Sustainability and ESG Crypto Connectedness

Auteur : Cange, Samuel
Promoteur(s) : Santi, Caterina
Faculté : HEC-Ecole de gestion de l'Université de Liège
Diplôme : Master de spécialisation en gestion des risques financiers
Année académique : 2024-2025
URI/URL : http://hdl.handle.net/2268.2/24906

Sustainability and ESG–Crypto Connectedness

Promoteur :
Santi Caterina

Lecteur :
Clerc Pierrick

Travail de fin d'études présenté par :
Samuel CANGE

En vue de l'obtention du diplôme de Master de spécialisation en gestion des risques financiers, option Sustainable and Climate Finance

Année académique 2024–2025

Dedication

To my parents for their love and sacrifices, to my family for their patience and faith, and to my friends for their unwavering encouragement. To Prof. Caterina Santi for her generous guidance and exacting standards. To the HEC Liège community—and to all who work to align finance with sustainability—this work is dedicated with gratitude.

Acknowledgements

I am deeply grateful to my supervisor, Prof. Caterina Santi, for her guidance, high standards, and steady support throughout this project. Her probing questions and clear feedback shaped both the scope and the rigor of this work.

My thanks extend to the examination committee for their careful reading and constructive suggestions, and to the faculty and staff of HEC Liège's Specialized Master in Sustainable Finance for an intellectually demanding and supportive learning environment. I also appreciate the insights shared by guest lecturers and practitioners who bridged theory and practice.

To my classmates in the Specialized Master in Sustainable Finance cohort: thank you for the collaboration, debates, and encouragement. Finally, I owe the greatest debt to my family and close friends for their patience, faith, and unwavering support. Any remaining errors are mine alone.

Abstract

This thesis examines the intersection of sustainable finance and digital assets, focusing on the dynamic linkages between Environmental, Social, and Governance (ESG) instruments and cryptocurrencies. Using a combination of time-varying connectedness models, spillover analyses, and environmental impact assessments, the research explores how shocks propagate between ESG asset classes and both mainstream and so-called “green” cryptocurrencies.

The study integrates environmental footprint estimates of blockchain protocols, contrasting energy-intensive Proof-of-Work systems with more sustainable consensus mechanisms such as Proof-of-Stake and Federated Byzantine Agreement. It also considers market microstructure frictions, including cross-venue segmentation and liquidity constraints, that affect the integration of cryptoassets into sustainable portfolios.

The findings contribute to the design of a policy-aligned allocation framework for sustainable cryptoasset exposure. This framework balances financial performance with ESG compliance, incorporates regime-dependent hedging strategies, and aligns with evolving disclosure standards such as ESRS, ISSB, and GRI. By bridging sustainable finance principles with the realities of the digital asset ecosystem, the thesis offers practical insights for institutional investors, policymakers, and asset managers navigating the emerging field of sustainable crypto finance.

Contents

List of Figures

List of Tables

List of Acronyms

Acronym	Definition
ARCH	Autoregressive Conditional Heteroskedasticity
BEKK	Baba–Engle–Kraft–Kroner (multivariate GARCH)
CoVaR	Conditional Value at Risk
CSRD	Corporate Sustainability Reporting Directive
DCC	Dynamic Conditional Correlation
dMRV	Digital Monitoring, Reporting, and Verification
EIP	Ethereum Improvement Proposal
ESG	Environmental, Social, and Governance
ESRS	European Sustainability Reporting Standards
ETP	Exchange-Traded Product
FBA	Federated Byzantine Agreement
FEVD	Forecast Error Variance Decomposition
GRI	Global Reporting Initiative
GARCH	Generalised ARCH
HHI	Herfindahl–Hirschman Index
IFRS	International Financial Reporting Standards
ISSB	International Sustainability Standards Board
LPPLS	Log-Periodic Power-Law Singularity
MEV	Maximal Extractable Value
MRV	Monitoring, Reporting, and Verification
PBS	Proposer-Builder Separation
PCI	Pairwise Connectedness Index
PPI	Producer Price Index
PoR	Proof of Reserves
PoS	Proof-of-Stake
PoW	Proof-of-Work
REC	Renewable Energy Certificate

SLA	Service-Level Agreement
TCI	Total Connectedness Index
TVP–VAR	Time-Varying Parameter Vector Autoregression
VaR	Value at Risk
WEEE	Waste Electrical and Electronic Equipment

Chapter 1

Introduction

1.1 Motivation and Context

Over the past decade, sustainable finance has undergone a marked transformation. Once confined to specialised funds and niche investor groups, environmental, social, and governance (ESG) principles now occupy a central role in global capital allocation. This mainstreaming is not merely a matter of investor preference; it reflects structural changes in regulation, market incentives, and societal expectations. Parallel to this evolution, digital assets—initially perceived as speculative novelties—have matured into a globally traded, high-volatility asset class with deep liquidity and increasingly sophisticated market infrastructure.

The intersection of these two domains is both inevitable and complex. Investors, asset managers, and regulators are increasingly faced with the challenge of integrating ESG considerations into portfolios that include cryptoassets and crypto-adjacent securities. This integration is not trivial: cryptoassets exhibit unique risk factors, such as high-frequency volatility clustering, market-structure segmentation, and technological evolution in consensus protocols, that interact in non-linear ways with traditional ESG metrics.

The regulatory environment reinforces this convergence. The European Union's Corporate Sustainability Reporting Directive (CSRD) and European Sustainability Reporting Standards (ESRS), the International Sustainability Standards Board's IFRS S1 and S2, and the U.S. Securities and Exchange Commission's climate-related disclosure proposals collectively elevate the standard of sustainability reporting. These frameworks do not simply demand more disclosure; they embed ESG criteria into the core of portfolio design, making non-compliance a potential source of both reputational and regulatory risk.

From a portfolio-management perspective, a credible framework for "sustainable crypto" investing must therefore rest on three mutually reinforcing pillars:

- Empirically robust evidence on the connectedness between ESG assets and cryptoassets, capturing both directionality and regime dependence.
- Rigorous environmental and governance measurement, with verifiable metrics on blockchain energy consumption, emissions, e-waste, and governance structures.
- Governance-aware risk controls capable of addressing market-structure frictions, such as persistent cross-exchange price differentials and jurisdiction-specific capital-flow restrictions.

Empirical evidence supports the need for such a framework. Time-varying parameter vector autoregression (TVP–VAR) models show that connectedness between ESG equity leaders and major cryptocurrencies is substantial and highly state-dependent, intensifying during systemic stress events such as the onset of COVID-19 or the Russia–Ukraine conflict. Negative shocks, in particular, propagate more forcefully than positive ones, eroding diversification benefits at precisely the moments when they are most needed, as Bibi et al. (2025) study. Disaggregated analyses reveal that ESG indices often act as net transmitters of return and volatility shocks, while green cryptocurrencies—defined as assets running on low-energy consensus mechanisms like Proof-of-Stake—tend to be net receivers in normal conditions. However, hedging effectiveness declines sharply during crises, and transaction costs rise as hedge ratios adjust more frequently, as Alharbi et al. (2025) and Ali et al. (2024) show.

Market-microstructure realities complicate implementation. Persistent cross-exchange price differentials, driven by capital controls, fiat on/off-ramp bottlenecks, and venue segmentation, introduce execution frictions that directly affect hedge feasibility and cost-efficiency, as Makarov and Schoar (2020) document. Without accounting for these frictions, portfolio models risk overstating achievable diversification.

The environmental dimension is equally central. Comparative life-cycle assessments show that Proof-of-Stake protocols consume orders of magnitude less energy than Proof-of-Work systems, with the latter also producing significant electronic waste and being prone to carbon leakage when mining relocates after adverse regulatory shocks. Standardised frameworks for assessing blockchain sustainability—including Wh per transaction, annual MWh consumption, CO_2 emissions, renewable-energy share, and e-waste intensity—are beginning to emerge, but methodological heterogeneity and inconsistent reporting practices hinder comparability, as Wendl et al. (2023) and Alzoubi and Mishra (2023) note.

Beyond asset-level sustainability, blockchain technology itself holds potential as “ESG infrastructure,” enabling real-time monitoring, reporting, and verification (MRV) as well as supply-chain traceability. Yet operationalising these capabilities within recognised frameworks such as ESRS, ISSB, or GRI remains inconsistent, as Mulligan et al. (2024) and Figueiredo et al. (2022) observe.

Finally, crypto-adjacent instruments such as green bonds may serve as more effective diversifiers of blockchain-equity risk than clean-energy equities, particularly in crisis regimes where the latter’s co-movement with crypto markets intensifies, as Mzoughi et al. (2024) and Ul Haq et al. (2023) find.

Taken together, these insights frame the motivation for this thesis: to synthesise the fragmented empirical literature into a policy-compliant, ESG-integrated risk framework for cryptoasset exposure that addresses environmental measurement, market-structure governance, and regime-dependent connectedness.

1.2 Problem Statement

The central research challenge addressed in this thesis is the formulation of a regime-resilient, empirically grounded, and ESG-aligned framework for allocating cryptoassets within sustainable investment portfolios. Such a framework must operate under three interrelated constraints:

(1) Regime-dependent and asymmetric ESG–crypto connectedness. Empirical studies consistently demonstrate that linkages between ESG assets and cryptoassets are both time-varying and shock-asymmetric. Diversification benefits often vanish in systemic stress regimes, as negative-shock connectedness exceeds positive-shock connectedness, as Bibi et al. (2025) show. While ESG indices and developed-market equities generally act as net transmitters of shocks under normal conditions, green cryptocurrencies often serve as net receivers, according to Alharbi et al. (2025) and Ali et al. (2024). In crisis episodes, these transmission roles may partially reverse, particularly during crypto-specific shocks such as exchange failures or regulatory interventions. Static-correlation assumptions in practitioner risk models fail to capture these dynamics, leading to underestimated drawdown risk.

(2) Heterogeneous and uncertain environmental footprints of blockchain protocols. Proof-of-Stake and other energy-efficient consensus mechanisms exhibit orders-of-magnitude lower energy consumption than Proof-of-Work systems. However, cross-study variability in reported figures—driven by differences in measurement scope, hardware assumptions, and network-usage profiles—introduces uncertainty into ESG scoring. This variability complicates compliance with disclosure frameworks such as ESRS, ISSB, or GRI, which require credible, verifiable, and comparable environmental metrics, as Wendl et al. (2023), Alzoubi and Mishra (2023), and Mulligan et al. (2024) emphasise.

(3) Market-structure frictions and policy constraints. Persistent cross-venue price segmentation, "Kimchi premiums," fiat-rail limitations, and jurisdiction-specific capital controls introduce basis risk that can erode hedging effectiveness, as Makarov and Schoar (2020) document. Moreover, environmental policy interventions—such as mining bans—can produce unintended consequences like carbon leakage, shifting high-emission mining operations to less-regulated jurisdictions. These frictions are rarely accounted for in operational portfolio frameworks, despite their material impact on execution costs, slippage, and overall governance compliance.

Addressing these constraints requires bridging identifiable gaps in both academic literature and practitioner guidance. While empirical finance has produced robust connectedness models, there remains a lack of integration with ESG assurance protocols and governance-aware portfolio rules. Similarly, environmental assessments of blockchain technology are often conducted in isolation from asset-allocation implications. Finally, the literature on market-microstructure effects in crypto markets rarely engages with sustainable-finance objectives, leaving open the question of how to design risk budgets that are both ESG-compliant and operationally feasible.

This thesis responds to these gaps by synthesising current empirical findings, mapping them explicitly to recognised ESG disclosure frameworks, and embedding them within a governance-aware, state-contingent portfolio-construction approach for cryptoassets.

1.3 Objectives and Research Questions

The overarching objective of this thesis is to consolidate fragmented empirical research on ESG–crypto connectedness, blockchain environmental performance, and market-structure governance into a unified, policy-aligned, and investable framework for sustainable cryptoasset allocation. This framework is designed to function across multiple market regimes, integrating empirical financial-econometrics with ESG assurance standards and execution-feasibility considerations.

Two primary objectives structure the research design:

(1) Evidence Synthesis. Conduct a systematic and critical review of peer-reviewed empirical studies covering: (a) the connectedness between ESG assets and cryptocurrencies, including directionality, asymmetry, and regime dependence; (b) environmental-footprint metrics of blockchain consensus mechanisms, incorporating measurement uncertainty and assurance protocols; (c) market-structure and policy frictions affecting hedging feasibility and execution costs; and (d) macro-financial sensitivities relevant to conditional risk management. The synthesis aims to identify patterns, contradictions, and methodological gaps that affect portfolio-construction decisions.

(2) Framework Design. Translate the synthesised evidence into a multi-layer portfolio-construction architecture comprising:

- ESG-aligned asset selection criteria incorporating environmental, governance, and macro-risk dimensions.
- State-contingent hedging rules informed by time-varying connectedness and quantile-dependent macro sensitivities.
- Governance mechanisms addressing execution frictions, policy risk, and environmental-impact verification.
- Risk-reporting protocols explicitly mapped to ESRS, ISSB, and GRI disclosure requirements to ensure policy-compliant transparency.

Within this structure, the research is guided by the following questions:

Connectedness Dynamics: What is the magnitude, directionality, and asymmetry of ESG–crypto connectedness under stable versus crisis conditions, and how should these dynamics inform portfolio weights, hedge ratios, and hedging effectiveness?

Environmental Footprints: Which consensus mechanisms and blockchain platforms meet credible environmental thresholds, and how should uncertainty in footprint measurement be incorporated into portfolio limits, ESG scoring, and public disclosures?

Market-Structure Frictions: How do capital controls, fiat-currency rails, and venue segmentation influence hedging feasibility, and what governance measures are effective in mitigating these execution and policy risks?

Macroeconomic Sensitivities: Which macroeconomic variables—such as producer price indices, energy prices, USD liquidity, and interest-rate dynamics—exert quantile-dependent effects on cryptocurrency returns, and how can these sensitivities be operationalised as dynamic risk triggers?

ESG-Aligned Diversifiers: Among ESG-aligned instruments, which assets best diversify crypto-linked exposure under systemic-stress regimes, and under what conditions do green bonds outperform clean-energy equities in this role?

This dual focus on synthesis and design ensures that the proposed framework is both empirically grounded and operationally implementable, enabling institutional investors to reconcile ESG objectives with cryptoasset exposure in a robust and transparent manner.

1.4 Scope, Positioning, and Contributions

This thesis employs an evidence-synthesis methodology grounded exclusively in secondary data. All empirical material is drawn from peer-reviewed studies, authoritative institutional reports, and reputable working-paper series. No new econometric estimations are performed; instead, the research integrates and critically analyses existing quantitative and qualitative findings to design a policy-aligned sustainable-crypto allocation framework.

