arise from model complexity. We also stress the underlying regulatory challenges.

The paper closes by discussing how a widespread adoption of AI models may affect financial stability. While concentration in the market for models and data could give rise to correlated behaviour of AI-powered market participants, these concerns may be ameliorated by market participants' incentives to use models or data whose output is not yet reflected in market prices (and thus represents profit opportunities). We also discuss how an increase in market concentration can be a double-edged sword from a financial stability perspective, and highlight challenges related to model risk, rogue algorithms, and changes to the business model of key financial intermediaries.

1 Introduction

Artificial intelligence (AI) is rapidly transforming the financial industry.¹ Financial intermediaries (e.g., asset managers and banks) are investing heavily in AI capabilities, including large-scale data infrastructures, machine learning models, computing resources, and human capital. Recent surveys (e.g., [ESMA, 2026](#)) indicate that AI adoption is widespread and accelerating across financial intermediaries, with particularly rapid diffusion among large firms and those operating in securities markets. Moreover, according to several recent studies, occupations in the financial industry are among those most exposed to AI-driven automation.²

The importance of AI for finance is not surprising given the central role of information in virtually every financial activity. Financial intermediaries exist largely to collect, process, and transmit information in order to allocate capital, manage risks, and mitigate informational frictions between market participants ([Levine, 1997](#)). By combining vast amounts of data with increasingly powerful algorithms and computing resources, AI potentially enables financial intermediaries to generate more accurate predictions, make better decisions, and ultimately allocate capital more efficiently at a lower cost.

These developments raise fundamental questions for economists and policymakers. Does AI improve the efficiency of financial markets by making prices more informative about fundamentals and improving capital allocation? Or does it instead strengthen the market power of a few technologically advanced intermediaries? Will AI reduce informational frictions by making financial information more accessible, or widen the infor-

¹For this paper, we adopt the [OECD \(2024\)](#) definition, which describes AI as “a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.”

²See [Acemoglu et al. \(2022\)](#) and [Eisfeldt et al. \(2023\)](#).

mational advantage of sophisticated institutions capable of investing in proprietary data, computing infrastructure, and advanced algorithms? More broadly, how will AI reshape financial intermediation, market liquidity, and financial stability?

The goal of this paper is to discuss these questions in light of recent academic research. We organize our discussion around the idea that financial intermediaries combine data and labor to produce information and make decisions, such as evaluating investment opportunities or selecting portfolios. AI transforms this production technology by increasing the productivity of data in information production through advances in algorithms and computing power. As a result, financial intermediaries can extract more information from increasingly diverse datasets while relying less on labor-intensive information processing.

This perspective provides a useful framework for understanding the economic effects of AI in finance. By changing the technology through which information is produced, AI affects not only the demand for labor, but also the efficiency of financial markets, the competitive structure of the financial industry, informational frictions among market participants, and ultimately financial stability.

We first examine the implications of AI for the production of information and the resulting implications for the role of labor in the financial industry (Section 2). We then discuss how AI technologies are being deployed across different types of financial intermediaries, with particular attention to proprietary trading firms and quantitative asset managers, which have emerged as “early adopters” (Section 3).

We then turn to the broader economic consequences of AI for financial markets (Sections 4 and 5). We argue that AI has the potential to generate substantial efficiency gains, but at the same time may also amplify existing and create new distortions. In the following, we highlight the key channels.

Efficiency and capital allocation. AI enables investors to form more accurate forecasts of cash flows and asset returns at lower cost. This can lead to improved asset valuations, better screening and monitoring of investment projects, and ultimately more efficient capital allocation. However, these gains are not guaranteed. Financial intermediaries can instead use AI to produce information that is privately valuable but socially wasteful. For example, AI can be used to trade rapidly on newly released information at the expense of accuracy or to extract signals from the trading footprints of others rather than from information about fundamentals. Such activities generate little social value and may even reduce incentives for producing fundamental information. As a result, AI could ultimately render asset prices less informative.

The nature of the underlying data also matters. Because AI relies heavily on historical data, it may perform poorly in novel environments or when evaluating projects whose payoffs materialize only over long horizons. As a result, AI may reject genuinely innovative projects or distort long-term investment decisions. In either case, this could lead to missed investment opportunities for society.

Market power and industry structure. AI has the potential to alter the competitive landscape in financial markets. On the one hand, it may strengthen economies of scale in data, computing, and model development and thereby increase concentration and market power. On the other hand, it could also facilitate the creation of startups and thereby increase competition, making this ultimately an empirical question. A brief analysis for Dutch supervisory data for proprietary trading firms (PTFs) paints a mixed picture, consistent with these opposing channels. While PTFs have gained market share on aggregate, within-sector concentration has remained rather flat, or even decreased. We also discuss how AI enables autonomous pricing algorithms that can sustain non-

competitive outcomes, intentionally or inadvertently, and thereby weakens competition in financial markets. Moreover, as predicting the behavior of autonomous algorithms is difficult, the adoption of such algorithms for securities trading can complicate policy interventions.

Informational frictions. AI can either reduce or amplify informational frictions. While it can make information more widely available, it can also increase the informational advantage of sophisticated intermediaries investing in proprietary data, algorithms, and computing infrastructure. As a result, AI could amplify informational asymmetries between sophisticated and unsophisticated investors, increasing adverse selection costs and trading frictions.

Financial stability. Finally, we discuss how these developments may affect financial stability. The widespread use of similar datasets and data-processing algorithms could increase the correlation in investors' signals, give rise to “crowded trades”, and amplify fire sales in times of distress. At the same time, incentives for acquiring signals not yet reflected in prices can act as a countervailing force. While greater concentration among financial intermediaries, resulting from the economies of scale discussed above, may further amplify these risks, it can also give rise to more prudent behavior due to “skin in the game”. We also discuss other financial stability aspects such as malicious behavior, model risk, and changes to the nature of financial intermediation.

Other recent papers also review different aspects of AI in economics and finance.³ Our survey differs from these contributions in two respects. First, rather than organizing the

³[Foucault *et al.* \(2025\)](#) provide a broad overview of the effects of AI across banking, insurance, corporate finance, and securities markets. [Babina \(2026\)](#) surveys firm-level evidence on AI adoption and its economic effects, with a particular focus on measuring firms' AI investments while [Eisfeldt and Schubert \(2025\)](#) focus on the implications of generative AI for finance. Policy reports by [International Monetary Fund \(2024\)](#) and [Aldasoro *et al.* \(2024\)](#) review the opportunities, risks, and policy challenges associated with AI adoption throughout the financial system.

discussion around a particular AI technology, financial sector, or market failure, we view AI as a technological innovation that transforms the production of financial information. This perspective provides a coherent framework for understanding how advances in data, algorithms, and computing power affect market efficiency, competition, informational frictions, and financial stability. Second, we place particular emphasis on securities markets, where AI adoption has been especially rapid and where many of the economic mechanisms generated by AI can already be observed.

2 New Inputs: Data, Computing, and Algorithms

2.1 Producing financial information

The financial industry produces and facilitates the trading of financial contracts to share risks, transfer resources, and provide incentives ([Philippon, 2015](#)). Information is a key input in this process. As emphasized by [Levine \(1997\)](#), financial markets and institutions arise to mitigate the frictions created by imperfect information. Information enables investors to assess the payoffs of financial contracts more accurately and make better investment decisions. It also reduces frictions arising from informational asymmetries, such as adverse selection and moral hazard.

Many finance jobs are therefore related to data collection and processing, enabling professionals to make predictions and act accordingly. For instance, securities analysts collect and process data to forecast firms' future cash flows, estimate their “fair” values, and make buy/sell recommendations. Similarly, active fund managers collect data about securities to forecast their returns and make buy/sell decisions based on these forecasts.

People in these occupations transform data into knowledge, i.e., information that can

be used for financial decisions. The more knowledge they accumulate, the better their decisions. For instance, with more data, an active asset manager can identify undervalued (overvalued) securities more accurately, overweight (underweight) them in her portfolio (the decision) and deliver better performance.⁴ Another example is that of a venture capitalist: more knowledge about the prospects of potential investment targets enables the venture capitalist to distinguish good from bad projects better and allocate capital accordingly.

Given the central role of information in finance, AI has profound implications for the financial industry. It is a technological innovation that enables financial intermediaries to transform data into valuable information more efficiently. As an illustration, let D_i be the amount of data and L_i be the amount of labor used by financial intermediary i to produce an amount of “knowledge” K_i according to the following Cobb-Douglas function for knowledge

$$K_i = A_i D_i^{\alpha(c,a)} L_i^{\beta(c,a)}, \quad (1)$$

similar to [Abis and Veldkamp \(2024\)](#).

The coefficient $\alpha(c, a)$ determines the elasticity of knowledge with respect to data, that is, the percentage increase in “knowledge” for a one-percent increase in data. In our discussion, the function α is increasing in computing power c and in a , which is the sophistication of the firm’s algorithms to process data and generate predictions. Similarly, the coefficient $\beta(c, a)$ represents the elasticity of knowledge with respect to labour, and A_i is a firm-specific productivity parameter. With better tools for data processing (more

⁴For example, [Dugast and Foucault \(2025\)](#) develop a model in which active asset managers search for high-precision signals via data mining (a systematic exploration of various datasets). In their framework, managers who obtain more precise signals make better portfolio allocation decisions and thereby achieve higher expected excess returns.

powerful algorithms) and more computing power to train and use these algorithms, a financial intermediary can extract more knowledge from a given stock of data.

In addition to improving predictions, AI can be used to make autonomous decisions. This form of AI can be thought of as a layer added to the production function of knowledge: instead of transforming inputs into predictions that are then used to take actions, it directly transforms inputs into concrete actions.⁵ For instance, asset managers may generate predictions and use AI to make buy and sell decisions autonomously based on these predictions.

As explained below, AI stems from advances in computing power and increasingly powerful algorithms for data processing and decision-making. As such, it represents a positive shock to α , increasing the share of data relative to labor in the production of knowledge and raising concerns about the future of jobs in information-intensive industries such as finance.⁶

In the rest of this section, we describe the forces shaping how AI affects the production of knowledge in finance in more detail: data (Section 2.2), algorithms (Section 2.3), and computing power (Section 2.4). We also discuss the potential implications of these forces for the future of jobs in the financial industry (Section 2.5). Throughout, we mostly focus on the “securities industry”—PTFs, asset managers, broker-dealers, and underwriters—because this segment of the financial industry has been at the forefront of AI adoption (as we explain below in Section 3).

⁵The distinction between prediction-generation and decision-making can be subtle. In some model architectures, such as Deep Q-learning, one can conceptually distinguish between a component that learns and stores information about the environment (the deep neural network approximating the Q-function) and a component that uses this information to select actions. In other models, such as basic Q-learning, this distinction is less obvious because the model has less room to learn rich representations of the environment. See Section 2.3 and Box 5 for further details.

⁶Holding the prices of data and labor constant, the data-to-labor ratio that maximizes knowledge production increases with α .

2.2 Data

Data have always played a central role in finance. Traditionally, finance professionals have relied on (i) firms' reports (e.g., regulatory filings, accounting data, etc.), (ii) macroeconomic data, (iii) market data (e.g., stock returns, order flows, quotes, etc.), and (iv) data obtained by intermediaries (e.g., banks and broker-dealers) about their clients as sources of information.

What is novel is the sharp increase in the volume and diversity of data available for forecasting future cash flows or returns. These data are often referred to as “alternative data” because they provide an alternative to the traditional sources mentioned above. Most alternative data can be grouped into one of three broad categories: (i) Data generated as a by-product of consumers' decisions (e.g., internet traffic on a firm's website, point-of-sale data, credit card data); (ii) Data generated through communications on social media (e.g., messages on StockTwits or X); (iii) Sensor data (e.g., satellite imagery or geolocation data).

Alternative data are often unstructured (e.g., text, images, or audio), and therefore often collected, cleaned and processed by data vendors before reselling them in a format suitable for analysis. The number of such providers has grown rapidly alongside the digitalization of the economy. For example, [Dessaint *et al.* \(2024\)](#) document an increase in the number of providers specializing in “finance” datasets from about 100 in 2000 to more than 500 in 2017 (see [Figure 1](#)). In addition, new data intermediaries such as Eagle Alpha, Databricks, and Datarade serve as marketplaces connecting buyers and sellers of alternative data. For instance, Datarade was listing over 3,500 alternative datasets as of 2025.




Figure 1: Estimated number of alternative data vendors by year based on the J.P. Morgan *2019 Alternative Data Handbook*. Reprinted from [Dessaint et al. \(2024\)](#).

Several of these intermediaries specifically cater to the finance industry. For instance, on its webpage, YipitData advertises itself as helping hedge funds, investment managers, mutual funds, private equity funds, and venture capitalists, while Eagle Alpha organizes webinars targeting fund managers. This is confirmed by many academic studies, which identify footprints left by alternative data users and show that alternative data can be used to produce information about future asset cash flows or returns for a wide range of datasets.⁷ Box 1 dives into more detail on research related to the use of satellite imagery.

⁷[Dessaint et al. \(2024\)](#) list 26 academic studies showing that various alternative datasets can be used to forecast stock returns and cash flows.

Box 1: Eye in the sky

Satellite imagery constitutes an important source of information about economic activity. It can be used to obtain data on traffic at shopping centers (proxied by parking-lot occupancy rates), industrial production (proxied by smoke emissions or nighttime light intensity), energy consumption (proxied by oil storage tank shadow analysis or coal stockpile size), and shipping and logistics activity. Many firms (e.g., Kayrros, RS Metrics, UP42) specialize in selling data derived from satellite imagery after processing and transforming raw observations into usable indicators.

Such data can be used by asset managers to extract signals about future asset returns and corporate cash flows. For instance, [Katona *et al.* \(2025\)](#) use satellite imagery documenting daily store-level parking-lot utilization rates for major U.S. retailers. They show that these data can be used to forecast retailers' quarterly sales and to build a portfolio strategy that exploits these signals by buying retailers' stocks with high traffic growth and selling (shorting) stocks with low traffic growth. Similarly, [Cao *et al.* \(2024\)](#) find that equity analysts' forecasts become more accurate when the firms they follow begin to be covered by alternative data obtained from satellite imagery. Finally, [Mukherjee *et al.* \(2021\)](#) show that weekly announcements of changes in U.S. oil inventories have a smaller impact on oil futures prices in clear-weather weeks than in cloudy weeks. This is consistent with market participants relying on satellite images of storage sites to forecast changes in oil inventories before these are announced.

2.3 Algorithms

Algorithms play a dual role: they improve the generation of knowledge by raising the productivity of data relative to labor, and they facilitate the use of this knowledge as input in financial decisions. In the first role, algorithms increase the quality and quantity of knowledge that can be produced from a given combination of data and labor; this is captured by a in the production function (1), where the elasticity $\alpha(c, a)$ of knowledge with respect to data also depends on computing power c . In the second role, firms develop and train an additional layer of algorithms that maps knowledge K_i , together with other inputs, into decisions such as the implementation of trading strategies or portfolio selection. In both roles, machine learning (ML) provides the methods used to develop and improve these algorithms. A single AI system may perform both roles. For instance, an agent based on a large language model (LLM) can extract information from filings and then act on it directly

For decades, algorithms have been central to the implementation of quantitative methods for forecasting and decision-making in financial markets. The large-scale deployment of data-driven models dates back to at least the 1980s, which coincides with the founding of firms such as Renaissance Technologies and D.E. Shaw (see Section 3.2 for details). Similarly, algorithmic trading has become prevalent over the past two decades and accounts for the majority of trading volume across many asset classes (equities, futures, options, and FX). For example, for Dutch blue-chip stocks, Figure 2 shows that algorithmic trading accounts for approximately 70% of total trading volume by value and an even larger share by number of trades. Algorithmic trading provides a natural technological infrastructure for deploying AI applications, but an order classified as algorithmic does not necessarily involve ML.




