

The Repricing of Compute

The Economics of Climate Risk and Resilience
for AI-Era Capital

Rémi Paccou & Thomas Epelbaum

June 2026

Foreword

Steve Wilhite

Executive Vice President, Energy & Sustainability Practice

SE ADVISORY
SERVICES



What is the climate price tag on the world's compute infrastructure?

The data center industry stands at an inflection point. Over the past three decades, we have built the physical layer of the digital economy—facility by facility, market by market—into an infrastructure class that now underpins every sector of the global economy. The scale of what comes next dwarfs what came before: the AI era demands not incremental expansion but a structural transformation of how we design, power, and operate these critical assets.

At Schneider Electric, we work alongside operators, investors, and grid planners who are making capital allocation decisions today that will determine the resilience of digital infrastructure through 2050. The question they face is not whether climate risk affects data centers—every operator who has managed through a winter storm, a heat-wave, or a grid emergency knows the answer. The question is how much, where, and what to do about it.

This paper provides the first quantitative answer. By translating physical and transition climate hazards into enterprise value loss at facility level—across the global installed base and the prospective AI pipeline—it makes visible a dimension of risk that capital markets have not yet priced. The supply chain analysis is particularly consequential: the finding that indirect business interruption exceeds direct exposure by an order of magnitude reframes how we think about facility resilience.

What gives this work practical weight is that it does not stop at diagnosis. The adaptation analysis demonstrates that structured investment reduces climate exposure by a third, with positive returns under every scenario tested. For an industry deploying trillions in long-lived assets, that is an actionable result.

The analytical framework developed here—and the Schneider Electric Research Institute's commitment to open, reproducible methodology—provides the foundation for integrating climate risk into the investment decisions that will shape our industry for the next quarter century.



Perspectives

Stéphane Voisin

Director of Sustainable Finance Research,
ILB & Secretary General, Fondation PARC

When climate enters the valuation equation

Climate-related risk is rapidly integrating into the financial landscape, shifting from an abstract concern to a tangible factor embedded in the design of risk models across institutions. The tools used to capture this very specific nature of environmental risk must evolve accordingly. Early initiatives—such as the Dutch Central Bank's *Underwater* study—marked an important step by introducing Discounted Cash Flow (DCF) analysis into climate stress testing. Yet these approaches largely treated climate shocks as external inputs applied onto otherwise unchanged financial frameworks. The present work contributes to advancing a more structural approach: a rethinking of valuation itself, where climate is no longer a scenario layered on top, but a variable within the system.

Infrastructure is certainly the asset class where this shift is most needed. Data centres, the silent engines of AI, are scaling at unprecedented speed, while remaining deeply exposed to physical climate constraints: heat, water stress, grid instability. Their long lifespans collide with uncertain climate trajectories, and yet their valuation still assumes a form of environmental neutrality. This paper brings that contradiction into focus—and, for the first time, prices it.

By embedding climate dynamics directly into valuation through a unified framework, it moves beyond stress testing to measure how risk propagates through costs, operations and asset value itself. It offers a bridge between climate science and financial economics, turning Climate Value at Risk from a theoretical construct into a metric that operators, investors, and insurers can act on today.

The question this work raises is no longer whether climate belongs in the valuation equation. It does. The question is how quickly the institutions that allocate capital will build that reality into their models—and how much value will have been destroyed in the meantime.

About the Authors



Rémi Paccou

*Schneider Electric Research Institute
Doctoral Researcher, CIRED*

Rémi Paccou is Director of Research at the Schneider Electric Research Institute and a doctoral candidate at CIRED, one of France's leading environmental economics laboratories. His research focuses on the macroeconomic and climate implications of digitalization, combining scenario modeling, life-cycle assessment, and energy economics to inform both corporate strategy and public policy. He leads the Institute's work on digitalization and energy systems, with publications covering AI electricity demand trajectories, net impact assessment of AI-powered energy management, and the thermodynamic boundaries of computational intelligence. Prior to his research career, Rémi spent over twenty years at Schneider Electric in operational and strategic leadership roles. He holds an engineering degree with a specialization in thermodynamics from Arts et Métiers ParisTech and executive education from UNC Kenan-Flagler and INSEAD.



Thomas Epelbaum

*SE Advisory Services
PhD, CEA Saclay (IPhT)*

Thomas Epelbaum is Head of Python Developments and Machine Learning at SE Advisory Services, Schneider Electric's global consulting practice (formerly EcoAct). He built from scratch the geospatial and computational backend powering the platform's physical and transition climate risk assessments—integrating hazard projections, geolocation algorithms, and machine learning into production-grade software serving infrastructure portfolios across energy, real estate, and digital sectors. Before joining EcoAct, Thomas held data science and ML roles at Shift Technology, designing research libraries, NLP systems, and ML monitoring tools. He holds a PhD in theoretical physics from CEA Saclay (IPhT), where he studied the thermalization of the Quark-Gluon Plasma, and was trained at École Polytechnique and École Normale Supérieure.

Executive Summary

Rémi Paccou & Thomas Epelbaum

Recognized, Not Yet Priced

Climate risk has become a financial variable for every physical asset—factories, ports, power grids, real estate, transport networks. Yet across asset classes it remains recognized but not priced: markets can name the hazard without quantifying what it costs, where it concentrates, or whether it can be reduced. This paper brings clarity to that gap. We focus on data centers: as the fastest-growing asset class of the AI era, they are a natural place to start—but the framework, and the problem it addresses, apply to any physical asset.

Data centers make the case concrete. For two decades this infrastructure scaled predictably, largely insulated from physical climate variability. Since 2022, AI workloads have changed the picture: energy systems face demand they were not designed to absorb, capital is being redirected at historic scale, and assets with 15–25 year horizons are being committed without pricing the climate trajectory they will operate under. The design choices made in the next five years will set their exposure for the following quarter century.

The financial signal is already visible. In 2025, global natural-catastrophe losses reached \$296 billion, of which insurers covered \$129 billion—leaving a protection gap of \$167 billion that widens 6 to 10% each year. Data center insurance costs already exceed \$1 billion annually and are projected to triple by mid-century. For a sector that requires trillions in new investment by 2030, this is less a warning than a call for better information: capital cannot be priced efficiently against a risk no one has measured.

The difficulty is straightforward. Operators, investors, and insurers lack the metrics to distinguish a resilient facility from an exposed one, to quantify how much value each hazard erodes, or to test whether adaptation generates a positive return. Without facility-level valuation, capital is allocated as if climate exposure were uniform—or negligible. The purpose of this paper is to replace that blind spot with clarity.

This paper brings that clarity.

No study currently quantifies how much climate risk the global data center sector carries, where it concentrates, through which channels it operates, and how it can be reduced. We measure this risk as *enterprise value loss at facility level* across the global portfolio, combining three novelties.

A *climate-adjusted valuation model* translates physical hazards and transition risks into discounted cash flow impacts across five economic channels—cooling, carbon, business interruption, physical damage, heat productivity—producing auditable, facility-level estimates of enterprise value loss. The model is deliberately a white box: every parameter, physical and financial, is transparent and can be customized to an operator's or an investor's own assumptions. Shapley decomposition attributes each channel's contribution to specific hazards; Leontief input-output analysis propagates business interruption through upstream supply chains. Neither tool has previously been applied to this sector.

A *geocoded database* of 8,572 installed facilities (90.8 GW across 120 countries) is augmented by prospective AI-era infrastructure under three growth scenarios—Limits to Growth (+63 GW), Sustainable AI (+121 GW), Abundance Without Boundaries (+296 GW)—covering both the existing stock and the infrastructure whose climate exposure remains a design choice.

A *scenario architecture* crossing warming pathways, transition policies, AI growth trajectories, and adaptation strategies identifies which interventions generate positive risk-adjusted returns, and under which conditions. Calibrated here on data centers, the framework is portable to any physical asset class.

Four findings.

#1 Climate risk represents a third of current data center enterprise value—US\$388 billion on the installed base, scaling to US\$1.0–3.7 trillion with AI infrastructure. Standard valuations leave this figure invisible; measuring it is the first step to managing it.

#2 A twofold dispersion in risk intensity across regions, driven by fundamentally different channel compositions, means clarity must be site-specific—uniform mitigation strategies leave value on the table.

#3 AI-era infrastructure carries 2.6× more physical risk per GW than the installed base, deepening a physical-leaning profile: carbon costs fall from 28% to just 6% on AI-era assets.

#4 Proactive adaptation—PPAs, liquid cooling, microgrids, climate-informed siting—restores 30 to 39% of exposed value at a positive return under every scenario tested. A structural residual of 61% marks the boundary of what facility-level action alone can achieve, and where systemic resilience must take over.

Three case studies—in the US South, France, and southern China—show how the framework decomposes each site's risk into its economic channels and traces it through upstream supply chains, turning portfolio-level findings into site-specific engineering priorities.

Key messages.

1. Physical climate risk deserves the same analytical attention—and the same clarity—that the financial sector has spent a decade building for carbon. No equivalent toolkit yet exists for physical risk.

2. Pricing physical risk into infrastructure valuation requires facility-level granularity, channel-specific diagnostics, and supply chain visibility that aggregate portfolio approaches cannot provide.

3. Financial resilience is a system property, not a site property. Infrastructure does not fail in isolation—it fails with grids, supply chains, and interconnected services. Risk management that stops at the facility perimeter misses the dominant transmission channel.

4. Adaptation is most valuable when anticipated at the front end. The design choices being made today set climate exposure for the next quarter century, and building resilience in is far cheaper than retrofitting it later. Anticipation creates value; delay destroys it.

Key Insights

Climate Value at Risk — Global Data Center Infrastructure (BAU Pathway, 8,572 Facilities, 120 Countries)

US\$388B

Climate risk on the installed base

38% of US\$1,012B aggregate enterprise value

38%

Average climate risk intensity

Range: 28% (United States) to 56% (China)

APAC 54%


EU 33%

US 27%

×11

Supply chain amplification

For every \$1 of direct disruption at a facility, \$11 propagate through its suppliers



30–39%

Risk reduction through adaptation

Net positive returns in every scenario | US\$150B–1.1T protected

Baseline

Selective

Proactive

Source: ClimVaR model V8 results. DCF enterprise values under BAU climate-economic pathway, MEAN estimates.

Implications

Capital markets & financial modelling.

Physical climate risk lacks the anchor metrics that carbon intensity provides for transition risk—and that absence is an opening. Facility-level ClimVaR supplies the missing metric in a transparent, white-box format compatible with DCF models and disclosure frameworks. For the first time, investors can price physical exposure into capital allocation rather than around it, differentiate assets that look identical on paper, and identify mispriced opportunities before the market reprices them.

Operators, system planners & insurers.

The same dollar of adaptation investment generates 2–3× different returns depending on location and channel composition—which means the upside of allocating it well is large. Channel-level attribution identifies which lever matters most at each site, turning a generic resilience budget into a targeted, return-maximizing one. The ×11 supply chain multiplier, once visible, becomes a lever in its own right: it reveals where systemic resilience pays off well beyond the facility perimeter.

AI developers & hyperscalers.

The further concentration of risk into physical channels means that sustainability strategies centered on renew-

able procurement address less than 10% of AI-era climate exposure. The other 90% is where the advantage lies: physical resilience, designed in at the front end, is both the binding constraint and the lever competitors are not yet pulling. Build it early and it is cheap; the same protection retrofitted later costs several times more.

Policymakers & regulators.

Data centers have shifted from commercial real estate to strategic infrastructure, but their climate exposure remains weakly integrated into permitting and prudential oversight. A facility-level metric makes targeted policy possible: climate-adjusted siting rules and infrastructure stress tests calibrated to actual financial exposure—steering capital toward resilient buildout rather than constraining it across the board.

Chief sustainability officers.

Climate reporting that stops at carbon emissions captures only a quarter of the exposure. Bringing the full five-channel picture into disclosure—cooling, business interruption, and physical damage alongside carbon—turns a reporting obligation into a strategic instrument: a credible, quantified basis for resilience investment that boards and investors can act on.



Table of Contents

Foreword — Steve Wilhite	2
Perspectives — Stéphane Voisin	3
About the Authors	4
Executive Summary	5

Chapter I Introduction

When Digital Measures Its Own Risk	7
The Infrastructure at Stake	8
The Research Gap	10
Contributions	11

Chapter II Methods & Data

The ClimVaR Model	13
Why Data Centers Are Different	14
Shapley & Leontief	15
Database	16
Scenario Architecture	17

Chapter III Results & Discussion

The Transmission Mechanism	20
Finding 1: Unpriced Exposure	21
Finding 2: Geographic Dispersion	23
Finding 3: AI Amplification	25
Finding 4: The Adaptation Dividend	27
Three Facility Case Studies	30
Discussion	32

Chapter IV Conclusions

Key Messages	33
Implications by Stakeholder	34
Limitations	34
Future Research	35

Appendices

References	37
Terminology	39
A. ClimVaR Methodology	42
B. Data & Hypotheses	62
C. Scenario Architecture	72
D. Literature Review	75

List of Exhibits

Exhibit 1	Four Converging Phenomena Reshaping Data Center Risk	7
Exhibit 2	The Scissors: Capital In vs. Protection Out	8
Exhibit 3	Literature Coverage Against Nine Analytical Dimensions	10
Exhibit 4	Analytical Scope: Existing Frameworks Versus This Paper	11
Exhibit 5	ClimVaR Risk Channels: Notation, Mechanisms, and DC Calibration	14
Exhibit 6	ClimVaR Framework: Five Risk Channels to Enterprise Value	14
Exhibit 7	AI Growth Scenarios: Capacity Trajectories	17
Exhibit 7a	Facility Coverage by Region and AI Growth Scenario	17
Exhibit 8	Scenario Architecture: Four Dimensions of Uncertainty	18
Exhibit 9	Physical Climate Scenarios and DC Implications	18
Exhibit 10	Panoramic Results: ClimVaR Across Asset Classes and Scenarios	19
Exhibit 11	The Transmission Map: Hazards to Enterprise Value Loss	20
Exhibit 12	ClimVaR Decomposition by Risk Channel	21
Exhibit 13	US\$388B Decomposed: Four Regions, Five Channels	22
Exhibit 14	Regional ClimVaR Intensity and Adaptation Headroom	23
Exhibit 15	Four Fingerprints: Regional Channel Composition	24
Exhibit 16	Channel Decomposition: Installed vs. AI-Tier Facilities	25
Exhibit 17	ClimVaR Trajectories by AI Growth Scenario	26
Exhibit 18	Adaptation Value Creation by Region	27
Exhibit 19	The Adaptation Effect: \$150B in Net Protected Value	27
Exhibit 20	Channel-by-Channel Waterfall: Baseline to Proactive	29
Exhibit 21	Marginal Abatement Cost of Climate Risk	29
Exhibit 22	Three Prospective AI Facilities: ClimVaR Decomposition	30
Exhibit 23	Hazard Attribution (Shapley) and Propagation (Leontief)	31
Exhibit 24	Three Sites, Three Risk Profiles: Channel Fingerprints	31



Chapter I

Introduction

Data centers are the physical infrastructure of the digital economy—and the only major asset class whose climate exposure has never been financially quantified.

When Digital Measures Its Own Risk

In January 2026, a winter storm paralyzed the US East Coast's data center corridor—home to the world's largest concentration of digital infrastructure. Wholesale electricity prices surged from \$200 to \$1,800 per MWh overnight (Reuters, 2026). The PJM¹ grid approached 147 GW of demand, shattering its previous winter record (Connell, 2026). The Department of Energy issued emergency orders requiring data centers to prepare backup generation (Stapczynski, 2026).

This was not an isolated event. In the same winter, a major Atlantic storm knocked out power to over one million customers across Western Europe, disrupting pan-European cloud workloads (Gallagher Re, 2026). Months earlier, a major typhoon flooded a coastal province in southern China—one of Asia's densest concentrations of digital infrastructure—causing cascading grid failures that took 72 hours to resolve (Gallagher Re, 2026). In the Gulf, temperatures exceeding 50° C pushed cooling systems to their design limits. Each episode exposes a different failure mode—flooding, fuel gelling, cooling exceedance, cascading grid collapse. The common denominator is financial: physical climate hazards now impose direct costs on the

infrastructure that supports the digital economy (Forzieri et al., 2018).

These are not marginal assets. Data centers are capital-intensive facilities with 15–25 year operational horizons, whose performance depends on energy systems, physical site conditions, and insurance markets—all of which climate change is repricing simultaneously. Governments designate them as critical infrastructure. Capital markets direct trillions toward them. Between 40% and 60% of 2035 global capacity will be built in the next five years (IEA, 2025). Once commissioned, these facilities cannot relocate. Every siting decision made today embeds a quarter-century bet on climate trajectories.

And yet: the sector that hosts the computational infrastructure used for climate projections, that stores the data underpinning TCFD² and CSRD³ disclosures, that prices risk for every other industry—cannot price its own. No valuation framework quantifies how much climate risk a data center carries. No disclosure regime requires the number. No insurance product is calibrated to it. The risk is not unrecognized. It is unpriced.

¹PJM Interconnection: the regional transmission organization coordinating wholesale electricity in 13 US states and the District of Columbia.

²Task Force on Climate-related Financial Disclosures.

³Corporate Sustainability Reporting Directive (EU).

The Infrastructure at Stake

Data centers have transitioned from commercial real estate to strategic national infrastructure. The scale of capital committed to this transition has no precedent in the sector's history.

Compute as strategic infrastructure

In June 2025, the United States designated hyperscale data centers as Critical Digital Infrastructure—enabling expedited permitting and eligibility for federal risk-sharing programs (Tyron, 2025). Industry estimates place global economic activity directly dependent on data center operations at \$8.75 trillion by 2027 (Equinix, 2025)—a figure difficult to verify independently, but whose order of magnitude is consistent with the sector's role as transactional backbone of the digital economy. Private investment in US data centers reached record highs in 2024, with the three largest hyperscalers committing approximately \$210 billion in AI-related capital expenditure (Allianz Commercial, 2025; Deloitte, 2025). Morgan Stanley projects nearly \$3 trillion in global data center spending by 2029; McKinsey estimates cumulative outlays approaching \$7 trillion by 2030 (Allianz Commercial, 2025).

Geopolitical competition reinforces the shift: at the February 2025 AI Action Summit, the EU announced InvestAI, a EUR 200 billion initiative (S&P Global, 2025). China operates at comparable scale through state-backed programs, with over 25 GW concentrated in the Yangtze Delta, Pearl River Delta, and Beijing-Tianjin corridor. China's "East Data West Compute" initiative is redistributing capacity toward western provinces—a strategy that reduces latency costs but increases exposure to seismic and water-stress hazards in regions with less resilient grid infrastructure. The Gulf states deploy sovereign AI infrastructure in regions where ambient temperatures already exceed cooling design thresholds.

A Decade of Climate Events

The vulnerability of digital infrastructure to physical climate hazards is documented across four distinct failure modes—each mapping onto the economic channels quantified in Chapter II.

Flooding strikes foundations. Hurricane Sandy (2012) submerged basement generators at major East Coast facilities under five feet of water (Data Center Dynamics, 2012). Hurricane Harvey (2017) tested Gulf Coast facilities at design limits with over 40 inches of rainfall (TRG Datacenters, 2024). In 2024, severe flooding across Bangladesh and Myanmar disrupted submarine cable landing stations serving regional cloud infrastructure.

Power and grid failure operates differently. Winter Storm Uri (2021) drove temperatures 30–40° F below seasonal norms, causing generator fuel to gel, batteries to fail, and cooling

systems to malfunction simultaneously. NERC⁴ concluded that the ERCOT⁵ grid came within minutes of uncontrolled cascade failure (Data Center Knowledge, 2021; NERC, 2025). The January 2026 US East Coast storm confirmed the pattern: wholesale electricity spiked ninefold overnight while the grid approached 147 GW of peak demand.⁶

Cooling failure is subtler but no less consequential. The July 2022 UK heatwave pushed ambient temperatures above 40° C for the first time in recorded history; Multiple hyperscaler and enterprise facilities in the UK reported outages—not because equipment failed, but because ambient conditions exceeded the thermal envelope within which mechanical cooling was designed to operate (Data Center Knowledge, 2021). A December 2025 analysis found that nearly 7,000 of the world's 8,808 operational data centers now sit outside optimal temperature ranges for cooling efficiency (Gault and Williams, 2025).⁷

Supply chain disruption completes the picture. A major typhoon caused cascading grid failures across one of Asia's densest concentrations of digital infrastructure in southern China. The 2024 Atlantic hurricane season produced Helene (\$30–47 billion in losses) and Milton (3 million customers without power) in rapid succession (Hexatronic, 2025). The Uptime Institute's 2024 Annual Outage Analysis confirmed that extreme weather was an increasing driver of serious outages industry-wide (TechTarget, 2024).

These episodes share a structural feature: data centers do not fail in isolation. They fail with grids, supply chains, and cooling systems. NERC warned that data center load growth "is contributing to continued risk of supply shortfalls" during extreme weather (NERC, 2025); Texas enacted legislation authorizing grid operators to curtail power to data centers during emergencies (Associated Press, 2025). Global data center electricity consumption reached approximately 450 TWh in 2024, growing at nearly 20% annually, with preliminary forecasts suggesting over 50 GW of peak capacity in PJM alone by 2030 (IEA, 2025; NRDC and Citizens Utility Board, 2025). At this scale, the interaction between demand growth and climate hazards is a permanent feature of the sector's operating environment.

The Insurance Protection Gap

Physical hazards damage assets. The investment problem arises from how risk is priced, absorbed, or refused by insurance markets. The global insurance protection gap stood at 56% in 2025: of \$296 billion in economic losses, \$167 billion went uninsured—a gap growing at 6–10% annually in real terms (Gallagher Re, 2026; Swiss Re Institute, 2025). In the EU, only 20% of natural catastrophe losses are covered (European Central Bank and EIOPA, 2024). In China, penetration remains below 15%.

⁴North American Electric Reliability Corporation.

⁵Electric Reliability Council of Texas: the independent system operator managing 90% of Texas's electric load.

⁶The structural dimension of grid stress is visible in PJM's capacity auction for the 2025–2026 delivery year, which cleared at \$269 per MW-day—an 833% increase from the prior year—with data center load accounting for \$6.5 billion of \$16.4 billion in total capacity costs (Howland, 2026; IEEFA, 2025).

⁷The 8,808 figure is from Gault and Williams 2025; this paper's database covers 8,572 facilities assembled from a distinct source set.

When insurance retreats, asset values follow. Evidence from the US residential market documents declines of 12–48% in property values following coverage non-renewal (WWF, 2025). More than 1.9 million home insurance policies have been non-renewed in the United States since 2018. For data centers, placement programs already exceed \$1 billion, with mega-campuses seeking up to \$3 billion syndicated across multiple insurers (Allianz Commercial, 2025). XDI⁸ projects that global insurance premiums for the sector could triple by mid-century (XDI, 2025).

When private insurance retreats, governments absorb part of the residual risk. The United States spent approximately \$1 trillion on disaster recovery and climate-related needs in the twelve months to May 2025 (WWF, 2025). France has raised catastrophe premium surcharges from 12% to 20%, with reinsurer CCR anticipating a 40% rise in claims by 2050 (Hogan Lovells, 2025).

The Design Window

Between 40% and 60% of 2035 global data center capacity will be built in the next five years (IEA, 2025). Once commis-


sioned, these facilities cannot relocate. The siting decisions made in this window will embed a quarter-century of climate exposure—or a quarter-century of climate resilience. The economics of that choice are asymmetric. In the United States, every dollar invested in climate resilience avoids up to \$13 in losses; in the United Kingdom, each GBP 1 spent on flood risk management avoids GBP 8 in damages (WWF, 2025). These are among the highest benefit-cost ratios in infrastructure economics. Yet current frameworks for data center siting optimize for power availability, fiber connectivity, and near-term costs. The global project pipeline concentrates investment in markets selected for land access and energy cost rather than climate resilience.

The bottlenecks already identified—grid connection timelines, transformer supply chains with 18–36 month lead times, cooling water availability—are the vectors through which climate hazards translate into financial risk. Every year of delayed adaptation narrows the design window and raises the retrofit cost. The window has not yet closed.

Exhibit 1. Four converging phenomena reshaping data center risk.

Phenomenon	Evidence	Scale	Geographic scope
1. Compute as strategic infrastructure	Critical infrastructure designations; sovereign AI programs; hyperscaler capex	\$3–7T cumulative investment by 2030	US, EU, China, Gulf, India, Brazil
2. Climate hazards disrupt operations	12+ major events across all DC markets since 2012	12+ major events since 2012	All major DC markets
3. Insurance retreat	Protection gap growing 6–10%/yr; coverage withdrawal; premium escalation	\$167B uninsured in 2025 (56% gap)	US, EU, Asia, Middle East
4. Design window closing	40–60% of 2035 capacity built in next 5 yrs; 15–25yr asset horizons	2×+ installed capacity in pipeline	Global

Exhibit 2. The scissors: capital flowing into data centers accelerates while insurance protection retreats.



