

Northern inSights: Geopolitical Risk from Finnish News Media

Gene Ambrocio, Zuzana Fungáčová, Joni Heikkinen,
Eeva Kerola, Iikka Korhonen, Anni Norring

Eesti Pank Research Seminar, 3 September 2026

The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Bank of Finland, the Eurosystem, or the IMF.

Background

- A growing literature highlights the importance of **geopolitical risk** for economic and financial outcomes
- A key first step is to **measure geopolitical risk (GPR)**
- The leading approach constructs indices from **news coverage**
  - based on keywords related to wars, terrorism, and international tensions
- Seminal GPR indices rely on **English-language newspapers**
  - United States, United Kingdom, and Canada
  - reflect a largely **global / Anglo-Saxon perspective**

From Global Attention to Local Perceptions

- A global perspective may be too blunt for **local geopolitical risk**
- Country-specific indices based on English-language media might reflect **global attention** to countries
- This approach has important limitations:
  - geopolitical risks abroad can affect the domestic economy even without country-specific events
  - local perceptions may differ from English-language reporting
  - locally salient events may receive limited international coverage
- These issues motivate the use of **local-language, local news media** to capture **local geopolitical risk perceptions**

From Global Attention to Local Perceptions

- Finland provides a useful case to study local perceptions of geopolitical risk
- Finnish-language news is largely consumed domestically
- Trust in news media is exceptionally high:
  - 71% in Finland vs. 48% EU average (Eurobarometer, 2023)
  - 69% in Finland vs. 40% global average (Reuters Institute, 2024)
- This allows local news to capture **local geopolitical risk perceptions**, rather than global attention

In This Paper

- Construct **FinnGPR** from Finnish-language Yle news (2011–2024) using **two approaches**:
  - **Dictionary-based**: keyword rules (Caldara & Iacoviello, 2022)
  - **Language model-based**: fine-tuned Finnish BERT classifier
- Both validated against **human evaluation**; we additionally explore a generative LLM (**Poro 2**)
- Compare with English-language Anglo-Saxon measures:
  - **Global GPR** and **GPRC-FIN** (Finland-related coverage)
- Main findings:
  - FinnGPR aligns with **global** GPR; GPRC-FIN behaves like **global attention to Finland**
  - The BERT classifier matches humans **far better** than the dictionary
  - FinnGPR shocks: economy **resilient**, but **systemic stress rises** and confidence falls

Related Literature

- News-based measures of geopolitical risk:
  - Caldara & Iacoviello (2022): global geopolitical risk and country indices from English-language media
- Local-language, local-media approach:
  - Bondarenko et al. (2024, 2025): local GPR indices (Russia, euro area) and a perceptions interpretation
  - Jung et al. (2021): local news-based GPR for South Korea
  - Alonso-Alvarez et al. (2025): bilateral and regional GPR indices
- AI / language-model approaches:
  - Iacoviello & Tong (2026): AI-GPR – an LLM *semantically* scores articles on a continuous 0–1 risk scale, moving beyond keyword counting
  - Ioannou et al. (2026): LGPT – euro-area local-language press; BERT screens, LLM classifies; separates geopolitical from *geoeconomic* tension

Outline

- Background
- **Constructing a Finnish GPR**
  - Dictionary index
  - Language model index
- Empirical Analysis
- Conclusions

Data Source: Yle Finnish News Archive

- Source: **Yle Finnish News Archive** (all online articles)
- Publisher: Yle (Yleisradio Oy), Finland's national public broadcaster
  - Television, radio, and online media
  - State-owned, with editorial independence
- Sample period: **2011–2024**
  - About **1 million articles**
  - ∼200 articles per day, ∼6,200 per month
- Yle consistently ranks among the **most read** news sources and receives the **highest trust ratings** among Finnish outlets (Reuters Institute, 2024)
- Online news matches how Finns actually consume news
  - Online: **90%** of population (TV: 63%) (Reuters Institute, 2024)

Dictionary Index: Keyword-Based Approach

- We follow the methodology of Caldara and Iacoviello (2022) Definition
- Basic idea: identify articles via keyword combinations in eight GPR topics

1. War threats
2. Peace threats
3. Military buildup
4. Nuclear threats
5. Terrorist threats
6. Beginning of war
7. Escalation of war
8. Terrorist acts

Not fictional, sports, or historical

- We flag an article if it contains the required **word combination** for at least one topic
- Mostly a same-sentence rule; some topics only require both words somewhere in the article
- Exclusions are always enforced

