Artificial Intelligence in Central Banking

Aura Elena GRIGORESCU
Bucharest University of Economic Studies, Romania
grigorescu_aura@yahoo.com

Abstract. *The paper uses qualitative research to investigate the potential uses of artificial intelligence in the field of central banking. The analysis shows that monetary policy, prudential supervision and the oversight of payments are the areas where the use artificial intelligence is most likely to bring benefits. Monetary policy calibration involves working with long time series of data for various parameters and making the necessary analysis and forecasts, an activity in which artificial neural networks may prove useful. Bank supervision can benefit from the use natural language processing algorithms that can read documents and extract the relevant information. Such algorithms can read all of the required documents (not just those that the supervisor selected) and return all of the sentences that contain a certain predefined expression. In the field of the oversight of payments, the capabilities of machine learning to identify new patterns or anomalies in the data that could indicate fraud or money laundering will boost the efforts to combat them. In terms of challenges associated with the use of artificial intelligence in central banking, perhaps the two biggest challenges are that some of the models do not allow for a reasonable level of explainability of the algorithm(s) through which they arrive at the result (especially relevant for bank supervision) and data availability. With respect to the latter, although the issue of quantity of data can be dismissed as a shortcoming given the huge amounts of data available, the issue of data quality seems to be more pronounced, as deficiencies such as data measured incompletely or incorrectly, scarcity and regulatory barriers that impede data sharing may be difficult to surpass.*

Keywords: artificial intelligence, central banking, monetary policy, oversight of payments, bank supervision

Introduction

Perhaps Artificial Intelligence (AI) became a topic of interest when Turing (1950) published his paper "Computing Machinery and Intelligence", which is probably today one of the most well known in this field. The paper revolved around the question "Can machines think?" to which the author answered in an affirmative way, by concluding that machines will eventually be able to compete with humans in all intellectual fields.

Was this assumption correct? Fast forward almost 50 years later, in May 11th 1997, Deep Blue, a chess-playing computer developed by IBM, defeated chess grandmaster Garry Kasparov, a result that was widely debated at that time. After another 25 years, ChatGPT was developed, as an AI system, which, among other things, can engage in a meaningful conversation with a human user. These developments show that Turing's (1950) vision of machines being able to compete with people in intellectual fields might not be as utopic as it may have been considered back when it was first stated.

This observation raises the question: what is AI? Defining it in a simple way, AI is a branch of computer science that builds machines capable of performing tasks that usually require human intelligence i.e. machines that use artificial intelligence to try to reproduce human behavior as accurate as possible. This is done by processing large amounts of data through advanced algorithms, with the aim of identifying and learning data patterns or other characteristics. On a

DOI: 10.2478/picbe-2024-0159

more specific level, these processes can be divided into various groups, some of which are mentioned below, without being an exhaustive enumeration:

- Machine learning (ML), means that a machine learns to make decisions based on experiences by using what it learnt from data (for example patterns). Most of the times, a human expert sets up the data and the objective(s), in order to guide the algorithm towards the desired result. However, there is also the possibility to let the algorithm work on its own when it comes to identifying the required models, thus removing the need for constant human intervention.
- Artificial Neural Networks (ANN), which are a series of algorithms that attempt to recognize relationships in a data set through a process that mimics how the human brain works. They are formed from artificial neurons, which signal to each other and are organized in layers, which are subsequently stacked.
- Deep Learning (DL), which means that a machine learns to process inputs through multiple layers in order to classify, infer and finally predict the outcome. These algorithms are based on ANNs with more layers.
- Natural Language Processing (NLP), which allows a machine to read and interpret a language, both written and spoken. Once the machine understands what the user wants to communicate, it can respond to his requests with high accuracy.

Judging by what these algorithms aim to do (and probably will be able to do very accurately in the future) and how to do it, AI looks suitable for being used in tasks that involve repetitive operations with a rather precise routine. This is relevant for the scope of this paper, especially since these algorithms can *(i)* identify patterns and characteristics of the data and, furthermore, make predictions and *(ii)* read and interpret languages, two features that may be useful, as I will show later on. Finally, given the rather low number of people familiarized with AI, a simpler way to looks at these algorithms can be that of Araujo et al. (2024) who identify machine learning as an evolution of classic statistical and econometric techniques, but with the difference that the former do not rely on a pre-specified models or other assumptions, tools on which central banks' activity relies heavily.