Investment data: IEA (2025), McKinsey, Morgan Stanley estimates. Protection gap: Gallagher Re (2026), Swiss Re Sigma. Projections beyond 2025 are indicative.

⁸Cross Dependency Initiative: a climate physical risk analytics firm that models asset-level structural damage across eight hazard types.

The Missing Measurement

These converging pressures raise a direct question: how much climate risk does the global data center sector actually carry? No existing analysis answers it.

Academic literature

Academic research quantifying climate risk for data centers is sparse. A systematic review identified only one sustained program directly addressing the topic—the work of Pasupuleti (2022, 2024, 2025)—compared to hundreds of studies for analogous infrastructure sectors (Dawson et al., 2018; Forzieri et al., 2018; Luo et al., 2023; Oughton et al., 2023). This program establishes foundations for flood-aware site selection but does not attempt portfolio-level assessment, does not cover the full range of risk channels, and does not address the AI-era pipeline.

Quantitative frameworks for integrating climate risk into infrastructure valuation exist for energy and transport (In et al., 2020, 2021, 2022; Pasha et al., 2025; Prime et al., 2018), but none has been developed or validated for data center assets. The specific vulnerabilities—cooling system dependencies on ambient temperature, power quality requirements, IT equipment sensitivity, and the asymmetric revenue impact of service interruptions relative to physical damage—remain unmodeled. On supply chain propagation, only conceptual treatments exist (Koks et al., 2022); no quantitative model links upstream hazards to data center exposure. Compound hazards combining extreme heat, grid stress, and water scarcity remain unaddressed (Karagiannakis et al., 2024).

Practitioner analyses

On the practitioner side, three analyses provide the most comprehensive treatments. XDI’s 2025 report (XDI, 2025) covers 8,868 facilities across eight hazards with engineering vulnerability functions; major hubs across the US Northeast, Northern Europe, East Asia, and Southeast Asia rank among the top 20 most climate-exposed locations by 2050 (Data Center Dynamics, 2025; XDI, 2025). Its strength is physical engineering precision and coverage of the installed base. By design, it excludes extreme heat—the dominant economic channel for facilities where cooling accounts for 30–50% of operating costs—and omits transition risk entirely. The report provides a physical hazard map; the translation into financial valuation and the extension to future infrastructure remain open.

The WEF, drawing on S&P data, produced the first sector-aggregate financial projection: \$3.3 trillion in cumulative climate costs by 2055 (World Economic Forum and S&P Global, 2025). This estimate established financial materiality at board level but operates without facility-level breakdowns or channel decomposition. The underlying methodology has not been published in sufficient detail to enable reproduction.

Verdantix benchmarked 19 physical climate risk vendors across their full capability range (Verdantix, 2024). Its conclusion is direct: adaptation planning and supply chain risk quantification are the two weakest capabilities across the entire market. The analytical dimensions most relevant to capital allocation are those least developed.

Exhibit 3. Academic literature coverage of climate risk dimensions for data centers.

Dimension	Coverage	Key Gap
Physical risk quantification	3 DC studies (Pasupuleti, 2022, 2024, 2025); analogous sectors (Dawson et al., 2018; Forzieri et al., 2018; Oughton et al., 2023)	No sector-wide assessment; cooling and grid interdependencies unmodeled
Insurance & insurability	1 study, energy focus (Kerkhofs et al., 2025)	No DC-specific analysis; protection gaps undocumented
Valuation frameworks	Multiple for energy/transport (In et al., 2020, 2021, 2022)	None developed or validated for data centers
Supply chain propagation	Conceptual only (Koks et al., 2022)	No quantitative models linking upstream hazards to DC exposure
Compound hazards	Limited (Karagiannakis et al., 2024)	Heat + drought + grid stress scenarios unmodeled

Exhibit 4. Analytical scope: existing frameworks versus this paper.

	XDI	Pasupuleti	WEF/S&P	Verdantix	This paper
Data Center installed base	8,868	—	—	—	8,572
AI pipeline (prospective)	—	—	—	—	615–2,939
Investment-grade valuation	Partial	Partial	Partial	—	ClimVaR % EV
Thermal / cooling risk	—	—	Partial	—	γ channel
Transition risk (carbon)	—	—	—	—	CC^{1,2,3} + PPA
Channel attribution	—	—	—	—	Shapley
Supply chain propagation	—	—	—	Gap ID	Leontief I-O
Scenario architecture	RCP 8.5	Partial	High only	—	RCP × NGFS
Adaptation cost-benefit	1 archetype	Site only	—	Gap ID	3 pathways

Full = explicit methodology and sector-wide application; Partial = aggregate or single-site only; — = not addressed. Gap ID = gap identified but not quantified.

Three structural gaps

Three gaps explain why climate risk remains unpriced. First, existing analyses measure physical hazard intensity or replacement cost—not enterprise value loss, the metric that governs capital allocation, TCFD disclosure, and insurance underwriting. Second, all existing analyses characterize the installed base; none models the prospective AI-era

pipeline that represents over 80% of projected capacity additions. Third, no framework decomposes aggregate risk into actionable economic channels. Without knowing whether a facility's exposure is driven by cooling costs, carbon pricing, physical damage, or supply chain disruption, adaptation budgets cannot be allocated with analytical rigor. Verdantix documents this gap across all 19 vendors it benchmarks; none closes it.

What This Paper Does

This paper addresses all three gaps. It asks: *What is the Climate Value at Risk for global data center infrastructure, how does it vary across locations and scenarios, and which adaptation strategies generate positive risk-adjusted returns?*

It makes four contributions. The ClimVaR methodology (Epelbaum, 2025b) is adapted to data center infrastructure, embedding physical and transition risks into discounted cash flow analysis through five sector-specific channels: cooling stress, carbon pricing, business interruption, physical damage, and heat productivity loss. Shapley decomposition (Shapley, 1953) attributes each channel's contribution to specific hazards; Leontief input-output analysis (Miller and Blair, 2009b) propagates business interruption through upstream supply chains. Neither tool has previously been applied to this sector. The geocoded

database covers 8,572 installed facilities and up to 2,939 prospective AI-era facilities under three growth scenarios. The scenario architecture crosses warming pathways, transition policies, AI growth trajectories, and adaptation strategies to identify which interventions generate positive returns under which conditions. The adaptation business case is then quantified: proactive strategies reduce aggregate ClimVaR by 30–39% with positive net present value under every scenario tested—and confirms that all ten do.

The valuation framework and the data it requires are the subject of Chapter II. Chapter III applies this framework to four questions in sequence: how much risk the sector carries, where it concentrates, how the AI build-out alters the profile, and whether adaptation generates positive returns. The answers are not reassuring on the first three counts. They are on the fourth.

Chapter II

Methods & Data

The ClimVaR Model

Climate Value at Risk measures the present value of climate-induced losses over an asset's lifetime—the gap between what a facility is worth ignoring climate, and what it is worth pricing it in (Epelbaum, 2025b). The framework rests on three analytical layers: a climate-adjusted DCF model, Shapley decomposition for hazard attribution, and Leontief input-output analysis for supply chain propagation. Both mathematical tools are grounded in foundational contributions to economics—Shapley (Nobel 2012) for cooperative game theory, Leontief (Nobel 1973) for input-output analysis—but neither has been applied to data center infrastructure.

Climate-adjusted DCF

Infrastructure valuation conventionally relies on discounted cash flow analysis, estimating present value

as the sum of future cash flows discounted for time and risk. ClimVaR embeds climate effects directly into these cash flows, holding the discount rate constant. The result is the difference between a facility's climate-naïve value (V_0) and its climate-adjusted value (V^{cc})—cumulated over a 25-year horizon and discounted at the facility's weighted average cost of capital:

$$\text{ClimVaR} = V_0 - V^{cc} = \sum_{t=1}^{T+1} \frac{FCF^{(t)} - FCF^{(t,cc)}}{(1 + \omega)^t}$$

This makes ClimVaR a conservative lower bound on total value erosion: a climate-adjusted WACC⁹ would amplify the results. The full model aggregates revenue losses, COGS¹⁰ increases, and OPEX¹¹ increases across five channels.

$$\text{ClimVaR} = \sum_{t=1}^{T+1} \frac{(1 - \tau)}{(1 + \omega)^t} \left[op \cdot \underbrace{\Delta Y^{(t)}}_{\beta} + Y^{(t,cc)} \left(\underbrace{s_o \xi^{(t)}}_{\xi} + \underbrace{e_c \gamma^{(t)}}_{\gamma} \right) + \underbrace{\sum_{i=1}^3 CC^{(t,i)}}_{CC^{1,2,3}} + \underbrace{(1 + \iota) \cdot AV \cdot \delta^{(t)}}_{\delta} \right]$$

τ : tax rate | ω : WACC | op : OPEX ratio | $\Delta Y^{(t)}$: revenue loss from business interruption (β) | $\xi^{(t)}$: heat productivity loss | $\gamma^{(t)}$: cooling cost multiplier | $CC^{(t,i)}$: carbon costs (Scopes 1–3) | $\delta^{(t)}$: physical damage rate | ι : insurance loading factor | AV : asset value. Each underbrace maps to a risk channel from Exhibit 5. Complete derivation in Appendix A.

⁹Weighted Average Cost of Capital.

¹⁰Cost of Goods Sold: direct costs attributable to service delivery, including energy and cooling.

¹¹Operating Expenditure: recurring costs of running a facility, including maintenance, insurance, and staffing.

What Makes Data Centers Different

Data centers exhibit characteristics that make climate risk assessment both urgent and methodologically distinctive. Four features drive the calibration of ClimVaR channels to this sector.

Cooling intensity. Data centers consume 30–50% of facility energy for cooling, making the γ channel unusually prominent. Each 1°C increase in ambient temperature raises cooling costs by 2–4%, degrading Power Usage Effectiveness (PUE) and directly eroding operating margin.

Electricity dependency. Energy intensity 10–100× that of office buildings means carbon cost channels (CC^{1,2,3}) dominate transition risk. Scope 2 (location-based, purchased electricity) typically represents 70–85% of to-

tal emissions. The highly variable carbon intensity of electricity across jurisdictions—from near-zero in hydro-dominated grids to coal-dependent systems—creates substantial geographic variation in transition exposure.

Uptime requirements. 99.99% availability Service Level Agreements mean business interruption (β) translates directly into contractual penalties and customer churn. A single hour of unplanned downtime can cost \$100,000–\$1M in SLA penalties, lost revenue, and reputational damage.

Supply chain concentration. Critical dependence on geographically concentrated semiconductor fabrication (Taiwan 54%, South Korea 17%) and power equipment manufacturing (transformer lead times 12–24 months) amplifies the Leontief propagation effect far beyond what facility-level hazard screening captures.

Exhibit 5. ClimVaR risk channels: notation, mechanisms, and data center calibration.


Symbol	Channel	Climate driver	DC-specific mechanism
γ	Cooling costs	Temperature increase	30–50% of energy; each 1°C adds 2–4% cooling cost
CC ^{1,2,3}	Carbon costs	Carbon pricing trajectory	Electricity intensity 10–100× office; Scope 2 dominant
β	Business interruption	Extreme weather events	Grid stress, supply chain disruption; 99.99% SLA exposure
δ	Physical damage	Floods, storms, wildfire	Structural damage; insurance cost escalation
ξ	Heat productivity	Sustained heat waves	Outdoor labor and equipment efficiency ¹²
σ	Customer shift	Not parameterized — insufficient calibration data	
ρ	Nature capital	Not parameterized — relevant for evaporative cooling in water-stressed regions	

The five parameterized channels capture distinct transmission mechanisms. Cooling costs (γ) and heat productivity (ξ) respond to gradual temperature trends. Carbon costs (CC) depend on policy trajectories. Business interruption (β) and physical damage (δ) respond to dis-

crete extreme events. This heterogeneity matters: channels driven by trend variables can be forecast with reasonable confidence; channels driven by extreme events carry irreducible uncertainty that adaptation can reduce but not eliminate.

¹²Equipment efficiency loss under sustained heat is not yet parameterized in the current model version and merits future parameterization, as thermal derating of outdoor electrical and mechanical systems may be material at high-density sites.

Exhibit 6. ClimVaR framework: climate drivers propagate through five risk channels to affect enterprise value.



Adapted from Epelbaum (2025). Complete derivation in Appendix A.

The general ClimVaR framework defines seven risk channels. Five are parameterized in this analysis. Two—customer shift (σ , capturing demand migration away from climate-exposed providers) and nature capital (ρ , capturing ecosystem service degradation)—remain in the mathematical architecture but are not activated due to insufficient sector-specific calibration data. Their exclusion makes the estimate conservative: both channels, if parameterized, would increase total ClimVaR. The ρ channel is potentially material for facilities using evaporative cooling in water-stressed regions—a gap identified for future parameterization.

Shapley decomposition

Business interruption aggregates exposure to multiple hazards that interact non-additively: a facility facing both drought and heat experiences correlated effects exceeding the sum of individual impacts. Simple additive decomposition produces attributions that either fail to sum to total risk or depend arbitrarily on ordering.

The Shapley value—from cooperative game theory (Nobel Prize 2012)—resolves this by averaging each hazard’s marginal contribution across all possible orderings of hazard inclusion. With five hazards, exact computation requires $2^5 = 32$ coalition evaluations per facility—tractable for the full database of 8,572 sites. The decomposition satisfies four axioms simultaneously: efficiency (attributions sum to total risk), symmetry, null player (irrelevant hazards receive zero), and additivity (Shapley, 1953). For a hazard h

in the set H of all hazards, the Shapley value is:

$$\phi_h(v) = \sum_{S \subseteq H \setminus \{h\}} \frac{|S|! (|H| - |S| - 1)!}{|H|!} [v(S \cup \{h\}) - v(S)]$$

where $v(S)$ is the business interruption cost when only hazards in coalition S are active. The bracketed term is the marginal contribution of hazard h given that hazards in S already operate. Shapley decomposition applies within the β and δ channels, where hazard interactions are significant. Across channels, ClimVaR components aggregate additively through distinct cash flow lines. The formal derivation is in Appendix A.

Supply chain propagation (Leontief)

Facility-level analysis misses indirect exposure through upstream supply chains. A flood affecting semiconductor fabrication in Taiwan, a hurricane damaging transformer manufacturing on the Gulf Coast, or a drought constraining equipment production—all translate into business interruption for facilities thousands of kilometers distant.

Leontief input-output analysis (Nobel Prize 1973) traces this propagation using OECD¹³ Multi-Regional Input-Output Tables (2023 vintage, reference year 2020) (Miller and Blair, 2009b). The Leontief inverse captures the full cascade—not just immediate suppliers, but suppliers of suppliers. The empirical multiplier is $\times 11$: for every dollar of direct business interruption, eleven dollars of supply chain value are at risk. At every facility examined, 98–99% of business interruption risk originates upstream, not on site.

¹³Organisation for Economic Co-operation and Development.

At every facility examined, 98--99% of business interruption risk originates upstream in the supply chain, not on site. A facility's apparent geographic safety may tell little about its actual interruption exposure.

8,572 Facilities, 120 Countries

The core database comprises 8,572 geocoded facilities totaling 90.8 GW across 120 countries, assembled from industry databases, regulatory filings, operator disclosures, and satellite imagery validation. Each facility carries 29 documented financial parameters (revenue, asset value, cost ratios, PUE, carbon intensity, insurance). Financial calibration uses regional and tier-specific benchmarks rather than facility-specific data. Asset values range from \$5M/MW for legacy colocation in emerging markets to \$42M/MW for hyperscale AI facilities in the United States. This archetype-based approach enables full database coverage while introducing uncertainty proportional to individual facilities' divergence from their regional benchmark.

Physical risk variables are extracted from bias-corrected, statistically downscaled CMIP projections at facility-level coordinates: temperature, precipitation, wind speed, and flood depth across five hazard types (tropical wind, river flood, coastal flood, wildfire, extreme heat). The total dataset comprises approximately 50 million climate observations across two RCP pathways and the 2026–2050 horizon. Projections are sourced from multi-model ensembles; the Q25/MEAN/Q95 reporting structure in Chapter III reflects variation from spatial heterogeneity, model sensitivity, and stochastic event occurrence.

Supply chain propagation uses the OECD Multi-Regional

Input-Output Tables (2023 vintage, 75 countries × 45 sectors), with eight components extracted for the data center value chain. Carbon pricing trajectories are drawn from NGFS¹⁴ scenarios (2023), differentiated by region and policy pathway. Grid emission factors use IEA¹⁵ country-level data (2024), with forward projections aligned to each transition scenario.


Prospective capacity represents the AI-era pipeline through 2050, modeled under three growth scenarios from Paccou and Wijnhoven (2024): **Limits to Growth** (LTG, +63 GW)—where energy and resource constraints slow AI expansion; **Sustainable AI** (SAI, +121 GW)—growth within efficiency guardrails; and **Abundance Without Boundaries** (AWB, +296 GW)—compute demand scales with minimal constraint. All three are nested by construction: LTG ⊂ SAI ⊂ AWB. Facilities are drawn from four source pools ordered by decreasing certainty: Epoch AI GPU cluster database (62 facilities, 6.3 GW), Tier 2 announced projects (60 facilities, 6.0 GW), and two probabilistic pools allocated via regional demand models (IEA, 2025; Paccou and Wijnhoven, 2024). Major identified projects include large-scale AI training facilities in the US South (up to 2,800 MW), next-generation GPU clusters, and sovereign AI deployments. Actual buildout will differ in timing, scale, and location; ClimVaR estimates for prospective infrastructure are indicative of the risk profile under each scenario, not point predictions.

Exhibit 7a. Facility coverage by region and AI growth scenario.

Region	Installed		+ Limits to Growth		+ Sustainable AI		+ Abundance	
	Sites	GW	Sites	GW	Sites	GW	Sites	GW
United States	2,383	30.3	2,620	54.0	2,831	75.1	3,464	138.4
China	1,284	25.1	1,395	36.2	1,522	48.9	1,902	86.9
Europe	2,019	15.4	2,101	23.6	2,170	30.5	2,376	51.1
Asia-Pacific	1,593	11.6	1,692	22.9	1,804	34.3	2,145	68.5
Rest of World	1,293	8.5	1,379	17.1	1,440	23.2	1,624	41.6
Total	8,572	90.8	9,187	153.7	9,767	212.0	11,511	386.4

¹⁴Network for Greening the Financial System: a consortium of central banks and supervisors developing climate scenario frameworks.

¹⁵International Energy Agency.

Exhibit 7b. AI growth scenarios: capacity trajectories with sensitivity to system dynamics parameters.

Existing climate risk assessments for data centers focus exclusively on the installed base—the 8,572 facilities already operating. This analysis is the first to extend ClimVaR to prospective infrastructure: the 615 to 2,939 facilities (depending on scenario) whose design has not yet been finalized. This matters because prospective capacity will exceed installed capacity by 2035 under all three scenarios, and because design-stage decisions carry fundamentally different economics than retrofit decisions.

The distinction matters for adaptation economics. Legacy facilities face locked-in exposure reducible only through costly retrofits (40–60% of CAPEX). Prospective facilities can integrate climate resilience at the design stage—

typically at 15–20% of CAPEX. Each year without action converts a greenfield opportunity into a future retrofit cost. The window is narrow: facilities entering design in 2025–2030 will operate through 2050–2055, spanning the period when climate impacts accelerate most sharply under both RCP pathways.

Of the 29 financial parameters, 48% are calibrated from high-confidence public sources, 34% from mixed sources, and 17% from expert judgment. Sensitivity analysis confirms robustness to $\pm 20\%$ variation in expert-judgment parameters; the dominant uncertainty source is physical hazard projection, not financial calibration. The complete data architecture is documented in Appendix B.

Four Dimensions of Uncertainty

The analysis crosses four dimensions of uncertainty. The first two—physical pathways and transition policies—represent external uncertainty that operators cannot individually influence. The third—AI growth—determines the asset base. The fourth—adaptation—is the strategic variable.

Exhibit 8. Scenario architecture: four dimensions of uncertainty.

Dimension	Source	Options	What it determines
Physical pathway	IPCC (RCP)	2.6 / 8.5	Temperature trajectories, extreme event frequency, flood probabilities
Transition policy	NGFS	BAU / Net Zero	Carbon price trajectories, grid decarbonization pace
AI growth	Paccou & Wijnhoven (2024)	LTG / SAI / AWB	Scale and geography of infrastructure deployment
Adaptation strategy	Operator choice	Baseline / Selective / Proactive	Technology, siting, and retrofit decisions

Physical pathways

Representative Concentration Pathways (RCPs)¹⁶ determine the physical hazard environment—temperature trajectories, extreme event frequencies, flood and storm probabilities—which operators cannot influence individually. Two endpoints bracket the range from Paris-aligned mitigation to unabated emissions.

Exhibit 9. Physical climate scenarios and data center implications.

Scenario	Warming by 2100	Implied policy	DC-relevant implications
RCP 2.6 / SSP1-2.6	+1.5–2.0°C	Paris Agreement met	Cooling costs +15–25%; extreme events +20–40%
RCP 8.5 / SSP5-8.5	+4.0–5.0°C	No mitigation	Cooling costs +35–50%; extreme events +100–200%

RCPs drive physical risk channels: γ (cooling), β (extreme events), δ (damage), and ξ (heat exposure). Downscaled from CMIP ensemble (Eyring et al., 2016).

Transition and Adaptation

Transition policies cross physical pathways with carbon pricing trajectories. Under BAU, carbon prices follow current legislated trajectories. Under Net Zero Transition (NZT), prices rise to levels consistent with NGFS net-zero-by-2050 scenarios (NGFS, 2023). NZT increases installed base ClimVaR (from 38% to 83%) because higher carbon prices amplify the CC channel—a counter-intuitive result with direct implications for operators whose sustainability strategy relies on transition risk declining over time.

Three adaptation strategies are modeled. **Baseline:** current practice, no additional climate investment. **Selective:** targeted intervention on the highest-Shapley-value channel per facility—the “one champion action” approach. **Proactive:** comprehensive investment across all five channels. Risk reductions across channels are multiplicative, not additive, because the channels interact through the Shapley structure.

Three AI growth scenarios bracket deployment uncertainty

(Paccou and Wijnhoven, 2024). **Limits to Growth (LTG):** energy and resource constraints slow AI expansion (+63 GW). **Sustainable AI (SAI):** growth within efficiency guardrails (+121 GW). **Abundance Without Boundaries (AWB):** compute demand scales with minimal constraint (+296 GW). All three are nested by construction: $LTG \subset SAI \subset AWB$.

Intra-scenario uncertainty

Within each scenario combination, results are reported at three statistical levels to characterize the distribution of outcomes: **Q25** (25th percentile—optimistic estimate), **MEAN** (expected value across the facility distribution—the central estimate used for most results), and **Q95** (95th percentile—tail risk estimate). This allows distinguishing between expected outcomes and tail risks that drive insurance pricing and capital allocation decisions. The Q25–Q95 range reflects variation from spatial heterogeneity, firm-level differences in carbon intensity, and stochastic event occurrence. Unless noted otherwise, all results in Chapter III use the BAU/MEAN pathway.

¹⁶RCPs are greenhouse gas concentration trajectories adopted by the IPCC for climate modelling. RCP 2.6 is consistent with the Paris Agreement; RCP 8.5 represents unabated emissions.

Chapter III

Results

The repricing is already here

ClimVaR does not describe a distant tail risk. It reveals a present mispricing of digital infrastructure value under climate-naïve assumptions. Unless otherwise specified, all results throughout this paper are reported under the BAU/MEAN pathway (RCP 8.5, business-as-usual transition policy, mean climate projection).

Exhibit 10. Panoramic results: ClimVaR across asset classes, AI scenarios, and adaptation strategies (BAU/MEAN).

Portfolio	V ₀ (\$B)	Baseline	Selective	Proactive	Max red.
Installed base	1,012	388 (38%)	305 (30%)	238 (24%)	39%
United States	393	108 (28%)	89 (23%)	72 (18%)	33%
Europe	201	107 (53%)	81 (40%)	60 (30%)	44%
China	198	111 (56%)	86 (44%)	67 (34%)	40%
APMEA ¹⁷	220	62 (28%)	49 (22%)	39 (18%)	37%
Installed + LTG	1,992	1,008 (51%)	821 (41%)	669 (34%)	34%
Installed + SAI	2,849	1,669 (59%)	1,366 (48%)	1,134 (40%)	32%
Installed + AWB	5,590	3,671 (66%)	3,006 (54%)	2,536 (45%)	31%

ClimVaR% in parentheses = ClimVaR as % of V₀. Max reduction = (Baseline – Proactive) / Baseline. Under NZT, installed base ClimVaR rises to \$845B (83%), with Proactive reducing it by 52%.

Aggregate mispricing at portfolio scale

The installed base of 8,572 data centers carries **\$388 billion** in climate-related value at risk—38.3% of enterprise value under the BAU pathway. With prospective AI infrastructure, exposure rises to **\$1.0–3.7 trillion**, depending on the growth scenario. Even the most conservative pathway exceeds \$1 trillion. To put this in perspective: the installed-base figure alone is larger than the annual revenue of every listed data center operator combined.

