

The economic exposure to wildfires: a firm-level analysis based on innovative hazard mapping

Riccardo Russo
Research Data Centre, Banca d'Italia

13th biennial IFC Conference
Basel, 20-21 August 2026


Wildfires: a major source of physical risk

- Wildfires constitute a major source of physical risk arising from climate change
- 608 fatalities and 18.8 billion dollars damages across EU-27 countries between 2000 and 2025 (EM-DAT)
- 2026 fire season is projected to be one of the worst on record

<https://www.emdat.be>

<https://forest-fire.emergency.copernicus.eu/apps/effis.statistics/seasonaltrend>

EFFIS Weekly Cumulative Burnt Areas



- Unlike floods and landslides, wildfires are not targeted by specific regulations such as the Floods Directive (2007/60/EC)
- ECB's climate-related stress tests indicate that 22% of loans granted within the euro area are directed to firms facing high physical climate risks, **more than half due to wildfires**

Research gap: detailed, firm-level measure of the current and future exposure of Italian businesses to wildfire hazard.

- The core contribution is a spatial match between business locations and high-resolution wildfire hazard maps




- The Business Register provides establishment-level data for our analysis, including:
 - **address information**, which is geocoded to obtain precise geographic coordinates
 - **number of employees** at each establishment, used as a weighting factor to assign relative importance in aggregation procedures

Algorithm: **Random Forest Classifier**

- Inputs: Corine Land Cover, ERA-5, Copernicus
- Target variable: burned area (EFFIS wildfire register)
- Output: **susceptibility** class that is then intersected with fuel type (**fire intensity**)



HAZARD MATRIX		FUEL TYPE			
		1 Grasslands and croplands	2 Broadleaves	3 Shrublands	4 Coniferous and Eucalyptus
SUSCEPTIBILITY	1 Low	1 low intensity surface fires with low likelihood	4 medium intensity forest fires with low likelihood	7 high intensity bushfire with low likelihood	10 very high intensity forest fires with low likelihood
	2 Medium	2 low intensity surface fires with medium likelihood	5 medium intensity forest fires with medium likelihood	8 high intensity bushfire with medium likelihood	11 very high intensity forest fires with medium likelihood
	3 High	3 low intensity surface fires with high likelihood	6 medium intensity forest fires with high likelihood	9 high intensity bushfire with high likelihood	12 very high intensity forest fires with high likelihood



Wildfire exposure

- Not all classes correspond to significant hazard
- Only sites falling into classes 3, 5, 6, 8, 9, 11, 12 are considered at risk
- High hazard: classes 11 and 12

HAZARD MATRIX		FUEL TYPE			
		1 Grasslands and croplands	2 Broadleaves	3 Shrublands	4 Coniferous and Eucalyptus
SUSCEPTIBILITY	1 Low	1 low intensity surface fires with low likelihood	4 medium intensity forest fires with low likelihood	7 high intensity bushfire with low likelihood	10 very high intensity forest fires with low likelihood
	2 Medium	2 low intensity surface fires with medium likelihood	5 medium intensity forest fires with medium likelihood	8 high intensity bushfire with medium likelihood	11 very high intensity forest fires with medium likelihood
	3 High	3 low intensity surface fires with high likelihood	6 medium intensity forest fires with high likelihood	9 high intensity bushfire with high likelihood	12 very high intensity forest fires with high likelihood



1.9%

productive establishments in
wildfire-prone areas

1.4%

overall employment exposed

1.4%

firms with at least one
employee exposed

$\frac{2}{3}$

of exposed firms have their
entire workforce exposed


- Share of exposed firms (1.4%) is lower compared to floods (27%) and landslides (9%)
- However, there are pockets of higher exposure (e.g. Enna province in Sicily with 20% of overall employment subject to wildfires)

Results not shown today:

- regression analysis → identify driver of wildfire exposure (number of sites, tangible assets...)
- survival analysis → firms affected by wildfires in 2021-2023 do not seem to exhibit lower survival w.r.t. control