Dictionary Index: Keyword-Based Approach

- We follow the methodology of Caldara and Iacoviello (2022) Definition
- Basic idea: identify articles via keyword combinations in eight GPR topics

1. War threats
2. Peace threats
3. Military buildup
4. Nuclear threats
5. Terrorist threats
6. Beginning of war
7. Escalation of war
8. Terrorist acts

Not fictional, sports, or historical

- We flag an article if it contains the required **word combination** for at least one topic
- Mostly a same-sentence rule; some topics only require both words somewhere in the article
- Exclusions are always enforced

Word Lists Used: One Topic in Full

Word list 1 war, conflict, uprising, coup, warfare, hybrid influence, cyber attack ...

sota* soda* sotia sotaa konflikt* vihollisu* vallankumou* kapina* kansannous* vallankaappau* geopoliitti* sotatoim* yhteenot* hybridivaikutta* kyberhyök*

in the SAME SENTENCE as

Word list 2 risk, warn, fear, danger, threat, crisis, tension, imminent ...

riski* varoitta* pelko pelottaa vaara* uhka uhkaa uhata epäilys kriisi* vaikeu* kiista* huoli jännite* välitön* väistämä*

and NOT movie, museum, sport, history, tax, football ...

elokuv* muse* vuosipäiv* taide* kirja* `hintaso*` peli* histori* urheil* musiik* syöpä* mafia* vero* jääkiek* jalkapal* `sodankyl*`

This is **one** of eight topics. Each has its own word pair; the exclusion list is shared.

Dictionary Index: Finnish Language Adaptation

- We adapt Caldara & Iacoviello's (2022) GPR terms to Finnish-language news ▸ English Language Word List
- We iteratively validate the word list:
  - compare algorithmic classifications with human evaluations for a subset of articles
  - compare keywords with actual language used in news articles
- We add locally relevant terms: **airspace violations, hybrid influence, cyber attacks**

Dictionary Index: Finnish Language Adaptation

- We adapt Caldara & Iacoviello's (2022) GPR terms to Finnish-language news ▸ English Language Word List
- We iteratively validate the word list:
  - compare algorithmic classifications with human evaluations for a subset of articles
  - compare keywords with actual language used in news articles
- We add locally relevant terms: **airspace violations, hybrid influence, cyber attacks**

Dictionary Index: Finnish Language Adaptation

- We adapt Caldara & Iacoviello's (2022) GPR terms to Finnish-language news ▸ English Language Word List
- We iteratively validate the word list:
  - compare algorithmic classifications with human evaluations for a subset of articles
  - compare keywords with actual language used in news articles
- We add locally relevant terms: **airspace violations, hybrid influence, cyber attacks**

Dictionary Index Validation: Design

- Does the dictionary actually work? **We validate against a human evaluator**
- But geopolitical articles are rare – only about **1.4%** of the corpus
  - random sampling alone would yield too few positive cases
- We therefore use a validation set that combines:
  - random sampling
  - targeted sampling around peak-risk days and major events
- We iteratively compare the algorithm's classifications with human evaluations

Building the Dictionary: An Iterative Loop

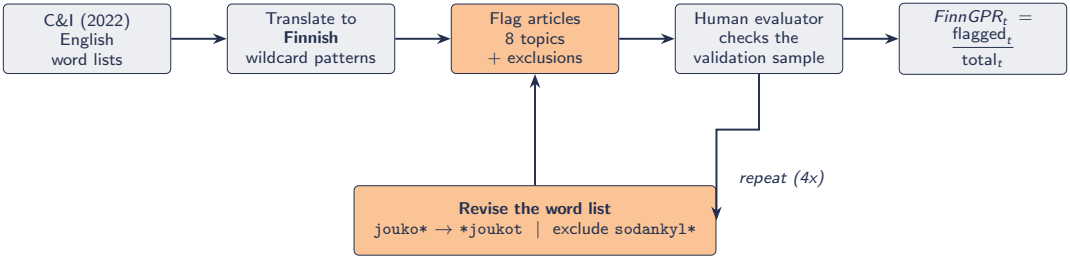

- `jouko*` was **too broad**: *joukot* (troops), but equally *joukossa* ("among") ⇒ moved the wildcard to the front, `*joukot`, which also captures *venäläisjoukot*
- `sodan*` ("of the war") also matches **Sodankylä**, a municipality in Lapland ⇒ it goes on the **exclusion** list