Figure 2: Share of Algorithmic Trading. Percentage of transactions executed via an algorithm over time. Data include all large-cap (AEX) stocks traded on Euronext Amsterdam during 2020–2026. Source: MiFID II data.

Arguably, the modern era of AI began with [Krizhevsky *et al.* \(2012\)](#), who show that deep neural networks can dramatically outperform previous models on the ImageNet visual recognition task, demonstrating that deep learning is practical at scale. The ML algorithms that have since been adopted in finance are novel in both roles. In the first, they can handle a large number of predictors relative to the number of observations (Section [2.3.2](#)). In the second, they can learn how to make autonomous decisions without relying solely on pre-set rules (Section [2.3.3](#)).

2.3.1 Artificial intelligence and machine learning algorithms

While AI refers an entire machine-based system (recall the definition in footnote 1), ML encompasses methods that use data or interactive feedback to improve predictions and de-

cisions.⁸ For predictive applications, an algorithm learns statistical relationships that map inputs into forecasts, classifications, or representations, typically in high-dimensional settings where regularization methods are used to improve out-of-sample performance. For decision-making, an algorithm learns a policy that maps information about the current state into an action, taking into account the rewards generated by that action. This distinction matters because predictive learning generally treats the surrounding environment as given, whereas decision-making frequently occurs in an environment that is at least in part shaped by the algorithm’s own actions. In financial markets, individual actions affect the environment and the behavior of other participants, so the algorithm must continue to learn and update in tandem with changes in the market environment. For example, state-of-the-art algorithms for executing large orders in electronic markets are designed to learn how the market reacts to the execution algorithm’s own actions.

In discussing ML algorithms, it is useful to distinguish a model’s *learning paradigm* from its *architecture*. The learning paradigm describes what the model learns from and what it is trained to optimize; Box 2 presents the three main categories, namely supervised learning, unsupervised learning, and reinforcement learning (RL). On the other hand, the architecture describes the functional structure through which inputs are represented and processed. The same architecture can be trained under different learning paradigms, and the same learning paradigm can be implemented using different architectures.

⁸There is no accepted definition of the term machine learning (Gu *et al.*, 2020); the term was first used in Samuel (1959). A commonly used definition is that “a computer program is said to learn from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T, as measured by P, improves with experience E”, which dates back to Mitchell (1997).

Box 2: Learning paradigms: How a model learns

1. *Supervised learning* uses data that include pairs of inputs and corresponding labels, and the model learns direct mappings from the inputs to labels by minimizing some error function. This paradigm is suited to classification and regression tasks with clearly defined labels. A financial application is forecasting future corporate earnings, cash flows, or asset returns.
2. *Unsupervised learning* relies on data without labeled outputs. Rather than predicting predefined targets, the model identifies hidden structures or statistical regularities in the data. Typical tasks include clustering, which groups observations with similar characteristics, and dimensionality reduction, which summarizes high-dimensional data while retaining key information. In finance, unsupervised learning can be used to extract latent factors from large datasets, cluster firms or investors into groups, or reduce a large collection of predictors to a smaller set of informative representations, which is a clear example of how raw data is used to produce useful knowledge; see, e.g., [Cartea et al. \(2025b\)](#).
3. *Reinforcement learning* trains models through interaction with an environment, using feedback in the form of rewards or penalties associated with model actions. The objective is to maximize cumulative rewards over time, which makes this paradigm particularly well-suited to sequential and dynamic decision-making problems ([Sutton and Barto, 1998](#)). RL in finance can be used for optimal trade execution, where a large order is strategically split into smaller pieces and executed over time to minimize costs or market im-

pact (Cartea *et al.*, 2015).^a Because the actions selected by an RL agent can affect its future environment, RL is particularly consequential when several algorithms interact strategically.

These paradigms serve different purposes, and they typically map onto the two roles of algorithms. Supervised and unsupervised learning produce knowledge from data, for example by forecasting stock returns or extracting latent factors. RL selects actions and learns from the rewards they generate, for example when an agent learns how to execute a large order. However, the distinction is not absolute because modern AI systems are often built as pipelines that combine several paradigms at different stages of development or deployment, rather than relying on a single learning method throughout.

^aRL does not necessarily entail online learning. Models are often trained offline, for example in simulated environments, and deployed without further updating. The RL framework can also describe decision-making at a higher level of aggregation. A firm that selects which algorithm to deploy in each state of the market, and updates this choice based on realized immediate and future rewards, is effectively implementing an RL policy over algorithms.

Model architectures define the set of relationships that a model can represent and the degree of flexibility with which it can approximate patterns in the data. Several model architectures are commonly used in financial applications. Linear models are relatively simple and interpretable, and are popular in financial domains where speed and explainability are important. Tree-based models, including decision trees, random forests, and gradient boosting methods, can learn nonlinear relationships and retain a partially interpretable form, and are used in, for example, credit scoring (Dastile *et al.*, 2020). Neural network architectures are a flexible class of function approximators, but are generally considered less interpretable than linear models.

These distinctions matter from a policy perspective because they give rise to different regulatory considerations. Model architectures largely determine the range and complexity of patterns that can be extracted from data, the associated requirements for compute and specialized labor, and the difficulty of interpreting outputs (Dixon *et al.*, 2020). By contrast, learning paradigms determine whether AI is used primarily to generate predictions or to make decisions through feedback, and concerns about unintended or emergent outcomes are usually associated with the paradigm rather than the architecture.

2.3.2 Making predictions

One key strength of ML algorithms relative to more traditional predictive tools (e.g., OLS regressions) is their ability to handle a very large number of predictors relative to the number of observations, a property referred to as “complexity” by Kelly *et al.* (2024). This makes them particularly well-suited to harness the informational content of large and unstructured data, which often characterize the alternative datasets discussed in Section 2.2. However, the marginal value of data and of model complexity is not unbounded. In terms of (1), diminishing returns would require $\alpha(c, a)$ to be less than one. Moreover, greater algorithmic sophistication (higher a) does not always yield more knowledge. Cartea *et al.* (2025c) show that an increase in model complexity is not always optimal when model features are observed with noise. Specifically, they show that the marginal benefits of increasing model complexity will, at some point, start to diminish and become negative as the size of the model increases due to measurement noise in the learned parameters used by the complex model.⁹

Many academic studies show that ML algorithms improve forecasting power relative to

⁹See also Nagel (2025), Buncic (2025), Fallahgoul (2025), and Cartea and Shi (2025).

more traditional methods. For instance, [Gu *et al.* \(2020\)](#) use more than 900 predictors to compare the out-of-sample predictive performance of various ML algorithms for monthly excess stock returns from 1957 to 2016. They show that the use of ML algorithms leads to a significant improvement in return predictability due to their ability to account for non-linear relationships and their ability to handle a far larger number of predictors than traditional methods.

For similar reasons, several studies document the superiority of ML algorithms in predicting firms' earnings. For instance, [Cao and You \(2024\)](#) compare the predictive performance of decision trees and neural networks for corporate earnings for U.S. firms from 1975 to 2019 to the performance of standard linear forecasting models (e.g., a random walk model or an AR(1) model). They train ML algorithms to predict yearly earnings one, two, and three years ahead using 60 features obtained from firms' financial statements (e.g., balance sheets or income statements), and observe that ML models achieve higher out-of-sample forecasting accuracy than simpler models.

These improvements in out-of-sample forecasting performance do not arise from algorithms alone; they depend on complementary investments in data, computing resources, and specialized human capital. Because acquiring and developing these inputs may involve substantial fixed costs, access to the most sophisticated predictive models may differ across financial intermediaries.

2.3.3 Towards autonomous decision-making

Besides improving forecasting, ML can also aid decision-making. For example, RL algorithms can learn to select the optimal action in a given state of the world. This learning paradigm is especially well-suited to problems that involve sequential decision-making be-

cause it takes into account both the immediate reward from an action and its consequences for future states (and thus future rewards).¹⁰

In the financial industry, RL has been applied to problems such as market making, optimal trade execution, and portfolio optimization (Bai *et al.*, 2025). For example, Ning *et al.* (2021) train a Deep Q-learning model to optimize the timing of an investor seeking to liquidate its inventory and show that it outperforms a standard benchmark for most stocks in their sample.¹¹

Another development is the emergence of generative AI (GenAI), which refers to AI systems designed for the generation of new content (e.g., text, images, code, audio, or video). LLMs are a popular subset of GenAI systems that specialize in the generation, interpretation, and processing of text, and are foundational to tools such as ChatGPT, Claude, and Copilot. The development of modern GenAI was greatly accelerated by the introduction of the Transformer architecture (Vaswani *et al.*, 2017), which relies on attention mechanisms and became the foundation of modern LLMs. GenAI is best understood not as a learning paradigm in its own right, but as a class of models commonly developed through training pipelines that may combine self-supervised and supervised learning, and RL-based methods.¹² GenAI can support the predictive role of algorithms by interpreting documents, extracting information from unstructured data or generating summaries. It can also support decision-making by issuing recommendations, formulating plans, writ-

¹⁰See Hamblly *et al.* (2023) for an overview of the theory of RL models.

¹¹Recent advances in portfolio optimization also showcase the potential of RL algorithms for dynamic asset allocation, see, e.g., Choudhary *et al.* (2025).

¹²Modern GenAI models are typically trained through a multi-stage process. First, they are trained on large volumes of unlabeled data via self-supervised learning, where the labels are derived from the data itself (e.g., next-word prediction in text). This stage is commonly followed by supervised learning on human-curated datasets to improve performance on specific tasks and to align the model outputs better with intended use cases. Finally, many GenAI systems incorporate RL from human feedback, in which models are optimized using evaluative feedback or preference rankings provided by human reviewers to encourage desirable properties such as accuracy, coherence, and safety (Zhang *et al.*, 2023).

ing or executing code, and interacting with external tools, although a human typically remains in the decision loop.¹³

Both RL and GenAI feed into the emergence of agentic AI, which refers to autonomous systems designed to pursue complex goals with minimal human intervention ([Acharya et al., 2025](#)). Unlike traditional AI, which depends on structured instructions, agentic AI demonstrates a sense of autonomy that is reflected in its adaptability, advanced decision-making capabilities, and self-sufficiency.¹⁴ Potential use cases in finance range from autonomous trading agents that monitor markets and decide when and what to trade to banks deploying agents that combine data sources in real time and dynamically adjust the credit limits of clients. An agentic system therefore unites prediction- and decision-making.

2.4 Compute

Training and using ML algorithms require a very large number of computations. For this reason, the third driver of the rise of AI in finance is innovation in chip design, which allows a greater number of calculations per unit of time at lower cost. As a result of these innovations, computing power (c in the production function (1)) has become cheaper over time; that is, \$1 invested in chips buys more computations per unit of time (see [Nordhaus \(2021\)](#) and [IMF \(2024\)](#)).¹⁵

Without these advances, it would be excessively costly and time-consuming to train

¹³Early empirical work corroborates the view of a significant uptake of GenAI among investors. For instance, [Bertomeu et al. \(2025\)](#) find that the number and accuracy of equity analysts' forecasts declined during Italy's temporary ban on ChatGPT in 2023. [Blankespoor et al. \(2026\)](#) provide descriptive and survey evidence of retail investors' adoption of GenAI tools.

¹⁴See, e.g., [Aldridge et al. \(2026\)](#) for a survey of agentic AI in finance.

¹⁵A chip's computing power is expressed in floating-point operations per second (FLOPS). A floating-point operation refers to a single computation performed on floating-point numbers (i.e., numbers with decimal points, expressed in scientific notation). The NVIDIA A100 Tensor Core GPU can perform 312×10^{12} FLOPS.

the ML algorithms described in the previous section.¹⁶ In fact, many ML algorithms are not new; rather, improvements in computing power have enabled their practical use. At the same time, the development of such algorithms has been a strong driver of demand for chips over the past decade and has spurred further innovation in chip design, contributing to even lower computing costs.¹⁷

Anecdotal evidence suggests that the financial industry is investing heavily in computing power, either by expanding in-house capacity or by renting computing resources from cloud service providers.¹⁸ For instance, the PTF XTX is building a large data center in Finland, precisely to increase its computing power, while Citadel Securities (another PTF) moved its data and algorithms to Google cloud computing services to train and test their algorithms.¹⁹ Other intermediaries are developing research teams to take advantage of quantum computing for risk management and pricing of complex products.²⁰

2.5 Labor

Changes in the availability of data and algorithms for their processing affect the productivity of workers in the production function (1), and hence the demand for labor in the financial sector. Figure 3 shows that job ads requiring AI-related skills (“AI-intensity of

¹⁶For instance, Brogaard and Zareei (2023) use ML algorithms to search for profitable trading rules based on past price movements (technical analysis). They note that “the average time needed to find the optimum trading rules for a diversified portfolio of 10 NYSE/AMEX volatility deciles for the 40-year sample using a computer with an Intel Core (TM) CPU i7-2600 and 16 GB RAM is 459.29 days (11,022.97 hours).” To overcome this issue, they have to run their analysis on “a supercomputer with an Intel Xeon E5-2690v4 CPU and 128 GB memory per node using parallel computing procedures.”

¹⁷Demirer *et al.* (2025) illustrate the speed of change in the market for LLMs and document that the “price of intelligence” has already fallen by a factor of 10 between 2023 and 2025.

¹⁸As an explanation of this move, Citadel’s head of research said: “The data that we work with is massive, the computations even more so. It includes data from the exchanges we work with but we also subscribe to additional, external types of data that researchers use to build these models.” See Financial Times, April 9, 2024.

¹⁹See “High-speed trader XTX Markets to build vast data centre in Finland”, Financial Times, April 12, 2024 and “Citadel Securities moves data and algorithm testing to Google Cloud”, Financial Times, April 9, 2024

²⁰See “Wall Street banks ramp up research into quantum finance”, Financial Times, January 2020.

job ads”) have been steadily increasing since 2010 and accelerated recently with the rise of GenAI. Moreover, the growth in the “AI intensity” of job ads has been particularly strong in the financial sector, which has increased from around 1% in 2012 to around 3.5% in 2024. These numbers are consistent with AI affecting the production function of financial services, and changing the demand of financial firms for labor.

A more refined analysis is provided by [Eisfeldt and Schubert \(2025\)](#), who measure the exposure of finance occupations to GenAI, relative to other occupations. Using the O*NET (Occupational Information Network) database, they break down each occupation into a number of different tasks with a description, and then use an LLM to assess the extent to which each task could be performed by GenAI.²¹ This method assigns an “AI exposure” to each occupation, which can be aggregated to measure the exposure of workers in different sectors of the economy. The authors show that the exposure index of finance occupations is twice that of the U.S. average, second only to IT jobs.

²¹The O*NET database is the primary U.S. database describing the tasks, skills, knowledge, abilities, work activities, and work context associated with occupations. It is maintained by the U.S. Department of Labor.




Figure 3: AI intensity of job ads across different sectors. Source: PwC AI Jobs Barometer. See [PwC \(2025\)](#) for methodological details.

The extent to which advances in AI affect the demand for labor depends on the properties of the production function (1). [Abis and Veldkamp \(2024\)](#) provide a structural estimate of a variant of (1) using measures of data (D) and labor (L) from the U.S. asset management industry over 2010-2018. They find that the amount of labor required to convert a given amount of data into knowledge has fallen over this period (e.g., fewer analysts are necessary for each dataset).²² This development may decrease the demand for human labor via a substitution effect. However, it could also induce firms to get more data and hire additional workers in an effort to take advantage of expanded production possibilities.

