

Leveraging Artificial Intelligence in Financial Market Supervision: Applications, Challenges, and Insights from the Czech National Bank

Vladimir TOMSIK^{1*}, Jan BRIDZIK²

^{1,2} Czech National Bank, Na Prikope, Prague, Czech Republic

*Corresponding Author Email: vladimir.tomsik@cnb.cz

Abstract

The rapid development of Artificial Intelligence (AI), particularly Large Language Models (LLMs), presents significant opportunities and challenges for financial supervisors globally. The Czech National Bank (CNB) is actively exploring the potential of AI to enhance the efficiency and effectiveness of its licensing and supervisory functions, especially concerning the evolving regulatory landscape including MiCA and DORA. This paper presents practical insights from the CNB's pilot testing, discussing a broad range of applications including market integrity monitoring (e.g., influencers, sanctions screening) and automated document analysis for licensing and compliance checks. It provides a particular deep dive into a project focused on automating the assessment of licensing documentation, which demonstrated the potential for a dramatic reduction in processing time from an average of 80 working days to single-digit days. The paper highlights crucial lessons from this process, including the necessity of a "human-in-the-loop" approach to mitigate AI inaccuracies and a shift towards structured inputs (e.g., JSON) to overcome the limitations of simple prompt engineering. The findings underscore that while AI offers substantial efficiency gains, its successful integration requires a strategic, process-oriented approach and cannot replace expert human judgment.

Keywords

Artificial Intelligence, AI, Financial Supervision, Licensing, Central Banking, MiCA, DORA, AI, Artificial Intelligence, Capital Markets, CASP (Crypto-Asset Service Providers), Central Banking, ChatGPT, DORA, Financial Supervision, Finfluencers, Human-in-the-Loop, Large Language Models (LLM), Licensing, MiCA, Process Automation, Retrieval-Augmented Generation (RAG), Risk Management, Sanctions Screening

INTRODUCTION

The recent surge in Artificial Intelligence (AI) capabilities, driven largely by advancements in Large Language Models (LLMs) and accessible platforms like ChatGPT and Microsoft Copilot, has spurred financial institutions and regulators worldwide to explore its potential applications. For central banks and supervisory authorities, AI offers the prospect of significantly enhancing operational efficiency and analytical capacity in the complex domain of financial market oversight. The Czech National Bank (CNB) recognizes this potential and has undertaken initiatives to test and evaluate the practical use of AI within its internal processes, particularly in the demanding areas of licensing and ongoing supervision of financial market entities.

This exploration is particularly relevant given the evolving regulatory framework, including the implementation of the Markets in Crypto-Assets Regulation (MiCA) [1] and the Digital Operational Resilience Act (DORA) [2], alongside the continuous oversight of established capital markets. These regulations often involve the assessment of large volumes of complex documentation and require robust monitoring capabilities, areas where AI could potentially offer substantial support. Indeed, recent research highlights the significant potential of technologies such as Retrieval-

Augmented Generation (RAG) for automating demanding tasks like financial document processing and regulatory compliance reporting [3].

This paper aims to share the CNB's initial experiences and findings from pilot projects focused on leveraging AI in its financial market supervision activities. It deliberately focuses on the practical applications demonstrated, the limitations encountered, and the risks identified during this exploratory phase. While the CNB is also developing an internal governance structure for AI adoption, this paper concentrates on the 'what' and 'how' of AI application in supervision, rather than the organizational 'who'. The following sections will detail specific use cases tested, the types of AI tools employed, the challenges faced, and conclude with reflections on the future role of AI as a tool for financial supervisors.

PRACTICAL APPLICATIONS OF AI IN CNB SUPERVISION

The CNB's testing revealed several promising areas where AI can assist in licensing and supervisory tasks, often yielding significant efficiency gains.

Document Analysis and Information Extraction

A core function of supervision involves analyzing vast amounts of documentation submitted by financial entities. AI tools have shown considerable promise in automating parts of this process:

1. *Licensing Documentation Review (MiCA and DORA Example)*: One of the most laborious yet crucial activities in supervision is the assessment of complex licensing documentation. Given the expected enormous increase in license applications for Crypto-Asset Service Providers (CASP) in connection with the MiCA regulation [1], we decided to implement a pilot project in this area to verify the possibilities of automating and streamlining the entire process.