This is why ClimVaR should be read as a measure of systematic mispricing. A traditional DCF assumes that cooling costs, carbon prices, outages, and physical damages remain within the historical range embedded in past cash flows. Climate-adjusted DCF removes that assumption. The 38.3% gap is therefore not a tail event; it is the median

repricing implied by the BAU/MEAN pathway.

Intensity rises with AI scale

Exposure is not evenly distributed. Europe and China reach 53–56% ClimVaR intensity, while the United States stands at 28%—a twofold gap that reflects not just geography, but different channel structures. Business interruption dominates in the United States, carbon in Europe, and physical damage in Asia-Pacific. A single global reading of “data center climate risk” therefore obscures the real levers of action.


Just as importantly, **risk intensity increases with AI scale**. Installed-base intensity is 38%; with AI infrastructure, it rises to 51% (LTG), 59% (SAI), and 66% (AWB). This is not a simple scaling effect. It reflects a deeper shift in portfolio composition, where physical exposure becomes increasingly dominant as compute expands.

¹⁷ Asia-Pacific, Middle East, and Africa.

The Transmission Mechanism

How hazards propagate through risk channels to erode enterprise value

Exhibit 11. The transmission map: hazards to enterprise value loss.



Seven hazards propagate through four principal risk channels to erode enterprise value. Ribbon width proportional to US\$ value. Shapley attribution for hazard-to-channel mapping. BAU/MEAN, installed base.

From hazards to financial channels

Exhibit 11 makes visible the mechanism that the aggregate figure of \$388 billion tends to obscure. Climate risk does not arrive at the balance sheet in the abstract. It passes through a sequence: hazards strike physical systems, physical systems disrupt operations, and those disruptions are translated into discounted cash-flow losses. What the transmission map shows is therefore not an illustration but the architecture of the loss itself.

On the left, seven hazards are ranked by contribution. Tropical wind dominates, followed by extreme heat and river flood, while carbon pricing appears as a transition hazard with a weight comparable to the largest physical drivers. This ordering is already instructive. A naive reading of data-center exposure might expect flood to dominate, because flood is the most visible and the most intuitively destructive hazard. Yet tropical wind ranks first because its effects travel further: not only through direct site exposure, but through semiconductor fabrication, electrical equipment manufacturing, logistics corridors, and other supply-chain dependencies whose disruption propagates across the portfolio.

Why the map changes the diagnosis

On the right, hazards are aggregated into risk channels. Business interruption is the largest channel, ahead of carbon costs, physical damages, and cooling penalties. This matters because the ranking of channels is not the ranking of hazards: a single hazard can contribute to several channels simultaneously, and a single channel can be fed by several hazards at once. Tropical wind, for instance, contributes both to business interruption and to physical damage; extreme heat drives cooling stress through rising ambient temperatures. The map therefore formalizes the non-additivity of climate risk.

This is also why Shapley decomposition is essential to the framework. Without it, one would either double-count the contribution of a hazard that affects several channels, or collapse distinct mechanisms into a single undifferentiated estimate—both analytically fatal for adaptation planning. The model does not merely say that climate risk is large; it shows how it is transmitted, through which levers it becomes financially material, and therefore where adaptation can act with the greatest force.

Finding 1: Unpriced Exposure

Climate risk erodes 38% of current data center asset value

ClimVaR as structural repricing of the asset

The installed base of 8,572 data centers, valued at US\$1,012 billion under standard discounted cash-flow assumptions, carries US\$388 billion in climate-adjusted value erosion—38.3% of enterprise value. No current valuation framework captures this figure, no disclosure regime requires its reporting, no insurance product prices it. US\$388B exceeds the combined annual revenue of every publicly listed data center operator. It is a lower bound.

ClimVaR can be interpreted in the same way as a credit impairment or a permanent margin compression: not as the probability of a catastrophic event, but as a structural repricing of the asset over its 15–25 year lifetime. A facility with a climate-naïve valuation of US\$100 million and a ClimVaR of 38% does not face a 38% chance of total loss; it faces a future in which, on the median pathway, cooling penalties, carbon costs, outages, and damage combine to erode US\$38 million of its discounted cash flows.

Five channels, three central observations

The aggregate figure combines five risk channels, each operating through a distinct financial mechanism and each responding to different adaptation levers. Three observations structure the decomposition that Exhibit 12 presents.

First, business interruption leads at 31% of total ClimVaR, but 91% of it propagates through supply chains rather than through direct facility disruption—a ratio of eleven to one between indirect and direct exposure. An operator monitoring only on-site hazards misses the primary risk vector. Second, carbon costs represent 28% and span all three scopes, but are almost entirely Scope 2 (location-based): US\$106B of US\$107B originate in grid carbon intensity, not in on-site emissions. Scope 3 is included in the model but negligible at this scale. The PPA¹⁸ is therefore the primary lever for transition risk. Third, the physical channels—business interruption, damages, cooling—together account for 72% of total ClimVaR on the installed base; carbon accounts for 28%. This ratio inverts when AI infrastructure enters the portfolio, as Finding 3 documents. From an investment perspective, the distinction between channels matters more than the aggregate number. A credit analyst can treat ClimVaR as an additional spread over the climate-naïve discount rate, but an engineer or operator needs to know whether that spread comes from outage risk, from cooling penalties, from structural damage, or from carbon. The same dollar of adaptation budget does not buy the same reduction in each channel. Exhibit 12 is therefore not a purely descriptive breakdown; it is the bridge between the macro figure (US\$388B) and the micro-level levers that later sections quantify.

Exhibit 12. ClimVaR decomposition by risk channel—installed base, BAU/MEAN.

Risk Channel	\$B	Share	Mechanism
Business interruption (β)	121.4	31%	Extreme events, grid stress, downtime
Carbon costs ($CC^{1,2,3}$)	107.1	28%	Grid emissions $\times$ carbon price
Physical damages (δ)	84.1	22%	Floods, storms, structural damage
Cooling penalty (γ)	72.6	19%	Ambient temperature $\rightarrow$ energy cost
Heat productivity (ξ)	2.4	1%	Worker/equipment efficiency loss
Total ClimVaR	387.5	100%	

Tropical wind versus flood: severity and reach

Why does tropical wind rank first among hazards, ahead of flood? The answer lies in the distinction between severity and reach. River flood inflicts catastrophic damage at exposed sites but affects only facilities in floodplains. Tropical wind propagates through supply chains serving the entire portfolio: a typhoon disrupting semiconductor fabrication in Taiwan degrades every facility depending on those inputs, regardless of location. At portfolio level, a hazard touching many facilities indirectly outweighs one devastating a few directly. The Sankey in Exhibit 11 makes this visible: the tropical-wind ribbon is wide because its consequences travel further through the network.

What \$388B does not capture

The framework does not capture nature-related risk—water stress, biodiversity constraints, ecosystem degradation around cooling sources—nor systemic externalities: grid strain from concentrated AI load, community displacement, or cascading service interruptions across interconnected facilities. Each omitted channel would widen, not narrow, the exposure estimate. The \$388B figure is therefore the floor of an exposure range whose ceiling remains unquantified. For investors and regulators: climate risk on digital infrastructure is already material at the lower bound; model refinements are more likely to move the estimate upward than downward.

¹⁸Power Purchase Agreement: a long-term contract to buy electricity directly from a renewable energy producer, bypassing grid carbon exposure.

Finding 1: Unpriced Exposure

Four regions, five channels, and the geography of \$388 billion

Exhibit 13. US\$388 billion decomposed: four regions, five channels, one treemap.



Area proportional to ClimVaR in US\$B. Outer rectangles = regions; inner tiles = risk channels. BAU/MEAN, installed base. The dominant tile in each region identifies the primary adaptation lever.

One map, two dimensions of exposure

The treemap encodes the same US\$388B from two angles simultaneously: region and channel. The area of each tile is proportional to dollar exposure, so that a single glance shows three things: where the money is, which channel drives it, and how the composition shifts by geography. Outer rectangles show regions; inner tiles show the five risk channels. The dominant tile in each region is therefore the first-order adaptation lever. Note that business interruption (91% supply-chain-driven) and carbon costs (99% Scope 2) are both quasi-exclusively upstream exposures: the treemap therefore also maps supply chain vulnerability, even though it is organized by financial channel.

Four fingerprints, four risk structures

The United States concentrates 28% of global exposure—US\$108B—but at the lowest intensity (27.6%). Business interruption fills nearly half the US rectangle, a consequence

of dependence on East Asian semiconductor supply chains and grid stress during extreme heat events. Europe shows a different signature: US\$45B in carbon costs alone, driven by higher carbon pricing under the EU ETS¹⁹ and residual grid emission intensity. China carries a triple threat of comparable magnitude—carbon (US\$36B), physical damage (US\$33B), and business interruption (US\$31B)—so that intervention on any single channel leaves roughly two-thirds of the exposure intact. APMEA is the smallest rectangle but the most evenly distributed across channels, a profile that resists single-lever intervention.

The asymmetry between these four fingerprints is the diagnostic: the US\$388B headline is not a single problem but a family of regional problems that share a magnitude and differ in mechanism. A uniform adaptation strategy systematically underinvests in the dominant channel. The Shapley decomposition quantifies this fingerprint from the global portfolio down to individual facilities.

¹⁹European Union Emissions Trading System: the world's largest carbon market, covering approximately 40% of EU greenhouse gas emissions.

Finding 2: Geographic Dispersion

Unevenly distributed, differently composed _____

The same asset class, four different risk worlds

The global headline conceals a decisive fact: there is no single geography of data-center climate risk. The installed base spans 8,572 facilities across four major regions, but the exposure profile is not merely a matter of different magnitudes; it is a matter of different structures. Europe and China reach ClimVaR intensities of 53–56%, while the United States stands at 28%. That twofold spread cannot be explained by hazard maps alone. It reflects the interaction between regional climate conditions, electricity systems, industrial dependencies, and carbon-policy regimes.

This distinction matters because adaptation is not purchased against “climate risk” in the abstract. It is purchased against specific transmission mechanisms. A portfolio exposed primarily through business interruption requires different interventions from one exposed primarily through carbon, cooling, or direct damage. Once risk is decomposed by geography and channel, the idea of a universal adaptation playbook begins to collapse.

Regional intensity is only the first layer of the problem

The first reading of the evidence is intuitive: some regions are simply riskier than others. But this is only the first layer. The deeper point is that regional risk intensity and regional risk composition do not move together in a simple way. The United States carries the largest absolute share of enterprise value and a large share of total ClimVaR, yet does so at lower relative intensity. Europe carries less enterprise value but much higher relative erosion because carbon pricing and grid exposure weigh more heavily on long-term cash flows.

China combines high intensity with a more diversified set of channels, making it less responsive to narrow interventions.

In other words, a region can matter for three different reasons: because it is large, because it is intense, or because its internal composition makes adaptation harder. Those three logics do not always coincide. This is precisely why a portfolio-wide average is analytically weak: it merges scale, intensity, and composition into a single figure and thereby obscures the actual economics of intervention.

Why a global average misallocates adaptation capital

A global average encourages the wrong instinct: spread adaptation capital evenly, or prioritize the largest regions first. Yet neither rule follows from the evidence. If one region is dominated by carbon exposure, another by business interruption, and a third by physical damages, then the same dollar does not buy the same protection in each geography. Capital allocation must therefore be differentiated not only by hazard level, but by the structure of loss transmission.

This is one of the central analytical contributions of the framework. The model does not stop at saying that geography matters; it specifies **how** it matters. It identifies where the portfolio is most eroded, through which channels that erosion occurs, and therefore where targeted adaptation is likely to produce the greatest reduction in value at risk. Geography is not a backdrop to the model. It is one of its principal explanatory dimensions.

Exhibit 14. Regional ClimVaR intensity and adaptation headroom across the installed base.

Region	V ₀ (\$B)	Baseline	Selective	Proactive	Max red.
United States	393	108 (28%)	89 (23%)	72 (18%)	33%
Europe	201	107 (53%)	81 (40%)	60 (30%)	44%
China	198	111 (56%)	86 (44%)	67 (34%)	40%
APMEA ²⁰	220	62 (28%)	49 (22%)	39 (18%)	37%

Adaptation value also varies by region

Exhibit 14 adds a second layer to the diagnosis: not only are regional risks uneven, but the value of adaptation is uneven as well. Europe shows the highest proportional reduction under the Proactive strategy, with a 44% reduction from baseline. China follows at 40%, then APMEA at 37%, while the United States reaches 33%. The implication is important: the most adaptation-responsive regions are not

necessarily the largest, but they may offer the greatest risk reduction per unit of intervention.

This is where the model becomes operationally powerful. It makes visible not only where risk is located, but where adaptation has the strongest leverage over future value erosion. For a global operator, this is the difference between adaptation as generic resilience spending and adaptation as disciplined capital allocation.

²⁰Asia-Pacific, Middle East, and Africa.

Finding 2: Geographic Dispersion

Why adaptation must be targeted, not standardized

Exhibit 15. Four fingerprints: regional channel composition as radar overlay.



Each polygon is one region's channel fingerprint, expressed as % of regional ClimVaR. The asymmetry between shapes shows why a uniform adaptation strategy fails. BAU/MEAN, installed base.

What the radar chart reveals

The radar shows what tables cannot: regions differ not only in *how much* ClimVaR they carry, but in *how they carry it*. Each geography has a distinct risk fingerprint. The United States is pulled outward by business interruption, Europe by carbon, China by a multi-channel structure, and APMEA by a flatter profile. These shapes matter because adaptation strategies act on shapes, not just on totals.

The table says that Europe is highly exposed and highly responsive to proactive adaptation; the radar explains why. Carbon dominates its risk architecture, and carbon is a channel where targeted interventions have strong effect. China, by contrast, shows no single dominant channel, so narrow measures have limited leverage unless they are combined.

Four regions, four adaptation logics

The United States calls for a *continuity* logic: with business interruption dominant, resilience hinges on supply-chain robustness, backup systems, and reduced downtime. Europe calls for a *decarbonization-and-procurement* logic: carbon is central to value erosion, making power sourcing and grid exposure decisive. China calls for a *portfolio* logic, since exposure is spread across several large channels. AP-

MEA calls for a *balancing* logic, as its more even profile resists single-lever intervention.

Regional fingerprints shift the question from "Where is risk high?" to "What kind of intervention is economically rational here?" They preserve analytical rigor while making visible the use value of the framework for capital planning, siting, resilience, and energy strategy.

Geography determines what counts as adaptation

The strongest implication of Finding 2 is not just that regions should be ranked differently. It is that geography changes the meaning of adaptation. In one region, adaptation mainly means reducing outage risk; in another, lowering carbon exposure; in a third, managing physical damage and cooling stress together. Treating all of these under a single adaptation label is conceptually convenient but economically misleading.

The framework matters beyond the headline estimate because it offers a white-box view of regional heterogeneity—transparent, decomposable, and usable for scenario design. It shows what black-box risk scores often miss: adaptation is not simply a volume of spend, but a geography-specific strategy whose effectiveness depends on **matching interventions to the regional structure of loss**.

Finding 3: AI Amplification

AI-era infrastructure carries $\times 2.6$ the physical risk per GW

From carbon-driven to physically-driven risk

The AI build-out does not replicate the risk profile of the installed base—it inverts it. On legacy facilities, carbon costs represent 28% of total ClimVaR and physical channels 72%. On AI-era facilities, physical channels account for 93–94% while carbon drops to 6–7%. The inversion is not gradual: even the most conservative growth scenario (Limits to Growth, adding 63 GW) pushes the physical share above 93%.
This has a direct consequence that no sustainability strategy focused on renewable procurement can resolve. Net-zero commitments address the carbon component—which on AI infrastructure represents less than one-tenth of the exposure. For operators who have built their climate strategy around Scope 2 reductions, the centre of gravity of risk has moved somewhere their existing toolkit barely touches.

A portfolio that adds capacity and multiplies physical exposure

From the installed base (90.8 GW) to the Abundance Without Boundaries portfolio (386 GW), capacity grows by a factor of 4.3. Physical risk grows by a factor of 11—which yields an amplification of $\times 2.6$ per GW added. Three mechanisms drive this non-linearity.

First, *geographic concentration*: AI facilities cluster in regions combining cheap power with fast permitting—the US South, the US East Coast, parts of China and Southeast Asia—the same regions with elevated physical climate exposure. The installed base is geographically diverse; AI additions concentrate risk. Second, *power density*: AI racks at 30–100+ kW generate more waste heat per square meter, so that a cooling system failure at a 5 kW/rack legacy facility is an inconvenience, while at a 60 kW/rack AI cluster it triggers thermal shutdown within minutes. Third, *supply chain criticality*: AI training runs are non-interruptible workloads where a single disruption can invalidate weeks of compute, creating losses that far exceed the downtime itself.
Under the Abundance Without Boundaries scenario, several regional portfolios approach conditions where ClimVaR exceeds V_0 —that is, where climate-adjusted residual value approaches zero or turns negative under tail conditions. Europe reaches 93.5% ClimVaR intensity; China 79.5%; the United States 64.3%. Only APMEA, at 33.5%, remains in moderate territory. The channels that are hardest to mitigate—physical damage, cooling stress, supply chain interruption—are the channels that grow fastest with AI scale. Carbon, the most responsive channel, shrinks to irrelevance in relative terms. The toolkit that worked for the installed base does not scale to the AI era.

Exhibit 16. Channel decomposition: installed base versus AI-tier facilities (BAU/MEAN).

Channel	Installed (%)	AI-LTG (%)	AI-SAI (%)	AI-AWB (%)
Business interruption	31%	35%	33%	33%
Physical damages	22%	29%	35%	36%
Cooling	19%	29%	25%	25%
Carbon	28%	7%	6%	6%
Heat / productivity	1%	1%	1%	1%
Physical (total)	72%	93%	94%	94%

What the $\times 2.6$ factor really means


The $\times 2.6$ amplification factor is not a single effect; it is the product of several shifts that all point in the same direction. The move from carbon-dominated to physically-dominated risk means that every additional GW of AI capacity is structurally more exposed than the last. The move from diffuse to concentrated siting means that local shocks now propagate more widely through the portfolio. The move from interruptible to non-interruptible work-

loads means that the same outage now destroys more economic value.
For investors and operators, the implication is straightforward. A climate strategy that assumes AI infrastructure can be de-risked by repeating the installed-base playbook—signing more PPAs, buying more renewable certificates, leaning harder on Scope 2 metrics—will produce a growing mismatch between where exposure lies and where adaptation budget is spent. Physical resilience, not energy sourcing, becomes the binding constraint.

Finding 3: AI Amplification

Across all AI scenarios, proactive adaptation bends but does not erase the risk curve

Exhibit 17. ClimVaR trajectories by AI growth scenario, with and without proactive adaptation (installed base + AI-tier).



Baseline (solid) vs. Proactive adaptation (dashed) under three AI growth scenarios: Limits to Growth, Sustainable AI, Abundance Without Boundaries. Values cumulative over asset lifetimes.

The AI build-out does not replicate the installed base's risk profile---it inverts it. Net-zero commitments address 28% of the exposure on legacy assets but less than 7% on AI infrastructure. Physical resilience, not energy sourcing, is the binding constraint.

A single starting point, three diverging futures

Exhibit 17 shows how the combined portfolio evolves once AI growth is layered on top of the installed base. All scenarios start from US\$388B of ClimVaR on 90.8 GW and diverge sharply by 2050: about US\$1.0 trillion under Limits to Growth, US\$1.7 trillion under Sustainable AI, and US\$3.7 trillion under Abundance Without Boundaries, with proactive adaptation bending each curve downwards by hundreds of billions.

Two features matter. The curves are convex: most of the increase in ClimVaR occurs between 2030 and 2045, when AI capacity scales fastest. And as scenarios become more aggressive, the gap between Baseline and Proactive widens: in Abundance Without Boundaries, proactive measures avoid over US\$1.1 trillion of cumulative ClimVaR by 2050, yet about US\$2.5 trillion of residual exposure remains,

underscoring that adaptation cannot restore a pre-AI risk profile.

A second message is methodological. The $\times 2.6$ amplification factor mixes physical amplification (higher power density, more climate-exposed sites, concentrated supply chains) and a pure valuation effect (US\$28–42M per MW versus US\$9M for the installed base). The Net Zero Transition pathway illustrates the stakes: carbon costs on the installed base multiply by 6.2, raising total ClimVaR from US\$388B (BAU) to US\$845B (NZE, 83%), so transition-only strategies (PPAs, grid decarbonization, Scope 2 reduction) can increase total exposure if not paired with physical resilience. Finding 4 shows how a channel-by-channel, region-by-region adaptation portfolio reshapes these trajectories and where the limits of that reshaping lie.

Finding 4: The Adaptation Dividend

Adaptation always pays: net exposure falls by 30--39% _____

The first three findings quantify a problem; this fourth measures what happens when operators act. Three strategies are compared: **Baseline** (no climate investment), **Selective** (one champion action per channel on new builds), and **Proactive** (systematic adaptation across all assets). All figures are net of adaptation cost.

Exhibit 18. Adaptation value creation—installed base, BAU/MEAN pathway.

Region	CV Baseline	CV Selective	Δ Sel.	CV Proactive	Δ Proact.	Reduction
US	108	89	19	72	36	34%
Europe	107	81	26	60	46	43%
APMEA	62	49	13	39	23	37%
China	111	86	24	67	44	40%
Global	388	305	83	238	150	39%

Exhibit 19. The adaptation effect: US\$150 billion in net protected value, channel by channel.




Left = Baseline; right = Proactive residual. Central badges = net reduction. BAU/MEAN, installed base.

Proactive adaptation reduces ClimVaR from US\$388B to US\$238B—39% globally. Europe benefits most (43%), driven by carbon channel leverage. Carbon compresses fastest; physical channels dominate the residual. Selective cap-

tures 56% of Proactive’s value for 15–20% of CAPEX—the low-regret floor. Proactive adds legacy retrofits and geographic rebalancing, indispensable for the 60–70% of capacity operating beyond 2040.

Exhibit 20. Channel-by-channel waterfall from Baseline to Proactive.




Each step = one channel's contribution to the US\$150B reduction.

Carbon delivers the largest step. Cooling and business interruption follow through liquid cooling, microgrids, and supply chain diversification. Physical damage and heat set

the *irreducible floor*: once sited in a floodplain, hardening helps but cannot move the river. Operators who target the tallest steps first live with permanently lower residual risk.

Exhibit 21. Marginal abatement cost of climate risk: all ten actions below break-even.



All ten actions generate net positive returns. PPAs (US\$0.04/\$ saved), thermal shift, and liquid cooling at design stage are greenfield interventions where resilience is designed in, not retrofitted. The most expensive lever still costs US\$0.65 per dollar of ClimVaR reduced—well below break-even. Action costs calibrated to industry benchmarks.

Finding 4: The Adaptation Dividend

From diagnosis to strategy: scenario sensitivity and the logic of proactive action _____

What the four findings say together

Read in sequence, the four findings form a single argument. Finding 1 establishes that US\$388 billion of enterprise value is currently unpriced. Finding 2 shows that this exposure varies by a factor of two across regions and, more consequentially, that it differs in composition—so that a dollar of adaptation does not buy the same protection everywhere. Finding 3 demonstrates that AI infrastructure amplifies physical risk by a factor of 2.6 per GW while shifting the portfolio from carbon-dominated to physically-dominated exposure. Finding 4 closes the loop: proactive adaptation reduces installed-base ClimVaR by 39%, and every one of the ten measured interventions creates net value.

The question is no longer whether climate risk is material for data center portfolios. It is how fast operators convert diagnosis into capital allocation.

Adaptation as a real option on Net Zero

The value of adaptation is not fixed; it increases with climate policy ambition. Under the BAU pathway, Proactive adaptation reduces installed-base ClimVaR by about 39%. Under Net Zero Transition, the same set of actions cuts exposure by more than 50%, because the carbon channel—where PPAs and energy efficiency are most effective—accounts for a larger share of total risk when carbon prices rise. An operator who signs a long-term PPA in 2026 is not predicting a specific carbon price in 2040. That operator is purchasing upside in high-ambition worlds while still earning positive returns if policy remains weaker and prices stay

lower.

This asymmetry gives adaptation the payoff structure of a real option. The cost is paid once, at current prices; the benefit scales with the stringency of future regulation. In a world where the direction of climate policy is clear but the speed is not, this convexity is worth more than a point forecast.

Matching strategy to portfolio: Selective versus Proactive

The choice between Selective and Proactive depends primarily on portfolio composition and transition speed. For hyperscalers building mostly new capacity, Selective captures the majority of the adaptation dividend because it targets greenfield assets where marginal cost is lowest and design choices are still reversible. Selective achieves 56% of Proactive's total risk reduction for 15 to 20% of capital expenditure—the low-regret floor.

