21

Artificial Intelligence in Financial Services

New Risks and the Need for More Regulation?

*Matthias Paul**

I. INTRODUCTION

The financial services industry has been at the forefront of digitization and big data usage for decades. For the most part, data processing has been automized by information management systems. Not surprisingly, Artificial Intelligence (AI) applications, capturing the more intelligent ways of handling financial activities and information, have increasingly found their way into the financial services industry over the last years; from algorithmic trading, smart automized credit decisions, intelligent credit card fraud detection processes, personalized banking applications and even into areas like so-called robo-advisory services and quantitative investment and asset management more recently.[1]

The financial industry has also been one of the most regulated industries in the world. In particular, since the collapse of Lehmann Brothers in 2008, leading into one of the most severe financial crises in history, regulation efforts of all kinds of finance-related activities and financial organizations as a whole by the different regulators around the world have significantly increased. In general, most regulations relating to the financial industry, in particular those put into place after the financial crisis in 2008, have focused on areas like safeguarding the financial institutions themselves, safeguarding the customers of financial institutions, and making sure the institutions comply with general laws overall and on a global scale, given the truly global nature of the financial industry.

More recently, authors have argued that with the emergence of AI-based applications in the financial industry, new kinds of risks have emerged that require additional regulations.[2] They have pointed for instance to increased data processing risk, cybersecurity risks, additional challenges to financial stability, and even to general ethical risks stemming from AI in financial services. Some regulators like the Monetary Authority of Singapore (MAS) have proposed an AI

* I want to thank Silja Voeneky for many insightful discussions of the topic of AI, for sharing and exchanging many ideas, and also for her comments on an earlier draft version of this chapter.

[1] See C Chan and others, 'Artificial Intelligence Applications in Financial Services – Asset Management, Banking and Insurance' (*Oliver Wyman Research Report*, 2019), www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2019/dec/ai-app-in-fs.pdf; for an overview, T Boobier, *AI and the Future of Banking* (2020), and also T Guida, *Big Data and Machine Learning in Quantitative Investment* (2019) (hereafter Guida, *Big Data*) for more recent developments in quantitative investment.

[2] See D Zetzsche and others, 'Artificial Intelligence in Finance – Putting the Human in the Loop' (2020) University of Hong Kong Faculty of Law Research Paper No. 2020/006 (hereafter Zetzsche and others, 'Artificial Intelligence in Finance'), Guida, Big Data (n 1), or the recent regulatory proposals from the Monetary Authority of Singapore (2019).

governance framework for financial institutions.[3] The EU also explored this topic and published a report on big data risks for the financial sector, including AI, stressing appropriate control and monitoring mechanisms.[4] Scholars have developed this topic further by adopting so-called personal responsibility frameworks to regulate any new emerging AI-based applications in the financial industry.[5] In its recent draft regulation, the EU has presented a general risk-based regulatory approach of AI which regulates and even prohibits certain so-called high risk AI system; and some of them can supposedly also be found in the financial industry.[6]

This chapter will explore this entire topic of AI in the financial industry (which will also be referred to as robo-finance) further. One focus of the article will be on whether AI in the financial industry gives rise to new kinds of risks or merely increases existing risks already present in the industry. Further, the article will review one prominent general regulatory approach many scholars and regulators have put forward to limit or mitigate these alleged new risks, namely the so-called (personal) responsibility frameworks. In the final section of this chapter, a different proposal will be presented on how and to what extent best to regulate robo-finance, which will take up key elements and concepts from the recent Draft EU AIA.[7] To lay the groundwork for the discussion of these topics, the nature of AI, in particular as a general-purpose technology, will be explored first. In addition, an overview of the current state of AI applications in financial services will be given, and the different regulatory layers or focus areas for regulations that are present in the financial industry today will be presented. Based on these introductory discussions, the main topics of the chapter can then be spelled out.

II. AI AS A NEW GENERAL PURPOSE TECHNOLOGY

Electricity is a technology or technology domain which came into life more than 150 years ago, and it still drives a lot of change today. It comprises different concepts like electrical current, electrical charge, electric field, electromagnetics etc. which have led to many different application areas in their own right; from the light ball to electrical telegraphs or to electric engines, to mention only a few. It is fair to say that electricity as a technology field or domain has revolutionized the world in many ways, and it still does. And it has changed and transformed whole industries as it transforms the automotive industry with the transition from combustion engines to electric cars.

Given its wide range of underlying concepts with multiple specific application areas of their own right, several authors have referred to electricity as a general-purpose technology (GPT).[8]

[3] See the so-called IAC (Individual Accountability) guidelines by the Monetary Authority of Singapore, MAS, 'Guidelines on Individual Accountability and Conduct' (MAS, 10 September 2020) www.mas.gov.sg/-/media/MAS/MPI/Guidelines/Guidelines-on-Individual-Accountability-and-Conduct.pdf.

[4] See the joint report of the European Supervisory Authorities EBA, ESMA, EIOPA on the use of big data, including AI, by financial institutions, December 2016, JC/2016/86.

[5] See for instance Zetzsche and others, 'Artificial Intelligence in Finance' (n 2).

[6] The EU published the General Regulation on a European Approach for Artificial Intelligence in 2021, which regulates the financial industry in some areas of AI applications as well, European Commission, Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts COM (2021) 206 final (hereafter Draft EU AIA). Since this draft was published after the writing of this article, its impact will and can be discussed only to a smaller extent in this paper.

[7] See Zetzsche and others, 'Artificial Intelligence in Finance' (n 2).

[8] See E Brynjolfsson and A McAfee, 'The Business of Artificial Intelligence' (*Harvard Business Review*, 18 July 2017) https://hbr.org/2017/07/the-business-of-artificial-intelligence 3 (hereafter Brynjolfsson and McAfee, 'The Business of Artificial Intelligence'), see also the interview with *Andre Ng* in M Ford, *Architects of Intelligence – The Truth about AI from the People Building It* (2018) 190 (hereafter Ford, Architects of Intelligence)

https://doi.org/10.1017/9781009207898.028 Published online by Cambridge University Press

What is characteristic of GPTs is that there exists a wide range of different use cases in different industries, thus GPTs are not use-case-specific or industry-specific technologies but have applications across industries and across many types of use cases. Other examples of GPTs which scholars have identified are the wheel, printing, the steam engine, and the combustion engine, to mention a few.[9] As such, GPTs are seen as technologies that can have a wide-ranging impact on an entire economy and, therefore, have the potential to drastically alter societies through their impact on economic and social structures.[10]

Several authors have claimed or argued in recent years that AI can or should also be considered a GPT, 'the most important one of our era' in fact.[11] Or as *Andrew Ng* says: 'AI will transform multiple industries'.[12] AI's impact on societies as a whole is seen as significant, for instance, changing the way we work, the way we interact with each other and with artificial devices, how we drive, how wars might be conducted, etc. Further, like in the case of electricity, there are many different concepts underlying AI today, from classical logic or rule-based AI to machine learning and deep learning based AI, as employed so successfully today in many areas. Some hybrid applications combine both concepts.[13] These concepts have allowed for many new types of AI applications, similar to the case of electricity, where different concepts have been merged together as well.