Dictionary Index Validation: Results

Algorithm / Human	Flagged	Not Flagged	Total
Flagged	147 (94%)	10 (6%)	157
Not Flagged	141 (16%)	725 (84%)	866
Total	288	735	1023

- Overall accuracy: **85%**
- Errors are asymmetric:
  - fewer false positives than false negatives
  - the dictionary index is **conservative** relative to a human evaluator

From Articles to the FinnGPR Index

$$GPR_{day} = No.flagged_{day} \; / \; Tot.articles_{day}$$

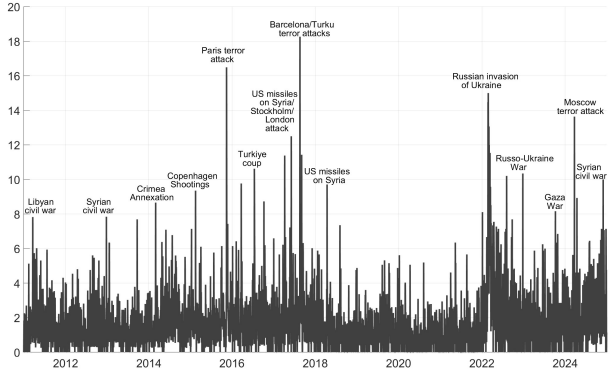

FinnGPR Sub-Indices: Dynamics Over Time

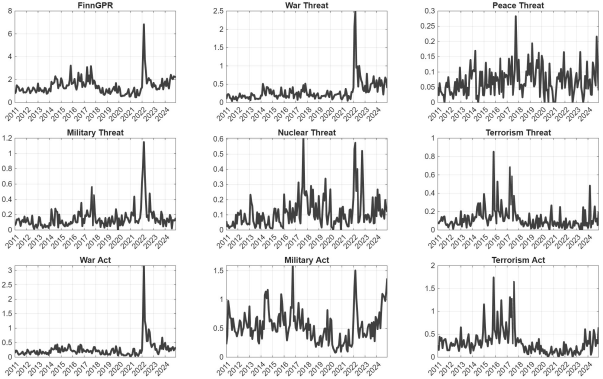

- We also examine the eight FinnGPR sub-indices separately
- Their dynamics differ over time
- War- and military-related topics drive major spikes
- Terrorism-related components peak in the mid-2010s

Outline

- Background
- **Constructing a Finnish GPR**
  - Dictionary index
  - **Language model index**
- Empirical Analysis
- Conclusions

What Is a Language Model?

- Trained on **one task**: predict a missing word

"She took the train from Helsinki to _____"

- Language models get this ability from **pre-training**: the missing word is the answer, so **no human has to label anything**
- Doing it well requires internally representing **grammar, morphology and context**
- That representation is what makes the model useful for *other* tasks

Two Kinds of Language Model

BERT *Bidirectional Encoder Representations from Transformers* Google, 2018	**GPT** *Generative Pre-trained Transformer* OpenAI, 2018; GPT-3 2020; ChatGPT 2022
Reads the whole text at once	Reads left to right
Fine-tuned to output a **label**	Prompted to output **text**
Smaller ($\sim$100M)	Larger (billions) – the **"L"** in LLM
Must be **taught** a task	Can be **told** a task

We use a BERT-type model: our task is **one yes/no question, asked a million times.**

Our Classifier: Fine-Tuned FinnBERT

- BERT model **pre-trained on Finnish** (Virtanen et al., 2019): it arrives knowing the language
- **Context comes with it**: *war* in *price war* is not the same as *war* in *war close to the border*
- But it knows **nothing about geopolitical risk**: our definition has to come from somewhere
- **Fine-tuned** on the **700** articles our evaluator had already labeled
  - we add a **yes/no classification layer** on top of the pre-trained model
  - it learns to imitate the **evaluator**, not our keyword list