For colocation providers and enterprise operators with large installed bases, Proactive is required to address the 60 to 70% of capacity that Selective leaves untouched and that will still be operating into the 2040s. Every year of delay converts a greenfield opportunity—where resilience is a line in the design brief—into a future retrofit where it becomes a construction project with higher cost, more downtime, and narrower technical options. The marginal cost curve confirms that even retrofit actions remain below break-even, but the gap between designing in and bolting on is the gap between US\$0.04 and US\$0.65 per dollar of risk saved. Timing is not a secondary consideration; it is the primary determinant of adaptation economics.

All ten adaptation levers create net value. The cheapest resilience is designed in, not retrofitted. Every year of delay narrows the technical options and raises the cost per dollar of risk reduced.

Three Case Studies

Translating aggregate patterns into site-level decisions

Three prospective AI data centers produce ClimVaR intensities ranging from 37.6% to 167%, a 4.5× dispersion at site level. The Shapley decomposition and Leontief analysis produce their most concrete results at this scale. Site A (US) and Site B (France) are representative facilities; Site C

(Southern China) is a representative 100 MW allocation reflecting projected AI deployment in the region—its results are indicative of the regional risk profile rather than a specific site assessment.

Exhibit 22. Three prospective AI facilities: site characteristics and ClimVaR decomposition (BAU/MEAN).

	Site A (US South)	Site B (France)	Site C (S. China)
Operator / Project	Major AI hyperscaler	Sovereign AI project	Major coastal delta
Capacity	2,800 MW	100 MW	100 MW
Enterprise value (V ₀)	\$1.9B	\$2.1B	\$1.4B
PUE	1.15	1.12	1.28
Business interruption	\$384M (43%)	\$374M (11%)	\$276M (53%)
Cooling	\$336M (37%)	\$459M (13%)	\$96M (18%)
Physical damage	\$143M (16%)	\$2,678M (75%)	\$29M (6%)
Carbon	\$28M (3%)	\$42M (1%)	\$113M (22%)
Heat productivity	\$11M (1%)	\$6M (<1%)	\$4M (1%)
Total ClimVaR	\$902M	\$3,560M	\$517M
Intensity	46.5%	167%	37.6%

V8 model results. BI includes direct + supply chain (Leontief). Enterprise values at facility level. Site C is a probabilistic allocation; results are indicative before site-specific due diligence. The same table structure can be replicated for any facility once site-level financial, hazard, and engineering data are available.

Site A (US) — The Supply Chain Problem

The largest announced AI facility carries 46.5% ClimVaR intensity. Business interruption (43%) and cooling (37%) dominate; carbon and heat are marginal. The Shapley decomposition shows 56% of BI risk comes from tropical wind propagating through semiconductor supply chains. A facility in the US South derives most BI exposure from cyclones thousands of kilometers away. Direct on-site hazards: <1% of BI.

Site B (France) — The Flood Trap

A temperate European site carries 167% ClimVaR intensity—risk exceeds enterprise value. River flood (\$2.7B, 75%) dominates. Without granular attribution, generic assessments classify this as low-risk. The 167% is a screening signal: flood exposure merits engineering assessment with site-

specific data on elevation, defenses, and drainage.

Site C (Southern China) — The Triple Threat

Site C’s 37.6% intensity masks complexity: BI (53%), carbon (22%), cooling (18%) form a three-pronged profile. No single lever reduces more than a quarter of exposure. Effective risk reduction requires a *multi-channel* package: microgrids for BI, PPAs for carbon, high-efficiency cooling for thermal loads.

From portfolio to site-level decisions


The same ClimVaR, Shapley, and Leontief machinery applies at site level. For Site B, local flood defense data can recalibrate the 167% signal. For Site A, supply chain audits refine Leontief coefficients. The framework answers: *which channel dominates, which hazards drive it, which interventions work fastest?*

Exhibit 23. Hazard attribution (Shapley) and propagation pathway (Leontief) at three sites.

	Site A (US South)	Site B (France)	Site C (S. China)
BI: hazard attribution (Shapley)			
Tropical wind / cyclone	56%	52%	60%
Storm (convective)	24%	22%	27%
Drought	15%	23%	8%
BI: propagation pathway (Leontief)			
Direct operations	<1%	2%	2%
Supply chain (indirect)	>99%	98%	98%
Damage: dominant hazard			
	Fire (97%)	River flood (100%)	Fire (79%)

Shapley: 32 coalitions per facility. Leontief: OECD MRIOT 2023.

Exhibit 24. Three sites, three risk profiles: channel decomposition at facility level.



Radar overlay: each polygon is one facility's channel fingerprint (% of site-level ClimVaR). Site A stretches toward BI and Cooling; Site B collapses onto Damage; Site C is balanced. Overlaying profiles makes the contrast between risk architectures immediately visible.

Discussion

Interpretation and systemic implications

The \$388 Billion Baseline in Context

The US\$388 billion ClimVaR on the installed base (38.3% of enterprise value) can be situated against three reference points. The WEF/S&P aggregate of US\$3.3 trillion by 2055 operates at a different analytical level, but the order of magnitude is consistent when discounting and scope differences are accounted for (World Economic Forum and S&P Global, 2025). XDI covers comparable facility counts but reports structural damage metrics rather than enterprise value loss; our framework extends this by adding cooling stress, carbon pricing, and supply chain propagation—channels that account for 65–75% of total ClimVaR (XDI, 2025). Dietz et al. (2016) estimated 1.8% of global financial assets at risk; the 38% intensity for data centers reflects the sector's higher exposure through energy costs, uptime requirements, and supply chain concentration.

The \$388B figure represents a lower bound on total economic exposure. The framework quantifies private financial risk but does not capture systemic externalities: grid strain from concentrated AI load, water resource competition between data centers and agriculture, community displacement from extreme events, or the cascading economic effects of simultaneous service interruptions across interconnected facilities. These excluded channels would increase, not decrease, the total.

The gap between 38% enterprise value erosion and zero in standard DCF valuations quantifies the mispricing. Concurrent work in portfolio management confirms this gap from the asset allocation side: Luciani et al. (2026) show that physical climate risk lacks the standardized anchor metric equivalent to carbon intensity, and that reducing portfolio-level physical exposure beyond 10% destabilizes diversified portfolios—the facility-level ClimVaR framework developed here is designed to provide this missing metric.

The Supply Chain Multiplier

The $\times 11$ ratio of supply chain to direct business interruption is the most consequential finding for risk management practice. At all three case study sites, 98–99% of business interruption propagates through supply chains

rather than from direct on-site hazards. Taiwan accounts for over 60% of advanced semiconductor fabrication; large power transformers carry 12–24 month lead times; high-bandwidth memory production concentrates in South Korea. A single typhoon disrupting TSMC operations would propagate through every data center in the global portfolio via the Leontief inverse (Verdantix, 2024).

The Leontief framework assumes fixed technical coefficients—firms cannot reroute sourcing under stress. Long-duration disruptions are partially attenuated by adaptive behavior not captured in the model. The $\times 11$ multiplier is an upper bound for prolonged disruptions and a reasonable estimate for acute shocks. More fundamentally, the approach relies on aggregate multi-regional input-output tables rather than on the actual supply chain of each facility. Where facility-specific procurement data are available—bills of materials, qualified vendor lists, logistics routing—they should replace the MRIOT coefficients, yielding sharper estimates of both the multiplier and its geographic decomposition. The present results establish the order of magnitude and the direction; facility-level supply chain audits are the natural next step.

Physical Risk is Systemic

The residual 61% of ClimVaR that proactive adaptation cannot eliminate is not a failure of the framework—it reflects a structural reality. Luciani et al. (2026) reach an analogous conclusion from the portfolio management side: “physical risk is largely systemic, and systemic risks are difficult—if not impossible—to diversify.” Their analysis shows that constraining physical risk exposure by even 10% reduces investable portfolios to fewer than 200 stocks, with tracking error volatility rising by an order of magnitude compared to equivalent carbon constraints. Our facility-level results and their portfolio-level results converge on the same conclusion: the fraction of physical climate risk that individual actors can manage is bounded, and the remainder requires collective instruments—public risk-sharing, infrastructure stress tests, and regulatory frameworks calibrated to actual financial exposure.



Chapter IV

Conclusions

Key Messages

Key message 1: The scale of exposure

The installed base of 8,572 data centers carries US\$388 billion in climate-related value at risk—38% of enterprise value. With prospective AI infrastructure, total exposure rises to \$1.0–3.7 trillion depending on growth scenario. This is the first sector-wide quantification of climate risk as enterprise value loss for digital infrastructure. The gap between 38% enterprise value erosion and zero in standard DCF valuations quantifies the mispricing. Correcting it requires a metric. ClimVaR provides one.

Key message 2: Geography determines the risk, not scale

A twofold dispersion in risk intensity (28% in the US to 56% in China) and distinct channel signatures by region mean that a single global adaptation strategy systematically underinvests in the dominant risk channel. Business interruption dominates in the US (44%), carbon pricing in Europe (43%), physical damage in Asia-Pacific (39%), and China carries compound exposure across three channels. Two identical facilities, operated by the same company, face different climate-adjusted valuations depending on where they sit. Channel attribution is not an analytical refinement—it is the precondition for any adaptation strategy that does not systematically misallocate capital.

Key message 3: AI infrastructure inverts the risk profile

AI-era infrastructure does not replicate the installed base's risk profile—it amplifies it. On the installed base, carbon costs represent 28% of ClimVaR and physical channels 72%. On AI-tier facilities, physical channels account for 94% while carbon drops to 6%. Net-zero commitments address the carbon component but leave 82–88% of AI-era risk untouched. Each additional GW of AI capacity amplifies physical exposure through three mechanisms: geographic concentration in climate-vulnerable locations, higher power density per rack, and supply chain dependence on a narrow set of component suppliers. Sustainability strategies centered on renewable procurement address less than 10% of this exposure. The other 90% requires a different strategy: cooling architecture, geographic diversification, and supply chain redundancy.

Key message 4: Adaptation pays for itself

Proactive adaptation reduces aggregate ClimVaR by 30–39% with positive net present value under every scenario tested. All ten adaptation actions fall below the break-even line. Carbon is the most responsive channel (–63% via PPAs). But the residual 61% defines the limit of facility-level action: extreme physical events, supply chain concentration (geographically concentrated semiconductor fabrication, power transformer lead times), and thermodynamic limits on cooling in the hottest climates cannot be engineered away. Every year of delayed adaptation converts greenfield opportunity into future retrofit cost, permanently.

Implications for Stakeholders

Operators and system planners

ClimVaR provides a metric compatible with the DCF models that govern investment decisions. The $2\times$ dispersion between regions makes site selection a risk management decision, not merely a cost optimization. The Shapley decomposition identifies the highest-return adaptation lever at each site: identical capital deployed at Site A/US (supply chain dominated), Site B/France (flood trapped), and Site C/China (triple threat) generates fundamentally different risk reductions—invisible without channel attribution. Every data center built today embeds assumptions about climate futures. The framework makes those assumptions explicit and their financial consequences quantifiable.

Investors and insurers

The installed base's loss experience understates the prospective portfolio's risk profile. The $\times 2.6$ physical risk amplification on AI infrastructure and the $\times 11$ supply chain multiplier explain why facility-level loss ratios underestimate systemic exposure—at every site examined, 98–99% of business interruption originates upstream, not on site. The Shapley + Leontief decomposition offers a structure for differentiating premiums by channel composition rather than by geography alone. A portfolio with disclosed ClimVaR of 38% should command a different price from one at 15%—yet no disclosure framework currently enables this distinction. Building that framework is now a competitive advantage for the institutions that move first.

Policymakers and regulators

Data centers have shifted from commercial real estate to strategic infrastructure, but their climate exposure remains weakly integrated into permitting, prudential oversight, and critical-infrastructure policy. The 38% intensity on over

\$1 trillion of installed assets exceeds the scope of most national infrastructure protection programs. The concentration of AI investment in climate-vulnerable locations—the US South, the US East Coast, southern China, Southeast Asia—represents a market failure requiring climate-adjusted siting rules, infrastructure stress tests, and public risk-sharing mechanisms calibrated to the actual financial exposure. The analytical basis for those mechanisms now exists. Integrating it into regulatory practice is the next step.

AI developers

The sustainability agenda for AI infrastructure has been built around the wrong constraint. Physical resilience, not energy sourcing, is what determines whether a \$10 billion training cluster survives a climate event. The structural inversion from carbon-dominated to physically-dominated risk means that sustainability strategies focused on renewable procurement address less than 10% of AI-era climate exposure. Cooling architecture, geographic diversification, supply chain redundancy, and operational continuity planning are the levers that matter most for protecting multi-billion dollar training investments from climate disruption.

Chief sustainability officers

Climate risk reporting that stops at carbon emissions misses three quarters of the exposure. CSOs overseeing data center portfolios face a disclosure gap: current frameworks (CDP, TCFD, CSRD) emphasize transition risk metrics—Scope 1/2/3 emissions, renewable energy share, carbon intensity—while physical risk remains unquantified. The ClimVaR framework provides the missing metric. Integrating it into sustainability reporting would align disclosed risk with actual financial exposure—and position early adopters ahead of the disclosure curve, before regulators make the choice for them.

Limitations

The \$388 billion figure represents a lower bound on total economic exposure. Several factors not captured in the model would increase it.

Climate projection uncertainty

Physical climate projections carry inherent uncertainty at facility-specific spatial scales. The scenario structure (RCP 2.6/8.5 $\times$ NZT/BAU) spans a wide portion of the assessed range, but tail risks—tipping point cascades, correlated infrastructure failures, compound events exceeding historical analogs—may exceed modeled bounds. The January 2026 US East Coast winter storm illustrates the type of correlated event the framework captures directionally but cannot predict with precision.

Archetype-based calibration

All facilities within a region and tier share financial parameters (PUE, PPA factor, WACC, cost ratios), reflecting industry benchmarks from JLL, Synergy Research, CBRE, and IEA. This does not capture facility-specific characteristics: construction materials, elevation, cooling system vintage, backup generator capacity. The framework is most reliable for portfolio-level and regional comparison; individual facility estimates carry uncertainty proportional to their divergence from the archetype.

Prospective infrastructure

The prospective facilities (615–2,939 depending on growth scenario) are based on announced, planned, and specula-

tive projects. Actual buildout will differ in timing, scale, and location. ClimVaR estimates for prospective infrastructure should be read as indicative of the risk profile under each scenario, not as point predictions.

Adaptation modeling

The Selective and Proactive plans apply uniform reduction rates by channel. In practice, adaptation effectiveness varies by facility age, operator capability, and local regulatory context. The $\times 2.6$ amplification finding combines genuine physical amplification with a valuation effect (AI facilities carry higher asset value per MW); separating these two components matters for adaptation strategy, since actions that reduce physical exposure differ from those that address valuation premiums.

Scope of risk

The framework quantifies private financial risk through five channels. It does not capture systemic externalities: grid strain from concentrated AI load, water resource competition, community displacement, or cascading economic effects of simultaneous service interruptions across interconnected facilities.

Nor does it capture several risk categories that are likely to become material over the asset lifetime. **Nature-related**

risk: water stress, biodiversity constraints on siting permits, and ecosystem degradation around cooling water sources. **Reputation and litigation risk:** public opposition to data center siting, environmental litigation, and greenwashing claims against operators whose sustainability reporting omits physical exposure. **Technology risk:** obsolescence of cooling architectures, stranded assets from chip architecture shifts, and lock-in effects from early design choices. **Investor sentiment risk:** portfolio-level repricing as climate disclosure frameworks mature and physical risk metrics become standardized. Firm-level ClimVaR outputs will likely remain proprietary—granular site valuations carry material disclosure risk and competitive sensitivity. Yet capital markets will progressively internalize physical climate exposure through rising risk premia, rating adjustments, and stranded-asset repricing, regardless of voluntary disclosure. The question is not whether repricing occurs, but whether it proceeds in an orderly or disruptive fashion. **Electricity price risk:** the feedback loop between data center demand growth, grid congestion, and wholesale electricity prices—particularly in markets where AI load growth outpaces generation capacity. Each of these channels would increase, not decrease, the total exposure estimate.

Future Research

Several extensions would strengthen and generalize the analysis.

Empirical calibration

The January 2026 US East Coast winter storm provides a natural experiment for calibrating the β channel against observed losses. Systematic collection of post-event operational and financial data—downtime duration, revenue impact, insurance claims, recovery costs—would enable model validation against real-world outcomes. Collaboration with operators and insurers would be required.

Water-energy-climate nexus

Liquid cooling reduces energy consumption but increases water consumption. In water-stressed regions—the US Southwest, parts of India, northern China—this substitution may exchange one climate vulnerability for another. Integrating water availability projections with cooling technology trajectories would close a gap in the current analysis and is particularly relevant as AI facilities adopt liquid cooling at scale.

Real options for adaptation timing

The current framework treats adaptation as a binary decision (invest or not). A real-option extension would value the flexibility to delay while preserving the ability to act—relevant for legacy facilities where retrofit timing is uncertain and technology is evolving rapidly. The 15–20% versus 40–60% CAPEX differential between greenfield and retrofit suggests significant option value in early commitment.

Expected Shortfall and tail risk metrics

The current analysis reports ClimVaR at Q25, MEAN, and Q95 percentiles. Extending this to Expected Shortfall (CVaR)—the conditional expectation of losses beyond a given quantile—would provide a risk metric better suited to insurance pricing and prudential regulation, where tail losses dominate decision-making. The mathematical foundations are developed in Appendix A.

Cross-sector generalization

The combination of entity-level ClimVaR, Shapley channel attribution, and Leontief supply chain propagation provides a transferable methodology. Manufacturing, energy, and transportation infrastructure face analogous challenges of long-lived capital assets under climate uncertainty. Extension to multi-sector infrastructure portfolios—particularly at points of interdependency (data centers $\times$ power grids $\times$ telecommunications)—would capture systemic risks that sector-specific analysis underestimates.

System Dynamics Modeling

The current model compares outcomes across fixed scenario endpoints. A natural extension would make the pathways themselves time-varying: carbon prices responding to policy signals, operator strategies adapting to observed losses, and MRIOT coefficients reflecting supply chain restructuring over time. Dynamic multi-regional input-output modeling along these lines is being developed in parallel work (Beaufils, 2025).

Appendices

The background of the page is a blurred photograph of a server room. It shows multiple rows of server racks stretching into the distance. The racks are filled with electronic equipment, and numerous small, colorful lights (red, green, yellow, blue) are visible, indicating active systems. The perspective is from a low angle, looking down the aisle between the racks, creating a sense of depth. The overall color palette is dominated by cool blues and greys, with the warm colors of the indicator lights providing contrast.

Appendix Overview

References		<i>253 entries</i>
	Bibliography	Full reference list in APA format
Terminology		<i>Key definitions</i>
	Glossary	Abbreviations, risk channels, scenario labels, financial metrics
Appendix A — Complete ClimVaR Methodology		<i>13 sections</i>
A.1–4	Foundations	Notation, financial concepts, climate risk taxonomy, MRIOT framework
A.5	ClimVaR model	Entity, company, portfolio, site, product-level formulas
A.6–7	Loss terms	Physical ($\beta, \delta, \xi, \gamma$ + Shapley), Transition (CC), Nature (ρ)
A.8	Scenarios	Physical $\times$ transition, intra-scenario uncertainty, Expected Shortfall
A.9	Adaptation	Formal treatment of adaptation within the framework
A.10–13	Extensions	Reputation/litigation risk, spatial estimation, ES derivation
Appendix B — Data and Hypotheses		<i>7 sections</i>
B.1	Model architecture	17 input components, version history, data flow
B.2	Facility database	8,572 installed + 2,939 prospective; 120 countries; 64 attributes
B.3	Financial calibration	29 parameters, asset valuation matrices, 20 markets
B.4	Physical risk data	CMIP at facility coordinates; 4 variables; 5 hazard types; $\sim$ 50M obs
B.5	Supply chain	OECD MRIOT (75 countries $\times$ 45 sectors); 8 components
B.6	Output structure	Aggregate (288×14), annual (288×305), sites (72×592)
B.7	Confidence	48% high / 34% mixed / 17% expert; sensitivity bounds
Appendix C — Scenario Trajectories		<i>4 sections</i>
C.1	Efficiency	PPA factor, PUE, carbon ratio trajectories by region (2024–2050)
C.2	Carbon prices	BAU and NZT by region ($\$/\text{tCO}_2$, 2025–2050)
C.3	Adaptation	Reduction rates by channel, strategy definitions
C.4	Risk modeling	Cross-reference to A.6–A.9
Appendix D — Systematic Literature Review		<i>3 sections</i>
D.1	Search strategy	Semantic search + structured screening across three research questions
D.2	Results	Physical risk (Q1), insurance (Q2), valuation frameworks (Q3)
D.3	Gap analysis	Coverage mapping against 9 dimensions; three structural gaps

References

- T. Adenot, M. Briere, P. Counathe, et al. Cascading effects of carbon price through the value chain, 2022.
- Allianz Commercial. The data center construction boom: Risks and insurance implications. Technical report, Allianz Commercial, September 2025.
- P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.
- ASHRAE. Cooling degree days methodology (base 18c), 2024.
- Associated Press. Memories of deadly winter blackouts haunt the texas data center boom. *The Spokesman-Review*, August 2025.
- G. Azzone et al. Physical climate risk assessment for infrastructure assets, 2024.
- J.-N. Barrot and J. Sauvagnat. Input specificity and the propagation of idiosyncratic shocks in production networks. *Quarterly Journal of Economics*, 131(3): 1543–1592, 2016.
- L. Baumstark, N. Bauer, F. Benke, et al. REMIND2.1: Transformation and innovation dynamics of the energy-economic system, 2021.
- T. Beaufils. Dynamic multi-regional input-output modelling for climate risk propagation. Working paper, in preparation, 2025.
- M. Bernhofen et al. A first global flood hazard assessment at facility level, 2022.
- B. Caldecott, N. Howarth, and P. McSharry. Stranded assets in agriculture, 2013.
- CDP. Climate change questionnaire and scoring methodology, 2024.
- CISL. Handbook for nature-related financial risks, 2021.
- J. Connell. PJM’s updated 20-year forecast continues to see significant long-term load growth. PJM Inside Lines, 2026. Record winter peak 143,336 MW on Jan 22, 2025; 10-year growth now 3.6%/year.
- J. D. Cummins. Multi-period discounted cash flow rate-making models in property-liability insurance. *Journal of Risk and Insurance*, 57(1):79–109, 1990.
- A. Damodaran. Investment valuation (3rd ed.), 2012.
- A. Damodaran. Country risk premiums, 2024a.
- A. Damodaran. Equity risk premiums, 2024b.
- Data Center Dynamics. Hurricane sandy: How data centers weathered the storm. Data Center Dynamics, 2012. Accessed January 2026.
- Data Center Dynamics. Over 6% of data centers face high physical climate risk, report finds. *Data Center Dynamics*, January 2025.
- Data Center Knowledge. Most texas data centers weathered the storm. it didn’t go smoothly. *Data Center Knowledge*, March 2021.
- R. Dawson, D. Thompson, D. Johns, R. Wood, G. Darch, L. Chapman, P. N. Hughes, G. Watson, G. Shercliff, and J. Hall. A systems framework for national assessment of climate risks to infrastructure. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376 (2121):20170298, 2018.
- F. de Maximy, V. Pouderoux, and T. Le Guenedal. Modeling direct and indirect climate-related physical risks, 2024.
- Deloitte. Comparison of significant sustainability-related reporting requirements, 2025.
- Deloitte. Can us infrastructure keep up with the ai economy? Technical report, Deloitte Center for Energy and Industrials, June 2025.
- E. Di Febo. Transition risk in climate change: A literature review, 2025.
- J. P. Dunne, R. J. Stouffer, and J. G. John. Reductions in labour capacity from heat stress, 2013.
- T. Epelbaum. Business interruption days methodology (ClimVaR framework), 2025a.
- T. Epelbaum. ClimVaR: Discounted Cash Flow approach on climate risks. Technical report, Schneider Electric Advisory Services, Climate Data Analytics, 2025b.
- Equinix. Global tech trends survey 2025. Technical report, Equinix, 2025.
- European Central Bank and EIOPA. Policy options to reduce the climate insurance protection gap. Technical report, ECB/EIOPA, April 2024.
- European Commission. Sustainability-related disclosure in the financial services sector (SFDR), 2021a.
- European Commission. European climate law (regulation (EU) 2021/1119), 2021b.
- European Commission. Omnibus simplification package for sustainability reporting, 2025.
- EXIOBASE. Multi-regional environmentally extended supply-use and input-output database v3, 2022.
- V. Eyring et al. Overview of cmip6. *Geoscientific Model Development*, 9:1937–1958, 2016.
- FDIC. Statement of principles for climate-related financial risk management, 2024.
- G. Forzieri, A. Bianchi, F. B. e. Silva, M. Marín Herrera, A. Leblois, C. Lavalle, J. C. Aerts, and L. Feyen. Escalating impacts of climate extremes on critical infrastructures in Europe. *Global Environmental Change*, 48:97–107, 2018. Projects infrastructure damages 10x present levels by end of century.
- Gallagher Re. Natural catastrophe and climate report. Gallagher Re, 2026. January 2026.
- M. Gault and J. Williams. We mapped the world’s hottest data centers. *Rest of World*, December 2025.
- S. Giglio, T. Kuchler, J. Stroebe, and O. Wang. Nature and biodiversity loss: A research agenda for financial economics, 2024.
- Hexatronic. Hurricane season 2024: Impact on data center and fiber infrastructure. Hexatronic Group, 2025.
- Hogan Lovells. France reforms natural catastrophe insurance regime. *Hogan Lovells Insurance Hub*, January 2025.
- E. Howland. Data centers were 40% of PJM capacity costs in last auction: Market monitor. *Utility Dive*, January 2026. Data center load accounted for \$6.5 billion of \$16.4 billion in auction costs.
- IEA. Electricity 2024. Technical report, International Energy Agency, 2024.
- IEA. Energy and AI. Technical report, International Energy Agency, April 2025. URL <https://www.iea.org/reports/energy-and-ai>. Special Report.
- IEEFA. Projected data center growth spurs PJM capacity prices by factor of 10. Technical report, Institute for Energy Economics and Financial Analysis, 2025. Capacity prices increased from \$28.92 to \$329.17/MW-day.
- IIASA. Global biosphere management model (GLOBIOM) documentation, 2023a.
- IIASA. MESSAGE-GLOBIOM integrated assessment model, 2023b.
- S. Y. In, J. Weyant, and B. Manav. Pricing climate-related risks of energy investments. *Social Science Research Network*, 2020.