In fact, because the use cases for AI technologies are so enormous today, companies like Facebook have created their internal AI labs or what they have called their 'AI workshop' where many different applications of AI technologies, in particular machine learning applications, get explored and developed.[14] The underlying assumption of such companies is that AI can be applied to so many different areas and tasks that they need to find good ways to leverage their technological expertise in all such different areas.

Clearly, AI is still in its early stages of technological development, with fewer implementations in widespread operation than in the case of electricity. But there have been language and speech processing applications, visual recognition applications like face recognition in smartphones, photo optimization algorithms in digital cameras, many kinds of big data analytics applications, etc. AI technologies have also changed the interface between humans and machines, some turn machines into helpful assistants, others allow for intelligent ways of automating processes and so on. The applications of AI are already widespread today, and we seem to be just at the beginning of a long journey of bringing more applications to life.[15]

In the following, we will look at the financial industry as one major application area for AI as a general-purpose technology. The financial industry is interesting in so far as it is heavily regulated on the one hand, but also highly digitalized and technologically advanced on the other hand, with many kinds of AI use cases operational already today.

[9] See R Lipsey and IC Kenneth, *Economic Transformations: General Purpose Technologies and Long Term Economic Growth* (2005) (hereafter Lipsey and Kenneth, Economic Transformations) for a broader discussion of different GPTs and their role for economic development and the transformation of societies as a whole.

[10] Besides Lipsey and Kenneth, Economic Transformations (n 9) see also TF Bresnahan and M Trajtenberg, 'General Purpose Technologies "Engines of Growth"?' (1995) 65(1) *Journal of Econometrics* 83 for another interesting article on the wider topic of the role and impact of GPTs.

[11] See Brynjolfsson and McAfee, The Business of Artificial Intelligence (n 8) 4.

[12] See the interview with *Andrew Ng* in Ford, *Architects of Intelligence* (n 8) 190 *et seq*.

[13] See ibid.

[14] See J Candela and S Berinato, *Artificial Intelligence: Insights You Need from Harvard Business Review* (2019)

[15] It is worth noting that today it is not entirely clear which direction AI as a technology will go over the next years. Despite the enormous success of machine learning as an AI concept or paradigm, several authors have pointed to its limitations – see for example the interview with *Barbara Grosz* in Ford, *Architects of Intelligence* (n 8) 333–356.

III. ROBO-FINANCE: FROM AUTOMATION TO THE WIDE-SPREAD USE OF AI APPLICATIONS

The financial industry has been one of the most data-intensive and digitized industries for decades. In 1973, SWIFT was founded and launched, the so-called Society for Worldwide Interbank Financial Telecommunication, bringing together 239 banks from 15 countries worldwide with the aim of handling the communication of cross-border payments. The main components of the original system included a computer-based messaging platform and a standard message system.[16] This system disrupted the manual processes of the past, and today more than 11,000 financial institutions from more than 200 countries are connected through SWIFT's financial global technology infrastructure. Nasdaq, to give another example, the world's first electronic stock exchange, began its operations even earlier, in 1971, leading the way to fully digitized exchanges for the trading of any kinds of financial securities, which are the standard and norm today. And real-time financial market data and news, probably the first big data sets used in history, were made available in the early 1980s by companies such as Thomson Reuters and Bloomberg through their market data feeds and terminal services.[17]

In the years to follow, the financial industry has been at the forefront of leveraging information (management) systems to manage and process the vast amounts of data and information available.[18] In fact, today, many financial institutions resemble technology companies more than traditional banking houses, and it is no surprise that companies like Paypal or, more recently, many new fintech players were able to further transform this traditional industry by leveraging new technologies like the Internet or mobile services, platforms, and infrastructures.[19]

This development of digitizing financial information and financial transactions has made the automation of data handling and processing, not just a possibility but rather a necessity to maintain and defend one's competitiveness and to deal with and manage the various kinds of risks inherent in the financial industry. The execution of payments within the international banking system or the execution of buying or selling orders on the exchanges can be fully automized today based on simple parameters (such as dates and amounts, or stop-loss orders to manage risk, etc.). Clearly, these ways of automizing financial transactions and processes are in no way intelligent, nevertheless they have helped the investment banks and other actors in the financial industry tremendously to increase process speed, accuracy, and also improve

[16] See www.swift.com/about-us/history for more details on the introduction of the SWIFT system.

[17] Commonly big data sets are defined by the so-called 4 Vs: volume (the amount of data), velocity (the speed in which new data get created or are generated), variety (different kinds of data types from different data sources, in particular, often a mix of structured and unstructured data), and veracity (discrepancies, errors, and gaps in data sets). Typical market data feeds in the financial industry fulfill at least three of these criteria, namely volume, velocity and veracity, as the data feeds deliver fairly structured data sets. This might change in the future when data feeds might also include other kinds of data such as press releases or social media posts as it is the case already with so-called sentiment feeds including sentiment data. See B Marr, *Big Data: Using SMART Big Data, Analytics and Metrics to Make Better Decisions and Improve Performance* (2015) for a more general introduction in the area of big data, and Guida, *Big Data* (n 1) for more insights into big data in areas of financial information.

[18] In broader terms an information (management) system is simply defined as a set of interrelated components consisting of an application system, and an interface for human interaction to define the tasks for the system and retrieve information. The application system consists of hardware, software, data, and a network connection. For more details see K Laudon and J Laudon, *Management Information Systems: Managing the Digial Firm* (15th ed. 2018).

[19] Examples in the payment sector are WeChat Pay, Alipay or Apple Pay, new competitors to the established credit card payment services. In fact, today many big tech companies are moving into financial services with their own finance applications, often in areas like payments, as Apple Pay or Google Pay.

risk management.[20] Therefore, it is no surprise that many actors in the industry have constantly searched and tried to develop more sophisticated processes, which has opened the doors for AI applications in the financial industry.

Today, there is a wide range of AI applications present in the financial industry of which the following are just key application areas with multiple kinds of use cases:[21]

(1) *Customer Related Processes:*
   a. new ways of segmenting customers based on the use of so-called cluster algorithms or analyses,[22]
   b. personalized banking services and offers based, for instance, on profiling algorithms,[23]
   c. robo-advisory services replacing human financial advisory with machines,[24]
   d. intelligent chatbots advising or providing information to clients in different areas of their financial decision making.[25]

(2) *Operations and Risk Management:*
   a. underwriting automation in credit decisions and algorithmic credit scoring,[26]
   b. automized stress testing.

(3) *Trading and Investment Management:*
   a. algorithmic trading – from simple rule-based AI to more sophisticated machine learning based algorithms,[27]
   b. automatic portfolio rebalancing in asset management adjusting the portfolio to the predefined asset allocation scheme based on simple rule-based algorithms,

[20] They operate more like a thermostat for a heating system, setting thresholds for certain actions to take place, like selling a stock position based on a predefined stop-loss order. The system will automatically initiate the transaction, but it is solely based on predefined parameters.