From One Million Articles to a Monthly Index

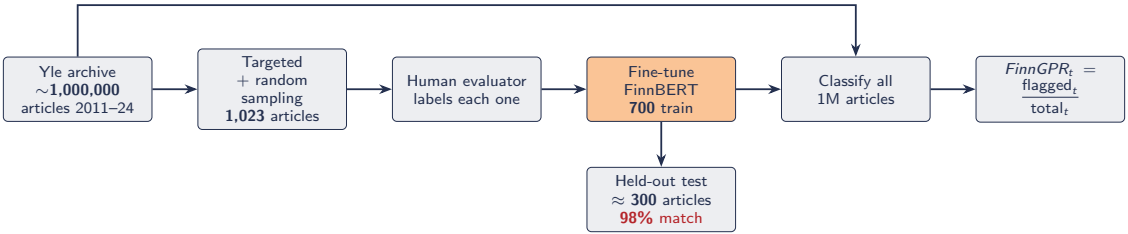

- Same index formula as the dictionary – only the classifier changes

Same Article, Two Readings — Actual Cases

Language model	Dictionary
Paris night turned into a bloodbath (14 Nov 2015) At least 120 died in the terrorist attacks in Paris. Terrorists struck six targets. A state of emergency declared in France. **Geopolitical** ✓ agrees with human	At least 120 died in the terrorist attacks in Paris. Terrorists struck six targets. A state of emergency declared in France. **Not geopolitical** × no listed word pair
White House condemned the Istanbul attack (1 Jan 2017) The attack on New Year's revellers shows the perpetrator's brutality. Gunmen killed at least 35 people. **Geopolitical** ✓ agrees with human	The attack on New Year's revellers shows the perpetrator's brutality. Gunmen killed at least 35 people. **Not geopolitical** × no listed word pair
Police separate rival demonstrators in Turku (19 Aug 2017) Two demonstrations in the market square, site of yesterday's terror attack. Police keep the groups apart. **Geopolitical** × borderline aftermath coverage	Two demonstrations in the market square, site of yesterday's terror attack. Police keep the groups apart. **Not geopolitical** ✓ agrees with human

Notes: actual Yle articles, translated and abridged. ✓/× = agreement with the human evaluator. Finnish compounds count as one word: "terrorist attack" is *terrori-isku*.

Can We See Why? Looking Inside the Classifier

- We have **some tools** for inspecting how the classifier reaches a decision
- **Attention weights**: when a word is ambiguous, they show what the model looked at to resolve it
  - seeing *war*, the model looks around: are the nearby words *prices* and *chains*, or *troops* and *border*?
  - we can read those weights off and see **what it looked at**
- **Leave-one-out**: drop a word and see how the prediction moves
  - directly measures what the classification *depends on*

⇒ The shading on the previous slide is our reading of the context; these tools let us check it against the model

▸ Inside BERT

But Why Not Just Use ChatGPT?

- A fair question: a large generative model can be **told** the task rather than taught it
- **Zero-shot**: the instruction only
  - *"Does this article concern geopolitical risk? Answer yes or no."* + article
- **Few-shot**: the instruction *plus worked examples*
  - *"... Here are five articles and the correct answers: ..."* + article
  - the model infers the task from the examples; **nothing is retrained**
- Few-shot and fine-tuning both use labeled examples. The difference is **where they live**: in every prompt, or in the weights

Why Not an LLM? Finnish Is the First Problem

- Commercial API models are not available to us: we run **open-weight** models locally
- We tested six LLMs incl.: **Gemma**, **Qwen3**, **Phi-4** and **Poro 2**, a model trained specifically on Finnish
- The general multilingual models struggle with Finnish
  - asked for four case forms of *talo*: **Poro 2** got all four, **Gemma 3** got two
- Only **Poro 2** was reliable enough to use, and it is also the **slowest** of the six

▸ Model comparison

Why Not an LLM? Accuracy and Cost

Classifier	Match w/ human
Dictionary	∼88%
Poro 2, few-shot	∼90%
Fine-tuned FinnBERT	∼**98%**

On the held-out test sample

- Zero-shot was **markedly worse**; few-shot brought Poro 2 to ∼90%
- But Poro 2 is **8B** parameters against FinnBERT's **110M**: far slower and heavier per article
- And few-shot prompts are **long**: the worked examples are re-read on *every* article

⇒ Too costly to run over **1M** articles. Poro 2 is a **benchmark, not the classifier**; the index is built with FinnBERT