Scope. The analysis encompasses:

- ESG indices in both aggregate and disaggregated forms (environmental, social, governance subcomponents).
- Green cryptocurrencies, defined as digital assets operating under energy-efficient consensus protocols such as Proof-of-Stake (PoS) or Federated Byzantine Agreement (FBA), with Proof-of-Work (PoW) assets retained for comparative benchmarking.
- Crypto-adjacent securities, including blockchain-thematic equities and green bonds.
- Relevant macroeconomic and policy variables such as energy prices, producer price indices, USD exchange rates, and sovereign yields.

Positioning. The research situates itself at the intersection of:

- (i) Sustainable finance, by embedding environmental and governance constraints into asset selection and allocation rules;
- (ii) Digital-asset portfolio management, through the application of time-varying connectedness models, quantile-based sensitivity analysis, and hedging-effectiveness metrics; and

- (iii) Policy-oriented ESG assessment, by mapping all proposed metrics and governance mechanisms to internationally recognised reporting frameworks such as ESRS, ISSB, and GRI.

Contributions. The thesis advances the literature and practitioner toolkit in four primary ways:

- **Integrative Connectedness Mapping**—synthesises empirical findings on ESG–crypto linkages across normal and crisis regimes, highlighting asymmetric contagion patterns and proposing macro-conditioned overlays to account for quantile-dependent sensitivities.
- **Consolidated ESG Metric Set**—develops a unified set of environmental and governance indicators (e.g., Wh per transaction, annual MWh consumption, CO_2 emissions, renewable-energy share, e-waste intensity, validator concentration) accompanied by uncertainty bands and third-party assurance requirements.
- **Governance-Aware Portfolio Playbook**—outlines execution strategies that address market segmentation, basis risk, and capital-flow restrictions, incorporating a shock taxonomy distinguishing endogenous from exogenous events to inform drawdown controls.
- **Policy-Consistent Diversifier Selection**—defines criteria under which green bonds provide superior hedging benefits relative to clean-energy equities in blockchain-thematic portfolios, while identifying crisis conditions where the latter may amplify downside risk.

Through this scope and positioning, the thesis aims to deliver a framework that is simultaneously empirically grounded, operationally implementable, and policy compliant, enabling institutional investors to reconcile ESG commitments with cryptoasset exposure while managing regime-specific risks.

1.5 Conceptual Framework (Thesis Lens)

The conceptual framework underpinning this research is conceived as a multi-layer architecture integrating empirical evidence into a coherent portfolio-construction model for sustainable cryptoasset exposure.

At the market level (Layer A), time-varying connectedness studies reveal that inter-asset relationships in returns and volatility are asymmetric and regime-dependent.

Negative shocks typically induce stronger co-movement than positive shocks, and systemic stress—such as the COVID-19 onset or the Russia–Ukraine escalation—produces marked spikes in ESG–crypto and G7–green-crypto connectedness, elevating hedge ratios and reducing diversification benefits, as Bibi et al. (2025), Alharbi et al. (2025), and Ali et al. (2024) show.

The macroeconomic conditioning layer (Layer B) incorporates evidence from quantile-regression analyses showing that sensitivities of cryptoassets to producer price indices, energy prices, USD/CNY exchange rates, and sovereign yields vary across market states, often reversing in sign between low- and high-volatility regimes. These results justify the inclusion of state-contingent macro overlays in hedging rules, as Lin et al. (2025) demonstrate.

At the ESG-measurement and assurance layer (Layer C), environmental footprints differ substantially by consensus protocol, validator hardware, and geographic location. Proof-of-Stake systems consistently display orders-of-magnitude lower energy use per transaction than Proof-of-Work systems. Measurement variability across studies demands explicit reporting of uncertainty ranges and third-party verification. Governance structures—permissioning models, validator concentration, and on-chain voting mechanisms—must be assessed in conjunction with environmental metrics, as Wendl et al. (2023), Alzoubi and Mishra (2023), and Mulligan et al. (2024) report.

The market-microstructure and policy layer (Layer D) addresses persistent cross-border price differentials, such as the "Kimchi premium," arising from segmentation, capital controls, and fiat-rail constraints, which erode hedge feasibility. Policy interventions, including mining bans or moratoria, may exacerbate environmental risks via carbon leakage and relocation of energy-intensive operations, as Makarov and Schoar (2020) and Wendl et al. (2023) discuss.

Finally, the portfolio-construction layer (Layer E) integrates these insights into allocation and hedging rules that are connectedness-aware, ESG-tilted, and governed by macro-conditioned triggers. Green bonds are identified as effective diversifiers of blockchain-equity risk under stress regimes, while clean-energy equities may co-move more strongly with crypto markets during crises, thereby amplifying downside exposure, consistent with Bibi et al. (2025), Mzoughi et al. (2024), and Ul Haq et al. (2023).

Chapter 2

Literature Review

This chapter does more than review prior work: it connects what the literature *measures* to what an investor must *do*. We take the main empirical signals—time-varying connectedness, horizon-specific co-movement, and tail risk—and turn them into decision variables the portfolio can act on: regime flags, thresholds that trigger hedges or caps, and rules for how much protection to buy and when.

Two gaps motivate this approach. First, many studies stop at description: they report higher connectedness in crises and stronger downside spillovers, but do not say which levels should change weights or limit exposures, or how to penalise systemic linkages in an optimiser. Second, feasibility is often ignored. Real trading faces venue segmentation, basis risk, custody and disclosure constraints, and uneven data quality. If a signal cannot be executed, audited, or verified, it should not drive risk.

Our response is pragmatic. We extract regime indicators and transmitter/receiver roles to initialise triggers; use environmental evidence to set eligibility screens and budgets (e.g., footprint caps, reporting requirements); and map microstructure realities into turnover, liquidity, and jurisdictional constraints. Where inputs are weak or unverified, the default is conservative sizing or ineligibility. The rest of the chapter follows this logic: it distils the evidence into actionable signals and guardrails, which later chapters embed in the allocation and hedging framework with explicit thresholds, persistence filters, and reporting rules.

2.1 Overview and Organisation

This chapter synthesises four strands of literature relevant to sustainable allocation of cryptoassets: (i) theoretical and empirical work on cross-market connectedness; (ii) environmental performance of blockchain consensus mechanisms and measurement/assurance practices; (iii) market-microstructure frictions and policy constraints in digital-asset markets; and (iv) the role of ESG-aligned instruments as diversifiers for blockchain-thematic exposures. Throughout, we emphasise regime dependence, shock asymmetry, and the implications for portfolio construction, risk reporting, and compliance with recognised sustainability frameworks.

2.2 Theoretical Foundations: Risk, Diversification, and Connectedness

Classical portfolio theory motivates diversification across imperfectly correlated assets (Markowitz, 1952, 1959), with performance evaluated using mean–variance and reward-to-variability metrics (Sharpe, 1994). For multi-asset systems subject to state dependence and nonlinearity, impulse-response analysis in linear and nonlinear multivariate settings clarifies transmission channels (Koop et al., 1996; Pesaran & Shin, 1998). In crypto-focused applications, time-varying parameter vector autoregressions (TVP–VAR) and related connectedness measures capture evolving shock spillovers and their asymmetries (Antonakakis et al., 2020; Chatziantoniou & Gabauer, 2021; Gabauer, 2021).

Empirical results consistently show that ESG leaders and cryptocurrencies exhibit substantial, state-dependent connectedness, with negative shocks propagating more strongly than positive ones. For example, Bibi et al. (2025) document asymmetric TVP–VAR connectedness between ESG leaders and major cryptocurrencies, with diversification benefits eroding during systemic stress. Studies concentrating on "green" cryptocurrencies (operating under low-energy consensus) report that these assets tend to be net receivers of shocks in tranquil regimes but provide limited hedging efficacy during crises (Alharbi et al., 2025; Ali et al., 2024).

2.3 ESG–Crypto Connectedness: Evidence and Implications

Table 2.1 summarises prominent empirical contributions, including sample periods, methods, and findings on directionality, asymmetry, and regime dependence. A unifying pattern is that connectedness rises sharply around market-wide disruptions (e.g., pandemic onset, geopolitical escalations), compressing cross-asset hedging gains precisely when they are most valued. This implies that static correlation assumptions are untenable for risk budgeting in sustainable portfolios that include cryptoassets.

Table 2.1: Selected studies on ESG–crypto connectedness and spillovers

Study	**Assets / Sample**	**Methodological Core**	**Main Findings**
Bibi et al. (2025)	ESG leaders vs. major cryptocurrencies; multi-year panel	TVP–VAR, asymmetric connectedness, regime analysis	Strong, state-dependent connectedness; negative-shock dominance; diversification fades in crises.
Alharbi et al. (2025)	ESG equities vs. “green” cryptocurrencies	Spillover indices, hedging effectiveness	ESG indices are net transmitters in normal times; green cryptos net receivers; hedging worsens in stress.
Ali et al. (2024)	G7 equities vs. green cryptocurrencies	Return/volatility spillovers	Intensified spillovers during turmoil; clean-energy equities comove with crypto under stress.
Gabauer (2021)	Major cryptocurrencies	TVP–VAR connectedness	Time-varying spillovers prominent; role switching across regimes.
Antonakakis et al. (2020)	Cross-market systems	Refined TVP–VAR connectedness	Method improves inference under time variation; supports regime-aware hedging.
Chatziantoniou and Gabauer (2021)	Connectedness metrics	Monte Carlo reassessment of TCI bounds	Adjusted TCI recommended for accurate scaling and interpretation.

2.4 Environmental Performance of Blockchain Protocols

2.4.1 Consensus Mechanisms and Footprint Metrics

Life-cycle and system-boundary choices drive considerable heterogeneity in reported energy and emissions metrics for blockchain networks. Systematic reviews converge on the conclusion that Proof-of-Stake (PoS) protocols consume orders of magnitude less energy than Proof-of-Work (PoW), with additional concerns for PoW regarding e-waste and carbon leakage via relocations after policy shocks (Alzoubi & Mishra, 2023; Wendl et al., 2023). Integrating these metrics into investment processes requires verifiable, comparable indicators (e.g., Wh/tx, annual MWh, CO_2e, renewable share, e-waste intensity) and explicit uncertainty disclosures.

Table 2.2: Qualitative comparison of environmental and governance attributes by consensus type

Attribute	**PoW**	**PoS**	**FBA / Permissioned BFT**
Relative energy intensity	High	Very low	Very low
CO_2e exposure to grid mix	High (location-sensitive)	Low–moderate	Low
E-waste potential	Significant (hardware churn)	Minimal	Minimal
Validator concentration/governance risk	Variable	Variable	Depends on permissioning
MRV[a] readiness (on-chain/off-chain)	Emerging	Emerging	High (enterprise implementations)

[a] MRV = monitoring, reporting, and verification. Synthesis based on Alzoubi and Mishra (2023), Mulligan et al. (2024), and Wendl et al. (2023).

Table 2.2 summarises environmental and governance attributes across PoW, PoS, and permissioned BFT variants.

2.4.2 Measurement, Assurance, and Policy Alignment

Methodological variance across studies—differences in hardware assumptions, network utilisation, and boundary definitions—creates nontrivial uncertainty for ESG scoring and disclosure (Wendl et al., 2023). Policy-oriented reviews emphasise the role of blockchain as sustainability infrastructure (e.g., MRV systems and traceability), though operationalisation within ESRS/ISSB/GRI remains uneven (Figueiredo et al., 2022; Mulligan et al., 2024). For investment governance, uncertainty bands and third-party assurance should accompany environmental metrics to preserve comparability and auditability.

Table 1
Estimates of energy consumption of PoW cryptocurrencies (own representation).

Author	Publishing Date	Focus	Methodology	2014	2015	2016	2017	2018	2019	2020	2021	2022
Digiconomist	Daily calculation since 2017; Data for 2022 retrieved on June 19, 2022	Bitcoin Bitcoin Cash	Top-down & Economic				**LB 7,37 UB 69,01**	**LB 40,42 UB 122,91**	**LB 43,97 UB 93,72**	**LB 49,13 UB 95,98**	**LB 45,69 UB 332,91**	**LB 65,32 BE 130,24**
CBECI	Daily calculation since 2011; Data retrieved from 2014 until June 19, 2022	Bitcoin	Hybrid top-down	**BE 3,76**	**BE 2,46**	**BE 8,24**	**BE 16,92**	**BE 45,43**	**BE 57,1**	**BE 68,51**	**BE 104,89**	**BE 107,55**
Küfeoglu and Özkuran (2019)	June 2018	Bitcoin	Hybrid top-down					**LB 15.47 UB 50,24**				
De Vries (2018)	May 2018	Bitcoin	Top-down &Economic				**LB 22 UB 67**					
De Vries (2019)	April 2019	Bitcoin	Top-down & Economic					**LB 40 UB 62,3**				
De Vries (2020)	December 2020	Bitcoin	Economic						**BE 87,1**			
Stoll et al., 2019a	July 2019	Bitcoin	Top Down & Economic					**BE 48,20**				
Li et al., 2019	February 2019	Monero	Extrapolation					**BE 0,65**				
Krause and Tolaymat (2018)	November 2018	Bitcoin, Ethereum, Litecoin, Monero	Top-down			**BE 2,85**	**BE 20,28**	**BE 44,01**				
Köhler and Pizzol (2019)	November 2019	Bitcoin	Top-down					**BE 31,29**				
Seidlmeier et al., 2020a	February 2020	Bitcoin, Ethereum	Top-down &Economic							**LB 60 UB 125**		
O'Dwyer and Malone, 2014	January 2014	Bitcoin	Top-down	**LB 0,88 UB 87,6 BE 26,28**								
McCook, 2018	July 2018	Bitcoin	Top Down					**BE 105,82**				

Estimates in TWh - LB = Lower-bound estimate - UB = Upper-bound estimate - BE = Best estimate.

Figure 2.1: Estimated annual energy consumption of major PoW cryptocurrencies.

Note. Adapted from Wendl et al. (2023). Estimates are expressed in terawatt-hours (TWh). LB = Lower Bound; UB = Upper Bound; BE = Best Estimate.

As shown in Figure 2.1, PoW networks exhibit orders-of-magnitude higher annual energy consumption than PoS platforms, motivating protocol-level screening.

2.5 Market-Microstructure and Policy Frictions

Crypto markets exhibit persistent cross-venue segmentation and fiat on/off-ramp frictions that generate basis risk and complicate execution. Makarov and Schoar (2020) document durable cross-exchange price differentials attributable to capital controls and venue segmentation, challenging assumptions of frictionless hedging. Such frictions intensify during regulatory interventions or exchange-specific shocks, with implications for hedge timing, slippage, and capital efficiency.