[21] Zetzsche and others, 'Artificial Intelligence in Finance' (n 2) present a similar classification of the AI application present today in the financial industry. See also T Boobier, *AI and the Future of Banking* (2020). For discussion of several of the application areas discussed here, as well as a recent leadership paper by the consultancy firms Oliver Wyman, Marsch, BCLP and Hermes, C Chan, and others, 'Artificial Intelligence Applications in Financial Services – Asset Management, Banking and Insurance' (Oliver Wyman Research Report, 2019), www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2019/dec/ai-app-in-fs.pdf.

[22] See M Hassan and M Tabasum, 'Customer Profiling and Segmentation in Retail Banks Using Data Mining Techniques' (2018) 9(4) *International Journal of Advanced Research in Computer Science*.

[23] See R Ragotan, 'AI Has Changed the Way Banks Interact with Their Customers' (*Fintech News*, 5 February 2020) www.fintechnews.org/ai-has-changed-the-way-banks-interact-with-their-customers. For a discussion of some of the applications and service providers.

[24] So-called robo-advisors or advisory solutions like Betterment, Wealthfront, and Vanguard Digital, to give a few examples from the more advanced US robo advisory market, have in recent years been launched in competition with traditional human banking or financial advisors. These solutions automize and digitalize the advisory process in wealth management and private banking, thereby lowering the asset under management threshold for private investors for accessing high quality advisory solutions. Although some of the new players have also automized the asset management process itself, the primary focus of these solutions is enhancing the advisory process by replacing the human banking advisor with a machine or AI-based interface. In this regard they are classified here under customer related solutions and not under AI-based trading and portfolio management solutions as done by Zetzsche and others, 'Artificial Intelligence in Finance' (n 2) which is rather misleading.

[25] As pointed out in a study by the consulting firm McKinsey & Company (2018), data analytics applications often using AI techniques are most widespread in sales and marketing areas of businesses, that is, areas which try to generate and develop new customer relationships and transaction.

[26] See an interesting article by N Aggarval, 'The Norms of Algorithmic Credit Scoring' (2021) 8(2) *The Cambridge Law Journal* 42 on the norms of algorithmic credit scoring.

[27] See M Lewis, *Flash Boys: A Wall Street Revolt* (2014) (hereafter Lewis, *Flash Boys*) or S Patterson, *Dark Pools: The Rise of the Machine Traders and the Rigging of the U.S. Stock Market* (2012) (hereafter Patterson, *Dark Pools*) for good non-expert introductions into this area, for a more systematic and scientific account see R Kissell, *Algorithmic Trading Methods: Applications Using Advanced Statistics, Optimization, and Machine Learning Techniques* (2021).

c. big data and machine learning–based (assisted or fully automized) asset management.[28]

(4) *Payment Processes:*
fraud detection algorithms in credit card payments using big data analytics and learning algorithms.[29]

(5) *Data Security and Cybersecurity:*[30]
a. data security – algorithms protecting the data from inside a financial institution,
b. cybersecurity – algorithms protecting the data from outside attacks.[31]

(6) *General Regulatory Services and Compliance Requirements:*[32]
a. Anti Money Laundering (AML) automation and protection algorithms helping to identify politically exposed people (so-called Peps) or criminals involved in certain financial transactions,
b. detection of compliance breaches in case of insider trading etc.

As shown here, AI is already employed today in many areas of the financial industry, and new applications are emerging every day. The question is whether additional or increased risks stem from these applications, which might require additional regulations, as argued by some authors.[33] This line of argument will be reviewed in more detail in the following sections. But first, it is important to understand from a high-level perspective the main areas and layers of regulations in the financial industry today.

IV. A SHORT OVERVIEW OF REGULATION IN THE FINANCIAL SERVICES INDUSTRY

The financial services industry is probably one of the most regulated industries. Regulation of trading practices for instance dates back to the seventeenth century when in 1610 in Holland, some first forms of *short selling* became prohibited.[34] At the same time, the first central banks were created, such as the Swedish Riksbank in 1668, to regulate payment transactions on a national level and establish national currencies by issuing banking notes. Some of the early

[28] See Guida, *Big Data* (n 1) on a recent collection of articles on this new emerging and developing field. So far, AI applications and tools have mainly been used in assisting fund managers in the asset allocation process, but it is possible that there will be fully AI-based fund management in the future. Some authors like E Syrotyuk, 'State of Machine Learning Applications in Investment Management' in T Guida (ed), *Big Data and Machine Learning in Quantitative Investment* (2019) seem to be more sceptical in regard to fully automized asset management because of the more erratic nature of financial markets.

[29] Companies like Teradata (teradata.com) and Datavisor (datavisor.com) provide AI-based financial fraud detection solutions. Datavisor, for instance, claims that their solution can detect 30% more frauds with 90% accuracy. Their solutions are mainly based on machine learning algorithms according to own research.

[30] See A Bouveret, 'Cyber Risk for the Financial Sector: A Framework for Quantitative Assessment' (2018) International Monetary Fund Working Paper 18/143 for a thorough overview and analysis of cyber security risk in the financial industry by sectors and countries/regions.

[31] See J Li, 'Cyber Security Meets Artificial Intelligence: a Survey' (2018) 19 Frontiers of Information Technology & Electronic Engineering for a more detailed analysis of the potential of using AI systems in preventing or reducing cyberattacks. The article also highlights the fact that AI systems might be used in facilitating cyber security attacks, as will be discussed also later in the article.

[32] A new sector has emerged in recent years often referred to as RegTech – see Zetzsche and others, 'Artificial Intelligence in Finance' (n 2) – using technology to help financial institutions to comply with the various regulatory requirements. Quite a few regtech solutions have increasingly made use of AI technologies; for a good overview see 'AI in RegTech: a quiet upheaval' (*Chartis*, 2018) www.ibm.com/downloads/cas/NAJXEKE6.

[33] See Zetzsche and others, 'Artificial Intelligence in Finance' (n 2) as a recent example.

