An Early Warning System for Fintech Supervision

FIRSA Working Paper No. 2025/01

Alfredo Hernández Sánchez, *PhD*

Vilnius University

October 2025

Understanding FinTech Regulatory Sandbox Development in Europe (FIRSA)

Website: https://firsa.eu

Funded by the European Union

VU TSPMI

FIRSA Working Papers

This working paper series publishes early research results and policy papers connected to the EU Horizon Marie Skłodowska-Curie Actions ERA Fellowship *FIRSA*. For any comments, please contact:

Alfredo Hernández Sánchez
Institute of International Relations and Political Science
Vilnius University
Vilnius, Lithuania
alfredo.sanchez@tsmpi.vu.lt

Acknowledgements

This paper was produced as part of the course *Practical Data Science in Financial Supervision* hosted by the Cambridge Centre for Alternative Finance. The author would like to thank Rweyemamu Barongo from the AI and Data Innovation Hub of the Bank of Tanzania for helpful comments and suggestions. All errors are my own.

Disclaimer

FIRSA is funded by the European Union under grant agreement No. 101180601 titled: *Understanding FinTech Regulatory Sandbox Development in Europe*. Views expressed are those of the author(s) and do not necessarily reflect those of Vilnius University or the European Union.

Table of contents

Early Warning Proposal

This proposal is directed towards stakeholders in Fintech regulatory sandboxes and innovation hubs. These innovation facilitators are commonly adopted by many jurisdictions wishing to promote financial innovation whilst safeguarding financial stability and consumer protection (Raudla et al. 2025, 2024; Philippon 2017; Cornelli et al. 2024).

Regulatory sandboxes create structured environments where firms can experiment with new financial products and services under regulatory observation. However, as innovation accelerates, many supervisors lack the tools to maintain real-time oversight over new actors entering the ecosystem.

To address this challenge, I propose a system that enables regulators to track, map, and classify emerging Fintech firms and products through a live, automated workflow. Using text scraped from various publicly available sources —such as business directories, news media, and social platforms— the system processes entries to extract firm names, categorize products, assess sentiment, and determine jurisdictional relevance. This early detection allows regulators to engage with firms proactively, offering guidance where needed and enabling oversight without stifling innovation.

Expected Impacts and Benefits

Fintech innovation facilitators currently operate with a limited capacity to detect new entrants until those firms actively seek engagement (e.g. license applications) or until they contravene public interest in a notorious way (e.g. fraudulent behavior). This creates a gap in the regulatory pipeline: firms and products may begin operating or gaining traction well before a supervisor becomes aware of them. The baseline remains reactive, leaving innovation hubs exposed to reputational and operational risks.

This proposal explores how textual data from online sources can be used to identify new Fintech firms, assess risk signals, and support early regulatory engagement. By transforming unstructured information into structured, triage-ready signals, the system offers supervisors a way to understand what is emerging in the fintech landscape, who is behind it, and whether their activity may require contact or support. It is particularly useful in dynamic sectors such as crypto, P2P lending, and crowdfunding—where risks can scale quickly if left unchecked.

This project aligns with the strategic goals of innovation facilitators to foster safe experimentation while minimizing exposure to regulatory blind spots. The system provides supervisors with a practical means of maintaining a live view of market developments, helping them identify where risks may emerge and where firms might need outreach or support—even before those firms realize it themselves.

Rather than relying on passive reporting, the system actively monitors developments across a wide spectrum of information sources. While some of these are open, others require specific access or compliance with terms of service, which is acknowledged in the system's flexible and modular design. The intent is not just to improve oversight but to improve communication, ensuring firms are not unnecessarily penalized, and foster a better guidance of responsible financial innovation.

Objectives of the Project

The objective of the project is to build a scalable, data-driven mechanism that helps regulators detect and categorize new fintech firms and products, especially those operating near the edges of formal oversight. The current prototype demonstrates this through a simulated pipeline that collects short-form text, extracts key data points (such as firm name, product type, risk keywords, and sentiment), and evaluates entries against basic thresholds to determine if contact is warranted.

The ultimate vision is to deploy this as a user-facing dashboard, where regulators can intuitively explore flagged entries, view emerging trends, and adjust detection parameters over time. The prototype shows how core functions can be performed in an interpretable and modular way. Regulators will be able to tweak parameters inside the proposed dashboard in order to better identify firms that need early guidance or risk management assistance. While still in proof-of-concept form, the system is designed to evolve into an institutional tool that assists supervisors in making informed, timely decisions.

