NUS National University of Singapore | Sustainable and Green Finance Institute

A Multi-Provider Assessment of Asset-Level Flood Risk in the Banking Sector

Lessons from Indonesia for credit risk and supervision

Zhang Feimo

Sevi W. Perwitasari

Chen Mengzhu

Asda J. Pandiangan

Johan Sulaeman

The Sustainable and Green Finance Institute:

The Sustainable and Green Finance Institute (SGFIN) is a research institute established by the National University of Singapore (NUS). SGFIN aims to develop deep research capabilities in sustainable and green finance with a focal point on Asia, and to provide thought leadership and shape sustainability outcomes in policymaking across the financial sector and the economy at large. Supported by exceptional domain experts across NUS, SGFIN will equip businesses with critical cross-disciplinary knowledge, training, and toolkits to integrate sustainability dynamics into their business strategies and investment decisions to better quantify the environmental and social impacts of their business developments, operations, products, and services. In essence, SGFIN aims to help companies embed sustainability as a key pillar in their business decisions.

Cover Image Acknowledgement:

AI-generated content is used in the creation of this cover.

Recommended citation:

Zhang, F., Perwitasari, S. W., Mengzhu, C., Pandiangan, A. J., & Sulaeman, J. (2026). *A Multi-Provider Assessment of Asset-Level Flood Risk in the Banking Sector*. SGFIN.

Keywords: *Climate Risk, Physical Risk, Flood Risk, Pluvial Flood, Fluvial Flood, Climate Stress Testing, Credit Risk, Asset-level Analysis, Hazard Data Comparison, Loan Books, Syndicated Loan, Indonesian Banking*

Sustainable and Green Finance Institute (SGFIN)

National University of Singapore

Innovation 4.0, 3 Research Link

#02-02 Singapore 117602

Editorial Correspondence should be addressed to SGFIN@nus.edu.sg.

Key Takeaways:

1. ***Flood risk is credit-material, especially when hazard overlaps with capital-intensive borrower exposure and regional financial density.***

Power generation and mining assets are major contributors to projected flood-related losses. The key driver is not only flood intensity and frequency, but their interaction with asset value, capital intensity, operational criticality, bank exposure, and the regional clustering of assets and credit. For project finance and asset-backed lending, flood risk is directly relevant to collateral recoverability, business interruption, and obligor concentration. This is especially important in economically dense regions such as Java and Bali, where corporate assets, commercial credit, and economic activity are highly concentrated.

2. ***Asset-level analysis reveals facility-level risks that generate wider and more decision-relevant loss distributions.***

Company-level estimates compress the geographic and capital-intensity heterogeneity that drives projected flood losses. After applying high-resolution JBA hazard data to asset-location data and aggregating the results into company-level projected loss estimates, the interquartile range of company-level projected losses is 0.52 percent of loan value, almost triple the corresponding company-level range from S&P, at 0.18 percent. This suggests that company-level views based on coarser hazard inputs may understate loss variation, especially for capital-intensive borrowers.

3. ***Flood-related losses are more concentrated than the underlying credit exposure, and the degree of concentration is provider-sensitive.***

The top 5 percent of loss-driving assets account for more than 90 percent of projected system loss, far exceeding their share of credit exposure. Flood risk therefore does not simply mirror the loan book's existing concentration structure; it amplifies it. Because this amplification varies across hazard providers, single-provider analysis may understate the concentration of tail risk.

4. ***No single hazard dataset provides a complete view of tail-risk concentration.***

The three hazard data providers show limited agreement on which assets are most exposed. Out of nearly 6,000 assets flagged as high risk by at least one provider, only fewer than 7 percent, or 391 assets, are flagged as high risk by all three. This divergence carries through to materially different dollar loss estimates, meaning single-provider screening may miss up to two-thirds of the pooled upper tail.

5. ***Multi-provider analysis identifies blind spots while stress-testing the same core exposed banks.***

Although provider views differ at the asset level, the main absorbing institutions remain relatively stable. The four large domestic commercial banks absorb around 40 percent of top-quartile losses under all three views, making multi-provider analysis useful for identifying asset-level blind spots while stress-testing the same core group of exposed institutions.

Foreword

Uncertainty is not an obstacle to physical climate risk analysis. It is the subject of it. Banks and supervisors working to incorporate flood risk into credit frameworks do not face a shortage of data. They face a more difficult problem. Multiple datasets can each be internally consistent, grounded in defensible methodology, and yet produce materially different answers about where tail risk sits and how large it is. The practical question is not which estimate is correct, but how to make that uncertainty visible, structured, and usable for credit decisions and supervisory monitoring.

This study is our attempt to answer that question for the Indonesian banking system. It builds directly on our previous whitepaper, which established that flood-related losses are material at the portfolio level and that concentration can amplify physical risk exposure in ways that aggregate metrics do not fully capture. That analysis drew on company-level estimates. The natural next step was to ask what happens when the same portfolio is examined at the asset level, where individual facilities, not corporate averages, become the unit of analysis, and when the hazard inputs are varied across independent providers.

The results illustrate both the value and the limits of any single provider view. Under three different hazard dataset providers, system-wide mean projected losses are broadly comparable across providers, ranging from 0.33 percent to 0.50 percent, but this convergence at the mean masks more consequential divergence in the tail. Each provider individually captures only 27 to 47 percent of the pooled tail loss estimate of USD 920 million for Indonesian banks' portfolios. The top-quartile loss region is dominated by power generation and mining assets under all three views, but the degree of concentration differs materially across providers. In some views, a small number of assets drive the tail. In others, the tail is spread across a broader set of assets. The projected losses will be mainly absorbed by large domestic commercial banks, as shown by all data providers.

These are not simply findings about which dataset is better or worse. They are findings about the structure of physical climate risk uncertainty and what it implies for how banks should screen exposures, how supervisors should challenge internal models, and how policymakers should think about data standards and disclosure. The framework developed here does not resolve that uncertainty. It makes it legible. We hope this study is useful to banks incorporating asset-level physical risk into credit workflows, to supervisors developing expectations around provider sensitivity and stress testing, and to the broader community working to build more robust climate risk infrastructure for the financial system.

Prof. Sumit Agarwal

Managing Director, SGFIN

Low Tuck Kwong Distinguished Professor of Finance

Professor of Economics and Real Estate

President of Asian Bureau of Finance and Economic Research

Executive Summary

Flood risk is already a credit-risk issue for Indonesian banks, not only a long-term sustainability concern. Because corporate assets, economic activity, and bank credit are concentrated in flood-prone regions, physical flood exposure can affect collateral values, business continuity, project-finance performance, correlated borrower exposures, and portfolio credit quality.

This study builds on our earlier whitepaper on climate risk in Indonesian bank loan portfolios. This paper moves the earlier portfolio-level analysis to the asset level. By linking large-corporate borrower exposure to georeferenced assets and comparing flood hazard outputs from three data providers (S&P, JBA Risk Management, and Swiss Re), the study identifies where borrower-level averages may hide facility-level risk. The purpose is not to select a single "correct" flood-loss estimate, but to show how provider choice, spatial resolution, and aggregation assumptions change the identification of risky assets, borrowers, and banks. The resulting framework helps identify asset-level blind spots, testing sensitivity to provider choice, and making physical climate risk more usable for credit risk assessment, supervision, and targeted action. Asset-level resolution identifies the specific exposures driving risk, rather than implying blunt and often infeasible sector-wide reductions in lending.

The main analysis focuses on inland flood risk, namely fluvial and pluvial flooding, because these hazards form the core comparable dataset across providers. Coastal flooding is also examined as an exposure concern but is not converted into loss estimates in this paper. Other hazards, such as drought, wildfire, and typhoons, are outside the empirical scope, although the same asset-level framework can be extended to them.

A central challenge is that hazard providers and banks hold different pieces of the risk equation. Hazard providers document flood hazard and often report modelled financial impact, but banks hold the credit exposure, collateral, and borrower-level information needed to translate physical damage into credit risk. This study bridges this gap by constructing transparent analytical proxies: vulnerability functions mapping hazard scores to indicative asset-level losses, and relative asset-value weights that aggregate those losses to borrowers and bank portfolios. These analytical proxies are designed for stress testing and comparison, not as engineering damage functions.

The exposure results show a clear pattern: inland flood risk rises over time and is geographically concentrated. Under a high-emission pathway, the number of current assets located at high inland flood-risk areas would almost double from roughly 2,500 today to about 4,700 by 2100. Java and Bali contain the largest numbers of high-risk assets, while Kalimantan has the highest exposure rate and the greatest per-asset risk intensity across flood types. Coastal flooding, although excluded from loss estimates, is also material: around 1,850 current assets exposed to high coastal flood risk today, concentrated along the developed northern coast of Java and Bali.

The most important differences across providers appear in both the level and distribution of projected losses. System-wide mean projected losses range from 0.33 to 0.50 percent of loan books under the baseline climate scenario, meaning the highest provider estimate is roughly 50 percent larger than the lowest. The differences become even more pronounced in the dispersion of company-level losses. Company-level estimates based on higher-resolution asset-level hazard data reveal wider and more decision-relevant loss distributions than company-level estimates based on coarser hazard data: the interquartile range is 0.18

percent using S&P's coarser hazard data, but nearly triples when higher-resolution hazard data are applied to asset locations and then aggregated to the company level, reaching 0.52 percent under JBA and 0.55 percent under Swiss Re. For large electricity and mining borrowers, some individual assets sit well above the corporate mean, indicating that borrower-level averages may understate facility-level loss potential in capital-intensive sectors.

Provider choice materially affects which assets are identified as high risk. For each provider, we define high-risk assets as those in the top 10 percent of that provider's flood-risk distribution. Using this provider-specific cutoff, the two high-resolution datasets (JBA and Swiss Re) agree only moderately on which assets are high risk, whereas the third, coarser-resolution dataset (S&P) diverges more substantially. Only 391 assets are flagged as high risk by all three providers, a small fraction of the nearly 6,000 assets flagged as high risk by at least one provider. This disagreement carries through to tail-loss estimates: single-provider tail-loss estimates range from USD 245 million to USD 432 million, less than half of the pooled-provider upper-bound estimate of USD 920 million for Indonesian banks' portfolios. Yet, the main institutions absorbing the projected losses remain relatively stable: the four large domestic commercial banks jointly account for around 40 percent of top-quartile losses under all three provider views.

The findings reinforce our earlier whitepaper's conclusion that portfolio concentration can amplify physical risk, while showing the mechanism at the asset level. Flood exposure becomes credit-relevant when high-hazard locations overlap with large borrower exposures and capital-intensive asset clusters, especially in economically dense regions such as Java and Bali. For banks, our proposed framework can serve as a second-layer screen for project finance, collateral-backed lending, and large corporate exposures. For supervisors, the key question is not whether a bank has chosen the "correct" provider for climate risk analysis, but whether it understands how risk rankings and tail-risk estimates change under alternative providers and aggregation assumptions. The practical value of this framework is to make physical climate-risk uncertainty visible enough to manage, comparable across provider views, and actionable for credit risk management, stress testing, and supervisory monitoring.

About the Authors

Zhang Feimo is a Research Associate at the Sustainable and Green Finance Institute (SGFIN). His research interests include climate risk assessment, impact assessment for the built environment and the renewables industry. He holds a Bachelor of Business Administration in Business Analytics and Finance and an M.Sc. in Finance, both from NUS Business School.

Sevi Wening Perwitasari is a Visiting Research Associate at the Sustainable and Green Finance Institute (SGFIN). Her research interests span climate-risk and impact assessment, scenario analysis, and sovereign transition financing, with a particular focus on emerging-market policy environments. She holds a Bachelor of Applied Accounting from the Polytechnic of State Finance STAN and a Master of Science in Sustainable and Green Finance from NUS Business School.

Chen Mengzhu is a Research Fellow at the Sustainable and Green Finance Institute (SGFIN). Her research focuses on climate risk modelling, with particular emphasis on flood hazards, extreme rainfall analysis, and the integration of forward-looking climate scenarios into financial risk assessment. She holds a PhD in Hydrology from the University of Liverpool and an MSc in Water and Environmental Management from the University of Bristol. Before joining the SGFIN, she was a postdoctoral research fellow at NUS Department of Civil and Environmental Engineering. Her work has been published in peer-reviewed journals on hydrology and natural hazards.

Asda J. Pandiangan is a Research Associate at the Sustainable and Green Finance Institute (SGFIN). His research interests include sustainability reporting frameworks, social impact valuation, and credit risk modelling in agriculture. He holds a Bachelor of Accounting from the University of Indonesia and an M.Sc. in Applied Finance from Singapore Management University (SMU).

Johan Sulaeman is the Director of SGFIN. He is also a Dean's Chair and a Professor in the Department of Finance at the NUS Business School. He also serves as a Scientist at AIDF (Asian Institute of Digital Finance). He has been involved in designing and delivering executive programmes for large companies and financial institutions in the region, primarily on topics related to sustainability, finance, and technology. His current research focuses on corporate social and environmental performance. His research has appeared in top economics and business journals, and been covered in various international publications, including The Wall Street Journal and The New York Times. He also contributes opinion pieces to the Straits Times and Channel News Asia.

Sustainable and Green Finance Institute

Contents

List of Figures

List of Appendix Figures

List of Tables

List of Abbreviations

Abbreviation	Full Term
AAL	Average Annual Loss
ADB	Asian Development Bank
AIDF - CRI	Asian Institute of Digital Finance - Credit Research Initiative
BCA	Bank Central Asia
BCBS	Basel Committee on Banking Supervision
BICS	Bloomberg Industry Classification Standard
BNI	Bank Negara Indonesia
BNPB	Badan Nasional Penanggulangan Bencana (National Disaster Management Agency)
bps	Basis Points
BPD	Bank Pembangunan Daerah (Regional Development Bank)
BRI	Bank Rakyat Indonesia
CCS	Carbon Capture and Storage
CDB	China Development Bank
CMIP6	Coupled Model Intercomparison Project Phase 6
CRST	Climate Risk Stress Testing
DBS	Development Bank of Singapore
DFI	Development Finance Institution
EAD	Exposure at Default
EBA	European Banking Authority
ECB	European Central Bank
EL	Expected Loss
ENSO	El Niño Southern Oscillation
GCM	Global Climate Model
GDP	Gross Domestic Product
HHI	Herfindahl-Hirschman Index
ICT	Information and Communication Technology
IMF	International Monetary Fund
IPCC	Intergovernmental Panel on Climate Change
IPP	Independent Power Producer
IQR	Interquartile Range
IRBI	Indeks Risiko Bencana Indonesia (Indonesia Disaster Risk Index)
ISIN	International Securities Identification Number
JBIC	Japan Bank for International Cooperation
JETP	Just Energy Transition Partnership
KEXIM	Export-Import Bank of Korea
LGD	Loss Given Default
LNG	Liquefied Natural Gas
LTS-LCCR	Long-Term Strategy for Low Carbon and Climate Resilience
MJO	Madden-Julian Oscillation
MRI-ESM2-0	Meteorological Research Institute Earth System Model version 2.0
MSCI	Morgan Stanley Capital International
MUFG	Mitsubishi UFJ Financial Group
NACE	Nomenclature of Economic Activities
NDC	Nationally Determined Contribution

NGFS	Network for Greening the Financial System
NUS	National University of Singapore
OCBC	Oversea-Chinese Banking Corporation
OECD	Organisation for Economic Co-operation and Development
OJK	Otoritas Jasa Keuangan (Financial Services Authority)
OPEC	Organization of the Petroleum Exporting Countries
PD	Probability of Default
PLN	Perusahaan Listrik Negara
PPA	Power Purchase Agreement
PRA	Prudential Regulation Authority
RCP	Representative Concentration Pathway
RUKN	Rencana Umum Ketenagalistrikan Nasional (National General Electricity Plan)
RUPTL	Rencana Usaha Penyediaan Tenaga Listrik (Electricity Supply Business Plan)
SGFIN	Sustainable and Green Finance Institute
SIC	Standard Industrial Classification
SMBC	Sumitomo Mitsui Banking Corporation
SNDC	Second Nationally Determined Contribution
SSP	Shared Socioeconomic Pathway
TPI	Transition Pathway Initiative
TRS	Transition Readiness Scorecard
UOB	United Overseas Bank
USD	United States Dollar
UNFCCC	United Nations Framework Convention on Climate Change
WRI	World Resources Institute

Glossary

Climate Scenario

A structured set of assumptions describing possible future pathways for climate policy, emissions, technology, physical hazards, and macroeconomic variables. Scenarios (e.g., NGFS, IPCC) are used to assess how different futures would affect financial risks.

Coastal Flood

Flooding caused by abnormal sea-level conditions, including storm surge, sea-level rise, or extreme tidal events in coastal areas. It primarily affects low-lying zones near the shoreline.

Credit Risk

The risk of financial loss arising from a borrower's failure to meet its obligations in full and on time. Credit risk depends on the borrower's probability of default, the exposure at default, and the loss given default.

Dollar Loss

Monetary loss on an asset: impact × asset value. Aggregates to firm- and portfolio-level dollar losses.

(Flood) Hazard

The physical flood condition affecting an asset location, typically described by the frequency, severity, and type of flooding. In this study, flood hazard is represented through provider-specific hazard scores for specific flood types, scenarios, and time horizons.

(Flood) Impact

See Projected Asset Loss

(Flood) Risk

Expected loss from flooding: Hazard Score × Vulnerability × Asset Value. Equals the Projected Asset Loss in percentage, or Dollar Loss in levels.

Fluvial Flood

Flooding that occurs when rivers or streams overflow their banks due to heavy rainfall or upstream runoff.

Gini Coefficient

A concentration metric ranging from 0 to 1, where 0 indicates an even distribution and 1 indicates maximum concentration. In this study, the Gini coefficient is computed for both credit exposure and projected flood losses across the asset inventory. A loss Gini that exceeds the exposure Gini indicates that projected flood losses are more narrowly concentrated than the underlying credit exposure.

Hazard Score (Flood Score)

Provider-specific measure of flood hazard at an asset's coordinates, by flood type, scenario, and year, incorporating both frequency and severity, at normalised scales. Supplied by S&P Sustainable1, JBA, and Swiss Re.

HHI (Herfindahl-Hirschman Index)

A measure of concentration computed as the sum of squared portfolio shares, multiplied by 10,000. Higher values indicate greater concentration. Its reciprocal, scaled by 10,000, gives the effective asset count. In this study, HHI is applied to both credit exposure and projected flood losses to assess whether loss concentration exceeds exposure concentration across providers.

Inland Flood

Flooding driven by river overflow (fluvial flooding) or surface water accumulation from intense rainfall (pluvial flooding), occurring in areas away from the coast.

Interquartile Range (IQR)

The difference between the 75th and 25th percentiles of a distribution, representing the spread of the middle 50 percent of observations. In this study, IQR is used to assess heterogeneity in projected loss estimates across assets, borrowers, sectors, or providers. A wider IQR indicates greater dispersion; a narrower IQR indicates more homogeneous outcomes.

Loss

In this paper, refers to either percentage or dollar loss.

Lorenz Curve

Plot of the cumulative share of a quantity (y-axis) against the cumulative share of units ranked smallest to largest (x-axis); the gap from the 45-degree equality line measures inequality of the distribution and defines the Gini coefficient.

Physical Risk

Climate-related risk arising from the interaction of climate hazards with the exposure and vulnerability of assets, operations, or borrowers. Physical risk can result from acute events, such as floods, storms, and heatwaves, or chronic changes, such as sea-level rise and long-term temperature increases. In this study, physical risk is measured either as an annualised percentage loss of asset value, referred to as Projected Asset Loss, or as a monetary loss, referred to as Dollar Loss.

Pluvial Flood

Surface flooding caused by intense rainfall that overwhelms drainage systems, even when rivers are not overflowing. Often occurs in urban areas with poor drainage.

Probability of Default

The likelihood that a borrower will fail to meet its debt obligations over a given time horizon (for example, within one year). It is usually expressed as a percentage (e.g. 2 percent PD).

Projected Asset Loss

The annualised projected loss on an asset due to flooding hazards, expressed as a percentage of asset value, for specific checkpoint years (e.g., 2050, 2080) under defined SSP scenarios. At the asset level it is the vulnerability function applied to the hazard score; at the borrower level it aggregates the borrower's assets weighted by relative asset value; at the sector and bank-portfolio levels it is the exposure-weighted average of borrower-level losses. (S&P Sustainable1 reports this quantity as "financial impact").

Relative Asset Value

The estimated value of an individual asset proxied on the asset (sub)-type level, expressed as a multiple of the asset value of a headquarters office. Converts a percentage impact into a monetary loss or is used to aggregate asset-level metrics to company-level by means of a weighted average.

Supervisors

Regulatory and oversight authorities responsible for the prudential management of financial institutions. In this study, the term encompasses microprudential supervision, carried out by Otoritas Jasa Keuangan (OJK), and macroprudential oversight, carried out by Bank Indonesia (BI). Where used without qualification, it refers to both functions collectively.

Vulnerability (Function)

Asset-type-specific, generally non-linear mapping from a hazard score to a fractional loss of asset value.

Introduction

I. Introduction

1. Flood risk as a Material Physical Credit Risk in Indonesia

Physical climate risk is no longer a distant threat and has become an active credit risk factor for banks. Floods, storms, and extreme heat impair borrowers' operations, erode asset values, and weaken debt repayment capacity, transmitting physical damage directly onto lender balance sheets. [1] The scale of this exposure is substantial. MSCI Institute estimates that, under the Network for Greening the Financial System (NGFS) Current Policies scenario (approximately 3°C warming), average annual loss (AAL) across 9,409 companies in the MSCI Climate Change Metrics universe was approximately USD 1.2 billion in 2024 and rises to USD 4.6 billion in 2050, a nearly four-fold increase. [2] Firms exposed to higher climate risk may face higher interest rates, stricter collateral requirements, and elevated default probability. [3] Banks have also been shown to price flood exposure specifically into lending conditions, charging higher interest rates on loans to firms in flood-prone areas. [4] At the same time, climate risk reduces banks' willingness to lend, compressing loan supply particularly for institutions in high-exposure areas. [5]

Indonesia presents an acute case of flood-related credit risk. The country ranks among the most exposed globally, with a large share of its population located in low-elevation coastal zones [6] and significant economic activity overlapping with both coastal and riverine flood-prone areas. [7] For the Indonesian banking system, flood risk becomes a structural exposure embedded across credit portfolios. Yet the extent to which flood risk is reflected in credit pricing and portfolio risk management across Indonesian banks remains unclear, underscoring the need for frameworks that can make this exposure legible and actionable.

Prior portfolio-level evidence corroborates this assessment. Our previous research on Indonesian banks, which considered broader flood-related impacts including coastal exposure, indicates that projected asset loss due to floods is not isolated shocks, but a persistent, forward-looking stress factor embedded within bank exposures. [8] Under the baseline climate scenario, large domestic commercial banks face projected asset losses of approximately 0.4 percent annually by 2060, with risk accumulation intensifying thereafter (See Figure 1).

Figure 1 Projected asset loss trajectory for Indonesian banks' portfolios under SSP2-4.5

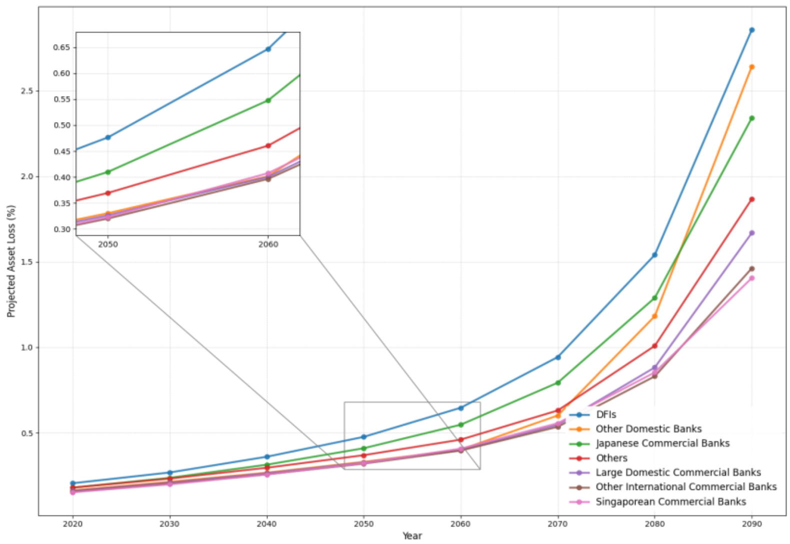

Note: DFIs refer to JBIC, ADB, KEXIM, and CDB. Japanese Commercial Banks are SMBC, MUFG, and Mizuho. Large Domestic Commercial Banks are Mandiri, BCA, BNI, BRI. Singaporean Commercial Banks refer to OCBC, DBS, and UOB.