The aggregate impact of data and AI on labor demand is likely to mask important heterogeneity across different types of jobs. To study this heterogeneity, a growing liter-

²²In addition, the importance of data relative to labor in knowledge production varies across asset managers and is higher for those that rely more heavily on data-intensive processes or advanced information technologies.

ature goes beyond production functions with aggregate labor as in (1) and instead relies on the “task-based framework” pioneered by Autor *et al.* (2003) and developed by Acemoglu and Autor (2011).²³ The core idea is that production requires the execution of a certain number of tasks, each of which can be allocated either to human workers (who may have different skills) or a “machine” (at a certain cost). In this environment, technological progress implies that machines become cheaper and can be given tasks that used to be performed by humans. A robust insight from this literature is that workers with skills substitutable by machines will suffer (a substitution effect), while workers with skills complementary to the machines will benefit. Importantly, the task composition of jobs evolves as retained workers focus on tasks where their skills provide them with a comparative advantage over the machine, such as non-routine, judgment, interpersonal, or contextual tasks. The predictions of this framework for the impact of advances in AI on labor demand depend at least in part on whether algorithms are used for generating predictions or for autonomous decision-making (Section 2.3).

In predictive applications, algorithms transform data into signals (i.e., produce knowledge K_i in (1)), while humans retain responsibility for converting those signals into decisions. Such applications may therefore complement workers whose skills lie in interpretation and decision-making (Agrawal *et al.*, 2019). Empirically, Cao *et al.* (2024) find evidence for such complementarity between human analysts and AI.²⁴ Aragon *et al.* (2025) show that the availability of more alternative data benefited discretionary investment funds relying on humans with AI skills, relative to quantitative investment funds.²⁵

²³See, e.g., Acemoglu and Restrepo (2018) for a survey of this literature and a recent formal framework.

²⁴However, there is also evidence for a displacement of analyst jobs, see Grennan and Michaely (2020).

²⁵Consistent with this view, the Head of AI research at PTF XTX Markets recently stated that “even in quant and trading, there’s a layer AI hasn’t cracked. Generating a truly novel hypothesis or understanding a truly unseen situation, that’s still human.” See <https://www.linkedin.com/pulse/lecture-harvard-ai-finance-bigger-questions-atlas-wang-5wple/>

Both papers thus provide evidence that machines and humans are complementary when machines are used for forecasting. This can be the basis of an optimistic view of the future for finance workers: a change in the task content of finance jobs, with a reallocation from more “routine” tasks towards more creative ones, and an increase in the total number of jobs due to general gains in productivity (e.g., [Brynjolfsson *et al.*, 2018](#); [Aghion *et al.*, 2025](#)).

According to this view, technological progress in finance could just follow a trend towards the automation of an increasing number of tasks, a trend that dates back to at least the Industrial Revolution and has affected all industries. Historically, this process has created new jobs and destroyed existing ones ([Acemoglu and Restrepo, 2019](#)), benefited skilled labor in the aggregate, but harmed displaced skilled workers ([Kogan *et al.*, 2023](#)). In recent decades, automation seems to have increasingly affected workers at the upper end of the skill distribution ([Autor, 2015](#)), with an increased exposure for workers in finance since the 1980s ([Kogan *et al.*, 2023](#)).

However, this optimistic view for finance workers rests on the assumption of complementarity between human skills and AI. Instead, autonomous and agentic systems may take over the decision layer described in Section 2.1, and combine information collection, prediction, and decision-making end-to-end. As a result, they could displace entire bundles of tasks rather than individual ones. For instance, [Cong *et al.* \(2024\)](#) challenge the idea that AI will focus on forecasting and humans on decision-making by training a deep learning algorithm to select profitable investment strategies directly instead of simply forecasting returns. Empirically, their algorithm, which does not rely on human decision-making, generates an impressive alpha of 13%.

GenAI has the potential to further tip the balance towards AI being a substitute rather

than a complement to human labor. However, the analysis in [Eisfeldt and Schubert \(2025\)](#) paints a mixed picture in this regard for the financial sector. In some occupations, the GenAI exposure comes mostly from tasks that are not critical to the occupation, showing a potential for complementarity. For instance, managers may have more time to focus on managing their subordinates, and financial advisors to interact with their clients. In other occupations, however, critical or “core” tasks are more exposed, such as for financial analysts or accountants.

It is still early to tell whether the impact of AI on finance occupations and finance workers will be more in line with the complementarity view, or on the contrary will be more disruptive. As underlined by [Mihet and Philippon \(2019\)](#), it takes significant time for firms to adapt to technological revolutions and exploit the new possibilities they offer, so early evidence is at best an imperfect guide to changes to come. One topic that has received some attention lately is the potentially differential impact of AI on junior vs. senior job positions, see Box 3 for a discussion.

Box 3: Will junior positions disappear?

The impact of AI is unlikely to be even across different positions within a firm's hierarchy, which may require changing the organization of work. The possibility that the number of junior positions will shrink is an important concern in many areas of finance. For example, many investment banking tasks currently performed by junior analysts can possibly be allocated to GenAI.^a

This is not the only conceivable scenario though, as GenAI can also be used to record and codify the tacit knowledge of experienced workers, making their experience less valuable (Brynjolfsson *et al.*, 2025). This possibility is indeed observed, for instance, in coding occupations. Peng *et al.* (2023) show that more junior coders benefited more from the introduction of the AI tool GitHub Copilot relative to more experienced ones.^b

Ide and Talamàs (2025) study these forces in an equilibrium model and show that the impact of AI will depend on how exactly it enters the production function. On average, across industries, the introduction of ChatGPT is associated with a decline in job postings, particularly for entry-level positions (Liu *et al.*, 2025). An important concern is that reducing entry-level positions makes production more efficient in the short run, but at the cost of reducing the pool of experienced workers down the road, which may make the industry less efficient in the long run. See Ide (2025) for a formal analysis, and Acemoglu (2024) for a discussion of the relevant inefficiencies.

^aSee, e.g., “AI forecast to put 200,000 European banking jobs at risk by 2030”, Financial Times, 31 December 2025.

^bYe (2025) also finds evidence for a signaling mechanism: as standard coding tasks are easier to do with the help of GitHub Copilot, the most able coders signal their ability by doing more exploration and more creative work.

3 The Diffusion of AI in the Securities Industry

Information production is central to all segments of the financial industry, including financial markets, banking, and insurance. However, AI adoption has been particularly rapid in securities trading and asset management. Survey evidence confirms that adoption is widespread, but uneven. For instance, the European Securities and Markets Authority (ESMA) reports that 55% of 728 financial firms (mostly investment managers) surveyed in 2024 have some AI use case in production or in experimentation ([ESMA, 2026](#)). Moreover, 70% of these firms anticipate an increase in their AI-related investment over the period 2025–27. Interestingly, there is significant heterogeneity, with larger firms exhibiting faster adoption (more than 75% with active AI deployment) and more significant investment. These broad findings are echoed in other surveys (e.g., [Financial Conduct Authority, 2024a](#)).

It is therefore instructive to zoom in on market participants that can be considered “early adopters” because it provides a glimpse at likely future developments. [Figure 4](#) plots the AI-intensity of job ads for different types of financial intermediaries, based on web-scraped data and broadly following the methodology of [Abis and Veldkamp \(2024\)](#) (see [Appendix A](#) for details). As can be seen, the AI intensity of job offers posted by banks and large general-purpose asset managers (e.g., BlackRock, Amundi) is broadly in line with the most recent point estimate for the aggregate financial sector in [Figure 3](#) at around 3%. However, so-called “pod shops”—multi-manager hedge funds—exhibit an AI intensity that is more than twice as high, at around 7.5%. Yet, this high level is still easily exceeded by PTFs, whose job offers have an AI intensity of more than 10%. We discuss these two types of market participants below.




Figure 4: AI Intensity of job ads for different financial firms. The underlying data are job ads collected via web-scraping from company websites, and the classification methodology follows [Abis and Veldkamp \(2024\)](#). See Appendix A for details.

3.1 Proprietary trading firms

As their name suggests, PTFs solely trade with their own capital and do not manage client assets.²⁶ Much of their trading activity focuses on market making and (statistical) arbitrage in exchange-traded assets such as equities, derivatives, and ETFs. Several of today's large PTFs were founded by options traders in the 1980s and 1990s (e.g., Optiver, IMC, Susquehanna, DRW). Their growth accelerated significantly in the 2000s when most trading activity moved from physical trading pits to electronic limit order books, and regulatory reform (Reg NMS in the U.S., MiFID in Europe) paved the way for competition among trading venues and facilitated automated order routing.

Together with advances in IT, these structural changes enabled PTFs to compete

²⁶Virtually all major PTFs are privately held firms. Notable exceptions include Virtu Financial LLC and Flow Traders N.V.

directly with traditional broker-dealers in liquidity provision. Over time, many banks reduced or exited their proprietary trading activities, particularly following the implementation of regulatory reforms in the aftermath of the 2007-08 financial crisis such as the Volcker Rule.

PTFs traditionally operate at the technological frontier. They employ automated trading systems that analyze market data, generate trading signals, and execute orders with minimal human intervention. For much of the 2000s and early 2010s, competitive advantage was defined by speed (“latency”), and firms invested billions in fiber-optic cables and microwave towers to shave microseconds off round-trip times between Chicago and New York.²⁷ This period also coined the term “high-frequency trading”.

More recently, however, anecdotal evidence suggests that competition among trading firms has shifted away from the latency arms race to the concentrated use of AI.²⁸ Consistent with this view, the share of patents related to AI and ML in the context of high-frequency or algorithmic trading has increased sharply since the mid-2010s (IMF, 2024). Moreover, evidence from the labor market is also consistent with this view. Figure 5 shows that the AI intensity of job ads posted by PTFs has increased strongly from 2022 to 2025.

²⁷Famously, the book “Flash Boys” by Michael Lewis details the story of the construction of an 827-mile fiber-optic cable—cutting directly through mountains—to reduce the signal transmission time from Chicago to New Jersey by 4 milliseconds.

²⁸See the discussion in “Odd Lots: How Hudson River Trading Actually Uses AI”, Bloomberg Podcasts, 31 October 2025.




Figure 5: AI Intensity of job ads posted by PTFs. The underlying data are job ads collected via web-scraping from company webpages, and the classification methodology follows [Abis and Veldkamp \(2024\)](#). See Appendix A for details. Information prior to 2025 was obtained via the Wayback Machine.

While PTFs are famously secretive, a study by the Dutch Authority for Financial Markets (AFM) provides some insight into the type of tools used in the industry ([AFM, 2022](#)). Virtually all firms use ML for their trading algorithms and rely on limit order book data as the primary source of information. While firms report that they prefer relatively simple models, industry accounts suggest that neural networks have been used since at least 2017.²⁹ Finally, the AFM study suggests that trading algorithms based on RL were not widely used at the time of writing. Consistent with this, an earlier OECD reported that ML models in high-frequency trading are predominantly employed offline for model calibration or the refinement of decision rules ([OECD, 2021](#)).

Taken together, available evidence suggests that PTFs mainly deploy ML in its predictive role, while autonomous decision-making remains comparatively limited. However,

²⁹See “Odd Lots: How Hudson River Trading Actually Uses AI”, Bloomberg Podcasts, 31 October 2025.

this assessment is based on survey evidence from 2021-2022, so the sharp rise in AI-related hiring documented above suggests that the practice may have changed since.

3.2 Asset management

The asset management sector is diverse, spanning mutual and hedge funds, sovereign wealth funds, family offices, and insurance companies. Moreover, a significant share of the sector outside the United States is integrated into universal banks ([Ferreira *et al.*, 2018](#)). Given that these tend to be large organizations with substantial headcount, there is a diverse range of AI use cases, ranging from investment-related tasks (e.g., exploration of new asset classes, extraction of trading signals from large/complex datasets, portfolio optimization and asset allocation) to back-office functions (document management, workflow optimization) and customer-facing activities (robo advising, marketing).

Matters are more clear-cut for quantitatively focused asset managers. Given limited organizational overhead, they are poised to devote most AI-related efforts to return prediction and trade execution. While the origins of quantitative investing can be traced back to the development of basic finance models (1950s to 1970s), the industry only began to reach scale in the 1980s as the use of models became practical with advances in computing power and increased data availability. While early quantitative approaches were often “structural” (i.e., grounded in theory), this period also laid the foundation for a more data-driven approach spearheaded by hedge funds such as Renaissance Technologies (see Box 4) or D.E. Shaw, largely run by mathematicians and computer scientists. Since then, quantitative investing has grown tremendously, and spilled over from institutional into retail products.³⁰

³⁰ [Abis \(2022\)](#) shows that around 7% of all active mutual funds can be classified as “quantitative” as of 2018.

Box 4: Renaissance Technologies

In 1978, Jim Simons left a distinguished career in academia to establish the hedge fund Monometrics, drawing on his background in mathematics and physics. The firm was renamed Renaissance Technologies in 1982 and quickly became a pioneer in quantitative trading.

From the start, Renaissance distinguished itself through its focus on hiring scientists—particularly mathematicians, physicists, statisticians, and computer scientists—instead of traditional traders and economists. Although the fund’s early days predate the current terminology of AI and ML, its approach—using algorithms and models trained on historical data to generate and execute trades—represents an early large-scale deployment of predictive, data-driven models in financial markets (Zuckerman, 2019). Simons was skeptical of the Efficient Market Hypothesis, which motivated his search for statistical regularities that could be monetized. The firm maintained a constant focus on identifying new sources of alpha, because it recognized that most inefficiencies are short-lived.

Renaissance Technologies’ flagship vehicle, the Medallion Fund, is widely regarded as the most successful hedge fund in history. Between 1994 and mid-2014, Medallion generated an average annual return of 71.8% before fees. It has been closed to external investors since 1993 and is available only to current and former employees and their families. Renaissance offers other funds to outside investors, but these vehicles rely on strategies that differ from those employed by Medallion.^a

^aSee “Renaissance’s shrinking hedge funds”, Financial Times, 20 September 2024.

Given the long-standing use of sophisticated predictive models by quantitative hedge

funds, competition dictates that successful firms must remain at the technological frontier. This includes the use of state-of-the-art prediction models (including deep neural networks), the use of alternative data (several hedge funds have dedicated business units³¹), and vast computing resources. Naturally, hedge funds are also embracing GenAI, including for the generation of investment signals. *Sheng et al. (2026)* report that about 60% of surveyed hedge funds in their study (45 U.S. funds) use GenAI tools for investment purposes, and usage is associated with superior performance.

Given the significant costs of technology investments, the rise of AI is also transforming existing business models. One example is the strong growth of multi-manager hedge funds—so-called “pod shops”—since the early 2010s. In essence, some of the larger hedge funds (e.g., Citadel, Millennium Management) have moved towards a platform structure where infrastructure (IT, data, compliance, etc.) is shared, while investment decisions are made in a decentralized way by splitting funds across a large number of individual teams. These “pods” operate their own trading strategies and P&L subject to strict risk limits. This enables pod shops to reap significant scale economies by leveraging resources (in particular data, due to its non-rival nature), diversify individually volatile strategies (a key drawback of traditional hedge funds), and maintain a competitive market for talent.

While the growth of pod shops has exceeded that of the remaining hedge fund industry, they account for only about 8% of total hedge fund assets. However, they employ significant leverage and therefore exhibit a disproportionate financial market footprint. Goldman Sachs estimates that their share among U.S. hedge fund equity holdings increased from 14% in 2014 to 27% in 2022.³²

³¹Examples include the Data Strategies Group (DSG) at Citadel or the Market Intelligence Unit at Point72 (see <https://point72.com/market-intelligence>).

³²See “Are hedge fund pioneers facing the end of a golden era?”, *Financial Times*, 24 August 2023.