- S. Y. In, B. Manav, C. Venereau, L. E. Cruz Rodríguez, and J. Weyant. Pricing climate risks of energy investments: A comparative case study. *Energy Economics*, 2021.
- S. Y. In, B. Manav, C. Venereau, L. E. Cruz R., and J. Weyant. Climate-related financial risk assessment on energy infrastructure investments. *Social Science Research Network*, 2022.
- IPCC. AR6 WG1 chapter 11: Weather and climate extreme events, 2021a.
- IPCC. AR6 WG1 SPM figure SPM.6 dataset, 2021b.
- IPCC. AR6 WG2 chapter 16: Key risks across sectors and regions, 2022a.
- IPCC. AR6 WG3 annex III: Scenarios and modelling methods, 2022b.
- ISO. Wet-bulb globe temperature for occupational heat stress (ISO 7243), 2017.
- A. Jean, T. Roncalli, and A. Usardi. Natural capital and economic growth, 2025.
- P. Jorion. *Value at Risk: The New Benchmark for Managing Financial Risk*. McGraw-Hill, 3rd edition, 2006.
- G. Karagiannakis, M. Panteli, and S. Argyroudis. Fragility modeling of power grid infrastructure for addressing climate change risks and adaptation. *WIREs Climate Change*, 15(2):e866, 2024.
- R. Kerkhofs, M. V. Bernhofen, M. Borsuk, M. Baer, N. Ranger, W. Schoutens, and G. Shrimali. An asset-level analysis of financial tail risks under extreme weather events. *Environmental Research: Climate*, 4(1):015001, 2025.
- T. Kjellstrom et al. Heat index methodology for occupational health, 2016a.
- T. Kjellstrom et al. Heat, human performance, and occupational health, 2016b.
- E. Koks, D. Le Bars, A. Essenfelder, S. Nirandjan, and P. Sayers. The impacts of coastal flooding and sea level rise on critical infrastructure: A novel storyline approach. *Sustainable and Resilient Infrastructure*, 8(sup1):237–255, 2022.
- S. Kolaric. The impact of climate litigation and activism on stock prices, 2024.
- KPMG. Guide to annual financial statements: IFRS and US GAAP, 2024.
- L. Kruschwitz and A. Loeffler. Discounted cash flow: A theory of the valuation of firms, 2006.
- T. Lorans, J. Priol, and V. Bouchet. Beyond carbon price: Scenario-based quantification of portfolio financial loss, 2025.
- A. Luciani, E. Music, and T. Mertens. From transition to physical risk: Rethinking portfolio management. Technical report, Amundi Institute, 2026.
- T. Luo, Y. Cheng, J. Falzon, J. F. Kölbel, L. Zhou, Y.-L. Wu, and A. Habchi. A framework to assess multi-hazard physical climate risk for power generation projects from publicly-accessible sources. *Communications Earth & Environment*, 4:117, 2023.
- D. H. Meadows, D. L. Meadows, J. Randers, and W. W. Behrens. The limits to growth, 1972.
- R. E. Miller and P. D. Blair. Input-output analysis: Foundations and extensions, 2009a.
- R. E. Miller and P. D. Blair. *Input-Output Analysis: Foundations and Extensions*. Cambridge University Press, Cambridge, UK, 2nd edition, 2009b.
- MSCI. Climate value-at-risk methodology, 2023.
- NERC. 2025–2026 winter reliability assessment. Technical report, North American Electric Reliability Corporation, November 2025.
- NGFS. Climate scenarios for central banks – phase II, 2021.
- NGFS. Climate scenarios for central banks – phase III, 2022.
- NGFS. Ngfs scenarios for central banks. Technical report, Network for Greening the Financial System, 2023.
- NGFS. Climate scenarios – technical documentation v5.0, 2024.
- W. D. Nordhaus. Projections and uncertainties about climate change (DICE-2016R2), 2017.
- NRDC and Citizens Utility Board. Rising demand from data centers driving reliability, cost concerns. Technical report, Natural Resources Defense Council, October 2025. Data centers could require 50+ GW peak capacity by 2030 in PJM.
- OECD. Multi-regional input-output tables (MRIOT), 2023a.
- OECD. Trade in value added (TiVA) database, 2023b.
- E. Oughton, T. Russell, J. Oh, S. Ballan, and J. W. Hall. Global vulnerability assessment of mobile telecommunications infrastructure to climate hazards using crowdsourced open data. *Scientific Data*, 2023. Analysis of 7.6 million telecom assets globally.
- R. Paccou and F. Wijnhoven. Artificial intelligence and electricity: A system dynamics approach. Technical report, Schneider Electric Sustainability Research Institute, 2024.
- M. D. S. Pasha, G. G. Dirgantara, S. K. Wiliandi, and E. T. Budianto. Flood damages in an infrastructure climate risk stress test: A case study of a solar power plant project. *Journal of Infrastructure Policy and Management*, 2025.
- M. K. Pasupuleti. Climate and flood risk in data centre strategy: Integrating hydrologic simulation (HEC-HMS) into location and capacity decisions. *International Journal of Academic and Industrial Research Innovations (IJAIRI)*, 2022.
- M. K. Pasupuleti. Climate and flood-aware data centre siting: A GIS and hydrologic simulation-based framework for risk-aware cloud infrastructure. *International Journal of Academic and Industrial Research Innovations (IJAIRI)*, 2024.
- M. K. Pasupuleti. GIS-driven AI for climate-resilient data centre siting and operation. *International Journal of Academic and Industrial Research Innovations (IJAIRI)*, 2025.
- D. Popp. Innovation and climate policy, 2010.
- T. Prime, K. Morrissey, J. Brown, and A. Plater. Protecting energy infrastructure against the uncertainty of future climate change: A real options approach. *Journal of Ocean and Coastal Economics*, 5(2):Article 1, 2018.
- Reuters. Power prices surge as winter storm spikes demand in U.S. data center alley. *CNBC*, January 2026. Wholesale electricity prices topped \$1,800/MWh in Virginia’s Dominion territory.
- K. Richardson et al. Earth beyond six of nine planetary boundaries, 2023.
- R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–42, 2000.
- J. Rockstrom et al. A safe operating space for humanity, 2009.
- T. Roncalli. Climate risk and asset management, 2023.
- T. Roncalli. Lecture notes on biodiversity, 2025.
- T. Roncalli and R. Semet. The economic cost of the carbon tax, 2024.

- M. Sato, G. Gostlow, C. Higham, et al. Impacts of climate litigation on firm value, 2024.
- L. S. Shapley. A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317, 1953.
- S&P Global. Biodiversity and nature-related financial risks, 2024.
- S&P Global. Geopolitics of data centers: An ai showdown that will reshape the world. Technical report, S&P Global Market Intelligence, December 2025.
- S. Stapczynski. US calls on data centers to prepare backup power ahead of storm. *Bloomberg*, January 2026. DOE emergency orders for 35 GW of backup capacity.
- Swiss Re Institute. Sigma 1/2025: Natural catastrophes – insured losses on trend to USD145 billion in 2025. Technical report, Swiss Re, 2025. URL <https://www.swissre.com/institute/research/sigma-research/sigma-2025-01-natural-catastrophes-trend.html>.
- TCFD. Final report: Recommendations on climate-related financial disclosures, 2017.
- TechTarget. Uptime institute 2024 annual outage analysis: Weather-related failures on the rise. TechTarget / Uptime Institute, 2024.
- TNFD. Recommendations of the taskforce on nature-related financial disclosures, 2023.
- TNFD. Evidence review on the financial effects of nature-related risks, 2025.
- TRG Datacenters. Lessons from hurricane harvey: Data center resilience in houston. TRG Datacenters, 2024.
- F. Tyron. Data centers: Critical infrastructure in a volatile world. *The Data Scientist*, September 2025.
- United Nations. International standard industrial classification (ISIC) rev. 4, 2008.
- Verdantix. Smart innovators: Physical climate risk solutions. Verdantix Ltd., 2024. December 2024.
- World Bank. State and trends of carbon pricing 2024, 2024.
- World Economic Forum and S&P Global. Data centers and the climate challenge: A \$3.3 trillion question. Technical report, World Economic Forum, 2025.
- WRI/WBCSD. GHG protocol: Scope 1 emissions, 2004.
- WRI/WBCSD. GHG protocol: Scope 2 emissions, 2015.
- WWF. Tackling the insurance protection gap: Leveraging climate and nature to increase resilience. Technical report, World Wildlife Fund, November 2025.
- XDI. 2025 global data centre physical climate risk and adaptation report. Technical report, Cross Dependency Initiative, 2025.
- M. Zhao et al. Global multi-model projections of local urban climates, 2021.

Terminology

Acronyms

APMEA	Asia-Pacific, Middle East, and Africa
AWB	Abundance Without Boundaries (AI growth scenario)
BAU	Business as Usual
BESS	Battery Energy Storage System
BI	Business Interruption
CAPEX	Capital Expenditures
CC	Carbon Cost (Scopes 1, 2, 3)
CDD	Cooling Degree Days
ClimVaR	Climate Value at Risk
CMIP	Coupled Model Intercomparison Project (multi-model ensemble)
COGS	Cost of Goods Sold
CSRD	Corporate Sustainability Reporting Directive (EU)
DC	Data Center
DCF	Discounted Cash Flow
ERCOT	Electric Reliability Council of Texas
ETS	Emissions Trading System
EV	Enterprise Value
GW	Gigawatt
I-O	Input-Output (Leontief analysis)
IEA	International Energy Agency
LTG	Limits to Growth (AI growth scenario)
MACC	Marginal Abatement Cost Curve
MRIOT	Multi-Regional Input-Output Tables
MW	Megawatt
NERC	North American Electric Reliability Corporation
NGFS	Network for Greening the Financial System
NPV	Net Present Value
NZT	Net Zero Transition
OPEX	Operating Expenses
PJM	PJM Interconnection (US Mid-Atlantic grid operator)
PPA	Power Purchase Agreement
PUE	Power Usage Effectiveness
RCP	Representative Concentration Pathway
SAI	Sustainable AI (AI growth scenario)
SLA	Service Level Agreement
TCFD	Task Force on Climate-related Financial Disclosures
UPS	Uninterruptible Power Supply
WACC	Weighted Average Cost of Capital

Channel Symbols

γ	Cooling cost channel — increased energy consumption from rising temperatures
$CC^{1,2,3}$	Carbon cost channels — Scope 1 (direct), Scope 2 (purchased electricity), Scope 3 (supply chain)
β	Business interruption channel — operational downtime from extreme events and grid failures
δ	Physical damage channel — structural damage from acute events (floods, storms, seismic)
ξ	Heat productivity channel — outdoor labor efficiency loss from sustained heat exposure
σ	Customer shift channel — revenue loss from customer migration to lower-carbon providers
ρ	Nature capital channel — revenue loss from ecosystem degradation (water stress, biodiversity)

Financial Notation

V_0	Baseline enterprise value (DCF, no climate adjustment)
V^{cc}	Climate-adjusted enterprise value
ClimVaR	$V_0 - V^{cc}$ — present value of climate-induced value erosion
ClimVaR%	ClimVaR / V_0 — climate risk intensity as percentage of enterprise value
$FCF^{(t)}$	Free cash flow in year t
ω	Weighted average cost of capital (WACC)
τ	Corporate tax rate
T	Explicit forecast horizon (terminal year)
AV	Asset value (book value of physical infrastructure)

Scenario Notation

RCP 2.6 / 8.5	Physical climate pathways (Paris-aligned / high emissions)
BAU / NZT	Transition policy pathways (current policies / Net Zero by 2050)
LTG / SAI / AWB	AI growth scenarios (constrained / sustainable / unconstrained)
Baseline / Selective / Proactive	Adaptation strategies (no action / new builds only / full portfolio)
Q25 / MEAN / Q95	Statistical reporting quantiles within each scenario

Key Definitions

Installed base	8,572 operational data center facilities (90.8 GW), pre-2025
Prospective	AI-era facilities planned or projected through 2050 (615–2,939 depending on scenario)
Adaptation wedge	Difference in ClimVaR between Baseline and Proactive strategies
Residual risk	ClimVaR remaining after Proactive adaptation (61% of total on installed base)
Channel fingerprint	The distinctive risk channel composition of a region or facility
Shapley value	Fair attribution of aggregate risk to individual contributing hazards
Leontief inverse	$(I - DA^T)^{-1}$ — captures direct and indirect supply chain propagation
Pass-through rate	Fraction of upstream disruption transmitted to downstream operations



Appendix A

Complete ClimVaR Methodology

Lloyd Shapley (1923–2016)

Each player receives a payoff equal to his average marginal contribution over all possible orderings of the players.

Climate risk for data centers is not a single hazard—it is a portfolio of five interacting channels. Because these channels interact non-additively (cooling stress amplifies business interruption; carbon pricing compounds with physical damage), simple summation double-counts risk.

The Shapley value from cooperative game theory provides a principled decomposition. For $n = 5$ risk channels, it evaluates all $2^5 = 32$ coalition combinations per facility, attributing each channel's marginal contribution to total ClimVaR.

The decomposition satisfies efficiency, symmetry, and the null player property.

Contents

- A.1 Complete ClimVaR methodology (Epelbaum 2025)
- A.2 Shapley value formulation for risk attribution
- A.3 Computational implementation (32 coalitions)
- A.4 Expected Shortfall estimation

ClimVaR Methodology

Introduction

Climate change is one of the central problem humanity has to tackle in the 21st century. The frequency of extreme climate events is already increasing, and this tendency can only worsen if no action is taken to limit anthropic carbon emissions. In recent years, several climate/financial regulation frameworks [European Commission \(2021a\)](#)-[Deloitte \(2025\)](#) were developed, aiming at collecting climate related data and be somewhat prescriptive [European Commission \(2021b\)](#) on how to decarbonate. Nevertheless, the current accounting norms (GAAP, IFRS...) do not even fully capture simple materiality aspects, making relevant financial decisions at public and private level difficult.

As we will see in section [A.2](#), several studies have been done in the past years on the financial impact caused by some climate risks. But most of these studies focus on one aspect of the problem: a specific transition risk [Di Febo \(2025\)](#) like carbon cost [Adenot et al. \(2022\)](#)[Roncalli and Semet \(2024\)](#) or litigation [Sato et al. \(2024\)](#)[Kolaric \(2024\)](#), nature risk [TNFD \(2025\)](#)[CISL \(2021\)](#)[Giglio et al. \(2024\)](#)[Roncalli \(2025\)](#)[Jean et al. \(2025\)](#), or physical climate risk [de Maximy et al. \(2024\)](#).

A unifying approach is therefore somewhat lacking, though most of the tools needed to get there have already been developed.

In this paper we will use a standard Discounted Cash Flow (DCF) approach to cover all the major climate risks to be faced in the years/decades to come.

Our aim is to help public and private organisms measure the economic viability of systems. Five types of systems will be studied:

- A portfolio of financial products.
- A company.
- An entity within a company.
- A site of an entity²¹.
- A company product (depending on several sites). A product is produced on several sites, and one site can produce several products.

Having access to better financial assessments of systems should allow for a better resilience of the economic actors with respect to climate change, and should ease the transformations needed to fight it: decarbonate the supply chain, build floodgates, avoid brown sectors when it makes sense...

Our objective will be reached if regulators, auditors, public and private actors can start using the white box formalism that we (though obviously building a lot on existing works) develop in the next sections. This should help them really assess the economic profitability of the projects/systems they supervise by duly accounting for climate change.

As already stated, we will adapt a standard DCF approach to estimate the future valuation V of a system. In all the following, the system can be one of the five already stated. We will augment the traditional DCF approach by considering variations related to climate. These variations can be related to

- physical risk: flood, drought, wildfire [IPCC \(2021a\)](#)[IPCC \(2021b\)](#)...
- transition risk: carbon tax, customer shift [Adenot et al. \(2022\)](#)[Roncalli and Semet \(2024\)](#)[Lorans et al. \(2025\)](#)[NGFS \(2024\)](#)..
- nature risk: loss of nature capital [Jean et al. \(2025\)](#)

We call our central DCF formula embedding climate risks ClimVaR (for Climate Value at Risk)

The rest of this note is devoted to deriving ClimVaR for the five systems already mentioned. It is organized as follows. Section [A.2](#) discusses related work, insisting on the fact that the current paper builds on existing (but siloed) studies. Section [A.3](#) regroups all notations. The tools developed in this paper are not advanced, but put together they cover a large range of concepts. The reader should therefore refer to this section to understand the approach and all symbols. Section [A.4](#) presents the central result of this work: a valuation formula encompassing physical and transition climate-related effects. Section [A.5](#) summarizes the derivation of individual loss terms, referring the reader to [Epelbaum \(2025b\)](#) for the complete mathematical treatment. Section [A.6](#) elaborates on the scenario and uncertainty topic. The link with the usual Value At Risk (VaR) is also made. The financially knowledgeable reader pondering on the use of the ClimVaR term will benefit from [A.6.2](#). Adaptation strategies and their economic merit are discussed in the main text (Chapter III, Finding 4). Extensions to additional transition risks (reputation, litigation, technology, investor sentiment) are deferred to subsequent work.

Related work

Relating climate risks to financial impacts is not a new topic, but it has seen a rise of interest in recent years.

The climate/economy relationship could be traced back to the Meadows Report [Meadows et al. \(1972\)](#), which initiated the works on Integrated assessment Models (IAMs), coupling classical macro-economic equations with exogenous (climate, but not only) economic data. Following its steps, numerous IAMs [Nordhaus \(2017\)](#) [IIASA \(2023a\)](#) [Baumstark et al. \(2021\)](#) [IIASA \(2023b\)](#) were developed in the past 50 years, adding among other things local (by opposition to world or large regions) and sectoral modelling.

Climate risks could loosely be divided into 3 different categories:

- Physical risk [IPCC \(2022a\)](#) correspond to the impact of climate events on human activity: a strong wind damaging a building, a flood interrupting business activities... Several papers [de Maximy et al. \(2024\)](#) [Epelbaum \(2025a\)](#) [Bernhofen et al. \(2022\)](#) [Azzone et al. \(2024\)](#) have investigated the impact of these physical events onto financial activities.
- Nature (biodiversity) risk is a more recent center of interest. It analyzes the effect of nature/ecosystem/biosphere degradation onto economic activities [Jean et al. \(2025\)](#) [S&P Global \(2024\)](#) [TNFD \(2023\)](#). As we are approaching several planetary limits [Rockstrom et al. \(2009\)](#) [Richardson et al. \(2023\)](#), this topic should see intense activities in the years to come.
- Transition risk represents the human adaptation to climate change, in the forms of laws, norms, bans, taxes, individual consumer behaviours... It may be the topic that have been the focus of the most attention in the past decade [Roncalli and Semet \(2024\)](#) [Caldecott et al. \(2013\)](#) [FDIC \(2024\)](#) [TCFD \(2017\)](#), but recent changes in international geopolitics [European Commission \(2025\)](#) could make this risk less pregnant than the physical one in the years to come.

Some paper on Climate Value At Risk [Roncalli \(2023\)](#) [MSCI \(2023\)](#) already exist, but they do not yet offer a unifying, quantitative, white box analysis encompassing all the risks previously mentioned. The present work aims to do just so.

Notation

Generic concepts

- t : times. In the present paper year is the time unit, unless specified otherwise.
- ν : Scenario (more on scenarios in section [A.6.1](#))
- λ : Intra scenario uncertainty (more on it and link with VaR in section [A.6.2](#)).
- r : Supply chain rank. 0 means direct operation (i.e. a system directly operated by the project leader). 1 means direct supplier (meaning supplier to which the project leader directly buys). 1+ means all suppliers (direct and indirect. Indirect suppliers are suppliers of suppliers, suppliers of suppliers of suppliers... And so on). 0+ means all contributions, direct and suppliers.

Company related concepts

We will denote

- FP : Financial Portfolio.
- C : Company (possibly part of for a financial portfolio).
- E : Entity of a company.
- S : Site of an entity. We will interchangeably use assets a .
- P : Product.
- c : World countries.
- s : Activity sectors. we take the ISIC standard [United Nations \(2008\)](#) for encoding the activity sectors, and consider the 45 ISIC sectors used in the 2023 OECD work [OECD \(2023a\)](#) [OECD \(2023b\)](#) unless stated otherwise.
- $S^{(1)}$: Carbon equivalent emissions within a year corresponding to direct system activities, called Scope 1 [WRI/WBCSD \(2004\)](#) in the litterature.
- $S^{(2)}$: Carbon equivalent emissions within a year corresponding to system electricity consumptions, called Scope 2 [WRI/WBCSD \(2015\)](#) in the litterature.

Financial concepts

We will note

- V : Valuation.
- VaR: Valuation At Risk.
- FCF : Future Cash Flow. .
- T : Terminal time value.

- g : Perpetual growth rate.
- $g^{(t)}$: Yearly growth rate.
- Y : System gross revenues.
- ω (or $WACC$): Weighted Average Cost of Capital. ω will be preferred in equations. $WACC$ variation is deferred to future work.
- DCF : Discounted Cash Flow. $DCF^{(t)} = FCF^{(t)} / (1 + \omega)^t$.
- $COGS$: Cost Of Goods Sold.
- c : $COGS$ to Revenue ratio.
- $OPEX$: Operating Expenses.
- o : $OPEX$ to Revenue ratio.
- s : Salary to Revenue ratio.
- P : Absolute (monetary) system purchases will be noted as P .
- p : system purchases ratio will be noted as p , and we consider $p = c - s$.
- so : Outdoor salary mass to Revenue ratio.
- si : Indoor salary mass is not used in the present analysis. Obviously $s = si + so$.
- e : Electricity expenses to Revenue ratio.
- ec : Electricity cooling expenses to Revenue ratio.
- τ : system (federal) tax rate.
- CFO : Cash Flow from Operation.
- $CAPEX$: CAPital Expenditures.
- NWC : for Net Working of Capital.
- DA : for Depreciation (and Amortization)²².
- k : for $CAPEX$ over Revenue ratio.
- n : for NWC over Revenue ratio.
- d : for DA over Revenue ratio.
- op : Operating Profit over Revenue Ratio. $op = 1 - c - o$.
- i : Insurance premium.
- AV : Asset Value (in the sense of hard assets: buildings...).

Climate risk related concepts

- cc : Climate change (by contrast with quantities computed without considering it).
- h : Climate hazard (drought, flood, wild fire...).
- σ : Customer shift.
- β : Business Interruption Days (BID) (as a fraction of year) at country sector level.
- χ : Business Interruption Days (BID) at asset site level.
- ρ : Degradation of renewable nature capital (as a % of revenue).
- δ : Damages on assets (as a % of asset value).
- $CC^{(i)}$: Carbon costs related to scope i .
- ξ : Heat Productivity Loss Kjellstrom et al. (2016a)ISO (2017)Kjellstrom et al. (2016b)Dunne et al. (2013)Zhao et al. (2021).
- γ : Cooling cost increase ASHRAE (2024) (from a reference to a future time period).

Multi Regional Input Output table (MRIOT) related concepts

MRIOT is a rich topic on its own, that we only minimally cover in this paper to the extent of our need.

- $Z_{cs,c's'}$: Multi Regional Input Output Table (MRIOT). Following Miller and Blair (2009a), we adopt the convention where rows of Z correspond to suppliers and columns to clients. $Z_{cs,c's'}$ is therefore a financial flux within a year²³ between sector s' of country c' , buying goods from sector s of country c .
- X_{cs} : Total financial output of sector s of country c .
- $A_{cs,c's'}$: A Technical coefficient matrix. Coefficients as $A_{cs,c's'} = Z_{cs,c's'} / X_{c's'}$.
- $L_{cs,c's'}$: Leontief matrix. $L = \mathbb{I} + A^T + (A^T)^2 + \dots = (\mathbb{I} - A^T)^{-1}$.