The prior analysis further suggested that more adverse climate pathways further amplify these losses nonlinearly, and distributional analysis reveals that losses are concentrated among a subset of obligors exposed disproportionately to flood hazards, even when mean portfolio losses appear moderate. [9] This implies that credit stress is more likely to emerge through localised borrower clusters than through broad sectoral deterioration, a pattern with direct implications for credit concentration monitoring and stress testing.

Taken together, the evidence positions flood risk not as an environmental externality, but as a structural credit risk channel requiring proactive integration into bank risk frameworks.

2. Limitations of Sectoral Aggregation in Physical Risk Assessment

Traditional portfolio stress testing often evaluates climate exposure through sector-level averages. While sector classification provides a useful high-level screening tool, it is structurally insufficient for capturing the distributional mechanics of physical risk. Flood exposure is fundamentally location- and asset-specific, meaning that firms within the same sector can face dramatically different hazard intensities depending on where their operations are situated. Sectoral aggregation can compress this heterogeneity and, in doing so, understates the role of borrower concentration in shaping portfolio outcomes. The academic literature is consistent on this point: assessing physical climate risk requires granular information on asset location and geography. [10]

Evidence from our prior portfolio-level research also supports this interpretation. Its sectoral variance decomposition showed that mining and electricity contributed materially to portfolio-loss dispersion, but differences across major domestic banks could not be explained by sector composition alone. Given their broadly similar sectoral mixes, the dispersion was interpreted as reflecting obligor-level concentration across geolocations with varying flood exposures. [11] This is consistent with recent UBS asset-level evidence showing that physical-risk dispersion becomes more pronounced as sector classifications become more granular, with broad sector averages obscuring substantial variation across sub-industries and asset-intensive business models. [12]

Furthermore, tail analysis sharpens this finding. As shown in Figure 2, the sectoral composition of tail losses appears broadly similar across many banks, yet their overall tail loss exposure differs substantially. [13] Sector classification alone cannot explain this divergence: what determines whether a bank faces higher or lower tail losses within the same sector is the geographic distribution of its borrowers' underlying assets, not the sector label itself. Sectoral aggregation obscures the underlying geographic heterogeneity, systematically underestimating extreme but plausible stress scenarios. Effective physical risk assessment therefore requires moving beyond sectoral proxies towards asset-level exposure mapping, where hazard intensity, borrower characteristics, and credit concentration can be evaluated jointly.

Recent methodological surveys of Climate Risk Stress Testing (CRST) corroborate this finding at the system level. Reinders, Schoenmaker, and Van Dijk (2025) identify overreliance on macro models with low sectoral and spatial granularity as one of the central shortcomings of current practice, noting that both transition and physical risks have asymmetrical sectoral and spatial impacts that macro-level approaches cannot adequately resolve. [14] The asset-level approach developed in this study is designed to address this methodological gap.

Figure 2 Tail loss composition in Indonesian banks' portfolios

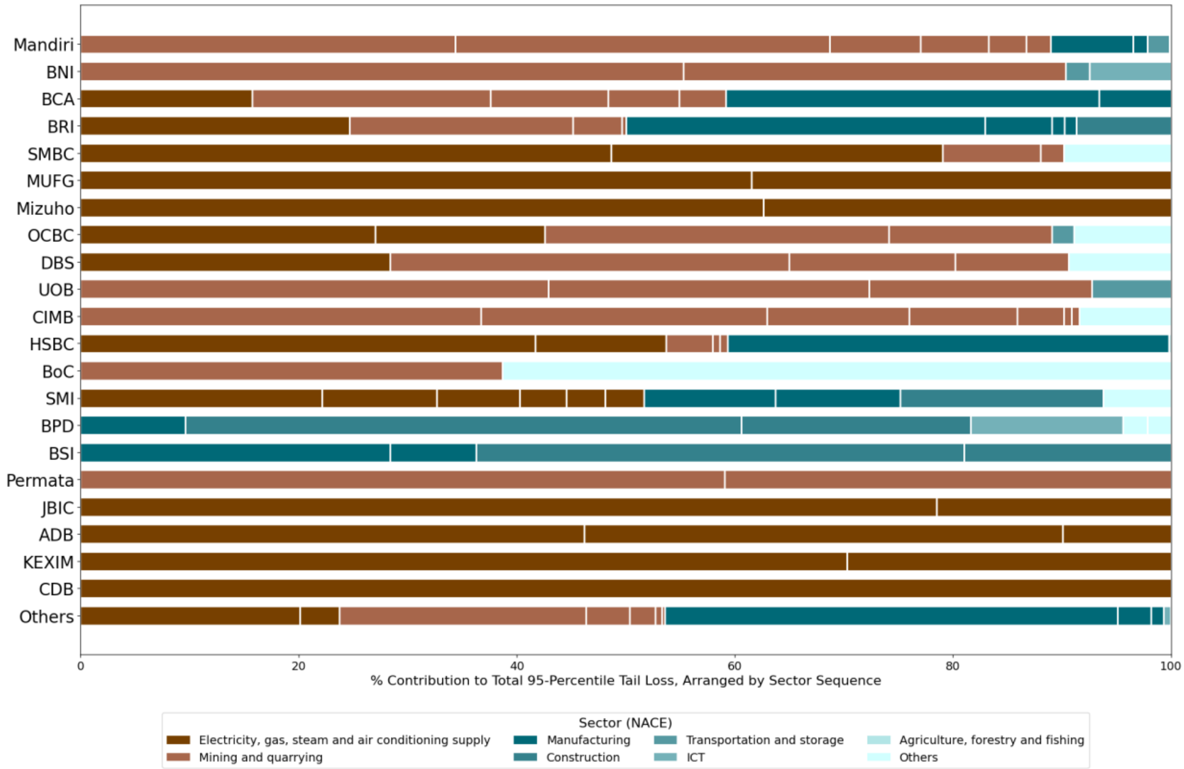

3. Granular Asset-level Analysis for Flood Risk Management

The limitations of sectoral aggregation establish a clear methodological requirement: physical risk frameworks must be capable of linking hazard directly to the assets underpinning borrowers. Asset-level analysis provides this capability by mapping flood hazards to individual borrower operations and aggregating the resulting impacts into portfolio metrics. Infrastructure dependencies, hazards, and the interaction between hazard and vulnerability assumptions, produce differentiated credit stress estimates rather than assumed sector-uniform outcomes.

This granular approach improves risk visibility along three dimensions. First, it identifies concentration risk arising from clusters of borrowers operating in hazard-prone regions, allowing banks to distinguish between diversified sector exposure and spatially correlated vulnerability. Second, it reveals tail-driven portfolio dynamics, where a small number of high-impact assets disproportionately influence stress outcomes. Identification on the asset level allows for targeted actions for risk reduction, as compared to blunt and often infeasible sector-level reductions. Third, it supports forward-looking scenario analysis by tracing how hazard escalation alters borrower-level financial resilience over time, enabling earlier recognition of accumulating risk than is possible under static sector-level assessments.

Importantly, asset-level modelling does not replace sector analysis but refines it. Sector classification remains useful for screening and capital allocation decisions, while granular exposure mapping ensures that sector-level signals are interpreted through the lens of borrower- and asset-specific risk. This integration reduces the likelihood of underestimating extreme outcomes and strengthens the analytical foundation for provisioning, portfolio steering, and supervisory stress testing. In doing so, it aligns flood risk assessment with established credit principles: exposure concentration, obligor sensitivity, and scenario-based stress evaluation. This approach embeds flood hazard metrics within the credit risk frameworks that banks already use.

However, asset-level mapping alone does not resolve all measurement challenges. Physical flood risk estimates remain sensitive to the hazard datasets used, because providers differ in spatial resolution, hydrological assumptions, climate scenario treatment, and calibration choices. They also depend on how hazard intensity is translated into credit-relevant loss metrics, particularly when commercial vulnerability estimates are produced through proprietary models. This makes provider sensitivity central to asset-level flood risk assessment. A credible framework should therefore make explicit how differences in hazard datasets and loss-scaling assumptions affect borrower rankings, tail-risk estimates, and portfolio concentration signals, without treating any single provider output as definitive.

4. Contribution of This Study

This study contributes to the emerging literature and supervisory practice on physical climate risk assessment by developing an asset-level, multi-dataset framework for integrating inland flood risk into bank credit risk analysis in Indonesia. While physical climate risk is increasingly recognised as financially material, its incorporation into credit risk frameworks remains constrained by three practical limitations: the aggregation of risk at sectoral or national levels, reliance on single-source hazard datasets — partly reflecting the cost and operational complexity of integrating multiple commercial providers — and limited transparency over vulnerability assumptions embedded in commercial physical risk models.

The first contribution of this study is to demonstrate why asset-level analysis is necessary for bank credit risk assessment. Flood risk is inherently spatial, and its financial implications depend not only on borrower sector or domicile but also on the precise location, asset type, and concentration of physical assets. Sector- or country-level approaches can compress this geographic heterogeneity and may create a misleading impression of homogeneity across borrowers. By constructing a georeferenced asset inventory for Indonesian-domiciled companies using S&P Sustainable1 asset-level data, this study enables borrower risk to be assessed at the level where physical hazards are actually realised and provides a basis for evaluating how asset-level analysis changes the distribution of estimated flood-related credit losses relative to more aggregated approaches.

The second contribution is to introduce a multi-dataset approach for assessing the robustness of flood risk estimates. Commercial and specialist hazard providers differ in modelling assumptions, spatial resolution, hazard calibration, and treatment of climate scenarios. This study overlays the same asset inventory with inland flood hazard data from multiple providers, including S&P Sustainable1, JBA Risk Management and Swiss Re, and evaluates the consistency of high-risk asset identification across providers, scenarios, and time horizons. By comparing provider outputs, the study provides a basis for assessing whether reliance on a single hazard source may create blind spots in the identification of high-risk assets, obligors, or regional concentrations.

The third contribution is methodological. This study develops a transparent, asset-type-specific loss-vulnerability framework that bridges externally sourced flood hazard scores and bank-specific credit exposure values. Because commercial providers embed vulnerability assumptions within their integrated outputs without disclosing them as standalone parameters, this study derives study-specific proxies from observable hazard-impact relationships in licensed outputs. These proxies are not intended to replicate any provider's proprietary model. Rather, they provide a replicable analytical structure that can be applied consistently across alternative hazard datasets, allowing provider divergence to be isolated and interpreted without dependence on any single provider's integrated output.

The fourth contribution is to connect physical flood hazard assessment with regional concentration risk in the Indonesian banking context. Building on the portfolio-level evidence in Zhang et al. (2025), this study examines this issue through an asset-level and multi-provider lens. The empirical application allows the study to examine how flood hazard overlaps with the geographic concentration of corporate assets. This is particularly relevant for economically dense regions such as Java and Bali, where physical exposure and financial concentration may reinforce each other. The study therefore frames physical climate risk not only as an environmental hazard, but also as a potential source of regional concentration risk within bank loan portfolios.

This study extends the physical risk analysis presented in Zhang et al. (2025), which estimated baseline flood-induced asset losses for Indonesian bank portfolios using S&P Sustainable1 firm-level impact data directly. In contrast, this study estimates asset-type-specific vulnerability parameters from observed hazard-impact patterns and applies them to independent flood hazard datasets. As a result, the estimates are not intended to be directly comparable with the previous whitepaper. Differences in results reflect differences in methodology, hazard sources, spatial treatment, and the use of inferred vulnerability parameters rather than direct firm-level impact estimates.

Taken together, these contributions provide Indonesian banks, regulators, and supervisors with a replicable framework for incorporating physical flood risk into credit risk assessment. The framework is operationally relevant because it uses available asset and hazard data, recognises uncertainty across providers, allows application to bank-specific exposure values, and highlights the regional concentration channels through which physical climate risk can affect credit portfolios. While the empirical application focuses on inland flood risk, the broader framework can be extended to other hazards, including drought, wildfire, coastal inundation, and extreme heat, where these hazards are material to specific sectors or regions of the Indonesian archipelago.

NUS National University of Singapore | Sustainable and Green Finance Institute

Conceptual Framework

II. Conceptual Framework: From Assets to Portfolio Risk

Translating physical hazard exposure into portfolio-level credit risk requires bridging three analytical domains that are often treated separately in conventional credit analysis: the geographic and operational characteristics of borrower assets, the spatial distribution of flood hazard under climate scenarios, and the financial mechanisms through which asset impairment propagates to loan performance.[15] The framework proposed here integrates these domains into an asset-to-portfolio risk chain, enabling physical hazard intensity to be traced from individual assets to borrower-level exposure and, ultimately, to portfolio-level credit outcomes.

The central methodological challenge is that standard credit exposure measures, such as loan amounts and collateral values, capture financial magnitude but not the spatial and physical conditions that determine flood-related loss potential. The same loan extended to two borrowers in the same sector can carry materially different flood risk depending on asset location and facility type, a distinction that aggregate financial data cannot resolve. Addressing this requires augmenting credit exposure records with georeferenced asset information and flood hazard intensity estimates, and a study-specific vulnerability proxy. This allows indicative loss to be approximated as a function of hazard, asset-type vulnerability, and exposure value. Figure 3 presents the three-layer framework that operationalises this approach.

Figure 3 Flood risk assessment conceptual framework

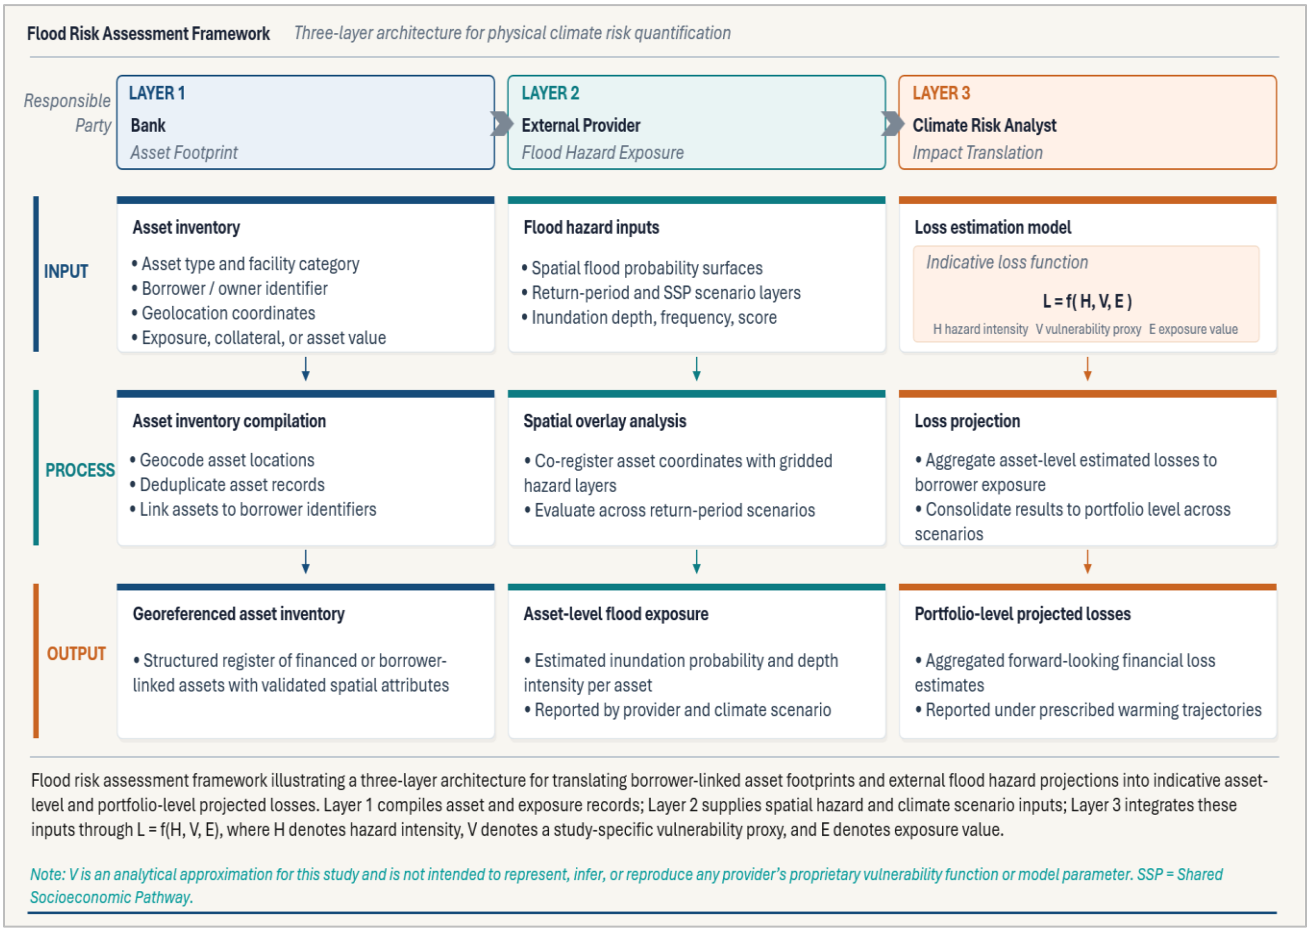

1. Asset Footprint

The asset footprint layer establishes the georeferenced inventory of bank-financed or borrower-linked assets that serves as the foundational input for all subsequent risk assessment steps. This layer answers three essential questions: what assets the bank is exposed to, how much they are worth and where those assets are located. Sector-level assessment is insufficient for capturing physical risk because hazard exposure depends fundamentally on geographic location. Regulatory stress-testing exercises emphasise that physical risks cannot be adequately represented through sectoral averages alone, as hazard varies materially across asset locations. [16] For each asset, the framework requires four categories of information that jointly determine its relevance for asset-level flood risk assessment (See Table 1). Together, these attributes form the minimum information set required to connect physical hazard exposure with borrower-level financial exposure. Once the asset inventory is established, the next step is to characterise the flood hazard conditions associated with each asset location under the relevant climate scenario and hazard dataset.

Table 1 Asset information categories

Data Element	Description	Importance
Asset Type	Classification of the physical asset or facility, such as a power plant, factory, office, warehouse, or port facility.	Supports the assignment of asset-type loss sensitivity or vulnerability proxy.
Asset Owner	Legal entity that owns, operates, or is financially linked to the asset.	Enables aggregation from asset level to borrower or obligor level.
Asset Geolocation	Precise latitude and longitude coordinates of the asset.	Enables spatial overlay with flood hazard maps.
Exposure Value	Exposure value, asset value, replacement cost, collateral value, outstanding loan amount, or another credit exposure proxy.[1]	Provides the financial weight used to scale indicative asset-level loss estimates.

[1] *Where direct asset valuations are unavailable, exposure value may be proxied using replacement cost, collateral valuation, outstanding loan amount, or another bank-specific exposure measure, subject to data availability and modelling purpose.*

2. Flood Hazard Exposure

2.1. Types of flood hazards

Indonesia experiences some of the highest rainfall levels globally due to its tropical maritime climate and strong ocean–atmosphere interactions across the Indonesian Maritime Continent. Annual precipitation across much of the archipelago commonly exceeds 2,000 mm, with national averages approaching 3,000 mm per year, although substantial spatial variability exists between regions [17]. Rainfall patterns are strongly influenced by the Asian–Australian monsoon system, which drives seasonal wet and dry periods across the country. During the boreal winter monsoon, moisture transport from the surrounding warm tropical oceans leads to intensified convection and widespread rainfall across western and central Indonesia [18]. In addition to monsoon dynamics, rainfall variability is strongly modulated by large-scale climate drivers such as the Madden–Julian Oscillation (MJO) and the El Niño–Southern Oscillation (ENSO), both of which can significantly influence the frequency and intensity of extreme precipitation events across the region [19].

These climatic conditions contribute to frequent flood hazards across the country. In combination with steep river catchments, rapid urbanisation, land-use change, and limited drainage capacity in many cities, extreme rainfall events often translate into widespread flooding across Indonesian river basins and urban areas. As a result, flood hazards represent one of the most significant natural disaster risks in Indonesia. In this paper, flood hazards are characterised by three primary mechanisms: fluvial flooding, pluvial flooding, and coastal flooding, which together represent the dominant pathways through which climate-related hazards translate into physical and economic losses.

Fluvial flooding (river flooding) occurs when rivers or streams overflow their banks following prolonged or high-intensity rainfall, often influenced by upstream land-use change, watershed degradation, and basin-scale hydrological processes. These events typically develop at the scale of river basins and may affect extensive downstream floodplains. Global risk assessments indicate that Indonesia faces a high exposure to river flooding, and climate change is expected to increase this risk due to more frequent heavy rainfall days and more intense precipitation extremes [20]. Fluvial floods pose a direct threat to economic assets located along major river systems, including manufacturing clusters, agro-processing facilities, and resource extraction sites, increasing the potential for physical asset damage, inventory loss, and business interruption.

Pluvial flooding (surface water flooding) occurs when intense rainfall exceeds the infiltration capacity of soils or the capacity of urban drainage systems, leading to water accumulation on the land surface independent of river overflow. This type of flooding typically develops rapidly during short-duration rainfall events and is often highly localised. Modelling results from the World Bank indicate that Indonesia is also highly exposed to urban flood risks. Disaster records from Indonesia's National Disaster Management Agency (BNPB) show that major cities such as Jakarta, Bandung, Makassar, Surabaya, Palembang, Samarinda, and Denpasar frequently experience pluvial flood events. Because these floods often affect dense commercial districts, industrial parks, and residential areas, they can lead to repeated short-cycle disruptions to economic activity and local infrastructure.

Coastal flooding occurs when abnormal rises in sea water levels, driven primarily by tropical cyclone winds and pressure gradients, inundate low-lying coastal areas beyond the reach of inland river systems. As the world's largest archipelago with a coastline spanning approximately 54,000 kilometres across more than 17,000 islands, Indonesia faces extensive exposure to this hazard, particularly across low-lying deltaic plains and densely populated coastal cities. Although Indonesia lies largely outside the primary tropical cyclone tracks, it is regularly affected by cyclones forming in the surrounding Western Pacific and Indian Ocean basins, which can drive significant storm surge events along exposed coastlines, further compounded by astronomical tides and wave run-up. Major coastal cities including Jakarta, Semarang, Surabaya, and Makassar are particularly vulnerable, as rapid land subsidence driven by excessive groundwater extraction and urban development further amplifies effective sea level at the local scale. Economic assets concentrated in coastal industrial zones, port facilities, and export-oriented manufacturing clusters face direct exposure to these compounding risks.

Considering Indonesia's tropical rainfall regime, extensive coastline, and the spatial distribution of economic assets, all three flood types are included in this analysis to provide a comprehensive characterisation of the physical flood risk landscape affecting asset locations across the country.

2.2. Metrics for flood hazard assessment

Flood risk assessments commonly describe flood hazard conditions using two key metrics: flood intensity and flood frequency. Flood intensity describes the severity of a flood event and is commonly represented by flood depth or water level. Flood frequency refers to how often a flood of a given magnitude is expected to occur and is typically expressed using return periods. Together, these metrics provide a standardised way to quantify the potential flood hazard affecting a specific location.

Flood depth represents the height of water above ground level during a flood event. It is one of the most important indicators of flood intensity because physical damage to buildings, infrastructure, and equipment generally increases with water depth. Flood depth is typically expressed in metres and is estimated for specific flood events with defined probabilities of occurrence.

Flood frequency is commonly expressed using return periods, which represent the average intervals between flood events of a given magnitude. For example, a 100-year return period refers to a flood event that has a 1 percent probability of occurring in any given year. Flood hazard maps often provide estimates of flood depth for several return periods, such as 20, 50, 100, 200, or 500 years. These estimates help analysts understand how flood intensity changes as the rarity of events increases.

While flood depth and return periods describe individual flood events, many climate risk assessments also use composite flood hazard indicators that combine information about flood frequency and intensity into a single metric. In this report, **Flood Scores** are used to summarise flood hazard exposure at a location by integrating information on both the likelihood and severity of flooding. This composite indicator provides a simplified representation of flood exposure and enables rapid comparison of flood hazards across a large number of asset locations.

2.3. Flood risk under climate change

Flood risk is expected to evolve over time as climate change alters precipitation patterns and the frequency of extreme rainfall events. In tropical regions such as Indonesia, climate projections indicate that warming temperatures may intensify the hydrological cycle, leading to an increase in heavy rainfall events and potentially greater flood hazards in many areas [21]. These changes may affect both fluvial flooding, through higher river discharge during extreme rainfall events, and pluvial flooding, through more intense short duration rainfall that can exceed the capacity of urban drainage systems [22].

