INTERNATIONAL JOURNAL OF PUBLIC ADMINISTRATION,
MANAGEMENT AND ECONOMIC DEVELOPMENT

The use of AI in market news sentiment analysis and its impact on investment decisions: Evidence from the Polish sWIG80 index in the first quarter of 2026

Diana Kral-Grabka
University of Economics in Katowice, 50 1 Maja Street, 40-287 Katowice, Poland, E-mail: diana.kral-grabka@uekat.pl, ORCID: 0000-0001-5568-6335

Krzysztof Podgórski
University of Economics in Katowice, 50 1 Maja Street, 40-287 Katowice, Poland, E-mail: krzysztof.podgorski@uekat.pl, ORCID: 0000-0002-7549-1997

Abstract

The aim of this study is to analyse the impact of sentiment surrounding extreme geopolitical news on the behaviour of the Polish small-cap index sWIG80 in the first quarter of 2026. This study fills a research gap concerning the reaction of the small-cap segment, dominated by emotionally driven retail investors, by testing the hypothesis of an immediate sell-off in the index in response to extremely negative news shocks. The research process employed a classical event analysis methodology for a narrow, symmetric window [-1; +1]. To assess the sentiment of 1,098 headlines sourced from Google News channels, the Gemini Pro large language model (LLM) was used in a zero-shot architecture, foregoing traditional dictionary-based methods. Subsequently, 13 extremely negative trading days were identified, constituting the lower quintile of the calculated net sentiment index. Market reaction was measured using average (AAR) and cumulative (CAAR) excess returns, verified using Welch's t-test. The analyses conducted showed that the research hypothesis was not confirmed. The cumulative average abnormal return (CAAR) in the three-day event window was +0.13% with a complete lack of statistical significance (p=0.851), whilst on the day of the shock itself (t=0) a negligible change in AAR of -0.04% was recorded. It has thus been demonstrated that sWIG80 share prices proved highly resilient to superficial geopolitical noise. This study provides unique empirical evidence for the functioning of the semi-strong form of the efficient market hypothesis in the domestic small-cap segment and highlights the clear advantage of contextual inference by LLM models over lexical methods. It is also recommended that future research include an analysis of intraday quotes to examine intraday dynamics and the implementation of domain-specific models, such as Polish RoBERTa.

Keywords: sentiment analysis; artificial intelligence; sWIG80; event analysis; efficient market hypothesis

JEL classification: G14; G41; C34

1. Introduction

Today's financial markets are undergoing advanced digitalisation, with artificial intelligence (AI) and high-frequency trading (HFT) being the main drivers of change. In an era of growing importance of Big Data, relying solely on traditional technical and fundamental measures no longer guarantees a long-term investment advantage. Instead, the focus is shifting towards the exploration of unstructured data, with particular emphasis on textual market information (Seidler-de Alwis, 2018). As the latest literature on the subject demonstrates, machine learning algorithms are currently used not only to estimate future prices, but primarily for the automated, split-second assessment of sentiment from a massive influx of news reports (Tkachyk and Kvyetnyy, 2022). The full automation of this mechanism enables the ability to react instantly within intraday trading, which consequently redefines the structure of liquidity and price volatility on modern stock exchanges (Anese et al., 2023).

The existing literature on the subject provides mixed evidence regarding the effectiveness of sentiment analysis in forecasting market behaviour, particularly with regard to specific emerging markets. Despite the existence of numerous global studies (Ranjbaran et al., 2024), the Polish capital market remains an area with a clear dichotomy in results. On the one hand, studies on the WIG-banks and WIG-energy sub-indices have shown that press headlines lack significant explanatory power for intraday price movements (Polak, 2021; Wojarnik, 2022). On the other hand, the integration of text data from Monetary Policy Council statements allowed for a statistically significant reduction in the forecast error of the WIG20 index returns by 2.1% (Kostyra, 2025). The Polish RoBERTa model is currently considered the most effective architecture for classifying sentiment in Polish-language financial texts (Koptyra et al., 2024).

However, a key research gap in the current state of knowledge remains the complete lack of studies analysing the impact of artificial intelligence on the small-cap segment, such as the sWIG80 index. This is a critical point, as the structure of this index, dominated by retail investors who are highly reactive to emotional stimuli, makes it significantly more susceptible to information shocks than the largest companies in the WIG20. Furthermore, the phenomenon of delayed price discovery in the small-cap segment suggests that the market's absorption of geopolitical news may take longer, creating exploitable market anomalies. Filling this gap by examining the sWIG80's reaction to extreme geopolitical information shocks is essential for understanding the mechanisms by which market sentiment propagates under conditions of heightened volatility.