Literature review

Due to the fact that the subject of AI use in the field of central banking doesn't seem to be yet widely debated, I searched mostly for official views on this topic of various institutions. From the international financial institutions, the views of the Bank for International Settlements (BIS) and the Federal Reserve Bank (FED) were considered to be the most relevant for this paper. These were complemented by opinions from private researchers and case studies.

As an example, the BIS, through the work of Araujo et al. (2024) signaled that central banks can be seen as early adopters of machine learning techniques in areas such as statistics, payment systems oversight and supervision. This adoption is considered to be a successful one, albeit with some challenges, as it will be discussed later in the paper. An interesting idea is highlighted as a possible solution to (partially) overcome these challenges in the form of a collaboration between central banks, with the aim to share knowledge and practical experiences in this field. The BIS also took some significant steps in this area, through Project Aurora, which explores the use of machine learning and network analysis for detecting money laundering activities. This project demonstrates the advantages and potential of using payments data in combination with AI algorithms in order to detect (complex) money laundering schemes. FED's Mullin (2023) states that AI's role in commercial banking has been growing and this trend will continue. From this perspective, central banks must monitor commercial banks' use of AI.

Attention was also put on the works of the academia. A practical example comes from Tan et al. (2022), who study the use of AI by the Central Bank of Malaysia in order to improve the supervisory communication with commercial banks. The authors, in an effort to improve both the efficiency in the writing process as well as the consistency in the communication tone of supervisory letters, propose an application that allows for the calibration of the tone used in the supervisory letters, on a scale from "Neutral" to "Forceful" and also for the searching for similar sentences from past letters, with the aim to speed up the writing process of the current letters. Buckmann et al. (2023) offer a good example of how a neural network can break down services inflation into different components (past price increases, inflation expectations, the output gap and international prices), thus allowing a determination of the contribution to inflation of each of these components. Danielson & Uthemann (2023) consider that AI may not help in the case of extreme events, when what should have been the relevant data is not known until the event has occurred, highlighting some of the shortcomings of AI.

Methodology

The methodology used was qualitative research, a technique I considered feasible because there is still little (official) data regarding the potential usage of AI in central banks' activities. Therefore, a quantitative approach didn't seem fit for the purpose of this paper. Consequently, I preferred to conduct the research in a manner of blending the information I was able to find on AI with my own professional experience in the field of central banks' activities. I used sources from various authors and entities (such as described in the literature review section) and I tried to identify those areas where a consensus was reached, at least at a conceptual level, on the genuine potential uses of AI. Afterwards, I used these ideas to try to sum up the benefits, but also the challenges, of using AI in central banking. All of the data used was publicly available data, most of it being also accessible online.

Results and discussions

Uses of artificial intelligence in central banking

Commercial banks, which are subject to regulation and supervision by the central banks, have started to use AI in various ways. Perhaps one of the most known area is that of customer support, where robots have started to replace humans for basic tasks (Saurabh, 2024), leaving the option to engage with a bank employee only for those tasks that need a certain type of tailoring to the customer's needs. However, more relevant is the use of AI for risk management purposes such as fraud detection (Juneja, 2023), compliance with regulatory requirements and various analytics that allow, for instance, the computing of the likelihood of defaults.

Consequently, central banks cannot ignore these developments, not only because they will (have to) to use themselves AI, but also because they have to understand how the AI used by the commercial banks works. Consequently, in order to identify what could be the current uses of AI in central banks, the best starting point would be the central bank's responsibilities. For instance, if we consider the National Bank of Romania (NBR), we see that, according to The National Bank of Romania Act (Law No.312/June 2004), the NBR is primarily responsible with ensuring and maintaining price stability. In this respect, the main tasks of the NBR are the following:

- to define and implement the monetary policy and the exchange rate policy;

- to conduct the authorization, regulation and prudential supervision of credit institutions;
- to oversee the smooth operation of payment systems;
- to issue notes and coins to be used as legal tender on the territory of Romania;
- to manage the official reserves of Romania.