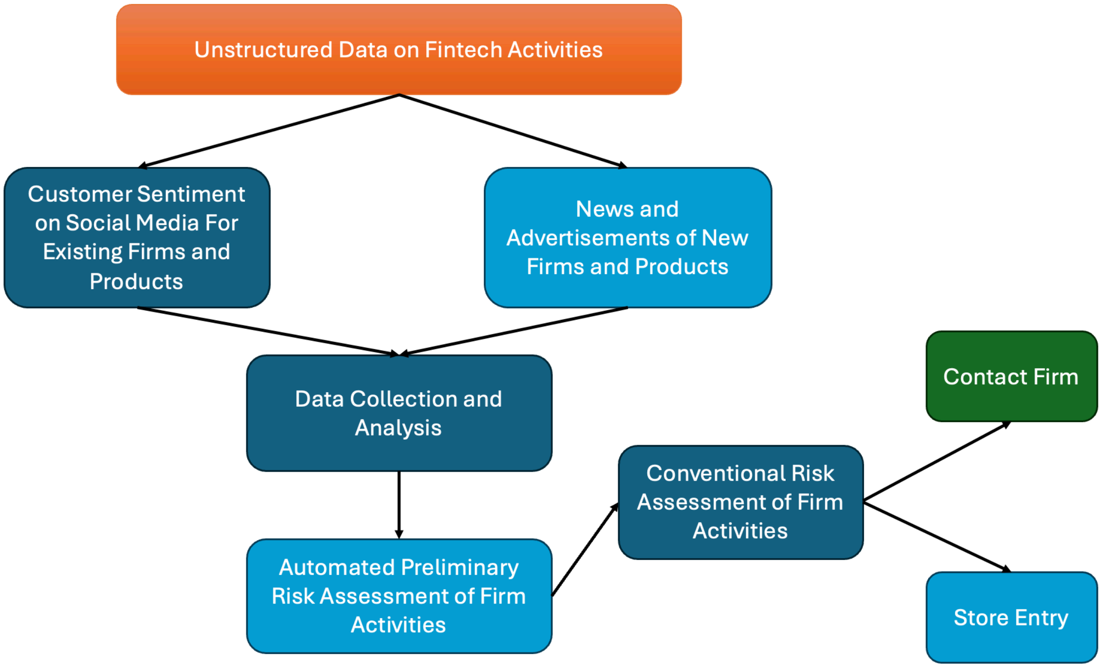

Figure 1: Flowchart of the proposed solution.

The system functions by scanning externally available text entries and processing them through an analytical pipeline. It identifies the main firm referenced, classifies the associated product or service, assigns a risk level based on keyword patterns, and flags entries with strong negative sentiment or jurisdictional relevance. The result is a structured registry of emerging actors and activities.

Rather than rely on heavy infrastructure or black-box models (LLMs), the solution prioritizes interpretability and modularity. All data is processed into a clean, tabular format that could be served to regulators through a simple web-based dashboard. Each record in the system is actionable: it tells the supervisor who is involved, what kind of financial product they are working on, and whether the system suggests further attention. The approach is designed for gradual scale-up, with new data sources and refinements easily added as capacity and need grow.

Risk Assessment

Several risks must be addressed in deploying this solution. First, sourcing data from scraped or API-based platforms requires adherence to legal and ethical standards, including data minimization and compliance with terms of service. Any operational version of the system must be built around responsible sourcing practices and, where necessary, obtain formal data access agreements.

There is also a risk of false positives or classification errors. Because the system uses lightweight NLP models to extract names and assign categories, it may occasionally mislabel firms or products. This can be mitigated by updating keyword libraries, allowing human review of flagged entries, and incorporating supervisor feedback directly into the model refinement process.

A further risk is redundant contact as regulators may reach out to firms already under review or in communication. The system must therefore be linked to internal firm registries to avoid unnecessary outreach. Finally, though not yet implemented, language limitations may become a concern in jurisdictions like Lithuania where content may appear in both English and Lithuanian. Multilingual adaptation will be a necessary future step to ensure the tool's effectiveness across linguistic boundaries.

Insight-Driven Strategy

The prototype developed demonstrates the technical feasibility of building an automated early warning system tailored to fintech innovation hubs. The most immediate output is a structured, filtered registry of firms and products that can support prioritization and outreach. The longer-term goal is a full-featured dashboard – built with tools like Streamlit – where supervisors can explore live data, adjust thresholds, and export findings for internal use. This dashboard would act as an operational interface for a living dataset that evolves with the ecosystem it monitors.

Next steps include refining the classification logic, testing the tool against more jurisdiction-specific data, and connecting it to internal supervisory systems. Key performance indicators for success will include the number of relevant entries flagged that would otherwise go undetected, positive feedback from supervisory users, and the tool's capacity to incorporate new categories and sources over time. The system is designed to grow, adapt, and ultimately strengthen the supervisory function of innovation hubs by offering earlier, more meaningful insight into where guidance is most needed.