The aim of this article is to analyse the relationship between geopolitical news sentiment and the behaviour of the sWIG80 index in the first quarter of 2026. The study seeks to answer the question of whether AI models are capable of effectively translating complex and multi-faceted geopolitical messages into measurable market signals within short time windows.

Based on the principles of behavioural finance and algorithmic trading mechanisms, a research hypothesis was formulated, positing that: Geopolitical news with strongly

negative sentiment triggers an immediate sell-off in the sWIG80 index, anticipating the reaction of retail investors.

The thesis is divided into seven parts. Following this introduction, which provides an overview of the impact of sentiment on financial markets, Chapter Two presents a detailed review of the literature on the influence of sentiment on investment decisions, taking into account the specific characteristics of the Polish market. Chapter three describes the methodology and research assumptions, including the procedure for data collection and processing, and the construction of event windows. Chapter four presents the results of the empirical research, including an analysis of the AAR and CAAR indicators and statistical significance tests. Part five discusses the findings in relation to those of previous researchers. Part six summarises the research findings, research limitations and proposed directions for future research. The final, seventh part contains information on the sources of funding for the publication.

2. Literature review

Contemporary financial markets are evolving towards an information-based economy, where the speed and precision of interpreting textual data using advanced transformer models are becoming key determinants of competitive advantage (Fazlija and Harder, 2022; Pasupuleti, 2025). The following review systematises the state of knowledge regarding the use of artificial intelligence (AI) in sentiment analysis, with particular emphasis on the specific characteristics of small-cap companies, which, due to high retail investor participation, exhibit a stronger reactivity to emotional stimuli (Sriasih et al., 2025). The study also takes into account the dynamics of the Polish capital market, where the first successes in forecasting volatility were recorded thanks to the use of precise NLP tools to analyse central bank communications and business texts (Koptyra et al., 2024; Kostyra, 2025).

The transition from traditional dictionary-based methods to advanced transformer models represents a strategic breakthrough in market modelling, as these models (e.g. BERT, RoBERTa) consistently achieve classification accuracy in the range of 89–90%, significantly outperforming earlier CNN or LSTM architectures (Anjum et al., 2024; Kamal et al., 2022). Static keyword analysis (bag-of-words) is giving way to Transformer-type architectures, which, thanks to attention mechanisms and the ability to encode long-range dependencies, enable a comprehensive, contextual understanding of financial text (Pasupuleti, 2025). The key challenge in quantitative terms is no longer merely prediction, but interpretability. The use of Explainable AI (XAI) techniques, such as SHAP or LIME (Pasupuleti, 2025; Ranade et al., 2026), enables precise sentiment attribution at the level of individual tokens. This allows the 'black boxes' of deep learning models to be unmasked, which is essential for analysts to understand the logical underpinnings of investment signals and to build retail investors' trust in AI-based systems (Saleh, 2025).

Performance analysis demonstrates that domain-fine-tuned models exhibit a significant advantage over lexical tools (VADER, TextBlob). The FinBERT model achieves an F1 score of 89.2% on the metric, with a correlation with future returns (R^2) of 0.76 (Pasupuleti, 2025). By comparison, lexical tools in studies on the FTSE 100 index achieved an accuracy of just 57% (Jin et al., 2023), rendering them ineffective in an environment with a high

level of information noise. A comparison of the effectiveness of selected AI models is presented in Table 1.

Table 1. Comparison of the effectiveness of selected AI models

Study	Model	Market studied / Dataset	Result (Accuracy / F1)	Key findings
Pasupuleti (2025)	FinBERT	S&P 500	F1: 89.2%	Highest accuracy; use of SHAP/LIME for XAI.
Kamal et al. (2022)	BERT / RoBERTa	S&P 500 / Nifty 100	Acc: 90% (BERT)	Clear dominance over CNN (80.8%) and LSTM (84%).
Anjum et al. (2024)	BERT	US banking sector	Acc: 89.4%	More effective than SVM (80.2%) and Naive Bayes (76.4%).

Source: own analysis.

Analysis of the collected data indicates that these models consistently achieve results in the range of 89–90% in terms of both accuracy and F1 score. Moving away from general-purpose models in favour of domain-specific adaptation is essential for achieving a sustainable information advantage, which becomes particularly critical when analysing the specific conditions of local markets.

