The Shape of Things to Come

Probabilistic emissions forecasting for climate risk

Dr. Ben McNeil
Dr. Nicholas Pittman

EMMI

12 DECEMBER 2025

Executive Summary

Institutional climate risk analysis depends critically on how future emissions are represented. Most current approaches rely on static, business-as-usual baselines, despite clear evidence that corporate emissions are dynamic and often follow persistent upward or downward trajectories.

Our analysis shows that treating emissions as flat systematically distorts forward-looking climate risk metrics, overstating exposure for companies already reducing emissions and understating it for those on an upward path. In practice, emissions follow recognisable patterns, and extending flat emissions into the future ignores observable behavior that contains predictive information.

Forward-looking emissions are therefore essential inputs to credible climate risk analysis, particularly for financial institutions assessing portfolio-level exposure.

In Emmi's approach to emissions forecasting, we introduce a Monte Carlo Momentum (MCM) framework that uses recent emissions behavior to generate forward-looking distributions for each company and Scope. Rather than a single deterministic path, the model simulates thousands of plausible futures and summarizes them with a median forecast and an explicit uncertainty range, consistent with probabilistic emissions approaches in the academic literature (Srikrishnan et al., 2022). Forecasts trajectories that imply implausibly large increases by 2030 are removed by excluding cases where the forecast standard deviation exceeds three times the historical mean.

This provides a more realistic business-as-usual baseline for forward-looking climate risk analysis, particularly for companies without explicit Science-Based Targets, and supplies a reference trajectory against which target-aligned pathways can be compared. By anchoring forecasts in observed corporate behaviour while explicitly recognising uncertainty, the framework supports more credible scenario analysis, portfolio risk assessment and engagement.

The key design choices are:

- Only data from 2020 onwards is used.
- At least three years of valid reported emissions are required per entity-scope.
- The default forecast horizon extends to 2030.
- Each entity-scope is simulated with 1,000 Monte Carlo paths.
- Forecasts are rejected if the 2030 median exceeds the historical mean plus three times the historical standard deviation.

In universes such as the S&P 500, the majority of companies display sufficiently stable emissions behavior to support useful forecasts, while a smaller group remain too volatile or too thinly disclosed to forecast reliably. This is precisely why strict quality controls are applied. Acknowledging this distinction is essential to interpreting momentum results and understanding where forward-looking estimates genuinely add value. For investors, this matters because static emissions baselines can distort assessments of future climate-related exposure, while forward-looking forecasts provide a more accurate basis for portfolio risk assessment and engagement.

This approach aligns with growing supervisory expectations for forward-looking, uncertainty-aware emissions analysis, as reflected in recent NGFS and TCFD guidance.

Contents

Static Emissions in a Dynamic World

Most transition risk analysis assumes that corporate emissions remain unchanged into the future. This is simple, but it overlooks the fact that emissions rarely remain stationary. Some companies have been reducing emissions for years, while others have been steadily increasing them, and treating all companies as flat disregards information contained in observed emissions behavior, an issue that has also been examined using Monte-Carlo analysis in national inventory studies (Ramírez et al., 2008). Yet emissions trajectories are not always clear or consistent. Some companies show relatively stable year-to-year changes and can be projected with confidence, while others are highly volatile or unevenly disclosed, making extrapolation unreliable. A single static assumption masks these differences and can distort future risk estimates.

Two practical questions therefore arise:

1. what happens if companies simply continue along their recent trajectory, and
2. how should we treat companies whose emissions are too volatile to model usefully?

Without addressing these questions, forward-looking analysis risks overstating exposure for companies that are already moving in the right direction and understating it for those on an upward path.

Incorporating recent behavior into forecasting allows analysts to move from a present-day emissions footprint to a more realistic view of future pathways, improving the credibility of scenario analysis, valuation and portfolio risk assessment. Treating all of these companies as flat through time has two consequences. First, it can misstate future risk. A firm with a long record of falling emissions may look as risky as one with rising emissions if both are assumed to remain constant. Second, it hides uncertainty. A single point estimate for future emissions masks the fact that some firms' trajectories are reasonably predictable while others are inherently volatile.

A Monte-Carlo Momentum (MCM) approach tackles both issues. It asks a simple question: given how a company's emissions have behaved in recent years, what are the plausible futures consistent with that behavior? The model estimates a central tendency (the drift in emissions) and the typical volatility around that drift, and then simulates many possible futures consistent with those parameters. The result is not a single forecast line, but a distribution: a median path and an explicit range of uncertainty that can be carried through into climate risk analysis, consistent with stochastic emissions approaches in recent studies (Jamatutu et al., 2024).

Data Foundation

Source data and filtering

The MCM forecasts rely on the same emissions and Partnership for Carbon Accounting Financials score (PCAF) information used across Emmi climate analytics. We use annual emissions by Scope, expressed in tonnes of CO_2e, and the associated PCAF score that reflects the quality of each observation. To ensure that forecasts reflect recent business conditions, we limit the historical window to emissions data reported from 2020 onwards. Reporting practices have improved substantially in recent years, and older data often reflects different operational boundaries or accounting methods that can distort trend estimation.

A minimum of three valid post-2020 emissions reports are required for each entity and Scope. This is the smallest sample that can meaningfully support both a directional estimate and a volatility estimate. Missing values are removed, non-positive values are excluded, and the resulting series must remain chronologically ordered. If fewer than three valid years remain after cleaning, the entity-scope is not forecast because it is pragmatic to acknowledge the lack of usable history rather than to force an extrapolation that is inherently unreliable.

Scope-level forecasting

MCM is estimated separately for Scope 1, Scope 2 and Scope 3 emissions. These Scopes often follow different operational drivers, rely on different reporting sources and exhibit different levels of volatility. Treating them independently preserves these distinctions and aligns the forecasts with how Science-Based Targets are handled elsewhere in Emmi analytics. Scope-level modelling also allows the uncertainty of each Scope to be carried through to downstream risk measures rather than blending them into a single number.