Table 2.3: Taxonomy of microstructure and policy frictions relevant to sustainable crypto allocation

Friction	**Mechanism / Portfolio Impact**	**Representative Evidence**
Venue segmentation & capital controls	Persistent cross-exchange premia (e.g., "Kimchi premium"); hedging basis risk; execution delays	Makarov and Schoar (2020)
Fiat rail bottlenecks	Slow settlement; unstable liquidity; higher slippage in rebalancing	Makarov and Schoar (2020)
Policy shocks (e.g., mining bans)	Relocation of hash power; potential carbon leakage; footprint uncertainty	Wendl et al. (2023)
Exchange-specific outages/failures	Transient dislocations; impaired hedging; gap risk	Fry and Cheah (2016)

Notes: Effects often intensify in stress regimes; governance controls should reflect execution feasibility and reporting obligations.

2.6 Macroeconomic Sensitivities and State Dependence

Quantile-based evidence indicates that macro determinants of cryptocurrency returns (e.g., producer prices, energy benchmarks, USD liquidity, and sovereign yields) exert state-dependent effects that can switch sign across volatility regimes. Lin et al. (2025) use quantile regression to show such conditional sensitivities, supporting macro-conditioned overlays in risk controls. These results dovetail with connectedness evidence, suggesting that both cross-asset spillovers and macro betas strengthen under stress—a dual mechanism of diversification decay.

2.7 ESG-Aligned Diversifiers for Blockchain-Thematic Risk

The literature comparing green bonds and clean-energy equities as diversifiers for crypto-linked exposures points to regime-specific performance. Evidence suggests that green bonds can provide more reliable downside protection during systemic stress, whereas clean-energy equities may co-move more tightly with crypto and broader risk assets in crises (Mzoughi et al., 2024; Ul Haq et al., 2023). Portfolio implementations should therefore condition diversifier selection on market regime, liquidity, and issuance characteristics.

2.8 Methodological Considerations and Best Practices

Reliable inference in connectedness studies hinges on model specification, stability diagnostics, and shock identification. TVP–VAR frameworks with robust connectedness metrics (Antonakakis et al., 2020; Chatziantoniou & Gabauer, 2021) and impulse-response analysis attuned to nonlinearity (Koop et al., 1996; Pesaran & Shin, 1998) are now standard. While classical diagnostics (e.g., normality of residuals) are not always strictly required for consistent connectedness estimation, transparent reporting of distributional features remains good practice (Anscombe & Glynn, 1983; Jarque & Bera, 1980). Hedging-effectiveness metrics should explicitly incorporate transaction costs and execution frictions to avoid overstating achievable performance (Ederington, 1979).

2.9 Synthesis and Identified Gaps

Three gaps recur across the literature:

1. **Integration of ESG Measurement with Portfolio Rules.** Environmental footprint metrics (and their uncertainty) are rarely mapped to binding portfolio limits, despite the auditability requirements of ESRS/ISSB/GRI (Mulligan et al., 2024; Wendl et al., 2023).

2. **Governance-Aware Execution.** Microstructure frictions documented by Makarov and Schoar (2020) are infrequently embedded in connectedness-driven hedging rules, leading to a gap between ex-ante design and ex-post feasibility.

3. **Regime-Conditioned Diversifier Choice.** The relative efficacy of green bonds versus clean-energy equities as hedges for blockchain-thematic risk is contingent on market states, yet regime-switching selection rules are seldom formalised (Mzoughi et al., 2024; Ul Haq et al., 2023).

These gaps motivate the framework in Chapter 3, which operationalises state-contingent allocation and hedging overlays, integrates verifiable environmental metrics with uncertainty bands, and introduces governance controls aligned with policy-oriented disclosure standards.

Chapter 3

Methods Used in the Literature

This chapter is organized not as a catalogue of techniques but as a bridge from empirical signal extraction to rule–based portfolio decisions under auditability and implementation frictions. The guiding premise is that statistical outputs only become *policy–relevant* when they are mapped to explicit state variables, thresholds, and feasibility constraints that a portfolio can obey in real time. Accordingly, methods are presented in terms of what they measure, why those measurements are economically meaningful in ESG–crypto settings, and how they translate into allocation, hedging, and risk–control actions that remain defensible under reporting and assurance regimes.

The first task is to clarify the inferential status of dependence measures and the conditions under which they should move capital. Time–varying connectedness—estimated from state–space representations with generalized variance decompositions—provides a time–indexed map of shock transmission across ESG and crypto assets (Antonakakis et al., 2020; Chatziantoniou & Gabauer, 2021; Gabauer, 2021; Koop et al., 1996; Pesaran & Shin, 1998). Rather than stopping at description, we convert total and directional connectedness into state variables that (i) enter the objective as penalties on systemic linkages, producing allocations that reduce exposure to dominant transmitters, and (ii) trigger budget and turnover gates when transmitter concentration exceeds calibrated bounds. Time–frequency analysis then localizes co–movement to the relevant investment horizon: short–scale in–phase episodes warn against tactical hedges that will be quickly arbitraged away, whereas persistent low–frequency anti–phase structure motivates strategic diversifiers (Ul Haq et al., 2023). Tail–risk toolkits—copulas and conditional risk measures such as CoVaR—identify regimes in which average spillovers materially understate joint downside, thereby elevating the priority of downside overlays (Mzoughi et al., 2024). Finally, quantile regression binds these regimes to observables (macro factors and market proxies), so that weights and hedge coverage are *conditioned* on lower–tail exposures, not just unconditional variances (Lin et al., 2025).

Feasibility is made explicit. Even if dependence diagnostics favour hedging, execution can fail in segmented crypto markets where cross–venue premia, fiat–rail bottlenecks, and exchange outages are recurrent (Makarov & Schoar, 2020). To avoid pro–forma recommendations, signals are filtered by persistence and half–life; hedge ratios from conditional covariance models are scaled by liquidity tiers and jurisdictional routing; turnover penalties rise mechanically with measured segmentation; and all statistical improvements are traced through to economic value net of transaction costs, slippage, and margin requirements à la Ederington (1979). In parallel, ESG measurement and assurance are treated as investability gates: location–versus market–based carbon reporting, provenance of digitally monitored and verified (dMRV) data, and proof–of–reserves cadence determine whether an exposure is eligible or capped, regardless of what the econometrics alone might suggest (Alzoubi & Mishra, 2023; Mulligan et al., 2024; Wendl et al., 2023). Where assurance falls short, conservative limits or ineligibility are enforced.

Because estimates are sensitive to modelling choices—forecast horizon in variance decompositions, functional form for tails, and sub–sample segmentation—the chapter pairs each method with a cross–validation and uncertainty protocol. Connectedness results are stress–tested against correlation–based multivariate volatility models and Bayesian VARs; tail dependence is contrasted with empirical copulas and evaluated via differences in conditional risk and distributional tests; time–frequency patterns are checked for scale–grid stability and cone–of–influence artefacts; and rolling designs report bootstrap bands to reflect parameter and window–length uncertainty. These diagnostics are not relegated to appendices: they feed directly into the *policy* layer as tolerance bands, persistence filters, and guardrails on when a signal is actionable and when it should be ignored as statistical noise.

The mapping from method to decision is therefore explicit. Time–varying connectedness produces state variables that penalise systemic linkages and trigger hedging overlays only when threshold breaches are persistent. Time–frequency evidence determines whether diversification is tactical or strategic given the rebalancing clock. Tail metrics and conditional quantiles define left–tail states in which risk is reallocated toward assets with low lower–tail dependence (for example, green bonds) and hedge budgets are scaled up as stress intensifies. Conditional covariance models yield implementable hedge ratios whose realised hedging effectiveness is evaluated after costs and segmentation frictions. Every statistical recommendation is thus routed through feasibility (liquidity, turnover, governance) and assurance (measurement provenance and disclosure alignment), ensuring that the allocation framework developed in later chapters is both econometrically grounded and operable within contemporary market and reporting constraints.

3.1 Data Preparation and Pre-tests Used in the Literature

Empirical studies on ESG–cryptocurrency linkages typically begin with careful construction of return series and diagnostic testing (Alharbi et al., 2025; Ali et al., 2024; Bibi et al., 2025). This section lays out the core transformations and pre-tests reported in the literature and clarifies their implications for downstream inference.

3.1.1 Return construction and sampling choices

Let P_t denote the closing price (or index level) at time t. Two transformations dominate:

Logarithmic returns

$$r_t = \ln\left(\frac{P_t}{P_{t-1}}\right),$$

which are time-additive across horizons and more robust to heteroskedastic scaling in trending markets. Log returns are the default input to linear and state–space models because they simplify aggregation and often stabilise the variance over moderate horizons.

Simple percentage returns

$$r_t = \frac{P_t - P_{t-1}}{P_{t-1}},$$

which preserve proportional changes in levels and can be convenient for direct economic interpretation in certain portfolio applications. In high-volatility regimes, however, the log formulation is typically preferred.

Sampling frequency varies by question: daily data remain standard in connectedness studies (Alharbi et al., 2025; Bibi et al., 2025), whereas intraday data can be informative for microstructure questions but introduce market–closure misalignment (crypto trains continuously; equities and ESG indices close by exchange), calendar effects, and stronger microstructure noise. Weekly aggregation can attenuate noise but risks obscuring short-lived contagion.

3.1.2 Currency, calendar, and venue harmonisation

Cross-asset studies must address (i) base-currency consistency (e.g., converting crypto/USD and equity/EUR to a common base); (ii) holiday/calendar mismatches (crypto trades on weekends, most equity/ESG indices do not); and (iii) venue effects for crypto (exchange selection, stablecoin vs. fiat pairs). Misalignment here mechanically induces spurious lead–lag structure and bias in connectedness. Studies therefore synchronise by: removing non-overlapping dates, forward-filling only for index closings (not prices), and restricting crypto quotes to fiat pairs and major venues to limit segmentation artefacts (Makarov & Schoar, 2020).

3.1.3 Asymmetric decomposition of returns

To study sign–dependent spillovers, many papers split returns into positive and negative components (Alharbi et al., 2025; Bibi et al., 2025):

$$r_t^+ = \mathbb{1}_{\{r_t>0\}}\, r_t, \qquad r_t^- = \mathbb{1}_{\{r_t<0\}}\, r_t.$$

Estimating models separately on $\{r_t^+\}$ and $\{r_t^-\}$ reveals whether "bad news" propagates more strongly than "good news"—a recurring finding in ESG–crypto systems.

3.1.4 Stationarity and order of integration

Levels of asset prices are usually $I(1)$. The Elliott–Rothenberg–Stock (ERS) DF–GLS test is frequently applied to levels to confirm unit roots, while returns are treated as $I(0)$ inputs to VAR and related frameworks (Elliott et al., 1996). When long spans include major breaks (e.g., early 2020; early 2022), unit-root tests are interpreted alongside break segmentation (below) to avoid size distortions.

3.1.5 Distributional diagnostics and conditional heteroskedasticity

Return series systematically deviate from normality. Jarque–Bera, D'Agostino (skewness), and Anscombe–Glynn (kurtosis) tests commonly reject Gaussianity (Anscombe & Glynn, 1983; D'Agostino, 1970; Jarque & Bera, 1980). ARCH-type volatility clustering motivates time-varying parameterisations downstream (e.g., TVP–VAR, DCC–GARCH) and informs bootstrap inference rather than closed-form asymptotics.

3.1.6 Structural break segmentation and event windows

Because connectedness is regime-dependent, samples are segmented around known events—commonly the COVID-19 onset in Q1 2020 and the Russia–Ukraine escalation in Q1 2022—where sharp spikes in spillovers are observed (Alharbi et al., 2025; Ali et al., 2024; Bibi et al., 2025). Sub-sample analyses and rolling-window designs are then used to isolate crisis vs. calm behavior.

3.1.7 Asset universe and ESG classification

ESG universes in this literature often comprise broad "leaders" indices as well as E/S/G sub-indices; crypto sets include BTC, ETH, and energy-efficient ("green") Proof-of-Stake assets such as ADA, XLM, MIOTA, and NANO (Alharbi et al., 2025). To mitigate methodology-induced ESG measurement error, some studies cross-validate results with alternative ESG providers; where not feasible, sensitivity tests restrict attention to environmental sub-sleeves to align with green-crypto thematics.

Descriptive statistics of return series.

Statistics	WTI	Brent	Natural Gas	Coal	Cardano	Ripple	IOTA	Nano
N	1358	1358	1358	1358	1358	1358	1358	1358
Mean	0.00100	0.00096	0.00133	0.00121	0.00541	0.00321	0.00207	0.00642
SD	0.03694	0.03653	0.03851	0.03156	0.09552	0.07850	0.07901	0.10495
Median	0.00098	0.00142	0.00000	0.00000	−0.00056	−0.00102	−0.00091	0.00000
Min	−0.51232	−0.53823	−0.20508	−0.41543	−0.39571	−0.42328	−0.41919	−0.45912
Max	0.35017	0.50987	0.39480	0.38571	1.42717	0.87138	0.90345	1.21001
Range	0.86248	1.04810	0.59988	0.80115	1.82288	1.29466	1.32264	1.66913
Skewness	−0.43330	−0.00015	0.81890	0.51608	5.62042	2.84989	1.52420	3.13561
Kurtosis	47.00	73.26	10.76	55.95	72.44	27.45	16.37	25.41
SE	0.00100	0.00099	0.00104	0.00086	0.00259	0.00213	0.00214	0.00285
j-b statistic	125,427	304,668	6729	177,779	305,070	44,620	15,761	38,891
p-value	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
Arch(20)	165.11	113.25	178.39	211.82	438.55	719.32	659.01	370.77
p-value	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
ADF	−33.85	−28.43	−22.37	−25.64	−27.35	−29.68	−17.25	−25.96
p-value	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01

Figure 3.1: Descriptive statistics of energy commodities and selected cryptocurrencies.

Note. Adapted from Ali et al. (2024). The table reports sample size (N), mean, standard deviation (SD), median, minimum, maximum, range, skewness, kurtosis, standard error (SE), Jarque–Bera statistic with p-values, ARCH(20) statistic with p-values, and Augmented Dickey–Fuller (ADF) unit root test results. All p-values below 0.01 indicate strong rejection of the null hypotheses for normality, homoskedasticity, and unit roots.

3.2 From TVP–VAR to Connectedness: The Standard Pipeline

Time–varying parameter VAR (TVP–VAR) models have become the workhorse for measuring evolving spillovers in ESG–crypto systems (Alharbi et al., 2025; Ali et al., 2024; Bibi et al., 2025). They allow coefficients and shock covariances to evolve, capturing slow-moving structure and crisis-time surges that static models would average away. The average connectedness matrix in Table 3.1 illustrates how variance shares are allocated across ESG and crypto assets, while Figure 3.2 below shows the time profile of net spillovers.