To assess potential future changes in flood hazard, climate risk assessments typically use climate projections derived from Global Climate Models (GCMs) participating in the most recent Coupled Model Intercomparison Project Phase 6 (CMIP6). These models simulate future climate conditions under different assumptions about global socioeconomic development, greenhouse gas emissions, and climate policy.

The climate scenarios in CMIP6 are represented through Shared Socioeconomic Pathways (SSPs). Each SSP describes a different pathway of global development, including assumptions about population growth, economic development, energy use, and climate mitigation policies. These pathways lead to different greenhouse gas emission trajectories and therefore different levels of global warming by the end of the century [23].

This study has considered flood hazard across four CMIP6 scenarios spanning a range of potential future climate outcomes: SSP1-2.6, SSP2-4.5, SSP3-7.0, and SSP5-8.5. This paper focuses on SSP2-4.5 and SSP5-8.5, which represent a moderate and a high-end warming pathway respectively, as these two scenarios are most widely used in climate risk assessments and together bracket the plausible range of future flood risk outcomes relevant to financial decision-making.

Table 2 Overview of CMIP6 Shared Socioeconomic Pathways used in this analysis

Scenario	Approximate Global Warming by 2100	Socioeconomic Pathway	Emission Level
SSP1-2.6	~2°C	Sustainability pathway with strong mitigation and low emissions	Low
SSP2-4.5	~2.5–3°C	Middle-of-the-road development with moderate emissions	Medium
SSP3-7.0	~3–4°C	Regional rivalry with slower economic and technological development	High
SSP5-8.5	>4°C	Fossil-fuel intensive growth with very high emissions	Very high

(Adapted from IPCC Sixth Assessment Report [24])

The scenarios shown above do not represent predictions of future climate conditions. Instead, they provide a framework for exploring how climate risks may evolve under different development and emissions pathways. Flood hazard datasets used in this report incorporate these scenarios by adjusting the frequency and intensity of flood events based on projected changes in precipitation and hydrological conditions derived from climate model simulations. While the specific modelling approaches differ between data providers, the use of common CMIP6 scenarios provides a consistent framework for evaluating how flood hazards may evolve under future climate conditions.

3. Impact Translation

The final layer translates asset-level hazard exposure into financially meaningful credit-risk metrics. Once an asset is associated with a given level of flood hazard, the key analytical question is how that exposure translates into indicative financial loss, and how such losses propagate through borrower-level and portfolio-level credit assessments.

In conventional credit-risk modelling, expected loss is driven by three parameters: probability of default (PD), loss given default (LGD), and exposure at default (EAD). Physical climate risk transmits to credit risk through two of these parameters, PD and LGD. Flood-induced asset impairment can reduce operating capacity and revenue generation, thereby weakening debt-servicing ability and increasing PD. Simultaneously, physical damage may erode collateral value or impair asset recoverability, influencing LGD. Even in the absence of immediate default, heightened earnings volatility and refinancing pressure can alter forward-looking credit assessments.

The framework acknowledges these transmission channels but does not attempt to estimate separate PD and LGD adjustments. Instead, it produces an indicative annual flood-induced loss estimate that can inform credit risk assessment, stress testing, and exposure concentration analysis. This approach is appropriate for the present application because the flood hazard inputs already combine information on hazard frequency and intensity.

At the asset level, this loss estimate is obtained by applying a vulnerability function to the hazard score. The vulnerability function is an asset-type-specific, generally non-linear mapping that converts a given level of hazard into a fractional loss of asset value, expressed as

$$\mathrm{I} = f(\mathrm{H}, \mathrm{V}),$$

where H is the hazard score at the asset's location and V the associated vulnerability response. Asset-level impact therefore depends not only on the severity of the hazard but also on how sensitive a particular asset type is to it.

Determining appropriate vulnerability proxies for different asset types therefore represents a central modelling challenge, one that the broader stress testing literature has explicitly acknowledged but not yet resolved for bank lending portfolios.[25] Assets with comparable exposure values may exhibit materially different sensitivity to flooding depending on construction standards, operational flexibility, supply-chain dependencies, and recovery capacity. The framework addresses this through transparent, study-specific assumptions across asset categories, developed in the vulnerability discussion below.

Translating these asset-level impacts into borrower- and portfolio-level metrics requires a further aggregation step. Banks typically observe exposure values at the borrower or facility level, so the conversion requires explicit assumptions about relative asset value, asset-type sensitivity, and the role of each asset in borrower cash-flow generation. As illustrated in Figure 4, asset-level impacts are first combined into borrower-level projected losses using relative asset values as weights, and borrower-level losses are then aggregated to the portfolio level using credit exposure as weights.

Figure 4 Asset to portfolio flood impact aggregation

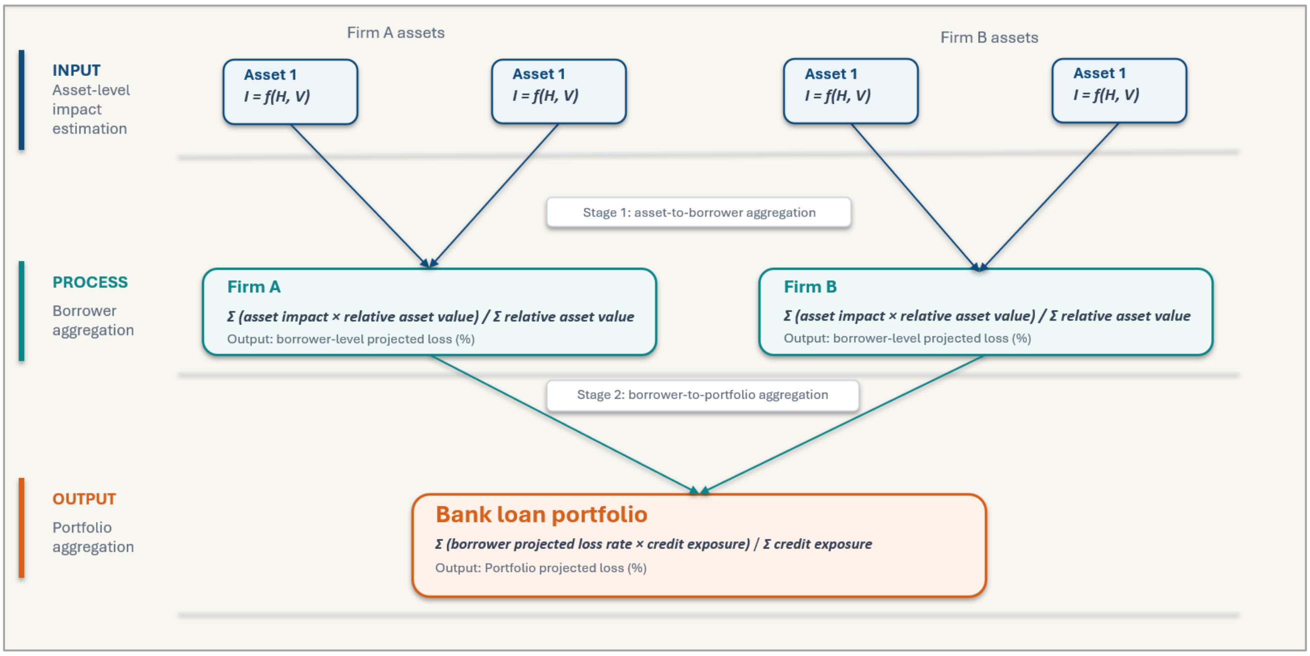

NUS National University of Singapore | Sustainable and Green Finance Institute

Asset Data Overview

III. Asset Data Overview

1. Ideal Asset-level Data

As discussed in the asset footprint section, effective physical risk assessment at the borrower level requires asset-level data that satisfies four minimum conditions. First, asset geolocation is essential for spatial overlay with flood hazard maps. Without location data, hazard intensity cannot be assigned to individual assets, and the analysis collapses to sector or country averages. Second, asset type informs the vulnerability proxy applied to the asset: a power plant, warehouse, and retail outlet located at identical coordinates may respond differently to the same flood hazard, and this differentiation is only possible when asset classification is known. Third, asset owner enables loss estimates to be aggregated from the asset level to the borrower level, and subsequently to the portfolio level. Fourth, asset value provides the financial scaling parameter used to translate indicative impairment rates into monetary loss estimates.

In practice, bank loan documentation typically captures borrower identity, collateral location, facility type, and financial exposure value. The binding constraint is therefore not data availability alone, but the need for assumptions that translate hazards into asset-specific financial loss estimations. Knowing that a borrower operates a warehouse in a flood-prone area does not, by itself, yield a loss estimate. It requires an additional assumption about how flood hazards may translate into indicative impairment and business disruption for that type of asset. S&P Global Sustainable1 describes physical risk financial impact as combining a hazard exposure score with an asset type-specific loss response profile and for company-level impact, aggregating through assumed asset value.[26] Building on this general structure, this study develops study-specific vulnerability function proxies for analytical comparison, loss analysis, and stress-testing purposes.

2. Asset Dataset from S&P Sustainable1

This study draws on the S&P Sustainable1 dataset as the primary source of asset-level data. The S&P Sustainable1 Physical Risk dataset assigns scores for climate physical hazards to more than 6.8 million corporate assets and assesses financial impact across nine hazards, including coastal flood, fluvial flood, and pluvial flood. The dataset used in this study provides georeferenced asset locations linked to corporate owners and asset-type classifications across sectors including industrial, energy, mining, transport, and urban environment facilities.[27]

While banks may hold many of the minimum asset-level data requirements within their internal credit and collateral records, actual asset valuation data are not consistently available through external commercial sources, including financial data providers and physical risk vendors. What S&P commercial datasets do provide is a broad coverage, georeferenced, asset-type-classified information for Indonesian-domiciled companies, including asset-level flood score and flood financial impact estimates. After applying relevant filters, the dataset used in this study contains 27,572 [a] assets from 948 companies at 26,841 unique coordinates, spanning 34 different asset types and 59 different subtypes (See Figure 5).

The dataset is employed for two purposes. First, it is used to construct a georeferenced asset inventory that can be overlaid with flood hazard maps, supporting the asset footprint and flood

[a] The number of assets, 27,572, measures all the assets used for vulnerability function proxy, relative asset weight estimation. Part VI Portfolio Inland Flood Risk Analysis uses a subset of 27,572 which is covered by all data providers. More details can be found in VI.1

hazard exposure layers of the framework. Second, consistent with S&P's description that firm-level financial impact combines asset location hazard scores, asset type-specific loss response profiles, and asset value for firm-level aggregation, observable relationships in the licensed outputs are used to construct study-specific vulnerability function proxies for analytical comparison and stress-testing purposes. These proxies are designed to be applied consistently across alternative hazard datasets and, where available, to bank-specific exposure values or a broader multi-bank borrower universe to produce system-level portfolio loss estimates. The resulting proxies should not be interpreted as estimates of any provider's proprietary vulnerability function or model parameter.

Figure 5 Share of asset location by broad category

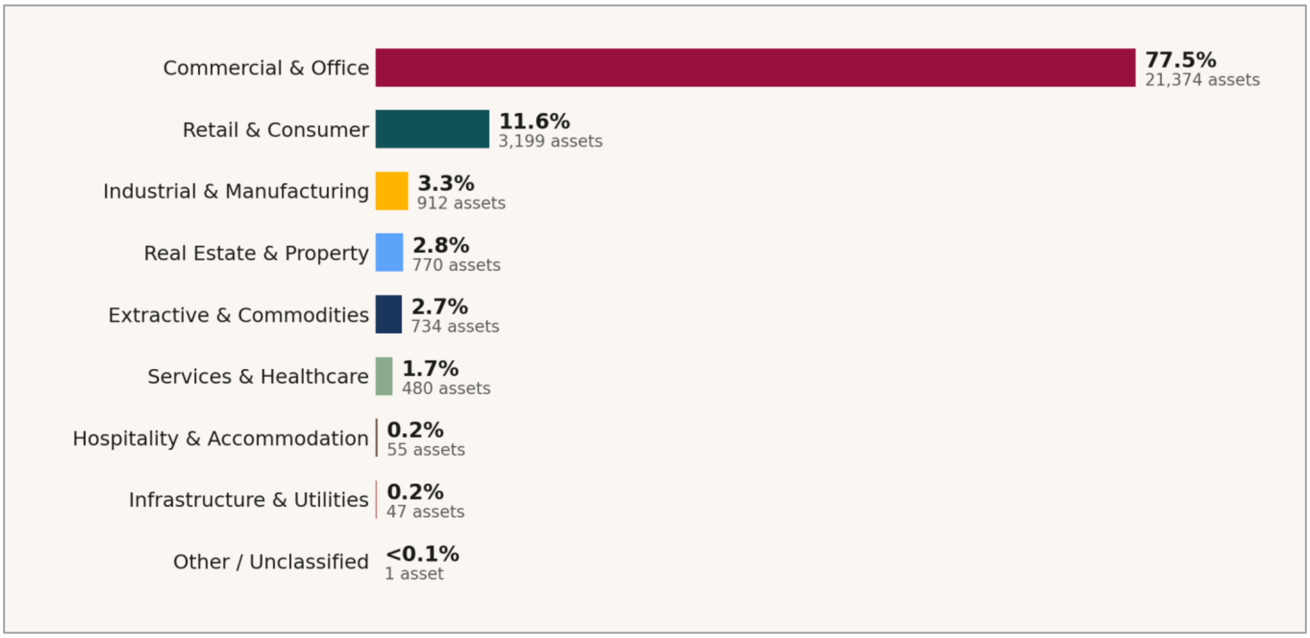

Note: Asset categories are grouped from the underlying asset-type classifications based on asset function and sectoral relevance. Percentages show the share of asset locations in the analytical sample.

As shown in Figure 5, Commercial and Office assets account for the largest share of recorded locations in the analytical sample. This likely indicates both the prevalence of service-sector companies in the S&P coverage and the relative ease of identifying these types of assets from public disclosure. Retail and Consumer assets form the second-largest group, while physically intensive categories, including industrial, infrastructure, utilities, and extractive assets, represent smaller shares of recorded locations. Their lower location share should not be interpreted as lower financial materiality, as these assets may exhibit higher flood-related loss sensitivity, greater operational criticality, or larger exposure values.

Low-frequency asset categories are therefore retained in the analytical dataset rather than excluded solely on the basis of location count. Commercial physical risk datasets typically apply different modelling treatments depending on the depth and completeness of available asset-level information. In this study, low-frequency categories are retained where they contribute to the observed asset composition of a corporate entity and are treated with appropriate data-quality controls. Additional filters are applied to reduce the risk of over-interpreting sparse observations, as discussed in Section 3.

The dataset's primary strength for this application lies in its georeferenced, asset-type-classified inventory, structured at a scale and level of standardisation that enables systematic portfolio-level analysis across a broad borrower universe. The coverage relies on S&P's proprietary asset database, which draws on company filings, regulatory disclosures, and sector-specific datasets across multiple industry sources. As with any commercially assembled dataset, coverage completeness cannot be guaranteed, particularly for smaller private companies with limited public disclosure. The results should therefore be interpreted as conditional on the asset coverage available within the S&P database. In the Indonesian context, this caveat is particularly relevant for domestic companies and state-owned enterprises whose asset disclosures may be limited relative to their economic significance within bank loan portfolios.

NUS National University of Singapore | Sustainable and Green Finance Institute

Flood Hazard Data

IV. Flood Hazard Data

Assessing physical climate risk at the asset level requires flood hazard data that is spatially precise, methodologically transparent, and capable of representing both present-day conditions and future climate trajectories. This study draws on flood hazard data from three leading global providers: JBA Risk Management (JBA), Swiss Re, and S&P Global Sustainable1 (S&P) to assess flood risk exposure across a portfolio of Indonesian assets. Each provider offers both present-day baseline flood hazard data and forward-looking projections under a range of climate scenarios aligned with the IPCC Sixth Assessment Report, covering fluvial (river), pluvial (surface water), and coastal flood types. While all three datasets address the same underlying physical hazards, they differ in their modelling approaches, spatial resolution, output metrics, and scenario coverage. Table summarises the key characteristics of the three datasets.

Table 3 Overview of flood hazard datasets

	JBA	Swiss Re	S&P
Flood Types	Fluvial, Pluvial	Fluvial, Pluvial, Coastal	Fluvial, Pluvial, Coastal
Score Range	0–20	0–10	0–100
Spatial Resolution	30 m (fluvial and pluvial)	30 m (fluvial, pluvial and coastal)	1 km (fluvial); 25 km (pluvial); 90 m (coastal);
Flood Defence	Included	Not included	Not included
Suitable for (Use case)	Asset-level analysis	Asset-level analysis	Portfolio-level screening

1. JBA Risk Management

JBA Risk Management is a specialist flood risk consultancy with deep expertise in global flood modelling. For this study, two JBA products are used: the Global Flood Data (GFD), which represents present-day flood conditions, and the Global Climate Change Flood Data (GCCFD), which projects how flood hazard may change under future climate scenarios. Together, they cover fluvial and pluvial flooding at a high spatial resolution of 30 metres, providing a detailed picture of flood exposure at the individual asset level. Present-day flood hazard is derived from a combination of hydrological modelling, which estimates flood flows and rainfall intensities at a series of return periods, and 2D hydraulic modelling, which simulates the spatial extent and depth of flooding across the landscape. The resulting flood maps are validated against recorded flood events to ensure reliability. The key output used in this study is the Flood Score, a single numerical index (scored 0–20 per flood type, or 0–40 combined) that summarises the overall flood hazard at a given location, taking into account the full range of possible flood intensities. Higher scores indicate greater flood hazard. Notably, JBA is the only dataset in this study to incorporate flood defence information: both undefended and defended Flood Scores are available for fluvial flooding, with the defended score accounting for the standard of protection offered by existing flood defence structures at each location.

Future projections are produced by applying climate-driven adjustment factors, derived from the MRI-ESM2-0 CMIP6 global climate model to the baseline flood maps rather than re-running the full model chain. This approach ensures that present-day and future Flood Scores remain directly comparable. Projections are available across five SSP scenarios (SSP1-1.9 through SSP5-8.5) and multiple time horizons out to 2100. For Indonesia, the dataset projects a general increase in both river and surface water flood hazard under warming scenarios, consistent with the expected intensification of the regional monsoon system.

2. Swiss Re

Swiss Re is one of the world's leading reinsurance companies, with extensive expertise in natural catastrophe risk modelling. For this study, two Swiss Re products are used: Global Flood Zones, which characterise present-day flood hazard, and Climate Risk Scores, which quantify projected changes in flood hazard under future climate scenarios. These layers are combined into a Future Hazard classification, which provides a qualitative rating from Very Low to Extreme based on both current flood exposure and projected climate change effects.

The Swiss Re dataset covers three flood types relevant to Indonesia: fluvial flooding, pluvial flooding, and storm surge, with the latter used in this study as a proxy for coastal flood hazard. All layers are produced at 30-metre spatial resolution. Present-day flood zones delineate flood extents for return periods ranging from 50 to 500 years (and up to 1,000 years for storm surge), and all layers adopt an undefended view, meaning flood protection infrastructure is not assumed to reduce exposure, ensuring a consistent assessment of the underlying physical hazard. Each flood type is modelled using a different approach suited to its physical characteristics: a geomorphological regression method for fluvial flooding, a hydrodynamic simulation for pluvial flooding, and a probabilistic tropical cyclone model for storm surge.

The Climate Risk Scores are expressed on a standardised scale of 0–10, derived from an ensemble of 6–10 CMIP6 global climate models, and are available for four SSP scenarios (SSP1-2.6 through SSP5-8.5) at regular intervals from 2030 to 2085. Model uncertainty is reported as the spread across the ensemble, providing a measure of confidence alongside the central estimate. The combined Future Hazard classification integrates present-day exposure with projected climate change to provide an intuitive, asset-level summary that is well-suited to portfolio-level risk screening.

3. S&P Global Sustainable1

S&P Global Sustainable1 is a leading provider of sustainability data and analytics for financial institutions. Its climate physical risk dataset covers a broad range of climate hazards globally and is specifically designed to support financial risk assessment, including integration into credit risk and portfolio analysis workflows. For this study, three flood-related components are used: fluvial flooding, pluvial flooding, and coastal flooding, applied to Indonesian asset locations.

Flood hazard is modelled using a combination of global climate model output and established hydrological datasets. Fluvial flood projections are derived by combining statistical relationships between river discharge and climatological variables, driven by an ensemble of 35 CMIP6 models downscaled to approximately 25 km resolution, with projected flood extent data from the WRI Aqueduct dataset at approximately 1 km resolution. Pluvial flood hazard is similarly derived from the same CMIP6 ensemble. Coastal flood hazard is modelled by combining historical storm-tide levels with sea-level rise projections, using topographic data at 90-metre resolution for Indonesia.

The primary output used in this study is the Physical Risk Exposure Score, a normalised index from 1 to 100 representing a location's flood exposure relative to global conditions, where a score of 100 reflects the highest level of exposure observed anywhere in the world. Scores are available across four SSP scenarios (SSP1-2.6 through SSP5-8.5) and for each decade from the 2020s through to the 2090s, enabling long-range scenario analysis. In addition to exposure scores, the S&P dataset also provides Financial Impact estimates, expressed as a percentage of asset value, which are particularly valuable for translating physical flood risk into terms directly relevant to loan portfolio management and credit risk assessment.

However, the S&P dataset has important limitations for asset-level flood hazard assessment. Its coarse spatial resolution, particularly for fluvial (~1 km) and pluvial (~25 km) flood components, limits the ability to represent localised flood conditions at individual asset locations. In the context of this study, where assets are densely distributed across Indonesian cities and industrial zones, multiple assets may fall within a square kilometre or larger grid cell, resulting in identical exposure scores regardless of their actual position within that cell. For this reason, S&P's flood data is more suitable for high-level portfolio screening, benchmarking, and financial impact interpretation, rather than as the primary dataset for differentiating flood hazard at individual asset locations.

4. Cross-dataset Consistency Assessment

To assess whether the three datasets provide consistent spatial signals of flood hazard, Spearman rank correlation coefficients were computed for all pairwise dataset combinations across each scenario and time horizon. Spearman correlation compares the relative ranking of assets across datasets rather than their absolute flood scores, making it suitable for cross-provider comparison, given that each dataset uses a different scoring scale. The analysis focuses on fluvial flood scores, as fluvial flooding represents the dominant flood risk in Indonesia and accounts for the largest share of exposed assets across all datasets and scenarios.

Table 4 Spearman rank correlation between flood hazard datasets

Scenario	Time Horizon	Swiss Re vs JBA	Swiss Re vs S&P	JBA vs S&P
Baseline	Historical	0.45	-	-
SSP2-4.5	2030	0.43	0.21	0.23
SSP2-4.5	2050	0.43	0.22	0.24
SSP2-4.5	2080	0.43	0.22	0.24
SSP5-8.5	2030	0.43	0.22	0.24
SSP5-8.5	2050	0.43	0.22	0.24
SSP5-8.5	2080	0.42	0.22	0.24

The results in Table 4 show that Swiss Re and JBA provide the most consistent hazard signals. Their Spearman rank correlations are moderate and remain stable across scenarios and time horizons, ranging from 0.42 to 0.45. This indicates that, despite differences in modelling approach and score construction, the two high-resolution datasets tend to rank Indonesian assets in a broadly similar order of flood hazard. The moderate level of correlation also suggests that the datasets are not interchangeable, but they provide a reasonably consistent indication of relative flood exposure.

By contrast, S&P shows weaker agreement with both Swiss Re and JBA. The correlation between Swiss Re and S&P ranges from 0.21 to 0.22, while the correlation between JBA and S&P ranges from 0.23 to 0.24. This weaker consistency is likely related to S&P's coarser spatial resolution, 1 km × 1 km for fluvial flooding. At this spatial resolution, localised variations in flood hazard may be smoothed, reducing the dataset's ability to differentiate flood exposure among closely located assets.

The correlations remain relatively stable across scenarios and time horizons, with minimal variation between SSP2-4.5 and SSP5-8.5, and across the 2030, 2050, and 2080 time horizons. This stability suggests that cross-dataset disagreement is driven mainly by differences in baseline methodology, spatial resolution, and hazard metric design, rather than by differences in climate adjustment methodologies.