Early Warning Prototype

This notebook is a prototype for a fintech early warning system designed to assist regulators in identifying and monitoring emerging fintech activity in their jurisdiction. It simulates a workflow where data is collected from publicly available sources (e.g., Crunchbase, news, social media), processed using natural language techniques, and assessed for potential supervisory concern.

The workflow below:

- Collects data from example fintech-related headlines or announcements
- Extracts named entities (firms, founders, locations) using spaCy
- Identifies main organizations that are central to the announcement
- Extracts keywords and assigns a simple risk score based on suspicious or relevant terms
- Performs sentiment analysis using FinBERT to detect negative language
- Classifies topics with zero-shot learning into fintech categories (e.g., crypto, p2p lending)
- Flags entries that are within jurisdiction (e.g., Lithuania, Vilnius)

- Advises regulators to contact firms if risk or sentiment exceed certain thresholds

Environment Setup

This prototype serves as a foundation for developing a live dashboard or automated alert system that can be used by regulators or innovation hubs to reach out to innovative firms that may incur in risky activities.

```
# Set up environment
!pip install keybert
import pandas as pd
import re
import spacy
from keybert import KeyBERT
import requests
from bs4 import BeautifulSoup
from transformers import pipeline

# Load NLP models
nlp = spacy.load("en_core_web_sm")
kw_model = KeyBERT()

# Load zero-shot classification pipeline
zshot_classifier = pipeline("zero-shot-classification",
model="facebook/bart-large-mnli")

# Load sentiment analysis pipeline using FinBERT
sentiment_pipeline = pipeline("sentiment-analysis",
model="ProsusAI/finbert")
```

The next step is defining the labels used for the zero-shot classification based on fintech-related activity. Keywords that may indicate risky activities should also be defined.

```
# Define fintech labels
fintech_labels = [
    "other",
    "payment systems",
    "peer to peer lending",
    "crowdfunding",
    "decentralized finance (defi)",
    "digital assets",
    "cryptocurrencies",
    "traditional finance"
]

# Define risk keywords
risk_keys = [
    "crypto",
    "KYC",
    "unregulated",
    "yield",
    "token",
    "startup"
]
```

The code below defines the classification functions for the dictionary and classification pipelines.

```
# Function to classify text into fintech categories
def classify_fintech_category(text):
    result = zshot_classifier(text[:512], fintech_labels)
    return result["labels"][0], result["scores"][0]

# Function to classify sentiment
def classify_sentiment(text):
    result = sentiment_pipeline(text[:512])[0]
    return result["label"], result["score"]
```

Data Collection

The next step is to collect data. Below is sample data from multiple sources, including simulated Crunchbase entries and scraped news headlines. The pipeline then searches for named entities using spaCy, identifying organizations, people, and locations. From these, the main organization mentioned in each text is selected.

```
# Function to simulate Crunchbase-style fintech data
# NOTE: Real Crunchbase API requires an API key (not free...)
def fetch_crunchbase_simulation():
    return [
        {"source": "Crunchbase", "text":
          "Fintech startup PayFlex raises Series A to expand AI-driven payment systems"},
        {"source": "Crunchbase", "text":
          "BlockChainPay enters Lithuanian market with unregulated crypto lending"},
        {"source": "Crunchbase", "text":
          "NeoInsure launches embedded insurance products targeting neobanks"}
    ]

# Function to scrape news headlines from Finextra (example of 10 only from landing page)
def fetch_news_scrape():
    url = "https://www.finextra.com/news/latestnews.aspx"
    headers = {"User-Agent": "Contact: regulator@agency.com"} # User agent contact details
    response = requests.get(url, headers=headers)
    soup = BeautifulSoup(response.content, "html.parser")
    cards = soup.select(".h3")
    headlines = []
    for card in cards[:10]:
        headline_tag = card.find("a")
        if headline_tag:
            headlines.append({"source": "Finextra", "text": headline_tag.get_text(strip=True)})
    return headlines

# Example master data, this will be replaced by a live document.
static_data = [
    {"source": "Twitter", "text":
      "New crypto wallet launched by FinTrust Labs with zero KYC in Kaunas"},
```

```
    {"source": "News", "text":
      "Neobank startup GreenPay enters Lithuania with digital lending solution"},
    {"source": "LinkedIn", "text":
      "Jane Doe, founder of StableYield, joins Fintech Europe accelerator"}
]

# Combine master data with simulated Crunchbase entries and scraped news (update)
data_update = static_data + fetch_crunchbase_simulation() + fetch_news_scrape()

# Execute warning pipeline on new object DF
df = pd.DataFrame(data_update)
```