The growing dominance of AI-powered funds naturally raises the question of whether more traditional fund managers are at risk of being displaced. [Bonelli and Foucault \(2025\)](#) show that the performance of traditional (discretionary) fund managers declines when stocks become covered by alternative data and that this decline is stronger for fund managers that rely on traditional sources of expertise (e.g., specialized industry knowledge or geographical location). They also find that these fund managers divest from stocks after they become covered by alternative data providers. Consequently, if discretionary fund managers can produce signals that cannot be replicated by machines (e.g., because they specialize in predictions that cannot rely on past precedents), the type of information reflected in asset prices could change (see Section [4.1](#)).³³

4 Efficiency Gains or Market Distortions?

As explained in previous sections, the adoption of AI in finance enables finance professionals to form more precise forecasts and to automate decisions with self-learning algorithms. A key issue is whether this evolution improves overall welfare by enhancing efficiency and resource allocation, or whether it introduces costs, externalities, and risks that could reduce social welfare.³⁴

In this section, we analyze these issues. One possible benefit of AI in finance is that it improves the accuracy of investors' forecasts of future cash flows, thereby reducing the

³³There is evidence that AI-driven investing is also making inroads into private markets. [Bonelli \(2025\)](#) documents a strong increase in the number of venture capital firms hiring data scientists to build infrastructure for pre-investment screening. These data-driven firms appear to become more effective at screening startups similar to existing firms for which data is available. Conversely, humans retain an edge in evaluating more novel business ideas.

³⁴The answer to this question is not obvious. In fact, long-run progress in information technologies does not seem to have reduced the cost of financial intermediation. [Philippon \(2015\)](#) shows that the intermediation cost has remained roughly constant at about 2% of intermediated asset per year from 1886 to 2012, despite considerable progress in information technologies over the same period.

gap between asset values and fundamentals and improving capital allocation in financial markets. We discuss this possibility in Section 4.1. Another issue is whether AI weakens or strengthens market power for intermediaries, a point that we examine in Section 4.2. Last, informational rents arising from adverse selection or agency issues are important in the financial industry. We discuss whether AI mitigates or amplifies these rents in Section 4.3.

4.1 Price informativeness

The combination of ML algorithms and big data enables financial professionals to obtain more accurate forecasts of assets' future cash flows. Intuitively, this should bring asset prices closer to their “fundamental value”. If so, asset prices should be more informative about future cash flows, thereby improving capital allocation (Bond *et al.*, 2012).³⁵

Consistent with this conjecture, several empirical studies (e.g., Zhu, 2019; Mukherjee *et al.*, 2021; Katona *et al.*, 2025; Bonelli and Foucault, 2025) provide evidence that the availability of alternative data reduces the impact of public news on asset prices, which suggests improved price informativeness. Intuitively, alternative data enable sophisticated investors to access information before they are publicly disclosed. As they trade on this information, prices become more informative about future disclosures.

In addition, several studies (e.g., Brogaard *et al.*, 2014; Barbopoulos *et al.*, 2023) show that high-frequency traders, who use computing power and algorithms to react quickly to new information, make prices more informationally efficient in the sense that prices reflect available information more rapidly. Overall, these findings suggest that the rise of AI-powered intermediaries can enhance price informativeness.

³⁵Bai *et al.* (2016) and Farboodi *et al.* (2021) find that long-run trends in price informativeness have been heterogeneous across groups of stocks over the last 50 years.

There are potential obstacles, however. First, there is an important difference between trading on information that will eventually be disclosed (e.g., trading a few seconds or a few weeks before a macroeconomic announcement) and trading on information that would never be revealed in the absence of the effort expended to discover it. [Hirshleifer \(1971\)](#) refers to the former activity as “foreknowledge” and the latter as “discovery.” Both types of information generate private profits for those obtaining it, but the social value of the former is smaller than that of the latter and can even be negative because it destroys risk-sharing opportunities and creates adverse selection. Arguably, using algorithms to react fast to the information in public news (e.g., a regulatory filing by a firm or a macroeconomic announcement) falls more into the category of “foreknowledge” than “discovery”.³⁶

A second issue is that AI tools can be used to infer other traders’ information from their trades. That is, instead of collecting data and using ML tools to discover new information useful for forecasting fundamentals, some traders collect market data (limit order book data, trade and quote data, etc.) to extract information about other traders’ fundamental information.³⁷ This reduces the expected profits of traders focused on discovering new fundamental information and therefore weakens the incentives to discover such information. Consistent with this idea, [Weller \(2017\)](#) finds that algorithmic trading has a negative effect on price informativeness.

This behavior can also push fundamental investors to shift their trading to trading mechanisms that are less exposed to information leakage, such as dark pools ([Neumeier et al., 2023](#)) or closing auctions, and thus lead to a fragmentation of trading. Consistent

³⁶See [Brogaard et al. \(2014\)](#) for evidence that high-frequency traders place buy and sell orders immediately following macroeconomic announcements. See [Barbopoulos et al. \(2023\)](#) for evidence that algorithms are increasingly used to extract information from the text of new regulatory filings immediately after they become publicly available.

³⁷See [Farboodi and Veldkamp \(2020\)](#) or [Yang and Zhu \(2019\)](#) for theoretical formalizations of this scenario and [Van Kervel and Menkveld \(2019\)](#) for empirical evidence.




Figure 6: Percentage of non-PTF volume traded in the LOB. Data include all transactions in large-cap shares on Euronext Amsterdam. See Appendix B for the classification of PTFs and non-PTFs.

with this argument, Figure 6 shows that the share of non-PTFs (e.g., brokers executing client orders) that is executed in the electronic limit order book (continuous trading) has declined quite sharply over recent years as more and more of their trading activity has migrated to the closing auction. While part of this development may be fueled by the rise of passive investing, the magnitude of this change (a decline from 65% to less than 45%) suggests that other forces are at work, and buy-side market participants are actively avoiding the limit order book and interactions with PTFs.

Traders using algorithms to act swiftly on signals generated from news increase the speed at which prices reflect information and therefore make prices more informationally efficient. However, there is a trade-off between speed and accuracy. Signals extracted quickly from just-released announcements are likely to be less accurate than those that have undergone time-consuming cross-validation. This is a concern because fast trading based on imprecise yet informative signals can compete away the returns to investors who

invest in producing more precise signals, which are more costly to generate and therefore can only be exploited with a delay. [Dugast and Foucault \(2018\)](#) show that this mechanism can ultimately lead to prices being less informative than they would be in the absence of fast trading on imprecise signals.

A third issue concerns the *nature* of information. As mentioned earlier, many academic studies show that various alternative datasets can be used to forecast corporate earnings and stock returns. However, as pointed out by [Dessaint *et al.* \(2024\)](#), the evidence is primarily for short-term earnings and returns. For this reason, the availability of alternative data can shift forecasters' (e.g., equity analysts) attention toward the short term at the expense of the quality of their long-term forecasts. In fact, [Dessaint *et al.* \(2024\)](#) document that over the last 30 years, equity analysts' forecasts have become more informative about short-term earnings (up to two years) and less informative about long-term earnings (up to five years). These trends are more pronounced in industries for which analysts rely more frequently on alternative data for forming their forecasts.

Lastly, ML algorithms learn from *past* data. Thus, when data are scarce or the environment is non-stationary, their forecasting power and usefulness decline. In contrast, humans may still be able to form forecasts in such situations by relying on “models of the world.” As machines replace humans in investing and trading in securities markets, one risk is that the ability of markets to forecast cash flows for innovative projects (for which, by definition, there are few past examples to learn from) declines. If so, they might be priced less accurately by the market, potentially leading to a decline in investment in such projects.³⁸

In sum, the nature of the information produced with AI matters in finance. Because

³⁸Managers are more likely to underinvest in projects whose net present value takes more time to be reflected in prices (see [Fishman and Hagerty \(1989\)](#) and [Dessaint *et al.* \(2025\)](#)).

there is no mechanism to align the private and social benefits of the information produced in financial markets, the risk is that AI is used to produce information with low social value and even discourages the production of information that has the highest value for society.

4.2 Market power

AI is likely to change the competitive environment in which financial firms operate. The underlying mechanisms can be different, depending on whether AI is used to improve forecasting or for the automation of decision-making.

4.2.1 Improvement in forecasting ability

The increased use of data in the economy as a whole seems to have led to higher market concentration. [Farboodi *et al.* \(2019\)](#) study a model where firms accumulate data about their customers, which helps them to produce better products and increase sales. This gives rise to a positive feedback loop between data and sales, so that economic activity becomes concentrated among a few large firms. Furthermore, [Begenau *et al.* \(2018\)](#) show that it is easier for financial markets to assess the value of data-rich companies, which lowers their cost of capital and helps to further fuel their growth. [Babina *et al.* \(2024\)](#) show that the same dynamics remain true with the adoption of AI specifically over the period 2010-2018.

However, these general mechanisms are not directly applicable to financial markets because transaction data are typically collected and sold by exchanges to other market participants. This implies that entrants and smaller firms can easily access vast amounts of data to train algorithms, and AI may lower the cost of doing so. Indeed, some researchers

argue that AI may in general facilitate the creation of new start-ups because broad access to off-the-shelf models, cloud computing and public data will lead to a decrease in firms' entry costs. Consistent with this view, [Bena *et al.* \(2026\)](#) and [Cai *et al.* \(2025\)](#) show that the diffusion of GenAI has led to an increase in startup formation for the overall economy (i.e. across all sectors). Thus, improvements in data and technology may lower the barriers to entry in the financial sector and reduce market concentration.

Nevertheless, there are several countervailing channels through which concentration in the finance industry may ultimately increase. First, trading firms retain some proprietary information because market data are typically anonymized. Accordingly, they are the only ones to know which trades they were responsible for, and what motivated them. For some transactions, in particular off-exchange, the executing firm also knows the identity of its counterparty, while this information remains hidden from the rest of the market. Large traders thus have more information about who conducted transactions, giving them a strategic advantage over competitors.³⁹ Consistent with this view, recent survey evidence ([ESMA, 2026](#)) shows that large investment firms exhibit a greater reliance on internal data for model training and are adopting AI to a greater extent than smaller competitors.

Second, market data are typically sold by exchanges and trading platforms themselves, each of which has de facto monopoly power over data on trades occurring on its own platform ([Glosten, 2020](#)). This can lead to inefficiently high prices for such data and thereby create barriers to entry. The economics of data sales by exchanges are complex, though, because exchanges are selling both data and trading services. [Maurin and Fou-](#)


³⁹David Shaw, quoted in [Lo \(2017\)](#), states: “We were taking profit and paying for experimentation. We were able to run randomized controlled trials, for example, in which we could compare two models or parameter values to see which one performed better in actual trading. Analyzing the results of live trading taught us things that couldn’t be learned by studying historical data. We were doing a lot of trading, and the data we accumulated during one round of trading was helping us to increase our returns in the next round.”

[cault \(2026\)](#) show that investors can be better off overall when an exchange can monetize its data because it incentivizes the exchange to lower its trading fees. Moreover, they find that data sales can increase investors participation (even from those not buying data) because exchange data reduces market makers' exposure to adverse selection. [Hendershott *et al.* \(2025\)](#) provide evidence for such externalities of data sales for the New York Stock Exchange.


Third, AI may give rise to significant economies of scale because operating at the technological frontier requires significant fixed-costs investments into prediction models.⁴⁰ This development may dampen, or even overturn the usual diseconomies of scale that arise as price impact harms the profitability of trading strategies when they are scaled up ([Berk and Green, 2004](#); [Pastor *et al.*, 2015](#)). Consistent with this view, [Farboodi *et al.* \(2024\)](#) study a model where investors' private valuations of data increase with size, but the effect is dampened by price impact. [Baron *et al.* \(2019\)](#) provide evidence that even small technological advantages (trading speed) among PTFs can lead to a high level of market concentration due to a “winner-takes-all” environment.

Thus, the overall effect of AI on financial market concentration is ultimately an empirical question. To illustrate this issue in the European context, Figure 7 plots the evolution of PTFs' market shares (left-hand side) during recent years for different asset classes. PTFs significantly increased their market share in Futures and ETF trading, which contrasts to a more modest increase for options and a roughly unchanged picture for cash equities.

⁴⁰See [De Ridder \(2024\)](#) for a macroeconomic model where fixed IT investment costs act as a barrier to entry.



(A) Futures markets



(B) Cash and options markets

Figure 7: PTF market shares and concentration. PTF market shares are defined as $(PTF \text{ euro volume} / \text{Total euro volume}) \times 100$. Top-3 within PTF share is $(\text{Top-3 PTF euro volume} / \text{PTF euro volume}) \times 100$. Source: MiFID II data, own calculations. See Appendix B for details.

The Figure also illustrates the development of market concentration within the PTF sector, represented by the relative market share of the 3 largest firms. We only observe an increase in market concentration for options, while the other four series exhibit a constant level, or even a decline in concentration. Together, the evidence paints a mixed picture. The PTF sector has grown on aggregate and expanded into new asset classes, which is suggestive of both economies of scale and scope. At the same time, there is no evidence that the sector itself is becoming more concentrated, suggesting that it remains overall contestable. Accordingly, technological improvement appear to be having both pro- and anti-competitive effects, and their relative strengths may differ across markets.

Figure 8 studies the question whether technology is a driving force towards market concentration more directly. If technology is important, we would expect the best firms to capture a higher market share on more volatile days, when having faster and more accurate signals is more valuable. The figure shows a scatter plot of market concentration (measured by the Herfindahl-Hirschman Index) among PTFs against price volatility, at the stock-day level. Within the PTF sector, volume is indeed concentrated among fewer firms when a stock is more volatile on a particular day. The slope of the regression line is 0.0104, meaning that, on average, a doubling of stock-level volatility is associated with an increase in market concentration of about 4 per cent (based on its sample mean of 0.197). An interesting open question is whether current advances in AI will benefit the firms that already have better technology, which would reinforce this effect further, or on the contrary will level the playing field.




Figure 8: PTF market concentration and volatility. This figure depicts a bin scatterplot of the relationship between PTF market concentration (measured by the Herfindahl-Hirschman index) and volatility (measured in logs) at the stock-day level. Data include all trading during 2020–2026 on Euronext Amsterdam in Dutch large-cap (AEX) and mid-cap (AMX) securities. Source: EMIR data, MiFID II data, own calculations. See Appendix B for details.

4.2.2 Autonomous decision-making

AI-powered algorithms for autonomous trading may directly affect competition by profoundly changing the strategic interaction between market participants. A growing literature studies the interaction between price-setting algorithms and the implications for the nature of competition.

One strand of the literature shows that simple algorithms can coordinate on collusive equilibria that human agents would find difficult to reach (Calvano *et al.*, 2020), with the risk that ever-more sophisticated algorithms may be able to reach such outcomes in complex real-world environments (see Keppo *et al.* (2025) for recent experiments with LLMs of different sophistication). More generally, Cartea *et al.* (2026b) show that learning algorithms can achieve any payoff profile that is feasible and individually rational, meaning that algorithms can learn to collude where the collusive outcome is sustained by a reward-punishment mechanism.

Another strand of the literature points out that, to the contrary, the lack of sophistication of algorithms can lead them to “mistakenly” coordinate on noncompetitive prices (see, e.g., Asker *et al.* (2024), Abada *et al.* (2024), and Zhao and Xu (2024)). In fact, even simple and reasonable design choices for algorithms such as reacting to competitors’ price changes are sufficient to prevent competition from working as expected (Brown and MacKay (2023), Wang *et al.* (2022)).

Importantly, in many of these environments it may be a best response for a firm to keep a simple algorithm leading to a noncompetitive outcome, rather than using a more sophisticated one that would eventually lead to lower prices. Thus, both sophisticated pricing algorithms (via sophisticated collusive strategies) and naive algorithms (via lack

of competitive aggressiveness) may ultimately lead to less competitive prices.⁴¹ Empirically, [Assad *et al.* \(2024\)](#) provide clean causal evidence that the introduction of pricing algorithms in the German retail gasoline market led to higher prices.

Some recent papers study the interaction of algorithms in financial markets. [Colliard *et al.* \(2026\)](#) run experiments with simple RL algorithms in a market-making game. They find that the algorithms only partially learn to compete despite playing for many rounds, and do not reach competitive spreads. While this is generally in line with the “lack of sophistication” view, the authors suggest that financial markets may be special due to the strength of adverse selection and noisy payoffs. Intuitively, noise leads to more competitive outcomes by making it more difficult for algorithms to learn to undercut their competitors, whereas adverse selection reduces noise and thus leads to more competitive outcomes.