Figure 6 High-risk asset overlap across flood hazard datasets (Pairwise-only overlaps are shown separately from the three-way overlap)

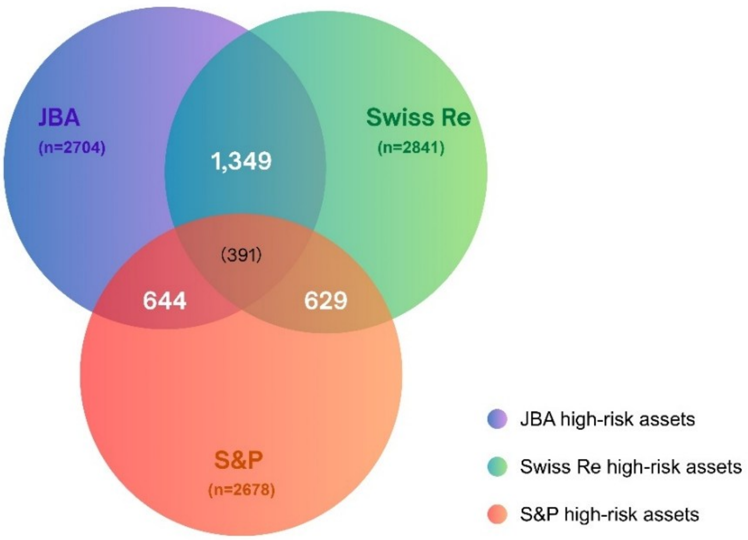

While rank correlation captures agreement across the full distribution of flood scores, a question of more direct relevance to financial risk management is whether the three datasets identify the same assets as high risk. To examine this, this study further assesses cross-dataset agreement by comparing the overlap of assets classified as high-risk across providers.

High-risk assets are defined as those within the top 10 percent of fluvial flood scores in each dataset. In principle, this threshold would correspond to approximately 2,684 assets, based on a total portfolio of 26,841 assets. However, due to tied scores at the percentile threshold, the number of high-risk assets varies across datasets. The high-risk classification agreement between datasets is measured using the *Jaccard Index*, calculated as the ratio of the number of assets classified as high-risk by both datasets (the intersection) to the total number of assets classified as high-risk by either dataset (the union). A value of 1.0 would indicate perfect overlap, meaning both datasets identify exactly the same set of high-risk assets, while a value of 0 would indicate no overlap at all.

$$Jaccard\ Index = |A \cap B| \,/\, |A \cup B|$$

where A and B denote the sets of high-risk assets identified by the two datasets being compared.

Table 5 Pairwise high-risk asset agreement based on the Jaccard Index.

Dataset Pair	Intersection	Union	Jaccard Index
JBA vs Swiss Re	1,349	4,196	0.32
JBA vs S&P	644	4,738	0.14
Swiss Re vs S&P	629	4,890	0.13
JBA vs Swiss Re vs S&P	391	5,992	0.07

As shown in Table 5, JBA and Swiss Re show the highest level of agreement, with 1,349 common high-risk assets and a Jaccard Index of 0.32. This indicates a moderate degree of consistency between the two high-resolution datasets. In contrast, S&P shows much lower overlap with both JBA and Swiss Re, with Jaccard Index values of 0.14 and 0.13, respectively. This low overlap reinforces the conclusion drawn from the correlation analysis: S&P's coarser spatial resolution produces a different picture of which specific assets are most exposed, even where the broad directional trends in flood risk are similar.

Notably, only 391 assets are classified as high-risk by all three datasets simultaneously. Relative to the combined set of 5,992 assets identified as high risk by at least one provider, this corresponds to a three-way Jaccard Index of 0.07. Despite the methodological differences between providers, these assets are consistently identified as the most exposed across all three independent assessments, providing a high-confidence priority set for immediate risk management attention.

These results indicate that the choice of dataset can materially affect which assets are classified as high risk. If a single dataset is used, the selection of high-risk assets may be strongly influenced by provider-specific modelling assumptions, spatial resolution, and score construction. For example, S&P overlaps with only 644 of the 2,704 assets identified as high risk by JBA, meaning that approximately 76 percent of JBA's high-risk assets would not be flagged if an institution relied solely on S&P for screening. This could have material implications for risk prioritisation, stress testing, and regulatory disclosure.

Taken together, the Spearman correlation and high-risk overlap analyses point to the same conclusion: dataset selection is not a neutral technical choice. It materially affects both the relative ranking of assets by flood exposure and the identification of which assets breach a high-risk threshold, with significant downstream consequences for credit risk assessment, stress testing, and regulatory disclosure. These findings underscore the value of adopting a multi-dataset approach in physical climate risk assessment for loan portfolios.

5. Projected Flood Hazard Exposure Across Indonesia

Building on the cross-dataset consistency assessment in the last section, this section presents the projected exposure estimates for the Indonesian asset portfolio under both present-day conditions and future climate scenarios. This section is organised by hazard type: inland flood exposure (fluvial and pluvial), coastal flood exposure, and the geographic distribution of high-risk assets.

5.1. Inland flood exposure (fluvial & pluvial)

Figure 7 presents the number of assets classified as exposed to high inland flood risk under Swiss Re and JBA across all scenarios and time horizons. Two independent assessments reveal consistent qualitative patterns alongside meaningful differences in absolute scale. Fluvial flooding is the dominant source of inland flood exposure in both datasets. Under present-day conditions, Swiss Re identifies 2,841 assets as exposed to high fluvial flood risk against 1,164 for pluvial, while JBA shows a stronger contrast of 4,296 against 395.

This reflects the physical character of flood hazard in Indonesia, where extensive river basins and dense asset concentrations along floodplains make fluvial flooding the principal driver of inland exposure. Both datasets indicate that inland flood exposure increases over time and intensifies under higher emissions. Under Swiss Re, fluvial exposure rises modestly from 2,841 assets at present to 2,990 by 2080 under SSP2-4.5, and to 3,727 under SSP5-8.5. JBA shows a parallel trajectory, with fluvial exposure rising from 4,296 at present to 4,992 under SSP2-4.5 and 5,691 under SSP5-8.5 by 2100. Pluvial exposure follows a similar upward trend, although from a substantially lower base. The directional consistency between two methodologically distinct providers strengthens confidence in the qualitative climate signal: a larger share of the Indonesian asset portfolio will face high inland flood risk under continued warming, and the rate of increase scales with emissions intensity.

Figure 7 The number of assets exposed to high inland flood risk across Indonesia

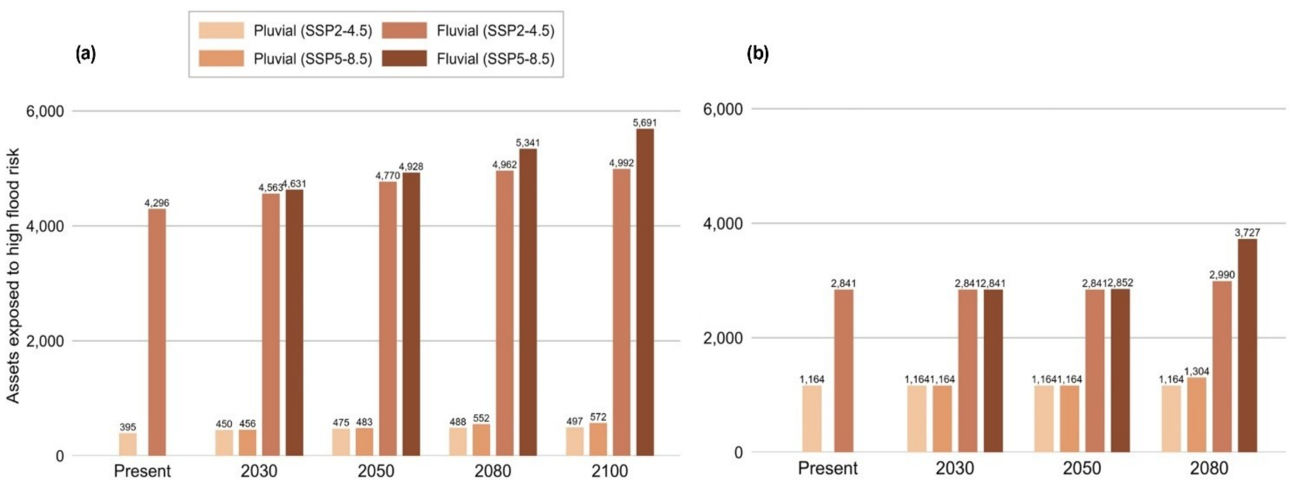

under different climate scenarios, as estimated by JBA (a) and Swiss Re (b)

Despite this directional consistency, the two datasets diverge substantially in the absolute scale of exposure. JBA consistently identifies more fluvial-risk assets than Swiss Re, with the present-day difference of 4,296 against 2,841 amounting to approximately 51 percent, a gap that persists across all scenarios and time horizons. This is consistent with the moderate but imperfect rank correlation reported in Section 4: although both providers operate at 30-metre resolution, differences in hydrodynamic modelling, return period definitions, and high-risk threshold mapping produce systematically different asset counts. The divergence reinforces the case for using multiple high-resolution datasets in parallel rather than relying on a single provider for asset-level risk identification.

The gap between the two climate scenarios also widens over time. At 2030, the scenarios are nearly indistinguishable: Swiss Re reports 2,841 fluvial high-risk assets under both, and JBA reports 4,563 against 4,631, a difference of only 68 assets. By 2080, Swiss Re fluvial exposure reaches 3,727 under SSP5-8.5 against 2,990 under SSP2-4.5, a difference of 737 assets, or 25 percent, with JBA showing a comparable widening through 2100. Scenario choice therefore has limited consequences for portfolio risk over the next decade but becomes a material driver over the multi-decadal horizons relevant to long-term lending and capital provisioning, underscoring the financial value of emissions mitigation.

5.2. Coastal flood exposure

Coastal flood exposure is assessed using Swiss Re storm surge data, which is the only high-resolution coastal flood dataset for the Indonesian portfolio that we can currently access. At the national level, 1,853 assets, approximately 7 percent of the total, are exposed to high coastal flood risk under present-day conditions. This figure remains unchanged under SSP2-4.5 through to 2080 and rises only modestly to 2,093 assets under SSP5-8.5, equivalent to 7.8 percent of the portfolio. Coastal flood exposure is therefore broadly stable over the projection period, with sensitivity to the emissions pathway emerging only under the high-emission scenario.

Figure 8 Regional distribution of high coastal flood risk exposure across Indonesia

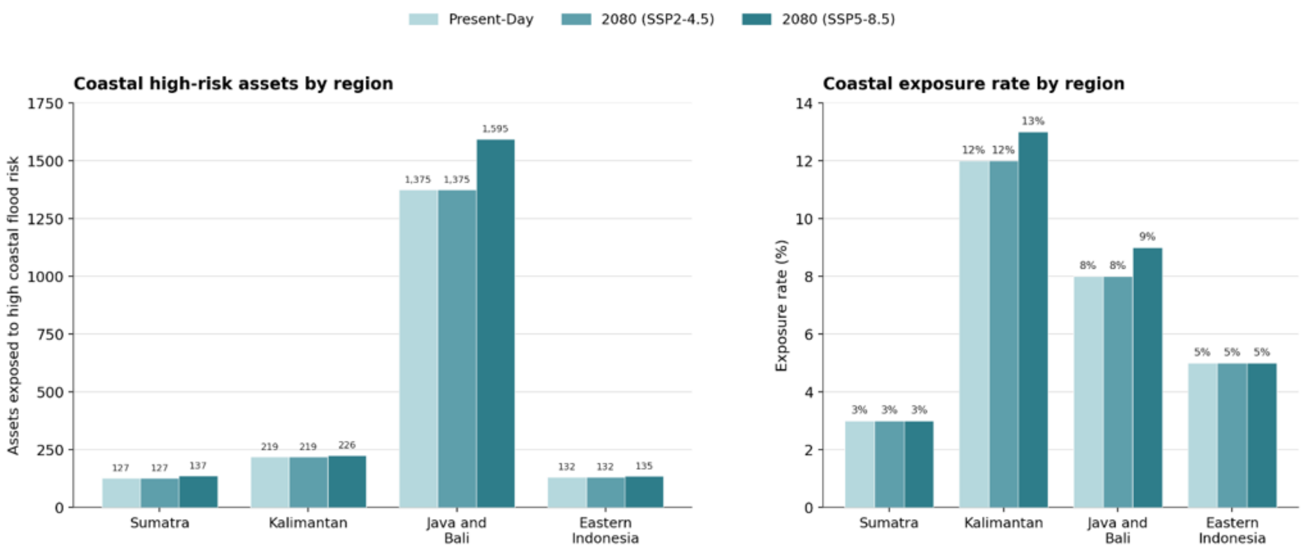

Based on Swiss Re storm surge data. Left: absolute number of high-risk assets. Right: coastal exposure rate by region

Figure 8 shows that coastal flood exposure is highly concentrated geographically. Java and Bali alone account for 1,375 of the 1,853 high-risk assets nationally, or approximately 74 percent, reflecting its extensive and densely developed northern coastline. As with inland flood hazard, however, the absolute distribution masks an important difference in proportional terms. Kalimantan has the highest coastal exposure rate of any region at 12 percent, compared with 8 percent for Java and Bali, 5 percent for Eastern Indonesia, and 3 percent for Sumatra. The two regions that dominate coastal flood risk therefore do so for distinct reasons: Java and Bali through the sheer number of exposed assets, and Kalimantan through the high proportion of its portfolio that is exposed. This mirrors the pattern identified for inland flood hazard in Section 5.1. and reinforces the importance of considering both absolute and proportional measures in coastal risk management.

The temporal behaviour of coastal flood exposure differs markedly from that of inland flood hazard. Whereas fluvial exposure rises steadily over time (Section 5.1), coastal exposure is essentially constant from the present day through 2080 under SSP2-4.5, and increases only marginally under SSP5-8.5, with the national total rising by 240 assets and the largest single increase occurring in Java and Bali, from 1,375 to 1,595. This relative stability should be interpreted with caution. It may reflect the structure of the storm surge modelling approach and the limitation of the projection horizon to 2080, rather than an absence of longer-term coastal risk dynamics such as progressive sea-level rise. The reliance on a single dataset further limits the confidence that can be placed in these projections.

5.3. Regional distribution of flood exposure

To examine the geographic distribution of flood risk across Indonesia, assets were grouped into four regions: Sumatra, Kalimantan, Java and Bali, and Eastern Indonesia. This analysis uses JBA defended flood scores, which combine fluvial and pluvial hazard and account for existing flood defence infrastructure, providing the most representative estimate of residual asset-level risk. Table 6 summarises the regional breakdown under present-day conditions and under both climate scenarios at 2100 (see Appendix for region-level trajectories across all time horizons).

Table 6 Regional distribution of high-risk assets across Indonesia.

Region	Total Assets	High-Risk Assets (Present-Day)	Exposure Rate (Present-Day)	High-Risk Assets (2100, SSP2-4.5)	Exposure Rate (2100, SSP2-4.5)	High-Risk Assets (2100, SSP5-8.5)	Exposure Rate (2100, SSP5-8.5)
Sumatra	4,410	481	10.9%	579	13.1%	630	14.3%
Kalimantan	1,803	583	32.3%	629	34.9%	655	36.3%
Java & Bali	17,784	1,287	7.2%	2,335	13.1%	3,153	17.7%
Eastern Indonesia	2,844	173	6.1%	211	7.4%	252	8.9%
Total	26,841	2,524	9.4%	3,754	14.0%	4,690	17.5%

Flood risk exposure is projected to increase across all four regions under both scenarios. The total number of high-risk assets rises from 2,524 at present to 3,754 by 2100 under SSP2-4.5, an increase of 49 percent, and to 4,690 under SSP5-8.5, an increase of 86 percent that approximately doubles present-day exposure. The portfolio-wide exposure rate correspondingly rises from 9.4 to 17.5 percent under the high-emission pathway.

The distribution of this exposure is highly uneven, and the picture differs depending on whether risk is measured in absolute or proportional terms. Java and Bali are by far the most exposed region in absolute terms, accounting for 1,287 of the 2,524 high-risk assets at present (51%) and rising to 3,153 of 4,690 by 2100 under SSP5-8.5 (67%). This concentration reflects both the region's dominant share of the asset portfolio (17,784 of 26,841 assets) and its dense, low-lying coastal geography. Notably, Java and Bali also show the steepest growth in proportional terms, with its exposure rate rising from 7.2 percent at present to 17.7 percent by 2100 under SSP5-8.5. The combination of the largest absolute concentration of risk and the fastest-deteriorating exposure rate makes Java and Bali the highest-priority region for climate risk management.

Kalimantan presents a contrasting risk profile. It has the highest present-day exposure rate of any region, with 583 of its 1,803 assets (32.3%) already classified as high-risk before accounting for future climate change, reflecting the region's extensive low-lying floodplains and peatland hydrology. However, its exposure rate increases only modestly under climate change, reaching 36.3 percent by 2100 under SSP5-8.5. This relative stability suggests that flood risk in Kalimantan is already largely realised under current conditions, in contrast to Java and Bali, where a substantial share of the risk is yet to materialise. Although Kalimantan's absolute exposure is considerably smaller, its consistently high exposure rate indicates that a large proportion of its assets face material flood risk today.

Sumatra and Eastern Indonesia show the lowest exposure on both measures, with present-day exposure rates of 10.9 percent and 6.1 percent respectively, but neither region is risk-free. Both follow the same upward trajectory under climate change, with exposure rates rising to 14.3 percent and 8.9 percent by 2100 under SSP5-8.5.

These patterns carry direct implications for portfolio risk management. Java and Bali represent the largest aggregate concentration of flood risk, the greatest source of potential portfolio loss, and the most rapidly worsening exposure, and therefore warrants priority attention on all three counts. Kalimantan's elevated but stable exposure rate signals high per-asset risk that is already present rather than emerging, which may call for region-specific underwriting or pricing adjustments distinct from the forward-looking provisioning appropriate to Java and Bali. A risk framework informed only by absolute counts would understate the intensity of present risk in Kalimantan, while one based solely on exposure rates would understate both the aggregate materiality and the future trajectory of Java and Bali.

Given their distinct risk profiles, Java and Bali and Kalimantan warrant closer examination. The two regions are the most consequential in the portfolio, for different reasons: Java and Bali combines the largest concentration of assets with the greatest absolute number of high-risk assets, while Kalimantan has the highest exposure rate of any region and is also the location of Indonesia’s new capital city, Nusantara, currently under construction in East Kalimantan and intended to assume national capital functions from Jakarta over the coming years. As Kalimantan's current exposure rate reflects an asset base concentrated in flood-prone riverine and coastal zones, the build-out of Nusantara may gradually shift the regional asset mix towards less exposed locations, though how the region's risk profile evolves remains to be observed. Examining the composition of flood hazard in these two regions, rather than aggregate exposure alone, reveals two dimensions that the combined high-risk classification in Table 6 does not capture: the contribution of coastal flooding, and the contrast between absolute and proportional exposure.

Figure 9 Composition of high flood risk exposure by hazard type in Java, Bali, and Kalimantan.

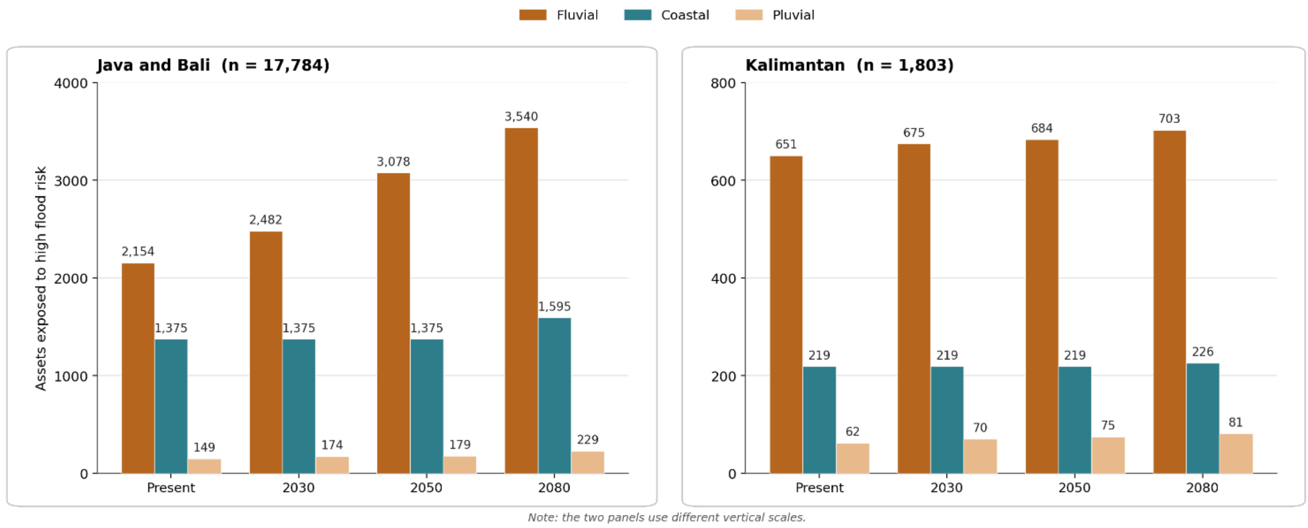

Under SSP5-8.5 (Present to 2080). Inland flood (fluvial and pluvial) is based on JBA defended scores; coastal flood is based on Swiss Re storm surge data, which is available only through 2080

Figure 9 shows high flood risk exposure in the two regions by hazard type. To incorporate coastal flooding, which is absent from the JBA defended scores used elsewhere in this section, coastal exposure is drawn from Swiss Re storm surge data, while inland fluvial and pluvial exposure remains on the JBA basis.

The decomposition shows that coastal flooding is the second-largest source of flood exposure in both regions, exceeding pluvial flood risk by a wide margin. In Java and Bali, 1,375 assets are exposed to high coastal flood risk under present-day conditions against 149 for pluvial, a difference of approximately ninefold; in Kalimantan, the corresponding figures are 219 against 62. This is a material finding, because the combined fluvial and pluvial classification used in Table 6 omits coastal hazard entirely and therefore understates total flood exposure in both regions. The prominence of coastal risk reflects their geography, with Java and Bali's densely developed northern coastline and Kalimantan's low-lying coastal and deltaic zones being particularly exposed to storm surge. For Kalimantan, this is especially relevant given that the capital relocation is concentrating major new strategic assets in a region that already exhibits the highest flood exposure rate in the portfolio.

The three hazard types also differ in their response to climate change. Fluvial exposure rises steadily over time, increasing in Java and Bali from 2,154 assets at present to 3,540 by 2080 under SSP5-8.5, whereas coastal exposure remains nearly constant, holding at 1,375 assets until a modest rise to 1,595 by 2080; a similar contrast holds in Kalimantan. This relative stability in coastal exposure should be interpreted with caution, as it may reflect the structure of the storm surge modelling approach and the limited projection horizon rather than an absence of future coastal risk dynamics, particularly since coastal estimates currently rely on a single provider.

Figure 10 presents the same exposure expressed as a rate, defined as the share of each region's total assets classified as high-risk, on a common vertical scale. Viewed this way, the relationship between the two regions reverses. Although Java and Bali contain far more high-risk assets in absolute terms, Kalimantan faces a higher exposure rate for every hazard type at every time horizon. Its fluvial exposure rate rises from 36.1 percent at present to 39.0 percent by 2080 under SSP5-8.5, against 12.1 percent to 19.9 percent for Java and Bali; its coastal rate holds near 12 percent against 8 to 9 percent for Java and Bali; and even its pluvial rate, the smallest hazard in both regions, is roughly three to four times that of Java and Bali. This reversal sharpens the distinction drawn earlier in this section: Java and Bali are the dominant drivers of aggregate portfolio exposure and potential loss, whereas Kalimantan carries the most intense per-asset risk, with a substantially larger share of its portfolio exposed across all flood types. The two regions therefore call for different risk management responses, with Java and Bali warranting attention proportionate to the scale of assets at stake, and Kalimantan warranting region-specific underwriting or pricing that reflects its elevated per-asset exposure.