Our approach was evolutionary, progressing from the initial manual processing (until Q3/2024), through the testing of four progressively evolving semi-automated AI platforms (from Q4/2024 to Q4/2025), to the design and testing of a fully automated solution based on direct API integration. In this context, a 'platform' is defined as a combination of a process workflow and the utilization of a specific version of a cloud-based language model (e.g., ChatGPT, Copilot) for individual analysis steps. Throughout this development, we monitored three key parameters: error rate, time consumption, and technical processing limits, as detailed in the following figures.

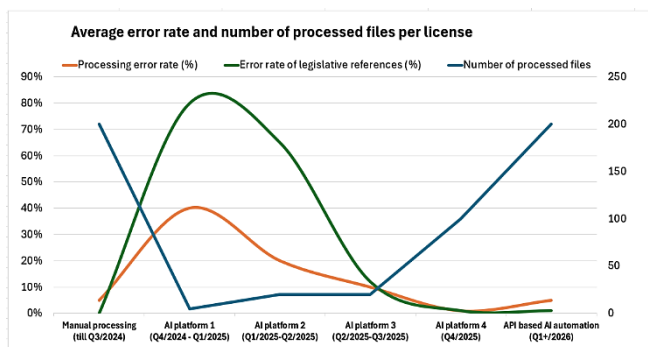

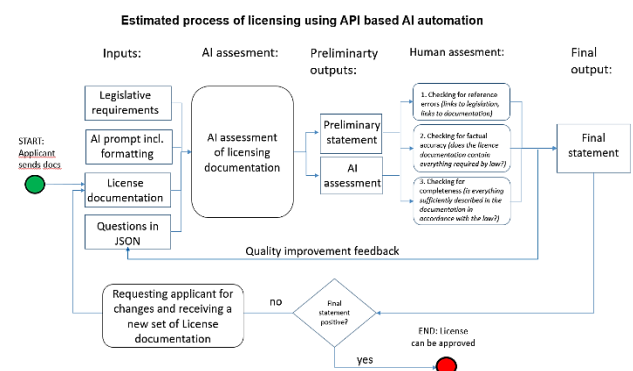


Figure 1 illustrates the evolution of the error rate and the limitations on processed files. We tracked two types of errors. The overall processing error rate paradoxically increased after the introduction of the first AI platform. This was caused by both the limited capabilities of AI at the time and our inexperience with prompt engineering. With subsequent platforms, the error rate gradually decreased, with the near-zero values for platform 4 being achieved partly as a side effect of semi-automation—the expert, while manually copying partial outputs from the AI into the target document, detected and corrected errors and model "hallucinations". With the final, fully automatic API solution, there is a slight resurgence in the error rate, as this element of continuous human micro-correction is eliminated. The second monitored

parameter, the error rate of legislative references, shows a steady downward trend after an initial problem, which we attribute primarily to the improving capabilities of the models and the increasing context window size, which allows for a better understanding of connections within extensive legal texts. The third curve highlights a significant technical barrier of the early stages, namely the limitation on the number of processed files. This limit had to be circumvented by manually merging or restricting input documents, which complicated the process. The proposed API solution is already capable of natively processing hundreds of files in a single batch.



Automation had an even more dramatic impact on the average net processing time for a single license, as shown in **Figure 2**. We divided the entire process into two phases: an initial completeness assessment (whether the applicant submitted all required documents and information) and a subsequent, more detailed substantive assessment (evaluating the content and meaning of the documentation). While the manual processing of a single application, which typically contains 100-300 files with an inconsistent structure, took an average of 80 working days, the target API solution reduces this time to a matter of days. The graph clearly demonstrates how, with the gradual improvement of AI accuracy and the overcoming of file processing limits, the process bottleneck shifted from the analysis itself to the manual consolidation of outputs in platform 4, and was finally eliminated in the fully automated solution.



The target state, which we are currently testing, is depicted in **Figure 3**. This process is based on a synergy between artificial intelligence and expert oversight. As input, the AI receives not only the license documentation itself but also a set of legislative requirements, a specific AI prompt defining the format and tasks, and, crucially, a structured file of questions in JSON format. The AI then performs an initial assessment. Its outputs ("Preliminary statement" and "AI assessment") are handed over to an expert, whose role is no longer the analysis itself, but a targeted validation in three areas: checking for reference errors, factual accuracy, and content completeness. The process includes two key feedback loops. The internal loop serves for the continuous improvement of output quality by refining the set of questions based on the expert's findings. The external loop then illustrates the communication with the applicant, which is not part of the automation — in the case of a negative assessment, a final report (locally termed 'Výzva' [Call for remedy]) is formally sent to them, and the process is repeated with a new set of documents. Furthermore, this process framework is universal and can be applied with the same logic for both the completeness and substantive assessments, with the only difference being the content of the input questions and prompts.