Running the Analysis

The code below classifies each entry into a fintech category using zero-shot learning and run sentiment analysis with FinBERT to detect negative tone. It also checks for jurisdictional relevance by scanning for Lithuanian-related terms in both text and entities.

```
# Function to extract named entities
def extract_entities(text):
    doc = nlp(text)
    return [(ent.text, ent.label_) for ent in doc.ents if ent.label_ in ["ORG", "PERSON", "GPE"]]

# Function to extract keywords and risk score (simple version)
def extract_keywords_and_score(text):
    keywords = kw_model.extract_keywords(text, top_n=3)
    keyword_list = [kw[0] for kw in keywords]
    # Simulated basic risk scoring
    risk_keywords = risk_keys
    score = sum(1 for word in keyword_list if word.lower() in risk_keywords)
    return keyword_list, score

# Define the extract_main_organization function (longest)
def extract_main_organization(entities):
```

```
    orgs = [ent for ent, label in entities if label == "ORG"]
    if not orgs:
        return None
    return max(orgs, key=len)

# Process each entry
df["entities"] = df["text"].apply(extract_entities)
df["keywords"], df["risk_score"] = zip(*df["text"].apply(extract_keywords_and_score))
df["risk_level"] = df["risk_score"].apply(lambda x:
  "high" if x >= 2 else ("medium" if x == 1 else "low"))
df["main_entity"] = df["entities"].apply(extract_main_organization)

# Apply zero-shot classification for fintech category
df["fintech_category"], df["category_confidence"] = zip(*df["text"].apply(classify_fintech_category))

# Apply sentiment analysis
df["sentiment_label"], df["sentiment_score"] = zip(*df["text"].apply(classify_sentiment))

# Create jurisdiction indicator
df["in_jurisdiction"] = df.apply(
    lambda row: int(
        any(re.search(r"lithua|vilniu|kaunas|klaipeda",
        ent[0], re.IGNORECASE) for ent in row["entities"])
        or bool(re.search(r"lithua|vilniu|kaunas|klaipeda",
        row["text"], re.IGNORECASE))
    ), axis=1
)

# Create firm contact flag (updated logic: OR condition)
df["contact_firm"] = df.apply(lambda row: int(
    (row["sentiment_label"] == "negative" and row["sentiment_score"] > 0.75)
    or row["risk_level"] in ["medium", "high"]), axis=1)
```

Inspecting the Results

In the final stage, firms are flagged for potential regulatory contact (e.g. a formal information requirement) if the sentiment is strongly negative or if the risk level is medium or high. The notebook ends by listing the firms that meet these conditions. A live version of this prototype can be found in this Google Colab Project.

```
# Output contact recommendations
contact_list = df[df["contact_firm"] == 1]["main_entity"].dropna().unique().tolist()
if contact_list:
    print("\nRegulatory Advisory: Please consider contacting the following firms based on risk
     ↪  indicators:")
    for firm in contact_list:
        print(f"- {firm}")
else:
    print("\nNo firms require contact at this time based on current risk thresholds.")

# Select main columns relevant to the project
main_columns = [
    "source",
    "text",
    "main_entity",
    "fintech_category",
    "category_confidence",
    "risk_level",
    "sentiment_label",
    "sentiment_score",
    "in_jurisdiction",
    "contact_firm"
]

# Display the table
df[main_columns].head(20)
```

References

Cornelli, Giulio, Sebastian Doerr, Leonardo Gambacorta, and Ouarda Merrouche. 2024. "Regulatory Sandboxes and Fintech Funding: Evidence from the UK." Edited by Christine Parlour. *Review of Finance* 28 (1): 203–33. https://doi.org/10.1093/rof/rfad017.

Philippon, Thomas. 2017. "The FinTech Opportunity." https://www.bis.org/publ/work655.pdf.

Raudla, Ringa, Egert Juuse, Vytautas Kuokstis, Aleksandrs Cepilovs, and James W. Douglas. 2024. "Regulatory Sandboxes and Innovation Hubs for FinTech: Experiences of the Baltic States." *European Journal of Law and Economics*, December. https://doi.org/10.1007/s10657-024-09830-y.

Raudla, Ringa, Egert Juuse, Vytautas Kuokštis, Aleksandrs Cepilovs, Vytenis Cipinys, and Matti Ylönen. 2025. "To Sandbox or Not to Sandbox? Diverging Strategies of Regulatory Responses to FinTech." *Regulation & Governance* 19 (3): 917–32. https://doi.org/10.1111/rego.12630.