Figure 10 Share of high-risk exposure in Java, Bali and Kalimantan

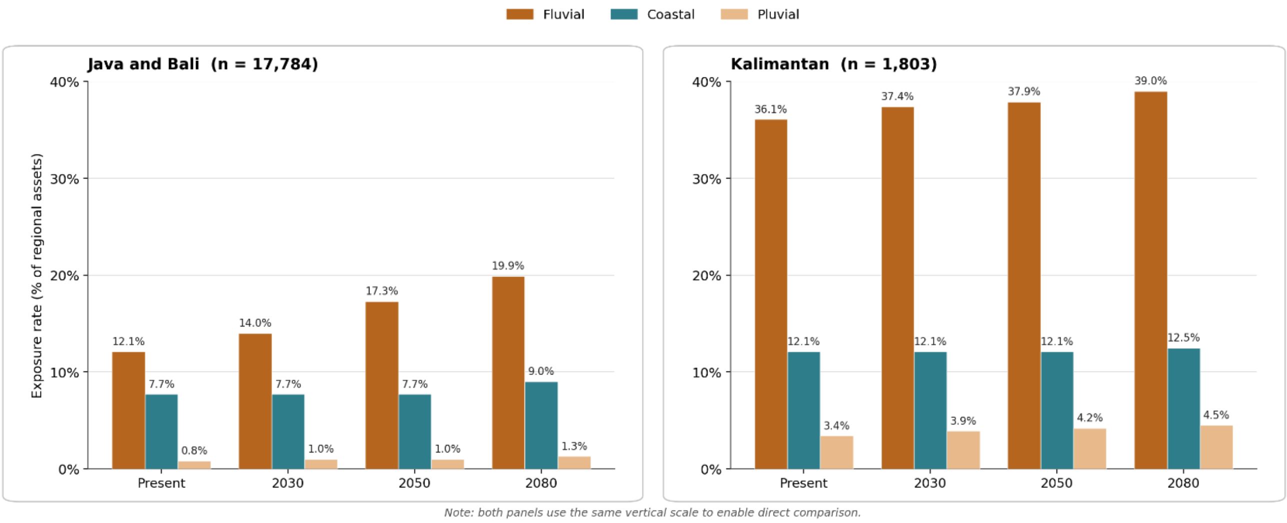

Defined as the share of each region's total assets classified as high-risk, by hazard type, under SSP5-8.5 (Present to 2080)

NUS National University of Singapore | Sustainable and Green Finance Institute

Translating Flood Exposure to Potential Bank Loss

V. Translating Flood Exposure to Potential Bank Loss

Having established the asset inventory and its coverage characteristics, the analytical challenge shifts from identifying where assets are located to assessing how flood exposure may translate into financial loss. The S&P Sustainable1 financial impact dataset provides a useful benchmark reference because it offers cross-company comparable physical risk impact estimates grounded in asset-level hazards and asset-type vulnerability.

1. Financial Impact Data from S&P Sustainable1 as a Benchmark Reference

1.1. Asset-level financial impact

Public S&P Global Sustainable1 materials describe asset-level financial impact as combining a physical hazard score, as mentioned in the Section 3, with an asset type-specific vulnerability function. The vulnerability profile considers the damage estimates in terms of increased operating expenses, lost revenues due to business interruption, physical damage, and repair costs. These costs are aggregated and expressed as a percentage of the value of each asset type, providing a normalised indicator of financial impact at the asset level. [28]

For this study, the relevant reference impacts are those associated with inland flood hazards, namely fluvial and pluvial flooding. The data are reported across climate scenarios and forecast years, allowing flood-related impact estimates to be compared over time and across assets. Coastal flood impacts are not used in the empirical application of this study, which is why the resulting estimates should not be interpreted as comprehensive all-flood loss estimates.

1.2. Company-level financial impact

At the company level, S&P's public materials describe financial impact as reflecting the weighted average financial impact of all assets linked to a company, weighted by assumed asset value assigned to an asset type globally. This aggregation structure is useful for this study because it links asset-level hazard exposure and asset-type vulnerability function to a company-level financial impact metric. It is worth noting that where asset coverage is limited, S&P relies on broader assumptions or fallback treatments; this reinforces the need to interpret company-level impact estimates in light of asset coverage and data-quality considerations. [29]

1.3. Role as benchmark reference

This study uses S&P financial impact data at the asset and company levels as a benchmark reference for constructing study-specific vulnerability function proxies and recovering the relative asset values used to aggregate them. The core analytical challenge is that firm-level financial impact metrics combine hazard scores, asset-type vulnerability, and estimated asset value into a single integrated output.

The analysis does not attempt to recover or reproduce any provider's proprietary vulnerability function, asset-value estimates, or model parameters. Instead, it uses observable patterns in the licensed outputs — together with the value-weighted aggregation structure that S&P describes in its public materials — to construct transparent, study-specific vulnerability function proxies and to estimate the relative asset-value weights used to aggregate asset-level estimates to the company level. These estimated weights are study-specific aggregation weights, not the provider's asset values. The proxies and weights allow flood-related loss estimates to be applied consistently across alternative hazard datasets and, where available, to bank-specific exposure values, so the resulting estimates should be interpreted as indicative flood-induced loss estimates for credit risk analysis rather than as replications of S&P financial impact outputs.

2. The Vulnerability Puzzle

Flood loss estimation requires more than knowing where an asset is located and how severely it is exposed to flooding. Two assets with identical financial values and identical flood hazard can sustain materially different losses depending on what they are: a warehouse, light manufacturing facility, and a retail outlet respond differently to the same depth of inundation, owing to differences in construction standards, operational dependencies, and recovery capacity. This asset-type-specific loss response is what we term ***vulnerability***: the generally non-linear mapping from a hazard score to a fractional loss of asset value. That fractional loss, the output of the mapping, is the Projected Asset Loss (%), which S&P Sustainable1 reports as the financial impact (%).

Conceptually, the indicative monetary loss on an asset combines a loss rate with the value at stake:

Dollar Loss = Projected Asset Loss × Asset Value.

The projected asset loss rate is itself a function of hazard and vulnerability, so the full decomposition is $L = f(H, V, E)$, where H is the flood hazard score, V a study-specific vulnerability function proxying asset-type loss response, and E the exposure value.

Each component is required to translate physical hazard exposure into economically meaningful loss estimates. The hazard score locates the asset and characterises how severe flooding is under a given climate scenario; the vulnerability function approximates how sharply projected loss responds to that hazard for the asset type; and the exposure value represents the financial value at stake, which may be represented by asset value, replacement value, collateral value, outstanding loan amount, or another bank-reported credit exposure measure.

Commercial physical risk data providers offer detailed estimates of hazard and projected asset losses across asset locations and asset types, but these outputs integrate hazard, asset-type vulnerability, and an estimated asset value within proprietary modelling frameworks. For a bank that adopts a single provider's hazard scores and projected asset losses directly, this integration poses no difficulty: the output is used as supplied. The difficulty arises because the projected asset loss depends on which provider produced it, and as shown later in Chapter VI, the choice of provider carries material consequences for the resulting loss estimates. The integrated figure does not expose the underlying asset-type vulnerability as a separate parameter, so it cannot be inspected, compared across providers, or carried over to a different hazard dataset.

This creates what we refer to as the “vulnerability puzzle”: The asset-type loss response linking hazard intensity to proportional asset impairment — the vulnerability function — is bundled within the provider's integrated output and is not directly observable as a standalone input. To interpret why losses estimate diverge across providers, and to deepen understanding of how those estimates are formed, this relationship must be approximated through transparent, study-specific assumptions.

This study addresses the vulnerability puzzle by constructing study-specific vulnerability proxies from observable relationships in the licensed S&P Global Sustainable1 outputs. These proxies are not intended to represent, infer, or reproduce any provider's proprietary sensitivity profile, vulnerability function, or model parameter; they are analytical approximations of asset-type vulnerability for use in comparison, loss analysis, and stress testing.

It should be noted that the asset value assumptions embedded in commercial physical risk datasets are not equivalent to bank-specific credit exposure so the resulting proxies should not be interpreted as pure engineering vulnerability parameters. They are better understood as study-specific translation weights that reflect the structure of the source data and are designed to support indicative loss estimation. When combined with bank-specific exposure values, they provide a practical basis for converting externally sourced normalised hazard scores into indicative asset-level losses.

Once constructed, the vulnerability proxies can be applied to normalised hazard scores from alternative providers — JBA Risk Management or Swiss Re — and to bank-reported asset values, allowing the same asset inventory to be assessed under different hazard datasets while holding the loss assumption fixed. With the proxy held constant, differences in estimated loss rates can be interpreted primarily as divergences in hazard inputs, subject to residual uncertainty from data coverage, spatial matching, and model specification. This supports a transparent comparison of provider divergence for portfolio stress testing.

3. Methodology for Vulnerability Function Proxies and Relative Asset Values Construction

This section sets out the technical procedure used to construct the study-specific vulnerability function proxies and relative asset-value weights introduced above. Readers concerned only with the portfolio results may proceed directly to Section VI.

The procedure has two stages. The first estimates, for each asset type, how a flood hazard score maps into an asset-level projected asset loss — the vulnerability function. The second estimates the relative asset-value weights that aggregate asset-level losses into firm-level estimates, which are then carried to the portfolio through credit exposures.

The asset-level loss is represented as a hazard–impact mapping,

$$I^f_{i,t,s} = \phi^f_k(H^f_{i,t,s}),$$

where i indexes the asset, t the forecast year, s the scenario, and f the flood type; $H^f_{i,t,s}$ is the hazard score; and $\phi^f_k(\cdot)$ is the vulnerability function for asset type k (modelled on S&P asset type – subtype level) and flood type f. Firm-level loss is the asset-value-weighted average of the losses on a firm's constituent assets,

$$I^f_{j,t,s} = \frac{\sum_{i\in A_j} w_{k_i}\, I^f_{i,t,s}}{\sum_{i\in A_j} w_{k_i}},$$

where A_j is the set of assets belonging to firm j, k_i the asset type of asset i, and w_{k_i} the relative asset-value weight assigned to that asset type. This mirrors the provider's stated company-level construction, in which firm impact is the value-weighted average of the impacts of its assets. The same value weights aggregate hazard, so firm-level hazard takes the same form,

$$H^f_{j,t,s} = \frac{\sum_{i\in A_j} w_{k_i}\, H^f_{i,t,s}}{\sum_{i\in A_j} w_{k_i}},$$

with the weights w_{k_i} common to both aggregations and estimated in Stage 2.

These relationships define the quantities to be estimated, not the estimators themselves: the vulnerability functions ϕ^f_k are fitted in Stage 1, and the weights w_k are estimated in Stage 2. The vulnerability functions are fitted separately for each flood type. The portfolio results reported in this study cover inland flooding — fluvial and pluvial; coastal flooding is fitted but is not carried into the reported results. The weight recovery in

Stage 2, by contrast, pools observations across all three flood types, including coastal, to identify a single set of asset-value weights.

The reported impact values used to fit these relationships have several characteristics that bear on estimation: missing impacts are recorded as zeros, the values have inconsistent, limited numerical resolution, and the hazard scale is normalised at 0-100, rendering impacts at hazard values of 100 uninformative about the underlying curves. These are handled directly when the vulnerability functions are fitted in Stage 1. They are also why the weights in Stage 2 are estimated from the asset-vs-firm hazard relation rather than the asset-vs-firm impact relation: hazard provides the cleaner signal for identifying the weights, which — common to both relations — are then applied to asset-level impacts to produce the firm-level estimates.

Two consistency rules are applied to the source data before either stage. Missing hazard observations are interpreted as the absence of modelled exposure and set to zero, consistent with existing observations on the dataset that $H^f_{i,t,s} = 0 \Rightarrow I^f_{i,t,s} = 0$ is always true. Observations reporting zero hazard with strictly positive impact violate this relationship and are removed. All remaining sample construction is stage-specific and is described under the relevant stage.

3.1. Stage 1: Vulnerability function estimation

1) Data treatment

Several characteristics of the hazard–impact pairs $\left(H^f_{i,t,s}, I^f_{i,t,s}\right)$ are addressed before ϕ^f_k is fitted.

For each asset subtype, the onset threshold

$$H^*_k = \min\left\{ H^f_{i,t,s} \mid I^f_{i,t,s} > 0 \right\}$$

is the lowest hazard score at which a positive impact is recorded. Observations with $H^f_{i,t,s} < H^*_k$ and $I^f_{i,t,s} = 0$ are excluded from the curve fitting but used to determine the threshold behaviours; subtypes for which impact is zero throughout provide no variation to identify ϕ^f_k and are dropped.

Each non-zero observation is classified as continuous or discrete according to how often the same impact value recurs over a range of hazard scores: an impact value that recurs is treated as a discrete level, and one that appears only once as a higher-resolution observation. Where a subtype has a sufficient number of continuous observations, the curve is fitted on those alone, since they carry the clearest slope signal. Otherwise, the curve is fitted on all observations, with a midpoint correction

$$\tilde{y}_i = y_i + \delta_k/2$$

applied to the non-zero discrete observations, where δ_k is the smallest interval between distinct impact values for the subtype, so that each discrete level is represented by the midpoint of its interval rather than its lower edge.

To equalise influence across the hazard range, curves are fitted on per-score means — the mean impact across all observations sharing a given hazard score — rather than on the raw observations, so that frequently occurring scores do not dominate the fit. Goodness of fit is always evaluated on the underlying observations, not on the per-score means.

The hazard score is expressed on a normalised scale from 0 to 100, where 100 represents the maximum modelled hazard. Observations at $H^f_{i,t,s} = 100$ therefore sit at the top of this scale and reflect their upper bound rather than the marginal response to increasing hazard; they are held out of curve fitting, and their treatment in prediction is set out below.

2) Functional forms and identification

Each vulnerability function is selected from a small set of candidate forms, defined on the hazard variable x — equal to H when the onset threshold is zero, and to $H - H_k^*$ otherwise:

$$\text{linear: } \phi = a\,x;$$
$$\text{power: } \phi = a\,x^b,\ 1 \le b \le 5;$$

$$\text{exponential: } \phi = a\big(e^{bx} - 1\big), 0.005\ \le b \le 0.06;$$
$$\text{quadratic: } \phi = c_2\,x^2 + c_1\,x,\ c_2 \ge 0.$$

Two structural constraints keep the fitted mappings consistent with the expected response to flooding. The origin constraint

$$\phi_k^f(0) = 0$$

requires zero impact at zero hazard, and each mapping is required to be increasing in hazard, since a higher hazard score should not correspond to a lower projected impact. The candidate forms are further restricted to be convex — through $1 \le b \le 5$ for the power form, $0.005\ \le b \le 0.06$ for the exponential form and $c_2 \ge 0$ for the quadratic — which rules out responses that would flatten at high hazard, contrary to the rising tail observed in the data; concave power and concave quadratic fits are therefore not admitted. The upper limit on the power exponent is set to the 99th percentile of the observed exponents fitted across the data-rich subtypes, which prevents isolated fits from extrapolating to implausibly large impacts at high hazard.

The functional form for each subtype is selected by adjusted R^2 evaluated on the underlying observations, which penalises the additional parameters of the more flexible forms and so guards against over-fitting.

3) Threshold regime

For subtypes whose impacts remain zero up to a positive onset threshold H_k^*, the fitted mapping is piecewise, and its form above the threshold depends on the empirical pattern of impacts. When those impacts are near-constant — identified by a low coefficient of variation among the non-zero values — the response is represented as a step:

$$\phi_k^f(H) = \begin{cases} 0, & H < H_k^* \\ c, & H \ge H_k^* \end{cases}$$

where c is the constant impact level above the threshold. Otherwise, a smooth form is fitted on the shifted hazard variable $H - H_k^*$, with an intercept at the threshold:

$$\phi_k^f(H) = \begin{cases} 0, & H < H_k^* \\ g(H - H_k^*) + c, & H \ge H_k^* \end{cases}$$

where $g(\cdot)$ is the selected functional form and c the minimum non-zero impact observed for the subtype. This captures the pattern, common across subtypes, in which modelled impact appears only once a hazard threshold is reached.

4) Ceiling treatment

Because the $H = 100$ observations are held out of fitting, the predicted impact at the top of the scale is set to the observed median at that level for each subtype,

$$\phi_k^f(100) = \mathrm{median}\{\, I_{i,t,s}^f \mid H_{i,t,s}^f = 100 \,\}.$$

A ceiling-gap diagnostic — the difference between this median and the value implied by extrapolating the fitted curve to $H = 100$ — is reported to gauge the sensitivity of each mapping to the upper bound.

5) Hierarchical fallback

Prediction uses a fallback hierarchy. For a given flood type and asset subtype, the subtype-level fitted curve is used where one exists; because a subtype is fitted only when it has enough usable observations, a subtype that was not fitted falls back to the curve fitted at its parent asset-type level, and, failing that, to a global curve fitted across all assets. At $H = 100$, the ceiling median at the matched level is applied. Each asset therefore receives a prediction from the most granular vulnerability function the data support, with the broader levels ensuring full coverage.

3.2. Stage 2: Relative asset-value weight recovery

The second stage recovers the asset-value weights w_k that aggregate asset-level impacts into firm-level outcomes through the weighted average defined above. The provider reports both asset-level and firm-level impacts but not the assumed asset values used in aggregation, so these weights must be estimated to obtain a consistent asset-to-firm rule that can subsequently be applied to bank exposures.

6) Estimation sample

Estimating the weights requires a clean mapping from asset-level observations to firm-level aggregates, which governs the sample. Only firms whose impacts derive from asset-based methodologies — the higher data-quality classes — are retained, so that firm-level outputs are genuinely built up from their assets rather than from geographic, revenue, or country proxies. Within this set, only firms with at least two distinct asset (type, subtype) combinations are kept, since a firm of uniform composition carries no information on relative weights. Missing hazards are set to zero as above. The remaining two filters apply at the level of each firm and flood type: an observation is dropped if the firm's reported firm-level hazard for that flood type lies outside the range spanned by its constituent asset hazards, or if that firm-level hazard value is zero for the flood type. These restrictions reduce the sample but improve the internal consistency of the estimation.

7) Estimation objective

The relative weights are recovered by matching the value-weighted asset hazard to the reported firm hazard — the hazard aggregation given at the outset — choosing a single set of weights to minimise the discrepancy, pooled across all three flood types:

$$\min_{w>0} \sum_{f} \frac{1}{\mathrm{Var}_f} \sum_{j,t,s} \left(\widehat{H}^{f}_{j,t,s}(w) - H^{f}_{j,t,s} \right)^2,$$

where $\widehat{H}^{f}_{j,t,s}(w)$ is the value-weighted asset hazard given at the outset, evaluated at the weights w; the sum over f runs over fluvial, pluvial, and coastal flooding; and Var_f is the variance of reported firm hazards for flood type f, included so that impacts for three different flood types, each of different magnitudes, do not dominate one another. Because the aggregation is a ratio, the weights are identified only up to a common positive scale; the scale is fixed by normalising the reference asset type to unity:

$$w_{\mathrm{ref}} = 1,$$
$$\mathrm{ref} = \text{Office} \;\|\; \text{Company Headquarters}.$$

The estimated weights are study-specific aggregation weights, not observed asset values, replacement costs, or provider model parameters. Because the impact and hazard aggregations share these weights, the recovered weights are applied to asset-level impacts to produce firm-level estimates; this assumption — that the value weights are common to both aggregations — is reasonable given the shared asset inventory and firm structure, though it is not separately identified within this exercise.

8) Estimation and stability

The problem is solved without regularization. An earlier specification regularised the weights towards external priors, but the solution proved sensitive to the strength of the penalty so the prior-free specification is preferred. Because the sample is limited and some asset types are sparsely represented, the unregularised problem can admit multiple near-equivalent solutions. The optimization is repeated across five random-seed sets with twenty restarts each, and the reported weights are the average across runs; the dispersion across runs is examined to confirm that the solution is stable.

Weights for sparsely represented asset types may be weakly identified and, in some cases, close to zero. These should be read as limited information in the estimation sample for separating an asset type's contribution to firm-level aggregation, not as evidence that the asset type lacks economic value or flood relevance. A larger global asset sample would sharpen these estimates in future work.

9) Applying the relative asset value weights

When the weights are carried beyond the estimation sample, assets whose (type, subtype) combination was not estimated directly are assigned a weight by the same hierarchical logic used for the vulnerability-function fallbacks. This ensures every asset receives a weight while preserving the most granular estimate the data support.

4. Findings

4.1. Vulnerability function results

The fitted study-specific vulnerability mappings show good empirical fit for asset subtypes with sufficient data coverage. The mappings are fitted at the flood-type × subtype level: of the 85 combinations with enough data to support a fit, 83 yield estimable mappings, and adjusted R^2 exceeds 0.9 in roughly four-fifths of these.

The fitted mappings are predominantly exponential or power in form, indicating a nonlinear relationship between flood hazard intensity and modelled financial impact. In practical terms, higher hazard scores are associated with disproportionately higher indicative losses, particularly in the upper range of hazard exposure. This accelerating profile matters for credit risk analysis because it implies that tail exposures may contribute disproportionately to portfolio loss, even when average hazard exposure appears moderate. Figure 11 illustrates the pattern, showing the estimated marginal response of modelled asset-level loss to changes in hazard intensity across selected asset subtypes. These estimates should be interpreted as study-specific indicators of loss response, not as engineering damage functions or provider model parameters.

4.2. Asset weight estimation results

Estimated over 52 asset (type, subtype) weights from 158 *firm x flood-type combinations* with sufficient asset diversity, the unregularised optimisation produces relative asset-value weights that exhibit an economically intuitive hierarchy across asset categories. Large, specialised, and capital-intensive facilities — including LNG facilities, steel plants, mining facilities, and other industrial assets — tend to occupy the upper tiers. Energy and extractive infrastructure, including power generation facilities, ports, and commodity storage assets, form a second tier. Commercial, institutional, office, and service facilities occupy lower tiers, reflecting their lower implied relative asset-value weights within the estimation sample (Table 7).

This ordering is broadly consistent with engineering cost intuition: large industrial and infrastructure assets typically require higher capital expenditure than office or branch facilities. The weights should not, however, be read as point estimates of replacement cost or market value. They are study-specific aggregation weights estimated from the available asset-to-firm structure and are best interpreted as ordinal indicators of relative asset-value contribution.

The recovered weights are stable: the hazard-matching fit reaches an R^2 of about 0.95, identical across initializations, and weight estimates vary by only a few percent across seeds at the median. Weights for sparsely represented asset types may be weakly identified and, in some cases, close to zero. Such values should be read as limited identification within the estimation sample rather than as evidence that the corresponding asset types have no economic value or flood relevance.

Combined with the study-specific vulnerability function proxies, the relative asset-value weights provide a consistent framework for translating asset-level flood exposure into borrower-level indicative losses. The framework can be applied to borrowers with known asset locations, asset types, and exposure values, and can be used across alternative hazard data providers subject to spatial matching and data-coverage constraints.

For portfolio analysis, the resulting borrower-level indicative losses are aggregated using each borrower's share of total loan exposure. This represents a conservative stress assumption in which flood-related asset impairment is assumed to translate into credit-relevant loss without offsetting effects from insurance, equity buffers, operational adaptation, or government support. The resulting estimates should therefore be read as unmitigated physical-loss stress metrics rather than as realised credit losses.

Figure 11 Rate of change of flood-induced impact (%) per unit increase in hazard score

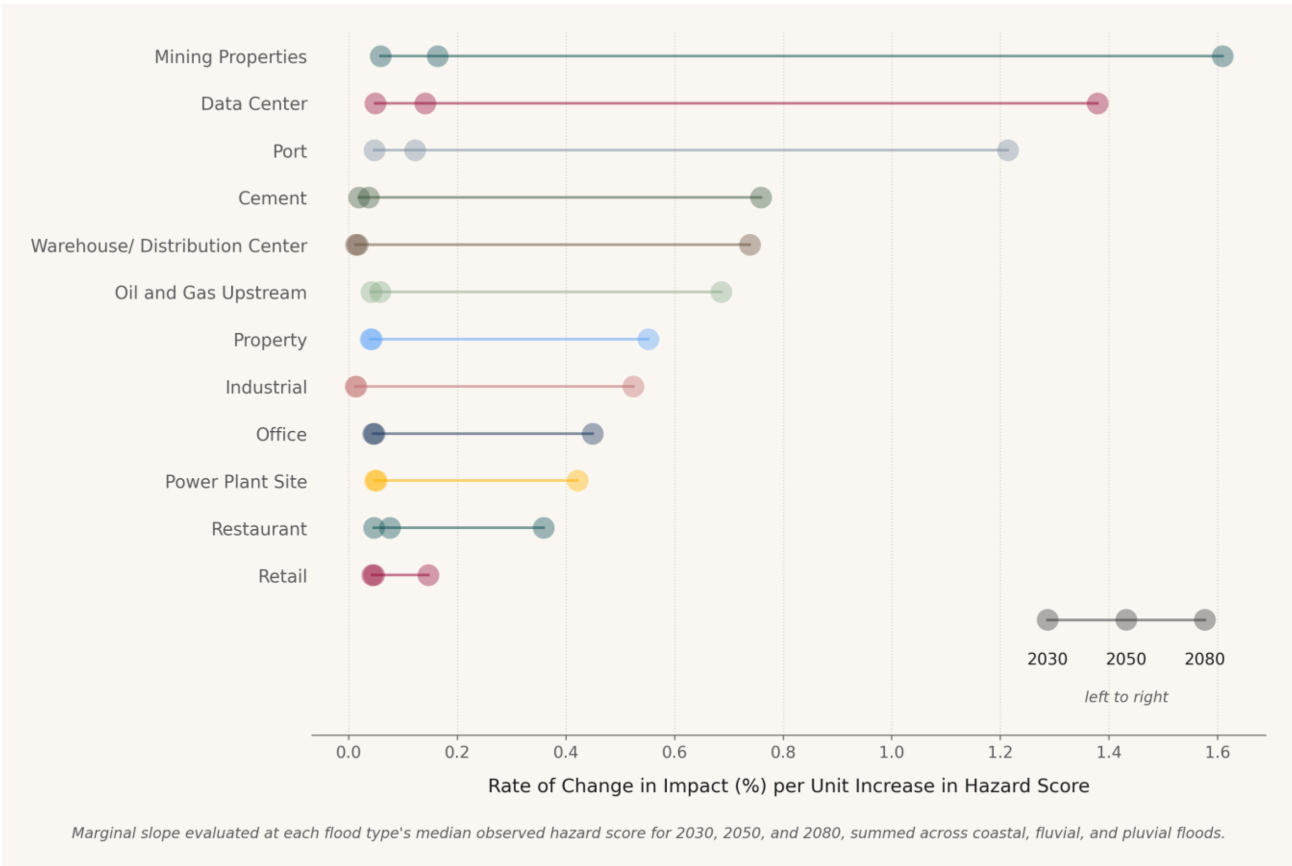

Table 7 Indicative relative asset value tiers (sample of assets)

Tier	***Asset Category***	***Examples***	***Relative Asset Value***
1	Large specialised industrial facilities	LNG facilities, steel plants, mining properties, cement plants	Very high
2	Energy and extractive infrastructure	Power generation facilities, ports, industrial facilities, commodity storage	High
3	Commercial and institutional facilities	Data centres, health care, education institutions, production facilities	Moderate
4	Office and service facilities	Offices, company headquarters, retail, residential, branches	Low

Note: Selected asset categories are shown for illustrative purposes. The full estimation covers 34 asset types across 59 subtypes. Relative weights are presented as ordinal tiers rather than point estimates and should be interpreted as study-specific aggregation weights inferred from the analytical sample, not as observed replacement costs, market values, or provider model parameters.