Now, based on what has been previously mentioned about the possibilities that AI offers and taking into consideration the specific requirements needed for fulfilling these tasks, it can be speculated that monetary policy, prudential supervision and the oversight of payments may be the areas where the use AI can bring benefits. The next section of the paper will emphasize on each of them.

Uses of artificial intelligence in monetary policy

From the NBR's dedicated section on its website, it can be drawn the conclusion that monetary policy decisions are based on analyzing data. At the core of the decision process sits the inflation forecast and the factors that affect it. The macroeconomic variables relevant for each decision seem to be the same, more or less, (inflation rate, interest rate, exchange rate, output gap, cyclical position of the economy, labor market conditions etc.), but it can be considered that their importance may be adjusted over time, taking into consideration the macroeconomic developments.

Based on those mentioned above, the fact that monetary policy calibration involves working with long time series of data for various parameters and making the necessary analysis, implies a certain degree of difficulty in identifying the relevant information in terms of predictions, correlations and contributions of each factor to the inflation forecast. In this situation, AI can prove be a valuable tool.

For instance, Buckmann et al. (2023) offer a good example of how an ANN can help with monetary policy analysis, by showing how it can break down services inflation into different components (past price increases, inflation expectations, the output gap and international prices), thus allowing a determination of the contribution to inflation of each of these components. To do this, they use an ANN to approximate the Phillips Curve for United Kingdom's services inflation. Its structure is presented in the graph below, where it can be seen that the Phillips Curve links the actual inflation to the four drivers mentioned earlier.

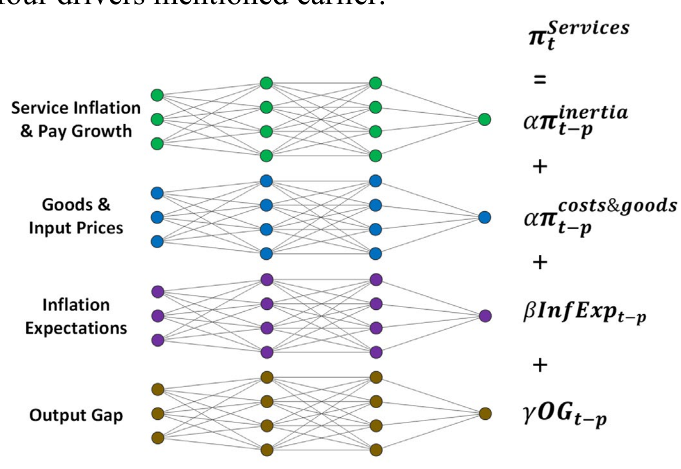

Figure 1. A neural network with Phillips Curve structure

Source: Buckmann et al. (2023).

The authors find that this algorithm can help explaining the drivers behind the rise in the United Kingdom's services inflation, but it has limited capabilities of differentiating between each driver and also that, like any empirical model, it has some uncertainty around its predictions.

Uses of artificial intelligence in prudential supervision

The activity of prudential supervision has three elements: risk identification, risk assessment and risk mitigation. Obviously, it involves dealing with large quantities of data. The supervisor deals with many numerical data, whether it is the credit portfolio of the supervised entity or other indicators that need to meet certain values, in order to comply with the legal requirements in force. Analyzing such huge volumes of data is time consuming and prone to operational error, given that, ultimately, the supervisor will not be able to cover all of the information and, therefore, face uncertainties in making the final assessment. In other words, supervisors may be prone to biases when it comes to judgment-based decision making, which may lead to inadequate risk identification and / or suboptimal decision making on risk mitigation measures.