Final remarks: a look in the future

- Source: *Climate Impact Explorer*, developed by Climate Analytics, ETH Zurich, NGFS and Potsdam Institute



+3%

exposed business sites
under +1.5°C scenario

+11%

exposed business sites
under +3°C scenario

Main message: wildfire exposure is currently modest, but concentrated enough to matter for risk monitoring and local adaptation in view of expected increases



EUROPEAN CENTRAL BANK

EUROSYSTEM

Navigating spatial scales in physical climate risk assessment - from address to regions

BIS IFC conference,
Basel

20-21 August 2026

The views expressed are those of the author and do not necessarily reflect those of the European Central Bank (ECB).

Malgorzata Osiewicz



Motivation for the study

Motivation:

- **Limited location granularity:** Key assets may be known at regional or country level rather than at the precise address, while complementary socio-economic variables such as population, GDP and employment are also typically reported regionally.
- **Mismatched spatial granularity:** for example, debtors may be geocoded at address level, while real-estate collateral or other exposures may only be available at postal-code or NUTS3 level.
- **High computational cost:** Address-level assessment requires spatial matching and processing of large geospatial datasets, which can be resource-intensive and difficult to implement at scale.

Research questions:

- **How well can physical risk at an exact location be approximated using regional aggregates?**
- **Can hazard or risk estimates constructed at different spatial resolutions - e.g. address, postal code and regional level - be combined without introducing systematic bias?**


Methodology for geospatial aggregations

Geospatial aggregations - methodology

The **modifiable areal unit problem (MAUP)** refers to the sensitivity of statistical results to the spatial units used to aggregate geographically referenced data [Openshaw (1984)] :

- **the scale effect:** results change as observations are aggregated into larger spatial units
- **the zoning effect:** areas of broadly similar scale are delineated using different boundaries.

In addition, **aggregation operator** used within them: a regional mean, upper percentile, maximum, or fraction of the area exceeding a hazard threshold can provide substantially different representations of the same underlying hazard measure.




Sources: Own calculations based on OpenStreetMap data for [Hessen](#); historical river-flood hazard maps for the 1979–2000 baseline period from Delft University of Technology (Paprotny, 2016); Eurostat [GISCO](#); and H3 level 7 spatial indexing (Uber Technologies, 2018). Expected annual loss (EAL) is estimated using flood damage functions based on Huizinga (2017).

Notes: Each panel shows average EAL, expressed as a share of exposed asset value, within the corresponding spatial boundaries. EAL is calculated using five flood return periods (10, 30, 100, 300, 1000). Because the 1-in-10-year flood is the most frequent event considered, the theoretical maximum contribution from this return period alone is approximately 0.1: a value of 0.1 would correspond to complete (100%) loss of the exposed asset value in a flood occurring, on average, once every 10 years.

Geographical boundaries: statistical, administrative regions and grids

Grid systems:


- square grid
- hexagons



Sources: Uber Technologies (2018). H3: A Hexagonal Hierarchical Geospatial Indexing System. Documentation available at <https://h3geo.org>



Statistical:

- NUTS3, LAU (Local Administrative Units)
- Administrative: postal codes, municipalities



Note: Flood in Valencia region – event's footprint, versus affected postal codes and NUTS3. Sources: [Copernicus EMS Rapid Mapping EMSR773](#) for event footprint. NUTS3 and postal codes are based on data from [Eurostat GISCO](#).

NUTS3 regions versus postal codes



Sources: Own computations based on Eurostat GISCO portal.

Notes: Box plots show the interquartile range (IQR; 25th–75th percentile), with the median as the horizontal line and average as red dot. Whiskers extend to 1.5×IQR.