PCAF attribution

Every forecast carries through the maximum PCAF score observed in the historical input for that entity-scope. This does not change the simulation, but it provides context for interpreting the results. An MCM forecast based entirely on PCAF 1 or 2 reported values naturally carries a higher degree of confidence than one that draws on weaker estimates. Providing this information alongside the forecast ensures that users can weigh model uncertainty and data quality together when incorporating results into investment or risk decisions.

From Historical Momentum to Probable Futures

Estimating momentum

The forecasting process begins by converting annual emissions into year-over-year percentage changes. These changes capture both direction and variability and provide the basis for estimating how emissions have moved in recent years. To give greater influence to recent behavior, we apply a simple exponential time decay. With a half-life, the most recent observation carries full weight, and weights continue to diminish as we move back in time. The intention is straightforward: recent behavior is more informative about the near future than older data that may reflect different reporting practices or business configurations.

Occasionally, emissions change dramatically from one year to the next. This can occur because of acquisitions, divestitures, operational disruptions, or reporting revisions. Rather than remove these observations entirely, we down-weight changes greater than one hundred per cent in absolute terms so that they contribute to the estimate without dominating it. Keeping the data point preserves information, while the lower weight reduces the risk of a single event overwhelming the trend.

The result of this procedure is a weighted estimate of the average year-to-year change, which we refer to as momentum, and a corresponding estimate of historical volatility. Momentum can be positive or negative, indicating a tendency toward emissions growth or decline, while the volatility provides a measure of how much that change has varied historically.

Momentum is defined as an exponentially weighted average of year-over-year percentage changes, ensuring that recent emissions behavior has the greatest influence on the forecast.

Monte-Carlo simulation

Using these parameters, we simulate a large number of plausible futures, typically one thousand for each entity-scope, following Monte-Carlo methods widely applied in emissions forecasting (Jamatutu et al., 2024). Each simulated path applies random year-to-year changes drawn from a normal distribution defined by the estimated momentum and volatility. By stepping these random samples forward through time, the model generates a range of possible future emissions outcomes instead of one deterministic trajectory. The final year is summarized by its median value and the standard deviation across simulations, which together indicate the central estimate and the uncertainty surrounding it. An example of this distribution for Fortescue Metals Group and BlueScope Steel is shown in Figures 1 and 2.

Emissions are projected using **1,000 Monte Carlo simulations** per entity-scope, yielding a distribution of plausible future outcomes.

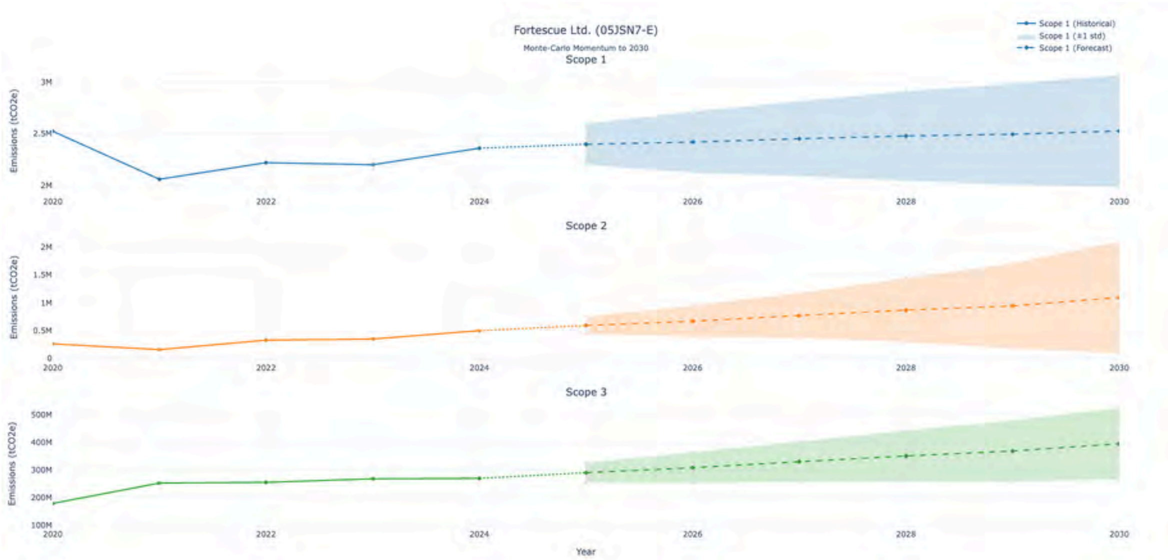

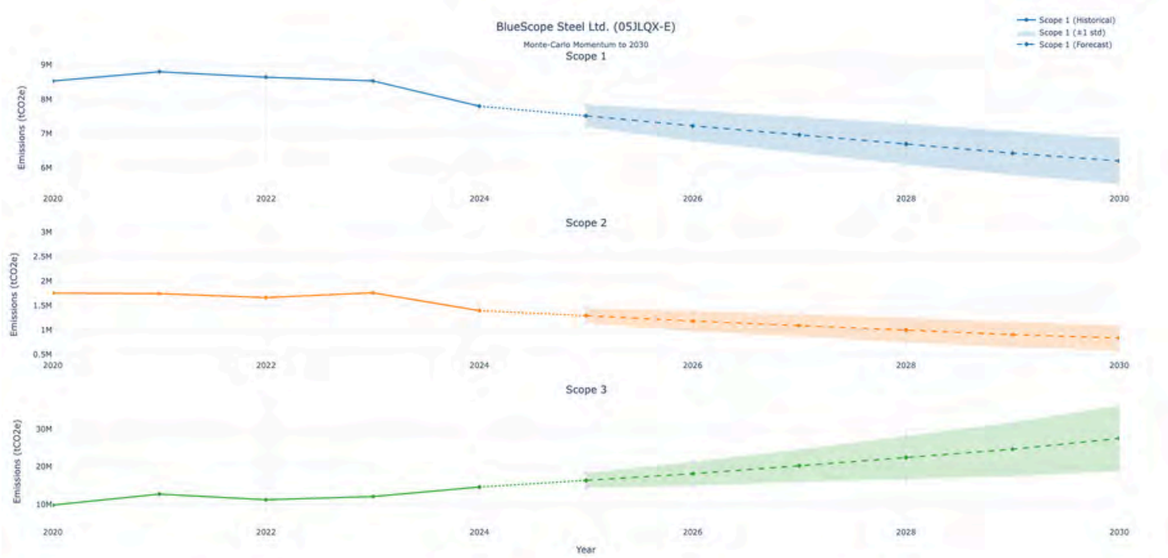

Forecast quality

To avoid physically unrealistic results, we place a lower boundary of zero on simulated emissions. Negative values are not meaningful, so any path that would fall below zero is set to zero. Because this adjustment is applied after evaluating uncertainty, it does not affect the measurement of historical volatility or the recognition of highly uncertain forecasts.