Carbon cost related concepts

We will use the following notations

²²In our approach we follow the Generally Accepted Accounting Principles (GAAP)KPMG (2024), which are for instance somewhat different from the French accounting formalism.

²³Usually in USD or M USD.

- $C_c^{(t)}$: "Raw" carbon price (as USD per tons of CO2 equivalent emissions, or TCO2e). Carbon prices [World Bank \(2024\)](#) are given at country level and are considered time dependent.
- $F_s^{(t)}$: Free allocation rates. These are sectoral allocations that allow sectors to discount part of their carbon costs. In our approach, Free allocations are always present in Business as Usual scenarios and removed after 2030 for other ones. More on scenarios in section [A.6](#).
- Passthrough P_s at sector level. Considered to be 1 for the remaining of this paper. This is a strong assumption that we will relax in future studies.
- $C_{cs}^{(t)}$: Combined "Raw" carbon price. $C_{cs}^{(t)} = F_s^{(t)} C_c^{(t)} P_s = F_s^{(t)} C_c^{(t)}$.
- $I_{cs}^{(i)}$: Carbon intensities (TCO2e per revenue (USD)) for scope i . When i is unspecified, it is for scope 1. These country sector intensities can be evaluated in several manners: having access to database of company emissions + financial data [CDP \(2024\)](#), or a MRIOT approach (like the one performed by Exiobase [EXIOBASE \(2022\)](#))...

ClimVaR model

We leave aside the notion of scenario and uncertainty for the time being. We will come back on these important topics in section [A.6](#). As is standard in valuation approaches, ClimVaR will take major financial information from income statements, balance sheets and cash flows to form the Discounted Cash Flow. ClimVaR novelty resides in how it computes and aggregates climate risk effects on different parts of the Discounted Cash Flow. As stated in section [A.1](#), we want to derive these effects for five systems (from the finer to the coarser level): products, sites, entities, companies, financial portfolios. As we will see in the following, financial portfolios can be seen as the aggregation of company valuation computations, and the same goes for entities with respect to companies and sites with respect to entities. The relation is not one to many but many to many for products to sites valuation: a product is generally produced by several sites, but sites usually also produce parts of several products at once.

We will start with the entity computation, that is the simpler to explain on its own. From it we will trivially derive ClimVaR formulas for companies and financial portfolios. We will then explain the adaptations needed on the entity formula to get the sites and finally the product ones.

Entity ClimVaR

A DCF model [Kruschwitz and Loeffler \(2006\)](#)[Damodaran \(2012\)](#) estimates the current value (at current time 0) of a system as the sum of its future Discounted Cash Flows. We first focus our attention on an entity, belonging to a sector s of a country c . Valuation reads there

$$V_{cs}^{(E)} = \sum_{t=1}^{+\infty} \frac{FCF_{cs}^{(t)}}{(1 + WACC_{cs})^t}, \quad (1)$$

where $WACC$ is assumed to be time/climate independant and equal to ω_{cs} .

It is good practice to take a terminal value year T (which we consider constant at company level), after which the system is experiencing fixed perpetual growth (Gordon Growth Model) g_{cs} .

$$V_{cs}^{(E)} = \sum_{t=1}^T \frac{FCF_{cs}^{(t)}}{(1 + \omega_{cs})^t} + \frac{FCF_{cs}^{(T)}}{(1 + \omega_{cs})^T} \frac{1 + g_{cs}}{\omega_{cs} - g_{cs}}. \quad (2)$$

To simplify the notation, we will take

$$FCF_{cs}^{(T+1)} = FCF_{cs}^{(T)} \frac{(1 + g_{cs})(1 + \omega_{cs})}{(\omega_{cs} - g_{cs})}, \quad (3)$$

so that

$$V_{cs}^{(E)} = \sum_{t=1}^{T+1} \frac{FCF_{cs}^{(t)}}{(1 + \omega_{cs})^t}. \quad (4)$$

In our model, we will be interested by the Climate Value at Risk (ClimVaR) [Kruschwitz and Loeffler \(2006\)](#)[Damodaran \(2012\)](#)[Damodaran \(2024a\)](#): namely the difference between a valuation without and with climate change. This is the central quantity that we want to estimate

$$\text{VaR}_{cs}^{(E)} = \Delta V_{cs}^{(E)} = V_{cs}^{(E)} - V_{cs}^{(E,cc)} = \sum_{t=1}^{T+1} \frac{FCF_{cs}^{(t)} - FCF_{cs}^{(t,cc)}}{(1 + \omega_{cs})^t} . \quad (5)$$

To be noted, we will also be interested by a δVAR , namely

$$\delta\text{VaR}_{cs}^{(E)} = \frac{\Delta V_{cs}^{(E)}}{V_{cs}^{(E)}} , \quad (6)$$

that estimates the share of the total value at risk²⁴.

The FCF at time t reads Damodaran (2024a)

$$FCF_{cs}^{(t)} = CFO_{cs}^{(t)} - CAPEX_{cs}^{(t)} - \delta NWC_{cs}^{(t)} , \quad (7)$$

In a first approach, we will assume that both NWC (corresponding to short term assets and liabilities cash fluxes) and $CAPEX$ do not depend on climate change (as we will see, we will consider fixed assets accelerated degradation because of climate change via the OPEX term in CFO. To be noted $CAPEX$ will play a role when taking into account climate actions), and that they represent a constant fraction of revenues Y :

$$\begin{aligned} \delta NWC_{cs}^{(t)} &= NWC_{cs}^{(t)} - NWC_{cs}^{(t-1)} \approx n_{cs} \left(Y_{cs}^{(t)} - Y_{cs}^{(t-1)} \right) , \\ CAPEX_{cs}^{(t)} &\approx k_{cs} Y_t . \end{aligned} \quad (8)$$

In that case

$$\text{VaR}_{cs}^{(E)} = \sum_{t=t_0+1}^{T+1} \frac{CFO_{cs}^{(t)} - CFO_{cs}^{(t,cc)}}{(1 + w)^t} . \quad (9)$$

CFO reads

$$\begin{aligned} CFO_{cs}^{(t)} &= NOPAT_{cs}^{(t)} + DA_{cs}^{(t)} \\ &= \left[Y_{cs}^{(t)} - COGS_{cs}^{(t)} - OPEX_{cs}^{(t)} - DA_{cs}^{(t)} \right] \left[1 - Tax_{cs}^{(t)} \right] + DA_{cs}^{(t)} , \end{aligned} \quad (10)$$

In our model, we assume that taxes are not affected by climate change²⁵ and are constant (τ_{cs}) through time. In addition, we also assume that DA is not affected by climate change (again, additional costs incurred to repair buildings will be considered as OPEX) and is a constant share of the total revenue

$$DA_{cs}^{(t)} \approx d_{cs} Y_{cs}^{(t)} . \quad (11)$$

We therefore focus our attention on the Operating profit

$$OP_{cs}^{(t)} = Y_{cs}^{(t)} - COGS_{cs}^{(t)} - OPEX_{cs}^{(t)} . \quad (12)$$

And our central climate VAR (5) becomes

$$\begin{aligned} \text{VaR}_{cs}^{(E)} &= \sum_{t=1}^{T+1} \frac{1 - \tau_{cs}}{(1 + \omega_{cs})^t} \left[OP_{cs}^{(t)} - OP_{cs}^{(t,cc)} \right] \\ &= \sum_{t=1}^{T+1} \frac{1 - \tau_{cs}}{(1 + \omega_{cs})^t} \left[Y_{cs}^{(t)} - Y_{cs}^{(t,cc)} - \left(COGS_{cs}^{(t)} - COGS_{cs}^{(t,cc)} \right) \right. \\ &\quad \left. - \left(OPEX_{cs}^{(t)} - OPEX_{cs}^{(t,cc)} \right) \right] . \end{aligned} \quad (13)$$

²⁴Keep in mind that the climate can generate value, in the case of customer shift for instance.

²⁵We will see later that Carbon Taxes are taken into account via the Revenue and the OPEX channel

We continue our modelling by assuming that $COGS$ and $OPEX$ are proportional to revenues [Damodaran \(2024b\)](#)

$$\begin{aligned} COGS_{cs}^{(t)} &\approx c_{cs} Y_{cs}^{(t)}, & OPEX_{cs}^{(t)} &\approx o_{cs} Y_{cs}^{(t)} \\ COGS_{cs}^{(t,cc)} &\approx c_{cs}^{(t,cc)} Y_{cs}^{(t,cc)}, & OPEX_{cs}^{(t,cc)} &\approx o_{cs}^{(t,cc)} Y_{cs}^{(t,cc)}, \end{aligned} \quad (14)$$

so that, recalling that

$$op_{cs} = 1 - c_{cs} - o_{cs}, \quad (15)$$

we get

$$\begin{aligned} \text{Var}_{cs}^{(E)} &= \sum_{t=1}^{T+1} \frac{1 - \tau_{cs}}{(1 + \omega_{cs})^t} \left[op_{cs} \left(Y_{cs}^{(t)} - Y_{cs}^{(t,cc)} \right) \right. \\ &\quad \left. + \left(c_{cs}^{(t,cc)} - c_{cs} \right) Y_{cs}^{(t,cc)} + \left(o_{cs}^{(t,cc)} - o_{cs} \right) Y_{cs}^{(t,cc)} \right]. \end{aligned} \quad (16)$$

Future system revenues will be taken with a time dependent cumulated growth rate $g_{cs}^{(t)}$, so²⁶

$$Y_{cs}^{(t)} = \left(1 + g_{cs}^{(t)} \right) Y_{cs}, \quad (17)$$

with Y_{cs} the current (at present time) revenues. Our model thus has currently the following unknowns:

- future revenues $Y_{cs}^{(t,cc)}$ taking into account climate change
- future COGS to revenue ratio $c_{cs}^{(t,cc)}$ taking into account climate change
- future OPEX to revenue ratio $o_{cs}^{(t,cc)}$ taking into account climate change

It is time to explicit these terms and how climate risks contribute to them.

Revenue change caused by climate effects We will consider three climate change effects on revenues:

- Customer shift (market transition risk) $\sigma_{cs}^{(t)}$
- total Business Interruption Days (BID) $\beta_{cs}^{(t)}$: from interruption within the entity and from its supply chain (physical risk).
- Loss of ecosystemic effect: degradation of renewable $\rho_{cs}^{(t)}$ nature capital (nature risk).

These terms are derived in detail in [Epelbaum \(2025b\)](#). We thus write

$$Y_{cs}^{(t,cc)} = \left(1 - \sigma_{cs}^{(t)} \right) \left(1 - \beta_{cs}^{(t)} \right) \left(1 - \rho_{cs}^{(t)} \right) Y_{cs}^{(t)}, \quad (18)$$

OPEX change caused by climate effects The only OPEX change we consider is related to damages on buildings. We take a simple model, where the insurance cost [Cummins \(1990\)](#) is increased due to future expected damages on buildings

$$o_{cs}^{(t,cc)} Y_{cs}^{(t,cc)} = o_{cs} Y_{cs}^{(t,cc)} + (1 + i_{cs}) AV_{cs} \delta_{cs}^{(t)}, \quad (19)$$

where AV_{cs} is the total (hard) assets owned by entity cs , i_{cs} is an insurance premium and $\delta_{cs}^{(t)}$ is the future damage fraction loss of the assets (see [Epelbaum \(2025b\)](#) for the full derivation). In the future we will also make use of asset value to reference revenue ratio

$$av_{cs} = \frac{AV_{cs}}{Y_{cs}} \quad (20)$$

COGS change caused by climate effects We will consider three climate change effects on COGS:

- Carbon cost, both from direct $CC^{(t1)}$ and indirect emissions: supply chain, from scope 2 and scope 3 emissions. We will respectively consider $CC^{(t2)}$ and $CC^{(t3)}$ carbon costs (policy transition risk).
- Heat Productivity loss $\xi_{cs}^{(t)}$ (physical risk).
- Cooling cost increase $\gamma_{cs}^{(t)}$ (physical risk).

To be able to apply them, we will need two specific parts of the COGS to revenue ratio

- outdoor salary mass to Revenue ratio so
- cooling costs to Revenue ratio ec

²⁶In practice We will presumably have access to yearly growth rates yg , so

$$g_{cs}^{(t)} = \prod_{t'=1}^t \left(1 + yg_{cs}^{(t')} \right) - 1$$

so that

$$c_{cs}^{(t,cc)} Y_{cs}^{(t,cc)} = c_{cs} Y_{cs}^{(t,cc)} + so_{cs} \xi_{cs}^{(t)} Y_{cs}^{(t,cc)} + ec_{cs} \gamma_{cs}^{(t)} Y_{cs}^{(t,cc)} + \sum_{i=1}^3 CC_{cs}^{(ti)} . \quad (21)$$

Combining all effects Taking into account the results of the three previous sections, our central equation becomes

ClimVaR central formula


$$\text{Var}_{cs}^{(E)} = \sum_{t=1}^{T+1} \frac{(1 - \tau_{cs})}{(1 + \omega_{cs})^t} \left[op_{cs} \left(Y_{cs}^{(t)} - Y_{cs}^{(t,cc)} \right) + Y_{cs}^{(t,cc)} \left(so_{cs} \xi_{cs}^{(t)} + ec_{cs} \gamma_{cs}^{(t)} \right) + \sum_{i=1}^3 CC_{cs}^{(ti)} + (1 + i_{cs}) AV_{cs} \delta_{cs}^{(t)} \right] . \quad (22)$$

This is the central result of this work. The first term in the bracket corresponds to revenue loss, the last one to OPEX increase and the others to COGS increases. Let us recall that $T + 1$ terms are the same than T modulo the prefactor shown in (3). Finally, since without climate change

$$V_{cs}^{(E)} = \sum_{t=1}^{T+1} \frac{(1 + g_{cs}^{(t)}) Y_{cs}}{(1 + \omega_{cs})^t} \left[op_{cs} (1 - \tau_{cs}) + \tau_{cs} d_{cs} - k_{cs} - n_{cs} \frac{g_{cs}^{(t)} - g_{cs}^{(t-1)}}{1 + g_{cs}^{(t)}} \right] , \quad (23)$$

we have all the ingredients needed to compute $\delta \text{Var}_{cs}^{(E)}$ from equation (6).

Fictitious example of ClimVaR waterfall losses



Equation (22) allows to estimate all climate related losses, and we showed in the above figure the kind of analysis allowed by it. Equation (23) allows to compare sectors between one another, and this could be done with a heatmap representations of all country sector combinations.

Company ClimVaR

We consider a company as the sum of its N_E entities (belonging to a certain country c /sector s). The company valuation is then just the sum of its entity valuations

$$V^{(C)} = \sum_{E=1}^{N_E} V_{cs}^{(E)} . \quad (24)$$

Financial portfolio ClimVaR

ClimVaR trivially extent to a portfolio of financial f products. In this context, the portfolio is just a weighted (weights being normalized sum invested in all N_C companies of the portfolio) sum of company Valuations. Calling $\theta_{FP,C}$ the ratios of

invested sums of financial portfolio FP into company C , financial portfolio valuation reads

$$V^{(FP)} = \sum_{C=1}^{N_C} \theta_{FP,C} V^{(C)} . \quad (25)$$

Site ClimVaR

Most of the analysis performed at entity level in section A.4.1 remain valid at site level. We consider an asset a site belonging to country c of sector s in the following. Changes are twofold.

On the financial input terms:

- Revenue is now specified at asset site level Y_a .
- Asset value is now obviously specified at asset site level AV_a .
- If it makes sense, growth rate g , tax rate τ , weighted average cost of capital ω , capex k , depreciation d , insurance premium i are also adapted at site level. If no further information is provided, these information are taken from the entity the site belongs to.
- If available, COGS to revenue ratio c , OPEX to revenue ratio o , salary to Revenue ratio s , outdoor salary mass to revenue ratio so and cooling costs ec are taken from site specific data. If not available, entity information is used.

On the risk loss terms:

- All physical losses are now obviously computed at site level: this means $\beta_a^{(t)}$, $\xi_a^{(t)}$, $\gamma_a^{(t)}$ and $\delta_a^{(t)}$ are now computed at site level
- Nature loss $\rho_a^{(t)}$ is also obviously taken at site level.
- For carbon cost, scope 1 and 2 emissions should be known at site level. If not available, they could be deduce from revenues with carbon intensities. Purchases could also be known, but this is generally aggregated by an Enterprise Resource Planning (ERP) system at entity level. From all this we deduce individual $CC_a^{(ti)}$

Future revenues taking into account climate change read²⁷

$$Y_a^{(t,cc)} = \left(1 - \sigma_{cs}^{(t)}\right) \left(1 - \beta_a^{(t)}\right) \left(1 - \rho_a^{(t)}\right) Y_a^{(t)} , \quad (26)$$

Our central equation (22) then becomes

$$\begin{aligned} \text{VAR}^{(S)} = & \sum_{t=1}^{T+1} \frac{(1 - \tau_a)}{(1 + \omega_a)^t} \left[op_a \left(Y_a - Y_a^{(t,cc)} \right) + Y_a^{(t,cc)} \left(so_a \xi_a^{(t)} + ec_a \gamma_a^{(t)} \right) \right. \\ & \left. + \sum_{i=1}^3 CC_a^{(ti)} + (1 + i_a) AV_a \delta_a^{(t)} \right] . \end{aligned} \quad (27)$$

Product ClimVaR

Carefully computing ClimVaR at product level would require among other things to perform a LCA analysis for getting the detailed product carbon footprint. By doing so, all emissions generated by the product creation are correctly identified and associated to the countries in which they have been emitted and to the sectors of activities that were parts of the product supply chain in these countries. This changes drastically the carbon cost computation.

But most of the companies are not yet able to associate their purchases to specific products produced in specific sites. For this first study, we will therefore adopt another angle of attack.

A product could indeed be seen as a distribution given on sites, in the sense that several sites are needed to produce a product (but obviously a site can produce several products). So what is needed to perform a ClimVaR at product level of analysis is

- The part of a site a revenue $\theta_{P,a}$ associated to the product P
- all these weights for all N_a sites contributing to producing P .

ClimVaR at product level then reads

$$\text{VAR}^{(P)} = \sum_{a=1}^{N_a} \theta_{P,a} \text{VAR}^{(a)} \quad (28)$$

²⁷Customer shift is the only information that has to be taken from entity computation.

And reinjecting (27) gives

$$\text{VAR}^{(p)} = \sum_{a=1}^{N_a} \theta_{P,a} \sum_{t=1}^{T+1} \frac{(1-\tau_a)}{(1+\omega_a)^t} \left[op_a \left(Y_a - Y_a^{(t,cc)} \right) + Y_a^{(t,cc)} \left(so_a \xi_a^{(t)} + ec_a \gamma_a^{(t)} \right) + \sum_{i=1}^3 CC_a^{(ti)} + (1+i_a) AV_a \delta_a^{(t)} \right]. \quad (29)$$

Refinement of the present approach will be the subject of future studies.

Derivation of individual ClimVaR loss terms

The detailed derivation of each individual loss term—cooling penalty (γ), carbon costs ($CC^{1,2,3}$), business interruption (β), physical damage (δ), and heat productivity (ξ)—follows the methodology established in [Epelbaum \(2025b\)](#) and is not reproduced here. The core result is the entity-level ClimVaR formula presented in Section A.4.1 and in the main text (Exhibit 6), where each channel maps to a specific cashflow line: γ , ξ , and CC affect COGS; β affects revenue; and δ affects OPEX via insurance costs.

For the full mathematical derivation including the Shapley decomposition of hazard contributions and the Leontief input-output propagation of supply chain business interruption, we refer the reader to [Epelbaum \(2025b\)](#), Sections 3–5.

Scenario

Up to now we have avoided touching on the concept of climate scenario. This is a rich and quite complicated topic [IPCC \(2022b\)](#) on its own, and we refer the reader to [IPCC \(2022b\)](#) [NGFS \(2024\)](#) to learn more on the matter. Here we will just focus on a simple combined physical/nature/transition scenario.

Combining physical and transition scenarios

Physical climate scenarios go along the names of RCP2.5, 4.5, 8.5, SSP... For the present work, we just consider a low carbon scenario (RCP2.6 to fix ideas) and a high carbon one (RCP8.5).

For transition risks, [NGFS \(2024\)](#) [NGFS \(2021\)](#) [NGFS \(2022\)](#) has several scenarios (delayed transition, fragmented world...). We focus in our present study on Net Zero Transition (NZT) and Current Policies (CP).

We then easily form two ClimVaR ν scenarios:

- An optimistic²⁸ one: Low Carbon for physical and NZT for transition
- A pessimistic one: High Carbon for physical and CP for transition

To be noted, one could consider a worst case scenario with NZT for transition but high carbon for physical: it would correspond to a world where even though ambitious carbon mechanisms are put in place in the years to come, the carbon quantity in the atmosphere does not decrease as fast as expected, due to high hysteresis physics that we do not fully grasp yet (a hint on that is the regular underestimation of climate models on temperatures).

Intra scenario uncertainty: link with usual VaR

While scenarios describe drastically different worlds, intra scenario uncertainty can be expressed in terms of the usual Value At Risk [Jorion \(2006\)](#) (VaR). Up to now we have been using the word ClimVaR somewhat loosely, but we now make the connection between VaR and ClimVaR more clear.

While VaR has a rich history, it has been somewhat superseded by Conditional Value At Risk, or Expected Shortfall (ES) [Artzner et al. \(1999\)](#)–[Rockafellar and Uryasev \(2000\)](#). Indeed ES better captures the full tail of the loss distribution. We will therefore derive the connection for both VaR and ES.

Physical uncertainty related to ES Let e denote an extreme climatic event happening on an asset a (flood, strong winds...). As briefly sketched in section A.5 physical quantities are derived from climate models²⁹. We will denote the climate models m . The future distribution of extreme events is given by the probability density function (pdf) $f(e | \theta_{am})$, parameterized by the model-specific vector $\theta_{am} = (\theta_{a1m}, \dots, \theta_{aIm})$ where I is the number of parameters of the pdf. Examples of parameters: scale and shape for a Weibull distribution, location and scale for a Gumbell distribution or mean and variance for a Normal distribution.

²⁸Optimistic for the planet and its human inhabitants, though it might induce higher costs for emission intensive activities due to (among other things) carbon taxes.

²⁹More on that in [Popp \(2010\)](#)

ClimVaR physical risks as Mean Expected Values The different losses related to physical climate risks can be expressed a little more rigorously as follows. The expected effect of event e at site a , applying the Law of the Unconscious Statistician (LOTUS), reads

$$\mathbb{E}[D_s] = \int D(e) f(e | \bar{\theta}_a) e, \quad (30)$$

where $D(e)$ denotes the asset-level effect function (could be a damage function, a business interruption days function, or a CDD/HPL function as detailed in [Epelbaum \(2025b\)](#)) and

$$\bar{\theta}_a = \frac{1}{|M|} \sum_{m \in M} \theta_{am}. \quad (31)$$

Strictly speaking, the site-level expected damage should average over model-specific quantities $g(\theta_{am}) = \int D(e) f(e | \theta_{am}) de$. However, the approximation in (30) (evaluating at the mean parameters) is consistent with a first-order delta method expansion around $\bar{\theta}_a$.

A country mean expected damage is then

$$\mu = \mathbb{E}[D] = \frac{1}{|A|} \sum_{a \in A} \mathbb{E}[D_a]. \quad (32)$$

μ is then directly linked to the $(\beta, \delta, \xi, \gamma)_{cs}$ that we derived in section [A.5](#).

Variance and Sources of Uncertainty Up to know the approach has been mostly deterministic, but different sources of uncertainties should be considered

- (LOTUS) uncertainty, due to inherent event randomness.
- Uncertainty from variability among climate models.
- Spatial sampling uncertainty, due to finite coverage of sites across the country.

LOTUS uncertainty is straightforward

$$\sigma_{1a}^2 = \mathbb{V}(D_a) = \int D^2(e) f(e | \bar{\theta}_a) e - [\mathbb{E}(D_a)]^2, \quad (33)$$

and veraging across assets leads to

$$\sigma_1^2 = \frac{1}{|A|} \sum_{a \in A} \sigma_{1a}^2. \quad (34)$$

For model uncertainty, we define the covariance matrix among climate model parameters

$$\bar{\Sigma}_{a,ij} = \frac{1}{|M| - 1} \sum_{m \in M} (\theta_{aim} - \bar{\theta}_{ai})(\theta_{ajm} - \bar{\theta}_{aj}). \quad (35)$$

Calling

$$g(\theta_a) = \int D(e) f(e | \theta_a) e, \quad (36)$$

and using the delta method, the variance contribution from model-to-model spread is approximated by

$$\sigma_{2a}^2 = \nabla g(\bar{\theta}_a)^\top \bar{\Sigma}_a \nabla g(\bar{\theta}_a), \quad (37)$$

where ∇g is the gradient column vector with respect to θ_a . Aggregating across assets gives

$$\sigma_2^2 = \frac{1}{|A|} \sum_{a \in A} \sigma_{2a}^2. \quad (38)$$

Spatial sampling uncertainty is straightforward to derive

$$\sigma_3^2 = \frac{1}{|A| - 1} \sum_{a \in A} [\mathbb{E}(D_a) - \mu]^2. \quad (39)$$

Assuming independence among the three uncertainty sources, we write the total uncertainty as

$$\sigma^2 = \sigma_1^2 + \sigma_2^2 + \sigma_3^2, \quad (40)$$

keeping in mind that in our context σ is really σ_{cs} .