3.2.1 State–space specification and estimation

Let $z_t \in \mathbb{R}^N$ denote returns of N assets. A TVP–VAR(p) is

$$z_t = \sum_{i=1}^{p} B_{i,t} z_{t-i} + u_t, \quad u_t \sim \mathcal{N}(0, \Sigma_t), \tag{3.1}$$

with stacked coefficients $\beta_t = \mathrm{vec}(B_{1,t}, \ldots, B_{p,t})$ following a random walk (or near-random walk) and, optionally, time-varying Σ_t. The observation and state equations are estimated via Kalman filtering and smoothing with diffuse priors or discount factors. Choice of discount factor governs responsiveness: lower values track faster regime changes but increase variance.

3.2.2 Lag order and forecast horizon

Daily returns usually favour parsimonious lag orders by BIC or equivalent criteria; $p = 1$ is frequently selected (Bibi et al., 2025). The connectedness horizon H (e.g., 10–20 days ahead) encodes the investor's risk horizon: short H captures immediate volatility transmission; longer H aligns with institutional rebalancing cycles.

3.2.3 VMA representation and generalised FEVD

After estimation, the model is recast in VMA form

$$z_t = \sum_{h=0}^{\infty} A_{h,t}\, u_{t-h},$$

and the H–step generalised forecast error variance decomposition (GFEVD) is computed (Koop et al., 1996; Pesaran & Shin, 1998):

$$\Psi^g_{ij,t}(H) = \frac{\sigma^{-1}_{jj,t} \sum_{h=0}^{H-1} (e_i' A_{h,t} \Sigma_t e_j)^2}{\sum_{h=0}^{H-1} \left(e_i' A_{h,t} \Sigma_t A_{h,t}' e_i\right)}. \tag{3.2}$$

Row-normalisation yields $\tilde{\Psi}^g_{ij,t}(H)$ with $\sum_j \tilde{\Psi}^g_{ij,t}(H) = 1$.

3.2.4 Connectedness indices

Define the Total Connectedness Index,

$$TCI_t = \frac{1}{N}\sum_{i=1}^{N}\sum_{j\neq i}\tilde{\Psi}^g_{ij,t}(H)\times 100,$$

and directional measures:

$$TO_i(t) = \sum_{j\neq i}\tilde{\Psi}^g_{ij,t}(H), \quad FROM_i(t) = \sum_{j\neq i}\tilde{\Psi}^g_{ji,t}(H), \quad NET_i(t) = TO_i(t) - FROM_i(t).$$

Pairwise connectedness between i and j is

$$PCI_{ij}(t) = \frac{\tilde{\Psi}^g_{ij,t}(H) + \tilde{\Psi}^g_{ji,t}(H)}{2}.$$

Finite-sample properties and scaling of TCI_t have been studied and refined; proper normalisation avoids misinterpretation of the index bounds (Chatziantoniou & Gabauer, 2021).

3.2.5 Shock-sign asymmetry

Estimating connectedness separately for $\{z_t^+\}$ and $\{z_t^-\}$ typically yields $TCI^- > TCI^+$, consistent with stronger propagation of negative shocks across ESG and crypto (Alharbi et al., 2025; Bibi et al., 2025). In portfolio terms, diversification benefits weaken precisely when protection is most valuable.

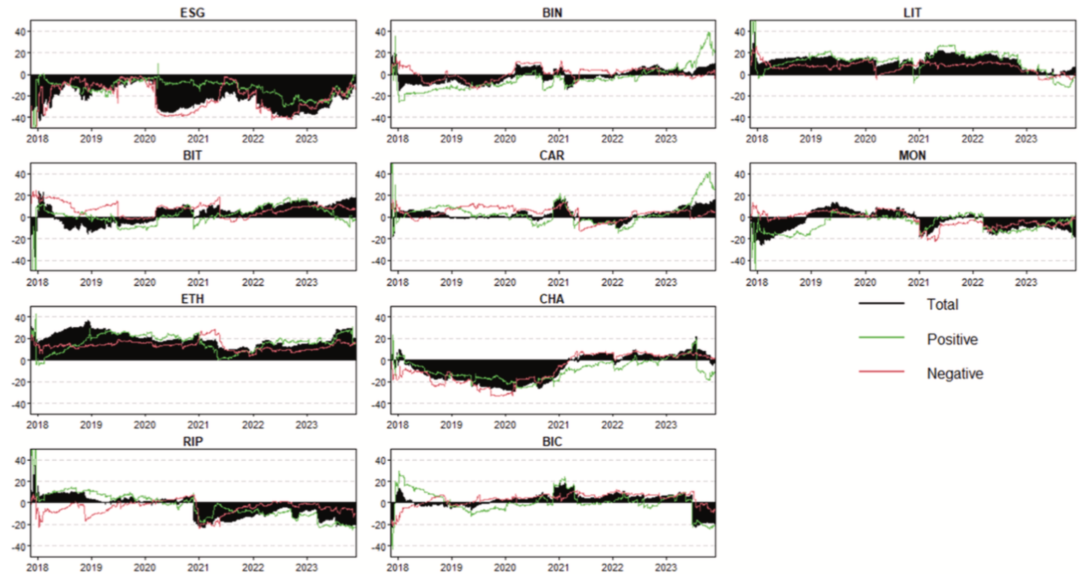

Figure 3.2: Dynamic net total directional connectedness across assets.

Note. Adapted from Bibi et al. (2025). The analysis covers ESG indices (ESG), cryptocurrency exchanges (BIN), lithium (LIT), Bitcoin (BIT), Cardano (CAR), Monero (MON), Ethereum (ETH), Chainlink (CHA), Ripple (RIP), and Bitcoin Cash (BIC). Estimates use a TVP–VAR model with one lag (BIC selection) and a 20-step-ahead generalized forecast error variance decomposition. The black shaded area depicts symmetric net total connectedness, while green and red lines indicate positive and negative net total directional connectedness, respectively. Notable shifts align with systemic stress episodes, such as the COVID-19 onset in early 2020 and the Russia–Ukraine escalation in 2022.

Table 3.1: Average Connectedness Table.

	ESG	BIT	ETH	RIP	BIN	CAR	CHA	BIC	LIT	MON	FROM
ESG	67.03	4.32	4.93	2.68	4.01	3.84	3.86	2.7	3.68	2.96	32.97
BIT	1.25	24.26	13.2	6.47	8.91	7.93	6.9	9.24	11.72	10.12	75.74
ETH	1.4	11.76	20.88	7.93	8.85	8.9	8.55	9.29	12.54	9.3	79.12
RIP	0.98	7.75	10.83	30.1	7.34	10.9	6.62	8.1	9.56	7.82	69.9
BIN	1.48	9.78	11.24	6.89	27.69	8.08	7.83	7.73	9.7	9.57	72.31
CAR	1.21	8.8	10.77	9.51	7.85	26.23	8.14	8.34	10.23	8.9	73.77
CHA	1.58	8.02	11.3	6.61	8.41	8.74	30.28	8.08	8.61	8.36	69.72
BIC	0.95	9.81	10.72	8.07	7.27	7.26	8.08	26.08	11.64	10.61	73.92
LIT	1.12	11.2	11.2	13.51	7.34	8.14	8.95	7.07	10.26	23.37	76.63
MON	1.03	10.71	11.05	6.97	8.95	8.48	7.52	10.42	9.78	25.09	74.91
TO	10.1	82.15	98.32	61.29	69.67	73.87	63.77	75.09	87.47	76.35	**69.90**
Inc.Own	78.03	106.41	119.19	91.39	97.35	100.1	94.05	101.18	110.85	101.44	**TCI**
NET	-21.97	6.41	19.19	-8.61	-2.65	0.1	-5.95	1.18	10.85	1.44	**69.90**

Note: A TVP–VAR model using a 20-step-ahead generalized forecast error variance decomposition and a lag length of order one (BIC) formed the basis of the results.
Source: Bibi et al. (2025).

3.3 Wavelet-Based Time–Frequency Analysis

Time-domain connectedness can conceal horizon-specific structure. Wavelet coherence uncovers scale-dependent co-movements, revealing, for instance, short-run contagion during crises and different long-run alignment in calmer regimes. ESG–crypto applications document that green bonds co-move with ESG indices at low frequencies while crypto linkages are more state- and horizon-dependent (Ul Haq et al., 2023).

3.3.1 Continuous wavelet transform and coherence

For series $x(t)$ and $y(t)$, compute the (complex) continuous wavelet transform $W_x(s,\tau)$ and $W_y(s,\tau)$. The squared wavelet coherence,

$$R^2_{xy}(s,\tau) = \frac{\left|S(s^{-1}W_{xy}(s,\tau))\right|^2}{S(s^{-1}|W_x(s,\tau)|^2)\ S(s^{-1}|W_y(s,\tau)|^2)},$$

with $W_{xy} = W_x W_y^*$ and smoothing operator $S(\cdot)$, lies in $[0,1]$ and localises correlation by scale s (investment horizon) and time τ. Phase arrows indicate sign (in-/out-of-phase) and lead–lag direction.

3.3.2 Portfolio interpretation

Short-horizon in-phase episodes signal transient contagion (limited diversification), whereas persistent out-of-phase structure at long horizons suggests strategic hedging potential. Green cryptocurrencies (e.g., ADA, MIOTA) often display stronger integration with environmental sub-indices than BTC/ETH in low-frequency bands, consistent with thematic alignment (Ul Haq et al., 2023).

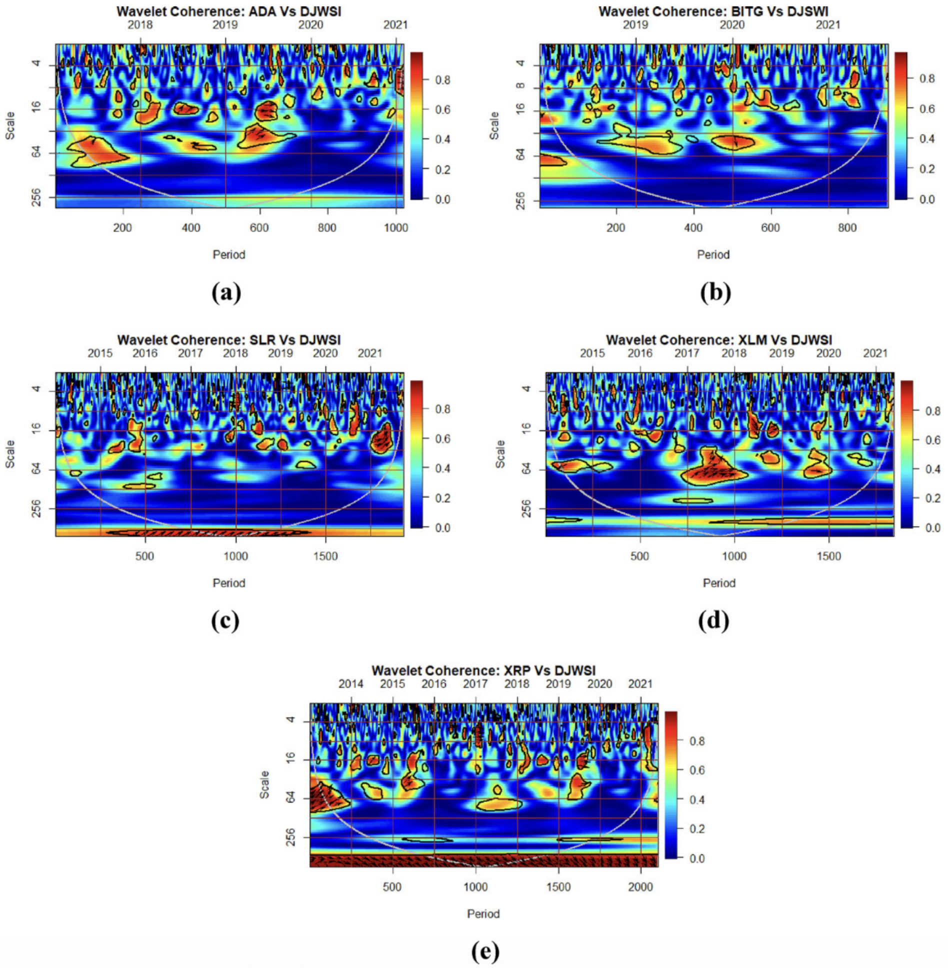

Figure 3.3: Wavelet coherence between sustainable cryptocurrencies and a global sustainability index.

Note. Adapted from Ul Haq et al. (2023). The analysis compares ADA, BTG, SLR, XLM, and XRP with the Dow Jones World Sustainability Index (DJWSI). Panels (a)–(e) show time–frequency localized correlations, where color intensity (blue = low, red = high) reflects coherence strength. The horizontal axis represents time (years) and the vertical axis represents scale (days), inversely related to frequency. Warm-colored regions with thick contours denote statistically significant coherence at the 5% level (Monte Carlo simulations). Arrows indicate phase differences: right = in-phase, left = anti-phase, up/down = lead–lag relationships between the cryptocurrency and the index.

Time–frequency co-movement patterns are visualised in Figure 3.3.

3.4 Copula and CoVaR Frameworks for Tail Risk and Systemic Spillovers

Average spillover metrics understate co-movement during extremes. Copulas model nonlinear, asymmetrical dependence in the tails; CoVaR quantifies the conditional downside of one asset given stress in another. ESG–crypto studies employing these tools find that green bonds typically contribute less to systemic downside than clean-energy equities when crypto sells off (Mzoughi et al., 2024; Ul Haq et al., 2023).

3.4.1 Copulas and tail coefficients

After filtering marginals (e.g., via GARCH), dependence is estimated with copulas that admit distinct tail features (e.g., Student-t symmetric tails; Clayton lower-tail; Gumbel upper-tail). Lower-tail dependence,

$$\lambda_L = \lim_{u \to 0^+} \Pr\Big(Y < F_Y^{-1}(u) \mid X < F_X^{-1}(u)\Big),$$

measures the probability of joint crashes. High λ_L for (BTC, clean-energy equities) in crises implies limited hedge value exactly when needed.

3.4.2 CoVaR and marginal systemic risk

Let VaR_j^{α} denote the α–quantile of asset j's return. The α–CoVaR of asset i conditional on j being at VaR_j^{α} solves

$$\Pr\Big(r_i \le CoVaR_{i|j}^{\alpha} \mid r_j = VaR_j^{\alpha}\Big) = \alpha,$$

with $\Delta CoVaR$ the deviation from median conditions. In ESG–crypto systems, $\Delta CoVaR$ spikes around systemic events and is materially smaller for green bonds than for clean-energy equities when conditioning on crypto stress (Mzoughi et al., 2024).