This in-depth analysis and testing confirm that the strategic implementation of AI can dramatically shorten the licensing procedure time from months to days, while maintaining or even increasing the quality and consistency of the assessment. The potential of this approach is also directly applicable to other areas, such as the analysis of crypto-asset white papers against MiCA requirements, checking marketing communications, or monitoring the disclosure requirements of Crypto-Asset Service Providers (CASPs).

2. *Prospectus Analysis (Capital Markets)*: For capital market supervision, AI (specifically a RAG approach) was tested to extract specific information, such as collateral types or security features, from multiple securities. This facilitated statistical analysis and trend identification. The RAG system achieved high accuracy (approx. 95%) and reduced the analysis time for the sample set from an estimated 2 hours manually to 7 minutes. This demonstrates AI's utility in handling structured information within lengthy documents common in capital markets. The exploration of AI for enhancing the analytical capabilities of supervisors in reviewing such capital market documentation is a subject of international focus, with bodies like the International Organization of Securities Commissions (IOSCO) actively examining various use cases and the associated regulatory considerations and challenges [4].
3. *Investment Fund Statutes Analysis (Capital Markets)*: Another capital markets application involved extracting

investment strategies and limits from the statutes of numerous special investment funds. Manually processing these documents across hundreds of funds is extremely time-consuming. AI tools, particularly Google's NotebookLM in this test case, demonstrated a strong ability to accurately identify asset types and associated limits from the statutes. The potential time saving across all funds was estimated to be in the hundreds of hours, freeing up capacity for more intellectually demanding supervisory activities. This capability is valuable for both ongoing supervision and the initial licensing/notification process.

4. *Annual Report Verification (Capital Markets)*: AI (ChatGPT) was used to assist in the annual review of investment firms' annual reports against requirements stipulated by the Capital Market Undertakings Act (ZPKT) [5] and accounting legislation. Using a supervisory checklist converted into prompts, the AI could identify whether required information was present and provide citations. This "machine pre-processing" was estimated to save between 2 to 8 hours per report, depending on complexity, allowing supervisors to focus on addressing identified deficiencies.

Compliance and Integrity Monitoring

AI offers potential in monitoring market integrity and ensuring compliance across various domains:

1. *Comparing Public vs. Reported Information*: An AI challenge explored using AI to monitor publicly available information (company websites, annual reports, press releases) and compare it with data formally reported to the CNB. The goal was to identify discrepancies, potentially misleading statements, or unsubstantiated claims that might warrant closer supervisory scrutiny. While AI systems struggled to reliably find specific documents online automatically, once provided with the documents, analysis (e.g., comparing financial figures, assessing consistency of statements) was feasible and demonstrated potential. This aligns with supervisory goals of ensuring market transparency and integrity.
2. *Monitoring Financial Influencers (Finfluencers)*: A significant modern challenge is monitoring the rapidly growing volume of content produced by financial influencers ("finfluencers") on social media to ensure compliance with regulations like the Market Abuse Regulation (MAR) [6], particularly concerning objective investment recommendations. Manually monitoring the sheer volume of text, audio, and video content is extremely resource-intensive. CNB explored a multi-stage, AI-assisted approach:
3. *Identification*: Leveraging native platform algorithms (e.g., X, YouTube) that suggest related content to identify

potentially relevant influencers.