The specific nature of the Warsaw Stock Exchange (WSE) presents unique research challenges that stem directly from limited liquidity and the dominance of retail investors in certain market segments (Polak, 2021; Sriasih et al., 2025). The existing literature presents a significant dichotomy in results, which — in a quantitative assessment — stems directly from the information density of the analysed text and the signal-to-noise ratio (Fazlija and Harder, 2022). Studies by Polak (2021) and Wojarnik (2022) on the WIG-Banks and WIG-Energy indices did not demonstrate significant predictive power of sentiment from newspaper headlines over a one-day horizon (Polak, 2021; Wojarnik, 2022). This contrasts with the findings of Kostyra (2025), who demonstrated that the use of Monetary Policy Council (MPC) statements — sources with high information density — allowed for a statistically significant reduction in the forecast error (RMSE) of 2.1% for the WIG20 index (Kostyra, 2025). This success suggests that the failures of earlier models stemmed from relying on generic, noise-laden headlines, whereas policy-specific texts carry a strong predictive load. This is also confirmed by research in mature markets, indicating that sentiment extracted from the full text of articles (content-based) has a significantly higher predictive value for price movements than the analysis of headlines alone (Anese et al., 2021; Fazlija and Harder, 2022). In the context of a small-cap segment such as the sWIG80, this relationship may be even stronger due to slower information absorption and lower informational efficiency in markets dominated by the emotional reactions of retail investors (Sriasih et al., 2025; Ren et al., 2025).

However, a key issue remains the ‘research gap’ concerning small-cap indices such as the sWIG80. Despite its strategic role, a comprehensive review of the literature covering the

years 2021–2026 lacks studies analysing the impact of AI on this specific segment. This is a critical point, as the structure of the sWIG80 makes it significantly more susceptible to information shocks than blue-chips, which opens up opportunities to exploit market anomalies arising from lower liquidity and a specific risk profile.

The small-cap segment is a unique area of research for behavioural finance. High retail investor participation, coupled with a lack of institutional research coverage, means that crowd psychology plays a dominant role here. These investors exhibit a stronger reactivity to emotional media stimuli, making sentiment analysis a tool with high predictive value. Studies in analogous emerging markets confirm these relationships: Sriasih et al. (2025): On the Indonesian market (IDX), a strong negative correlation (r = -0.61) was found between negative sentiment and returns, which was attributed to the affective decisions of retail investors. Ren et al. (2025): They identified the phenomenon of delayed price discovery in the small-cap segment. Using multimodal signals, including visual sentiment from 'New York Times' images, they demonstrated that the price reaction is delayed by 2–3 days. This temporary 'leak' of information makes the sWIG80 an ideal environment for verifying price anomalies, where the slow absorption of textual and visual signals allows for the generation of above-average returns.

Behavioural finance theory, based on loss aversion, points to a fundamental asymmetry in information processing. In the face of macroeconomic uncertainty, negative signals trigger more violent reactions than analogous positive messages. Empirical evidence for this asymmetry is unequivocal. Mendoza-Urdiales et al. (2022), using EGARCH models, confirmed that the intensity of the impact of negative information on share prices is statistically higher than that of positive news. Nandrajog et al. (2025) demonstrated a disparity in the reaction: a 3.4% decline for negative signals versus an increase of just 2.8% for positive ones. These findings directly support the research hypothesis regarding extreme geopolitical shocks in Q1 2026. According to the event analysis methodology, price anomalies are most pronounced in the so-called bottom quintile (the 20% lowest net sentiment readings), where shocks are absorbed rapidly due to: panic among retail investors and the lack of stabilising mechanisms, the limited depth of the small-cap market, which prevents the absorption of large supply, or an increase in systemic risk that paralyses the process of fundamental value discovery.

The use of advanced LLM models, such as Gemini Pro, allows these relationships to be uncovered in real time thanks to the ability for multidimensional contextual reasoning. Gemini Pro's advantage over dictionary-based methods lies in its ability to interpret 'entire phenomena' — the model does not simply count occurrences of the word 'sanctions', but analyses their actual impact on the supply chains of sWIG80 companies. This forms the direct foundation for modern research methodology and alpha-generation strategies in conditions of heightened geopolitical volatility.

3. Methodology and research assumptions

The aim of the study was to verify the reaction of the Polish capital market (represented by the sWIG80 index) to geopolitical and macroeconomic information shocks in the first quarter of 2026. This period was characterised by heightened volatility in global sentiment, which provided sufficient data density for analysis. Empirical verification was carried out using classical event study methodology.

The study was conducted according to the following procedure:

1. Obtaining daily closing prices for the sWIG80 index from the www.stooq.pl portal (Stooq, 20 April 2026).
2. Aggregation of the text news feed via RSS feeds from Google News (Google News, 20 April 2026) to monitor news headlines.
3. Evaluation and categorisation of the 1,098 collected headlines, bypassing static dictionary-based methods, using an AI model (Gemini Pro) with a dedicated sentiment-assessment prompt.
4. Assigning ranks to headlines: positive (+1), neutral (0), negative (-1).
5. Calculation of the net sentiment index for each session as the sum of the scores.
6. Application of an adjustment for the so-called 'weekend effect' – aggregation of news from non-trading days and assignment of these to the sentiment of the nearest trading session.
7. Setting a statistical cut-off threshold at the lowest quintile (bottom 20% of the distribution) of the net sentiment index for the entire quarter.
8. Identification of N=13 trading days representing extremely negative news shocks, which served as the reference point (day zero) for the analysis.
9. Adopting a narrow, symmetrical three-day event window: [-1; +1].
10. Adopting the simple daily return of the sWIG80 index as a direct measure of abnormal return. This model assumes, in simplified terms (the so-called constant mean return model), that for efficient markets the expected daily return in the ultra-short event window is close to zero.
11. Calculation of the average abnormal return (AAR) for each of the three days in the study window.
12. Calculation of the cumulative abnormal return (CAR) for each of the 13 events individually (using the compound interest method).
13. Determination of the final measure: the cumulative average abnormal return (CAAR) as the arithmetic mean of the CAR values.
14. Formulation of the null hypothesis (H_0) regarding the absence of a significant difference between the CAAR value and zero (which would imply effective absorption of the shock by the market).
15. Conducting a two-tailed Welch's t-test for a single sample (relative to an auxiliary control group with zero variance).

16. Assessing the significance of the results obtained at the standard confidence level of $\alpha=0.05$.

The study was based on two independent datasets: daily closing prices for the sWIG80 index obtained from the stooq.pl portal (Stooq, 20 April 2026) and a stream of text news items from aggregated via RSS feeds from Google News (Google News, 20 April 2026) obtained using the following keywords: growth, record, rebound, improvement, recovery, better than expected, profits rising, interest rate cut, stabilisation, crisis, recession, decline, collapse, crash, bankruptcy, sanctions, conflict, uncertainty, risk, rising inflation, interest rate hike. This allowed for the objective tracking of headlines that investors encounter on a daily basis.

Although the existing literature identifies the Polish RoBERT model as optimal for typical Polish financial texts and stock market reports, this study analysed strictly geopolitical headlines describing international events within a broader context. To this end, it was concluded that the ability of an LLM-class model such as Gemini Pro to generalise conclusions and understand global processes in context (so-called zero-shot reasoning) would be more suitable for the collected research sample than a specialised domain model.

To convert qualitative news headlines into quantitative values (an explanatory variable), static dictionary-based methods were abandoned in favour of contextual assessment using an LLM. The Gemini Pro model was used for categorisation, tasked with the role of an assistant assessing the potential impact of a geopolitical event on the sWIG80. Ranks (positive: 1, neutral: 0, negative: -1) were assigned automatically in a zero-shot model using the following prompt:

"You are a senior financial markets analyst and algorithmic expert. Your task is to assess the sentiment of the following geopolitical news headlines solely in terms of their potential impact on the Polish sWIG80 stock index.
Scoring rules: Return -1 if the news increases uncertainty, escalates conflict, introduces tariffs/sanctions or threatens trade (a signal for declines). Return 0 if it is a dry political fact, an announcement of talks without firm economic decisions, or information neutral to the Warsaw Stock Exchange. Return 1 if the news de-escalates tensions, removes trade barriers, or reports an inflow of capital/support from the EU (a signal for rises).
Return the result in table format: [Title] | [Rating]. Do not add any other text or comments."

A key advantage of the LLM-based approach (Gemini Pro) over dictionary-based methods was that the model categorised entire events and news headlines within their full geopolitical context, rather than merely evaluating sets of isolated keywords (the so-called bag-of-words approach-of-words approach), which allowed for a more accurate assignment of market sentiment rankings.

The number of headlines broken down into positive, neutral and negative is shown in Figure 1.

Fig. 1. Number of headlines broken down into: positive (1), neutral (0) and negative (-1)

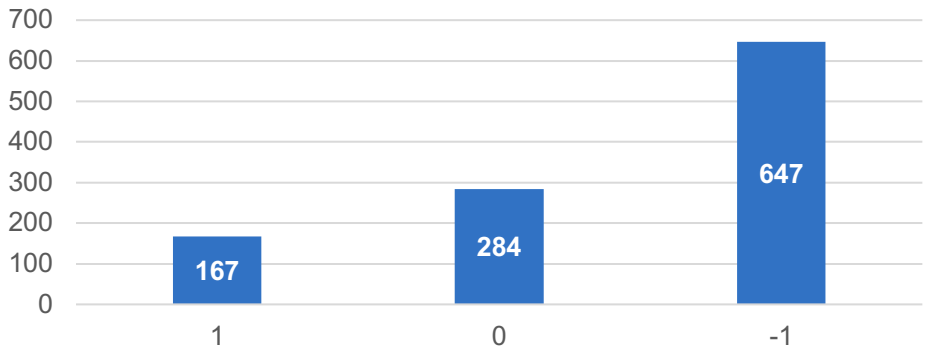

Source: own analysis.