Value at Risk (VaR) The average effect D of event e at country sector level can be approximated by a Lognormal distribution (a more natural choice than a Gaussian since this effect is positive): $D \sim \text{LogNormal}(\mu_{\log}, \sigma_{\log})$, where $(\mu_{\log}, \sigma_{\log})$ are computed by standard moment matching:

$$\mu_{\log} = \ln \left(\frac{\mu}{\sqrt{1 + \frac{\sigma^2}{\mu^2}}} \right), \quad \sigma_{\log} = \sqrt{\ln \left(1 + \frac{\sigma^2}{\mu^2} \right)}. \quad (41)$$

For a given quantile level $\lambda \in (0, 1)$, the Value-at-Risk is the $1 - \lambda$ quantile of the CDF of D and is therefore defined as

$$\text{VaR}_{\lambda} = F_D^{-1}(1 - \lambda). \quad (42)$$

For a Lognormal distribution, this leads to the standard formula:

$$\text{VaR}_{\lambda} = \exp(\mu_{\log} + \sigma_{\log} \Phi^{-1}(1 - \lambda)), \quad (43)$$

where Φ denotes the standard Normal cumulative distribution function (CDF).

Conditional Value-at-Risk (CVaR)/ Expected Shortfall (ES) ES Artzner et al. (1999)–Rockafellar and Uryasev (2000) extends VaR by averaging all realizations exceeding the VaR threshold:

$$\text{CVaR}_{\lambda} = \frac{1}{\lambda} \int_{1-\lambda}^1 F_D^{-1}(u) du = \frac{1}{\lambda} \int_0^{\lambda} F_D^{-1}(1 - u) du = \frac{1}{\lambda} \int_0^{\lambda} \text{VaR}_u du \quad (44)$$

For a Lognormal distribution, substituting (43) gives (This simple derivation is repeated in Appendix A.7)

$$\text{CVaR}_{\lambda} = \frac{\exp(\mu_{\log} + \frac{1}{2}\sigma_{\log}^2)}{\lambda} \Phi(\sigma_{\log} - \Phi^{-1}(1 - \lambda)). \quad (45)$$

ClimVaR physical risks could then be evaluated in this sense. The treatment of uncertainty under multiple climate models is deferred to subsequent work.

Transition uncertainty Up to now we only considered average country sectors carbon intensities I_{cs} . But these averages hide large variances: some virtuous actors have already started to reduce their emissions, change their energy mixes... While others still heavily relies on carbon intensive processes.

We thus consider I_{cs} to be a pdf with a certain mean and variance, and follows the same logic as in the previous section A.6.2 to derive $I_{cs}^{(\lambda)}$. This means that the carbon costs become³⁰

$$MC_{cs}^{(t\lambda)} = \left[L \times \left(\mathcal{C}^{(t)} \odot I^{(\lambda)} \right) \right]_{cs}, \quad EC_{cs}^{(t\lambda)} = \frac{\left[L \times \left(\mathcal{C}^{(t)} \odot I^{(\lambda)} \right) \right]_{cs}}{[L \times I]_{cs}}. \quad (46)$$

Revised ClimVaR formula

Formula (22) can then be revised by adding to it the (inter/intra/ES) scenario concepts everywhere it makes sense, and clearly expressing all terms.

Extended ClimVaR central $\text{VaR}_{cs}^{(E\nu\lambda)}$ formula with scenarios


$$\begin{aligned} & \sum_{t=1}^{T+1} \left[\underbrace{op_{cs} \sigma_{cs}^{(t\nu\lambda)} dY_{cs}^{(t)}}_{\text{Customer Shift}} + \underbrace{op_{cs} \left(1 - \sigma_{cs}^{(t\nu\lambda)} \right) \beta_{cs}^{(t\nu\lambda)} dY_{cs}^{(t)}}_{\text{Business interruptions}} \right. \\ & + \underbrace{op_{cs} \left(1 - \sigma_{cs}^{(t\nu\lambda)} \right) \left(1 - \beta_{cs}^{(t\nu\lambda)} \right) \rho_{cs}^{(t\nu\lambda)} dY_{cs}^{(t)}}_{\text{Renewable Natural Capital loss}} + \underbrace{(1 + i_{cs}) \delta_{cs}^{(t\nu\lambda)} dAV_{cs}^{(t)}}_{\text{Asset Insurance premium}} \\ & + \underbrace{s_{cs}^{(1)} \mathcal{C}_{cs}^{(t\nu\lambda)} dY_{cs}^{(t\nu\lambda)}}_{\text{Scope 1 CC}} + \underbrace{s_{cs}^{(2)} EC_{cs}^{(t\nu\lambda)} \left(1 + \varepsilon_{cs} \gamma_{cs}^{(t)} \right) dY_{cs}^{(t\nu\lambda)}}_{\text{Scope 2 CC}} \\ & \left. + \underbrace{p_{cs} \mathcal{M}_{cs}^{(t\nu\lambda)} dY_{cs}^{(t\nu\lambda)}}_{\text{Scope 3 CC}} + \underbrace{so_{cs} \xi_{cs}^{(t\nu\lambda)} dY_{cs}^{(t\nu\lambda)}}_{\text{Heat Productivity Loss}} + \underbrace{ec_{cs} \gamma_{cs}^{(t\nu\lambda)} dY_{cs}^{(t\nu\lambda)}}_{\text{Cooling Cost Increase}} \right], \quad (47) \end{aligned}$$

³⁰Please note that it is important to fix the denominator intensity for EC in this intra scenario setup, otherwise EC would not duly increase as the VaR/ES quantile considered increases.

where we have defined

$$dY_{cs}^{(t)} = \frac{(1 - \tau_{cs}) Y_{cs}^{(t)}}{(1 + \omega_{cs})^t}, \quad dY_{cs}^{(t\nu\lambda)} = dY_{cs}^{(t)} \frac{Y_{cs}^{(t\nu\lambda, cc)}}{Y_{cs}^{(t)}}, \quad dAV_{cs}^{(t)} = \frac{(1 - \tau_{cs}) AV_{cs}}{(1 + \omega_{cs})^t}. \quad (48)$$

In a nutshell, ClimVaR then looks as follow



Expected Shortfall for a lognormal distribution

Starting from (44) and using (43),

$$\begin{aligned} CVaR_{\lambda} &= \frac{1}{\lambda} \int_0^{\lambda} \exp(\mu_{\log} + \sigma_{\log} \Phi^{-1}(1 - \gamma)) d\gamma \\ &= \frac{1}{\lambda} \int_{1-\lambda}^1 \exp(\mu_{\log} + \sigma_{\log} \Phi^{-1}(\gamma)) d\gamma. \end{aligned} \quad (49)$$

With the variable substitution $\gamma = \Phi(x)$, we have $\gamma = \phi(x) x = \frac{1}{\sqrt{2\pi}} e^{-x^2/2} x$, leading to

$$\text{CVaR}_\lambda = \frac{e^{\mu_{\log}}}{\sqrt{2\pi}\lambda} \int_{\Phi^{-1}(1-\lambda)}^{+\infty} e^{\sigma_{\log} x} e^{-x^2/2} dx \quad (50)$$

$$= \frac{e^{\mu_{\log} + \frac{1}{2}\sigma_{\log}^2}}{\sqrt{2\pi}\lambda} \int_{\Phi^{-1}(1-\lambda) - \sigma_{\log}}^{+\infty} e^{-x^2/2} dx \quad (51)$$

$$= \frac{e^{\mu_{\log} + \frac{1}{2}\sigma_{\log}^2}}{\sqrt{2\pi}\lambda} \int_{-\infty}^{\sigma_{\log} - \Phi^{-1}(1-\lambda)} e^{-x^2/2} dx \quad (52)$$

$$= \frac{e^{\mu_{\log} + \frac{1}{2}\sigma_{\log}^2}}{\lambda} \Phi(\sigma_{\log} - \Phi^{-1}(1-\lambda)), \quad (53)$$

which confirms the closed-form expression in (45).



Appendix B

Data and Hypotheses

Wassily Leontief (1906–1999)

The data for a particular year, set out in the form of a rectangular input-output table, describe the flows of goods and services among all sectors of an economy.

The empirical foundation of this study rests on a geocoded database of 8,572 installed data centers across 120 countries (90.8 GW), linked to facility-level financial parameters, CMIP climate projections, and three AI-era growth scenarios. This appendix documents model architecture, database construction, financial calibration, physical risk data, supply chain propagation (Leontief input-output),

and output data—enabling full reproducibility.

Contents

- B.1 Model architecture and database construction
- B.2 Financial parameter calibration
- B.3 Physical risk data and supply chain data
- B.4 Output data and parameter confidence

Data and Hypotheses

Model Architecture

The analysis is built from a structured input system containing 17 interconnected components that form the complete hypothesis chain. The architecture is designed for transparency and reproducibility: every numerical assumption is stored in a single documented location, and modifying any parameter triggers automatic recalculation of all downstream outputs.

Table 1. Model input architecture — 17 components

Component	Category	Content
Scenario definitions	Documentation	Model version history (V5–V7), scenario structure, changelog
Data flow architecture	Documentation	Dependency map across all input and output components
Global parameters	Hypotheses	29 modifiable parameters: asset valuation, operating ratios, growth, depreciation
Carbon decoupling ratios	Hypotheses	Carbon/revenue trajectories by region × scenario (2025–2050)
Source traceability	Hypotheses	Literature references and confidence ratings for each parameter
Country financial data	Parameters	120 countries: grid carbon intensity (5 horizons), tax, COGS, OPEX, WACC
Facility database	Core data	13,205 geocoded facilities × 64 attributes
Efficiency matrices	Parameters	PPA factors and PUE by region × facility tier, with dynamic trajectories
Regional carbon projections	Trajectories	MW-weighted regional carbon intensity to 2050
Aggregate results	Outputs	ClimVaR by scenario, region, channel, adaptation pathway
Efficiency sensitivity	Outputs	Sensitivity of results to PPA and PUE parameter variation
Regional decomposition	Outputs	Portfolio, ClimVaR × region, facility type, channels, risk tiers
IEA reference mapping	Calibration	IEA source citations with section-level traceability
AI demand projections	Calibration	IEA electricity demand under three AI growth scenarios
Grid supply mix	Calibration	Regional generation mix projections to 2050
Emissions cross-check	Calibration	Modeled emissions validated against IEA benchmarks
Extended risk parameters	Calibration	Grid delay, electricity price multipliers, supply chain data

The model has undergone four documented version increments. Version 5 introduced dynamic IEA trajectories replacing static parameters. Version 6 added regionalized asset values, cooling ratios, and grid congestion modeling. Version 7 computed facility-level ClimVaR across all scenarios, producing 24 output columns per facility (4 scenario combinations × 6 channel metrics). Version 7.5 integrated three AI growth scenarios with binary inclusion flags, enabling portfolio-level analysis under LTG, SAI, and AWB assumptions from a single database.

Facility Database

Database Dimensions

The core facility database contains 13,205 rows and 64 columns. Each row represents a single data center facility; each column carries one of six categories of information:

- **Location** (3 columns): geocoded latitude, longitude, and ISO country code (120 countries represented).
- **Identity and capacity** (8 columns): unique identifier, IT capacity (MW), facility name, owner, provider type, sub-region, rack count, source database.
- **Financial structure** (14 columns): annual revenue, asset value, COGS ratio, OPEX ratio, tax rate, WACC, yearly growth, terminal growth, depreciation ratio, CAPEX ratio, insurance premium, electricity cooling ratio, outdoor salary ratio, purchase ratio.
- **Efficiency and emissions** (6 columns): Scope 1 emissions (tCO₂/yr), Scope 2 raw and PPA-adjusted, PUE, PPA factor, and terminal value year.
- **ClimVaR outputs** (24 columns): facility-level ClimVaR percentage and five channel decompositions ($\gamma, CC^2, \beta, \delta, \xi$), under 4 scenario combinations (BAU/NZT × MEAN/Q95).
- **Scenario flags** (3 columns): binary inclusion flags for LTG, SAI, AWB growth scenarios.

Installed Base: Sources and Coverage

The installed base comprises 8,572 facilities totaling 90,818 MW (90.8 GW), assembled from four complementary source types:

1. **Commercial databases:** Data Center Map, Cloudscene, and proprietary Schneider Electric client records provide the core inventory with location, operator, and capacity attributes. Cross-referencing eliminates duplicates.
2. **Regulatory filings:** planning permissions, environmental impact assessments, and grid connection applications in the US, EU, China, and Japan provide independent capacity verification for facilities above 10 MW.
3. **Satellite imagery validation:** building footprint analysis provides independent size estimates, calibrated against known facilities to derive MW capacity from floor area.
4. **Operator disclosures:** annual and sustainability reports from publicly listed operators (Equinix, Digital Realty, CyrusOne, NTT, Chindata, GDS) provide facility-level data for approximately 35% of total installed GW.

Capacity distribution. Facility IT capacities range from 15 kW (enterprise server rooms) to 300 MW (hyperscale campuses), with a median of 10.6 MW. The distribution is heavily right-skewed: 73.4% of facilities exceed 1 MW, but only 4.7% exceed 50 MW and 3.0% exceed 100 MW. The top decile by capacity accounts for over 60% of total GW.

Table 2. Installed base capacity distribution

Capacity band	Facilities	% of base	Cumulative GW
< 1 MW	2,283	26.6%	0.9
1–5 MW	3,221	37.6%	9.5
5–10 MW	906	10.6%	6.6
10–50 MW	1,755	20.5%	36.5
50–100 MW	152	1.8%	10.7
> 100 MW	255	3.0%	26.6
Total	8,572	100%	90.8

Provider types. The database classifies facilities into five provider categories reflecting operational characteristics and financial profiles.

Table 3. Installed base by provider type

Provider type	Facilities	% of base	Typical capacity
Retail colocation	5,662	66.1%	0.015–5 MW
Wholesale / large colocation	1,560	18.2%	5–50 MW
Telecommunications	1,068	12.5%	0.5–20 MW
Hyperscale (cloud)	149	1.7%	50–300 MW
Other (hosting, reseller, cloud)	133	1.6%	Variable

Country coverage. Table 4 presents the 20 largest data center markets by installed capacity. Together they account for 93% of global installed MW. The full database covers 120 countries.

Table 4. Top 20 data center markets by installed capacity

Country	DCs	MW	%	Country	DCs	MW	%
United States	2,383	30,329	33.4	Canada	272	1,327	1.5
China	1,284	25,076	27.6	Russia	255	1,054	1.2
United Kingdom	374	3,596	4.0	South Korea	94	1,023	1.1
Japan	447	2,940	3.2	Norway	27	896	1.0
India	206	2,603	2.9	Hong Kong	109	809	0.9
Germany	291	2,128	2.3	UAE	49	775	0.9
France	186	1,682	1.9	Israel	61	739	0.8
Australia	247	1,624	1.8	Italy	151	716	0.8
Singapore	115	1,547	1.7	Brazil	116	683	0.8
Ireland	63	1,542	1.7	Netherlands	207	1,434	1.6

Aggregate emissions. The installed base generates approximately 1.0 MtCO₂/yr in Scope 1 emissions (diesel backup) and 482.7 MtCO₂/yr in raw Scope 2 emissions (purchased electricity). After PPA adjustment, Scope 2 decreases to 400.0 MtCO₂/yr (17% reduction). Total annual revenue is \$147 billion; total book asset value is \$750 billion.

Prospective AI-Era Infrastructure

Prospective facilities are drawn from four source databases, ordered by decreasing certainty.

Table 5. Prospective facility source databases

Source	Facilities	GW	Mean MW	Certainty
Epoch AI GPU Clusters	62	7.8	125	High
Tier 2 Announced	46	4.6	100	Medium
Probabilistic LTG	2,316	231.3	100	Low
Probabilistic SAI	2,195	219.1	100	Low

The total unique database contains 4,633 prospective facilities (462.7 GW). Facilities are assigned to growth scenarios via binary inclusion flags. The effective count per scenario is: LTG = 615 (62.9 GW), SAI = 1,195 (121.1 GW), AWB = 2,939 (295.6 GW). Results in Parts IV–VI are reported by scenario, not for the raw total.

Tier 1 confirmed projects (sample). The Epoch AI GPU Cluster database includes named, publicly documented projects: Major AI hyperscaler campus (US South, phased to 500 MW), Microsoft AI campuses (UK, Sweden, Spain), Meta superclusters (multiple US locations), xAI Memphis (500 MW announced), NVIDIA Blackwell clusters (multiple sites), AWS sovereign deployments (Saudi Arabia, Malaysia, Thailand, New Zealand, Chile), and Google APAC expansion projects. These represent the highest-certainty prospective capacity—most are under construction or recently commissioned.

Tier 2 announced projects (sample). Includes sovereign AI campuses (Sovereign AI project, France, 250 MW, DataVolt NEOM SA 1.5 GW phased), large colocation expansions with planning permissions (Pure DC UK, Vantage \$2B deployment, Brookfield \$10B plan), and national programs (SK 3 GW plan South Korea, Scala 4.75 GW Brazil).

Scenario Composition

Table 6. Portfolio composition by AI growth scenario

Scenario	Total DCs	Total GW	AI added	AI GW
Installed only	8,572	90.8	—	—
+ Limits to Growth	9,187	153.7	615	62.9
+ Sustainable AI	9,767	211.9	1,195	121.1
+ Abundance Without Boundaries	11,511	386.4	2,939	295.6

Regional Distribution

Table 7. Facility coverage by region

Region	Installed base		Total (AWB)	
	DCs	GW	DCs	GW
United States	2,383	30.3	4,071	199.1
China	1,284	25.1	2,267	123.4
Europe	2,019	15.4	2,574	70.9
Asia-Pacific	1,593	11.6	2,495	101.2
Middle East	255	2.6	370	14.1
Rest of World	1,038	5.9	1,428	44.9
Total	8,572	90.8	11,511	386.4

Asia-Pacific excludes China; includes Japan, Australia, Singapore, South Korea. Rest of World includes India, Canada, Latin America, Africa, Russia.

Geocoding and Climate Linkage

All facilities are geocoded to coordinates enabling linkage to downscaled CMIP6 climate projections at $0.1^\circ \times 0.1^\circ$ resolution (~ 10 km). For installed facilities, coordinates are derived from address geocoding or satellite imagery. For prospective facilities, coordinates are assigned at the announced site location (Tier 1/2) or metropolitan area centroid (probabilistic). The climate linkage assigns four physical risk variables to each facility for each year (2025–2050) under each RCP pathway at three distribution quantiles (Q25, MEAN, Q95).

Financial Parameter Calibration

Global Parameters

Table 8. Global model parameters (29 documented hypotheses)

Parameter	Value	Unit	Source	Confidence
Asset value / MW (base)	9	\$M/MW	Schneider, CBRE 2024	High
PUE (global default)	1.40	—	Uptime Institute 2024	High
Backup ratio (diesel)	0.50	—	Industry standard	High
Diesel consumption	300	L/MWh	Generator spec	High
Diesel emission factor	2.68	kgCO ₂ /L	GHG Protocol	High
AI revenue multiplier	2.0×	—	McKinsey, Thunder Said	Mixed
Size mult. (retail)	1.0×	—	Baseline	High
Size mult. (wholesale)	0.87×	—	CBRE, Houlihan Lokey	High
Size mult. (hyperscale)	0.75×	—	CBRE, Synergy Research	High
Revenue growth	5%/yr	—	Industry consensus	Mixed
Terminal growth	2%	—	Standard DCF	High
Elec. cooling / revenue	5%	—	Uptime, IEA 2025	Mixed
Outdoor salary / revenue	0.4%	—	Expert estimate	Low
Purchase ratio (p)	15%	—	$p \approx c - s$	Mixed
Insurance premium (base)	5%	of damage	Industry standard	Mixed
Insurance premium (AI)	10%	of damage	Expert estimate	Low
Depreciation / AV (d)	8%	—	Damodaran, Equinix	High
CAPEX / revenue (base)	10%	—	CBRE mature DC	High
CAPEX / revenue (AI tiers)	20%	—	Equinix 32%, DLR 27%	Mixed

Asset Valuation by Region and Tier

Table 9. Asset value per MW (\$M/MW) by region and facility tier

Region	Base	Tier1	Tier2	Tier3	Rationale
US	9	40	25	30	Highest construction costs
Europe	9	38	24	28	~95% of US costs
China	7	28	18	22	~70% of US costs
Asia-Pacific	8	32	20	25	Mix SG/JP (high) + ID/TH (low)
Middle East	9	42	26	32	Extreme heat cooling premium
RoW	8	30	20	24	Mixed; LatAm ~85% of US

Sources: IEA (2025); Schneider Electric TradeOff Tool; CBRE 2024. Tier1 upper bound (\$42M/MW, Middle East) reflects IEA frontier AI estimates.

The base asset value of \$9M/MW produces a total installed book value of \$750 billion. Size multipliers and the AI revenue multiplier produce an effective range from \$6.8M/MW (large legacy wholesale) to \$42M/MW (Tier 1 AI, Middle East). The DCF valuation (V_0) exceeds book value by approximately $\times 1.35$ for the installed base (\$1,012B on \$750B), reflecting the going-concern premium of operational facilities.

Observed Parameter Ranges

Table 10 reports the actual ranges computed across all 13,205 facilities, validating that calibration produces realistic distributions.

Table 10. Observed parameter ranges across the database

Parameter	Min	Median	Max	Unit
IT capacity	0.015	11.2	300.0	MW
Asset value	\$120K	\$91M	\$8.0B	\$
Annual revenue	\$29K	\$18.8M	\$720M	\$/yr
COGS / revenue	35%	46%	50%	ratio
OPEX / revenue	22%	27%	30%	ratio
WACC	5.5%	7.0%	15.5%	ratio
Tax rate	0%	25%	35%	ratio
PUE	1.12	1.40	1.55	ratio
PPA factor	0.49	0.70	0.90	ratio
ClimVaR% (BAU/MEAN)	2.4%	8.8%	16.3%	% of V_0
ClimVaR% (NZT/MEAN)	3.0%	9.7%	17.0%	% of V_0

Country-Level Parameters

The model provides 14 parameters for each of 120 countries: grid carbon intensity at five time horizons (2024, 2030, 2035, 2040, 2050), tax rate, COGS ratio, OPEX ratio, WACC, and regional classification. Table 11 presents the 20 largest markets.

Table 11. Financial and carbon parameters — 20 largest data center markets

Country	Region	Grid carbon intensity (gCO ₂ /kWh)					Tax	COGS	WACC
		2024	2030	2035	2040	2050			
US	N. Am.	390	293	215	156	120	21%	48%	6.8%
CN	AP Em.	580	464	348	232	136	25%	38%	9.0%
DE	W. Eur.	380	247	160	106	46	30%	49%	6.5%
GB	W. Eur.	220	143	93	62	26	25%	47%	7.2%
JP	AP Dev.	470	329	235	188	141	30%	48%	6.5%
FR	W. Eur.	55	36	23	15	7	25%	44%	6.8%
SG	AP Dev.	420	294	210	168	126	17%	45%	6.0%
IN	S. Asia	710	533	391	284	178	25%	35%	11.5%
AU	AP Dev.	620	496	372	260	136	30%	44%	7.2%
NL	W. Eur.	310	202	131	87	37	26%	47%	6.5%
BR	Lat. Am.	90	74	57	45	25	34%	40%	11.5%
KR	AP Dev.	450	315	225	180	90	24%	44%	6.5%
IE	W. Eur.	300	195	126	84	36	12%	47%	7.0%
AE	M. East	400	352	288	232	140	9%	36%	8.0%
SE	W. Eur.	15	10	6	4	3	21%	47%	6.5%
NO	W. Eur.	10	7	4	3	2	22%	47%	6.5%
MY	AP Em.	550	413	303	220	110	24%	38%	8.5%
ID	AP Em.	700	560	420	280	168	22%	35%	10.0%
SA	M. East	500	440	360	290	175	20%	40%	8.5%
MX	Lat. Am.	440	330	242	176	110	30%	42%	10.0%

Grid carbon intensity: IEA (2025) regional projections. Tax: OECD (2024). COGS: Equinix, Digital Realty, CyrusOne reports. WACC: Damodaran (2024) REIT sector with country risk premiums. The 100 remaining countries follow the same structure.

Physical Risk Data

Climate Projection Source

Physical risk variables are extracted from bias-corrected, statistically downscaled CMIP6 projections (Eyring et al., 2016) at $0.1^\circ \times 0.1^\circ$ resolution (~ 10 km at mid-latitudes). Projections are generated at each facility's geocoded coordinates for 2025–2050 under RCP 2.6, 4.5, and 8.5, at three distribution quantiles (Q25, MEAN, Q95).