[34] See AM Fleckner, 'Regulating Trading Practices' in N Moloney, E Ferran, and J Payne (eds), *The Oxford Handbook of Financial Regulation* (2015) 597 (hereafter Fleckner, 'Regulating Trading Practices').

regulation was 'private self-regulation', in other words, bottom up norm creation,[35] as in the case of regulatory practices around many of the emerging exchanges, but some regulation was already at these early times government- or state-driven (top down) as in the case of the establishment of central banks and their key role in establishing standardized payments practices based on backed up currencies.[36]

Today the financial industry is heavily regulated by national or supranational bodies, for instance, by the ESMA[37] in the EU or by the SEC[38] in the US in regard to activities on the different financial markets. Some of the regulations are financial-industry-specific, others are general regulations that severely impact the financial industry. Overall, the different types or layers of regulations in the financial industry can be classified by their underlying aims, namely: (i) regulations meant to safeguard overall financial stability, (ii) regulations for the protection of consumers of financial services, and (iii) regulations that are meant to make sure financial services can operate in a challenging and diverse international environment with sometimes conflicting rules and principles.[39]

The following overview tries to capture the main regulation areas or layers and their specific purpose or aim as they are present in the financial industry today. Some of the layers directly link up to the categories just mentioned, some are cutting across the different categories, and some are also mirroring the classification of the previous Section of AI-impacted application domains in the financial industry:

(1) Equity and liquidity requirements for banks and financial institutions to adhere to minimum capital ratios and liquid asset holdings to prevent financial stress, improve risk management, and promote transparency. Examples are the Basel I, Basel II, Basel III regulations which are global voluntary regulatory frameworks adhered to by most financial institutions today;[40]
(2) Infrastructure regulations, many still in the proposal stage, to improve financial services firms' operational resilience (in case of major disasters, for instance), and their responses to cyberattacks;[41]
(3) Pre- and post-trading regulations to strengthen investor protection and improve the functioning of financial markets, making them more efficient, resilient, and transparent like banning certain trading practices or making kickbacks by product issuers transparent. The MiFID I and II regulations in the EU are examples of such kinds of regulations;[42]

[35] For the different meanings of the notion 'regulation', cf. T Schmidt and S Voeneky, Chapter 8, in this volume.
[36] For a thorough analysis of regulation of trading practices in the financial industry discussing both sides of regulation see the article by Fleckner, 'Regulating Trading Practices' (n 34).
[37] European Securities and Market Authority (ESMA), the EU's securities market regulator located in Paris, created in 2011 and replacing the Committee of European Securities Regulators (CESR).
[38] Securities and Exchange Commission (SEC), the independent agency of the US federal government, created in the early 1930s following the stock market crash in 1929.
[39] See the KPMG report, 'EU Financial Services Regulation – A New Agenda Demands a New Approach' www.kpmg.com/regulatorychallenges, for giving a good overview on the various regulatory perspectives of regulation of financial services in the EU.
[40] See for a concise and high-level summary of the Basel I–III regulations the article 'History of the Basel Committee' (*BIS*) bis.org/bcbs/history.htm.
[41] See the European Commission, 'Proposal for the Regulation of the European Parliament and of the Council on Digital Operational Resilience for the Financial Sectors' COM (2020) 595.
[42] See M Comana, D Previtali, and L Bellardini, *The MiFID II Framework: How the Standards Are Reshaping the Investment Industry* (2019) for a detailed analysis of the MiFID II regulations including a comparison with the MiFID I rules.

(4) Payment services regulations, like the PSD II directive (2015) in the EU, with the aim of creating more integrated payments markets, making payments safer and more secure and also protecting consumers, for instance, from the financial damage resulting from fraudulent credit card payments;
(5) Various kinds of compliance regulations, for instance anti-money-laundering or terrorist financing regulations etc., to ensure that financial institutions obey the treaties and laws and do not enter any illegal transactions or practices, also regarding cross-border transactions;[43]
(6) *General data privacy protections* like the GDPR[44] in the EU, which is highly relevant as financial transactions involve much sensitive personal data.

As we can see, there are no specific AI regulations of financial services – although many of the regulations will also impact AI-based financial services. In fact, there are even very few regulations regarding the underlying technologies in financial services, but most of them focus on the use cases or financial activities, processes, and on the outcomes themselves. Yet, recently some scholars and some regulators have argued that there might be new risks stemming from AI applications and technologies in financial services which require additional regulation. In the following, we will look at some of the alleged risks as pointed out by scholars in the field and explore to what extent they might be covered by the above regulations already or whether there is a need for new regulatory frameworks.

V. NEW RISK CATEGORIES IN ROBO-FINANCE STEMMING FROM AI?

Dirk Zetzsche and others, in their recent paper, have identified the following four risk categories or risk areas allegedly related to AI applications in the financial industry:

(1) Data risks
(2) Cybersecurity risks
(3) Financial stability risks
(4) Ethical risks.[45]

Although I agree with the authors that all these kinds of risks are related to AI applications in financial services, it appears that these risks already existed before the emergence of robo-finance, given the advanced stage of the industry in terms of digitization and data dependency and usage. In fact, some of these risks might even be reduced or vanish when AI comes into place. Let us look at the different risk areas one by one.

Firstly, starting with the data risks of AI applications, *Zetzsche* and others bring up the following more specific arguments: (i) Because the data quality might be poor, there can be deficiencies stemming from AI applications. As a matter of fact, data quality has often been poor in many parts of the financial industry, for instance, outages at the data centers of the exchanges or of the market data providers leading to the misstating of prices of securities, which can have

[43] The laws and regulations around data privacy protections can also be seen as falling into this category but it has been listed here separately given its recent prominence.

[44] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), OJ 2016 L 119/1.

[45] This classification mirrors or reflects to some extent on the classification by the French prudential regulatory authority within the bank of France. It has recently put forward the following four risk categories as allegedly stemming from AI applications in the financial industry: (1) data processing risk, (2) cybersecurity risk, (3) challenges to financial stability, (4) player's dependency and change in power relationships in the financial market. See 'Artificial Intelligence: Challenges for the Financial Sector' (*ACPR*, December 2018), acpr.banque-france.fr/sites/default/files/medias/documents/2018_12_20_intelligence_artificielle_en.pdf.

https://doi.org/10.1017/9781009207898.028 Published online by Cambridge University Press

negative effects on investors' decisions and the markets overall. AI could actually be used to deal with the data issues in terms of detecting and even resolving them.[46] (ii) Besides, they argue that data used for AI analyses might suffer from biases, for instance, relating to what they call 'oversight' in a financial organization. Again, biases have already been influencing decision making in the financial industry even before the emergence of AI applications, maybe not in the form of what they call data-biases, but biases residing more generally in human decision, for instance in the making of credit decisions or consumer lending.[47] AI application might free us from certain biases by providing a more neutral stance if programmed accordingly or at least be sensitive to such kinds of biases. (iii) And it is claimed that AI interdependency can lead to what they call 'herding', for instance all systems selling securities triggered by certain market events, which can lead to what has been referred to as 'flash crashes'.[48] Again 'herding' behavior has existed in the financial markets for a long time, and whether the emergence of AI in electronic trading systems has been the cause of what has been called 'flash crashes' seems rather questionable. Simple rule-based algorithms, which today by industry experts would rather not be classified as AI systems can give rise to such behavior in contrast to more sophisticated systems trained on historical data relating to such events.