Let us consider how a traditional supervision exercise is carried out. The supervisor asks and receives documents form the supervised entity. Most of these are reports, minutes of the board's meetings and so on of various lengths and formats. Let us also recall that, after the Great Financial Crisis of 2008-09, a lot of emphasis was put on the need that banks have to make public disclosures more than before, which basically means submitting more reports. After the requested data is submitted by the supervised entity, the supervisor will end up with a folder that will look, more or less, like the one in the picture below[1]:

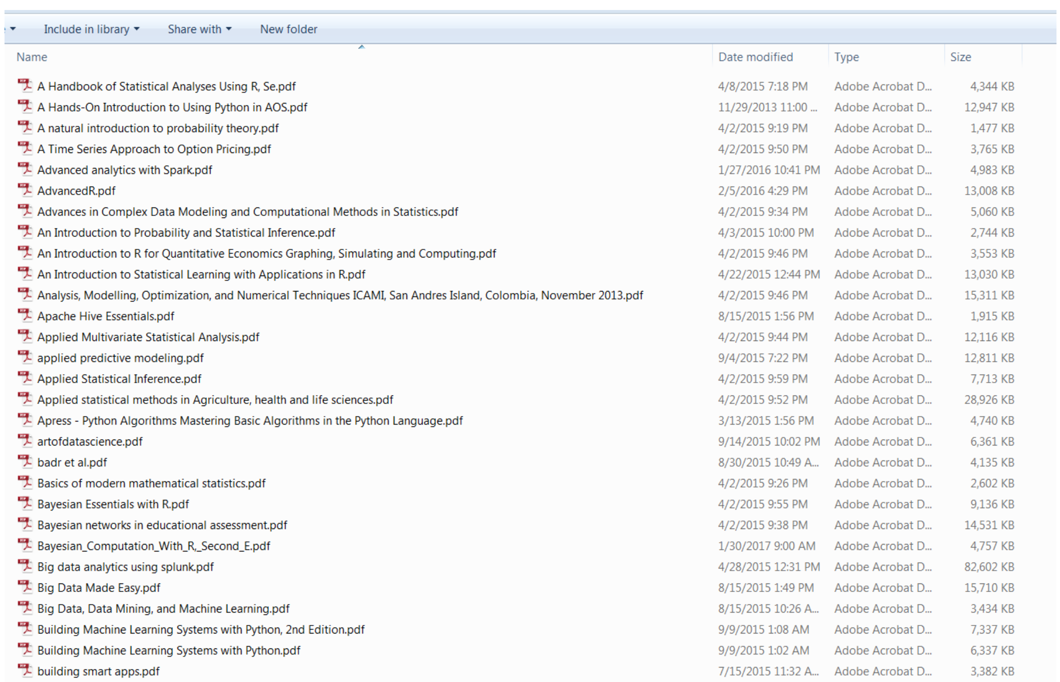

Figure 2. Supervisory documents received by bank supervisors

Source: the Internet.

[1] The purpose of the picture is to show to the reader the large number of the files and the diversity of the topics. Due to confidentiality, I did not use a folder with the actual files received from a supervised entity.

The supervisor will have to somehow sort all of these documents and choose which ones he/she will have to read. To make matters even more difficult, those documents are in various formats, ranging from texts to graphs and tables. Finally, when the supervisor starts reading the selected documents, all of the relevant information has to be manually extracted and, for instance, saved in a separate file. This process will repeat for each document, until the supervisor finishes reading the selected documents.

In this moment, let us recall from the Introduction section that jobs with repetitive operations and a rather clear routine can be replaced by a machine and also that some machines can read and interpret languages. Consequently, a NLP algorithm seems to fit perfectly on that part of the supervisor's job that requires reading documents and extracting the relevant information. For example, if the supervisor is interested in knowing the credit risk associated to the loan portfolio, such algorithms can read all of the documents (not just those that the supervisor selected) and return all of the sentences that contain the expression credit risk. Furthermore, the NLP algorithm can also return the details of where each sentence was found (document name, page etc.). It can be seen that, aside from the fact that such algorithm can cover all or, at least, more documents than the supervisor could, it can also find those sentences that the supervisor missed. Moreover, even if the algorithm returns some useless results, the supervisor will waste very little of his time analyzing and dismissing these.