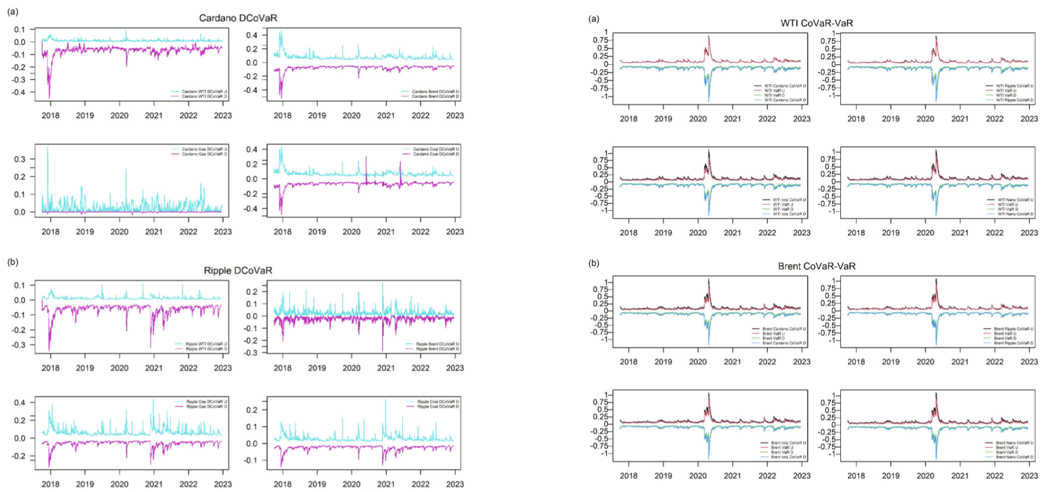

Figure 3.4: Dynamic Conditional Value-at-Risk (ΔCoVaR) for cryptocurrency–energy commodity pairs.

Note. Adapted from Mzoughi et al. (2024). The left panel shows time-varying upper-tail (U) and lower-tail (D) conditional risk spillovers from Cardano and Ripple to four major energy markets (WTI crude oil, Brent crude oil, natural gas, coal) from 2018 to 2023. The right panel depicts reciprocal effects, with CoVaR from energy commodities (WTI, Brent) to Cardano, Ripple, and Nano plotted alongside their own Value-at-Risk (VaR) benchmarks. Cyan lines = upper-tail CoVaR; magenta lines = lower-tail CoVaR; black/red lines = VaR levels. Spikes, especially in early 2020, coincide with global shocks such as the onset of COVID-19, reflecting heightened systemic risk and amplified downside spillovers.

Table 3.2: KS-test for equality of absolute DCoVaRs of cryptocurrencies.

	Cardano	**Ripple**	**Iota**	**Nano**
	DCoVaR D ≠ DCoVaR U	DCoVaR D ≠ DCoVaR U	DCoVaR D ≠ DCoVaR U	DCoVaR D ≠ DCoVaR U
WTI	0.95876 [0.00000]	0.95803 [0.00000]	1.00000 [0.00000]	0.08984 [0.00003]
Brent	0.11782 [0.00000]	0.02135 [0.91618]	0.03314 [0.44509]	0.09499 [0.00001]
Natural Gas	0.70471 [0.00000]	0.04124 [0.19847]	0.32622 [0.00000]	0.67526 [0.00000]
Coal	0.09794 [0.00000]	0.03976 [0.23323]	0.04050 [0.21532]	0.19882 [0.00000]

Note: This table presents the hypothesis tests for equality between downside and upside absolute DCoVaRs of clean cryptocurrencies conditional on dirty fuels. Columns 2–5 display the test statistics, with p-values in brackets. The K–S test is two-sided with bootstrapping.

Tail spillovers are documented in Figure 3.4; distributional asymmetries are tested in Table 3.2.

3.5 Quantile Regression and Regime-Specific Analysis

Quantile regression models heterogeneous sensitivities across the conditional return distribution, revealing that spillovers intensify in the left tail and amplify under crisis regimes (Lin et al., 2025). This directly informs state-contingent allocation and hedging.

3.5.1 Specification and interpretation

For ESG return $r_{i,t}$ and crypto return $r_{j,t}$,

$$Q_{r_{i,t}}(\tau \mid r_{j,t}, X_t) = \alpha(\tau) + \beta(\tau) r_{j,t} + X_t' \gamma(\tau),$$

with covariates X_t (e.g., macro controls or regime dummies). Large $|\beta(0.05)|$ relative to $|\beta(0.50)|$ signals heightened vulnerability under stress; symmetry of $\beta(\tau)$ across tails indicates co-movement in both sell-offs and rallies.

3.5.2 Findings and portfolio use

Lin et al. (2025) document stronger lower-tail sensitivities and regime amplification. Portfolio rules therefore scale down crypto exposure when tail risk indicators are elevated, or pair crypto with low-λ_L assets (e.g., green bonds) when $\beta(0.05)$ becomes large in magnitude.

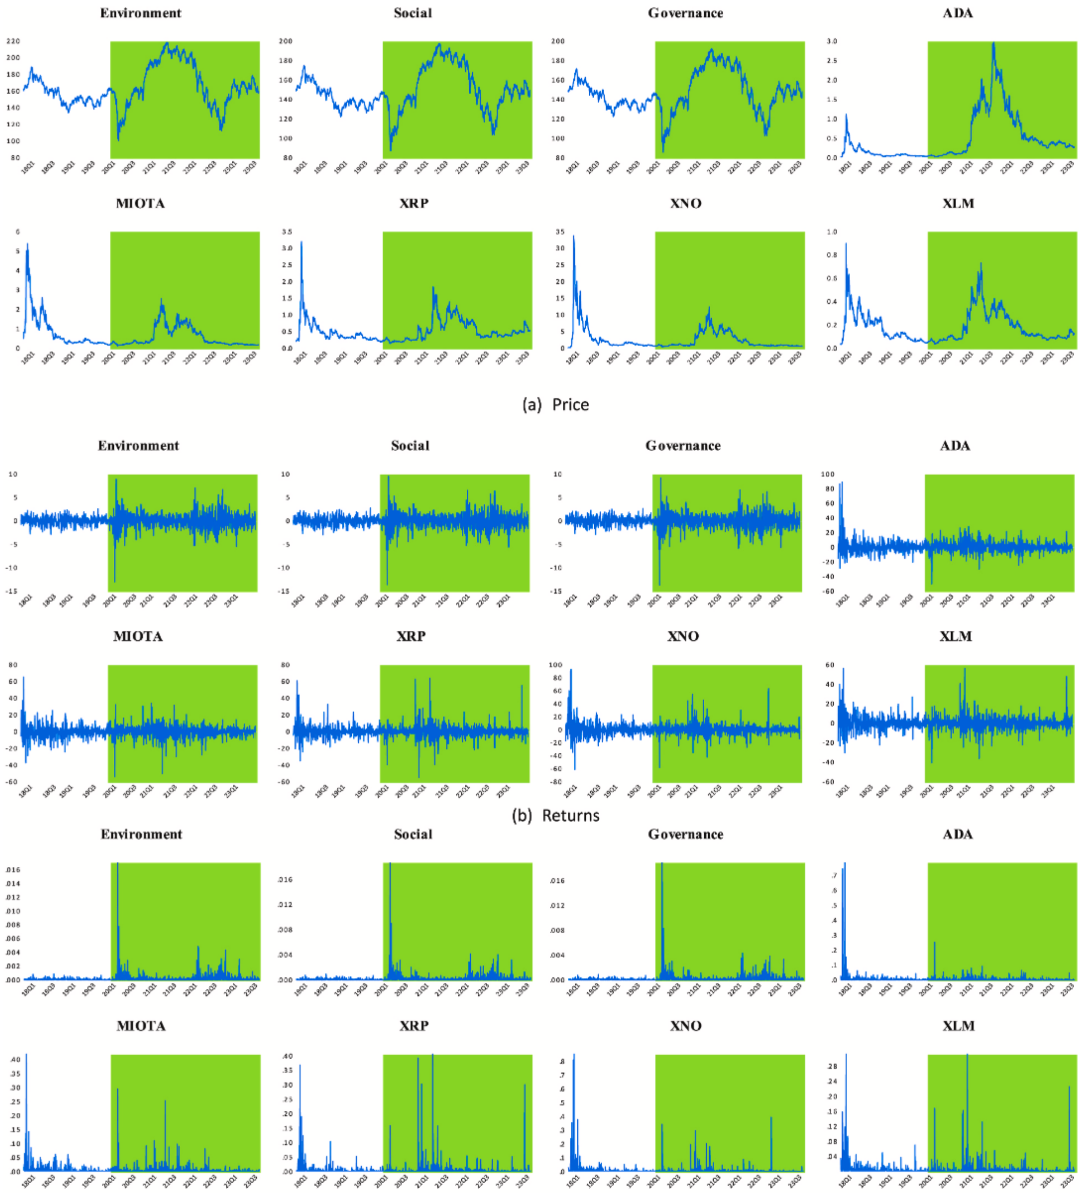

Figure 3.5: ESG pillar indices and sustainable cryptocurrency time series.

Note. Adapted from Alharbi et al. (2025). Panel (a) plots prices for ESG sub-indices (Environment, Social, Governance) alongside ADA, MIOTA, XRP, XNO, and XLM. Panel (b) shows the corresponding log-returns, with volatility spikes most evident between 2020 and 2021. Panel (c) displays volatility proxies (absolute and squared returns). The green shaded region denotes the subsample used in the empirical analysis.

3.6 Shock Classification and Bubble Diagnostics

Beyond dependence metrics, some studies classify the nature of shocks—fundamental, speculative (positive bubbles), or panic-driven (negative bubbles)—and link these to spillover regimes. The log-periodic power-law singularity (LPPLS) form is used to diagnose accelerating dynamics toward a critical time in crypto and, at times, ESG thematics (Fry & Cheah, 2016):

$$p(t) = A + B(t_c - t)^m + C(t_c - t)^m \cos\Big(\omega \ln(t_c - t) + \phi\Big) + \varepsilon_t,$$

with $0 < m < 1$ and sign of B indicating positive vs. negative bubbles. Identifying such phases informs pre-emptive de-risking when cross-market transmission historically intensifies near t_c.

3.7 Hybrid and Comparative Approaches

Because no single framework spans time variation, horizons, tails, and asymmetry, hybrid pipelines are common. A practical sequence is: TVP–VAR to flag high TCI_t; wavelet coherence to localise horizon-specific co-movement; copulas/CoVaR to quantify tail co-dependence; and quantile regression to map conditional sensitivities. Comparative work evaluates predictive accuracy, economic value (e.g., Sharpe improvement), and robustness across regimes (Mzoughi et al., 2024; Ul Haq et al., 2023).

Table 3.3: Optimal hedge ratios and hedging effectiveness of dirty fuels.

	Optimal Hedge Ratios				Hedging Effectiveness			
	WTI	Brent	Gas	Coal	WTI	Brent	Gas	Coal
Cardano	0.02941	0.03693	0.01876	0.01386	0.02274	0.02824	-0.00017	-0.00049
	(0.02400)	(0.03160)	(0.01030)	(0.01630)				
Ripple	0.02881	0.03871	0.01774	0.01791	0.01076	0.01823	-0.00065	-0.00096
	(0.03060)	(0.03550)	(0.01600)	(0.02020)				
IOTA	0.02345	0.02223	0.00858	0.00592	0.02637	0.01509	-0.00014	-0.00157
	(0.04880)	(0.02450)	(0.00510)	(0.00690)				
Nano	0.02177	0.02300	0.01364	-0.00110	0.01791	0.01921	0.00085	0.00064
	(0.03100)	(0.02560)	(0.00990)	(0.00140)				

Note: This table presents optimal hedge ratios when dirty fuels are hedged using clean cryptocurrencies. Optimal hedge ratios are obtained using an A-DCC-GJR-GARCH (1,1) model. Mean and standard deviation are in parentheses. Columns 6–9 present the hedging effectiveness index.

Table 3.4: Optimal hedge ratios and hedging effectiveness of clean cryptocurrencies.

	Optimal Hedge Ratios				Hedging Effectiveness			
	Cardano	Ripple	IOTA	Nano	Cardano	Ripple	IOTA	Nano
WTI	0.30202	0.23738	0.14474	0.23703	0.00576	0.00922	0.00653	0.00124
	(0.23928)	(0.20322)	(0.13917)	(0.20399)				
Brent	0.39229	0.33747	0.20799	0.28616	0.00375	0.00099	0.00149	-0.01362
	(0.32261)	(0.25789)	(0.13917)	(0.20399)				
Gas	0.10468	0.08203	0.04474	0.09013	0.00070	-0.00032	-0.00086	0.00017
	(0.07037)	(0.05718)	(0.02762)	(0.01670)				
Coal	0.23520	0.25593	0.09258	-0.02330	0.00163	0.00053	-0.00233	-0.00039
	(0.20747)	(0.23731)	(0.06362)	(0.01880)				

Note: This table presents optimal hedge ratios when clean cryptocurrencies are hedged against one another. Estimates are from an A-DCC-GJR-GARCH (1,1) model. Standard deviations are in parentheses.

3.8 Hedge Ratio Calibration and Hedging Effectiveness

Translating dependence into action requires hedge ratios and performance evaluation.

3.8.1 Minimum-variance hedge ratio

For ESG $r_{E,t}$ and crypto $r_{C,t}$,

$$h_t = \frac{\mathrm{Cov}(r_{E,t}, r_{C,t})}{\mathrm{Var}(r_{C,t})},$$

with dynamic versions obtained from rolling windows or conditional volatility/correlation (e.g., DCC). Directional connectedness can also inform sign and magnitude when transmissions are highly asymmetric.

3.8.2 Hedging effectiveness

Hedging effectiveness (HE) (Ederington, 1979) measures variance reduction:

$$HE = 1 - \frac{\mathrm{Var}(r_{P,h})}{\mathrm{Var}(r_{P,0})}.$$

Crisis-time HE often rises because correlations spike, but this may come with opportunity costs if reversals are swift. Strategy sets include Minimum-Variance (MVP), Minimum-Connectedness (MCP), and Minimum-CoVaR (MCoP); tail-aware portfolios can outperform MVP on risk-adjusted metrics in stress windows (Bibi et al., 2025).

3.8.3 Execution frictions and governance

Turnover constraints, liquidity screens, and venue-selection policies are required to avoid over-trading and to respect capital-market governance, especially in segmented crypto markets (Makarov & Schoar, 2020). Transaction costs and margin requirements materially alter net HE and should be incorporated into backtests.

3.9 Robustness Checks and Cross-Validation

Given model complexity and regime sensitivity, robust inference requires multiple diagnostic layers.