[Dou *et al.* \(2023\)](#) study the complementary case of algorithms used by informed investors. They show that these algorithms can learn collusive strategies when there is little noise in the market, or on the contrary may behave noncompetitively by mistake when there is too much noise. In general, it seems that outcomes crucially depend on the sophistication of algorithms relative to the complexity of the environment, which varies across assets and over time. Another approach to study the role of sophistication is to compare how different families of algorithms fare in the same environment, as in [Cartea *et al.* \(2022a\)](#) in a market-making context.⁴²

The rise of AI-powered traders raises new issues for policymakers. First, predicting and understanding the long-run behavior of even relatively simple RL algorithms is difficult, especially in environments in which several algorithms interact. This complicates

⁴¹See also [Harrington \(2022\)](#) for the interesting, and common, case in which a specialized firm designs and sells pricing algorithms to market participants.

⁴²Other recent papers in this fast-growing literature include [Banerjee and Szydlowski \(2025\)](#), [Dou *et al.* \(2025b\)](#), and [Gufler *et al.* \(2026\)](#).

policy intervention. The challenge does not lie in describing how algorithms learn from received feedback and adjust their actions accordingly, but rather in predicting the long-run behavior that emerges from the learning process (see Box 5). This can lead to unintended consequences such as algorithms learning to manipulate prices, or to outcomes that are difficult to anticipate such as convergence to noncompetitive pricing.

Box 5: “Simple” models and unintended consequences

An ML model may appear to have a complex objective but a simple structure, and vice versa. Consider a market maker that uses an RL algorithm to optimize profits. The market maker uses a technique called “Q-learning”, which estimates the value of “placing a bid” or “placing an offer” in a particular state. The objective (i.e., to make profits) is complex, in the sense that the model learns which action to take in a semi-autonomous manner via many interactions in an environment. However, the output of the RL algorithm can be represented in table-format, where each cell represents the expected cumulative reward from taking action a in state s (and then behaving optimally thereafter), which is a relatively interpretable mapping from inputs to decisions.

Compare to this to a deep neural network trained to predict future stock prices. The objective is relatively simple (i.e., minimize the error between label and prediction), but the output of the model is a collection of potentially millions of numbers, none of which individually give much insights into the “rationale” of the prediction.

A model’s learning paradigm also has implications for the morality of the behavior the model might learn. An RL algorithm takes into account both the immediate and future rewards of taking an action (e.g., sending a new bid or offer) in a particular

state. As a result, certain actions may be valuable not only because of their short-term payoffs, but because they place the algorithm on a profitable trajectory. It is precisely the capacity of RL algorithms to learn the relationship between present actions and future rewards that can give rise to unintended consequences such as tacit collusion.^a Importantly, such behavior is difficult to anticipate because it depends not only on the algorithm's own learning dynamics but also on the strategic interactions with other agents. By contrast, the objective of the apparently more complex deep neural network is purely predictive and does not explicitly link current actions to future payoffs, so it cannot learn the strategic unintended behaviors that may emerge via RL.^b

^aFor example, [Wang and Wellman \(2017\)](#) show that RL can learn to manipulate markets by submitting misleading orders. More generally, [Cartea *et al.* \(2025a\)](#) use a dynamic model of the limit order book to derive the conditions under which an RL algorithm will learn to manipulate the order book.

^bHybrid approaches, such as Deep Q-learning, combine reinforcement learning with deep neural networks. While powerful, these models may amplify concerns related to both interpretability and unintended or emergent behavior.

Second, there is still very little research on how regulatory frameworks should be adjusted to discipline trading algorithms instead of humans. In particular, one important question is how to design rules for markets populated by AI-powered agents. Some early work on this question includes [Banchio and Skrzypacz \(2022\)](#), who show that algorithms behave much more competitively under a second-price auction format than under a first-price auction format, and [Johnson *et al.* \(2023\)](#), who show that adjusting the rules of a marketplace can drastically change the likelihood that algorithms collude. For securities markets specifically, [Cartea *et al.* \(2022b\)](#) show that reducing the tick size can make algorithms more competitive, which shows that the design of securities markets must take into account the presence of algorithms.

4.3 AI and informational rents

4.3.1 Adverse selection costs

Asymmetric information is a key friction in financial markets. For instance, a hedge fund might have private information about a stock's incoming earnings announcement, which may induce it to buy or to sell the stock. In these situations, the profits of the informed party are obtained at the expense of the less informed party and are therefore a cost for the latter (adverse selection costs). To cover these costs, less informed investors are willing to buy (sell) assets only at a discount (markup). This phenomenon can explain why secondary markets are not perfectly liquid, why assets may be underpriced in primary markets, and why markets may break down.

The effect of AI on asymmetric information and adverse selection costs is ambiguous. On the one hand, AI could enable less informed investors to catch up with more informed investors, thereby reducing adverse selection costs. On the other hand, rather than leveling the playing field, AI can aggravate information asymmetries by increasing the informational edge of investors who make significant investments in datasets, algorithms, and computing power.

For example, the SEC's introduction of EDGAR in 1993 lowered the cost of accessing regulatory filings by requiring U.S. public companies to submit their information electronically to an online repository. [Gomez \(2024\)](#) shows that this reduced informational asymmetries between firm's managers and investors.⁴³ However, he also finds an increase in information asymmetries between sophisticated and unsophisticated investors. One possible reason is that sophisticated investors have greater information processing pos-

⁴³Using the same event, [Gao and Huang \(2019\)](#) document an increase in the informativeness of retail investors' trades.

sibilities. In fact, [Barbopoulos *et al.* \(2023\)](#) show that the number of downloads from EDGAR by machines has increased strongly over time and currently vastly outnumbers downloads by humans.

Similarly, [Zhao \(2023\)](#) documents that a technological improvement that increased the readability of U.S. firms' regulatory filings for machines benefited less sophisticated institutional investors because it leveled the playing field by reducing the cost of processing these filings for everyone. While AI could in theory have such an impact, the current evidence points that the net effect is in the opposite direction. [Sheng *et al.* \(2026\)](#) show that hedge funds have adopted AI more than the rest of the asset management industry, and this has increased the return gap between the two groups. This is also echoed by [Eisfeldt *et al.* \(2023\)](#), who show that markets expect firms with more data to benefit more from advances in AI.

Overall, these studies illustrate how the digitization of information, combined with AI tools, can have ambiguous effects on informational asymmetries. Interestingly, these findings echo those regarding the effects of algorithmic and high-frequency trading on informational asymmetries. Faster access to information enables PTFs both to pick off quotes that are not updated quickly enough when new information arrives and to update their own quotes before being picked off. The net effect on adverse selection, and ultimately illiquidity, is theoretically ambiguous (see, for instance, [Foucault *et al.*, 2016](#)).

Empirically, several studies show that technological upgrades enabling PTFs to become even faster generally lead to higher adverse selection costs, suggesting that these upgrades increase informational asymmetries ([Shkilko and Sokolov, 2020](#)). However, this does not imply that fast trading harms overall liquidity because it may simultaneously affect other determinants of market liquidity, e.g., inventory costs ([Foucault *et al.*, 2023](#)).

As mentioned previously, some PTFs often chase winner-take-all trading opportunities, which can lead to excessive investment in technologies that speed up access to these opportunities. Another cause of over-investment is the absence of incentives for traders to internalize adverse selection costs inflicted on others. Thus, the social costs of investment in algorithms, data, and computing power for the sake of obtaining an informational edge exceed the private costs, which naturally leads to excessive investment from a social viewpoint (Biais *et al.*, 2015; Hoffmann, 2014).⁴⁴

4.3.2 Moral hazard

An obvious difference between humans and AI-powered agents is that the former have their own agenda and must be incentivized, whereas the latter can be programmed (i.e., directly endowed with an objective function). Thus, one might expect that AI will be particularly useful in the financial industry because moral hazard is an important friction. However, there are opposing forces at play.

On the one hand, AI can reduce moral hazard by helping to monitor human employees. This can be done directly, e.g., by using AI-powered software to check that remote employees are working (Jiang *et al.*, 2025), or indirectly, by comparing an employee's forecast or performance to that generated by an AI tool. For instance, fund manager performance can be benchmarked against the performance of an AI agent.⁴⁵

On the other hand, humans may also be tempted to rely more on AI and exert less effort themselves. Colliard and Zhao (2025) study a model of task automation in the financial industry in which workers earn significant rents due to moral hazard (Philippon and

⁴⁴See Glode and Ordoñez (2024) for a generalization of the argument to industries in which investment in technology can be used to capture rents.

⁴⁵Cohen *et al.* (2026) show that a large share of fund manager trading is predictable, and this predictability is associated with poorer performance. Their approach can in principle be used to compute how much value a fund manager brings relative to an automated agent trained on past data.

[Reshef, 2012](#)). These rents imply that workers will appropriate part of the productivity increase generated by the machine. As workers can effectively “free-ride” on the machine, firms have lower incentives to automate tasks, leading to under-automation. However, as rents also make humans expensive, financial firms also hire too few humans in the model, to economize on rents, and may rely exclusively on machines instead, in which case there is over-automation. In this case an improvement in technology can paradoxically lead to lower production and welfare, as firms automate to recapture the workers’ rents despite production being less efficient.⁴⁶ This raises the question of how to allocate control optimally between human employees and AI systems to balance the different effects of AI on moral hazard ([Athey et al., 2020](#); [Zhong, 2025](#)).

In addition to affecting existing moral hazard problems, AI can also create new ones. As LLMs are becoming increasingly proficient at mimicking human behavior, the assumption that AI agents obey a well-known objective function may not hold. This is particularly true for algorithms with a complex objective function (see Box 5), or for third-party models where the source code and reward function are unknown.⁴⁷ [Fudenberg and Liang \(2025\)](#) provide a model where a human principal interacts with an AI agent that may potentially be “misaligned” and represent information in a misleading way. In finance, [Scheurer et al. \(2024\)](#) give the intriguing example of a simulation in which an LLM conducts insider trading despite being instructed not to do so, and then lies about its actions. [Cartea et al. \(2026a\)](#) show that LLMs participate in both rational

⁴⁶[Jiang et al. \(2025\)](#) study workers’ time allocation in the U.S. over 2004-2023 (across all industries) and show that workers who are more exposed to AI are indeed complementary to AI. As a result, employers make them work longer hours. However, the associated productivity gains are mostly appropriated by firms rather than by workers, so firms outside the financial sector have significant market power over their employees.

⁴⁷This is currently a widespread concern in cybersecurity, where AI agents have been found to autonomously try to hack into external systems. See this widely discussed case: AI Security Institute, “Incident Report: unsanctioned agent behaviour during cyber testing”, August 4, 2026.

and irrational bubbles, and are more likely to participate when they recognize that the trading environment is a speculative bubble.

More broadly, a risk with learning algorithms is that they may learn illegal or undesirable strategies unless their programmer had the foresight to explicitly forbid or penalize such strategies *ex ante*. Examples of RL algorithms learning to collude (e.g., [Calvano *et al.* \(2020\)](#)) or to manipulate the market ([Cartea *et al.*, 2025a](#); [Dou *et al.*, 2025a](#)) fall in this category.

5 Financial Stability

In this section, we discuss the potential implications of an increased use of AI models by market participants on financial stability.

5.1 Herding and crowding

Several research papers and policy reports (e.g., [Aldasoro *et al.*, 2024](#); [Financial Stability Board, 2024](#); [Danielsson and Uthemann, 2025](#)) forward the concern that commonality in either AI models or data (or both) may increase the risk of herding, or “crowded trades”.⁴⁸ In the context of equation (1) this would be reflected by intermediaries using similar algorithms (a) and overlapping datasets (D), which would lead to similar knowledge and, ultimately, decisions. As a consequence, even relatively small shocks could be amplified through substantial fire-sale externalities and generate market crashes as “everyone runs for the exit” ([Pedersen, 2009](#)). Indeed, there are several well-documented cases of crowded trades and their disorderly unwinding, such as the 1994 “Bond market massacre” ([Borio](#)

⁴⁸It is useful to distinguish herding/crowding from bubbles. While the former refers to behavior (investors mimicking others), the latter refers to an outcome (price far above the “fundamental” value of the traded asset).

and McCauley, 1996), the 2007 “Quant event” (Khandani and Lo, 2008), or the 2020 turmoil in the U.S. treasury market associated with the cash-futures basis trade (Barth and Kahn, 2025).

This argument is quite powerful. Indeed, all market participants share the same objective of predicting future prices, even if at different frequencies. Economies of scale imply that the market for off-the-shelf AI models (especially LLMs) is highly concentrated, and their widespread use for making investment decisions could indeed lead to the above-mentioned concerns. Similarly, the market for financial market data has also become increasingly concentrated over the past two decades as exchanges have been turned into for-profit enterprises that generate an increasing share of their revenues via data sales (see Chapter 4 in Duffie *et al.* (2022)) and a wave of mergers has led to a market with few leading data providers (e.g., Refinitiv, Bloomberg, S&P). Intuitively, the risk of a “data monoculture” appears particularly rampant in the domain of low-latency trading, where the use of co-location creates incentives to make predictions solely based on limit order book data. Finally, the biases of autonomous learning algorithms discussed in Section 4.2.2 may also favor herding. Indeed, Anand *et al.* (2026) show that AI agents can create new sources of financial instability, and the exact source depends on the type of AI architecture. Q-learning algorithms playing a mutual fund redemption game redeem more often than predicted by the efficient-equilibrium benchmark. This behaviour is due to a systematic learning bias, so that fragility emerges even in the absence of standard coordination problems. By contrast, LLMs playing the same game struggle to coordinate when the setting features multiple equilibria and coordination matters, thus making redemptions highly unpredictable.

However, there are important counterarguments based on the informational role of

prices and the incentive to acquire information. It is well known that the economic value of informative signals dissipates quickly as they become available to a larger number of agents and are reflected instantaneously in prices (e.g., [Grossman and Stiglitz, 1980](#)). This, in turn, creates incentives to look for other signals not yet reflected in market prices. Accordingly, market participants will find it profitable to invest in different (i.e. proprietary) models and acquire alternative datasets. The underlying mechanism is consistent with the model of [Wagner \(2011\)](#), where commonality in exposures gives rise to higher liquidation risk. Accordingly, investors aim to choose portfolios that differ from those of others, and optimally forego diversification benefits.

5.2 Competition and market concentration

As discussed previously, AI may increase market concentration in certain market segments such as market making and (statistical) arbitrage. This can affect financial stability through two channels, namely the rise of systemically important players and the competition-stability trade-off. While these aspects are not new, these channels may become amplified through the economies of scale inherent in the use of AI models in financial markets (see Section [4.2](#)).

First, consider the issue of systemic relevance. Typically, the “too-big-to-fail” problem centers around banks due to their central role for retail deposits, credit provision and payments. However, there were several instances in modern history (most notably LTCM in 1998, but more recently Lehman Brothers in 2008, Knight Capital in 2012, and Archegos in 2021) where large trading firms, and especially hedge funds, were perceived as systemic and became subject to a coordinated bailout/rescue. The concern typically centers on the indirect contagion effects resulting from large exposures and leverage. While these

concerns continue to apply, the AI-powered increase in the size (and scope) of PTFs may turn them into new systemic players. Since some of these firms are the main intermediaries in key markets (e.g., equities, futures, FX), their collapse could lead to a significant impairment of market functioning (see also [Yadav and Brogaard, 2026](#)).

Indeed, empirical evidence suggests that the failure of a major intermediary can create a “capacity gap” that cannot be filled instantaneously, but rather takes time to close. [Bogousslavsky *et al.* \(2021\)](#) study the 2012 failure of Knight Capital, a large market maker in U.S. equities (see Box 6 for details).⁴⁹ They show that the resulting decline in market liquidity was significant, especially for large institutions (as measured by implementation shortfall). However, the overall fallout of this episode was relatively contained as the firm was able to raise new capital within a week and was ultimately taken over by a competitor. In particular, the latter point suggests that a competitive ecosystem can be helpful in containing bankruptcy costs. However, this may become more difficult if economies of scale lead to increased market concentration.