Of all 1,098 headlines, the vast majority, 647 (59%), were negative, 284 (26%) were neutral, and only 167 (15%) were positive, confirming that negative sentiment dominated media coverage during the study period.

Each trading day was then assigned a net sentiment index, representing the sum of the previously determined rankings. To maintain informational continuity, a weekend effect aggregation was applied – information published on non-trading days was aggregated and assigned to the sentiment of the nearest trading session.

To avoid arbitrariness in defining an "information shock", the cut-off threshold was determined using a statistical method. Only those trading sessions in which the net sentiment index fell within the lowest quintile (the bottom 20% of the distribution) for the entire quarter were considered the event day (t=0). This rigorous selection process allowed us to identify N=13 extremely negative information shocks.

Due to the high density of events, and in order to avoid the problem of overlapping events, a narrow, symmetrical event window covering 3 trading days was used: [-1; +1]. This design made it possible to examine:

- the phenomenon of information leakage on day t=(-1),
- the immediate market reaction on the day of the shock, t=0,
- and the final absorption of information on day t=(+1).

For each day within the event window, the simple return on the sWIG80 index was calculated, which in this model served as the abnormal return (it was assumed that the expected daily return on the index in the ultra-short term tends towards zero). Next, the *average* abnormal return (AAR) was calculated for each day relative to the event:

$$AAR_t = \frac{1}{N}\sum_{i=1}^{N} R_{i,t}$$

where N is the number of events (N=13), and $R_{i,t}$is the return on day t for event i.

To assess the total impact of the geopolitical shock on the investor's portfolio across the entire window [-1; +1], the cumulative abnormal return (CAR) was calculated for each event i, using the compound interest formula:

$$CAR_i = \left(\prod_{t=-1}^{+1}(1 + R_{i,t})\right) - 1$$

The final research measure was the cumulative average abnormal return (CAAR), representing the arithmetic mean of all thirteen CAR_ivalues.

To verify the statistical significance of the observed price deviations, a null hypothesis (H_0) was formulated, assuming that the mean value of CAAR does not differ significantly from zero, which would indicate that the market had effectively absorbed the shock.

The verification was carried out using a two-tailed Student's t-test for a single sample. Due to limitations in the built-in functions of the spreadsheet's computational environment, this test was conducted using Welch's test algorithm, to which an auxiliary control group with zero variance was introduced. This procedure is mathematically sound and reduces directly to the classical formula for the one-sample Student's t-test (Aczel and Sounderpandian, 2018), allowing for a reliable verification of CAAR deviations from zero.

4. Research results

As part of the analysis, zero days (T=0) were identified from the entire period under review, representing the most severe geopolitical shocks affecting the Polish market. A rigorous cut-off threshold in the form of the lower quintile of the net sentiment indicator was used, which allowed for the identification of 13 independent events. Basic descriptive statistics for the identified sample are presented in Table 2.

Table 2. Descriptive statistics of the sample of geopolitical events (N=13)

Number of identified events (N)	13
Minimum net sentiment index (strongest shock)	-31.0
Average net sentiment index during the study period	-7.74
Average net sentiment index for identified events	-19.38

20% net sentiment percentile	-13.8
Event window used	[-1; +1]

Source: own analysis.

The collected data confirm that the selected sample is characterised by a high negative sentiment. These events represent days with extreme levels of market fear and pessimism ( ), which provides an optimal environment for testing the resilience (effectiveness) of the sWIG80 index.

A key element of the research verification was the calculation of average abnormal returns (AAR) for individual days within the event window and the cumulative average abnormal return (CAAR) for the entire study period. The results of the event analysis, together with statistical significance tests (Welch's t-test), are presented in Table 3.

Table 3. Results of event analysis for the sWIG80 index around geopolitical shocks

Day relative to the shock (t)	Average AAR / CAAR	t-statistic	p-value
-1 (Information leak)	+0.12%	0.40	0.694
0 (Shock day)	-0.04%	-0.16	0.877
+1 (Absorption)	+0.05%	0.13	0.902
CAAR [-1; +1]	+0.13%	0.19	0. ,851

Source: own calculations.

The data in Table 3 indicate that at no stage of the event window (days from t=(-1) to t=(+1)) did the rates of return differ in a statistically significant manner from zero (all p-values significantly exceed the α=0.05 threshold, ranging from 0.694 to 0.902).

The dynamics of geopolitical risk absorption by the capital market are illustrated by the shape of the CAAR curve, shown in Figure 2.