3.9.1 Parameter sensitivity

Vary lag order p, horizon H, and discount factors to test stability of TCI_t, directional measures, and hedge ratios. Adjust wavelet scale grids to ensure key patterns persist across frequency bands.

3.9.2 Alternative estimators

Benchmark TVP–VAR against DCC–GARCH or Bayesian VARs for spillovers; use empirical copulas for tail dependence to relax functional-form restrictions; compare shock-sign decomposition with quantile-based asymmetry. Convergent conclusions across estimators materially strengthen interpretation (Antonakakis et al., 2020; Chatziantoniou & Gabauer, 2021; Gabauer, 2021).

3.9.3 Sub-samples and event windows

Estimate models in pre-crisis, crisis, and post-crisis windows (e.g., 2019, 2020–2021, 2022–2023) to confirm regime dependence:

$$TCI_{\text{crisis}} > TCI_{\text{calm}}.$$

Event studies around policy shocks (e.g., PoW restrictions) identify local surges in spillovers with meaningful ESG implications.

3.9.4 Bootstrap inference

Residual or block bootstraps generate confidence intervals for TCI_t, $PCI_{ij}(t)$, and h_t. Uncertainty bands are essential when rolling windows reduce effective sample size.

3.9.5 Economic value tests

Complement statistical gains with portfolio metrics (Sharpe, Sortino, expected shortfall) and risk controls (VaR breach frequency). Gains that disappear after costs or leverage constraints should be treated as economically insignificant even if statistically detectable.

Table 3.5: Common robustness checks and their diagnostic purpose in ESG–crypto studies.

Check	Diagnostic purpose
Vary p and H	Memory vs. horizon dependence of spillovers
Alt. estimators (DCC, BVAR)	Model misspecification and linearity risk
Empirical copulas	Functional-form robustness for tail dependence
Sub-sample windows	Regime stability and event sensitivity
Bootstrap CIs	Parameter uncertainty in rolling designs
Econ. value tests	Practical relevance net of costs

Note: This table lists the most common robustness checks employed in ESG–crypto connectedness studies. Each check serves a distinct diagnostic purpose, ranging from model specification and parameter uncertainty to regime stability and the economic value of hedging results. *Source: The author.*

3.10 Limitations of Methods

Each framework has identifiable limits that inform cautious interpretation.

3.10.1 Specification risk and linearity

Linear TVP–VARs approximate complex dynamics; threshold effects, nonlinearities, and feedback from policy to markets can bias spillover measurement if unmodelled.

While more flexible models exist, they increase estimation variance and require longer spans.

3.10.2 Small-sample uncertainty

Rolling estimation and high dimensionality inflate sampling error; variance shares in GFEVD become noisy when N is large relative to window length. Scaling and normalisation choices affect the interpretation of TCI_t (Chatziantoniou & Gabauer, 2021; Pesaran & Shin, 1998).

3.10.3 Horizon and frequency sensitivity

Connectedness depends on the forecast horizon H; wavelet coherence depends on chosen scales. Using only one setting risks horizon-specific conclusions.

3.10.4 Tail estimation challenges

Parametric copulas can misrepresent empirical tails; nonparametric alternatives require large samples. CoVaR inference is sensitive to quantile estimation in sparse crisis data.

3.10.5 Asymmetry measurement

Binary sign-splits treat all negative returns as "bad news"; quantile methods provide richer diagnostics but are data-intensive and sensitive to outliers.

3.10.6 Execution frictions

Dynamic hedges derived from rapidly moving signals may over-trade; turnover, liquidity, and venue segmentation reduce realised HE, particularly in crypto (Makarov & Schoar, 2020). Governance constraints (e.g., venue whitelists, capital controls) further limit implementability.

3.10.7 Economic vs. statistical significance

Variance reductions that fail to improve Sharpe or expected shortfall net of costs have limited portfolio value. Report both statistical and economic metrics, and prefer designs whose gains survive frictions.

Chapter 4

Market Interconnectedness: ESG Equities, Green Assets, and Crypto

4.1 Purpose, Scope, and Contribution

This chapter develops a mechanism-first account of market interdependence between ESG equities, green assets (green bonds and clean-energy equities), and cryptocurrencies, and shows how the empirical signals documented in the literature become *operational* inputs to allocation and policy. The central claim is that the statistics produced by time-varying multivariate systems (connectedness matrices, directional roles, and their sign/horizon decompositions), tail-risk diagnostics (copulas and conditional systemic risk), and time–frequency tools (coherence) are not merely descriptive. When read through their structural determinants—innovation covariances and impulse-response propagation—they admit a direct mapping to portfolio states, trigger thresholds, and implementable constraints that remain valid under audit and trading frictions.

The organising perspective is that spikes in total and directional connectedness during crises are the *mechanical* consequence of two forces that tighten simultaneously: (i) correlations of structural innovations rise as funding conditions, volatility, and risk aversion become factor-dominated; and (ii) the impulse-response operator enlarges as intermediation capacity thins, raising pass-through. Under these conditions the connectedness mass shifts off-diagonal, transmitter/receiver roles polarise, and negative-shock decompositions dominate—regularities reported across independent studies of ESG–crypto systems (Alharbi et al., 2025; Ali et al., 2024; Bibi et al., 2025). Interpreting the statistics through this lens avoids over-attributing the results to model features: they reflect stress-regime economics—funding constraints, discount-rate channels, and benchmark/flow amplification—rather than artefacts of estimation.

A second strand concerns extremes and horizons. Average spillover metrics understate joint downside during turbulence; tail-dependence coefficients and conditional Value-at-Risk reveal that clean-energy and blockchain equities co-crash via shared growth/technology and discount-rate exposure, whereas green bonds retain low lower-tail linkage even when connectedness is elevated (Mzoughi et al., 2024). Time–frequency evidence complements this picture: in calm regimes, long-horizon anti-phase structure between crypto and broad ESG supports strategic diversification, while short-horizon in-phase bands emerge transiently in crises as a common liquidity factor dominates (Ul Haq et al., 2023). These findings justify, respectively, (i) tail-aware portfolio penalties and hedge selection rules; and (ii) horizon-matched rebalancing clocks that privilege signals persistent at investment, not trading, scales.

Third, implementability conditions matter. Cross-venue segmentation, capital controls, and fiat-rail bottlenecks create persistent basis premia that degrade hedge tracking even when return correlations look favourable (Makarov & Schoar, 2020). Any rule that keys off connectedness or tail risk must therefore be filtered by persistence and half-life to avoid reacting to statistical noise, scaled by liquidity tiers and jurisdictional routing, and evaluated on economic value net of costs and margin rather than on variance reduction alone. In parallel, environmental and governance disclosures (energy and emissions by location- vs. market-based accounting, e-waste stock–flow, validator concentration, proof-of-reserves cadence and scope) determine investability: where assurance is insufficient, exposure is capped or ineligible even if econometrics would otherwise support it.

With these principles, the chapter contributes in three ways. First, it formalises the network objects used in practice and derives their crisis behaviour from the underlying state-space representation (Section 4.2), then documents the empirical features that recur across stress episodes—spikes in total connectedness, negative-shock dominance, and role switching (Section 4.3). Second, it translates directionality, tail-risk, and horizon results into portfolio quantities—hedge ratios, hedge-coverage minima, and penalties on systemic linkages—that improve crisis performance relative to variance-only designs (Sections 4.4–4.7). Third, it embeds feasibility and disclosure into the same decision map: segmentation and execution constraints bound turnover and leverage (Section 4.10); quantile-dependent macro sensitivities condition exposures on observable states (Section 4.8); and reporting aligns with prevailing sustainability standards so that triggers, limits, and realised outcomes are auditable (Section 4.13).

Read pragmatically, the chapter provides a playbook. When sign-decomposed connectedness crosses persistent thresholds, portfolios rotate from minimum-variance toward minimum-connectedness designs and raise hedge budgets; when tail metrics widen, hedge instruments shift weight toward low lower-tail-linkage sleeves (e.g., green bonds) and away from thematically correlated clean-energy equities; when time–frequency structure concentrates at short scales, the framework resists strategic reallocation and opts for temporary overlays; and when segmentation intensifies, jurisdiction-matched execution and basis-aware penalties curb over-trading. The sections that follow develop the formal network statistics, document the crisis mechanics and directionality patterns, integrate tail and horizon diagnostics, and conclude with allocation, hedging, and disclosure implications calibrated to the empirical evidence and to real-market frictions.

4.2 A Network View of ESG–Crypto Interdependence

Let $\mathbf{y}_t \in \mathbb{R}^N$ collect returns on ESG leaders and sub-indices, green bonds, clean-energy equities, and selected cryptocurrencies. The time-varying parameter VAR(p) in (**??**) (defined in Chapter 3) induces a time-varying VMA,

$$\mathbf{y}_t \;=\; \sum_{h=0}^{\infty} A_{h,t}\,\boldsymbol{\varepsilon}_{t-h}, \qquad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \Sigma_t), \tag{4.1}$$

where $A_{0,t} = I$ and $A_{h,t}$ evolves with t. Generalised FEVD shares $\phi_{ij,t}^{(H)}$ (Koop et al., 1996; Pesaran & Shin, 1998) allocate the H-step forecast error variance of i to innovations in j and, after row normalisation, yield the *connectedness matrix* $\tilde{\Phi}_t^{(H)} = \{\tilde{\phi}_{ij,t}^{(H)}\}_{i,j}$ with row sums equal to one.

The Total Connectedness Index (TCI) in (**??**) is

$$\mathrm{TCI}_t^{(H)} \;=\; 100 \times \frac{\sum_{i\neq j} \tilde{\phi}_{ij,t}^{(H)}}{N} \;=\; 100 \times \left(\frac{1}{N} \sum_{i=1}^{N} \underbrace{\sum_{j\neq i} \tilde{\phi}_{ij,t}^{(H)}}_{\text{spillovers to } i} \right), \tag{4.2}$$

i.e., the average off-diagonal share of system-wide forecast uncertainty. Directional measures at t are

$$D^{(H)}_{\bullet\to i,t} = \sum_{j\neq i} \tilde{\phi}^{(H)}_{ij,t} \quad \text{(received by } i\text{)}, \tag{4.3}$$

$$D^{(H)}_{i\to\bullet,t} = \sum_{j\neq i} \tilde{\phi}^{(H)}_{ji,t} \quad \text{(transmitted by } i\text{)}, \tag{4.4}$$

$$\mathcal{N}^{(H)}_{i,t} = D^{(H)}_{i\to\bullet,t} - D^{(H)}_{\bullet\to i,t} \quad \text{(net role)}. \tag{4.5}$$

Mechanistic reading. $\mathrm{TCI}^{(H)}_t$ co-moves positively with (i) the off-diagonal elements of Σ_t (innovation correlations) and (ii) the operator norm of $A_{h,t}$ (impulse propagation). Stress regimes that elevate funding constraints, risk premia, and volatility therefore raise $\mathrm{TCI}^{(H)}_t$ *mechanically*, not just statistically.

4.3 Crisis Spikes and Sign Asymmetry: Evidence and Mechanisms

4.3.1 Empirical regularities

Bibi et al. (2025) report TCI in the ESG–crypto system oscillating between $\sim 55\%$ and $\sim 85\%$ over 2017–2023, with peaks at COVID-19 onset ($\approx 85\%$) and the Russia–Ukraine escalation ($\approx 82\%$). Decomposition by shock sign shows persistent *negative* dominance, $\mathrm{TCI}^{(-)}_t > \mathrm{TCI}^{(+)}_t$. Ali et al. (2024) document crisis-state increases in hedge ratios and intensified spillovers from G7 equities into green cryptocurrencies; Alharbi et al. (2025) find ESG sub-components frequently act as net transmitters to green cryptos. In parallel, tail dependence and ΔCoVaR widen between clean-energy and blockchain equities, while green bonds remain diversifying (Mzoughi et al., 2024).

4.3.2 Why crises elevate TCI

From (4.1)–(4.2), a shock in asset j affects i via $A_{h,t}\Sigma_t$; during stress, (i) *innovation correlations* rise as common liquidity and risk-aversion factors dominate (Σ_t tightens off-diagonals); (ii) *propagation* strengthens as dealers reduce intermediation capacity, increasing pass-through (larger $A_{h,t}$). Both raise the FEVD mass assigned to "others," lifting TCI.

4.3.3 Why negative shocks dominate

Downside asymmetry arises from three interacting channels:

1. **Deleveraging and fire-sales.** Losses trigger margin calls; correlated selling increases cross-series innovation covariance and amplifies $A_{h,t}\Sigma_t$.
2. **Convexity in risk limits.** VaR/ES constraints bind more tightly in drawdowns, forcing pro-cyclical deleveraging.
3. **Volatility feedback.** Higher volatility increases forecast error variance; the share attributable to cross shocks grows with volatility of the common factor.

Table 4.1: Crisis Timeline, Metrics, and Mechanistic Interpretation

Event	**Study**	**Metric movement**	**Mechanism**
COVID-19 (Q1 2020)	Bibi et al. (2025)	TCI $\approx$ 85%, $\mathrm{TCI}^{(-)} > \mathrm{TCI}^{(+)}$	Deleveraging; volatility factor; dealer capacity.
	Ali et al. (2024)	Hedge ratios $\uparrow$; G7 $\rightarrow$ green crypto	Funding/liquidity shock transmission.
	Mzoughi et al. (2024)	$\lambda_L \uparrow$, ΔCoVaR $\uparrow$ (CE$\leftrightarrow$BC Eq.)	Discount-rate/growth factor coupling.
Russia–Ukraine (Q1 2022)	Bibi et al. (2025)	TCI $\sim$ 82%; negative dominance	Energy/discount-rate channel.
FTX (Nov 2022)	Ali et al. (2024)	Spillovers $\uparrow$, hedge ratios $\uparrow$	Exchange/liquidity shock; trust contagion.

Note:

This table collates major stress episodes and associated movements in connectedness and tail-risk metrics reported in the literature. TCI = Total Connectedness Index; $\mathrm{TCI}^{(-)}/\mathrm{TCI}^{(+)}$ = sign-decomposed connectedness; λ_L = lower-tail dependence; ΔCoVaR = change in conditional Value-at-Risk; CE = clean-energy equities; BC Eq. = blockchain equities. Mechanisms summarise the dominant transmission channels during each episode. *Source: The author.*

4.4 Directional Transmission: Transmitters, Receivers, and Role Switching

4.4.1 Facts

Large-cap cryptocurrencies (BTC, ETH, LTC) are frequently *net transmitters*; ESG leaders tend to be *receivers* from large-cap crypto but *transmitters* to green cryptos; green cryptocurrencies (ADA, MIOTA, XLM, XNO, XRP) are persistent *receivers* (Alharbi et al., 2025; Ali et al., 2024; Bibi et al., 2025).