NUS National University of Singapore | Sustainable and Green Finance Institute

Portfolio Inland Flood Risk Analysis

VI. Portfolio Inland Flood Risk Analysis

Our previous whitepaper provided initial evidence that flood-induced asset losses are material across Indonesian bank loan portfolios, using company-level physical risk impact estimates from S&P Global Sustainable1. In that study, the projected asset loss distribution accounts for most cross-bank differences in physical-risk exposure.

This paper takes the same reconstructed loan portfolio and re-examines flood exposure at the asset level. The objective is not to reproduce the previous estimates directly, but to assess how the distribution of projected losses changes when borrower exposure is analysed through georeferenced asset locations and multiple flood hazard datasets.

The asset-level approach captures heterogeneity in flood risks within individual borrowers that is averaged away in the company-level view. Using S&P Sustainable1, JBA, and Swiss Re data in parallel allows the analysis to assess provider divergence and identify tail-risk assets that may be underweighted or missed under a single hazard source.

The comparative provider analysis is restricted to the 207 borrowers covered by both S&P asset inventory (as the input for georeferenced assets for all asset-level analysis regardless of the hazard vendors) and S&P company-level impact dataset at the 2050 horizon under SSP2-4.5. This common-sample restriction ensures that differences across providers are not driven by differences in borrower coverage.

Four series are used to distinguish aggregation level and hazard-source effects. The S&P Company series refers to company-level impact estimate outputs provided by the provider directly, used as the benchmark reference for the study. The S&P Asset series comprises the asset-level impacts reported on S&P's asset coordinates, with data integrity treatment by the SGFIN team. JBA and Swiss Re provide independent asset-level hazard scores, which we extracted and mapped to the same set of S&P assets. This design holds the borrower sample and asset inventory constant while varying the hazard source.

Projected asset losses in this section are defined as the sum of pluvial and fluvial flood impacts. Coastal flood is excluded from the comparative analysis and will be included in our future work. The estimates should therefore be interpreted as **inland flood loss estimates**, not comprehensive all-flood loss estimates.

1. Loan Book Coverage

Our analysis draws on a reconstructed loan portfolio as of January 2025, comprising 371 corporate borrowers with total loan exposure of approximately USD 185 billion. The portfolio is assembled from Bloomberg and Refinitiv syndicated loan records, harmonised at the tranche level and consolidated to the borrower-lender level, and is restricted to loans with contractual maturity on or after 1 January 2025. The construction methodology, including dataset harmonisation, exposure estimation, and lender consolidation follows our previous whitepaper Zhang et al. (2025), which provides full procedural detail. Because both source databases predominantly capture syndicated transactions, bilateral lending and lending to small and medium-sized enterprises are under-represented. The reconstructed portfolio should therefore be interpreted as a representative view of large-corporate flood exposure within the Indonesian banking system, rather than as a complete portfolio model.

Of the 371 borrowers, 207 borrowers are successfully matched to both the S&P Sustainable1 asset database and the company-level impact dataset, corresponding to 16,652 assets in our database. These matched borrowers represent 55.8 percent of companies and 63.8 percent of total loan value, equivalent to covered exposure of USD 118.1 billion. Matching is achieved primarily through name-based reconciliation and company identifier linkage.

Figure 12 shows that coverage is stronger in several asset-intensive sectors that are likely to be material for flood-risk analysis, including electricity, gas and steam supply, mining and quarrying, and manufacturing. These sectors achieve higher match rates by loan value, reflecting the depth of the S&P Sustainable1 dataset in picking up borrowers in asset-intensive industries with more extensive public disclosure. Coverage is comparatively weaker in agriculture and construction, where asset footprints may be more dispersed and less systematically documented in commercial databases. It is worth noting that a borrower is considered as covered as long as it maps to assets in the S&P database, which does not suggest the comprehensiveness of the list of the assets.

Because georeferenced asset locations and asset-type classifications are only available for borrowers matched to the S&P Sustainable1 asset database, the asset-level loss estimation is restricted to matched borrowers. Unmatched borrowers are excluded from the asset-level framework. Reported portfolio loss metrics should therefore be interpreted as conditional on matched coverage.

Figure 12 Loan book coverage for matched data

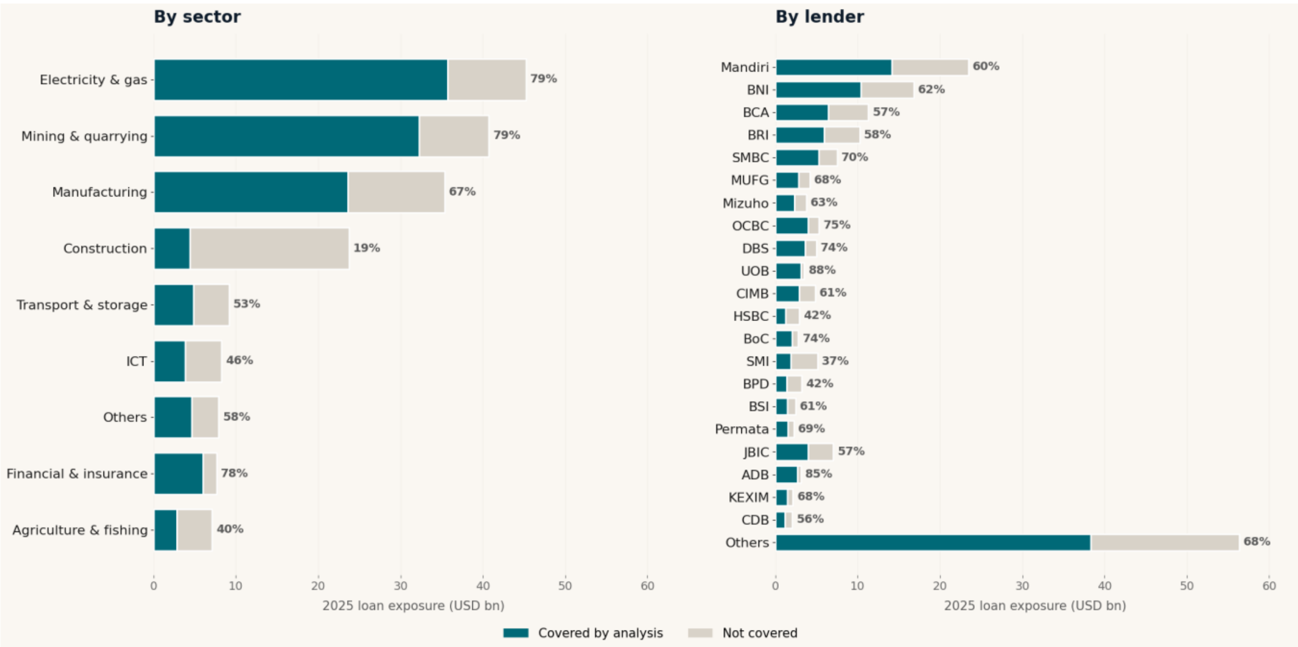

2. Asset to Loan Value Approximation

Loan value carries the most information for asset-level physical-risk analysis when it can be attached to the specific asset that the loan finances or secures. Under project finance and asset-backed lending, a facility is tied to an identifiable operating asset, so the loan value attributable to asset i of borrower j, written E_i, could in principle be read directly from the use of proceeds and the pledged collateral, with these asset-level loan values summing to the borrower's recorded total E_j.

In the reconstructed SGFIN loan book, however, loan value is recorded only at the borrower level — a single figure E_j per borrower — with no attribution to the borrower's constituent assets. General-purpose corporate lending is not earmarked to particular facilities, and the syndicated-loan records do not link a loan to the specific collateral behind it, so the true asset-level allocation is not observable.

We therefore apportion each borrower's loan value across its assets in proportion to their relative asset values, reusing the asset-value weights w_{k_i} recovered in Section V, where k_i is the asset type of asset *i* and the reference type is normalised to $w_{\text{ref}} = 1$. The loan value assigned to asset *i* is

$$E_i = E_j \cdot \frac{w_{k_i}}{\sum_{i' \in A_j} w_{k_{i'}}}, \quad i \in A_j$$

where the sum runs over the assets i' of borrower j, so that higher-value assets carry a proportionally larger share of the borrower's loan value and the assigned loan values recover the borrower total, $\sum_{i \in A_j} E_i = E_j$.

This apportioned loan value is the value base to which the projected loss rates $I^f_{i,t,s}$ of Section V are applied: the bank's projected dollar loss on asset *i* is $I^f_{i,t,s}\, E_i$, which aggregates to the borrower and portfolio totals used in the concentration and tail analysis below. This approximation depends on the comprehensiveness of the S&P Sustainable1 asset inventory: where a borrower's physical assets are only partially captured, loan value is apportioned across an incomplete asset set and the resulting asset-level figures inherit that coverage gap. Banks and supervisors with access to internal asset registers and facility-level collateral records should rely on those more complete sources in place of this proxy.

3. Loan Portfolio Profile

The reconstructed portfolio is not evenly distributed across banks, sectors, or borrowers. This matters for physical risk analysis because projected flood losses are exposure-weighted: the same hazard intensity produces different portfolio implications depending on where loan exposure is concentrated. As documented in our previous report, corporate loan exposures are concentrated among a few large domestic commercial banks, with Bank Mandiri, BNI, BCA, and BRI accounting for around one-third of reconstructed exposure, while the top 21 lenders account for roughly 70 percent. Sector exposures are similarly concentrated in electricity, mining, manufacturing, and construction, sectors that tend to involve large, fixed, and geographically specific physical assets.

For flood-risk analysis, this concentration structure matters because portfolio losses may be driven not only by average sector exposure, but by the overlap between large borrower exposures and high-hazard asset locations. A bank may appear diversified at the sector level while still carrying concentrated physical risk if several large borrowers hold assets in flood-prone areas. The asset-level analysis below therefore treats the loan book not simply as a sector allocation, but as a borrower- and asset-linked exposure structure.

4. Results

4.1. System-wide estimates are broadly consistent in magnitude, but not identical

Figure 13 reports the distribution of projected asset losses across the constructed portfolio in 2050 under SSP2-4.5. The interquartile ranges (IQR) and exposure-weighted means are also shown along with the distribution. At the company level, three providers return average estimates within a range of 0.33 to 0.50 percent. S&P yields the lowest estimate at 0.33 to 0.34 percent [b]. JBA produces a higher estimate of

[b] In principle, the S&P company-level series — which is itself a weighted aggregation of underlying asset-level data — should converge with the S&P asset-level series when both are aggregated to the same loan portfolio. The residual difference of 0.01% is

0.39 percent, while Swiss Re produces the highest estimate, 0.50 percent, which is about 50 percent more than the projected loss estimates from S&P.

At the asset level, the averages for JBA and Swiss Re are unchanged at 0.39 and 0.50 percent respectively, as these series are derived through SGFIN self-aggregation from the same asset-level inputs. The spread across these system averages, however, understates substantial differences in how each provider distributes losses across borrowers and assets. The remainder of the chapter examines this distributional structure.

Figure 13 Mean projected asset loss by provider

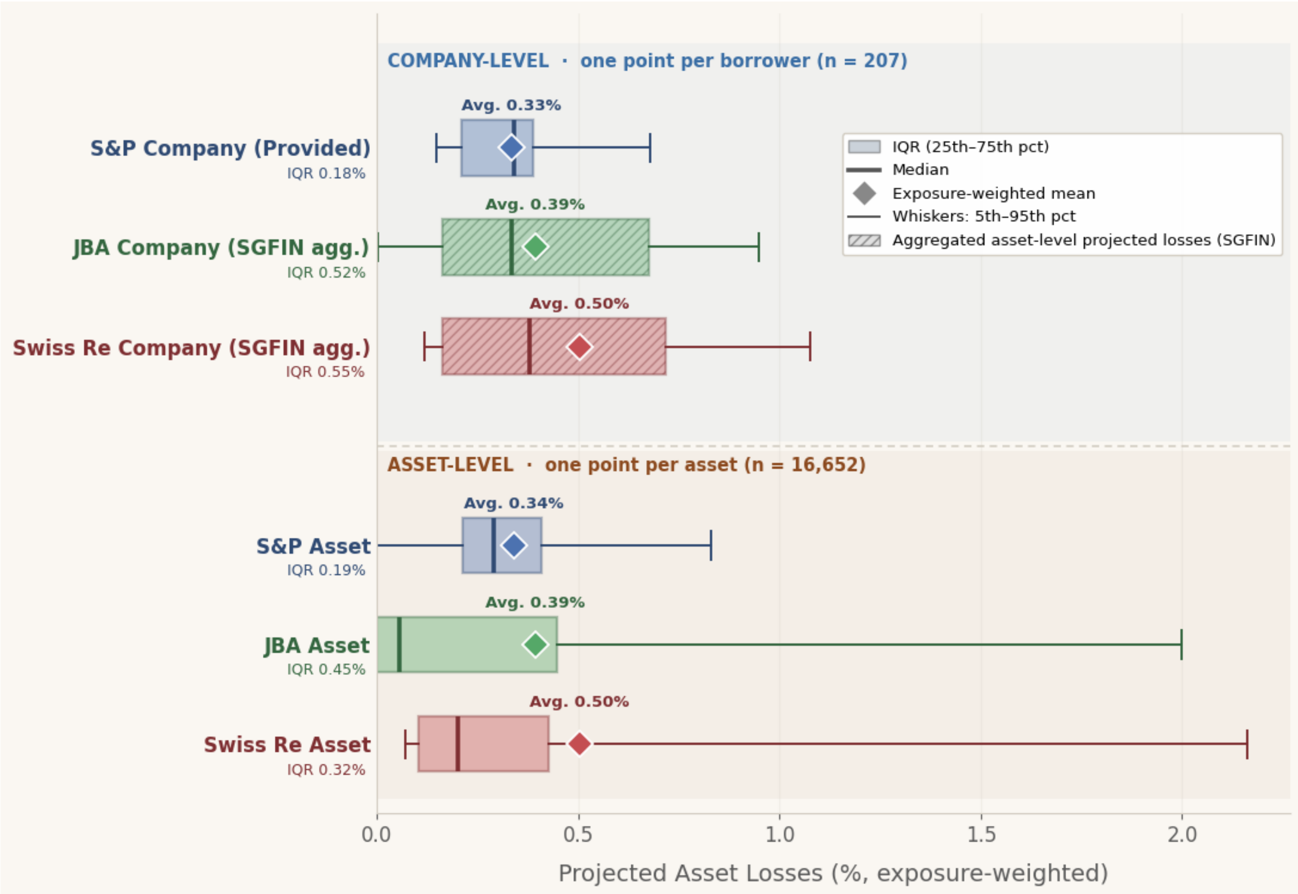

4.2. Dispersion differs more than the mean

While the system-level means are similar across the three providers, the spread of projected asset losses differs meaningfully. This is visible in Figure 13, where S&P Company series has the shortest whisker as compared to other asset-level series. The company-level whiskers are tighter than the corresponding asset-level distribution because the company-level series aggregates a borrower's high- and low-loss assets into a single per-firm number, dampening within-borrower variation in the process. The asset-level series, by contrast, preserve the individual asset observations and therefore reveal a wider spread across facilities.

Figure 14 reports the exposure-weighted interquartile range (IQR) which measures the difference in projected asset loss between the top 25-percentile risky exposure and the bottom 25-percentile risky exposure. The compression is

attributable to imputed values used to fill missing asset-level impacts.

visible in IQR, but more pronounced in the 95th-percentile tails.

For banks, the dispersion gap matters. Provisioning, capital allocation, and tail-risk monitoring depend not only on the portfolio mean, but also on the distribution of losses across obligors and assets. A loan book that appears tightly clustered around the mean under the company-level view exhibits substantially more heterogeneity once the underlying asset-level scores enter the assessment. We return to this point with bank-level distributions in a later section "4.4. Why can averages be misleading?".

Within the asset-level view, the three providers also return different IQRs. S&P returns the narrowest spread at 0.19 percent, partly reflecting the coarser hazard grid discussed earlier — assets within a larger spatial unit share the same hazard score, and therefore similar projected losses. JBA and Swiss Re use finer grids: JBA returns the widest IQR at 0.45 percent, more than twice the S&P Asset width, while Swiss Re sits between at 0.32 percent. Of the two, Swiss Re concentrates most assets in its lowest hazard band. The bulk of locations register as low risk, with only a handful of elevated hot spots, so the central mass of projected losses clusters tightly and the interquartile spread stays narrow even as the upper tail runs high.

All three asset-level views show materially wider dispersion than the company-level series, notwithstanding their differing IQRs. That this wider spread persists across providers using different hazard grids and scoring methods suggests it reflects genuine variation in the underlying assets rather than the modelling choices of any single vendor. Mean-based measures of physical-risk exposure will therefore understate this variation whichever provider supplies the scores.

Figure 14 Projected asset losses Interquartile Range (IQR) by provider

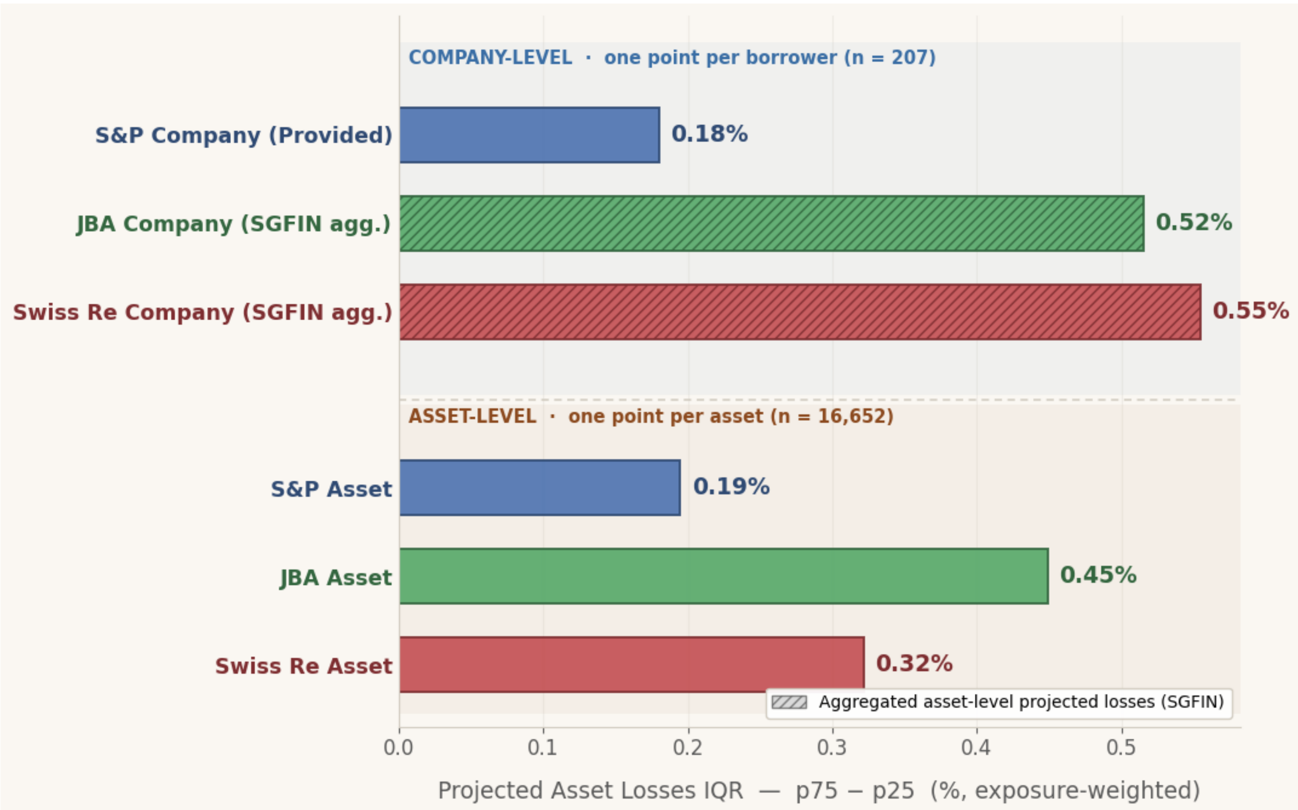

4.3. Distribution of asset-level projected losses

Figure 15 breaks the system-level distribution shown in Figure 13 into bank-level distributions. At the top of each panel, the Figure 13 system level distributions are labelled as "Overall". Left panel shows the company-level distribution under the three data sources. Right panel show the asset-level distribution under each of the three data sources, with each observation a single asset rather than a borrower.