Secondly, let us look at cybersecurity risks. Obviously, they have also existed before the arrival of AI, with most attacks initiated and conducted by human individuals directly or by simple processes, methods, or algorithms. Examples are emails carrying malware that, after it has installed itself on someone's computer, can silently send all sorts of confidential data from the computer or computer network to the attacker; a similar case of phishing attacks through links to websites – for instance, of online banks that mimic the log-in pages one is familiar with; or finally, simply the reuse of a user's credentials which the attackers have somehow got hold of – for instance, by one of the already mentioned measures or by simply spying on people in combination with our carelessness in setting passwords. That 'algorithms can be manipulated in an effort to transfer wealth' has nothing to do with the presence of AI systems because this could be done already before such systems were in place and it currently happens every day in many different ways within traditional information system environments.[49] It rather seems plausible that AI might provide some help in identifying and preventing cybersecurity attacks.[50]

[46] For instance, the construction of error correction codes can be used in handling issues in data transmission through noisy channels as for instance happens sometimes in the case of market data feeds. More recently AI techniques have been used in optimizing the design of error correction codes, see for instance L Huang and others, 'AI Coding: Learning to Construct Error Correction Code' (2019) 20(10) *IEEE Transactions on Communications* (hereafter Huang and others, 'AI Coding').

[47] In their interesting paper W Dobbie and others, 'Measuring Bias in Consumer Lending' (2021) *The Review of Economics Studies* https://doi.org/10.1093/restud/rdaa078, tried to measure the amount of bias in consumer lending decision. What they found is that in traditional non-AI-based lending decisions there is a significant bias against immigrant and older loan applicants.

[48] There has been a lot of debate around the so-called flash crash which happened on May 6, 2010, when the Dow Jones Index lost about a tenth of its value in just 36 minutes – see for instance A Kirilenko and others, 'The Flash Crash: The Impact on High Frequency Trading on an Electronic Market' (2017) 72 *The Journal of Finance* 967. In his recent article D Busch, 'MiFID II: Regulating High Frequency Trading, other Forms of Algorithmic Trading and Direct Electronic Market Access (2017) 2 *Law and Financial Markets Review* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3068104 (hereafter Busch, 'MiFID II') looks at how, by the MiFID II regulation, such flash crashes are meant to be banned by ruling out a technique of market manipulation referred to as 'spoofing'. This technique was allegedly used by a British stock market trader in 2010 when he tricked the market into believing that the prices were about to fall by placing huge amounts of sell orders which were later cancelled by him by his specially developed algorithms.

[49] See Zetzsche and others, 'Artificial Intelligence in Finance' (n 2) 21.

[50] See the discussion in Section II (5).

https://doi.org/10.1017/9781009207898.028 Published online by Cambridge University Press

Many services are offered today in this regard, and this seems to be one of the areas where the financial industry could benefit from employing AI-based solutions and thereby reduce potentially harmful cybersecurity risks.

Thirdly, when *Zetzsche* and others talk about financial stability risks, it is fairly unclear what they have in mind since they mention almost all areas of AI applications in financial services – as laid out above – from consumer facing and supporting applications, to trading and portfolio management systems, to general regulatory and compliance systems. Overall, their main concern here seems to be the emergence of 'additional third-party dependencies' to AI technology providers up to the 'level of oligopoly or monopoly'. Since many of these third-party technology providers are unregulated today, as they point out, and it might even be hard to regulate them as 'AI-related expertise beyond those developing the AI is limited', there appears to be a major risk. As they say, 'these third-party dependencies [. . .] could have systemic effects'.[51]

What can be said against this last part of their arguments is that many of the actors at the forefront of using AI in financial services today develop their applications inhouse, like the dominant hedge funds or algo trading shops set up by IT specialists.[52] AI-based technology has become such a core asset these days and a competitive factor of financial services that many financial institutions resemble IT companies more and more today to keep all that knowledge inhouse or with high IT expertise inside the organizations to manage their IT service providers or outsourcing partners – far from being entirely dependent or in the hands of monopolistic or oligopolistic structured IT providers.[53] Hence, their worry about systemic effects stemming from such dependencies seems to be overstated, at least in certain critical banking areas. Moreover, in many instances there are quite a few technology providers that offer similar services to the financial industries, for instance market data providers which increasingly have started to use AI technologies to organize and manage the quality of their market data feeds.[54] Financial institutions, at least in critical areas like trading, often make use of different providers at the same time, which also helps them to reduce their third-party-dependencies. Furthermore, with higher education AI or machine learning programs popping up at many educational institutions around the world, new graduates are also increasingly being educated and trained in these key areas. Thus, knowledge is building up quickly and will also be more widely available, reducing the fear of there being a kind of 'mystery science' only a few people have access to and can take advantage of.

Finally, let us focus on what *Zetzsche* and others refer to as new 'ethical risks' stemming from AI applications in financial services. The starting point of their argument is that algorithms do not feel anything, nor do they have values which the authors equate with a lack of ethical foundation in AI-decision making. For instance, they point out that such 'unethical' AI systems

[51] See Zetzsche and others, 'Artificial Intelligence in Finance' (n 2) 21.

[52] See Lewis, *Flash Boys* (n 27) and Patterson, *Dark Pools* (n 27) for a vivid description of the individuals, often IT experts or nerds, in setting up high frequency trading firms or the respective trading units at major banks. Also, major hedge funds with a quantitative focus like Renaissance Technologies, which had 133 billion USD under management as of November 2020, have a strong focus on developing their own mathematical models and algorithms.

[53] See the article by T York, 'Banks Becoming Technology Companies, Technology Companies Becoming Banks' (*San Diego Business Journal*, 30 September 2019) www.sdbj.com/news/2019/sep/30/banks-becoming-technology-companies-technology-com/; see also the recent BCG publication on this topic, J Erlebach and others, 'The Sun Is Setting on Traditional Banking' (*BCG*, 24 November 2020) www.bcg.com/publications/2020/bionic-banking-may-be-the-future-of-banking.

[54] For instance, market data providers like Bloomberg and Thomson Reuters have started to use AI methods and techniques, helping to digest larger data sets including unstructured data like texts from different sources, thereby delivering new kinds of analytics such as so-called sentiment analysis or feeds, trying to identify the sentiments in certain markets or regarding certain securities.

might nudge people to purchase unsuitable financial products, which might further be facilitated by the fact humans would easily develop a higher level of trust in the AI-based systems because with them, human–machine communication can nowadays be quite sophisticated. What this ultimately can and will lead to is reputational risk for the financial institution employing such systems, for example, when people are driven to make the wrong financial decisions and this becomes public or will be reported in the media or brought up to the courts.

There are quite a few problems with this line of reasoning as there are many financial institutions that do not have much direct interaction with human consumers, like mutual funds, hedge funds, credit card companies, etc. Besides, it is also conceivable that AI systems can have an ethical foundation, for instance thinking of utilitarian approaches which are less focused on being able to feel anything or have values. Such aligned AI systems might still be able to calculate the best outcome for society as a whole. But the main counter argument seems to be that the financial industry has not been a role model for ethical behavior to start with. Quite to the contrary, over many decades, financial institutions have been prone to all kinds of ethical misconduct. Just to give a few examples: (i) consumers have been pushed by financial advisors, humans with feelings and values, employed by financial institutions, to buy financial products which were often not suitable or beneficial to them, yet by selling them, the advisors were able to boost their commission payments, and the financial institutions could thereby boost their profits;[55] (ii) insider trading has happened frequently,[56] (iii) market manipulation has occurred, for instance in the case of the Libor scandal, and many other examples in different areas of the financial industry.[57] Thus, it is far from clear why AI-based systems and processes would make the industry less ethical than it has been in the past. In fact, the case could be made that AI-based systems and processes might allow society to create and control financial institutions and make them less driven by greed but more by higher motives to bring benefits to consumers and install fairness within the systems.