Results and conclusions

General framework and specifications

We follow the methodology underlying the **ESCB statistical climate indicators** (for details please see: [The Statistical Paper](#) April 2024 and [Technical annex](#)). The indicators use harmonised public climate datasets covering a broad range of hazards and combine firm-level information, climate hazard data and creditor portfolios. First, the impact of natural hazards on firms' financial health and debt-servicing capacity is assessed. These firm-level risks are then mapped to the portfolios of euro area financial institutions to measure climate-related financial exposure.

Two hazards with distinct spatial characteristics are selected:


- **river flooding** – assessed at 100 m resolution using expected annual loss (EAL) expressed as a share of exposed asset value. EAL is calculated using five flood return periods: 10, 30, 100, 300 and 1,000 years.
- **Consecutive Dry Days (CDD)** – a drought-related indicator available at approximately 12 km resolution, with hazard intensity expressed maximum number of consecutive dry days (with precipitation < 1mm per day) .

The sample covers indebted euro area firms based AnaCredit (covering loans above EUR 25,000).

Results – deviations of regional aggregates from address level estimate


Deviations: address - regional mean

River floods RCP8.5 2071-2100



- For floods, the region-mean proxy may over-estimate firm-level risk relative to the address estimate, as many assets within a region may lie outside the inundated area.
- Address-level precision is not always necessary, particularly for spatially smooth hazards such as temperature and for hazard maps that are themselves available only at coarse spatial resolution.

Consecutive Dry Days (CDD) RCP8.5 2041-2060




Sources: Own calculations based on sample of indebted on-financial corporations resident in the euro area in Register of Institutions and Affiliates Database (RIAD), AnaCredit at reference period December 2025; river flood hazard estimates from Delft University of Technology (Paprotny, 2016); and consecutive dry days (CDD) from the IPCC Interactive Atlas.

Notes: Hazard projections refer to the RCP8.5 climate scenario. CDD projections cover 2041-2060, while river flood projections cover 2071-2100. Box plots show the interquartile range (IQR; 25th–75th percentile), with the median as the horizontal line and average as red dot. Whiskers extend to p05-p95.

Results by country – deviations of regional aggregates from address level


Deviations: address - regional mean

River floods RCP8.5 2071-2100 (unit: EAL share)



- Bias varies substantially across countries and depends strongly on the size of NUTS3 regions and postal-code areas.
- For floods, aggregation can lead to both under- and overestimation. Underestimation is particularly evident in Slovenia, Austria and Ireland, including at postal-code level, while in many countries regional averages overstate risk, most notably for NUTS3 regions in the Netherlands.

Consecutive Dry Days (CDD) RCP8.5 2041-2060 (unit: number of days)



- For CDD, postal codes already provide a close approximation to address-level hazard. At NUTS3 level, the bias is generally limited to about one day, although larger underestimation occurs in several southern European countries.

Conclusions

Aggregation can provide a useful proxy for granular assessment:

- Regional aggregation can introduce bias for selected hazards but can still provide meaningful estimates where precise location data are unavailable.
- Approximation errors can be quantified, helping to identify when regional information is sufficient and when address-level data are needed, particularly for floods.
- It is important to disclose both the spatial precision of asset locations and the resolution of the underlying hazard maps.

Flood-risk overestimation may capture broader exposure channels:

- Regional aggregation may overstate direct flood exposure, as many assets lie outside the inundated area.
- Firms can also be affected indirectly - some overestimation may therefore capture broader regional effects through damage to infrastructure and other local disruptions.
- It should also be noted that current hazard models may underestimate risk, particularly where compound hazards and tipping points are not yet captured.

Addressing Data Gaps in Climate Change and Sustainability Statistics using Generative AI¹

Simon Oehler, Directorate General Data and Statistics, Deutsche Bundesbank

¹ We would like to express our sincere gratitude to the BIS Innovation Hub Eurosystem Centre (Project GAIA) and the Deutsche Bundesbank Directorate General Markets (ICAS) for their support and most fruitful exchanges throughout the course of this project.