Dictionary vs. Language Model: Matching Humans

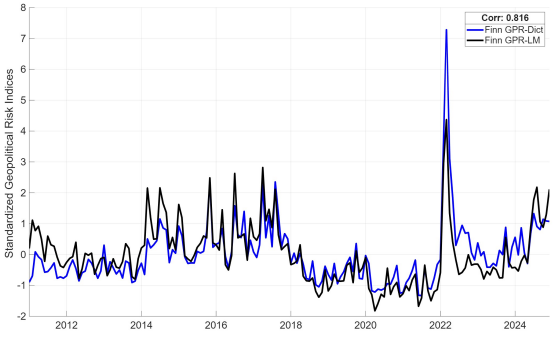

LM flags ∼**3.9%** of articles vs. ∼**1.4%** for the dictionary

	Match	False Pos.	False Neg.
Finn GPR-Dict	88.0%	0.7%	11.3%
Finn GPR-LM	97.5%	1.1%	1.4%

On ∼300 held-out human-labeled articles

- BERT matches humans on ∼**98%** vs. ∼**88%** for the dictionary
- LM error is small and **symmetric**
- Indices co-move (corr. **0.82**)
- Further check: **200 random articles**, never used in training or testing — **no false positives**

Comparing GPR Indices

- **FinnGPR**: geopolitical risk in Finnish-language local news (dictionary + LM)
- **Global GPR** (Caldara & Iacoviello):
  - geopolitical risk in **English-language global media**
- **GPRC-FIN** (Caldara & Iacoviello):
  - **Finland-related** geopolitical risk in English-language media
- **EA GPR** (Bondarenko et al. 2024/2025): euro-area local-language index
  - also compared: **Russia GPR** (Bondarenko et al. 2024)

Comparing GPR Indices: Correlations

	Finn GPR-Dict	Finn GPR-LM
GPRC-FIN	0.383	0.112
Global GPR	**0.798**	**0.578**
EA GPR	0.742	0.474
Russia GPR	0.493	0.201

Finn GPR-Dict and Finn GPR-LM correlate at **0.82**

▸ Full correlation table

- FinnGPR correlates most with **Global GPR**
- Much weaker correlation with **GPRC-FIN**
- Local perceptions are:
  - *close to* global perceptions of risk
  - *distinct from* global attention to Finland

Comparing GPR Indices

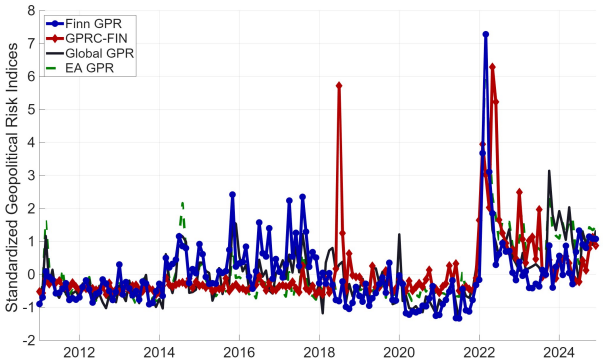

- FinnGPR reacts to major global and regional events
  - peaks at the **2022 Russian invasion of Ukraine**
- GPRC-FIN spikes when Finland is in the global spotlight
  - Trump-Putin Helsinki summit, July 2018
- **Local perceptions vs. international attention**

Outline

- Background
- Constructing a Finnish GPR
  - Dictionary index
  - Language model index
- **Empirical Analysis**
- Conclusions

Empirical Analysis: Local Projections

- Goal: quantify the effects of a **FinnGPR shock** on the Finnish economy and financial markets
- Method: **lag-augmented local projections** (monthly)
- Sample: **Jan 2011 – Nov 2024**
- Outcomes:
  - Industrial production growth, HICP inflation, unemployment
  - Consumer confidence, OMXH25 returns, **systemic stress (CISS)**
- Comparison: responses to **FinnGPR** vs. **Global GPR** vs. **EA GPR**

Empirical Results: Local Projections (Dictionary Index)

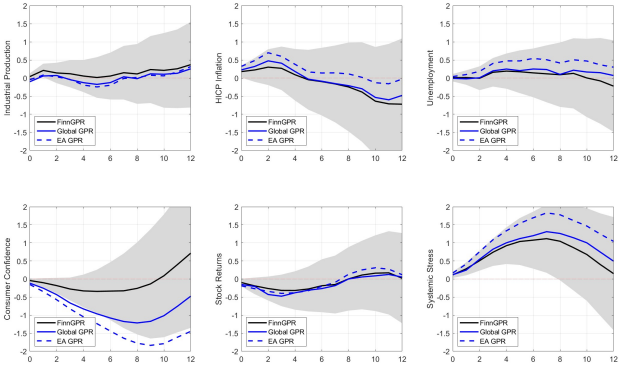

- Real activity: **muted responses**
- Inflation: small, short-lived
- Consumer confidence: declines
- **Systemic stress rises** significantly
- GPRC-FIN (not shown): small / insignificant