- *Data Extraction*: Using ChatGPT to generate Python scripts (run in CNB's Data Science Lab) for optical character recognition (OCR) and data scraping from screenshots of finfluencer profiles to collect key data (followers, post frequency, description), saving an estimated 10 hours of manual work for ~120 profiles.
 - *Content Processing*: Employing specialized AI services (like Descript) for automated speech-to-text transcription and translation (into English for better AI analysis) of audio/video content, tackling the challenge of analyzing hours of podcasts or videos and saving an estimated 70 hours compared to manual processing.
 - *Content Analysis*: Deploying a custom-built ChatGPT assistant ("Investment Recommendation Checker"), trained using inductive learning on ESMA Warning on posting Investment Recommendation on social media [7] and relevant RTS, to analyze the transcribed/translated content and detect potential direct or indirect investment recommendations requiring further supervisory review. This involved iterative tuning prioritizing sensitivity, combined with human validation ("human-in-the-loop").
 - This comprehensive approach enabled the monitoring of a previously intractable area, yielding total estimated time savings of around 100 hours (12.5 man-days) for the pilot project.
4. *Automating Sanctions List Screening*: Ensuring compliance with international sanctions is crucial. AI was used to automate the process of checking CNB clients against multiple international sanctions lists (OFAC, EU, UN, Czech DoS). ChatGPT generated VBA code for MS Excel that automates the comparison of client names (text strings) against downloaded sanctions lists and generates a compliance report. This replaces a highly manual, repetitive, and error-prone process (e.g., Ctrl+F searching across multiple PDF/Office files). The automated solution performs checks in seconds, regardless of the number of clients, offering significant efficiency gains (estimated 1 minute per client per list manually, potentially days for large lists) and providing a better basis for audit trails. A prerequisite for full automation is resolving network limitations on downloading external lists directly within the CNB environment.

Supporting Tools and General Efficiency Gains

Beyond specific supervisory tasks, AI provides general support:

1. *Summarization*: AI tools efficiently summarize long regulatory texts, research papers, or financial stability reports from other jurisdictions, enabling quicker

comprehension.

2. *Drafting and Editing*: Assisting staff in drafting various documents, including reports, official correspondence, and even blog posts, improving clarity, style, and grammar.
3. *Coding Assistance (SQL, R, VBA)*: Generating, optimizing, and debugging code used for data analysis and automation supporting supervisory tasks. This includes generating VBA macros for data processing (e.g., data formatting, sanctions checks) and assisting with Python code for test automation frameworks used for core systems, thereby indirectly supporting the reliability of systems used in supervision.
4. *Knowledge Enhancement & Search*: Acting as an efficient search tool for finding specific information within regulations or technical documentation, potentially faster than traditional search engines for complex queries.
5. *Related Application (Professional Competence)*: AI was successfully used to automate the validation of test questions for financial market professionals against updated legislation, saving an estimated 142 man-hours in one instance. This supports the overall quality of market participants.

AI TOOLS AND APPROACHES EXPLORED

The CNB's exploration involved several types of AI tools and technological approaches:

General Large Language Models (LLMs)

Commercially available LLMs like OpenAI's ChatGPT (including the Team version allowing custom GPT creation) and Microsoft's Copilot were frequently used.

- **Strengths**: Versatile for text generation, summarization, translation, answering general questions, basic document analysis, and coding assistance. These very tools formed the technological basis of the first "platforms" tested (see Figure 1 and 2) for the semi-automated analysis of licensing documentation, where the expert utilized their general capabilities for discrete, manually assigned tasks.
- **Weaknesses**: Can struggle with very large documents or comparing information across multiple files due to context window limitations. Accuracy requires careful prompting and validation. Difficulty reliably finding specific, up-to-date documents online.

Retrieval-Augmented Generation (RAG)

This technique enhances LLMs by first retrieving relevant information chunks from a specific set of documents and then feeding this context to the LLM for answer.

- **Strengths**: More effective than general LLMs for extracting specific, structured data points from large documents (e.g., prospectuses). Can improve accuracy by

grounding responses in provided texts.

- **Weaknesses:** Typically requires more technical setup and configuration, potentially involving coding (Python mentioned) or specialized platforms. Performance depends on the effectiveness of the initial retrieval step.

Specialized AI Tools

Tools designed for specific tasks were also encountered or utilized:

- *Google's NotebookLM:* Showed particular promise in analyzing investment fund statutes, potentially due to better handling of multiple sources or a larger effective context window, perhaps incorporating RAG techniques internally.
- *Speech-to-Text / Translation Services:* Essential for processing audio/video content from sources like YouTube or podcasts, enabling subsequent text-based AI analysis. Used effectively in the finfluencer monitoring project.
- *Potential for Legal AI:* The need for specialized tools, e.g., for legal research (case law analysis), was acknowledged, though not extensively tested in the provided documents.
- *Platform Recommendation Algorithms:* Utilized as a first step in identifying finfluencers by leveraging existing content-based recommendation systems on platforms like X and YouTube.

Cloud Platforms and Internal Environments (Microsoft Azure / M365 / DS Lab)

The CNB's IT strategy indicates a preference for leveraging the Microsoft cloud ecosystem (Azure, M365, Copilot suite) for AI deployment.