Agenda

Motivation

Climate change as macroeconomic and financial risk. Policy makers and researchers demand high-quality sustainability data.

The GIST project

„Greenhouse Gas Insights and Sustainability Tracking” (GIST) project - a collaboration between Ludwigs-Maximilians-Universität München and Deutsche Bundesbank

Results

Assuring data quality using gold-standard data, real world applications

Way forward

Opportunities and challenges ahead

Motivation

01

Climate change poses potentially severe macroeconomic and financial risks

02

Implications for core concerns of central banks (growth, inflation, output ...)

03

Central banks have begun integrating climate-related information, e.g. GHG emissions into climate stress tests etc.


04

Companies publish emissions data via company reports, i.e. text and varying formats

05

commercial data partly inconsistent between vendors

Figure 1: Comparison of scope 1 emissions data provided by ISS ESG and MSCI – reported vs. estimated values at the company level: noticeable differences between vendor's data while variation between data providers is strongly reduced for reported data in comparison to data that is estimated by the data providers.



The GIST project

The Greenhouse Gas Insights and Sustainability Tracking (GIST) Project is a collaboration between **Ludwig-Maximilians-University München** and **Deutsche Bundesbank**, initiated in 2023

- Software and tools for data extraction, preparation and analysis - narrow PDF reports to relevant pages using page-level embeddings and semantic retrieval, apply LLM-based extraction, and post-process the results into a structured database.
- Outputs: extractions of scope 1-3 GHG emissions (CO₂e) after unit normalization, benchmarked against gold-standard dataset, in structured format,
- Evaluation of different sources of errors (typology of data quality issues in reports),
- Provision of open-source Python package ([climatextract](#) · [PyPI](#) and [GitHub](#))

Figure 2: Overview of the RAG-based pipeline, dashed elements indicate optional components.




Table 1 - Examples of raw unit strings, normalized classes, and deterministic conversion factors to t CO₂e.

Raw unit string	Normalized class	Factor
"Tonnes CO ₂ e"	t CO ₂ e	1
"mt CO ₂ e"	t CO ₂ e	1
"kg CO ₂ equivalent"	kg CO ₂ e	0.001
"kilotonnes CO ₂ e"	kt CO ₂ e	1,000
"1,000 t CO ₂ "	kt CO ₂ e	1,000
"million t CO ₂ e"	Mt CO ₂ e	1,000,000
"US tons CO ₂ e"	USt CO ₂ e	0.907185
"kg CO ₂ /kWh"	Other	-

Results I

Evaluation of pipeline outputs (GHG emissions Scope 1-3) against a manually curated gold-standard data set*

Table 2 - Confusion matrix with rows indicating the model's extraction (value vs. NA) and columns indicating ground truth (value present vs. no value/different value present) for the best run with gpt-5.2 and question-style prompt and text context.

	Gold: value present	Gold: no value / different value
LLM extracts value	483	121
LLM extracts no value	50	4916

Table 3 - Averages and standard deviations of 5 runs of prompt and context experiments with gpt-5.2.

Prompt Type	Input Mode	Precision	Recall	F1 Score
Question	Text	0.779 (0.015)	0.878 (0.084)	0.823 (0.039)
Question	Text + table	0.777 (0.021)	0.903 (0.014)	0.835 (0.017)
Structured	text	0.760 (0.008)	0.860 (0.011)	0.807 (0.006)
structured	Text + table	0.755 (0.020)	0.830 (0.105)	0.787 (0.049)

* three-stage annotation procedure: non-expert annotation, expert annotation, and a final round of expert discussion.

Results II

Real world use case in collaboration with Markets Department at Deutsche Bundesbank

- Received set of pdf-files (yearly company reports) for about 2000 German groups and an accompanying dataset with manually extracted Scope1-3 GHG emission values.
- Run extraction pipeline and evaluate LLM extractions vs. manually extracted values:




Figure 3: Comparison of manually and LLM-extracted Scope 1 emissions at the company level.