Figure 15 Distribution of projected asset loss by provider

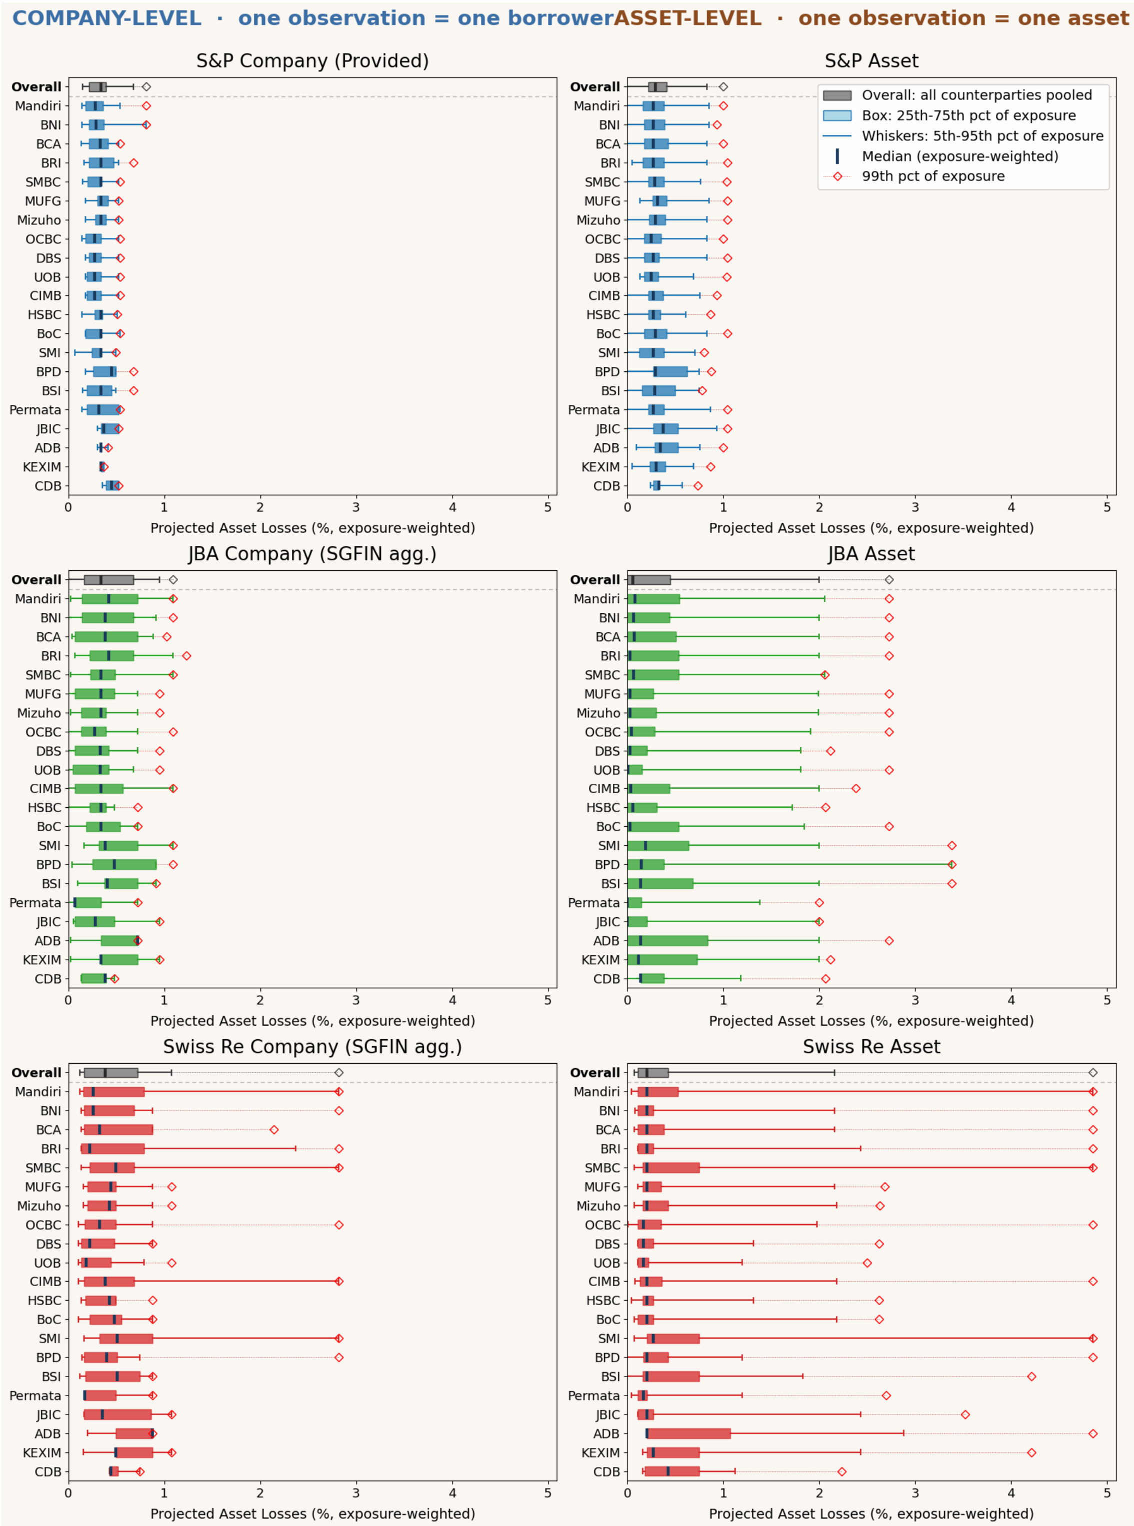

Comparing company level view (left panel) with asset-level view (right panel) produces visibly wider distributions for nearly every bank in the chart. The widening is most pronounced in the upper tail: the 99th-percentile thresholds in the asset-level panels reach several times the borrower-level values for many of the same banks. Borrower-level aggregates smooth over substantial within-firm dispersion that is material for tail-risk monitoring and stress testing.

The 99th-percentile thresholds are more sensitive to provider choice. In this common-sample analysis, Swiss Re returns the widest tails for these banks, S&P the narrowest, and JBA falls between the two. Whether across aggregation at the company level (left panel) or at the asset level (right panel), bank distributions look broadly similar near the centre and diverge most visibly in the tail. The differences sit predominantly in the part of the distribution that matters most for tail-risk monitoring.

4.4. Why can averages be misleading?

The previous sections showed that mean and median measures of projected asset loss smooth over substantial heterogeneity. It can therefore be misleading to rely on a single per-borrower number if within-borrower variation is material. Figure 16 illustrates the variation that this approach hides, showing the within-borrower distribution of asset-level losses for the ten largest borrowers, which jointly account for roughly one-third of loan portfolio. Each bubble is one asset, sized by its loan value; the vertical line is the borrower's asset value-weighted mean. Most of these borrowers are from electricity and mining industries, with the majority as SOE.

Figure 16 Within-borrower dispersion of asset-level losses

The distributions are wide and right skewed for most of these borrowers. For the largest electricity and mining borrowers in the chart, many individual assets sit well above the corporate average, in the cases of PLN and Kilang Pertamina BPN reaching multiple times the borrower-level mean under JBA and Swiss Re.

For some borrowers such as Pertamina Hulu and Pupuk Indonesia, some of the largest bubbles — assets that carry meaningful loan exposure — sit in the right tail rather than near the mean, indicating that the within-borrower spread is driven by assets that materially contribute to the bank's exposure to the borrower, not by minor outliers.

The company-level mean, by averaging across both groups, understates the projected loss at the borrower's most-exposed facilities. For banks lending against specific assets through project finance or asset-backed syndicated structures, the credit exposure attaches to the operating asset rather than to the parent entity's full balance sheet, and the loss-given-default-relevant view of the borrower is the asset distribution rather than the company mean. Loss provisions calibrated to the company-level number would understate the loss potential of the underlying collateral. For stress-testing and capital planning, the distribution assists with more informative interpretation of the data.

4.5. Loan portfolio is concentrated, but flood losses even more so

Risk concentration is central to prudential supervision: the Basel II framework identifies risk concentrations as *"arguably the single most important cause of major problems in banks"* [30]. Supervisors assess it with two measures in particular — the Gini coefficient and the Herfindahl–Hirschman Index (HHI) — both of which the European Banking Authority names as standard concentration indicators in its supervisory-review framework [31]. We apply both, together with the underlying Lorenz curve, at the asset level. The Lorenz curve [32] plots the cumulative share of a quantity — projected flood loss or exposure — against the cumulative share of the assets that carry it; the 45° diagonal denotes perfect equality, and the further the curve bows below it, the greater the concentration. The Gini coefficient [33] reduces that bow to a single number between 0 (perfect equality) and 1 (the entire quantity in one asset). These tools are standard in credit-portfolio analysis as mentioned in work for the Bank for International Settlements [34].

Both quantities entered into the Lorenz curves are measured at the asset level. The loan value on asset *i*, E_i — the quantity shown as the loan portfolio curve — is the amount apportioned to it from its borrower's loan value E_j in proportion to relative asset value,

$$E_i = E_j \cdot \frac{w_{k_i}}{\sum_{i' \in A_j} w_{k_{i'}}}, \quad i \in A_j$$

and is common to all three hazard providers. The projected dollar loss on the same asset, D_i, scales that loan value by the asset's projected loss rate $I^f_{i,t,s}$ of Section V,

$$D_i = I^f_{i,t,s} \cdot E_i$$

so that the dollar loss varies across the three providers through the hazard input, while the loan value does not.

Each Lorenz curve is constructed over the full set of *n* assets in the constructed loan book, separately for the two quantities of interest, projected dollar loss and loan portfolio, based on loan value apportioned to the assets as described in Section VI.2. For a given quantity, let $x_i \geq 0$ denote the amount carried by asset *i*; the assets are first ordered from smallest to largest by that quantity,

$$x_{(1)} \leq x_{(2)} \leq \cdots \leq x_{(n)}$$

where $x_{(i)}$ is the *i*-th order statistic. The curve then traces, across ranks $k = 0,1,\ldots,n$, the cumulative share of assets against the cumulative share of the quantity those assets account for,

$$p_k = \frac{k}{n} \quad , \quad L_k = \frac{\sum_{i=1}^{k} x_{(i)}}{\sum_{i=1}^{n} x_{(i)}}$$

anchored at the origin $(p_0, L_0) = (0, 0)$. The denominator $\sum_{i=1}^{n} x_{(i)}$ is the portfolio total: dividing the running sum by it expresses each point as a proportion on a common 0-to-1 scale, so that L_k is the fraction of total loss (or loan portfolio) borne by the smallest-ranked p_k share of assets, and the curves are directly comparable across hazard providers and between the two quantities.

The Gini coefficient summarises this curve as twice the area between it and the 45° line of equality, obtained directly from the ranked asset shares,

$$G = 1 - 2\int_0^1 L(p)\,dp = \frac{2\sum_{i=1}^{n} i \cdot x_{(i)}}{n\sum_{i=1}^{n} x_{(i)}} - \frac{n+1}{n}$$

By construction the index lies between 0 and $1 - 1/n$, approaching its upper bound as the quantity concentrates in a single asset.

Figure 17 Concentration of projected dollar loss versus loan portfolio, by hazard provider

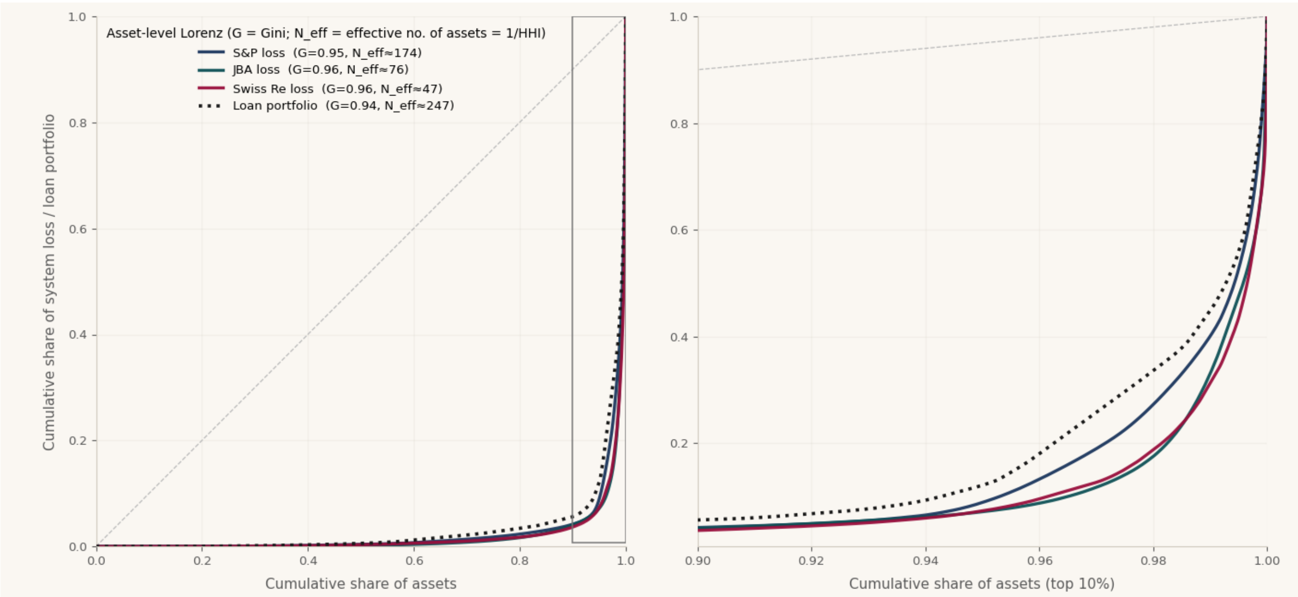

Asset-level Lorenz curves across the full loan book (2050 SSP 2-4.5), all counterparties. Solid = projected flood loss per vendor; dotted = loan portfolio (vendor-independent). Right panel enlarges the top decile of assets. The loss curves sit below the loan-portfolio curve -- loss is more concentrated -- so the top assets carry a larger share of loss than of loan portfolio.

We compute these measures at the asset level across the inner-set loan book (16,652 assets owned by the portfolio's borrowers, 2050 under SSP2-4.5, fluvial and pluvial flooding combined, aggregating all Indonesian counterparties). In Figure 17, exposure to each asset is independent of the hazard provider and therefore appears as a single reference curve, while projected loss is shown separately for S&P, JBA and Swiss Re. The loan portfolio is already highly unequal, with a Gini of 0.94: a small set of large borrowers and assets dominates the book. Projected loss is more concentrated still. Its Lorenz curve lies below the loan portfolio curve for every provider, with Gini coefficients of 0.95 (S&P), 0.96 (JBA) and 0.96 (Swiss Re). In other words, flood risk compounds the loan book asset concentration, loading a disproportionate share of projected loss onto an even narrower set of assets than loan values alone would suggest.

The size of that compounding is clearest when the asset set is held fixed. We rank assets by projected loss, take the top *N*%, and report that *same* set's share of total loss alongside its share of total the loan portfolio. The gap between the two shares is the concentration that flood risk adds beyond the lending pattern.

Table 8 Share of projected dollar loss versus loan portfolio held by the top dollar loss-ranked assets

	Loss attributed to riskiest 1% of assets (by $loss)	Loans attributed to riskiest 1% of assets (by $loss)	Loss attributed to riskiest 5% of assets (by $loss)	Loans attributed to riskiest 5% of assets (by $loss)	Loss attributed to riskiest 10% of assets (by $loss)	Loans attributed to riskiest 10% of assets (by $loss)
S&P	62	51	93	84	96	88
JBA	73	25	94	54	97	56
Swiss Re	69	38	93	84	96	91

As shown in Table 8 the top 1% of loss-driving assets account for 62–73 percent of projected system loss but only 25–51 percent of the loan portfolio; at the top 5% they carry 93–94 percent of loss against just 54–84 percent of the loan portfolio. The effect is most pronounced under JBA, where the top 1% of assets bear 73 percent of projected loss but only 25 percent of the loan portfolio, a 48 percentage-point gap.

In bank supervision, the HHI and the Gini coefficient are the standard measures of portfolio concentration. The Herfindahl–Hirschman Index captures the concentration through the dispersion of the asset shares, and, unlike the Lorenz construction, requires no ranking, being a symmetric function of the shares alone.

Retaining the quantity x_i carried by asset i, each asset's share of the portfolio total and the index are

$$s_i = \frac{x_i}{\sum_{j=1}^{n} x_j} \quad , \quad \mathrm{HHI} = 10{,}000 \cdot \sum_{i=1}^{n} s_i^2$$

where the multiplier 10,000 sets the conventional supervisory scale. Squaring the shares weights the measure towards the largest holdings, so the index rises steeply as the quantity concentrates in a few assets, attaining its maximum of 10,000 when one asset carries the entire total and falling to $10{,}000/n$ when all n assets are of equal size.

Because that lower bound depends on the number of assets [c], we therefore read it through its size-invariant numbers equivalent,

$$N_{\mathrm{eff}} = \frac{10{,}000}{\mathrm{HHI}} = \frac{1}{\sum_{i=1}^{n} s_i^2}$$

which expresses concentration as the number of equally sized assets that would reproduce the observed index and oris directly comparable across the two quantities and the three hazard providers.

As shown in Table 9, the loan portfolio has an HHI of 40.5, translating to N_{eff} of 247. A loan portfolio with only 247 assets, each taking equal share of the loan will produce the same HHI concentration measure. However, for projected asset loss, N_{eff} collapses to between 174 (S&P) and 47 (Swiss Re) — a striking contraction given a universe of more than sixteen thousand assets. It is worth noting that Swiss Re has the highest HHI, translating to an equivalent portfolio whose losses only distribute over 47 assets equally, only one-third of number of Loss - S&P and one-fifth of the loan portfolio.

[c] Supervisors typically apply HHI at the level of borrowers, sectors or regions — typically a few hundred units at most — whereas we compute them at the asset level, across 16,635 assets. A raw HHI is therefore not directly comparable to supervisory benchmarks calibrated to borrower or sector level granularity: on the 0–1 scale a perfectly even asset book's HHI already sits at $1/n \approx 6\times10^{-5}$.

Table 9 Summary concentration measures for flood loss and loan portfolio

	Gini	HHI (x10,000)	Effective no. of assets (N_{eff})
Loan portfolio	0.94	40.5	247
Loss – S&P	0.95	57.4	174
Loss - JBA	0.96	131.5	76
Loss – Swiss Re	0.96	213.7	47

Notably, the HHI and the Gini coefficient do not agree on which provider is the most concentrated. The HHI squares the asset shares and is therefore dominated by the few largest assets, ranking Swiss Re highest ($N_{\text{eff}} \approx 47$); the Gini reflects the full distribution and ranks JBA highest (0.96). Rather than a contradiction, this reflects the complementary information the two measures carry — Swiss Re's loss is driven by a small number of extreme assets, whereas JBA's is spread across a broader but still highly concentrated tail. Either way, both measures point to the same conclusion: projected flood loss is concentrated well beyond the exposure that underlies it. We turn next to what those few decisive assets represent in absolute dollar terms, and who bear the losses.

4.6. Dollar losses are concentrated in few capital-intensive assets

Figure 18 translates projected loss rates $I^f_{i,t,s}$ into estimated dollar losses D_i by combining each asset's projected loss rate with its estimated bank exposure E_i, where $D_i = I^f_{i,t,s} \cdot E_i$.

As previously defined, the loan value on asset i, E_i is the amount apportioned to it from its borrower's loan value E_j in proportion to relative asset value,

$$E_i = E_j \cdot \frac{w_{k_i}}{\sum_{i' \in A_j} w_{k_{i'}}}, \quad i \in A_j$$

and is common to all three hazard providers. The dollar loss varies across the three providers through the impact rate, while the estimated loan value for each asset does not.

Each rectangle is one asset, sized by its estimated loss D_i and coloured by asset type. The bold outline demarcates the assets with the largest dollar losses, that jointly account for the top quartile of estimated loss under each provider. The right panels show each bank's share of the top-quartile loss region.

All three providers point to capital-intensive assets as the dominant contributors to estimated portfolio loss, consistent with their steep damage functions, where even moderate hazard intensities translate into sizeable dollar losses once exposure values are large. The composition and concentration of the top quartile, however, differ across the three views. Under S&P, the top quartile spans 21 assets and is comparatively fragmented, with a substantial share from mining properties across several medium-sized sites. Under JBA and Swiss Re, it concentrates in just seven assets respectively, dominated by large power generation assets. The bank-level absorption of the top-quartile loss varies across providers, but the variation sits mainly within rather than across major banking groups. The four large Indonesian commercial banks jointly absorb a similar share under all three views — 37 percent under S&P, 45 percent under JBA, and 44 percent under Swiss Re — making this aggregate the most stable feature of the bank-level loss distribution. Japanese commercial banks see their combined share rise from 7 percent under S&P to 10 percent under JBA and 17 percent under Swiss Re, consistent with their deeper participation in the large power and resource projects that JBA and Swiss Re identify as the main loss drivers.

Singaporean commercial banks hold roughly steady (~6-8 percent), while development finance institutions vary from around 8 percent under S&P to 5 percent under JBA and Swiss Re. The same broad shortlist of absorbing institutions emerges across all three providers — the four large Indonesian banks together with Japanese commercial banks and DFIs concentrated in flood-exposed projects — but the relative weight on each group depends on the provider.

This matters for supervisors because the identity of major absorbing bank groups appears relatively stable, while the severity assigned to each group remains provider sensitive. From a macroprudential perspective, this suggests that supervisory monitoring may focus on a broadly consistent set of systemically relevant lenders and project-finance exposures. From a microprudential perspective, however, the assessment of specific obligors, collateral pools, and asset-backed exposures may vary materially depending on the hazard dataset embedded in a bank's risk framework.

.

Figure 18 Top-quartile dollar losses by individual asset

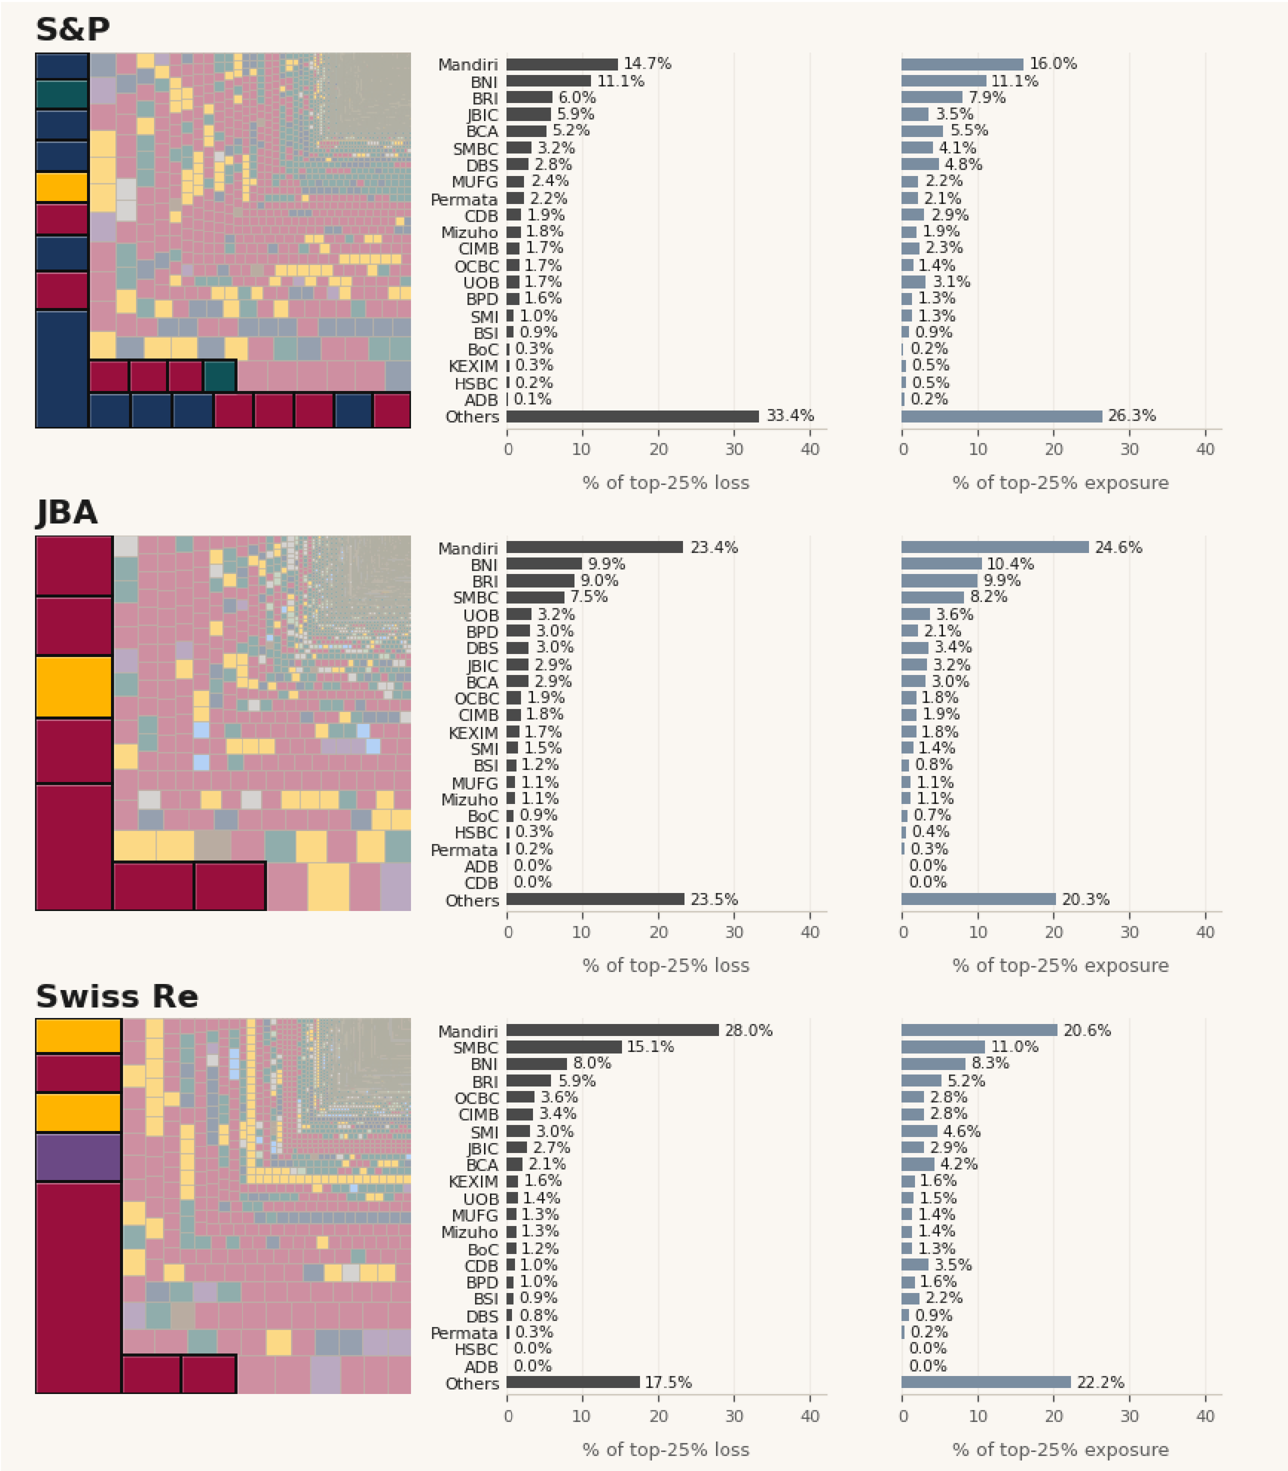

Asset type

- Power Plant Site
- Office
- Oil & Gas Upstream
- Mining Properties
- Port
- Retail Store
- Property
- Cement
- Others
- Beyond top 25%

Per-asset flood-induced loss across all Indonesian banks, 2050 SSP 2-4.5. Treemap rectangles (one per asset) are sized by the asset's projected dollar loss and coloured by asset type; bold outlines mark the top 25% loss-making assets (21 / 7 / 7 assets for S&P / JBA / Swiss Re). Middle bars show each lender's share of that top-25% loss, sorted high to low with Others last; right bars show each lender's share of total exposure to the same top-25% assets (both sum to 100%). A lender whose loss share exceeds its exposure share holds riskier-than-average assets among the worst losers.