We exclude simulations where the forecast standard deviation in 2030 exceeds three times the historical mean. Extreme upward paths that are inconsistent with observed corporate behavior and would dominate the uncertainty range in a way that is unlikely to be useful decision-useful. By removing implausible outcomes, our estimates remain anchored in realistic business patterns.

Not all entity-scopes support meaningful forecasting. To avoid presenting highly uncertain results as though they were reliable, we filter forecasts using the coefficient of variation, defined as the ratio of the standard deviation to the absolute value of the median in the final forecast year. When this ratio exceeds two, the distribution is so wide that the central estimate is not decision-useful. In these cases, the forecast is excluded altogether rather than reported with misleading precision. This approach makes uncertainty visible: some emissions histories are sufficiently stable to project, while others are not, and forecasting should respect that difference.

3× Historical Mean Std Exclusion Rule: When the standard deviation across Monte Carlo estimates exceeds three times the historical mean emissions, uncertainty is so extreme that the forecast is deemed implausible and removed.

Outputs

For each entity-scope and forecast year, the model returns the median emissions level, the associated standard deviation, the number of simulations, and the maximum PCAF score in the historical input. These results can be linked directly into scenario analysis, temperature alignment and transition risk metrics, providing a dynamic business-as-usual baseline that reflects observed behavior and explicitly recognizes uncertainty.

Why Monte-Carlo Momentum?

Alternative methods evaluated

During the development of our forecasting framework, we conducted extensive research and backtesting to evaluate multiple forecasting methodologies. The goal was to identify the approach that best balances accuracy, interpretability, and operational scalability for corporate emissions forecasting.

The following methods were evaluated:

- **Monte-Carlo Momentum (MCM):** Stochastic simulation based on historical year-over-year changes with uncertainty quantification.
- **Feature-Based Random Forest:** Machine learning model incorporating financial features (revenue, total assets, employees) alongside historical emissions.
- **Gaussian Process Regression:** Non-parametric probabilistic model with custom kernels combining RBF, periodic, and noise components for flexible trend capture.
- **Simple LSTM-style Weighted Average:** Weighted moving average approach with exponential decay, inspired by recurrent neural network memory mechanisms.
- **VAR-Simplified:** Vector autoregression-inspired method using correlation-weighted adjustments based on financial variable changes.

Backtesting results

Each method was rigorously backtested using historical corporate emissions data. The evaluation framework used training periods of at least 3 years (e.g., 2017-2019 or 2021-2023) to predict emissions in a target year (e.g., 2022 or 2024), allowing comparison against actual reported values.

When using 3-year training periods, the more complex machine learning methods (Random Forest, Gaussian Process Regression) achieved marginally better accuracy metrics in some test configurations. However, these gains were modest and came with significant practical trade-offs in terms of complexity, data requirements, and interpretability.

Critically, when training periods were extended to 5 years, the MCM approach demonstrated the highest accuracy among all methods tested. This finding is particularly relevant as longer historical windows become increasingly available as corporate emissions disclosure matures across markets.

> "MCM's combination of simplicity, transparency, and competitive accuracy makes it the optimal choice for production-scale emissions forecasting.

Rationale for selecting MCM

The MCM approach was selected as our production methodology based on its optimal combination of practical advantages:

Advantage	Description
Simplicity	Requires only historical emissions data. No external features, no complex model training, no hyperparameter tuning is required.
Transparency	The methodology is straightforward to explain to stakeholders, regulators, and clients. Each forecast can be traced directly to observed emissions behavior.
Probabilistic Output	Provides explicit uncertainty ranges through simulation, not just point estimates. This uncertainty quantification is essential for risk management applications.
Data-Driven	Forecasts are grounded in actual company behavior rather than assumptions about relationships between emissions and financial metrics.
Computational Efficiency	Fast execution enables forecasting for thousands of companies across multiple Scopes in production environments.
Accuracy	Most importantly, MCM does not sacrifice accuracy for simplicity. With 5-year training data, it achieves the highest accuracy among all methods evaluated.

Conclusion on method selection

The marginal accuracy improvements offered by more complex methods do not justify their additional operational burden. Machine learning approaches require ongoing feature engineering, model retraining, and careful monitoring for concept drift. They also introduce 'black box' elements that can be difficult to explain when forecasts are scrutinized by stakeholders.

MCM's combination of simplicity, transparency, and competitive accuracy makes it the optimal choice for production-scale emissions forecasting. The approach is robust across different market conditions, requires minimal maintenance, and produces results that can be clearly communicated and defended. As emissions disclosures continue to improve and longer time-series become available, MCM's relative accuracy advantage is expected to strengthen over time.