Empirical Results: Local Projections (Language Model Index)

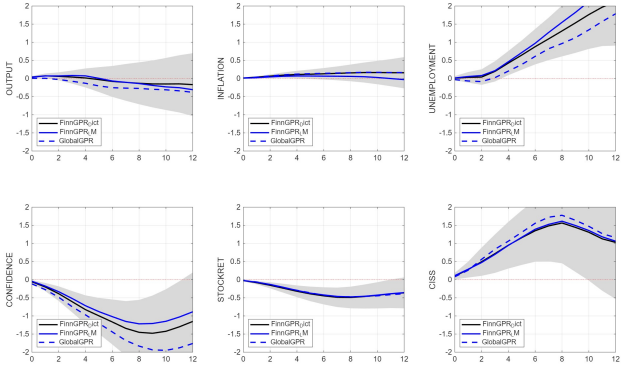

- We repeat the analysis with the **fine-tuned BERT** index
- Responses are **very similar** to the dictionary-based baseline
- Economy resilient; **systemic stress rises**; confidence falls
- Robust across specifications (LP levels/growth, factor-augmented) and BVAR

Empirical Results: Local Projections on Sub-Indices

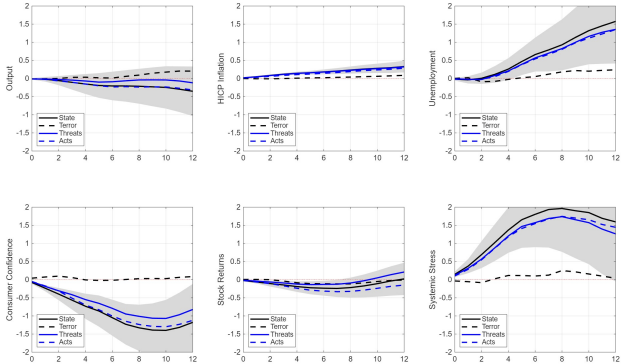

- Effects vary across **FinnGPR sub-indices**
- Threat- and Act-related news drives stress and confidence responses
- **Terror**-related news shows weak macro-financial effects

Interpretation

- Finding: the real economy is largely unaffected, while **systemic stress** rises and **consumer confidence** softens
- Possible explanations for the muted real-economy response:
  - a long history as a neighbour of, and intermediary with, Russia
  - exceptionally high trust in the armed forces (**97%** vs. 76% EU average)
  - strong willingness to defend the country ($\approx$4 in 5 Finns)
- Yet financial markets are **internationally integrated** $\Rightarrow$ not isolated from global contagion (stress rises)

Outline

- Background
- Constructing a Finnish GPR
  - Dictionary index
  - Language model index
- Empirical Analysis
- **Conclusions**

Conclusions

- We construct **FinnGPR**, a local Finnish-language GPR index, using a **dictionary** and a **fine-tuned BERT** model
  - the BERT classifier matches human judgment far better; both yield similar economic results
- FinnGPR differs from the English-language country index:
  - GPRC-FIN is informative about **international attention** to Finland
  - **FinnGPR** captures **local perceptions** of geopolitical risk
  - local perceptions align with **global** GPR and the **euro-area** index, not with GPRC-FIN
- Practical takeaway:
  - to study geopolitical risk *in Finland*, use our **FinnGPR**
  - **GPRC-FIN** is better suited for measuring **international salience**
- Finland's economy is **resilient** to geopolitical risk shocks, but **systemic stress increases**



Northern inSights: Geopolitical Risk from Finnish News Media

Gene Ambrocio, Zuzana Fungáčová, Joni Heikkinen,
Eeva Kerola, Iikka Korhonen, Anni Norring

Eesti Pank Research Seminar, 3 September 2026

The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Bank of Finland, the Eurosystem, or the IMF.