Figure 4: Comparison of manually and LLM-extracted Scope 1 emissions aggregated by scope.

Way forward

Three levels of uncertainty

- Varying accuracy of emissions information from corporate sustainability reports
- Not fully transparent third-party data provider's approaches
- New emerging risk topics not sufficiently represented in available datasets (e.g. biodiversity)

Way forward

- Broaden scope of companies
- Broaden scope of metrics, depending on availability in reports:
 - Water, electricity usage
 - Asset-level information
 - Biodiversity metrics
 - Transition plans
- Quality measures for data extractions

Thank you for your attention!

Contacts:



lisa.reichenbach@bundesbank.de, maurice.fehr@bundesbank.de,
simon.oehler@bundesbank.de



01 MOTIVATION

1

TRADE-OFF

Increasing demand for timely economic information
Official statistics suffer from substantial publication lags

2

MEANING THAT

In a fast-changing world, information can become outdated
before it becomes official!

3

SOLUTION

Produce timely estimates that remain conceptually consistent
with benchmark accounting statistics



02 CONCEPTUAL FRAMEWORK AND DATA SOURCES

- Timely indicator of corporate profitability: Gross Margin and Gross Operating Surplus:
 - $GM_t = T_t - IC_t$, where GM_t denotes Gross Margin, T_t denotes Turnover and IC_t denotes Intermediate Consumption
 - $GOS_t = GM_t - EE_t$, where GOS_t denotes Gross Operating Surplus and EE_t denotes Employee Expenses
- Econometric estimates for the Turnover, Intermediate Consumption and Employee Expenses used to calculate the Gross Margin and the Gross Operating Surplus

Invoicing Data from Tax Authority



External transactions and positions




Social Security Records




03 WHY ADMINISTRATIVE DATA SOURCES ARE NOT ENOUGH




Total NFCs



Electricity



Construction




KEY TAKEAWAY

Larger differences between administrative and benchmark data at disaggregated levels
Reflect structural differences in coverage, reporting practices and accounting concepts



04 METHODOLOGICAL FRAMEWORK

Five-step bridge-model workflow balances timeliness, benchmark coherence and forecast accuracy.






05 RESULTS

Bridge estimates track the official benchmark closely across all indicators.


Turnover

EUR million · 4-quarter cumulative · Total NFCs




Gross Margin

EUR million · 4-quarter cumulative · Total NFCs



Gross Operating Surplus

EUR million · 4-quarter cumulative · Total NFCs



— Sources — ITENF ● Flash estimation ◆ QoQ final estimation

Close benchmark fit

Estimated series follow the benchmark throughout the sample period.

Timeliness preserved

Bridge models keep the advantage of administrative data while remaining benchmark-consistent.

Flash and final estimates

Both estimates reproduce the evolution of the published ITENF series.




05 RESULTS

Sectoral evidence shows the bridge model adapting its correction to each discrepancy pattern.


Electricity

Gross margin / turnover · 4-quarter cumulative · %




Wholesale & retail trade

Gross margin / turnover · 4-quarter cumulative · %




Manufacturing

Gross margin / turnover · 4-quarter cumulative · %



Construction

Gross margin / turnover · 4-quarter cumulative · %



— Sources — ITENF ● Flash estimation ◆ QoQ final estimation



06 CONCLUSIONS AND FUTURE RESEARCH

MAIN CONCLUSIONS

1

EARLIER SIGNAL

Administrative data sources allow estimates of Gross Margin and Gross Operating Surplus around 1½ months before the official publication date.

2

COHERENT BRIDGE

The bridge-model architecture keeps timely administrative signals coherent with the benchmark ITENF series.

3

HYBRID STATISTICS

Short-term statistics should evolve into hybrid systems that combine surveys, administrative data and statistical models.

FUTURE WORK

- **Annual accounting information**
as an alternative benchmark
- **Production process**
further automation
- **Adaptive modelling**
to handle structural breaks