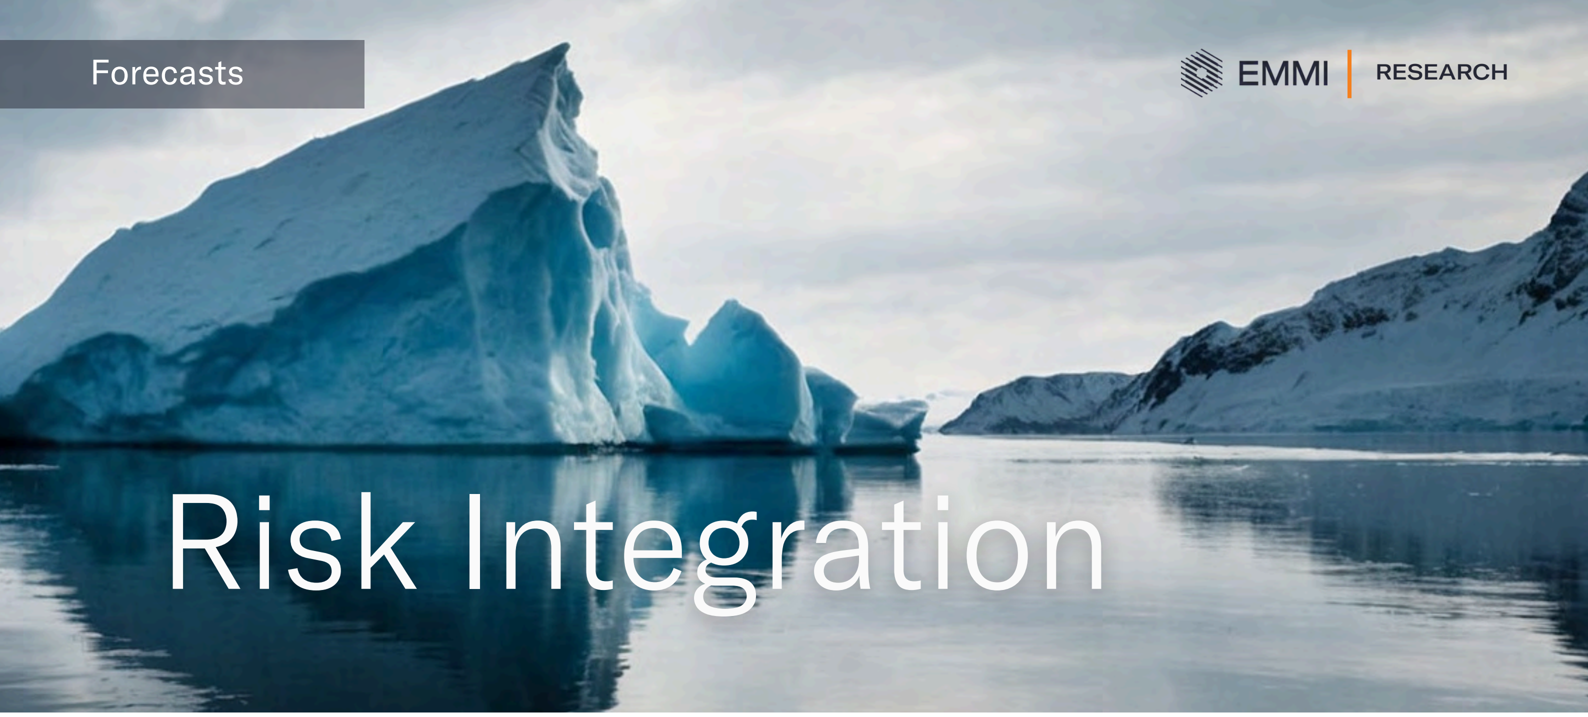

Risk Integration

MCM forecasts feed directly into Emmi's scenario-based transition risk framework. Rather than apply climate pathways to a static emissions baseline, we apply them to a forecast that reflects how each company has actually behaved in recent years. The result is a dynamic view of future emissions that changes how several core metrics are assessed:

Transition Value at Risk

Transition Value at Risk (TVaR) is sensitive to the level of emissions assumed in future years. When a company has been reducing emissions, the MCM forecast often sits below today's levels, which can reduce estimated future liability relative to a static baseline. Conversely, when emissions have been rising, MCM typically projects a higher future level than simple carry-forward, which can increase expected liability and bring attention to companies with growing exposure. In both cases, the risk calculation becomes more closely aligned with observable business dynamics and behavior.

Emissions reduction requirements

Under a given climate scenario, required reductions can be assessed against the MCM pathway rather than a flat line. A company already trending downward may have a smaller additional reduction to achieve, while an upward trend can imply a larger and more urgent transition requirement. This provides a clearer picture of where companies are already moving toward lower emissions and where significant structural change would be necessary under more ambitious climate pathways.

Temperature alignment

Temperature alignment metrics traditionally compare current emissions to scenario requirements. Using a dynamic baseline allows implied temperatures to reflect the likely shape of emissions through the forecast horizon. The result is often a better representation of whether a company is drifting toward or away from alignment, even before considering any explicit decarbonisation commitments.

Baseline comparison

All MCM-based metrics are calculated alongside the standard Emmi baseline, which assumes emissions remain constant. Comparing the two views makes it possible to understand how much of a company's future risk is driven by its recent emissions behavior. For companies with Science-Based Targets, the comparison also reveals how current behavior differs from stated commitments and whether targets imply a meaningful departure from recent trends. This side-by-side view gives investors a clearer sense of how behavior, commitments and policy interact in shaping future exposure.

Sensitivity and Validation

The MCM framework has been validated through extensive backtesting across multiple time periods and market universes. Key validation exercises include:

- Out-of-sample testing using historical training windows to predict known outcomes
- Comparison against naive baselines (flat emissions, simple trend extrapolation)
- Sensitivity analysis on key parameters including decay half-life and outlier thresholds
- Cross-validation across different market segments and geographic regions
- Evaluation of forecast stability across successive data releases

Results demonstrate that MCM consistently outperforms static baselines while maintaining appropriate humility about inherent uncertainty in emissions forecasting.

> “… MCM provides decision-useful forecasts across all three Scopes while appropriately reflecting the differing characteristics and data availability of each emission category.”

Forecast uncertainty and accuracy

Understanding the accuracy and uncertainty characteristics of MCM forecasts is essential for appropriate use in risk analysis. We report three key metrics from our backtesting framework:

- R^2 (Coefficient of Determination): Measures the proportion of variance in actual emissions explained by the forecast. Values closer to 1.0 indicate stronger predictive power.
- MDAPE (Median Absolute Percentage Error): The median percentage deviation between forecast and actual values. Using the median rather than the mean provides robustness against outliers and gives a realistic view of typical forecast accuracy.
- MAE (Mean Absolute Error): The average absolute deviation in tonnes CO_2e, providing a scale-aware measure of forecast error.