But this line of reasoning might sound overly naïve, given how many actors in the financial industry have successfully used technology over the last decades to their advantage, and to the disadvantage of other actors. One example has been the area of high frequency trading and the so-called dark pools where 'fast moving robot trading machines were front-running long term investors on exchanges'.[58] Darks pools are markets established by the financial actors themselves

[55] In Germany for instance the advisory services offered mostly by banks have been reviewed frequently by consumer protection agencies and independent bodies, and over many years the findings have been very disappointing with many banks not even fulfilling basic standards and requirements – see the magazine *Finanztest* 2/2016. In particular, elderly people have been frequently 'ripped off' and have been referred to internally as 'AD's (alt (old) and dumm (stupid)), to whom the advisors could sell products not suitable to the financial situation of the elderly or asking them to re-allocate their portfolio frequently mainly with the aim of generating extra commission fees on the triggered transaction, thereby exploiting their trust – see C Bauer, 'Banken zocken Senioren als "AD-Kunden" ab' *Westfaelische Rundschau* (9 July 2009) www.wr.de/wr-info/banken-zocken-senioren-als-ad-kunden-ab-id79712.html. In States like the US, where there has been a long tradition of investing in the financial markets also by private investors through their 401K pension plans with tax benefits, financial advisory services have been on higher professional levels. For a more thorough cross-country comparison see J Burke and A Hang (2015), 'Financial Advice Markets – A Cross-Country Comparison' (study by the Rand Corporation prepared for the US department of labor) www.rand.org/pubs/research_reports/RR1269.html.

[56] There has been a long history of insider trading; see the article by the New York Times, 'Dealbook – Timeline: A History of Insider Trading' *The New York Times* (6 December 2016), mainly focusing on cases in the US www.nytimes.com/interactive/2016/12/06/business/dealbook/insider-trading-timeline.html.

[57] Over many years traders had manipulated the banks' central lending rate, i.e. the LIBOR rate, to their benefit before it was discovered, see L Vaughan and G Finch, 'Libor Scandal: The Bankers Who Fixed the World's Most Important Number' *The Guardian* (18 January 2017) www.theguardian.com/business/2017/jan/18/libor-scandal-the-bankers-who-fixed-the-worlds-most-important-number.

[58] See Patterson, *Dark Pools* (n 27) 4, and also M Lewis, *Flash Boys* (n 27) for more details on this fascinating topic.

for trading securities outside of the exchanges, usually virtually unregulated. The benefits for the market actors were faster processing of orders with less or even no fees from the exchanges. But the real benefits for the involved high frequency trading firms were obviously financial: with their algorithms and high frequency trading infrastructures, they were able to read the directions markets were going, and being able to buy securities before the real investors could do it and then selling the securities back to them at a higher price only milliseconds after the initial orders by the investors had been made. This practice allowed them to make huge profits stealing away money from the long-term investors like pension funds, etc.

This is a very sophisticated version of an old, mostly considered illegal practice of so-called front-running – in other words, someone trading a stock or any other financial asset based on insider knowledge of a future transaction that is about to affect its price. One important point is that this practice has been around before the emergence of AI and the high frequency data processing infrastructures. But it needs to be acknowledged that the new technologies have allowed for a more sophisticated and harder-to-control form of front-running. Yet the problem here is not that the employed AI algorithms are unethical. The problem is that the actors have used the technologies in an unethical way which obviously needs to be prevented for the benefits of the wider investor community and for society as a whole. Again, this is not a new risk, but it shows that AI and technology can be an accelerator of existing risks inherent in the financial industry.

In sum, then, the arguments by *Zetzsche* and others that there are many new risks stemming from AI-based applications in financial services are not fully convincing. To the contrary, it is feasible that the employment of AI applications in the financial industry might provide a route of managing existing and inherent risks in a better way, or even being able to reduce or eliminate some of these risks.[59] But clearly, there are also cases like front-running based on algorithmic high frequency trading, where it seems obvious that through the employment of AI, existing inherent risks in the financial industry have increased and can cause additional damage. Therefore, it is also important to look at ways how such damages resulting from the use of the new technologies can be avoided. In this regard, in the following section one prominent regulatory approach, the so-called responsibility frameworks, will be discussed.

VI. RESPONSIBILITY FRAMEWORKS AS A SOLUTION FOR MANAGING AI RISKS IN FINANCIAL SERVICES?

In regulating the financial industry, many regulators have moved to so-called responsibility frameworks in recent years, like the EU's EBA/ESMA guidelines or the FCA in the UK.[60] The proposed measures focus on personal managerial responsibility, for example, the personal responsibility of directors, senior management, and individual line managers. Initially, such frameworks were meant to be applied to mitigate the risks of financial services in general, but recently authors have argued that they can also be applied to the emerging AI-based processes in the financial industry.[61]

[59] For another view, cf. T Schmidt and S Voeneky, Chapter 8, in this volume.

[60] See the report by European Banking Authority for example: EBA, 'Final Report on Guidelines on internal governance under Directive 2013/36/EU' (2 July 2021) EBA/GL/2021/05, 5-7 www.eba.europa.eu/sites/default/documents/files/document_library/Publications/Guidelines/2021/1016721/Final%20report%20on%20Guidelines%20on%20internal%20governance%20under%20CRD.pdf. For the UK, see the conduct rules as applied to the senior management functions as defined by the Bank of England report, Bank of England 'Senior Managers Regime: Approvals' www.bankofengland.co.uk/prudential-regulation/authorisations/senior-managers-regime-approvals.

[61] Cf. Zetzsche and others, 'Artificial Intelligence in Finance' (n 2).

https://doi.org/10.1017/9781009207898.028 Published online by Cambridge University Press

The responsibility-driven regulations by the EU, published by EBA and ESMA, focus mainly on the management bodies of financial institutions, in particular on their role in conducting their overall operational duties, but with a particular focus on risk management conduct. They are meant to ensure that a sound risk culture has been implemented in their respective organizations consistent with the individual risk profile and the overall business model of the institution. The UK's senior management regulatory framework for financial institutions has evolved from the overall EU framework but it has strengthened the establishment of clear conduct rules for senior managers. These rules specify in more detail the steps necessary to ensure that the business of the financial institution is controlled effectively and is in compliance with existing regulatory frameworks. Also, requirements are made on the delegation of responsibilities and on the disclosure of relevant information for the regulators. Other States like the US or Singapore have issued similar guidelines.[62]