Box 6: Knightmare on Wall Street

Knight Capital Group was a major U.S. market-making and electronic trading firm. The firm used high-frequency trading algorithms and became one of the most active equity traders, accounting for roughly 17% of NYSE and NASDAQ volume at its peak.

However, on the morning of 1 August 2012, a disaster unfolded. A software update inadvertently activated a trading algorithm that was never meant to run in pro-

⁴⁹Similarly, [Brewer and Jackson \(2000\)](#) provide evidence that the bankruptcy of investment bank Drexel Burnham Lambert had a profound and long-lasting effect on the U.S. junk bond market in which the firm commanded a 50% market share of total underwriting. Finally, [Di Maggio *et al.* \(2017\)](#) present evidence consistent with this channel from a large over-the-counter dealer collapse in 2008.

duction.^a The algorithm flooded the market with buy and sell orders, effectively buying high and selling low in dozens of stocks. Within 45 minutes, Knight had accumulated roughly \$440 million in losses, wiping out most of its capital.^b The firm never fully recovered and, after a rescue investment, was eventually acquired by rival Getco. Besides being disastrous for the firm itself, Knight's unintended trading activity significantly moved prices in many of the affected stocks.

European regulation on algorithmic trading aims to prevent incidents like this. For example, EU Regulatory Technical Standards require firms to test each material change to a trading algorithm before deployment, maintain kill switches that can immediately halt trading if an algorithm goes rogue, and apply robust pre-trade controls to block erroneous orders ([European Commission, 2017](#)).

More recently, the FCA fined Citigroup Global Markets Limited £27,766,200 for failures in its systems and controls ([Financial Conduct Authority, 2024b](#)). On 2 May 2022, a trader intended to sell a basket of equities worth \$58 million. However, the trader made an input error that created a basket valued at \$444 billion instead. In total, equities worth \$1.4 billion were sold across European exchanges before the order was canceled. The incident coincided with a short-lived, but material, drop in several European indices.

^aThe “Power Peg” was a test-program designed to buy high and sell low in order to verify the behavior of its other proprietary trading algorithms in a controlled environment. It was never intended for live trading.

^bSee “Knight Capital Says Trading Glitch Cost It \$440 Million”, The New York Times, 2 August 2012.

Second, we discuss the competition-stability nexus. It is not obvious that market concentration necessarily increases systemic risk. Famously, [Keeley \(1990\)](#) suggests that the prospect of steady future rents derived from market power can incentivize dominant

actors (in his case, banks) to behave prudently. By contrast, fierce competition erodes margins and therefore creates incentives for greater risk-taking. Naturally, this channel also applies to financial markets. If AI increases the ability of certain investors (large PTFs and hedge funds) to exert market power and extract rents, their incentives to behave prudently also increase.⁵⁰

5.3 Malicious behavior

While AI is a potential source of efficiency, it can equally be used for harmful activity. Interestingly, as explained previously (see Sections 4.3.2 and 4.2.2), this need not even be intentional.

One major concern is the use of AI for “disinformation”, i.e., the generation and subsequent propagation of false information. This is, of course, not a new phenomenon; an early well-known case is the “Great Stock Exchange Fraud of 1814”.⁵¹ However, the advent of GenAI means that the costs of this form of market manipulation have fallen sharply. Combined with market participants’ incentives for engaging in the use of alternative data (e.g., social media), this creates a perfect storm that may adversely affect markets.

One case in point is a 2023 deepfake of an explosion at the U.S. Pentagon, which subsequently went viral on social media and had a short-lived but significant effect on the U.S. stock market.⁵² The ability to use sophisticated AI models to generate new content may thus further amplify the mechanism explored in [Dugast and Foucault \(2018\)](#). As data

⁵⁰[Boyd and De Nicolo \(2005\)](#) develop an offsetting channel in a setting where asset risk is endogenous to market prices. While this mechanism is appealing in the case of banking, it has not been explored in financial market trading.

⁵¹On 21 February 1814, a group of fraudsters falsely claimed that Napoleon had been killed, which caused significant price disruptions until the rumors were debunked. Ultimately, eight people were convicted of conspiracy to defraud.

⁵²A somewhat related event, even before the age of AI, was the 2013 “Twitter Crash,” resulting from a hack of the AP Twitter account and the posting of a message about an explosion at the White House.

become abundant, investors are confronted with a dilemma: they must decide whether to use the data quickly at face value or delay trading to engage in time-consuming “due diligence”. In equilibrium, this can paradoxically reduce the informativeness of asset prices.

Another significant financial stability risk arises from model misalignment, as already discussed in Section 4.3.2. Profit-maximization by “rogue” algorithms may be individually rational, but turn into a destabilizing force at the aggregate level (Danielsson and Uthemann, 2025). While economically perhaps less interesting, the use of AI to disrupt systemic infrastructures such as exchanges or clearing houses is an even more pressing issue from a policy perspective. Rogue actors may develop AI agents to disrupt financial markets for private (e.g., ransom) or geopolitical motives. See IMF (2026) for a more detailed overview of recent developments.

5.4 Model and data risks

The complex nature of AI models makes it difficult to validate their outputs. This is particularly true for LLMs because their probabilistic nature favors statistical plausibility over factual correctness or logic. This can lead to “hallucinations”, i.e., factually incorrect outputs that appear credible to the observer. Accordingly, this feature can jeopardize effective (internal and regulatory) model oversight and risk controls, and lead to erroneous (investment) decisions.⁵³ Importantly, such issues become more rampant when relying on pre-trained models from third parties. While trading firms face legal obligations⁵⁴ aimed at ensuring model oversight (e.g., by employing qualified staff), including in the case of

⁵³One interesting anecdote is the 2023 case of two lawyers who submitted a legal brief that included fictitious cases generated by ChatGPT. See “New York lawyers sanctioned for using fake ChatGPT cases in legal brief”, Reuters, 26 June 2023.

⁵⁴See the formal requirements under MiFID II in European Commission (2017).

outsourcing, achieving compliance may become more challenging as model complexity increases.

New AI models (including, but not limited to, LLMs) facilitate the use of alternative data, including text and graphics. While this holds great promise in terms of increased predictive ability on average, unproven data quality and limited data histories can amplify the likelihood of overfitting, reduce generalizability, and render model predictions more volatile. However, such “model risk” is a general feature of any sophisticated statistical model.

5.5 The nature of intermediation

In the aftermath of the 2007-08 financial crisis, regulators moved to curtail banks’ proprietary trading activities (e.g., via the Volcker Rule or the supplementary leverage ratio). While these measures have likely contributed to making the financial system safer, it has been shown that they have also limited banks’ ability to act as intermediaries in financial markets ([Adrian *et al.*, 2017a](#); [Cenedese *et al.*, 2021](#); [Breckenfelder and Ivashina, 2021](#)).

Since PTFs and hedge funds have expanded to fill the void, this raises the question of whether the nature of liquidity provision has changed, and how this will be affected by the rise of AI-powered trading. For example, PTFs have historically aimed to intermediate by relying on technology, not balance sheet capacity. Together with a focus on less noisy short-term predictions and the minimization of clearing and settlement costs, PTFs typically end the trading day flat. Accordingly, some economists have questioned whether the contribution of PTFs to market liquidity is meaningful (e.g., [Haldane, 2012](#); [Stiglitz, 2014](#)).

An increasing reliance on AI may lead to an evolution of PTFs’ business model away




Figure 9: PTF Inventory Ratio. Source: MiFID II data and own calculations. See Appendix B for details.

from low-latency trading towards mid-frequency strategies. Media reports and interviews with insiders as well as the recent trading loss at Jane Street related to the failure of the hedge fund “Situational Awareness” appear to corroborate this development.⁵⁵ To examine this potential shift in PTF behavior empirically, we employ supervisory data on Dutch equities available to the AFM under MiFID II to compute the daily inventory ratio for a large sample of PTFs and aggregate the resulting numbers (across firms and days) at the quarterly level. A higher ratio means that PTFs keep their inventory position over longer time periods. Figure 9 shows that PTFs’ inventory ratio has increased significantly in recent years, consistent with a potential shift in their strategies.

In tandem with the growing reach of PTFs, the footprint of hedge funds has increased significantly in recent years, including in government bond markets.⁵⁶ However, the empirical evidence suggests that some of the underlying activity involves considerable leverage,

⁵⁵See See Financial Times, “Hedge funds and high-frequency traders are converging”, 29 September 2025 and “Jane Street’s \$15 Billion Loss Lays Bare Its Hedge Fund Side”, Bloomberg, 17 August 2026.

⁵⁶See the ECB Blog “Hedge funds: good or bad for market functioning?”, 23 September 2024.

and the disorderly unwinding of these positions can lead to severe market disruptions as evidenced by the 2020 Treasury Market Turmoil ([Barth and Kahn, 2025](#); [Kashyap *et al.*, 2025](#)).

Taken together, the evidence suggests that non-banks like PTFs and hedge funds have taken over a significant role in liquidity provision. While traditional dealer banks still play a significant role in many key markets (including FX, fixed income, credit, and OTC derivatives), their risk-bearing capacity has declined ([Adrian *et al.*, 2017b](#); [Duffie *et al.*, 2022](#)). Moreover, the rise of sophisticated competitors has led to pressure on intermediation margins, which creates incentives for using capital for more profitable activities. Accordingly, overall market liquidity has become more fragile and more fleeting.

6 Conclusion

We argue in this paper that AI can be understood as a technological shock that improves financial intermediaries' ability to leverage data for making better predictions and, ultimately, decisions. Framed this way, AI's consequences extend well beyond automating existing tasks: it reshapes the demand for labor, the competitive structure of the securities industry, the informativeness of prices, agency conflicts, and the stability of the financial system. Our discussion suggests that the ultimate impact of AI on financial markets is not a foregone conclusion in either direction. On the positive side, AI may make markets more efficient, help level the playing field, and ameliorate a variety of frictions. At the same time, it may also tilt market participants towards socially unproductive (or outright harmful) activities, give rise to market power and algorithmic collusion, create new forms of moral hazard, and jeopardize financial stability.

Our synthesis highlights several interesting avenues for future research. We highlight three examples. First, a significant body of research studies how technology improves predictions and the resulting implications for financial markets. By contrast, research on the effects of automated decision-making is naturally more limited due to its nascent nature. This is likely to prove as a fruitful area of research going forward. Second, the industrial organization of financial markets deserves more attention against the background of the mentioned ambiguous effects of AI on market concentration and competition. Potential challenges for pushing this agenda will be the availability of granular data and the development of suitable empirical methods. Third, further research could study financial stability issues, such as an analysis of the risks of “model and data monoculture” and their role for the amplification of shocks.

References

- ABADA, I., LAMBIN, X. and TCHAKAROV, N. (2024). Collusion by mistake: Does algorithmic sophistication drive supra-competitive profits? *European Journal of Operational Research*, **318** (3), 927–953.
- ABIS, S. (2022). *Man vs Machine: Quantitative and Discretionary Equity Management*. Working paper, University of Colorado Boulder.
- and VELDKAMP, L. (2024). The changing economics of knowledge production. *The Review of Financial Studies*, **37** (1), 89–118.
- ACEMOGLU, D. (2024). Harms of AI. In *The Oxford Handbook of AI Governance*, Oxford University Press.
- and AUTOR, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics, Handbook of Labor Economics*, vol. 4, 1st edn., Elsevier, pp. 1043–1171.
- , —, HAZELL, J. and RESTREPO, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, **40** (S1), S293–S340.
- and RESTREPO, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, **108** (6), 1488–1542.
- and — (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, **33** (2), 3–30.
- ACHARYA, D. B., KUPPAN, K. and DIVYA, B. (2025). Agentic AI: Autonomous intelligence for complex goals - a comprehensive survey. *IEEE Access*, **13**, 18912–18936.
- ADRIAN, T., BOYARCHENKO, N. and SHACHAR, O. (2017a). Dealer balance sheets and bond liquidity provision. *Journal of Monetary Economics*, **89**, 92–109.
- , FLEMING, M., SHACHAR, O. and VOGT, E. (2017b). Market liquidity after the financial crisis. *Annual Review of Financial Economics*, **9**, 43–83.
- AFM (2022). Machine learning in trading algorithms application by Dutch proprietary trading firms and possible risks. <https://www.afm.nl/en/sector/actueel/2023/maart/her-machine-learning>, accessed: 14/April/2023.
- AGHION, P., BUNEL, S., JARAVEL, X., MIKAELSEN, T., ROULET, A. and SNØGAARD, J. (2025). How different uses of AI shape labor demand: Evidence from France. *AEA Papers and Proceedings*, **115**, 62–67.
- AGRAWAL, A., GANS, J. S. and GOLDFARB, A. (2019). Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, **33** (2), 31–50.
- ALDASORO, I., GAMBACORTA, L., KORINEK, A., SHREETI, V. and STEIN, M. (2024). *Intelligent Financial System: How AI Is Transforming Finance*. Tech. Rep. 1194, Bank for International Settlements.
- ALDRIDGE, I., AN, J., BURKE, R., CAO, M., CHIEN, C.-Y., DENG, K., DENG, R., GAO, Y., GUO, O., HE, S. *et al.* (2026). Agentic artificial intelligence in finance: A comprehensive survey. *arXiv preprint arXiv:2604.21672*.

- ANAND, K., LEONELLO, A., PANETTI, E. and KAZINNIK, S. (2026). *Ex Machina: financial stability in the age of artificial intelligence*. Working Paper Series 3225, European Central Bank.
- ARAGON, G., KIM, S. and NANDA, V. (2025). *Will AI Replace or Enhance Human Intelligence in Asset Management?* Working paper.
- ASKER, J., FERSHTMAN, C. and PAKES, A. (2024). The impact of artificial intelligence design on pricing. *Journal of Economics & Management Strategy*, **33** (2), 276–304.
- ASSAD, S., CLARK, R., ERSHOV, D. and XU, L. (2024). Algorithmic pricing and competition: Empirical evidence from the german retail gasoline market. *Journal of Political Economy*, **132** (3), 723–771.
- ATHEY, S. C., BRYAN, K. A. and GANS, J. S. (2020). The allocation of decision authority to human and artificial intelligence. In *AEA Papers and Proceedings*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol. 110, pp. 80–84.
- AUTOR, D. H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, **29** (3), 3–30.
- , LEVY, F. and MURNANE, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, **118** (4), 1279–1333.
- BABINA, T. (2026). *Understanding Firms’ AI Efforts and Their Economic Impact*. Tech. Rep. 35123, National Bureau of Economic Research.
- , FEDYK, A., HE, A. and HODSON, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, **151**, 103745.
- BAI, J., PHILIPPON, T. and SAVOV, A. (2016). Have financial markets become more informative? *Journal of Financial Economics*, **122** (3), 625–654.
- BAI, Y., GAO, Y., WAN, R., ZHANG, S. and SONG, R. (2025). A review of reinforcement learning in financial applications. *Annual Review of Statistics and Its Application*, **12** (1), 209–232.
- BANCHIO, M. and SKRZYPACZ, A. (2022). Artificial intelligence and auction design. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, EC ’22, New York, NY, USA: Association for Computing Machinery, pp. 30–31.
- BANERJEE, S. and SZYDLOWSKI, M. (2025). *Trading against Algorithms: Price Dynamics and Risk-sharing in a Market with Q-learners*. Working paper.
- BARBOPOULOS, L. G., DAI, R., PUTNINS, T. J. and SAUNDERS, A. (2023). *Market Efficiency When Machines Access Information*. Working paper, Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3783221.
- BARON, M., BROGAARD, J., HAGSTROMER, B. and KIRILENKO, A. (2019). Risk and return in high-frequency trading. *Journal of Financial and Quantitative Analysis*, **54** (3), 993–1024.
- BARTH, D. and KAHN, R. J. (2025). Hedge funds and the treasury cash-futures basis trade. *Journal of Monetary Economics*, p. 103823.