The table below summarizes backtesting results for the MCM approach across all three emission Scopes:

Scope	R^2	MDAPE (%)	MAE (tonnes CO_2e)
Scope 1	0.919	22.87	490,732
Scope 2	0.862	34.04	131,662
Scope 3	0.853	31.67	5,439,266

These results confirm that MCM provides decision-useful forecasts across all three Scopes while appropriately reflecting the differing characteristics and data availability of each emission category. Users should interpret Scope 1 forecasts with higher confidence, while recognising that Scope 2 and Scope 3 forecasts carry greater inherent uncertainty that is explicitly captured in the simulation distributions.

Forecast accuracy by company size

An important finding from our backtesting is that forecast accuracy improves systematically with company size. Larger companies tend to have more stable emissions profiles, more consistent reporting practices, and fewer disruptive corporate actions relative to their overall emissions base. This translates directly into lower forecast uncertainty.

For Scope 1 emissions:

- Microcap and Smallcap companies exhibit MDAPE values of 26–40%, reflecting higher emissions volatility and less predictable year-to-year changes typical of smaller enterprises.
- Large and Megacap companies achieve substantially lower MDAPE values of 7–20%, indicating that forecasts for the largest emitters are considerably more reliable.

This pattern is intuitive: larger companies typically have more diversified operations, smoother emissions trajectories, and more mature reporting systems. Their emissions are less likely to experience the proportionally large swings that can affect smaller companies due to single facility changes, acquisitions, or operational disruptions.

Implications for users:

This size-based accuracy gradient has practical implications for portfolio-level analysis. Forecasts for large and mega-cap holdings, which often represent the majority of financed emissions in institutional portfolios, can be interpreted with higher confidence. For smaller holdings, the wider uncertainty bands produced by MCM appropriately reflect the greater difficulty in forecasting their emissions trajectories.

Users should note that while Scope 2 shows less consistent improvement across size segments (likely due to the influence of external factors such as grid mix changes and renewable energy procurement decisions), the overall pattern confirms that MCM forecasts are most reliable precisely where they matter most: for the large emitters that drive portfolio-level climate risk.

Limitations and Assumptions

MCM forecasting rests on the assumption that recent emissions behavior provides useful information about the near future. For many companies, this is reasonable, but sudden structural shifts such as acquisitions, divestitures or rapid electrification can make recent history a poor guide to what comes next. In such cases, MCM may lag reality until new data becomes available.

Annual percentage changes are modelled as normally distributed. Real emissions shocks can be asymmetric or extreme, and although outlier down-weighting and quality filters reduce the influence of unusual years, they cannot eliminate all extremes. MCM should therefore be viewed as a pragmatic approximation rather than a complete representation of uncertainty.

Data availability also shapes results. Requiring at least three valid post-2020 observations deliberately excludes entities with sparse disclosure, and forecasting is avoided where uncertainty would be misleadingly high. Similarly, Scopes are treated independently and do not capture potential shifts between them, such as electrification moving emissions from Scope 1 to Scope 2.

MCM does not replace explicit decarbonisation targets. Instead, it provides a behavioral baseline: a view of what the future might look like if recent patterns continue. When targets exist, MCM acts as a counterfactual against which ambition and potential risk reduction can be assessed.

"MCM does not replace explicit decarbonisation targets. Instead, it provides a behavioral baseline: a view of what the future might look like if recent patterns continue."

Applications

The primary purpose of MCM forecasting is to create a realistic business-as-usual view of future emissions, especially for companies without explicit decarbonisation targets. Many firms are already drifting upward or downward, and reflecting that behavior improves the credibility of transition risk estimates compared with a flat baseline.