4.4.2 Mechanism: precedence, depth, and benchmark flows

Role assignment reflects: (i) *shock precedence*: which market underwrites the common factor first (crypto idiosyncratic shocks ⇒ BTC/ETH transmit; macro/ESG shocks ⇒ ESG transmit); (ii) *liquidity depth*: thinner order books in green cryptos magnify pass-through; (iii) *benchmark/flow amplification*: ESG index shocks transmit to ESG-thematic crypto through cross-asset rebalancing.

Table 4.2: Directionality Map and Channels

Pair	**Role**	**Channel**	**Implication**
BTC/ETH → ESG Leaders	TX→RX	Global risk/funding factor	System-wide propagation in crypto shocks.
ESG (E/S/G) → Green Crypto	TX→RX	Benchmark flows, thematic premia	Green crypto absorbs ESG shocks.
ESG Leaders ↔ BTC/ETH	State-dependent	Precedence switches by regime	Role flip around crisis idiosyncrasies.

Note: This table summarises typical *directional transmission* roles between asset groups. TX = transmitter; RX = receiver; ESG = Environmental, Social, and Governance; BTC = Bitcoin; ETH = Ethereum. "Channel" indicates the dominant economic mechanism; "Implication" states the portfolio takeaway. *Source: The author.*

4.5 Portfolio Statistics: From Network Measures to Hedge Design

4.5.1 Hedge ratios and connectedness

For asset i hedged with j,

$$\beta_{i|j,t} = \frac{\mathrm{Cov}_t(r_i, r_j)}{\mathrm{Var}_t(r_j)} = \rho_{ij,t}\frac{\sigma_{i,t}}{\sigma_{j,t}}, \qquad \mathrm{HE} = 1 - \frac{\mathrm{Var}(r_P^{(H)})}{\mathrm{Var}(r_P^{(U)})}. \tag{4.6}$$

Because $\rho_{ij,t}$ increases with the off-diagonal FEVD mass (both respond to $A_{h,t}$ and Σ_t), *crisis* TCI spikes inflate $\beta_{i|j,t}$ mechanically. This is the principal reason hedges become larger and more costly during stress (Ali et al., 2024; Bibi et al., 2025).

4.5.2 Allocation rules and crisis performance

Bibi et al. (2025) show that the minimum-variance portfolio (MVP) often concentrates in ESG leaders (low volatility constituent), which can deliver *negative* HE for the ESG leg when downside co-movement lifts. Minimum-correlation (MCP) and minimum-connectedness (MCoP) portfolios diversify across assets and are more resilient when TCI is elevated, as MCoP explicitly penalises loadings on the spillover network.

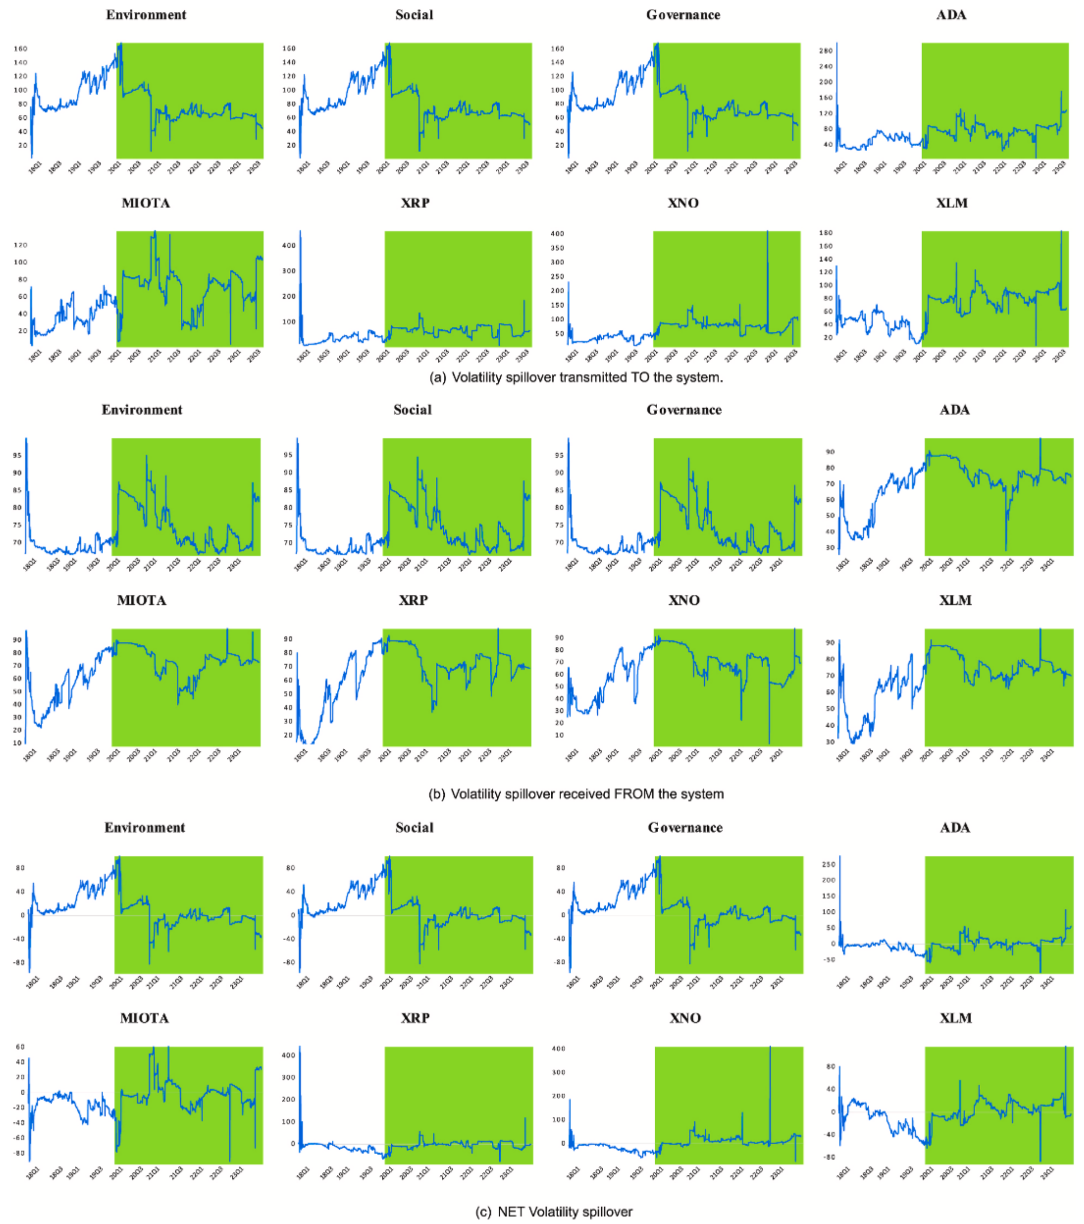

Figure 4.1: Volatility spillover dynamics between ESG indices and sustainable cryptocurrencies.

Note. Adapted from Alharbi et al. (2025). Panel (a) shows volatility spillovers transmitted **to** the system, Panel (b) displays volatility spillovers received **from** the system, and Panel (c) presents **net** volatility spillovers. Results are derived using the Diebold–Yilmaz connectedness framework on daily returns, where higher values indicate stronger directional volatility transmission.

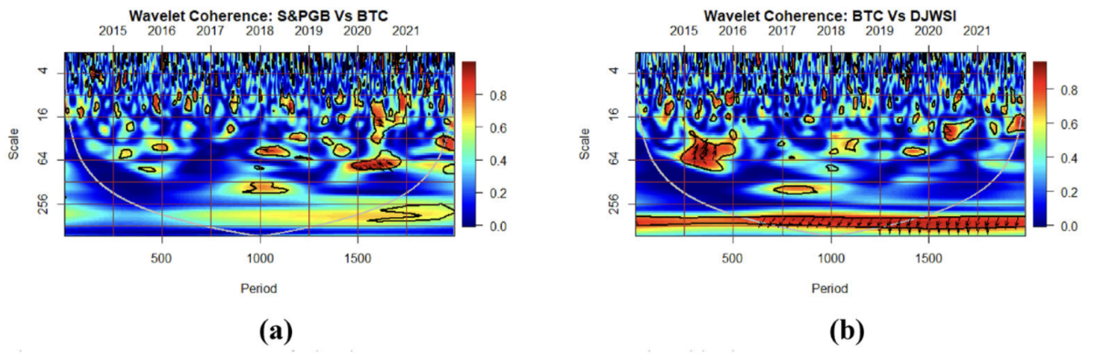

Figure 4.2: Wavelet coherence between Bitcoin, green bonds, and the Dow Jones World Sustainability Index.

Note. Adapted from Ul Haq et al. (2023). Panel (a) displays wavelet coherence between the S&P Green Bond Index (S&PGB) and Bitcoin (BTC), while Panel (b) shows wavelet coherence between BTC and the DJWSI. Warmer colors indicate higher coherence (red values approach 1). White contour lines denote the 5% significance level based on Monte Carlo simulations, and the cone of influence marks regions where edge effects are important.

4.6 Horizon Dependence: Time–Frequency Co-Movement

Wavelet coherence reveals that (i) green bonds and ESG indices are *in-phase* at medium/low frequencies; (ii) BTC and DJS World are *out-of-phase* at long horizons (structural diversification), with *in-phase* short-horizon spikes in crises (Ul Haq et al., 2023).

Mechanism. Long-horizon in-phase GB–ESG behaviour reflects common discount-rate and climate-policy cycles. BTC's long-horizon anti-phase vis-à-vis ESG equities is consistent with distinct growth/cash-flow versus adoption/liquidity drivers in calm regimes. Under stress, a short-horizon *liquidity* factor dominates both, producing temporary co-movement surges.

4.7 Systemic Risk and Tail Dependence

Let (X, Y) be returns on two assets with continuous marginals. Upper/lower tail dependence coefficients are

$$\lambda_U = \lim_{u \to 1^-} \Pr\Big(Y > F_Y^{-1}(u) \mid X > F_X^{-1}(u)\Big), \quad \lambda_L = \lim_{u \to 0^+} \Pr\Big(Y < F_Y^{-1}(u) \mid X < F_X^{-1}(u)\Big), \tag{4.7}$$

copula-invariant measures of joint extremes. Conditional systemic risk is assessed via CoVaR (**AdrianBrunnermeier2016**):

$$\Pr\Big(Y \leq \mathrm{CoVaR}_q^{Y|X} \;\Big|\; X = \mathrm{VaR}_q^X\Big) = q, \qquad \Delta\mathrm{CoVaR}_q^{Y|X} = \mathrm{CoVaR}_q^{Y|X=\mathrm{VaR}_q^X} - \mathrm{CoVaR}_q^{Y|X=\mathrm{Median}(X)}. \tag{4.8}$$

Facts and mechanism. Mzoughi et al. (2024) show that clean-energy and blockchain equities exhibit elevated λ_L and ΔCoVaR in crises (shared growth/technology and discount-rate exposure), whereas green bonds retain low tail dependence, preserving diversification. The *duration* and policy-support structure of green bonds breaks the equity co-crash channel.

Table 4.3: Tail Dependence and ΔCoVaR Across ESG–Crypto Adjacent Pairs

Pair	**Driver**	**Crisis outcome**	**Diversification**
Clean Energy ↔ Blockchain Eq.	Growth/discount rate	λ_L ↑, ΔCoVaR ↑	Low
Green Bonds ↔ Blockchain Eq.	Duration/policy support	λ_L low, ΔCoVaR modest	High

Note: This table summarises tail-risk co-movement between ESG-adjacent assets and blockchain equities. λ_L = lower-tail dependence (probability of joint downside extremes); ΔCoVaR = change in conditional Value-at-Risk when the conditioning asset is at its own VaR; "Eq." = equities. The Diversification column indicates relative usefulness under stress. *Source: The author.*

4.8 Quantile-Dependent Macro Sensitivities

Let r_t^{BTC} be BTC returns and $\mathbf{X}_t$ include changes in producer prices (PPI), USD/CNY, and the U.S. 10-year yield. The quantile regression

$$Q_\tau(r_t^{\mathrm{BTC}}|\mathbf{X}_t) = \alpha(\tau) + \beta_{\mathrm{PPI}}(\tau)\,\Delta\mathrm{PPI}_t + \beta_{\mathrm{USD/CNY}}(\tau)\,\Delta\mathrm{USD/CNY}_t + \beta_{10\mathrm{Y}}(\tau)\,\Delta y_t^{10\mathrm{Y}} + \boldsymbol{\gamma}(\tau)^\top \mathbf{Z}_t \tag{4.9}$$

shows *state-dependence*: PPI turns increasingly *negative* at upper quantiles; USD/CNY and 10Y flip sign from positive (lower quantiles) to negative (upper quantiles) (Lin et al., 2025). At lower quantiles (stress), dollar strength and rising yields can coincide with liquidity hoarding and reopening dynamics, supporting crypto demand; at upper quantiles (booms), discount-rate tightening and a stronger dollar weigh on risky assets. PPI proxies input/energy inflation; in booms it signals policy tightening, depressing high-quantile crypto returns.