- BEGENAU, J., FARBOODI, M. and VELDKAMP, L. (2018). Big data in finance and the growth of large firms. *Journal of Monetary Economics*, **97**, 71–87.
- BENA, J., BIAN, B. and GIANNETTI, M. (2026). Prompted to start: How generative ai is transforming entrepreneurship. *HKU Jockey Club Enterprise Sustainability Global Research Institute Paper*, (2025/182).
- BERK, J. and GREEN, R. (2004). Mutual fund flows and performance in rational markets. *Journal of Political Economy*, **112** (6), 1269–1295.
- BERTOMEU, J., LIN, Y., LIU, Y. and NI, Z. (2025). The impact of generative AI on information processing: Evidence from the ban of chatgpt in Italy. *Journal of Accounting and Economics*, p. 101782.
- BIAIS, B., FOUCAULT, T. and MOINAS, S. (2015). Equilibrium fast trading. *Journal of Financial Economics*, **116** (2), 292–313.
- BLANKESPOOR, E., CROOM, J. and GRANT, S. M. (2026). Generative AI and investor processing of financial information. Available at SSRN: <https://ssrn.com/abstract=5053905>.
- BOGOUSLAVSKY, V., COLLIN-DUFRESNE, P. and SAĞLAM, M. (2021). Slow-moving capital and execution costs: Evidence from a major trading glitch. *Journal of Financial Economics*, **139** (3), 922–949.
- BOND, P., EDMANS, A. and GOLDSTEIN, I. (2012). The real effects of financial markets. *Annual Review of Financial Economics*, **4** (Volume 4, 2012), 339–360.
- BONELLI, M. (2025). Data-driven investors. *The Review of Financial Studies*, p. hhaf078.
- and FOUCAULT, T. (2025). *Does Big Data Devalue Traditional Expertise? Evidence from Active Funds*. Working paper, HEC Paris.
- BORIO, C. E. and MCCAULEY, R. N. (1996). The anatomy of the bond market turbulence of 1994. In *Risk Management in Volatile Financial Markets*, Springer, pp. 60–84.
- BOYD, J. H. and DE NICOLO, G. (2005). The theory of bank risk taking and competition revisited. *The Journal of Finance*, **60** (3), 1329–1343.
- BRECKENFELDER, J. and IVASHINA, V. (2021). Bank balance sheet constraints and mutual fund fragility. Available at SSRN 3858543.
- BREWER, E. and JACKSON, W. E. (2000). Requiem for a market maker: the case of Drexel Burnham Lambert and junk bonds. *Journal of Financial Services Research*, **17** (3), 209–235.
- BROGAARD, J., HENDERSHOTT, T. and RIORDAN, R. (2014). High-Frequency Trading and Price Discovery. *The Review of Financial Studies*, **27** (8), 2267–2306.
- and ZAREEI, A. (2023). Machine learning and the stock market. *Journal of Financial and Quantitative Analysis*, **58** (4), 1431–1472.
- BROWN, Z. Y. and MACKAY, A. (2023). Competition in pricing algorithms. *American Economic Journal: Microeconomics*, **15** (2), 109–56.
- BRYNJOLFSSON, E., LI, D. and RAYMOND, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, **140** (2), 889–942.

- , MITCHELL, T. and ROCK, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, **108**, 43–47.
- BUNCIC, D. (2025). Simplified: A closer look at the virtue of complexity in return prediction. *Available at SSRN*.
- CAI, J. J., GU, X., SHENG, L., XIA, M., ZHAO, L. and ZHU, W. (2025). Ai as "co-founder": Genai for entrepreneurship. *arXiv preprint arXiv:2512.06506*.
- CALVANO, E., CALZOLARI, G., DENICOLO, V. and PASTORELLO, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, **110** (10).
- CAO, K. and YOU, H. (2024). Fundamental analysis via machine learning. *Financial Analysts Journal*, **80** (2), 74–98.
- CAO, S., JIANG, W., WANG, J. and YANG, B. (2024). From Man vs. Machine to Man + Machine: The art and AI of stock analyses. *Journal of Financial Economics*, **160**, 103910.
- CARTEA, Á., CHANG, P., CHEN, N. and ZHONG, M. (2026a). AI bubbles with large language models. *Available at SSRN*.
- , — and GARCIA-ARENAS, G. (2025a). *Spoofing and Manipulating Order Books with Learning Algorithms*. Working paper.
- , —, MROCZKA, M. and OOMEN, R. (2022a). AI-driven liquidity provision in OTC financial markets. *Quantitative Finance*, **22** (12), 2171–2204.
- , — and PENALVA, J. (2022b). Algorithms and supracompetitive prices in electronic markets: The impact of tick size. *Available at SSRN 4105954*.
- , —, PENALVA, J. and WALDON, H. (2026b). Algorithmic collusion and a folk theorem from learning with bounded rationality. *Games and Economic Behavior*, **157**, 1–21.
- , CUCURINGU, M. and KIM, D. (2025b). Learning firm characteristics for asset management. *Available at SSRN 5853103*.
- , JAIMUNGAL, S. and PENALVA, J. (2015). *Algorithmic and high-frequency trading*. Cambridge University Press.
- , JIN, Q. and SHI, Y. (2025c). The limited virtue of complexity in a noisy world. *Available at SSRN*.
- and SHI, Y. (2025). The mechanical virtue of random fourier features and model complexity. *Available at SSRN 5930257*.
- CENEDESE, G., DELLA CORTE, P. and WANG, T. (2021). Currency mispricing and dealer balance sheets. *The Journal of Finance*, **76** (6), 2763–2803.
- CHOUDHARY, H., ORRA, A., SAHOO, K. and THAKUR, M. (2025). Risk-adjusted deep reinforcement learning for portfolio optimization: A multi-reward approach. *International Journal of Computational Intelligence Systems*, **18** (1), 126.
- COHEN, L., LU, Y. and NGUYEN, Q. H. (2026). *Mimicking Finance*. Tech. rep., National Bureau of Economic Research.

- COLLIARD, J.-E., FOUCAULT, T. and LOVO, S. (2026). Algorithmic pricing and liquidity in securities markets. *Review of Financial Studies*, **Forthcoming**.
- and ZHAO, J. (2025). *Artificial Intelligence and the Rents of Finance Workers*. Working paper.
- CONG, L. W., TANG, K., WANG, J. and ZHANG, Y. (2024). *AlphaPortfolio: Direct Construction Through Deep Reinforcement Learning and Interpretable AI*. Working paper.
- DANIELSSON, J. and UTHEMANN, A. (2025). Artificial intelligence and financial crises. *Journal of Financial Stability*, p. 101453.
- DASTILE, X., CELIK, T. and POTSANE, M. (2020). Statistical and machine learning models in credit scoring: A systematic literature survey. *Applied Soft Computing*, **91**, 106263.
- DE RIDDER, M. (2024). Market power and innovation in the intangible economy. *American Economic Review*, **114** (1), 199–251.
- DEMIRER, M., FRADKIN, A., TADELIS, N. and PENG, S. (2025). *The Emerging Market for Intelligence: Pricing, Supply, and Demand for LLMs*. Working paper.
- DESSAINT, O., FOUCAULT, T. and FRÉSARD, L. (2024). Does alternative data improve financial forecasting? the horizon effect. *The Journal of Finance*, **79** (3), 2237–2287.
- , — and FRESARD, L. (2025). The horizon of investors' information and corporate investment. Available at <http://dx.doi.org/10.2139/ssrn.4276832>.
- DI MAGGIO, M., KERMANI, A. and SONG, Z. (2017). The value of trading relations in turbulent times. *Journal of Financial Economics*, **124** (2), 266–284.
- DIXON, M. F., HALPERIN, I., BILOKON, P. *et al.* (2020). *Machine learning in finance*, vol. 1170. Springer.
- DOU, W., GOLDSTEIN, I. and JI, Y. (2023). *AI-Powered Trading, Algorithmic Collusion, and Price Efficiency*. Working paper.
- , — and — (2025a). *Financial Market Fragility in the Era of AI Planning*. Working paper.
- , —, LI, Z. and YANG, L. (2025b). *Human Edge, Machine Limits: Algorithmic Herding and AI-Human Competition*. Working paper.
- DUFFIE, D., FOUCAULT, T., VELDKAMP, L. and VIVES, X. (2022). Technology, data and trading in securities markets. *CEPR Press Book Chapters*, **1**.
- DUGAST, J. and FOUCAULT, T. (2018). Data abundance and asset price informativeness. *Journal of Financial Economics*, **130** (2), 367–391.
- and — (2025). Equilibrium data mining and data abundance. *The Journal of Finance*, **80** (1), 211–258.
- EISFELDT, A. L. and SCHUBERT, G. (2025). Generative AI and finance. *Annual Review of Financial Economics*, **17**, 363–393.
- , — and ZHANG, M. B. (2023). *Generative AI and Firm Values*. NBER Working Papers 31222, National Bureau of Economic Research, Inc.

- ESMA (2026). AI adoption and trends in securities markets: EU evidence.
- EUROPEAN COMMISSION (2017). Commission Delegated Regulation (EU) 2017/589. Official Journal of the European Union, L 87, 31 March 2017, pp. 417–448.
- FALLAHGOUL, H. (2025). High-dimensional learning in finance. *arXiv preprint arXiv:2506.03780*.
- FARBOODI, M., MATRAY, A., VELDKAMP, L. and VENKATESWARAN, V. (2021). Where has all the data gone? *The Review of Financial Studies*, **35** (7), 3101–3138.
- , MIHET, R., PHILIPPON, T. and VELDKAMP, L. (2019). Big data and firm dynamics. *AEA Papers and Proceedings*, **109**, 38–42.
- , SINGAL, D., VELDKAMP, L. and VENKATESWARAN, V. (2024). Valuing financial data. *The Review of Financial Studies*, **38** (3), 938–980.
- and VELDKAMP, L. (2020). Long-run growth of financial data technology. *American Economic Review*, **110** (8), 2485–2523.
- FERREIRA, M. A., MATOS, P. and PIRES, P. (2018). Asset management within commercial banking groups: International evidence. *The Journal of Finance*, **73** (5), 2181–2227.
- FINANCIAL CONDUCT AUTHORITY (2024a). Artificial intelligence in UK financial services.
- FINANCIAL CONDUCT AUTHORITY (2024b). FCA fines CGML £27,766,200 for failures in its trading systems and controls.
- FINANCIAL STABILITY BOARD (2024). *The financial stability implications of artificial intelligence*. Report.
- FISHMAN, M. and HAGERTY, K. (1989). Disclosure decisions by firms and the competition for price efficiency. *Journal of Finance*, **44** (3), 633–646.
- FOUCAULT, T., GAMBACORTA, L., JIANG, W. and VIVES, X. (2025). *Artificial Intelligence in Finance, Barcelona Reports*, vol. 7. Paris and London: CEPR Press.
- , KOZHAN, R. and THAM, W. W. (2016). Toxic arbitrage. *The Review of Financial Studies*, **30** (4), 1053–1094.
- , PAGANO, M. and RÖELL, A. (2023). *Market Liquidity: Theory, Evidence, and Policy*. Oxford: 2nd edition, Oxford University Press.
- FUDENBERG, D. and LIANG, A. (2025). *Friend or Foe: Delegating to an AI Whose Alignment is Unknown*. Working paper.
- GAO, M. and HUANG, J. (2019). Informing the market: The effect of modern information technologies on information production. *The Review of Financial Studies*, pp. 1367–1411.
- GLODE, V. and ORDOÑEZ, G. (2024). Technological progress and rent seeking. *The Review of Financial Studies*, **38** (4), 1259–1289.
- GLOSTEN, L. (2020). *Economics of the Stock Exchange Business: Proprietary Market Data*. Working paper.

- GOMEZ, E. A. (2024). The effect of mandatory disclosure dissemination on information asymmetry among investors: Evidence from the implementation of the EDGAR system. *The Accounting Review*, **99** (1), 235–257.
- GRENNAN, J. and MICHAELY, R. (2020). *Artificial Intelligence and High-Skilled Work: Evidence from Analysts*. Working paper.
- GROSSMAN, S. J. and STIGLITZ, J. E. (1980). On the impossibility of informationally efficient markets. *The American Economic Review*, **70** (3), 393–408.
- GU, S., KELLY, B. and XIU, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, **33** (5), 2223–2273.
- GUFLER, I., SANGIORGI, F. and TARANTINO, E. (2026). *(Deep) Learning to Trade: An Analysis of AI Trading and Market Outcomes*. Working paper.
- HALDANE, A. G. (2012). The race to zero. In *The global macro economy and finance*, Springer, pp. 245–270.
- HAMBLY, B., XU, R. and YANG, H. (2023). Recent advances in reinforcement learning in finance. *Mathematical Finance*, **33** (3), 437–503.
- HARRINGTON, J. E. (2022). The effect of outsourcing pricing algorithms on market competition. *Management Science*, **68** (9), 6889–6906.
- HENDERSHOTT, T., RYSMAN, M. and SCHWABE, R. (2025). Stock exchanges as platforms for data and trading. *Journal of Financial Markets*, **75**, 100986.
- HIRSHLEIFER, J. (1971). The private and social value of information and the reward to inventive activity. *The American Economic Review*, **61** (4), 561–574.
- HOFFMANN, P. (2014). A dynamic limit order market with fast and slow traders. *Journal of Financial Economics*, **113** (1), 156–169.
- IDE, E. (2025). *Automation, AI, and the Intergenerational Transmission of Knowledge*. Papers 2507.16078, arXiv.org.
- and TALAMÀS, E. (2025). Artificial intelligence in the knowledge economy. *Journal of Political Economy*, **133** (12), 3762–3800.
- IMF (2024). *Chapter 3: Advances in Artificial Intelligence: Implications for Capital Market Activities*, Washington, DC: International Monetary Fund, p. CH003. Monetary and Capital Markets Department.
- IMF (2026). Artificial intelligence and cybersecurity in the financial sector.
- INTERNATIONAL MONETARY FUND (2024). *Global Financial Stability Report: Chapter 3—Advances in Artificial Intelligence: Implications for Capital Market Activities*. Washington, DC: International Monetary Fund.
- JIANG, W., PARK, J., XIAO, R. and ZHANG, S. (2025). *AI and the Extended Workday: Productivity, Contracting Efficiency, and Distribution of Rents*. Working paper.
- JOHNSON, J. P., RHODES, A. and WILDENBEESE, M. (2023). Platform design when sellers use pricing algorithms. *Econometrica*, **91** (5), 1841–1879.