- **Strengths:** Offers an integrated stack of tools – from basic Copilot Chat, through low-code agent building with Copilot Studio, to advanced custom solutions via Azure AI services and LLM APIs. Potential for enhanced security and compliance within a managed enterprise cloud environment, facilitating work with sensitive data. Integration with existing M365 tools (like SharePoint, Teams) could streamline workflows. Additionally, CNB's internal Data Science Lab (DS Lab) provides an environment for running AI-related code (e.g., Python scripts generated by ChatGPT for OCR/scraping) in a controlled setting.
- **Weaknesses:** Creates vendor dependency. Requires investment in licenses and potentially pay-as-you-go consumption costs. Requires building expertise in the Microsoft cloud platform. Full security benefits depend on implementing features like data classification and DLP, which may take time.

Custom AI Agents / GPTs

A recurring approach involved creating custom-configured AI assistants (often within ChatGPT's framework) tailored to specific tasks. Examples include the "Investment Recommendation Checker" trained on MAR/ESMA guidelines, assistants for summarizing Financial Stability Reports, assisting with SQL code generation, and generating specific VBA code for sanctions checks. This allows for embedding specific instructions, knowledge sources, and desired output formats, enhancing relevance and consistency for defined tasks.

Within the scope of the licensing documentation analysis project (see II.A.1), this approach was developed further into a more sophisticated, custom-built solution based on API integration. This solution is being implemented on a platform built with Google Agentspace technology, which proved to be more flexible and suitable during testing than initially considered alternatives like Copilot Studio. Unlike simple custom agents, this system does not operate merely with files and a general prompt; it requires highly structured inputs, such as legislative requirements and, crucially, a set of check questions in a machine-readable format (JSON), as depicted in Figure 3. This model implements the "human-in-the-loop" principle, where the AI performs the primary, time-consuming assessment and generates a structured preliminary output (the "AI assessment"), which is then passed to a human expert for targeted and rapid validation. A fundamental component of this approach is also the internal feedback loop, where insights from the human validation are used to refine the input questions, thereby allowing the system to learn and improve iteratively. This model thus represents a significant shift from using AI as a general assistant to deploying it as an integrated, process-driven tool for a specific and complex supervisory task.

IDENTIFIED LIMITATIONS AND RISKS

Despite the potential benefits, the CNB's exploration highlighted significant challenges and risks associated with using AI in supervision:

Accuracy and Reliability

AI models, particularly LLMs, are prone to generating incorrect or nonsensical information ("hallucinations"). Errors were observed in tasks ranging from extracting specific figures and identifying document locations to correctly interpreting complex rules like macroprudential policy settings. This necessitates rigorous validation of all AI-generated outputs by subject matter experts. AI should be viewed as an assistant providing a first draft or hypothesis, not a definitive answer.

Our own testing within the licensing documentation analysis project (see II.A.1) quantitatively confirms these findings. As shown in Figure 1, the introduction of the first AI platform led to a paradoxical increase in the overall error rate, which was a direct consequence of the technology's immaturity at the time. Although the error rate decreased with subsequent iterations, it must be emphasized that the near-zero error rate of platform 4 was also achieved thanks to a hidden human review — the expert subconsciously corrected minor errors while manually transferring the results. The slight increase in the error rate with the fully automated API solution clearly demonstrates the risk that eliminating even these seemingly minor human interventions can negatively impact the reliability of the final output. This finding underscores that even with the most advanced systems, the role of targeted expert validation is absolutely crucial.

Prompt Engineering

The usefulness of AI tools is highly sensitive to the way questions or instructions (prompts) are formulated. Crafting effective prompts that yield accurate and relevant results requires skill, domain knowledge, and often an iterative process of refinement. Poorly designed prompts lead to poor outputs.

Experience from our project shows that the solution to this challenge may not only be the formulation of better natural language prompts but a shift towards more systematic and structured inputs. While the initial phases of our testing struggled with issues arising from suboptimally designed prompts, the final API-based solution (see Figure 3) circumvents this problem by requiring a set of precisely defined questions in a machine-readable JSON format as a key input. This approach not only minimizes ambiguity but also allows for the systematic and iterative improvement of the inputs based on feedback from the validation process.