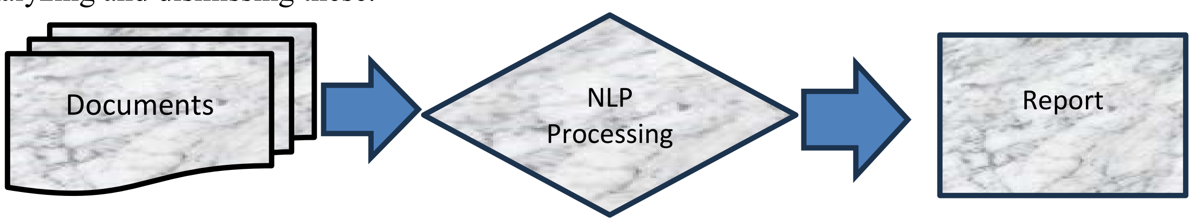

Figure 3. A basic scheme for a NPL algorithm

Source: author's' illustration.

A practical example is the use of AI by the Central Bank of Malaysia in order to improve the supervisory communication with commercial banks. According to Tan et al. (2022), the Central Bank of Malaysia has more than 120 institutions under supervision, which means that writing and reviewing supervisory letters requires a considerable amount of time and work. Therefore, the authors, in an effort to improve both the efficiency in the writing process as well as the consistency in the communication tone of supervisory letters, propose an application that allows for tone analysis and sentence search. The first feature allows a better calibration of the tone used in the supervisory letters, on a scale from "Neutral" to "Forceful", with two intermediary values "Cautious" and "Concerned". The second feature allows for searching for similar sentences from past letters, with the aim to speed up the writing process of the current letters.

Uses of artificial intelligence in the oversight of payments

Two of the main problems with payments are frauds and money laundering[2]. If in the case of supervision I showed the benefits of using NPL algorithms, in this case the attention will focus on machine learning, given its ability to learn from historical data and identify new patterns or anomalies that could indicate fraud. Recall that these algorithms are very good at finding patterns that have a high chance of going unnoticed by a human analyst.

[2] The process through which money obtained from illegal activities are hidden behind legitimate sources.

Now, one can argue that there is not a tremendous number of frauds happening, and that banks have already in place various security measures, but even so, the need to minimize the occurrence of this problem cannot be contested. The second issue, money laundering, is a global problem and amounts to an estimated 2 to 5% of global GDP, or between $2 trillion and $5 trillion, with a detection rate of less than 1% of those estimates (BIS, 2024).

An important role here is the calibration of the algorithm, given that a selection must be made from all of the data that is available about a transaction in order to assess whether it can be considered suspect or not. Commonly several attributes stand out i.e. size, location from where the payment was made, frequency, time etc. The algorithm uses these attributes to predict if a transaction has a fair degree of suspiciousness attached to it, given the usual behavior of that particular customer.

A good example in this area in Project Aurora, developed by the Bank for International Settlements Innovation Hub (2023), which explores the use of machine learning and network analysis for detecting money laundering activities. According to its creators, the project demonstrates the advantages and potential of using payments data in combination with AI algorithms in order to detect (complex) money laundering schemes.

Further examples of AI tools that are used by central banks can be found in the online annex to the work of Araujo et al. (2024), as well as in their previous work dedicated to machine learning applications in central banking (Araujo et al., 2022). The latter is particularly relevant, as it shows central banks' interest in machine learning, which highlights the potential that these applications may have in helping central banks to calibrate their decisions.

Challenges associated with the use of artificial intelligence in central banking

Perhaps the biggest challenge associated with the use of AI in central banking stems from the fact that some of the models do not allow for a reasonable level of explainability of the algorithm through which they arrive at the result, compared to traditional econometric techniques that are well understood. Ultimately, the central bank needs to explain how that result was computed, by identifying the parameters that were taken into consideration and by explaining the rationale behind their interpretation in an understandable manner. When it comes to bank supervision, this issue is even more pressing, as this activity is prone to some challenges when the commercial banks use AI models for assessing the credit risk of potential customers or portfolios (Mullin, 2023). In the second case, unfavorable outcomes may be more serious. If the supervisor is not be able to understand the models, given their characteristics, then it would be nearly impossible to assess if the model perpetuates any inaccuracy presented in the data on which it was trained or its compliance with the regulations in force. In a broader sense, this observation relates to the tradeoff between a model's capability to approximate different functions and its explainability. This issue was researched by Guidotti et al. (2019) and they found that there can be identified various levels of explainability, depending on the modelling stage. But even so, their comprehensive analysis of the available literature concluded that some important questions on this matter still remained unanswered. More specific are Boukherouaa et al. (2021), who identify several ways through which explicability can be enhanced: *(i)* by design, meaning that some models may be easier to interpret, given their type and algorithm employed; *(ii)* by control, meaning that during the models' development, various tests can be deployed to assess the models' viability; *(iii)* by using other models, simpler or complex, to explain the behavior of more complex models