- KASHYAP, A. K., STEIN, J. C., WALLEN, J. L. and YOUNGER, J. (2025). Treasury market dysfunction and the role of the central bank. *Brookings Papers on Economic Activity*, **2025** (1), 221–265.
- KATONA, Z., PAINTER, M. O., PATATOUKAS, P. N. and ZENG, J. (2025). On the capital market consequences of big data: Evidence from outer space. *Journal of Financial and Quantitative Analysis*, **60** (2), 551–579.
- KEELEY, M. C. (1990). Deposit insurance, risk, and market power in banking. *The American Economic Review*, pp. 1183–1200.
- KELLY, B., MALAMUD, S. and ZHOU, K. (2024). The virtue of complexity in return prediction. *The Journal of Finance*, **79** (1), 459–503.
- KEPPO, J., LI, Y., TSOUKALAS, G. and YUAN, N. (2025). *AI Pricing, Agent Heterogeneity, and Collusion*. Working paper.
- KHANDANI, A. E. and LO, A. W. (2008). *What happened to the quants in August 2007?: Evidence from factors and transactions data*. Tech. rep., National Bureau of Economic Research.
- KOGAN, L., PAPANIKOLAOU, D., SCHMIDT, L. D. and SEEGMILLER, B. (2023). *Technology and Labor Displacement: Evidence from Linking Patents with Worker-Level Data*. NBER Working Papers 31846, National Bureau of Economic Research, Inc.
- KRIZHEVSKY, A., SUTSKEVER, I. and HINTON, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, **25**.
- LEVINE, R. (1997). Financial development and economic growth: Views and agenda. *Journal of Economic Literature*, **35** (2), 688–726.
- LIU, Y., WANG, H. and YU, S. (2025). *Labor Demand in the Age of Generative AI : Early Evidence from the U.S. Job Posting Data*. Policy Research Working Paper Series 11263, The World Bank.
- LO, A. W. (2017). *Adaptive Markets: Financial Evolution at the Speed of Thought*. Princeton, NJ: Princeton University Press.
- MAURIN, V. and FOUCAULT, T. (2026). *The Price of Exchange Data*. Hec paris research paper, HEC Paris, available at SSRN: <https://ssrn.com/abstract=7072038>.
- MIHET, R. and PHILIPPON, T. (2019). The economics of big data and artificial intelligence. In *Disruptive Innovation in Business and Finance in the Digital World*, Emerald Publishing Limited.
- MITCHELL, T. M. (1997). *Machine Learning*. New York: McGraw-Hill.
- MUKHERJEE, A., PANAYOTOV, G. and SHON, J. (2021). Eye in the sky: Private satellites and government macro data. *Journal of Financial Economics*, **141** (1), 234–254.
- NAGEL, S. (2025). *Seemingly virtuous complexity in return prediction*. Tech. rep., National Bureau of Economic Research.
- NEUMEIER, C., GOZLUKLU, A., HOFFMANN, P., O’NEILL, P. and SUNTHEIM, F. (2023). Banning dark pools: Venue selection and investor trading costs. *Journal of Financial Markets*, **65**, 100831.

- NING, B., LIN, F. H. T. and JAIMUNGAL, S. (2021). Double deep q-learning for optimal execution. *Applied Mathematical Finance*, **28** (4), 361–380.
- NORDHAUS, W. D. (2021). Are we approaching an economic singularity? information technology and the future of economic growth. *American Economic Journal: Macroeconomics*, **13** (1), 299–332.
- OECD (2021). *Artificial Intelligence, Machine Learning and Big Data in Finance: Opportunities, Challenges and Implications for Policy Makers*. Report, OECD Publishing.
- OECD (2024). *Explanatory Memorandum on the Updated OECD Definition of an AI System*. Tech. Rep. 8, OECD.
- PASTOR, L., STAMBAUGH, R. F. and TAYLOR, L. A. (2015). Scale and skill in active management. *Journal of Financial Economics*, **116** (1), 23–45.
- PEDERSEN, L. H. (2009). When everyone runs for the exit. *International Journal of Central Banking*.
- PENG, S., KALLIAMVAKOU, E., CILON, P. and DEMIRER, M. (2023). The impact of ai on developer productivity: Evidence from github copilot.
- PHILIPPON, T. (2015). Has the us finance industry become less efficient? on the theory and measurement of financial intermediation. *American Economic Review*, **105** (4), 1408–38.
- and RESHEF, A. (2012). Wages and Human Capital in the U.S. Finance Industry: 1909–2006. *The Quarterly Journal of Economics*, **127** (4), 1551–1609.
- PORTER, M. F. (1980). An algorithm for suffix stripping. *Program*, **14** (3), 130–137.
- PwC (2025). The fearless future: 2025 global ai jobs barometer.
- SAMUEL, A. L. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of research and development*, **3** (3), 210–229.
- SCHEURER, J., BALESNI, M. and HOBBAHN, M. (2024). *Large Language Models can Strategically Deceive their Users when Put Under Pressure*. Working paper.
- SHENG, J., SUN, Z., YANG, B. and ZHANG, A. L. (2026). Generative AI and asset management. *The Review of Financial Studies*, **Forthcoming**.
- SHKILKO, A. and SOKOLOV, K. (2020). Every cloud has a silver lining: Fast trading, microwave connectivity, and trading costs. *The Journal of Finance*, **75** (6), 2899–2927.
- STIGLITZ, J. E. (2014). Tapping the brakes: Are less active markets safer and better for the economy? *Presented at the Federal Reserve Bank of Atlanta 2014 Financial Markets Conference*.
- SUTTON, R. S. and BARTO, A. G. (1998). *Reinforcement learning: An introduction*, vol. 1. MIT press Cambridge.
- VAN KERVEL, V. and MENKVELD, A. J. (2019). High-frequency trading around large institutional orders. *The Journal of Finance*, **74** (3), 1091–1137.

- VASWANI, A., SHAZEER, N., PARMAR, N., USZKOREIT, J., JONES, L., GOMEZ, A. N., KAISER, L. and POLOSUKHIN, I. (2017). Attention is all you need. *Advances in neural information processing systems*, **30**.
- WAGNER, W. (2011). Systemic liquidation risk and the diversity-diversification trade-off. *The Journal of Finance*, **66** (4), 1141–1175.
- WANG, Q., HUANG, Y. and SINGH, P. (2022). *Algorithms, Artificial Intelligence and Simple Rule Based Pricing*. Working paper.
- WANG, X. and WELLMAN, M. P. (2017). Spoofing the limit order book: An agent-based model. In *AAAI Workshops*.
- WELLER, B. M. (2017). Does algorithmic trading reduce information acquisition? *The Review of Financial Studies*, **31** (6), 2184–2226.
- YADAV, Y. and BROGAARD, J. (2026). The too few to fail traders in modern markets. *Vanderbilt Law Research Paper (forthcoming)*.
- YANG, L. and ZHU, H. (2019). Back-running: Seeking and hiding fundamental information in order flows. *The Review of Financial Studies*, **33** (4), 1484–1533.
- YE, S. (2025). *Task Efficiency and Signaling in the Age of GenAI: Effort Reallocation and Firm Value Effects*. Working paper.
- ZHANG, T., LIU, F., WONG, J., ABBEEL, P. and GONZALEZ, J. E. (2023). The wisdom of hindsight makes language models better instruction followers. In *Proceedings of the 40th International Conference on Machine Learning*, PMLR, *Proceedings of Machine Learning Research*, vol. 202, pp. 41414–41428.
- ZHAO, J. (2023). *Who Benefits (More) from the Growth of Data?* Working paper.
- ZHAO, W. and XU, Z. (2024). On mechanism underlying algorithmic collusion.
- ZHONG, H. (2025). *Optimal Integration: Human, Machine, and Generative AI*. Working paper.
- ZHU, C. (2019). Big data as a governance mechanism. *The Review of Financial Studies*, **32** (5), 2021–2061.
- ZUCKERMAN, G. (2019). *The Man Who Solved the Market: How Jim Simons Launched the Quant Revolution*. New York: Penguin Press.

A Measuring AI Intensity from Web-scraped Job Postings

This appendix describes the data and methodology used to measure the adoption of AI by PTFs, multi-manager hedge funds (“pod shops”), banks, and asset managers, as presented in Figures 4 and 5.

A.1 Data collection

We collect publicly available job postings from the career pages of 35 financial institutions, spanning four categories: 14 PTFs, 5 pod shops, 9 banks, and 7 asset managers (see Table A1 for the full list). For 2025, job postings were scraped directly from each firm’s career website in November 2025. For PTFs, we additionally retrieved archived versions of the same career pages from the Wayback Machine (<https://web.archive.org>) for the years 2022-2024 whenever possible. Not all firms have data available for every year; the coverage for each firm depends on whether the career page of interest was captured and preserved by the Wayback Machine at the relevant date.

Each job posting is compiled into a standardized record containing the job title, required skills, and described responsibilities. We apply a cue-word-based language filter to retain only English-language postings, thereby ensuring consistency in the text classification procedure described below.

A.2 AI intensity classification

To classify job postings by their technological orientation, we closely follow the methodology of [Abis and Veldkamp \(2024\)](#). The text of each posting (concatenating the title, skills,

PTFs	Pod shops	Banks	Asset managers
Citadel Securities*	Balyasny Asset Mgmt.	Banca Santander	Amundi
DaVinci Investments*	Citadel (Hedge Fund)	Barclays	APG Asset Mgmt.
DRW	Millennium Mgmt.	BBVA	BlackRock
Flow Traders	Point72 Asset Mgmt.	BNP Paribas	DWS Group
HRT	Schonfeld Strat. Adv.	Commerzbank	Fidelity Investments
IMC Trading*		Deutsche Bank	State Street Gl. Adv.
Jane Street*		HSBC	Vanguard Group
Mako		ING Group	
Optiver*		Société Générale	
Radix Trading			
TMG			
Virtu Financial*			
Webb Traders			
XTX Markets*			

Table A.1: Sample Firms by category. Firms marked with an asterisk have complete data for 2022-2025. All remaining firms only have data for 2025.

and responsibilities fields) undergoes a normalization pipeline: HTML tags are stripped, the text is lowercased, non-alphanumeric characters are replaced with whitespace, and each word is reduced to its root form using the [Porter \(1980\)](#) stemming algorithm.

We then count keyword matches against four curated dictionaries, each capturing a distinct skill domain: Data Management (91 keywords), General Analysis (47 keywords), Old Technology (47 keywords), and AI (48 keywords). The Data Management dictionary targets infrastructure and data engineering skills (e.g. “data warehousing”, “database administration”, “data engineering”). The General Analysis dictionary captures broad quantitative and financial analytics terms (e.g., “regression analysis”, “portfolio optimization”). The Old Technology dictionary covers classical statistical methods and legacy tools (e.g., “econometrics”, “Monte Carlo simulation”, “Black–Scholes”). The AI dictionary identifies ML and artificial intelligence methods and frameworks (e.g., “machine learning”, “neural network”, “deep learning”, “natural language processing”, “random forest”).

Classification proceeds in two steps. In the first step, we compute the share of Data Management keywords out of all matched keywords. If this share is at least 50%, the

posting is classified as “Data Management”; otherwise, it is classified as “Analysis”. Postings with no keyword matches in any category are labeled “Unclassified”. In the second step, which applies only to postings classified as “Analysis”, we compare the count of AI keywords to the count of Old Technology keywords. If the AI count strictly exceeds the Old Technology count and is positive, the posting is classified as an “AI” posting.

We define the AI job share for a given group of firms as:

$$\text{AI Job Share (\%)} = \frac{\text{Number of postings classified as “AI”}}{\text{Total number of postings}} \times 100$$

where the denominator includes all postings regardless of their first-step classification. This measure captures the prevalence of AI-related hiring activity within each firm category.

B Measuring PTF activity

This appendix describes the data and methodology used to measure PTF activity. Our sample of PTFs consists of firms that are members of either FIA EPTA or APT, two industry associations representing leading principal trading firms.

B.1 Share of trading

We measure the share of trading volume by PTFs on a quarterly basis over the period 2020–2026. We focus on four asset classes: futures, stocks, options, and exchange-traded funds (ETFs).

For each asset class and quarter, we define the PTF share of trading volume as

$$\text{PTF Share (\%)} = \frac{\text{PTF euro volume}}{\text{Total euro volume}} \times 100. \quad (\text{A.1})$$

For futures, we use the trading volume for the on-the-run contract of Euro Stoxx 50 and Bund Futures traded on Eurex. For stocks, we use the trading volume in euro-denominated large-cap stocks on Euronext Amsterdam. For options and ETFs, we do not use trading volume in euros because it can be misleading as a measure of trading activity.⁵⁷ For options, we use the quantity traded in euro-denominated put and call options on Euronext Equities, Indices and Derivatives (EQF). For ETFs, we use the trading quantity in euro-denominated ETFs on Euronext Amsterdam.

As a robustness check, we compare the volume traded by PTFs in stocks as reported in the MiFID II dataset with the volume traded by PTFs calculated using data from Euronext. The two measures coincide, indicating that the reported MiFID II data are of sufficiently good quality to reliably measure the PTF share of trading activity.

B.2 Concentration

We study how concentration within PTFs on a given stock-day varies with the stock-day's volatility. We measure concentration using the Herfindahl–Hirschman Index (HHI) based on PTF trading volume. The HHI reflects how unevenly total PTF volume is distributed across firms: higher values indicate that trading is concentrated among a few PTFs, while lower values indicate a more even distribution.

Let $s_{i,s,t}$ denote the share of total PTF trading volume in euros accounted for by PTF

⁵⁷Option values depend strongly on contract characteristics such as the strike price, and ETF prices can vary substantially across products. As a result, differences in euro volume may reflect price differences rather than actual trading activity.

i in stock s on day t . The HHI is defined as

$$\text{HHI}_{s,t} = \sum_i s_{i,s,t}^2. \quad (\text{A.2})$$

Alternatively, we assess concentration via the euro trading volume of the Top 3 PTFs relative to the euro trading volume of all PTFs. We measure volatility as the stock's intraday price range, defined as:

$$\text{Volatility}_{s,t} = \frac{\text{High}_{s,t} - \text{Low}_{s,t}}{(\text{High}_{s,t} + \text{Low}_{s,t})/2} \times 100, \quad (\text{A.3})$$

where $\text{High}_{s,t}$ and $\text{Low}_{s,t}$ denote the intraday high and low prices of stock s on day t . Volatility is expressed in percentage points.

Our sample includes all trading on Euronext Amsterdam in large-cap (AEX index) and mid-cap (AMX index) stocks that were constituents of these indices on 1 August 2025, and covers all trading days from 2020 to 2026. To mitigate the impact of potential data quality issues, we exclude the one percent of stock-days with the lowest and highest concentration within PTFs. The data for trading in futures covers 2018-2023.

B.3 Inventory ratio

We measure intraday position changes of PTFs using the inventory ratio. For each PTF i on day t , we compute the absolute net traded volume in euros, $|\text{BuyVol}_{i,t} - \text{SellVol}_{i,t}|$, and the total traded volume in euros, $\text{BuyVol}_{i,t} + \text{SellVol}_{i,t}$. We then aggregate these quantities across all PTFs and all trading days within a quarter and define the quarterly inventory ratio as

$$\text{Inventory Ratio}_q = \frac{\sum_{t \in q} \sum_{i \in \text{PTFs}} |\text{BuyVol}_{i,t} - \text{SellVol}_{i,t}|}{\sum_{t \in q} \sum_{i \in \text{PTFs}} (\text{BuyVol}_{i,t} + \text{SellVol}_{i,t})} \times 100. \quad (\text{A.4})$$

A higher inventory ratio indicates that PTFs, in aggregate, accumulate larger overnight positions relative to their total trading activity.

Our sample includes all transactions by PTFs in Dutch large-cap stocks (AEX constituents) across all trading venues reported in the MiFID II dataset over the period 2020–2026. We exclude off-exchange transactions, because these are not consistently reported.

As a robustness check, we compare the number of transactions and the value traded by PTFs in shares in our Euronext dataset with those reported in the MiFID II dataset. These quantities coincide, which confirms that the reported MiFID II data are of sufficiently good quality to calculate the inventory ratio.

Acknowledgements

The views expressed in this paper are the authors' and do not necessarily reflect those of the Autoriteit Financiële Markten (AFM), the European Central Bank (ECB) and/or the Eurosystem.

Álvaro Cartea

University of Oxford, Oxford, United Kingdom; email: alvaro.cartea@maths.ox.ac.uk

Jean-Edouard Colliard

HEC Paris, Paris, France; email: colliard@hec.fr

Thierry Foucault

HEC Paris, Paris, France; email: foucault@hec.fr

Peter Hoffmann

European Central Bank, Frankfurt am Main, Germany; email: peter.hoffmann@ecb.europa.eu

Rob Graumans

Autoriteit Financiële Markten, Amsterdam, The Netherlands; email: robert.graumans@afm.nl

Jean-David Sigaux

European Central Bank, Frankfurt am Main, Germany; email: jean-david.sigaux@ecb.europa.eu

© European Central Bank, 2026

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from ECB's website or from RePEc: Research Papers in Economics. Information on all of the papers published in the ECB Discussion Paper Series can be found on the ECB's website.

PDF

ISBN 978-92-899-8025-8

ISSN 2811-8820

doi: 10.2866/0635516

QB-01-26-228-EN-N