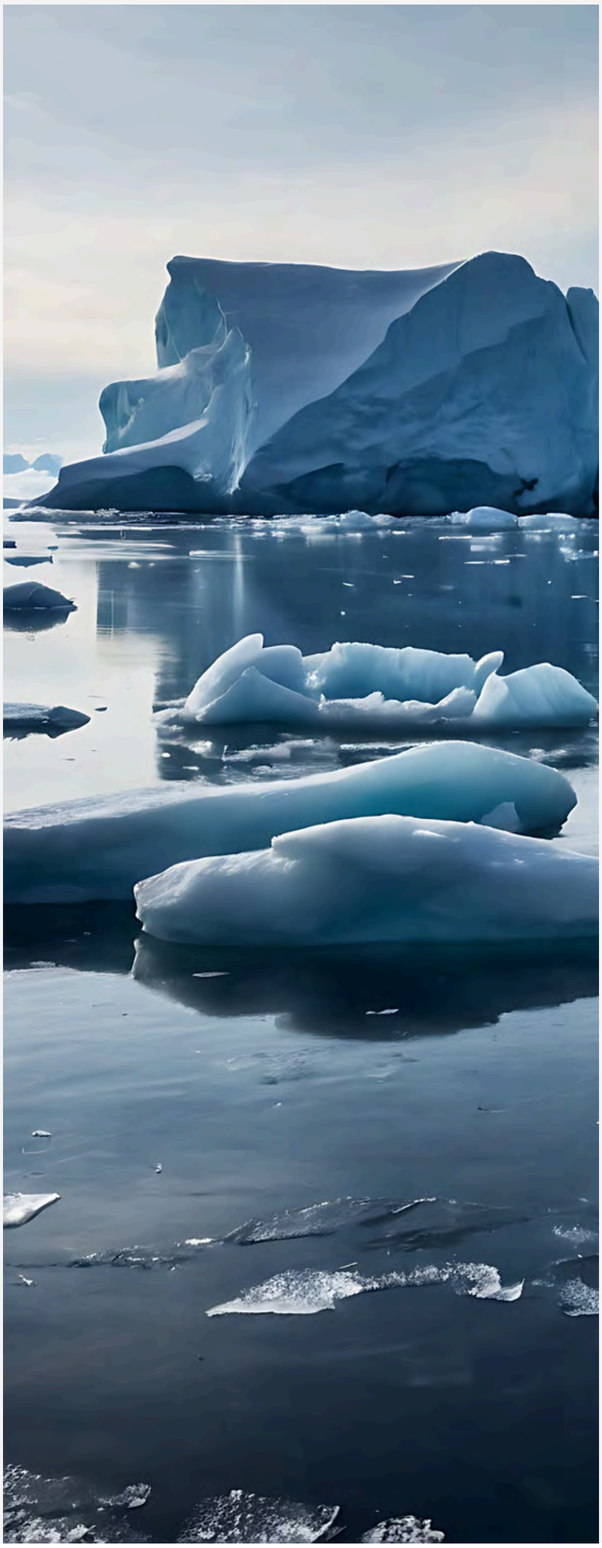

MCM also makes uncertainty visible. Some companies have relatively stable trajectories, while others remain highly volatile. Recognising this difference helps investors judge where forward-looking results are informative and where they are inherently uncertain, which can guide portfolio construction and engagement priorities.

Where targets exist, MCM forecasting provides a complementary view. It shows how far recent behavior is from the company's stated commitments and whether achieving those targets would require a significant shift from current trends. When used in climate scenarios and temperature alignment it supplies a distribution of plausible futures rather than a single starting point. Targets describe the future a company intends to reach, while MCM reflects the future implied by recent behavior, and comparing the two reveals whether current actions are consistent with stated ambition and where additional change may be required. Taken together, they provide a fuller picture of corporate decarbonisation and a more grounded basis for forward-looking climate analysis.

Data Access and Delivery

MCM forecasts are delivered through Emmi's standard distribution channels, including the FactSet Carbon Diagnostics Platform for interactive analysis and Snowflake for direct data access. Each forecast includes the median emissions estimate for the relevant years, the associated standard deviation, the number of simulations used, and the maximum PCAF score in the underlying historical data. These fields are provided at the entity and Scope level and can be linked directly to existing scenario and temperature alignment tools.

All outputs are versioned and tagged with methodology identifiers so that users can compare results over time and understand which parameters were in effect at each release. Because the forecasting engine is modular, parameters such as the decay half-life, outlier thresholds or volatility limits can be adjusted as evidence and user needs evolve. This flexibility allows MCM results to be interpreted alongside other forward-looking views, including target-based trajectories and scenario pathways, within a consistent analytical environment.

Each forecast includes the median emissions estimate for the relevant years, the associated standard deviation, the number of simulations used, and the maximum PCAF score in the underlying historical data.

Conclusion

Corporate emissions change over time, and forward-looking climate analysis needs a baseline that recognizes this. A Monte-Carlo Momentum (MCM) approach offers a simple and transparent way to reflect recent behavior and to quantify uncertainty, rather than assume emissions remain flat. By filtering out highly uncertain cases and constraining implausible outcomes, the method focuses on companies where forecasting is genuinely informative.

MCM is not a prediction in the strict sense, nor is it a substitute for explicit decarbonisation commitments. It is a behavioral view that shows where emissions appear to be heading if current patterns continue. When combined with Science-Based Targets and climate scenarios, it enables a richer understanding of whether companies are moving in the right direction, how ambitious their stated goals are, and where transition risk may be concentrated.

Our extensive research into alternative forecasting methodologies including machine learning approaches such as Random Forest and Gaussian Process Regression, confirmed that the additional complexity of these methods does not yield meaningful accuracy improvements over MCM, particularly when sufficient historical data is available. The MCM approach delivers the optimal balance of simplicity, transparency, and accuracy required for production-scale climate analytics.

In markets such as the S&P 500, the majority of companies display sufficiently stable emissions behavior to support useful forecasts, while a smaller group remain too volatile or too thinly disclosed to forecast reliably. This is precisely why strict quality controls are applied. Acknowledging this distinction is essential to interpreting MCM results and understanding where forward-looking estimates genuinely add value.

As disclosures lengthen and data quality improves, the robustness and explanatory power of MCM forecasts are expected to strengthen further.

References

Jamatutu, S. A., Abbass, K., Gawusu, S., Yeboah, K. E., Jamatutu, I. A. M., & Song, H. 2024. Quantifying future carbon emissions uncertainties under stochastic modeling and Monte Carlo simulation: Insights for environmental policy consideration for the Belt and Road Initiative Region. Journal of Environmental Management, 370, 122463. https://doi.org/10.1016/j.jenvman.2024.122463

Network for Greening the Financial System. 2023. NGFS Scenarios for Central Banks and Supervisors. NGFS. https://www.ngfs.net/ngfs-scenarios-portal/

Partnership for Carbon Accounting Financials. 2022. The Global GHG Accounting and Reporting Standard for the Financial Industry. PCAF. https://carbonaccountingfinancials.com/files/downloads/PCAF-Global-GHG-Standard.pdf

Ramírez, A., de Keizer, C., Van der Sluijs, J. P., Olivier, J., & Brandes, L. 2008. Monte Carlo analysis of uncertainties in the Netherlands greenhouse gas emission inventory for 1990–2004. Atmospheric Environment, 42(35), 8263–8272. https://doi.org/10.1016/j.atmosenv.2008.07.059

Srikrishnan, V., Guan, Y., Tol, R. S. J., & Keller, K. 2022. Probabilistic projections of baseline twenty-first century CO_2 emissions using a simple calibrated integrated assessment model. Climatic Change, 170(3), 37. https://doi.org/10.1007/s10584-021-03279-7

Task Force on Climate-related Financial Disclosures. 2021. Guidance on Metrics, Targets, and Transition Plans. TCFD. https://www.fsb.org/uploads/P141021-2.pdf

Climate risk, built for investors

Emmi provides comprehensive climate risk intelligence for investors. Our datasets cover emissions, transition and physical risk across public and private markets, covering all major asset classes and driving 100% portfolio coverage.

Built on a consistent methodology, Emmi delivers the transparency and customisation investors need to make better investment decisions, meet climate disclosure requirements, and align with regulatory and mandate expectations.