Although these responsibility frameworks have been very general in nature, meant to capture all kinds of aspects of risk management in financial institutions, Zetzsche has argued that they will give us the right framework to address and manage any new risks stemming from AI applications in financial services. They write: 'personal responsibility frameworks provide the basis for an appropriate system to address issues arising from AI in financial services'.[63] They have suggested the following three distinct instruments or measures for regulating activities related to the development and use of AI applications in financial services:

1. *AI Review Committees:* the installation of AI review committees is meant to address what they call the information asymmetry as to the function and limits of an AI system, namely the problem that third party vendors or inhouse AI developers understand the algorithms far better than the financial institutions that acquire and use them, and the supervisors of the institutions. These committees are meant to augment decision making and should not 'detract from the ultimate responsibility vested in management [. . .] regarding AI governance.[64]
2. *AI Due Diligence:* mandatory AI due diligence should be put in place, which should be done prior to any AI employment and should include what they call "a full stock of all the characteristics of the AI [. . .] in particular the mapping of the data set used by AI", including an analysis of data gaps and data quality.[65]
3. *AI Explainability:* the explainability requirement is proposed to be necessary as a minimum standard 'demanding that the function, limits and risks of AI can be explained to someone at a level of granularity that enables remanufacturing of the code'. And this someone 'should be a member of the executive board responsible for the AI'.[66]

Before we review this proposal, it is fair to mention that the authors themselves note a few limitations, of which I want to focus on the main one, namely the inability of their responsibility framework to control what they call 'autonomous AI'. What they mean by this are cases in which developers lose control over self-learning AI, not understanding anymore what the algorithms are doing.[67] What they propose is the concept of being able to always switch off the AI (as a kind of

[62] MAS, 'Guidelines on Individual Accountability and Conduct' (MAS, 10 September 2020) www.mas.gov.sg/-/media/MAS/MPI/Guidelines/Guidelines-on-Individual-Accountability-and-Conduct.pdf, para 3.3.
[63] Zetzsche and others, 'Artificial Intelligence in Finance' (n 2) 44.
[64] Ibid..
[65] Ibid.
[66] Ibid.
[67] Such a situation seems not so rare as also discussed in the recent documentary 'The Social Dilemma' (2020) on Netflix, in which many of the creators of the algorithms underlying the leading social media platforms like Facbook or Youtube discuss their inability to understand the content proposing aspects based on user profiling at a later stage of

human oversight) while the provided services would still be functioning. This seems, *prima facie*, to be a reasonable request, looking for instance at the example of self-learning AI application in payment fraud detection based on the analysis of large transaction data sets. The system might modify its outlier detection algorithm in a way which might force the financial institution to switch it off, maybe because fraudsters have fed the system with data to facilitate fraudulent transactions. In this setting, switching the AI system off would make sense, but the delivery of the basic payment services should not be impacted by this – for instance, in the case of a credit card company. Yet, there will be applications where such a switch-off mechanism might be more difficult to realize without causing any further damage, as in the case of trading financial securities where orders or transactions 'might get lost' by switching off applications.[68]

Overall, I agree with the authors that their approach is important and should be part of any software-based technology development in financial services. In fact, many elements have been in place in the industry for years already, for instance in terms of regular due diligence audits of the financial services' technology providers.[69] Hence, the financial industry is already prepared and experienced in conducting due diligence audits on a regular basis, and they do this frequently before the release of new technology and software systems or installations, irrespective of whether these systems would include AI technology or not.

Yet, on the other hand, the authors propose some specific requirements for their AI due diligence and also in regard to their explainability requirement for AI. As will be argued, these requirements are not entirely clear and potentially will also be hard to fulfill given the nature of many AI applications.

Firstly, they argue that any AI due diligence should comprise taking full stock of all the characteristics of the AI, 'in particular the mapping of the data set used by AI, including an analysis of data gaps and data quality. It is not clear what is meant by 'a full stock of all the characteristics of the AI'. Besides talking about the different functionalities of the AI system, they also seem to focus on the underlying data set used by it. The problem here is that data sets might be potentially infinite and/or not be fully determined at the outset. In the case of its employment, a self-learning AI might discover new data (sets), and as it is often the case with big data, there can be quality issues and gaps. What does this mean for the AI system: should it not be launched under such circumstances? Or is it just enough to be aware of such limitations?

Secondly, their explainability requirement seems even harder to deal with in the case of AI applications. Even in regard to existing non-AI applications, it is questionable whether this requirement can be met given the complexity of many software solutions in the financial sector with millions of lines of code and often old legacy systems.[70] In the case of AI-based application, the situation is even more complex because learning AI systems are less static but more dynamic in nature, which could mean that the system might even rewrite its code in the course of its operations. Making explainability a minimum standard in the sense defined above could be the

the operations of the algorithms. Essentially, the algorithms develop in their own way, which is hard to understand at later stages of their employment.

[68] The other two limitations they mention are overdeterrence – as long as the benefits are higher, I think this won't be such an issue – and the increased role of fintechs in developing AI applications which usually have less experienced managerial resources. Here, they propose that by suitable board structures this could be handled, a thought with which I agree.

[69] For instance, regarding the numerous cloud-based services in place today in areas like financial market data systems, trading terminals, or wealth management advisory solutions.

[70] See for instance the recent 2020 report by the consultancy firm Deloitte on this topic: 'Modernizing Legacy Banking Systems Practical Advice to Help Banks Succeed at Core and Application Modernization' (*Deloitte*, 2020) www2.deloitte.com/us/en/pages/financial-services/articles/modernizing-legacy-systems-in-banking.html.

https://doi.org/10.1017/9781009207898.028 Published online by Cambridge University Press

end of many applications in financial services, not only AI-driven ones. But what might be argued for is that a reduced version of this principle can be applied, not to require the remanufacturing of the actual code but, at a minimum, the possibility for the functions, limits, and risks of the AI systems to be understood on a higher functional level.

Thus, to conclude this section, the so-called responsibility frameworks can provide a basis for limiting the risks of AI applications in financial services, given that new risks really emerge. In a way, they have been present in the financial industry before the rise of AI applications, such as regular due diligence audits of technical systems and key software applications. But they also have their limitations, in particular, when certain requirements are taken in a very strict sense, as it was the case above in regard to the discussed explainability requirement. Yet reducing such requirements to a lower level raises the question to what extent the risk stemming from any kinds of new potential high risk AI applications in financial services can be contained. A slightly different approach will be presented in the final section, which builds on the approach put forward recently by the EU.

VII. STANDARDIZATION, HIGH RISK AI APPLICATIONS, AND THE NEW EU AI REGULATION

As noted before in this chapter, AI, being a general-purpose technology, will impact not just one industry but will also have many different kinds of use cases in different industries. There are already a lot of AI applications in the financial industry today, as pointed out earlier, and more are being added constantly by different financial institutions, their IT service providers, and innovative fintech companies. Many of the solutions might be simple or fairly basic, like the use of face recognition as an identification method giving users access to their financial accounts or applications. The others will be more complex, like developing so-called robo advisors, with AI systems engaging with users in natural languages trying to understand their financial needs and giving them suitable financial advice.