Data Security and Confidentiality

This remains arguably the most significant barrier to widespread AI adoption in supervision. Using external, cloud-based AI tools with sensitive, non-public supervisory data carries substantial risks of data breaches, unauthorized access, or misuse of data (e.g., for training third-party models). Current CNB rules and the lack of established secure frameworks limit the use of AI with internal documents. While anonymization was attempted in some tests, it is often impractical for complex analyses requiring interlinked information and does not eliminate all risks. Secure, controlled environments — potentially a well-configured private cloud instance within a trusted provider like Microsoft Azure — are seen as necessary prerequisites for broader use with sensitive data.

Technological Limitations

Current AI tools have inherent limitations:

- *Context Windows:* LLMs can only process a limited amount of information at once. In our practice, this limit manifested as a critical barrier, where early platform versions were unable to process complete licensing documentation consisting of hundreds of files. As shown in Figure 1, we were forced to circumvent this problem through time-consuming manual document merging or by limiting the scope of the analysis. Only the latest API-based solution is able to overcome this barrier and natively work with the full scope of the documentation.
- *Online Information Retrieval:* AI struggles to reliably find specific, up-to-date documents directly from the web.
- *Complex Reasoning:* Tasks requiring deep, multi-step reasoning or complex comparisons across disparate sources can exceed current capabilities.
- *Data Formats:* AI primarily excels with text; handling diverse data formats, especially non-standardized graphical or tabular data embedded in PDFs, remains challenging. This was noted as a limitation in monitoring open sources for crisis management, where diverse graphical formats hindered analysis.
- *Network / Integration Barriers:* Practical implementation can be hindered by internal network security policies restricting direct access to external resources needed by AI tools, such as automatically downloading updated sanctions lists or requiring complex troubleshooting for integrating tools on servers without internet access.
- *Language Dependence:* While powerful, LLMs often perform best in English due to the preponderance of English training data. For tasks requiring nuanced understanding (like detecting investment recommendations), translating source material (e.g., Czech influencer videos) into English before AI analysis was found to yield better results, adding a process step.

E. Resource Requirements

Effective AI implementation requires significant investment:

- *Skills:* Need for personnel skilled in AI, data science, prompt engineering, IT infrastructure management (especially cloud), and potentially AI-specific programming.
- *Technology:* Costs associated with licenses for commercial AI tools, cloud consumption (which can be unpredictable with pay-as-you-go models), and potentially hardware if considering on-premise options (though CNB strategy leans against this).
- *Time:* Significant time investment is needed for testing, prompt optimization, validation of results, and training

staff.

F. Regulatory, Ethical, and Other Risks

- *Compliance*: Ensuring AI use complies with emerging regulations like the EU AI Act [8] and existing ones like GDPR [9].
- *Over-reliance / Human Factor*: Risk of supervisors becoming overly dependent on AI tools, potentially leading to a decline in critical thinking and analytical skills ("deskilling" or "laziness"). It is for this very reason that our target process (Figure 3) was designed for the AI to serve as a tool for automating the routine and time-consuming part of the analysis, while leaving the final, critical judgment fully in the hands of the expert. The targeted human validation focuses on checking references, factual accuracy, and content completeness — areas that require deep domain knowledge that AI still lacks. Our results show that this "human-in-the-loop" approach is not just a safeguard but an absolutely essential component of a reliable and defensible supervisory process. The need for human validation ("human-in-the-loop") was explicitly highlighted in the finfluencer monitoring workflow to ensure accuracy despite prioritizing sensitivity in the AI model.
- *Vendor Lock-in*: Relying heavily on a single provider's ecosystem (e.g., Microsoft) creates dependency and potential future cost risks.
- *Rapid Evolution / Version Dependence*: The AI field changes extremely quickly, making long-term planning difficult and requiring continuous adaptation and learning. Experiences showed significant performance differences between AI model versions (e.g., ChatGPT 4 vs. 3.5), requiring users to adapt to updates and potentially revisit previous work.

CONCLUSION

The Czech National Bank's exploratory projects confirm that Artificial Intelligence holds significant potential to augment the capabilities of financial market supervisors. Our experiments have demonstrated tangible benefits across several areas, from monitoring market integrity and screening sanctions lists to supporting analytical tasks. However, the most conclusive results emerged from an in-depth, evolutionary project focused on automating the licensing documentation assessment process in the context of MiCA. This project revealed that the strategic deployment of AI can reduce the processing time for a single complex application from months (an average of 80 working days) to a matter of days, representing an order-of-magnitude increase in efficiency.

The path to this outcome, however, was not straightforward and clearly exposed key challenges and risks.