Variable Definitions

Table 12. Physical risk variables: definitions and channel mapping

Variable	Channel	Definition	Climate driver
COOLING	γ	Fractional increase in cooling energy vs. 2024	Cooling degree days (CDD, base 18°C, ASHRAE)
BID	β	Business interruption days/yr, fraction of year	Composite: flood, storm, drought, wild-fire, heat
DAMAGE	δ	Annual damage probability, fraction of AV	Flood depth, wind speed, seismic, fire
HPL	ξ	Fractional outdoor labor productivity loss	Days > 35°C / 40°C wet-bulb

Hazard Sub-Decomposition

For the business interruption and damage channels, the model further decomposes risk into five hazard types: drought (d), extreme wind (ew), riverine flood (r), storm surge (s), and tropical wind (tw). This decomposition provides the inputs for the Shapley attribution (Section 2.2.3) and is reported at facility level in the case study outputs. Each hazard type carries its own annual probability distribution at facility coordinates.

Dataset Scale

The complete projection dataset contains approximately 17 million observations: up to 11,511 facilities (AWB) $\times$ 25 years $\times$ 2–3 RCP pathways $\times$ 4 physical risk variables $\times$ 3 quantiles. Projections are computed for all facilities regardless of scenario inclusion, enabling flexible portfolio composition at analysis stage.

Supply Chain Data

Leontief Input-Output Tables

The technical coefficient matrix A is drawn from OECD Multi-Regional Input-Output Tables (MRIOT, 2019 vintage): 67 countries $\times$ 45 ISIC sectors. The data center sector maps to ISIC telecommunications, computer services, and electrical equipment manufacturing.

Supply Chain Concentration

Table 13. Critical supply chain components for data center infrastructure

Component	China share	Category	Substitut.	DC impact
GaN devices	99%	Semiconductors	Very low	Power conversion
High-purity Si (4N+)	95%	Raw material	Very low	Semiconductor base
Rare earth magnets	90%	Cooling	Very low	Fans, pumps, compressors
LFP battery cells	99%	Storage	Low	UPS, BESS
SiC substrates	50%	Semiconductors	Medium	Power supply
HBM memory	$\sim 40\%$ (KR)	Memory	Low	GPU bandwidth
Refined copper	40%	Cabling/cooling	Medium	Busbars, heat exchangers
Transformers	$\sim 35\%$	Grid equipment	Medium	DC power feed

Sources: IEA (2025) Table 4.1; IEA Critical Minerals; USGS.

Pass-Through Calibration

Table 14. Leontief pass-through rates by supply chain pathway

Pathway	Rate	Lead time	Rationale
Semiconductor fab	60–80%	6–12 mo	TSMC 54%, limited substitutability
Power transformers	40–60%	18–36 mo	Moderate concentration, long lead
Cooling components	20–40%	3–6 mo	More substitutable
Server/rack assembly	30–50%	3–9 mo	Component bottlenecks

Calibrated from: Barrot and Sauvagnat (2016), IEA (IEA, 2025), expert assessment. Fixed-coefficient limitation discussed in Section 5.10.

Output Data Structure

The model produces three output datasets, documented in Table 15.

Table 15. Model output datasets

Dataset	Rows	Cols	Content
Aggregate results	288	14	ClimVaR by scenario (BAU/NZT) × quantile (Q25/MEAN/Q95) × region (4) × asset class (4) × adaptation (3). Columns: V_0 , five channels, residual value, ClimVaR%, ClimVaR.
Annual decomposition	288	305	Same cross-section with annual resolution 2026–2050. 12 channel variables × 25 years + 3 annual summary metrics × 25 years. Enables trajectory analysis.
Facility case studies	72	592	Three sites (Site A, Site B, Site C) × 24 scenario combinations. 64 base columns + 528 time-series columns: 22 channel-hazard decompositions × 25 years.

The 12 channel variables decompose ClimVaR into: direct business interruption, supply chain interruption, physical damage (by 5 hazard types: drought, extreme wind, riverine flood, storm surge, tropical wind), heat productivity, cooling cost increase, and carbon costs (Scopes 1, 2, 3). The facility case study dataset further decomposes damage into the same five hazard types, enabling the Shapley attribution presented in Section 4.3.

Parameter Confidence and Sensitivity

Of the 29 documented global parameters, 14 (48%) are calibrated directly from published data with high confidence, 10 (34%) combine literature and expert judgment, and 5 (17%) rely primarily on expert estimates. The low-confidence parameters are minor contributors: a ±50% perturbation of all five changes aggregate ClimVaR by <3%.

Validation Cross-Checks

Industry benchmark. The WEF projects climate-driven costs for data centers of \$81 billion annually by 2035 (World Economic Forum and S&P Global, 2025). Our installed base ClimVaR of \$388 billion, annualized over 25 years at 7% discount rate, implies approximately \$34 billion/yr—lower than the WEF figure (which includes prospective infrastructure). Under Abundance Without Boundaries, our annualized equivalent rises to ~\$150 billion. Directional alignment confirms that ClimVaR values are conservative.

Sensitivity bounds. WACC ±2 pp shifts aggregate ClimVaR by ±15–20%; halving BI frequency reduces it by ~25%; ±50% on the five low-confidence parameters changes aggregate ClimVaR by <3%. The order of magnitude (~\$400B installed) is robust to reasonable parameter variation. Monte Carlo sensitivity analysis is planned for a subsequent version.



Appendix C

Scenario Architecture

Harry Markowitz (1927–2023)

Diversification is both observed and sensible; a rule of behavior which does not imply the superiority of diversification must be rejected both as a hypothesis and as a maxim.

Adaptation investment across a portfolio of data centers mirrors the logic of financial portfolio theory. Operators allocate finite capital across facilities with heterogeneous risk profiles, seeking the combination that maximizes risk reduction per dollar invested.

The framework models three strategies—Baseline, Selective, and Proactive—and demonstrates that risk reductions are multiplicative, not additive, because channels interact through the Shapley decomposition. The “one champion

action” approach (Selective) targets the highest-Shapley-value channel per facility.

Contents

- C.1 Climate-economic pathways (BAU, NZT)
- C.2 AI growth scenarios (LTG, SAI, AWB)
- C.3 Risk channel trajectories (2026–2050)
- C.4 Adaptation parameters and non-additivity

Scenario Architecture

Efficiency and Emissions Trajectories

PPA Factor Trajectories

PPA factors represent the fraction of grid emissions retained after renewable procurement (1.0 = fully grid-dependent; 0.0 = fully renewable). They decline over time as renewable procurement expands, at different rates by region reflecting market maturity and policy environment.

Table 16. PPA factor trajectories by region (base facilities, BAU pathway, 2024–2050)

Region	2024	2030	2035	2040	2045	2050
US	0.70	0.55	0.42	0.32	0.25	0.20
Europe	0.65	0.48	0.33	0.22	0.16	0.12
China	0.85	0.78	0.65	0.52	0.44	0.38
Asia-Pacific	0.80	0.68	0.55	0.42	0.35	0.30
Middle East	0.90	0.80	0.65	0.50	0.42	0.35
RoW	0.90	0.85	0.75	0.62	0.54	0.48

Tier multipliers: base $\times 1.0$, tier1 $\times 0.75$, tier2 $\times 0.80$, tier3 $\times 0.85$. Under NZT, decline is approximately 30% steeper. Sources: IEA, RE100, BNEF 2024. Validated against IEA (2025) §6: ~40% industry PPA coverage by 2027.

PUE Trajectories

Table 17. PUE trajectories by region (base / tier1)

Region	2024		2035		2050	
	Base	Tier1	Base	Tier1	Base	Tier1
US	1.45	1.15	1.24	1.10	1.12	1.06
Europe	1.40	1.12	1.20	1.08	1.10	1.05
China	1.50	1.18	1.28	1.11	1.14	1.06
Asia-Pacific	1.45	1.15	1.26	1.10	1.14	1.06
Middle East	1.55	1.20	1.32	1.12	1.16	1.07
RoW	1.55	1.20	1.32	1.12	1.16	1.07

Sources: Uptime Institute 2024 (industry avg = 1.56); Google sustainability reports (1.09); IEA (2025). China government mandates PUE <1.3 for new builds.

Carbon Ratio Trajectories

The carbon ratio decouples carbon cost growth from revenue growth, capturing the combined effect of grid decarbonization, renewable procurement, and efficiency improvement:

$$\text{carbon_ratio}(r, t) \approx \frac{I_{\text{carb}}(r, t)}{I_{\text{carb}}(r, 2024)} \times \frac{\text{PPA}(r, t)}{\text{PPA}(r, 2024)} \times \frac{\text{PUE}(r, t)}{\text{PUE}(r, 2024)}$$

Table 18. Carbon ratio trajectories (BAU / NZT)

Region	2030		2040		2050	
	BAU	NZT	BAU	NZT	BAU	NZT
US	0.92	0.82	0.75	0.48	0.62	0.25
Europe	0.85	0.75	0.58	0.35	0.42	0.15
China	0.95	0.88	0.78	0.55	0.62	0.32
Asia-Pacific	0.93	0.85	0.74	0.50	0.58	0.28
Middle East	0.96	0.88	0.76	0.52	0.56	0.28
RoW	0.97	0.92	0.84	0.65	0.68	0.40

Carbon Price Trajectories

Carbon cost trajectories follow IEA scenarios (IEA, 2024), differentiated by region and policy ambition.

Table 19. Carbon price trajectories by scenario and region (\$/tCO₂)

Scenario	Region	2025	2030	2035	2040	2050
BAU	Europe (ETS)	65	85	100	110	120
	China (ETS)	10	25	40	55	80
	US (shadow)	15	25	35	50	75
	RoW (effective)	5	12	20	30	50
NZT	Europe (ETS)	65	130	200	260	300
	China (ETS)	10	50	100	160	250
	US (carbon tax)	15	60	110	170	250
	RoW (effective)	5	30	65	110	200

BAU reflects current policy trends. NZT follows NGFS Net Zero 2050. Under NZT, carbon costs multiply $\times 3\text{--}6$ relative to BAU by 2050. Sources: NGFS Phase IV (2023), EU ETS market data, IEA WEO (2025).

Carbon costs are allocated to three scopes following GHG Protocol definitions: Scope 1 covers direct emissions from diesel backup generators; Scope 2 covers purchased electricity, adjusted for grid carbon intensity by country and PPA factor; Scope 3 covers upstream supply chain emissions using Leontief propagation coefficients.

Adaptation Parameters

Reduction Rates by Channel

Table 20. Adaptation reduction rates by channel and pathway

Channel	Selective	Proactive	Primary mechanism
γ Cooling	15–25%	24%	Liquid cooling adoption (immersion, DTC, CDU)
CC Carbon	40–50%	63%	Renewable PPAs + grid selection + on-site generation
β Business Interruption	10–20%	36%	Microgrid + BESS + predictive monitoring
δ Physical Damage	10–15%	24%	Structural hardening + climate-aware siting
ξ Heat Productivity	10–15%	28%	Automation + enclosed logistics + remote operations
Portfolio weighted	~21%	~38%	

Selective applies to new builds only (greenfield, +15–20% CAPEX). Proactive applies across the portfolio including legacy retrofit (+40–60% CAPEX). Reductions applied independently per channel. Interaction effects are partially captured but may be understated (see Section 5.10).

Adaptation Strategy Definitions

Three pathways span the range from inaction to systematic portfolio integration:

- Baseline:** No climate-specific adaptation. Air-cooled, grid-dependent, cost-optimized siting, no legacy retrofit. Represents the current industry default against which adaptation value is measured.
- Selective:** Climate-aware design for new builds only. Includes liquid cooling, climate-informed siting, and renewable PPAs for prospective facilities; legacy assets remain unchanged. Captures the “low-regret” adaptation available at greenfield stage. The Selective pathway captures approximately 56% of total Proactive value at lower capital outlay.
- Proactive:** Systematic adaptation across the portfolio. Includes all Selective measures plus microgrid retrofit on legacy assets, geographic rebalancing of new capacity, structured adaptation programs incorporating energy services and monitoring, and consideration of early retirement for highest-risk facilities. Combines infrastructure technology solutions, energy strategy decisions, and portfolio management choices.

Risk Channel Modeling

The mathematical derivation of how physical and transition risks enter the ClimVaR framework through the five channels (γ , $CC^{1:2:3}$, β , δ , ξ) is provided in Appendix B (Sections B.6–B.9), including the Shapley decomposition for hazard attribution within the β channel and the Leontief propagation for supply chain effects.

The key calibration result for transition risk: the combined PUE $\times$ PPA adjustment reduces aggregate Scope 2 emissions from 2,888 MtCO₂/year (raw) to 1,841 MtCO₂/year (adjusted)—a 36% reduction. Prospective AI facilities achieve approximately 50% reduction; the installed base achieves approximately 17%.

Physical and transition scenarios interact multiplicatively: RCP pathways set the hazard environment for physical channels; BAU/NZT pathways set carbon prices for transition channels; adaptation pathways modify both. The framework evaluates all combinations to identify robust strategies.



Appendix D

Literature Review

Gro Harlem Brundtland (b. 1939)

Sustainable development is development that meets the needs of the present without compromising the ability of future generations to meet their own needs.

The empirical foundation rests on a geocoded database of 8,572 installed data centers across 120 countries (90.8 GW), linked to facility-level financial parameters, climate projections, and three AI-era growth scenarios. This appendix documents the literature context and reviews existing approaches to climate risk assessment for data centers and digital infrastructure.

Contents

- D.1 Physical climate risks for data centers
- D.2 Financial valuation of climate risk
- D.3 Infrastructure-specific assessment frameworks
- D.4 Gap analysis and positioning

Literature Review

Results are organised along three guiding questions: (Q1) physical climate risks for data centers and digital infrastructure; (Q2) insurance, insurability and climate risk for critical infrastructure; (Q3) quantitative frameworks to integrate climate risk into infrastructure valuation.

Search Strategy and Screening Criteria

The search process combined a broad semantic query on climate risk for data centers and AI-related digital infrastructure with structured screening based on abstracts and, where available, full texts. Papers and reports were screened in when they met the following criteria:

- **Infrastructure focus:** data centers, digital infrastructure, AI infrastructure, or closely related critical infrastructure (energy systems, telecommunications, industrial facilities, utilities).
- **Climate risk analysis:** quantitative modelling or analysis of physical climate risks (heat, flooding, storms, sea level rise, drought) and/or their economic or operational impacts.
- **Publication quality:** peer-reviewed academic papers or high-quality institutional reports from central banks, regulators, international organisations, or major research institutes.
- **Financial risk assessment:** explicit analysis of climate-related insurance, risk transfer, or methods for integrating climate risk into investment and financial assessment.
- **Infrastructure scale:** large-scale critical infrastructure rather than purely residential or small commercial buildings.

For each included study, we extracted information on infrastructure scope, hazards, risk modelling approach, financial impacts, insurance treatment, valuation framework, main findings, and acknowledged limitations.

The remainder of this appendix summarises the evidence base by guiding question.

Q1 – Physical Climate Risks for Data Centers and Digital Infrastructure

The peer-reviewed literature directly quantifying physical climate risks for data centers is extremely limited: only three studies explicitly address data center siting and operations under climate change. A broader body of work on analogous infrastructure — power generation, telecommunications, and electricity grids — provides transferable methodologies. Flooding (riverine, pluvial and coastal) emerges as the most consistently analysed hazard, with heat extremes and storms also prominent. Most studies rely on RCP 4.5 and RCP 8.5 (or SSP-equivalent) scenarios, with horizons extending from the 2030s to 2100. Limitations repeatedly noted include coarse hazard-map resolution, generic vulnerability functions not calibrated to local conditions, and inadequate treatment of indirect impacts, cascading effects, and compound hazards.

Direct Data Center Studies

Table 21. Direct data center climate risk studies

Study	Scope	Key contributions and limitations
Pasupuleti (2022)	Data center siting and capacity decisions	Integrates HEC–HMS hydrologic simulations into data center location and capacity planning. Uses rainfall–runoff scenarios to estimate flood probabilities and downtime risk, combined with CloudSim models to quantify impacts on service availability and financial loss. Demonstrates a rare linkage from hydrological hazards to operational disruptions; does not yet include multi-hazard or insurance effects.
Pasupuleti (2024)	GIS-based siting under flood and climate risk	Develops a climate- and flood-aware GIS framework for data center site selection, incorporating riverine and pluvial flooding, sea-level rise, and extreme precipitation. High-resolution DEMs and multi-criteria decision analysis produce spatially explicit risk surfaces. Reports a 63% reduction in expected annual damage for hazard-aware siting relative to purely economic criteria; vulnerability remains represented by generic damage functions.
Pasupuleti (2025)	GIS–AI framework for resilient siting and operation	Combines climate datasets and machine learning models to build risk-informed suitability surfaces and predict outage probability and cooling energy demand. Shows that modest cost premiums for resilient locations can be offset by reduced disruption and lower energy expenditure over asset lifetimes. Focuses on siting and operation but does not quantify portfolio-level financial tail risks.

Analogous Infrastructure with Transferable Methods

Table 22. Representative climate risk studies for analogous infrastructure

Study	Infrastructure / hazards	Methodological insights
Dawson et al. (2018)	Energy, water, transport, ICT (UK)	Provides a systems framework for national climate risk assessment across infrastructure sectors, including ICT. Identifies flooding as the greatest risk to all sectors and projects a doubling of users reliant on electricity infrastructure at flood risk under 2 °C, tripling under 4 °C. Highlights interdependence between energy and ICT; notes limited asset-level performance data.
Forzieri et al. (2016)	Critical infrastructure in Europe; heat, drought, coastal floods	Projects that damages to European critical infrastructures could triple by the 2020s, increase six-fold by mid-century, and exceed 10 times present values by end of century. Identifies heatwaves, droughts, and coastal floods as showing the most dramatic rise. Uses pan-European hazard maps and generic vulnerability functions, with limited sector-specific calibration.
Oughton et al. (2023)	Global mobile telecommunications (7.6M assets); tropical cyclones, coastal flooding	Performs a global vulnerability assessment of mobile telecom infrastructure using crowdsourced asset data and climate projections under RCP 8.5. Estimates that a 0.01% annual probability tropical cyclone event could affect 2.26 million cells (USD 1.01 billion damage, 44% increase), and coastal flooding could affect 109.9 thousand cells (USD 2.69 billion damage, 78% increase). Demonstrates a scalable methodology directly transferable to digital infrastructure inventories.
Luo et al. (2023)	Thermal and hydro power plants; multi-hazard	Develops a multi-hazard risk framework for power generation projects across Eastern Europe, the Mediterranean, and Central Asia. Uses publicly available climate data and technology-specific vulnerability factors to estimate generation losses under RCP 4.5 and RCP 8.5. Highlights rising water temperatures as a key overlooked risk; does not explicitly monetise all operational impacts.

Cross-cutting Limitations

Across Q1 studies, several limitations are recurrent: coarse resolution of global hazard maps compared to facility-level siting decisions; reliance on generic damage functions not calibrated for data center equipment, cooling systems, or power infrastructure; limited treatment of indirect and cascading impacts (e.g. supply chain disruptions, grid failures); and under-representation of electric and telecommunications networks relative to other sectors. Data centers are rarely treated explicitly despite their growing systemic importance.

Q2 – Insurance, Insurability and Climate Risk for Critical Infrastructure

The literature on insurance and insurability in the context of infrastructure climate risk is sparse, and almost entirely absent for data centers and digital infrastructure. Of the 25 sources screened for insurance content, only one study (Kerkhofs et al., 2025) provides a substantive quantitative analysis of insurance mechanisms and their effect on portfolio tail risks. Most infrastructure climate risk studies do not address premiums, coverage terms, market withdrawals, or shifts of risk from private insurers to public budgets. No work was found that documents protection gaps, premium trajectories, or coverage restrictions specifically for data centers or comparable high-value digital assets.

Available Evidence on Insurance Mechanisms

Table 23. Insurance-related findings in reviewed studies

Study	Context	Insurance-related contributions
Kerkhofs et al. (2025)	Power firms (India, Thailand)	Develops an asset-level framework for physical risk assessment and translation to portfolio-level impacts, including explicit treatment of insurance. Uses Monte Carlo simulation with spatial correlations and Value-at-Risk metrics to quantify tail risks. Shows that insurance has a limited effect on average portfolio losses but materially affects tail risk distribution. Explores scenarios in which insurers withdraw from high-risk markets, increasing residual risk for asset owners.
Karagiannakis et al. (2024)	Power grid fragility	Discusses use of fragility curves by insurers for risk assessment and life-cycle costing of grid assets. Does not provide empirical analysis of premium trends, coverage restrictions, or protection gaps.
Other reviewed studies	Various infrastructure sectors	Do not include explicit insurance analysis. Insurance is occasionally mentioned qualitatively but without quantitative evidence on premiums, coverage, or protection gaps for critical infrastructure.

Identified Gaps for Data Centers

For data centers and digital infrastructure, the review did not identify any peer-reviewed work that:

- Quantifies protection gaps (insured vs. uninsured losses) at asset or portfolio level.
- Analyses premium dynamics (trend, volatility, deductibles) for climate-exposed facilities.
- Documents coverage withdrawals, non-renewals, or market capacity constraints for large digital campuses.
- Examines how climate-driven insurance conditions affect bankability, cost of capital, or asset valuation for data centers.

This absence is a critical limitation for any climate-economic framework that aims to integrate insurance dynamics into data center infrastructure valuation.

Q3 – Quantitative Frameworks to Integrate Climate Risk into Infrastructure Valuation

A growing literature proposes quantitative frameworks for integrating climate risk into infrastructure investment and valuation, but none are tailored to data centers. Existing approaches include probability-of-default and DSCR metrics for energy assets, real options analysis for flood protection, discounted cash flow models with climate adjustments, Expected Annual Damage coupled with stress testing, and storyline-based frameworks linking climate scenarios to asset damages.

Common elements are scenario-based hazard projections, hazard-to-impact translation via vulnerability functions, and financial metrics such as NPV, EAD, or VaR. Key methodological gaps for application to data centers include the absence of data center-specific vulnerability functions, incomplete representation of indirect and cascading impacts, and limited treatment of compound multi-hazard risks.

Representative Valuation Frameworks

Table 24. Representative quantitative valuation frameworks

Study	Framework type	Relevance and limitations
In et al. (2020, 2021, 2022)	Probability of default / DSCR for energy infrastructure	Develop a framework for pricing climate-related risks of energy investments, using DSCR and probability of default rather than NPV alone. Combine asset-specific physical and transition risks with financing structure to estimate timing and size of losses. Demonstrate that renewable projects are more resilient than fossil assets. Applied only to energy; does not consider digital infrastructure or supply-chain effects.
Prime et al. (2018)	Real options for flood protection	Applies real options analysis to flood defense investment around electricity infrastructure, using UKCP09 sea-level projections to 2100. Values flexibility in investment timing under climate uncertainty and shows which sites benefit from early vs. delayed adaptation. Could be adapted for data center flood defenses; currently focused on coastal electricity assets.
Neumann et al. (2015)	DCF-based integrated assessment across US infrastructure	Combines sector models and discounted cash flow (3% rate) to estimate climate impacts and adaptation benefits for roads, bridges, coastal development, and drainage. Finds that mitigation policies reduce cumulative impacts by 25–35%, with avoided costs exceeding \$460 billion through 2100. Electric and telecommunications networks are explicitly excluded.
Pasha et al. (2025)	Flood risk stress testing (EAD) for solar power	Implements a climate risk stress test based on the hazard–vulnerability–exposure approach, using global flood hazard data and depth–damage functions to derive Expected Annual Damage for a solar plant. Strength is transparency and replicability; limitations include coarse hazard resolution, generic damage functions, and omission of indirect financial impacts.
Koks et al. (2022)	Storyline framework for coastal infrastructure	Combines sea-level rise projections, inundation modelling, and infrastructure damage assessment into a storyline framework. Shows that coordinated adaptation can reduce coastal flood impacts by more than 70%. Limited to two adaptation strategies and coastal flooding, but illustrates how physical scenarios can be linked to discrete adaptation choices.

Implications for Data Center Valuation

Taken together, these frameworks demonstrate that climate risk can be integrated into infrastructure valuation through adjusted DCF, probability-of-default metrics, EAD-based stress testing, and real options. However, no framework in the reviewed literature has been designed for, or validated on, data center infrastructure. In particular, existing work does not capture:

- Cooling-system dependencies and power quality requirements specific to high-density compute.
- Operational and business-interruption impacts beyond physical asset damage, including service continuity requirements.
- Insurance and insurability dynamics as explicit financial channels.
- Cascading economic impacts of data center outages on downstream users and services.
- Multi-hazard compound scenarios (e.g. extreme heat, drought, and grid stress occurring jointly).



Schneider Electric™

Research Institute

Advancing Energy Tech

www.se.com/sustainability-research-institute