Fig. 2. Cumulative average excess return (CAAR) around day zero

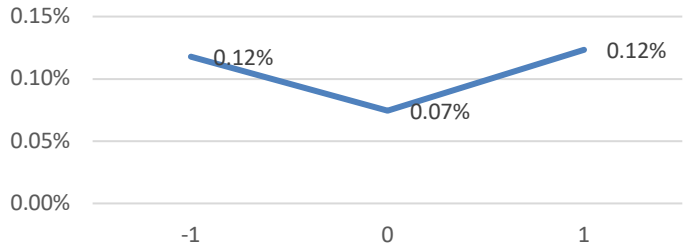

Source: own work.

Although a visual analysis of Figure 2 might suggest the occurrence of a symmetrical, short-lived correction (a decline on the shock day t=0 from 0.12% to 0.07% and a subsequent full rebound on day t=1), the lack of statistical significance dictates that these fluctuations should be treated solely as market noise. Furthermore, given the absence of a decline on day t=(-1), the available data cannot confirm the occurrence of information leakage, and the deviations themselves are too small to justify the presence of a full-fledged spring effect (Szyszka, 2009).

The key result determining the verification of the research hypothesis is the outcome of the test for the cumulative event window. Since the statistical test value was p=0.851, this result exceeds the standard rejection level of α=0.05 many times over. This indicates a complete lack of statistical significance for the observed changes in the three-day operational window.

The above result therefore confirms the validity of the semi-strong form of the Efficient Market Hypothesis (Fama, 1970) for the sWIG80 index in the first quarter of 2026. It has thus been demonstrated that, during the period under review, the sWIG80 index proved resilient to superficial geopolitical noise. The fact that the cumulative values remained consistently positive (above zero) throughout the entire event window ultimately proves that individual shocks of this type do not provide investors with the opportunity to systematically achieve above-average returns from short-selling strategies based on sentiment analysis in the case of Polish small-cap companies.

5. Discussion of results

The empirical results obtained, indicating a lack of statistical significance (p=0.851) for changes in the sWIG80 index in response to negative geopolitical shocks, are consistent with the clear dichotomy of results observed in the Polish literature on the subject. These results are consistent with the findings of Polak (2021) and Wojarnik (2022), who, in their studies of WSE sector sub-indices (WIG-Banks, WIG-Energy), did not demonstrate any significant explanatory power of sentiment extracted from newspaper headlines. At the same time, these results contrast with those of Kostyra (2025), who noted a statistically significant improvement in forecasts for the WIG20 index. This discrepancy may stem from a difference in the information density of the analysed texts; whereas this study analysed noisy newspaper headlines, Kostyra relied on substantive statements from the

Monetary Policy Council, suggesting that the sWIG80 market may be immune to general media coverage, reacting only to hard signals of a political and economic nature.

A comparison of the sWIG80 with similar emerging markets reveals significant differences in the mechanisms by which market sentiment is priced in. Although the structure of the Polish small-cap index is characterised by high retail investor participation, no such strong negative correlation between negative sentiment and returns was observed as that identified by Sriasih et al. (2025) in the Indonesian market ($r = -0.61$). Furthermore, the lack of a significant reaction in the [+1] window stands in contrast to the theory of delayed *price discovery* in the *small-cap* segment described by Ren et al. (2025). Instead of a 2–3-day delay in information absorption, the data obtained suggest behaviour consistent with the assumptions of the semi-strong form of the Efficient Market Hypothesis (Fama, 1970), where publicly available geopolitical information is discounted immediately or deemed irrelevant to the fundamental value of companies.

It is also worth referring to the theoretical asymmetry in information processing. Behavioural literature indicates that negative signals typically trigger more violent price reactions than positive announcements. Although the studies by Mendoza-Urdiales et al. (2022) and Nandrajoga et al. (2025) confirm this asymmetry in mature markets, the stability of the sWIG80 index in the face of prevailing media pessimism in Q1 2026 (59% negative headlines) may indicate a specific 'crisis resilience' or market saturation with negative reports. The use of an advanced LLM model (Gemini Pro) enabled the unmasking of these relationships through a contextual assessment of events rather than a simple counting of keywords. This confirms the methodological advantage of transformer models over lexical methods in identifying the actual impact of news on markets with lower liquidity.

6. Summary

Research conducted on the reaction of the sWIG80 index to extreme geopolitical shocks in the first quarter of 2026 has led to the following conclusions:

- **W_1 :** On the day before the occurrence of a negative information shock (t=-1), no 'information leak' or pre-emptive panic was observed. On the contrary, the average annualised return (AAR) was slightly positive (+0.12%), which, in the absence of statistical significance (p=0.694), demonstrates the high resilience of the small-cap market to early geopolitical signals.

- **W_2 :** The market does not exhibit a negative *post-event drift* following the occurrence of the event. On the day of the shock (t=0) and the following day (t=+1), the index fluctuated around zero (-0.04% and +0.05%), confirming the stability of sWIG80 valuations and the absence of a sustained reaction to the information stimuli under investigation.