Our data (Figure 1) showed that the initial deployment of AI can paradoxically lead to a temporary increase in the error rate and confirmed that the reliability of outputs depends not only on the model's capabilities but also on the quality of the inputs. This is precisely why a shift from general prompts to highly structured, machine-readable inputs (e.g., in JSON format)—which minimize ambiguity and allow for systematic improvement—proved to be a critical solution. It was also confirmed that the security of sensitive data remains the biggest barrier to widespread adoption and that technical limitations, such as the context window size, posed a fundamental obstacle in the project's early stages. Indeed, these findings are consistent with the broader academic and regulatory discourse, which consistently highlights significant regulatory and implementation hurdles in the integration of AI in financial services [10].

Our findings underscore that the successful implementation of AI in supervisory practice is not merely about acquiring a tool but about designing robust, process-driven systems. The target model we have designed and are currently testing (Figure 3) is based on synergy, not replacement. It implements the "human-in-the-loop" principle as an integral component, where the AI performs the primary, scalable analysis, and the expert focuses on targeted validation in areas requiring critical thinking and deep knowledge. Furthermore, this model includes an internal feedback loop, ensuring that the system iteratively learns and refines its accuracy with each processed case.

Moving forward, the CNB will therefore focus on developing and deploying such custom-built, integrated solutions, likely within a secure cloud environment (Microsoft Azure/M365/Google Agentspace). The key will not only be the further development of technological infrastructure but, crucially, the systematic cultivation of internal skills in areas such as process analysis for AI, structured input design, and the critical evaluation of outputs.

Ultimately, our experience confirms that Artificial Intelligence is becoming an invaluable assistant in the supervisor's toolkit, capable of handling scale and routine tasks to an extent previously unimaginable. However, it cannot, and should not, replace the nuanced judgment, ethical considerations, and ultimate responsibility of the human expert. Success in harnessing the potential of AI will depend on our ability to design and implement solutions that are not only technologically advanced but, above all, strategic, process-oriented, and fully risk-aware.

REFERENCES

- [1]. European Parliament and the Council of the EU, „Regulation (EU) 2023/1114 of the European Parliament and of the Council of 31 May 2023 on markets in crypto-assets and amending Regulations (EU) No 1093/2010 and (EU) No

- 1095/2010 and Directives 2013/36/EU and (EU) 2019/1937 (MiCA)“, 9 June 2023 (published in the Official Journal of the EU).
- [2]. European Parliament and the Council of the EU, „Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector and amending Regulations (EC) No 1060/2009, (EU) No 648/2012, (EU) No 600/2014, (EU) No 909/2014 and (EU) 2016/1011 (DORA)“, 27 December 2022 (published in the Official Journal of the EU).
- [3]. N. Malali, „The Role of Retrieval-Augmented Generation (RAG) in Financial Document Processing: Automating Compliance and Reporting“, *International Journal of Management Technology*, vol. 12, no. 3, pp. 26-46, Apr 2025.
- [4]. The Board of the International Organization of Securities Commissions, „Artificial intelligence in capital markets: Use cases, risks, and challenges“, *IOSCO Consultation Report*, CR/01/2025, Mar 2025.
- [5]. Parliament of the Czech Republic, „Act No. 256/2004 Coll., on Capital Market Undertakings (ZPKT)“, 22 April 2004 (date of initial promulgation; the Act has been continuously amended).
- [6]. European Parliament and the Council of the EU, „Regulation (EU) No 596/2014 of the European Parliament and of the Council of 16 April 2014 on market abuse (Market Abuse Regulation) and repealing Directive 2003/6/EC of the European Parliament and of the Council and Commission Directives 2003/124/EC, 2003/125/EC and 2004/72/EC (MAR)“, 12 June 2014 (published in the Official Journal of the EU).
- [7]. European Securities and Markets Authority (ESMA), „Warning on posting Investment Recommendation on social media“, 6 February 2023.
- [8]. European Parliament and the Council of the EU, „Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act)“, 13 June 2024 (published in the Official Journal of the EU).
- [9]. European Parliament and the Council of the EU, „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 – GDPR)“, 4 May 2016 (published in the Official Journal of the EU).
- [10]. D. B. Vuković, S. Dekpo-Adza, and S. Matović, „AI integration in financial services: a systematic review of trends and regulatory challenges“, *Humanities and Social Sciences Communications*, vol. 12, no. 1, pp. 1-29, Apr 2025.