4.7. Single-provider estimates capture only part of the pooled tail

Figure 19 shows the marginal information added by combining the three asset-level data sources. The lower three bars show the estimated dollar loss in each provider's own top quartile of loss assets. The top bar shows an upper-bound estimate constructed by pooling all assets that appear in any one of the three top quartiles, and applying the higher loss estimate per asset per flood type, then taking the sum of inland flood losses.

Across the universe of 16,652 exposed assets, each provider individually flags 4,163 top-quartile assets. The union of these three sets covers 8,072 assets — nearly twice any single provider's list — indicating that the three providers agree on only about half of the assets that fall in any of their respective top quartiles. In dollar terms, the pooled estimate reaches USD 920 million, compared with single provider estimates of USD 245 million under S&P, USD 402 million under JBA, and USD 432 million under Swiss Re. Each individual source captures between 27 and 47 percent of the pooled tail loss, leaving half to two thirds of the risk uncaptured. Combining with another source will increase the coverage to 71 percent to 75 percent.

For banks and supervisors, the practical implication is that physical-risk screens built on a single data source provide only a partial view of potential flood risky assets in the portfolio. Combining sources at the screening stage expands the visible set of high-exposure assets and reduces the chance that a concentrated exposure identified by one provider but not another is overlooked. Each provider applies an internally consistent methodology, and the providers are best treated as complementary rather than competing inputs. The implication is not to select a single "correct" provider, but to require banks to understand and disclose the sensitivity of their physical risk estimates to provider choice, particularly in the upper tail.

Figure 19 Single-provider and combined tail-loss estimates

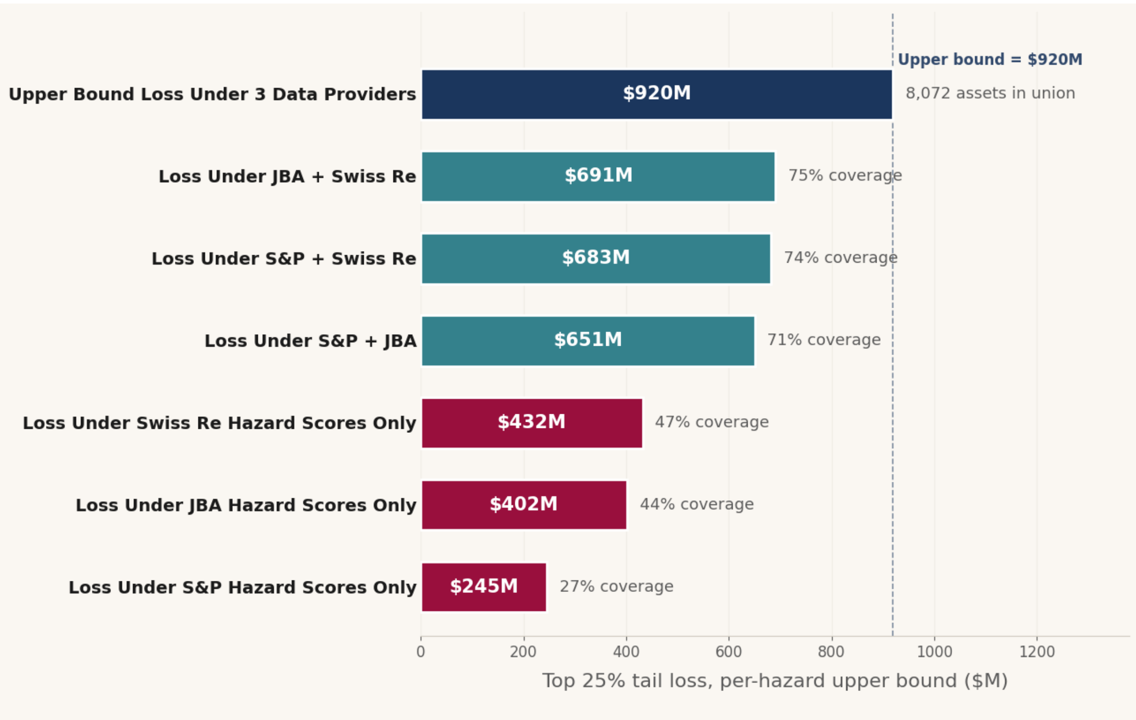

Notes: the universe is the 16,652 assets scored by all three data providers across the constructed loan book. Each provider's top-25% tail is the 4,163 assets it ranks highest by combined flood impact rate (fluvial + pluvial); each bar covers the union of its named providers' tails. Every asset's loss is calculated as the asset's estimated loan exposure multiplied by the sum of the inland-flood impact rates, with each flood's impact rate taken as the highest impact rate across the named providers.

VII. Implications for Bank Risk Management

1. Practical Use Cases

The framework can be used as a practical risk-triage tool within existing credit and supervisory workflows. Our framework helps identify borrowers, facilities, asset types, locations, and bank groups that warrant closer review because their risk profile becomes more visible at the asset level or under alternative hazard providers. For supervisors, the framework supports microprudential challenge of banks' physical-risk frameworks and macroprudential monitoring of system-wide exposure clusters.

1.1. For banks

For banks, the framework has three practical applications. The first is asset-level credit screening as a second-layer filter in loan origination, credit approval, and periodic review. Company-level physical risk metrics remain useful for initial triage, but they can smooth over material within-borrower heterogeneity. As shown in the results, the S&P company-level IQR is 0.18 percent, while asset-level IQRs range from 0.19 percent under S&P Asset to 0.45 percent under JBA. This suggests that borrower-level averages may be sufficient for broad screening, but less informative where large borrowers hold assets across different hazard zones.

The second application is collateral and project finance review. Once a loan has been identified as material or potentially sensitive, asset-level analysis provides a more transaction-specific lens for assessing collateral vulnerability, business interruption risk, insurance adequacy, and borrower adaptation readiness. This is particularly relevant in project finance and asset-backed lending, where credit exposure may be linked more closely to a specific operating asset than to the parent company's consolidated balance sheet. For this use case, the framework should be supplemented by borrower-specific information not available in this study, including collateral value, asset condition, insurance coverage, operational criticality, and adaptation measures.

The third application is portfolio concentration monitoring. The results show that flood-related losses are not evenly distributed across borrowers, assets, or bank groups. Each provider identifies a partially overlapping set of high-loss assets, and each individual provider captures only around 27-47 percent of the pooled tail-loss estimate. This makes the framework useful for constructing physical-risk watchlists, identifying recurring high-risk asset types or regions, and testing whether concentration risk remains visible under alternative hazard views.

1.2. For supervisors

For supervisors, the framework provides a practical way to move from climate-risk awareness to supervisory challenge. At the microprudential level, asset-level analysis can help supervisors assess whether banks' internal physical-risk frameworks capture facility-level heterogeneity, especially for large corporate, project finance, and collateral-backed exposures. The key question is not simply whether a bank has adopted a physical-risk data provider, but whether it understands how its results change when hazard inputs, aggregation levels, or vulnerability assumptions change.

At the macroprudential level, the framework can support monitoring of system-wide exposure clusters. The four large Indonesian commercial banks absorb around 40 percent of top-quartile losses under all three provider results, making multi-provider analysis useful for identifying asset-level blind spots while stress-testing the same core group of exposed institutions. At the same time, the allocation of tail-loss pressure is provider-sensitive, for example, Japanese commercial banks' share ranges from 7 percent under S&P to 17 percent under Swiss Re. This makes the framework useful not only for identifying which institutions warrant closer monitoring, but also for understanding how severity estimates shift under alternative provider assumptions.

Concentration diagnostics documented in the results provide supervisors with an additional quantitative lens. Gini coefficients and HHI metrics show that projected loss concentration is materially higher than exposure concentration across all three providers, and that the degree of narrowing differs by provider. These metrics are particularly useful for macroprudential monitoring because they translate the distributional findings into a single comparable indicator of how concentrated flood-related tail losses are relative to the underlying credit exposure.

2. Recommendation

2.1. For banks

Banks should strengthen the asset-level information base underlying physical-risk assessment. This study relies on proxy-based asset values and study-specific scaling assumptions because externally available data do not fully capture facility-level collateral values, asset condition, insurance coverage, adaptation measures, or operational criticality. Banks, however, are better positioned to obtain and maintain this information through credit origination, borrower due diligence, collateral review, and ongoing monitoring. For material exposures, asset-level analysis should therefore be supported by more complete information on geolocation, asset type, collateral linkage, exposure value, insurance coverage, operational role, and borrower adaptation plans.

Banks should also develop internal capacity to translate physical-risk signals into credit-relevant judgements. Vendor outputs can provide useful hazard and impact indicators, but they should not substitute for bank-specific assessment of asset value, collateral recoverability, business interruption risk, debt-servicing capacity, and borrower resilience. For material facilities, this implies the need for asset-level stress testing before credit approval and during periodic portfolio review, particularly where repayment depends on geographically fixed and capital-intensive assets.

Provider choice should be governed as a source of model risk. Banks do not necessarily need full multi-provider integration for every exposure, but they should understand how key outputs may change under different hazard datasets, scenarios, time horizons, aggregation methods, and loss-scaling assumptions. For material exposures, periodic cross-provider comparison, benchmark testing, or challenger-model checks could help identify where reliance on a single provider may create blind spots.

Banks should document the hazard provider or providers used, the scenarios and time horizons applied, the aggregation methodology, the treatment of asset values and collateral, and the sensitivity of key outputs to provider choice. Such documentation would strengthen internal model governance, credit committee review for material exposures, climate stress-testing records, and, where relevant, external comparability.

2.2. For supervisors

For microprudential supervisors, the relevant question is not simply whether a bank has adopted a physical-risk data provider, but whether the bank has the internal capacity, data quality, and governance processes needed to interpret provider outputs and translate them into credit-relevant judgements. Supervisory review could therefore examine whether banks' physical-risk frameworks capture asset-level heterogeneity, whether collateral and facility-level assumptions are adequately documented, and whether stress-test outputs are sensitive to alternative hazard datasets or aggregation assumptions.

Supervisors may also need to develop capacity to assess physical-risk vendors and modelling approaches. The objective is not to endorse a single preferred provider, but to understand how providers differ in spatial resolution, hazard coverage, hydrological assumptions, scenario treatment, calibration choices, and vulnerability translation. This would strengthen supervisory review of banks' vendor selection, model governance, and provider-sensitivity testing, while helping authorities assess which data sources are suitable for particular supervisory use cases in Indonesia.

A further step would be to require banks to report minimum asset-level data for material exposures, including geolocation, asset type, collateral linkage, exposure value, insurance coverage, and hazard-provider metadata. Without this information, supervisors cannot independently verify or challenge banks' physical risk assessments, limiting the effectiveness of supervisory review.

Macroprudential authorities may consider using multi-provider asset-level screening to monitor systemic flood-risk concentration. Under different providers' methodology, they need to identify geographic hotspots where flood hazards overlap with dense economic activity, critical infrastructure, large corporate borrowers, or concentrated bank exposures. These locations may warrant closer monitoring because correlated physical shocks could affect multiple borrowers, lenders, or economically important sectors at the same time.

Macroprudential monitoring should therefore focus on ranges, clusters, and robustness rather than point estimates alone. Stable risk concentration clusters can indicate areas for closer system-level monitoring, while material divergence across providers can itself provide useful information about uncertainty and potential blind spots. Concentration diagnostics such as Lorenz curves, Gini coefficients, HHI, effective number of assets, and top-loss asset shares can complement mean loss estimates by showing whether systemic flood-risk pressure is concentrated in a small number of assets, borrowers, regions, or bank groups.

Beyond identifying geographic hotspots, macroprudential authorities should assess whether the concentration of flood-exposed project finance among systemically important institutions creates potential contagion channels. Where systemically important banks hold concentrated exposures to the same high-risk assets or regions, correlated physical shocks could transmit losses simultaneously across multiple systemically relevant lenders. For economically necessary but high-risk projects, authorities may also consider whether existing risk-transfer and resilience arrangements, including insurance coverage, adaptation requirements, and risk-sharing structures, are sufficient to prevent correlated physical risk from becoming concentrated on a small number of lender balance sheets.

Addressing physical climate risk at scale requires coordination beyond individual banks and supervisors. The vulnerability puzzle in this study illustrates why such coordination is needed. Hazard providers holding location-specific physical-risk information and modelled asset sensitivity, while banks hold exposure values, collateral information, borrower relationships, and facility-level credit terms. Without a standardised bridge between these datasets, physical climate-risk assessment remains fragmented, provider-sensitive, and difficult to compare across institutions. Supervisors are well positioned to act as conveners, promoting coordination among banks, hazards modelling providers, and standard-setting bodies and developing the data and disclosure standards needed to make physical risk assessment more consistent and comparable.

A practical starting point would be minimum asset-level data standards for large corporate and project finance exposures. Disclosure expectations could similarly evolve to include the hazard provider or providers used, the scenarios and time horizons applied, and the aggregation methodology behind reported physical-risk estimates, enabling supervisors and market participants to determine whether differences in reported physical risk reflect genuine portfolio differences or methodological differences in provider choice, scenario selection, and aggregation level.

Appendix

Figure A- 1 Regional distribution of high flood risk exposure across Indonesia under climate change

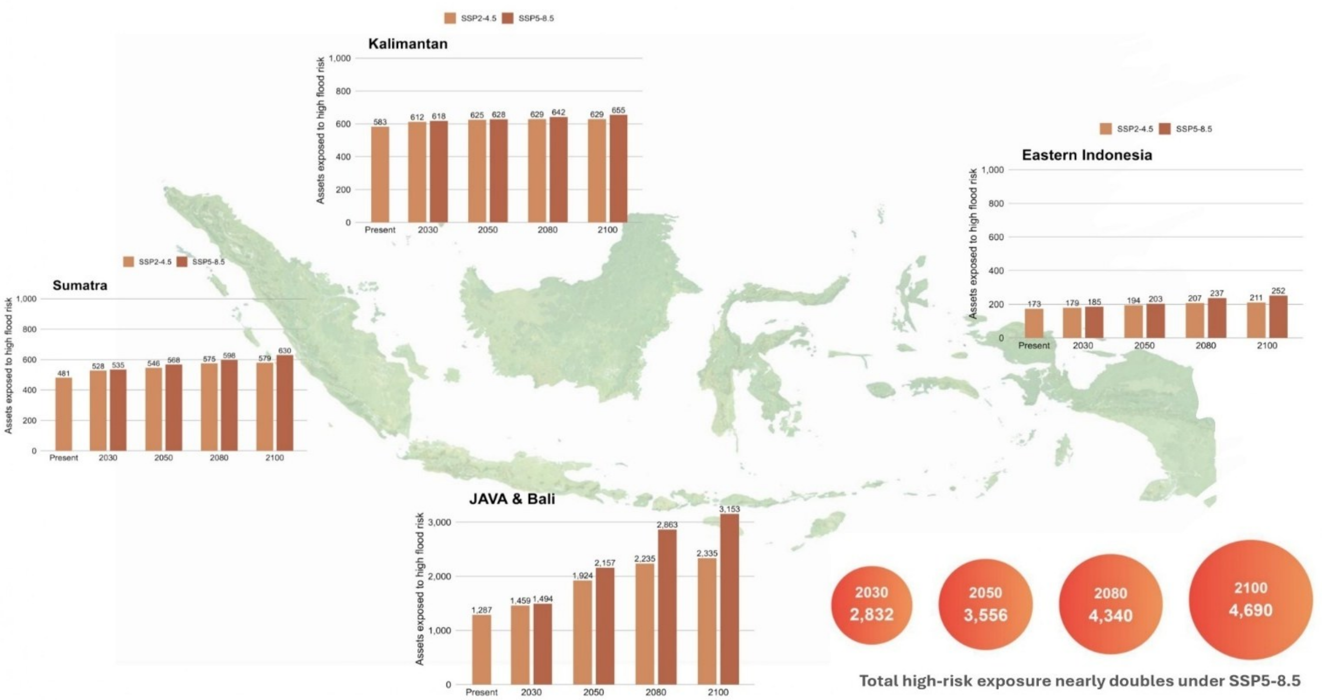

Each panel shows the number of assets exposed to high inland flood risk, defined on combined fluvial and pluvial hazard using JBA defended scores, in one of four regions (Sumatra, Kalimantan, Java and Bali, and Eastern Indonesia).

Reference List

[1] Vinzenz Peters, "How Banks Are Impacted by and Mediate the Economic Consequences of Natural Disasters and Climate Shocks: A Review," *De Economist* 173, no. 1 (March 2025): 47–85, https://doi.org/10.1007/s10645-024-09441-7.

[2] Rumi Mahmood, Pam Palena, and David Carlin, *What the Market Thinks: How Global Insurers Are Responding to Rising Physical Risk* (n.d.), accessed March 3, 2026, https://media.licdn.com/dms/document/media/v2/D4E1FAQH3VGO_h5Ardg/feedshare-document-pdf-analyzed/B4EZyuC0Q7IMAY-/0/1772446502040?e=1773273600&v=beta&t=Kx2jZmgPmgYOtVY9MK0oV1LES99YX7gwecCuwsyuN4w.

[3] Henry He Huang et al., "Firm Climate Risk, Risk Management, and Bank Loan Financing," *Strategic Management Journal* 43, no. 13 (December 2022): 2849–80, https://doi.org/10.1002/smj.3437.

[4] Luca Barbaglia, Serena Fatica, and Caterina Rho, "Flooded Credit Markets: Physical Climate Risk and Small Business Lending," *Mo.Fi.R. Working Papers*, Mo.Fi.R. Working Papers, July 2024, 186, https://ideas.repec.org//p/anc/wmofir/186.html.

[5] Shouwei Li and Xin Wu, "How Does Climate Risk Affect Bank Loan Supply? Empirical Evidence from China," *Economic Change and Restructuring* 56, no. 4 (August 2023): 2169–204, https://doi.org/10.1007/s10644-023-09505-9.

[6] World Bank Group and Asian Development Bank, *Climate Risk Country Profile* (World Bank, Washington, DC and Asian Development Bank, Manila, 2021), https://doi.org/10.1596/36379.

[7] World Bank, *Country Climate and Development Report*, with World Bank, Global Economic Prospects (The World Bank, 2023), 107–17, https://doi.org/10.1596/978-1-4648-0483-0_ch2_EAP.

[8] F. Zhang et al., *Managing Climate Risk Exposure: Transition and Physical Climate Risks in Bank Loan Portfolios Across the Indonesian Banking System* (SGFIN, 2025).

[9] Zhang et al.

[10] Giacomo Bressan et al., "Asset-Level Assessment of Climate Physical Risk Matters for Adaptation Finance," *Nature Communications* 15, no. 1 (July 2024): 5371, https://doi.org/10.1038/s41467-024-48820-1; Henk Jan Reinders, Dirk Schoenmaker, and Mathijs Van Dijk, "Climate Risk Stress Testing: A Critical Survey and Classification," *Journal of Climate Finance* 10 (March 2025): 100061, https://doi.org/10.1016/j.jclimf.2025.100061; NGFS, *Physical Climate Risk Assessment: Practical Lessons for the Development of Climate Scenarios with Extreme Weather Events from Emerging Markets and Developing Economies* (2022).

[11] Zhang et al., *Managing Climate Risk Exposure*.

[12] Phineas Glover et al., *How Can Investors Adapt to a +3oC Scenario?* (UBS, 2026); Bressan et al., "Asset-Level Assessment of Climate Physical Risk Matters for Adaptation Finance."

[13] Zhang et al., *Managing Climate Risk Exposure*.

[14] Reinders, Schoenmaker, and Van Dijk, "Climate Risk Stress Testing."

[15] Vasily Pozdyshev, Alexey Lobanov, and Kirill Ilinsky, *Incorporating Physical Climate Risks into Banks' Credit Risk Models*, 2025

[16] Nepomuk Dunz et al., "ECB's Economy-Wide Climate Stress Test," *SSRN Electronic Journal*, ahead of print, 2021, https://doi.org/10.2139/ssrn.3929178.

[17] World Bank Group (2021). Climate Risk Country Profile: Indonesia. https://doi.org/10.1596/36379.

[18] Aldrian, E., & Dwi Susaninto, R. (2003). Identification of three dominant rainfall regions within Indonesia and their relationship to sea surface temperature. *International Journal of Climatology: A Journal of the Royal Meteorological Society*, 23(12), 1435-1452.

[19] Muhammad, F. R., Lubis, S. W., & Setiawan, S. (2021). Impacts of the Madden–Julian oscillation on precipitation extremes in Indonesia. *International Journal of Climatology*, 41(3), 1970-1984.

[20] IPCC (2021). *Climate Change 2021: The Physical Science Basis.* Contribution of Working Group I to the Sixth Assessment Report. https://www.ipcc.ch/report/ar6/wg1/.

[21] Chen, M., Papadikis, K., & Jun, C. (2021). An investigation on the non-stationarity of flood frequency across the UK. *Journal of Hydrology*, 597, 126309.

[22] Rogers, J. S., Maneta, M. P., Sain, S. R., Madaus, L. E., & Hacker, J. P. (2025). The role of climate and population change in global flood exposure and vulnerability. *Nature Communications*, 16(1), 1287.

[23] O'Neill, B. C., Tebaldi, C., Van Vuuren, D. P., Eyring, V., Friedlingstein, P., Hurtt, G., ... & Sanderson, B. M. (2016). The scenario model intercomparison project (ScenarioMIP) for CMIP6. *Geoscientific Model Development,* 9(9), 3461-3482.

[24] IPCC, 2023: Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee

and J. Romero (eds.)]. IPCC, Geneva, Switzerland, 184 pp., doi: 10.59327/IPCC/AR6-9789291691647.

[25] Reinders, Schoenmaker, and Van Dijk, "Climate Risk Stress Testing."

[26] S&P Sustainable1, "For the World's Largest Companies, Climate Physical Risks Have a $1.2 Trillion Annual Price Tag by the 2050s," March 10, 2025, https://www.spglobal.com/sustainable1/en/insights/special-editorial/ceraweek-physical-risk.

[27] S&P Sustainable1.

[28] S&P Sustainable1.

[29] S&P Sustainable1.

[30] Basel Committee on Banking Supervision. (2006). *International convergence of capital measurement and capital standards: A revised framework, comprehensive version*. Bank for International Settlements.

[31] European Banking Authority. (2022). *Final report: Guidelines on common procedures and methodologies for the supervisory review and evaluation process (SREP) and supervisory stress testing under Directive 2013/36/EU* (EBA/GL/2022/03).

[32] Lorenz, M. O. (1905). Methods of measuring the concentration of wealth. *Publications of the American Statistical Association, 9*(70), 209–219. https://doi.org/10.2307/2276207

[33] Gini, C. (1921). Measurement of inequality of incomes available for purchase. *The Economic Journal, 31*(121), 124–125. https://doi.org/10.2307/2223319

[34] Ávila, F., Flores, E., López-Gallo, F., & Márquez, J. (2013). Concentration indicators: Assessing the gap between aggregate and detailed data. In Bank for International Settlements (Ed.), *Proceedings of the Sixth IFC Conference on "Statistical issues and activities in a changing environment"* (IFC Bulletin No. 36, pp. 542–559). Bank for International Settlements.