- **W_3 :** Despite the theoretically greater susceptibility of small-cap companies to emotional shocks, the sWIG80 index demonstrated resilience to superficial geopolitical noise during the period under review. This suggests that such

information is absorbed effectively or deemed irrelevant to the fundamental value of smaller-cap companies.

- **W_4 :** The results obtained confirm the operation of **the semi-strong form of the efficient market hypothesis** in the *small-cap* segment. It has been demonstrated that, based solely on negative sentiment from newspaper headlines, investors are unable to systematically generate above-average returns from short-selling strategies.

The research hypothesis posited that extremely negative information shocks would trigger a strong, multi-session downward reaction in the sWIG80 index. **Based on the empirical research conducted, this hypothesis was not confirmed.** The cumulative average abnormal return (CAAR) for the entire window [-1; +1] was +0.13% and proved to be completely statistically insignificant (p=0.851). This result repeatedly exceeded the adopted significance level of α=0.05, which unequivocally demonstrates the absence of a sustained price decline triggered by the factors under investigation.

When interpreting the results obtained, it is important to highlight the significant limitations of the study, the key one being that the analysis is based on daily closing prices, which may prevent the capture of the rapid dynamics of sentiment that often fades in algorithmic markets within just a few hours or minutes. Furthermore, the use of the Gemini Pro model's *zero-shot* evaluation, despite its effectiveness in understanding the broader geopolitical context, may be less precise than domain-specific transformer models such as Polish RoBERTa. It should also be noted that conducting sentiment analysis in isolation from technical analysis indicators (e.g. volume or RSI) has lower predictive power, as it is only the fusion of textual signals with market data that allows the full potential of AI models to be harnessed in forecasting stock market behaviour.

Proposed avenues for future research primarily involve extending the analysis to include high-frequency data (intra-session intervals), which would allow for a precise assessment of the speed at which information is absorbed in the less liquid segment, particularly as the market's reaction to sentiment may occur and subside exclusively on an intraday basis. The next significant step should be the implementation and fine-tuning of the Polish RoBERT model on specific announcements from sWIG80 companies (e.g. EBI/ESPI reports), which offer significantly higher information density than noisy press headlines and – as demonstrated in studies of the Polish market – constitute the most effective architecture for extracting emotions from business texts. Another promising avenue is the development of hybrid predictive systems that combine sentiment indicators with endogenous market variables and EGARCH volatility models, which, according to the literature, allows for a significant improvement in the accuracy of directional forecasts and a better interpretation of systemic risk. Finally, further verification of information asymmetry and delayed *price discovery* in various phases of the market cycle is essential, as the specific nature of the *small-cap segment* is associated with a price reaction to multimodal information shocks that may be delayed by 2–3 days, creating scope for exploiting persistent market anomalies.

Sources of funding

The project was co-financed by the Ministry of Science under the 'Regional Excellence Initiative' programme.

References

- Aczel A.D., Sounderpandian J., (2018): „Statystyka w zarządzaniu." *Wydawnictwo Naukowe PWN*, Warszawa.
- Anese G., Corazza M., Costola M., Pelizzon L., (2021): "Impact of public news sentiment on stock market index return and volatility." *Computational Management Science.* DOI: https://doi.org/10.1007/s10287-023-00454-2.
- Anese G., Corazza M., Costola M., Pelizzon L., (2023): "Impact of public news sentiment on stock market index return and volatility." *Computational Management Science*, nr 20(1)/2023. DOI: https://doi.org/10.1007/s10287-023-00454-2.
- Anjum N., Bony M.N.V.A., Alam M. i in., (2024): "Comparative Analysis of Sentiment Analysis Models on Banking Investment Impact by Machine Learning Algorithm." *International Journal of Computer Science & Information System.* DOI: https://doi.org/10.55640/ijcsis/volume09issue11-02.
- Fama E.F., (1970): "Efficient capital markets: A review of theory and empirical work." *Journal of Finance*, nr 25(2)/1970, s. 383-417. DOI: https://doi.org/10.1111/j.1540-6261.1970.tb00518.x.
- Fazlija B., Harder P., (2022): "Using Financial News Sentiment for Stock Price Direction Prediction." *Mathematics,* nr 10(13)/2022. DOI: https://doi.org/10.3390/math10132156.
- Google News, (2026): Usługa agregacji wiadomości prasowych (kanały RSS). https://news.google.com/ – data dostępu 01.04.2026.
- Jin C., Liu R., Tang B-Y, Cai B., (2023): "Predict FTSE100 Stock Movements Using Business News Sentiment and Machine Learning." *Theoretical and Natural Science.* DOI: https://doi.org/10.54254/2753-8818/2/20220148.
- Kamal S., Sharma S., Kumar V. i in., (2022): "Trading Stocks Based on Financial News Using Attention Mechanism." *Mathematics,* nr 10(12)/2022, nr art. 2001. DOI: https://doi.org/10.3390/math10122001.
- Koptyra B., Oleksy M., Dzięcioł E., Kocoń J., (2024): "Small Language Models for Emotion Recognition in Polish Stock Market Investor Opinions." *2024 IEEE*

International Conference on Data Mining Workshops (ICDMW), s. 463-470. DOI: https://doi.org/10.1109/ICDMW65004.2024.00065.