Emmi is founded on a simple idea: mobilising capital is the fastest path to decarbonisation. By quantifying climate risk exposure at scale, we enable the financial sector to allocate capital more efficiently toward climate-aligned outcomes.

To meet this need, we built Carbon Diagnostics - decision-useful climate insights, delivered at scale.

Authors

Dr Ben McNeil
Climate Scientist and Co-founder, Emmi
Adjunct Senior Fellow, Climate Change Research Centre, University of NSW, Sydney, Australia

Dr Nicholas Pittman
Machine Learning Engineer & Data Scientist
Doctor of Philosophy (PhD), Institute of Marine and Antarctic Studies, University of Tasmania

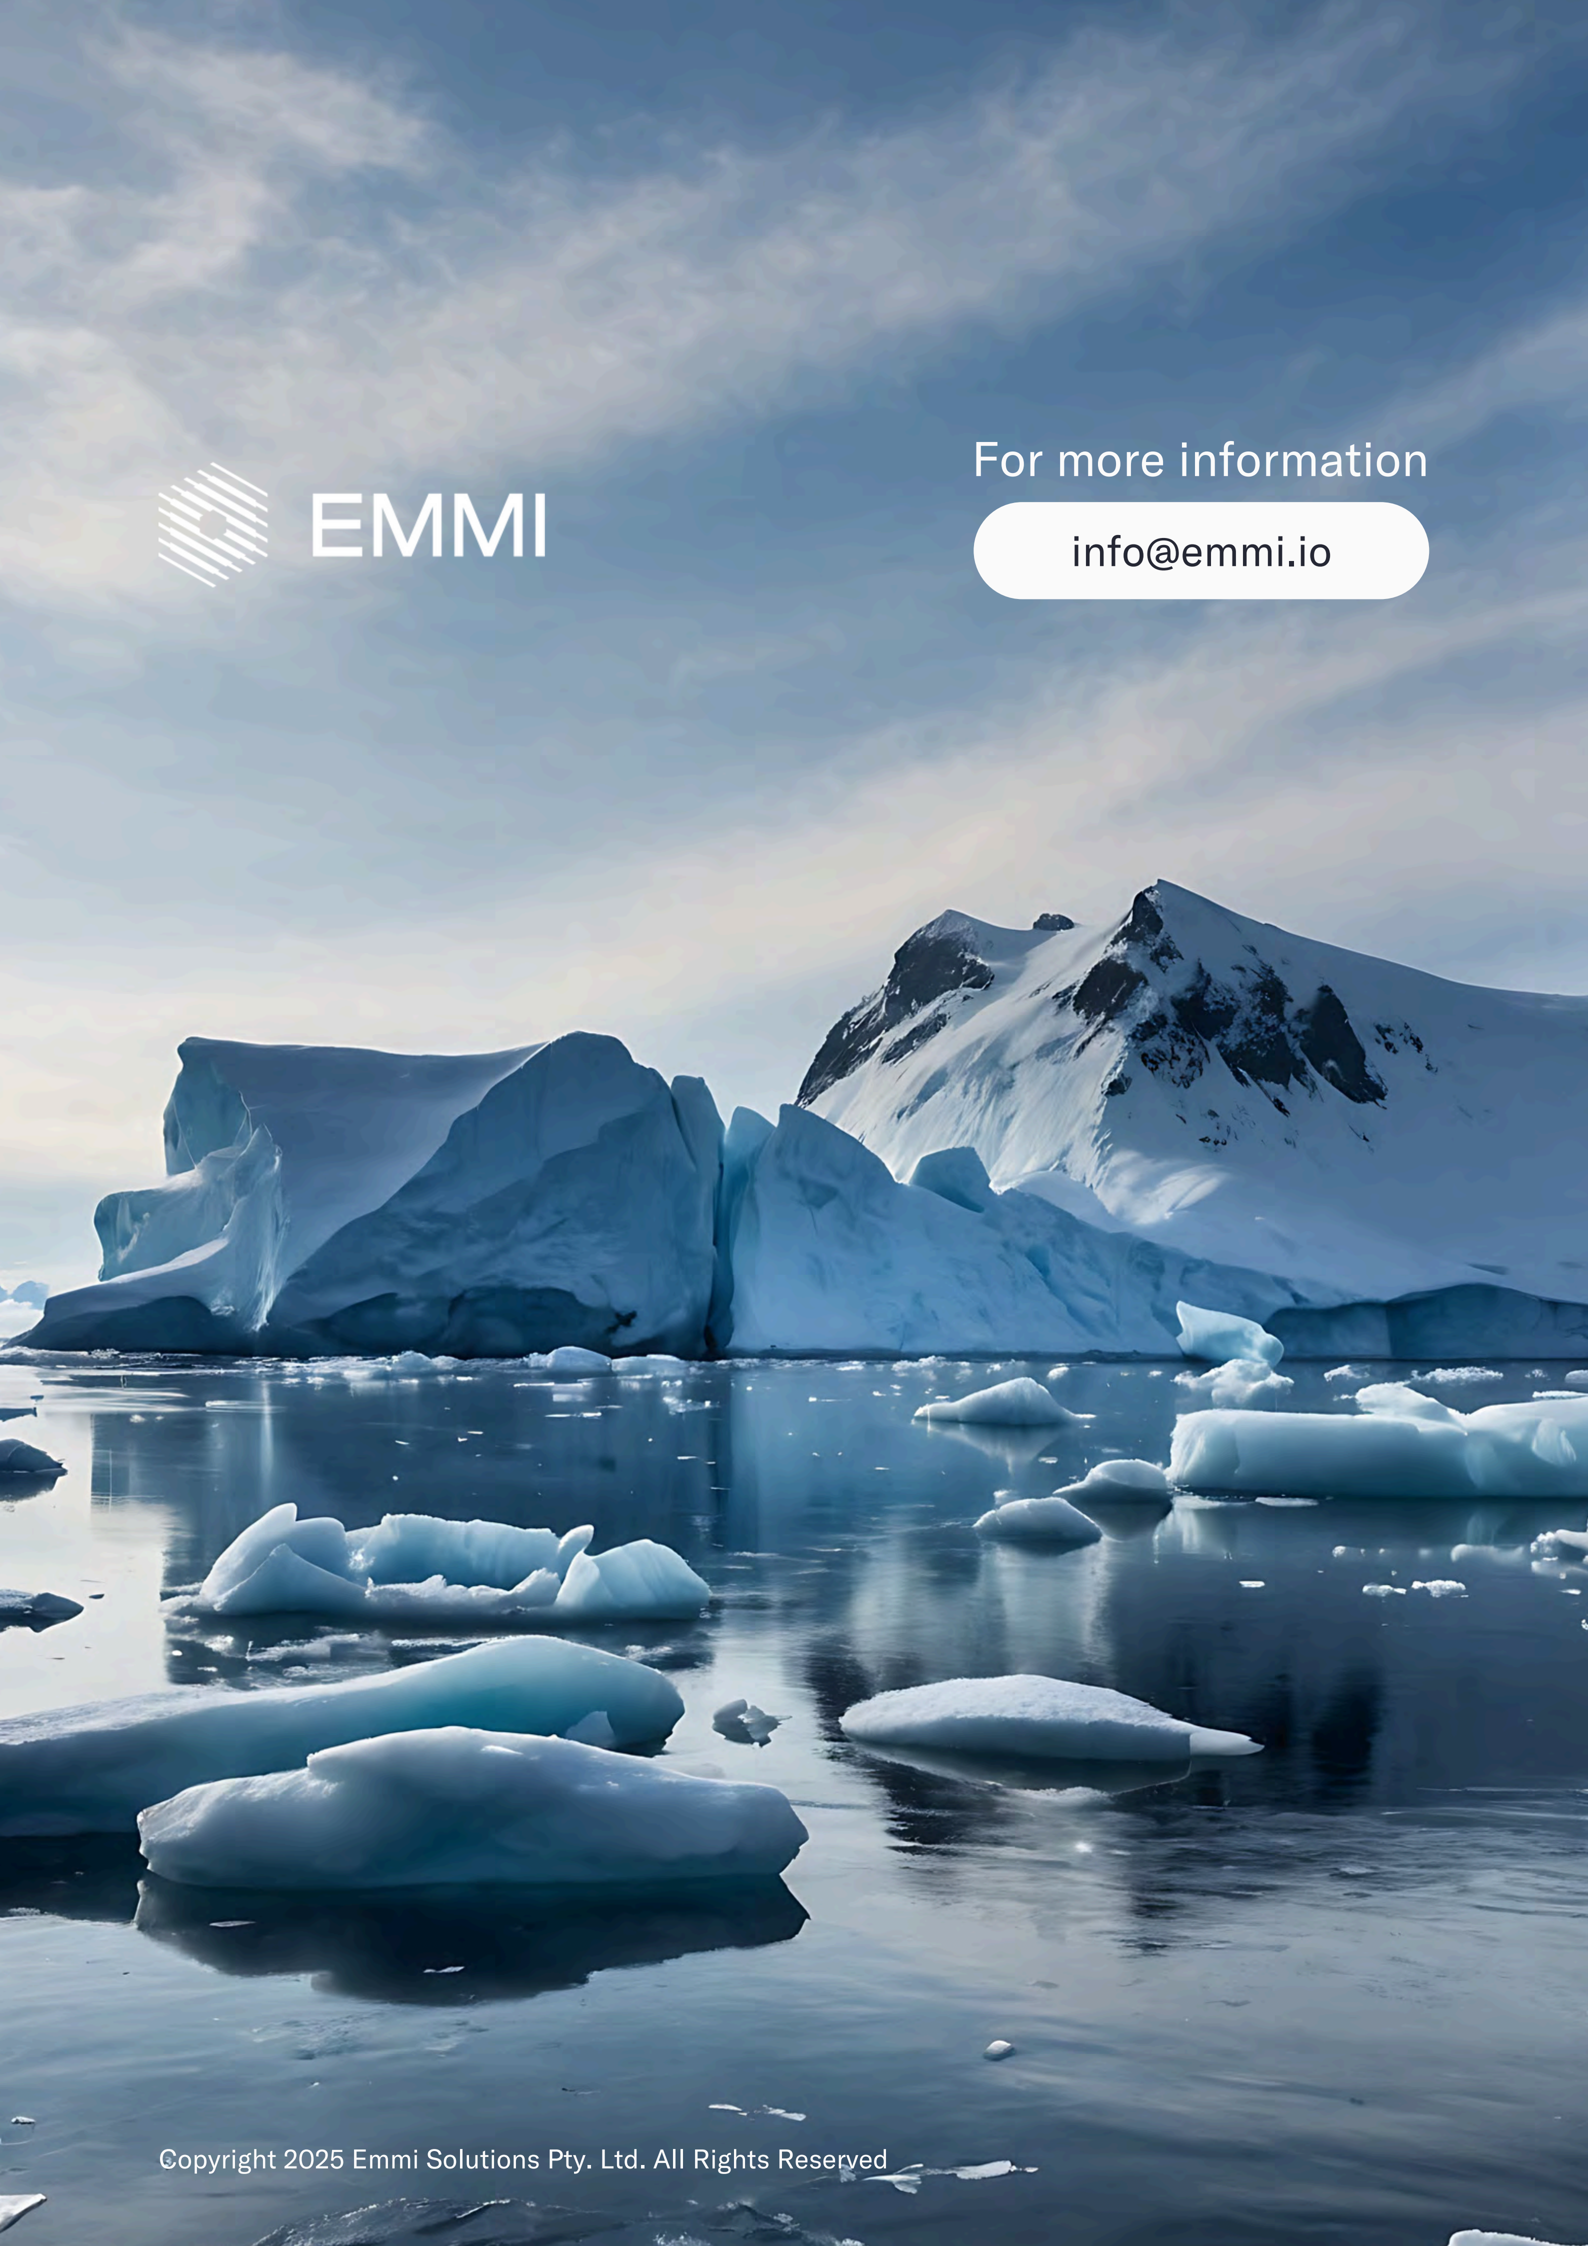

EMMI

For more information

info@emmi.io