Adaptation to Finnish News: Word Lists

Panel B. Search words

Topic sets	Phrases
War_words	**war OR conflict** OR hostilities OR revolution* OR insurrection OR uprising OR revolt OR coup OR geopolitical
Peace_words	**peace** OR truce OR armistice OR treaty OR parley
Military_words	**military OR troops** OR missile* OR "arms" OR weapon* OR bomb* OR warhead*
Nuclear_bigrams	**"nuclear war*"** OR "atomic war*" OR "nuclear missile*" OR "nuclear bomb*" **OR** "atomic bomb*" OR "h-bomb*" OR "hydrogen bomb*" OR "nuclear test" OR "nuclear weapon*"
Terrorism_words	**terror*** OR guerrilla* OR hostage*
Actor_words	**allie* OR enem* OR insurgen*** OR foe* OR army OR navy OR aerial OR troops OR rebels

Threat/act sets	Phrases
Threat_words	**threat* OR warn* OR fear* OR risk* OR concern*** OR danger* OR doubt* OR crisis OR troubl* OR disput* OR tension* OR imminen* OR inevitable OR footing OR menace* OR brink OR scare OR peril*
Peace_disruption_words	**threat*** OR menace* OR reject* OR peril* OR boycott* OR disrupt*
Buildup_words	**buildup* OR build-up* OR sanction*** OR blockad* OR embargo OR quarantine OR ultimatum OR mobiliz*
War_begin_words	**begin* OR start* OR declar*** OR begun OR began OR outbreak OR "broke out" OR breakout OR proclamation OR launch*
Actor_fight_words	**advance* OR attack* OR strike*** OR drive* OR shell* OR offensive OR invasion OR invad* OR clash* OR raid* OR launch*
Terrorism_act_words	**attack** OR act OR bomb* OR kill* OR strike* OR hijack*

Panel C. Excluded words

Exclusion words	movie* OR film* OR museum* OR anniversar* OR obituar* OR memorial* OR arts OR book OR books OR memoir* OR "price war" OR game OR story OR history OR veteran* OR tribute* OR sport OR music OR racing OR cancer OR "real estate" OR mafia OR trial OR tax

Return

Caldara & Iacoviello (2022) definition of Geopolitical Risk

- Caldara and Iacoviello (2022, p.1195) define geopolitical risk as the "threat, realization, and escalation of adverse events associated with wars, terrorism, and any tensions among states and political actors that affect the peaceful course of international relations."

Return

Full Correlation Table

	Finn GPR-LM	Finn GPR-Dict	GPRC-FIN	Global GPR	EA GPR	Russia GPR
Finn GPR-LM	1.000	0.817	0.112	0.578	0.474	0.201
Finn GPR-Dict	0.817	1.000	0.383	0.798	0.742	0.493
GPRC-FIN	0.112	0.383	1.000	0.510	0.594	0.537
Global GPR	0.578	0.798	0.510	1.000	0.893	0.655
EA GPR	0.474	0.742	0.594	0.893	1.000	0.749
Russia GPR	0.201	0.493	0.537	0.655	0.749	1.000

Sample period: January 2011 – December 2024.

Return

Inside BERT — Illustrative example

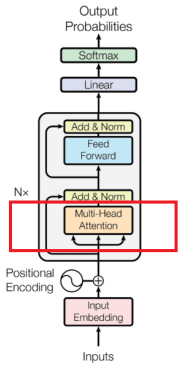

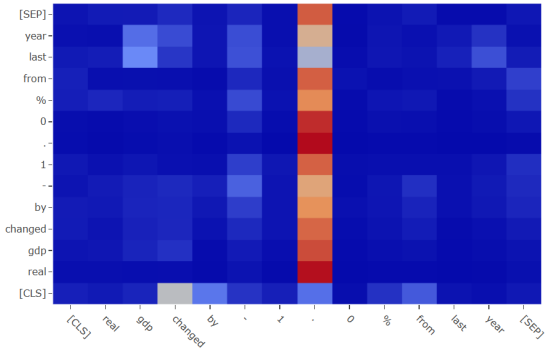

Illustration from a public BERT model

Return

Open-Weight Model Comparison

Model	Overall quality	Tokens/sec
Poro-2-8B-Instruct	**most reliable in Finnish**	7.7
Gemma-3-4B-IT	usable, more errors	13.0
Gemma-4-E4B-IT	weaker than Gemma 3	9.4
Qwen3-8B	unreliable	15.9
Qwen3-4B	unreliable	16.3
Phi-4-mini-instruct	unreliable	23.9

- **14 Finnish tasks** across 8 categories: general knowledge, concept explanation, domain reasoning, grammar, summarization, classification, instruction following, language control
- Scored by **keyword matching plus manual review**; keyword scoring alone was insufficient, since correct answers need not contain the expected terms
- Poro 2 and Gemma 3 tie on keyword score (0.52); Poro 2 wins on manual review

Return