- Kostyra T., (2025): "Predictive power of the sentiment of the Monetary Policy Council." *Bank i Kredyt*, nr 56(5)/2025, s. 543-556. DOI: https://doi.org/10.5604/01.3001.0055.3042.

- Mendoza-Urdiales R.A., Núñez-Mora J.A., Santillán-Salgado R.J., Valencia-Herrera H., (2022): "Twitter Sentiment Analysis and Influence on Stock Performance Using Transfer Entropy and EGARCH Methods." *Entropy*, nr 24(7)/2022, nr art. 874. DOI: https://doi.org/10.3390/e24070874.

- Nandrajog A.B., Kaur M., Kinra S. i in., (2025): "A Novel CNN-SVM Framework for Stock Trend Prediction Leveraging Historical Market Data and News Sentiment." *2025 5th International Conference on Intelligent Technologies (CONIT).* DOI: https://doi.org/10.1109/CONIT65521.2025.11167870.

- Pasupuleti M.K., (2025): "Explainable Sentiment Analysis for Financial News and Market Prediction." *International Journal of Academic and Industrial Research Innovations (IJAIRI).* DOI: https://doi.org/10.62311/nesx/rphcr10.

- Polak K., (2021): "The Impact of Investor Sentiment on Direction of Stock Price Changes: Evidence from the Polish Stock Market." *Journal of Banking and Financial Economics*, nr 2(16)/2021, s. 72-93. DOI: https://doi.org/10.7172/2353-6845.jbfe.2021.2.4.

- Ranade D.J., Bhaya S., Bhimakari S. i in., (2026): "On Explaining the Sentiments in Prediction of Stock Movement: An XAI-Based Analysis." *HCAIep.* DOI: https://doi.org/10.1145/3777490.3777510.

- Ranjbaran G., Moin M., Alizadeh S.H., Koochari A., (2024): "Investigating Sentiment Analysis of News in Stock Market Prediction Utilizing Machine Learning Technique." *International Journal of Information and Communication Technology Research*, nr 16(3)/2024. DOI: https://doi.org/10.61186/itrc.16.3.54.

- Ren X., Jiang W., Sun X., Wang S., (2025): "Investor sentiment from images: a few-shot learning investigation." *Journal of Accounting Literature*. DOI: https://doi.org/10.1108/jal-04-2025-0199.

- Saleh A., (2025): "Interpretable Multilingual NLP Models for Financial Decision-Making in Frontier Markets: A Case Study on the Pakistan Stock Exchange." *Pakistan journal of humanities and social sciences.* DOI: https://doi.org/10.52131/pjhss.2025.v13i4.3037.

- Seidler-de Alwis R., (2018): Markt- und Wettbewerbsanalyse in Zeiten von Big Data und Digitalisierung. W: Schade F., Georgy U. (red.), "Praxishandbuch Informationsmarketing: Konvergente Strategien, Methoden und Konzepte." *De Gruyter Saur*, Berlin, Boston, s. 59-70. DOI: https://doi.org/10.1515/9783110539011-005.

- Sriasih S.D., Razak F.A., Ikhsan H. al I., (2025): "AI-Driven Sentiment Analysis of Retail Investor Behavior during Market Volatility: A Study of Twitter Data in Southeast Asia." *Journal of Management and Informatics.* DOI: https://doi.org/10.51903/jmi.v4i1.179.

- Stooq.pl, (2026): Notowania historyczne indeksu WIG20. https://stooq.pl – data dostępu 01.04.2026.

- Szyszka A., (2009): „Finanse behawioralne. Nowe podejście do inwestowania na rynku kapitałowym." *Wydawnictwo Uniwersytetu Ekonomicznego w Poznaniu*, Poznań.

- Tkachyk D., Kvyetnyy R., (2022): "Sentiment analysis of the text for forecasting data in financial markets." *Information Technology and Computer Engineering*, nr 55(3)/2022, s. 51-58. DOI: https://doi.org/10.31649/1999-9941-2022-55-3-51-58.

- Wojarnik G., (2022): "Sentiment analysis towards listed companies in the energy sector on the example of the WIG-Energia index on the WSE." *Procedia Computer Science*, nr 207/2022, s. 2076–2085. DOI: https://doi.org/10.1016/j.procs.2022.09.265.