Another challenge comes from the data itself. The success of any model is based on both the quantity and the quality of data. The issue of quantity of data can be dismissed as a shortcoming given the huge amounts of data available and because even personal data tends to be more and more available, as people give their consent to its usage. Nevertheless, the way these models work i.e. by using data from multiple sources raises some questions in respect to data protection & privacy. While data collected by public institutions enjoy a clearly defined framework in terms of access and protection, gathering data from private sources, such as social media, is prone to poor data quality and uncertainty in terms of how and to what extent that data can be used. These questions are exacerbated by the fact that AI algorithms need more and more (personal) data for a better calibration of their algorithms and also because they can deduct some sensitive information about a specific person (personal information, such as habits, preferences etc.), based on that data. Tackling these concerns will require a joint effort by policy makers, mainly through revising the laws and establishing preemptive measures and AI developers and tech companies, who need to engage in self-regulation by creating models that respect user privacy. This can be done by implementing robust data protection measures or by reducing the amount of data required/used, albeit the latter might be deemed counterproductive, given that it may affect the models' capabilities. Finally, it is also worth mentioning here the risk of cyber-attacks. Comiter (2019) calls these risks „artificial intelligence attacks" and they basically involve altering the behavior of the model to serve a malicious end goal. In case of the financial sector, such an attack can undermine the capabilities of some of the participants to assess and manage risks, which could lead to the buildup of unobserved (systemic) risks. Consequently, given their potential impact on financial sector institutions, a proper regulatory framework is required as a mitigation measure. This framework could include compelling provisions regarding detection and reporting systems and strategies to secure model and data privacy. It should be noted that such a framework also needs to apply to the IT providers of these models, as they too can be subject to cyber-attacks.

The third challenge relates to the needs in terms of investments and human resources that a greater use of AI could require. While the central banks are is a position to make the necessary investments in terms of both hardware and software, compared to other state entities for instance, large costs may be difficult to explain to the public. Finding the right employees will prove, most likely, to be even more difficult. Given the nature of the central banks' activity, the ideal employee would need to be both a programmer and an economist, a set of skills that may be difficult to find or to retain, in the case of existing staff (Araujo et al., 2024). Salaries might also be a problem, as usually public institutions cannot compete with the private sector when it comes to the maximum levels they can offer.

Conclusion

In the field of central banks activities, monetary policy, prudential supervision and the oversight of payments are the areas where the use AI can bring benefits. Monetary policy calibration involves working with long time series of data for various parameters and making the necessary analysis, which implies a certain degree of difficulty in identifying the relevant information in terms of predictions, correlations and contributions of each factor to the inflation forecast. Here artificial neural networks may prove useful. Bank supervision can benefit from the use natural language processing algorithms that can read documents and extract the relevant information. Such algorithms can read all of the required documents (not just those that the supervisor selected) and return all of the sentences that contain a certain expression predefined by the supervisor. In the field of the oversight of payments, the capabilities of machine learning to identify new patterns or anomalies in the data that could indicate fraud or money laundering will boost the efforts to combat them.