What will be important going forward is that regarding some kinds of key AI applications – such as identification processes or human–machine interaction – there can and will be standards across industries which companies need to comply with, like there are standards and norms for the use of electricity irrespective of their specific domain of use.[71] Looking at the example of AI systems intended to interact with humans, providers of such systems might require that they be transparent to the user, that they are not communicating with humans but with a machine. Such notification obligations can then become part of the standard for such human–machine interaction enabling AI systems.[72]

Besides such general standard AI applications used across industries, there might also be ones very specific to certain industries like the financial industry which need to be dealt with outside of the model of standardization. In particular, when these specific applications give rise to higher or new risks, additional specific regulations might need to be put into place. For instance, there have been attempts to contain the risks of algorithmic trading applications in the financial industry, which can cause (and probably have already caused to some extent) significant financial damage in the form of leading markets to crash, thereby diminishing or blowing away

[71] See for instance all the different norms and standards defined by the VDE (the German association for electrical, electronic, and information technologies) over more than 100 years. In 1885, the first VDE regulation, the 'VDE 0100', was introduced, which regulated the safe construction of electrical systems. In 1904, the VDE published its first 'book for standards' comprising more than 17 provisions. Today, there exists a wide group of norms and standards ensuring the safety and well-functioning of all kinds of electrical systems.

[72] A similar obligation has been put forward by the EU in its recent Draft EU AIA, (n 6).

https://doi.org/10.1017/9781009207898.028 Published online by Cambridge University Press

investors' money in the course of seconds.[73] Such flash crashes have been at the center of some debates over the last years, and regulators such as the EU have tried to contain this risk by putting additional obligations into place through the MiFID II framework as discussed earlier.[74]

In its recent Draft EU AIA the EU has also distinguished between 'high risk' and non-high-risk (standard) AI applications.[75] The new proposed regulation starts with the assumption that AI applications are ultimately and potentially just tools to increase human well-being. Thus, the technology development of AI should not be hindered by any unnecessary constraints, but the rules should be balanced and proportionate. The regulation is centered on a 'risk based regulatory approach, whereby legal intervention should be tailored to those concrete situations where there is a justified cause of concern'.[76] A key distinction is made between the so-called high risk AI systems for which special requirements and obligations then apply and other AI systems with much more limited requirements and obligations. The classification of AI systems as high-risk is thereby mainly based on their intended purpose and their harmful impact on health and safety and human rights. High risk systems are more or less identified in a two-step process, namely whether they can cause certain harms to protected goods or rights and by the severity of the harm caused and the probability of occurrence.

Given this approach, it seems obvious that there cannot be one final list of high-risk AI applications because the technology is still emerging and new applications are being launched every day. The EU acknowledges this as it lists in its Draft EU AIA only a limited number of high-risk AI applications (Annex III). Further, it allows the EU Commission to amend this list over time based on criteria spelled out in Article 7.

Interestingly, many of the high-risk applications listed by the Draft EU AIA are not specific to one industry but are general AI applications that can be present in many industries. Examples are applications that embody what is called 'manipulative AI practices' and a second group with 'indiscriminate surveillance' practices. But there are also many very specific high risk AI applications listed in the draft. When it comes to high-risk AI applications in financial services, the EU draft of the regulation lists, prima facie, only one class, namely AI systems that evaluate the creditworthiness of persons (Annex III No 5 lit. b).[77] This class of applications is included in the high-risk list because of (i) possible discrimination of persons of certain ethnic or racial origin based on the potential perpetuation of historical patterns by the AI algorithms, and (ii) the potential severity of such acts of discrimination, as in the way such discriminating credit decisions can significantly affect the course of life of people.[78]

The second kind of AI application that can be associated with the financial services industry, listed in the Draft EU AIA, is the one written about above, namely AI systems intended to

[73] In the literature, there has been a long discussion of the so-called flash-crashes and the extent to which they have been caused by certain algorithmic trading practices. See Busch, 'MiFID II' (n 48) on this topic for a more detailed discussion regarding the recent MiFID II regulation and its impact on algorithmic trading practices. See also Huang and others, 'AI Coding' (n 46) for more details on this topic.

[74] For more details on this topic see Section IV and Huang and others, 'AI Coding' (n 46) of this chapter.

[75] For details cf. T Burri, Chapter 7, T Schmidt and S Voeneky, Chapter 8, and C Wendehorst, Chapter 12, in this volume.

[76] See the Draft EU AIA, (n 6).

[77] For requirements to be met by high-risk AI systems, cf. Article 8 et seq., Article 16 et seq., and especially the conformity assessment, Article 43.

[78] I assume this refers to the fact that simple learning algorithms might be trained on past credit decisions of financial institutions which might have embodied certain forms of discrimination. As has been pointed out before in Section IV, also before the arrival of AI in financial services, many credit decisions have been prone to discrimination. One solution could be that in training algorithms on making such credit decisions the training data could be prepared in a way that would make them bias free.

interact with natural persons or generate content consumed by such person. Such systems do not necessarily classify as high-risk systems, for instance they might just help someone to enter information or explain a product, but they pose the specific risk of impersonation and deception, and therefore they are subject to specific transparency obligations according to the Draft EU AIA that means that the natural persons have to be informed that they are interacting with an AI system (Article 52).

Overall, the risk-based regulatory approach which underpins the Draft EU AIA makes much more sense than any kind of generalized approach of regulating AI applications as a whole as embodied in the responsibility frameworks discussed above. As a general-purpose technology, there will be so many kinds of applications that not one standard set of rules can be applied across the board. A rigorous case by case approach is required, which also allows for amendments and revisions, as embodied in the outline of the EU regulation.

VIII. CONCLUSION

What has been shown in this chapter is first, that it is questionable whether there are many new additional risks stemming from AI applications in financial services today. The risks that have emerged recently, like data risks, cybersecurity risks, financial stability risks, and ethical risks have been inherent in the financial industry as a highly digitized and also complex global industry for decades. The author has taken the more positive view that by using AI these risks will not necessarily increase, but on the contrary, AI might help to mitigate and reduce them. Second, the responsibility frameworks as developed over the last few years, which are meant to deal with and limit the risks of AI in the financial industry overall, do not provide a suitable framework beyond what has been put in place already to manage the risks with more standard IT and software systems and applications in the financial industry. Furthermore, overseeing all AI applications in financial services will quickly become as complex as overseeing all types of applications in the area of electricity, to mention another general-purpose technology. What has been argued in this chapter is that for some kinds of key applications – like identification processes or human–machine interaction – there should be standards defined across industries with which companies need to comply. But for other very specific, potentially new high-risk financial AI applications, in case they emerge, there might be the need for additional very specific regulation, as in the case of certain algorithmic high frequency trading applications. But this will be less a regulation of the technology but more of the practices and intended uses of the technology, which has also been the core thinking underlying the recent Draft EU AIA of AI applications. In fact, this new EU regulation, like the GDPR a few years ago in regard to data privacy protection, in many ways points to the right direction of how to deal with AI and potential risks arising from it.