In terms of challenges associated with the use of artificial intelligence in central banking, perhaps the biggest challenge is that some of the models do not allow for a reasonable level of explainability of the algorithm(s) through which they arrive at the result. When it comes to bank supervision, this issue is even more pressing, because if the supervisor is not be able to understand the models, then it would be nearly impossible to assess if the model perpetuates any inaccuracy present in the data on which it was trained or if it complies with the regulations in force. Another challenge comes from the data availability. Although the issue of quantity of data can be dismissed as a shortcoming given the huge amounts of data available, the issue of data quality seems to be more pronounced. Some deficiencies such as data measured incompletely or incorrectly, scarcity of data, regulatory barriers that impede data sharing may be difficult to surpass. The third challenge relates to the needs in terms of investments and human resources that a greater use of AI could require. Given the nature of the central banks activity, the ideal employee would need to have a set of skills that may be difficult to find and salaries might also be a problem, as usually public institutions cannot compete with the private sector when it comes to the maximum levels they can offer.

Finally, it is worth mentioning that artificial intelligence excels is situations where both the task and the rules that need to be applied in order to fulfil that task are clear. In contrast, if the rules or the tasks becomes more nuanced, then the advantage of using AI instead of human judgement begins to fade. This could be the case when knowing how to fulfil the task is clear, but by doing so there is the risk of affecting other objectives (the classic example being that of raising the interest rate to fight inflation, assuming the potential cost of affecting the economic growth). Consequently, it does not seem that artificial intelligence will be able to fully replace human judgement.

References

Araujo, D., Doerr, S., Gambacorta, L., & Tissot B. (2024). Artificial intelligence in central banking. *BIS Bulletin*, no. 84.

Araujo, D., Bruno G., Marcucci J., Schmidt R. & Tissot, B. (2022). Machine learning applications in central banking. Retrieved from: https://www.bis.org/ifc/publ/ifcb57_01_rh.pdf

Athey, S., & Imbens, G. (2021). Machine learning methods that economists should know about. *Annual Review of Economics*, no. 11, 685–725.

BIS Innovation Hub (2023). Project Aurora: the power of data, technology and collaboration to combat money laundering across institutions and borders. Retrieved from: https://www.bis.org/publ/othp66.htm.

Boukherouaa, E., AlAjmi, K., Deodoro, J., Farias, A., & Ravikumar, R. (2021). Powering the Digital Economy: Opportunities and Risks of Artificial Intelligence in Finance. Departmental Papers, 2021(024), A001. Retrieved from https://doi.org/10.5089/9781589063952.087.A001

Buckmann, M., Potjagailo, G., & Schnattinger, P. (2023). Dissecting the use of AI service inflation via a neural network Phillips curve. Retrieved from: https://bankunderground.co.uk/2023/07/10/dissecting-uk-service-inflation-via-a-neural-network-phillips-curve/.

Comiter, M. (2019). Attacking Artificial Intelligence: AI's Security Vulnerability and What Policymakers Can Do About It. Belfer Center for Science and International Affairs, Harvard Kennedy School. Retrieved from: https://www.belfercenter.org/publication/AttackingAI

Danielson, J., & Uthemann, A. (2023). On the use of artificial intelligence in financial regulations and the impact on financial stability. Retrieved from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4604628#paper-references-widget.
Danielson, J. (2023). When artificial intelligence becomes a central banker. Retrieved from: https://cepr.org/voxeu/columns/when-artificial-intelligence-becomes-central-banker.
Guidotti, R., Monreale, A., Ruggieri S., Turini F., D. Pedreschi D., & F. Giannotti, F. (2019). A Survey of Methods for Explaining Black Box Models. *ACM Computing Surveys*, 51 (5), 1-42.
Juneja, P. (2023). Application of Blockchain in Commercial Banking. Retrieved from: https://www.managementstudyguide.com/application-of-blockchain-in-commercial-banking.htm.
Mullin, J. (2023). Artificial Intelligence and Bank Supervision. *Econ focus*, 2(2023), 8-11.
Saurabh, S. (2024). AI in banking – How Artificial Intelligence is used in banks. Retrieved from: https://appinventiv.com/blog/ai-in-banking/.
Tan, J., Shum, C., K,. & Amri, M., A. (2022). Supervisory letter writing app: expediting letter drafting and ensuring tone consistency. In Bank for International Settlements (ed.), *Machine learning in central banking*, volume 57.
Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, Volume LIX (1950), 433–460.