Table 4
Results of the Quantile regression test (*price*).

	q10	q25	q50	q75	q90
	btc	btc	btc	btc	btc
price	0.0000	−0.0001*	−0.0001*	−0.0000	−0.0002***
	(0.2445)	(-1.9445)	(-1.7152)	(-0.2157)	(-2.8494)
shin	0.5628	0.3993**	0.2359*	0.2199	−0.3307
	(1.1521)	(2.5496)	(1.6806)	(1.0624)	(-0.6805)
ag	0.6163	0.3763*	0.2551***	0.0931	0.4251
	(1.4764)	(1.8009)	(2.6788)	(0.8771)	(1.4999)
cyb	−0.1765	−0.0664	0.0056	0.0602	0.0655
	(-0.5790)	(-0.6036)	(0.0676)	(0.6133)	(0.2998)
au	−1.0757	−0.3129	−0.2150	0.2426	0.0724
	(-1.5874)	(-0.8070)	(-0.9948)	(1.0997)	(0.1295)
vix	−0.0008	−0.0003	−0.0001	0.0002	0.0007*
	(-1.6181)	(-1.5732)	(-0.9294)	(0.8874)	(1.8946)
stk	0.0019	−0.0014	−0.0137*	−0.0119	0.0105
	(0.0954)	(-0.1315)	(-1.7594)	(-1.0370)	(0.5649)
_cons	−0.0317**	0.0062	0.0146*	0.0178*	0.0644***
	(-2.1086)	−0.6688	−1.9034	−1.936	−3.2239

Table 5
Results of the Quantile regression test (*usd*).

	q10	q25	q50	q75	q90
	btc	btc	btc	btc	btc
usd	0.0335***	0.0156***	0.0002	−0.0143***	−0.0275***
	(3.1147)	(2.7723)	(0.0560)	(-3.3831)	(-4.8738)
shin	0.4459	0.4368***	0.2443**	0.0726	−0.5075
	(1.6419)	(2.8605)	(2.1207)	(0.4849)	(-1.0976)
ag	0.5716**	0.3805***	0.2608**	0.1460	0.3492
	(2.2013)	(2.9762)	(2.4750)	(1.4916)	(1.5538)
cyb	−0.1459	−0.0888	−0.0016	0.0848	0.1314
	(-0.7717)	(-1.0434)	(-0.0211)	(0.6707)	(0.6089)
au	−0.7748**	−0.3204*	−0.1735	0.1508	0.1286
	(-2.0038)	(-1.6627)	(-0.9971)	(0.6270)	(0.2384)
vix	−0.0004	−0.0003**	−0.0000	0.0003	0.0005
	(-0.8927)	(-2.1614)	(-0.5329)	(1.0780)	(1.0481)
stk	0.0175	−0.0023	−0.0125	−0.0091	0.0226
	(0.7916)	(-0.3117)	(-1.3019)	(-0.6139)	(0.8222)
_cons	−0.2621***	−0.1180***	−0.0002	0.1126***	0.2275***
	(-3.4310)	(-3.0394)	(-0.0077)	−3.9932	−5.171

Table 6
Results of the main regression test (*Int*).

	q10	q25	q50	q75	q90
	btc	btc	btc	btc	btc
Int	−0.0250***	−0.0064	0.0015	0.0066	0.0153***
	(-6.0647)	(-1.4977)	(0.8091)	(1.2604)	(2.7540)
shin	0.0684	0.4142***	0.2569	0.0605	−0.4225
	(0.2168)	(2.6525)	(1.5474)	(0.2299)	(-1.2362)
ag	0.7256**	0.3650**	0.2539**	0.1388	0.4658**
	(2.1684)	(2.1417)	(2.1060)	(0.9561)	(2.3629)
cyb	−0.0087	−0.0673	−0.0064	0.1128	0.0665
	(-0.0465)	(-0.6516)	(-0.0495)	(1.3141)	(0.4336)
au	−1.1548**	−0.3121	−0.1750	0.1528	−0.0264
	(-2.0030)	(-0.9862)	(-0.8005)	(0.5593)	(-0.0474)
vix	−0.0007*	−0.0004**	−0.0000	0.0003	0.0010***
	(-1.8214)	(-2.2614)	(-0.4266)	(1.3844)	(3.3850)
stk	0.0199	−0.0004	−0.0137**	−0.0114	0.0191
	(1.0545)	(-0.0367)	(-2.0174)	(-0.8478)	(0.6104)
_cons	0.0460***	0.0101	−0.0033	−0.0049	−0.0144
	(3.0601)	(0.7231)	(-0.5190)	(-0.2792)	(-0.6740)

6

Figure 4.3: Quantile regression results for Bitcoin and macroeconomic instruments.

Note. Adapted from Lin et al. (2025). Results are reported for BTC against macroeconomic variables across quantiles (τ): price (Table 4), USD/CNY exchange rate (Table 5), and 10-year government bond yields (Table 6). Coefficients are followed by t-statistics in parentheses, with significance denoted as * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$. Findings indicate that PPI becomes increasingly negative as $\tau \to 1$, while USD/CNY and 10Y yields shift from positive at low quantiles to negative at high quantiles, consistent with a liquidity-to-discount-rate dominance shift.

4.9 Shock Taxonomy: Exogenous Jumps vs. Endogenous Negative Bubbles

Fry and Cheah (2016) distinguish *exogenous* policy/technology shocks (e.g., 2013 China ban) from *endogenous* negative bubbles (2014–2015 BTC/XRP), using hazard-based diagnostics. Endogenous episodes exhibit accelerating drawdowns with rising crash hazard (self-excited order flow, leverage spirals) and documented contagion (XRP→BTC); exogenous shocks produce discrete jumps with limited pre-crash hazard build-up.

Governance cue. Exogenous shocks justify rapid but reversible hedging until policy clarity emerges; endogenous negative bubbles warrant structural de-risking and tighter leverage controls.

Table 4.4: Shock Types, Hazard Profiles, and Policy Response

Shock	**Hazard path**	**Contagion**	**Response**
Exogenous policy (2013)	Flat → jump	Limited	Temporary protection
Endogenous negative bubble (2014–2015)	Rising	XRP → BTC	Structural de-risking

Note: This table distinguishes *exogenous* shocks (policy/technology news that arrive as discrete jumps) from *endogenous* negative bubbles (self-reinforcing drawdowns with a rising crash hazard). "Hazard path" summarises the time profile of crash risk; "Contagion" indicates a typical transmission direction; "Response" states the portfolio playbook (short-lived hedges for exogenous events vs. more durable de-risking for endogenous episodes). *Source: The author.*

4.10 Market Microstructure and Segmentation

4.10.1 Facts

During 2017–2018, cross-border price gaps ("Kimchi premium") of 20–60% persisted between Korean and U.S. venues; within-country gaps were small. A common order-flow factor explains most daily return comovement across venues even when levels diverge (Makarov & Schoar, 2020).

4.10.2 Mechanism: arbitrage incompleteness and basis risk

Define the cross-venue premium $\pi_t^{(k,\ell)} = 100\,[P_t^{(k)} - P_t^{(\ell)}]/P_t^{(\ell)}$. Under capital controls, slow fiat rails, and settlement frictions, the arbitrage loop cannot be closed quickly; $\pi_t^{(k,\ell)}$ becomes *persistent*. A hedge placed in venue k against exposure priced in ℓ then inherits *basis risk*: high short-horizon correlation (due to a common order-flow factor) coexists with level divergence, degrading hedge tracking.

Execution governance. (i) Hedge within the same jurisdiction/capital-control regime; (ii) impose real-time basis caps that trigger hedge relocation or overlay spreads; (iii) size liquidity and margin buffers to settlement lags and stablecoin convertibility risk.

4.11 Robustness, Sensitivity, and External Validity

Parameter sensitivity. Connectedness magnitudes vary with lag length p and horizon H; crisis rankings are stable across $H \in [10, 30]$, but absolute levels shift. Too-short rolling windows inflate variance; too-long windows dilute regime shifts.

Method triangulation. TVP–VAR findings align with DCC-/BEKK-GARCH correlations (for hedge ratios), empirical copulas (for tails), and wavelet coherence (for horizon structure) (Ali et al., 2024; Mzoughi et al., 2024; Ul Haq et al., 2023). Cross-method convergence strengthens causal interpretation.

Economic, not only statistical, significance. Variance reduction without Sharpe improvement has limited value; turnover and transaction costs must be incorporated. In high-TCI states, MCoP's drawdown control often dominates MVP on crisis-period Sharpe (Bibi et al., 2025).

4.12 Limitations

(i) **Model specification risk:** linear VARs may understate threshold/nonlinear propagation; Markov-switching or nonlinear state-space models mitigate but increase complexity. (ii) **Small-sample bias:** short rolling windows inflate parameter uncertainty in FEVD shares. (iii) **Tail estimation:** parametric copulas can misspecify extremal dependence; nonparametric approaches require large samples. (iv)

ESG measurement heterogeneity: index construction differences and "green crypto" classifications vary across studies, affecting comparability.

4.13 Implications for Allocation, Hedging, and Disclosure

Allocation and hedging. (a) Avoid concentrated MVP in high-TCI states; (b) prefer MCoP/MCP with turnover controls; (c) calibrate hedge ratios with connectedness and macro quantiles (Sections 4.5 and 4.8); (d) prioritise green bonds over clean-energy equities when hedging blockchain-equity sleeves in stress (low λ_L, smaller ΔCoVaR).

Crisis playbook. When TCI$^{(-)}$ surges: (i) increase defensive overlays (green bonds) and reduce clean-energy equity exposure; (ii) tighten leverage caps and drawdown limits; (iii) activate basis monitors and jurisdictional matching (Section 4.10).

Policy/reporting. Map risk disclosure to ESRS/ISSB: report (i) TCI and sign-decomposed connectedness; (ii) regime-specific hedge ratios and HE; (iii) tail metrics (λ_L, ΔCoVaR); and (iv) environmental metrics for blockchain protocols (Alharbi et al., 2025; Mzoughi et al., 2024).

Chapter 5

ESG Pillars for Digital Assets: Measurement, Governance, and Integration

5.1 Purpose and Scope

This chapter develops an exhaustive, pillar-by-pillar treatment of the Environmental (E), Social (S), and Governance (G) dimensions as they pertain to digital assets and crypto-adjacent market infrastructure. Whereas earlier chapters concentrated on econometric interdependence, risk transmission, and allocation mechanics, the emphasis here is on *embedding* sustainability into portfolio construction and stewardship—not as an afterthought, but as a co-determinant of investment eligibility and ongoing oversight under ESG-aligned mandates.

The analysis synthesises the mainstream ESG literature together with a growing corpus focused on crypto ecosystems to propose a framework for *measurement*, *governance*, and *disclosure* of ESG attributes in portfolios that include digital assets. Concretely, we translate evidence from environmental lifecycle assessments (LCAs) of blockchain protocols, social-impact research on adoption pathways, and governance studies of decentralised networks into decision-useful indicators that map to prevailing disclosure regimes—notably the European Sustainability Reporting Standards (ESRS) under the CSRD, the ISSB's IFRS S1/S2 baseline, and the Global Reporting Initiative (GRI) standards.

Methodologically, the scope goes beyond qualitative discussion. For each pillar, we identify the channels through which crypto assets may *create* sustainability risks or *deliver* sustainability benefits and propose quantification methods that are both scientifically defensible and reporting-compatible. For the E pillar, this means disaggregating energy and carbon intensity by consensus algorithm, regional grid composition, and hardware life-cycle effects, including rebound and carbon leakage from miner/validator relocation. For the S pillar, we evaluate financial-inclusion and consumer-protection trade-offs, local community externalities from infrastructure siting, and the welfare implications of privacy/data governance. For the G pillar, we focus on decentralisation and capture-resilience, transparency of decision-making, MEV governance, and the operational governance of exchanges and custodians.

By situating ESG diagnostics in the same decision frame as connectedness and portfolio risk, the chapter reflects the reality that sustainability and risk–return are inextricably linked in long-horizon asset allocation. It also addresses a common failure mode in practice and scholarship: the siloed treatment of E, S, and G—for example, reporting environmental metrics divorced from governance quality, or asserting social outcomes without verifiable evidence or framework alignment. The structure is intentionally modular yet interlinked: each pillar is analysed in depth, with explicit cross-references to the others and to the connectedness evidence established in previous chapters whenever systemic linkages are present.

The objectives are twofold:

- Equip the reader with a rigorous understanding of how ESG pillars are *operationalised* in crypto, including measurement frictions and methodological debates that drive dispersion in reported outcomes.
- Provide a ready-to-implement mapping between internal ESG metrics for crypto portfolios and external reporting frameworks, ensuring disclosures that satisfy regulatory expectations and maintain analytical integrity.

5.2 Environmental Pillar (E)

Environmental impact is the most visible and contested aspect of crypto sustainability because *consensus mechanisms*—the rules that secure the ledger—differ radically in energy demand, carbon footprint, and hardware lifecycles. These impacts are not stationary: they vary with market prices, regulatory shocks, and protocol upgrades. Robust environmental assessment must therefore be *dynamic*, *geographically disaggregated*, and *methodologically explicit*.

5.2.1 Energy and Carbon Accounting: PoW vs. PoS

Consensus is the first-order driver of energy use. In **PoW** systems (e.g., Bitcoin; pre-Merge Ethereum), security is externalised to competitive computation. Difficulty adjusts to keep block intervals stable, creating a tight linkage between price, hash rate, and power draw. A practitioner identity that captures this intuition is:

$$\text{Miner revenue} \approx \Big(\text{block subsidy} + \text{fees}\Big) \times P_{\text{coin}}, \tag{5.1}$$

so that rising prices attract more hash power, difficulty ratchets up, and network electricity consumption scales accordingly.

By contrast, **PoS** and related designs (DPoS, NPoS, FBA) decouple validation from brute-force computation. Validators are selected probabilistically (often proportionate to stake), and operational energy resembles conventional server loads rather than dedicated mining farms. Empirical work indicates orders-of-magnitude lower *operational* energy intensity for PoS relative to PoW, with indicative values in the Wh/transaction range for PoS and the kWh/transaction range for large PoW networks (Wendl et al., 2023).

On "energy per transaction". In PoW, the oft-cited Wh/tx is a poor efficiency proxy because throughput is governed by protocol parameters and demand, not by marginal computation. A more robust measure is total operational energy over a window T,

$$E_{\text{op}}(T) = \sum_{t=1}^{T} P_{\text{net},t}\,\Delta t, \tag{5.2}$$

where $P_{\text{net},t}$ is instantaneous network power (MW).

Carbon with geographic heterogeneity. Let E_j be energy for site/node j and $CI_{\text{grid},j}$ its local grid carbon intensity (tCO_2/MWh). A portfolio-level intensity is

$$CI = \sum_{j=1}^{N} \left(\frac{E_j}{\sum_\ell E_\ell} \times CI_{\text{grid},j} \right), \tag{5.3}$$

highlighting that identical protocols can manifest very different footprints solely due to the geography of miners/validators (Mulligan et al., 2024).

Portfolio implications.

- Prioritise *PoS-heavy* exposure for ESG sleeves, given its sharply lower operational energy.
- If PoW assets are retained for liquidity or diversification, require region-specific carbon accounting and renewable attestations.

- Disclose *ranges* rather than point estimates to reflect methodological dispersion and sampling uncertainty (Mulligan et al., 2024).

Note. A more precise PoW revenue expectation is $\mathbb{E}[R_{i,t}] = s_{i,t}\,\lambda_t\,(BR_t + Fee_t)P_t$, with $s_{i,t}$ miner i's hash share and λ_t block frequency. The heuristic in (5.1) is retained as a policy-relevant simplification; the price–hash–energy feedback remains.

5.2.2 E-Waste and Hardware Turnover

Beyond electricity and carbon, *hardware churn* is a significant externality, especially under PoW. ASIC competitiveness depends on hash-per-watt and hash-per-dollar; as difficulty rises or price falls, older rigs become uneconomic and are decommissioned. A device-level profitability screen is

$$\Pi_{\text{ASIC}}(t) = \left(\frac{H_{\text{ASIC}}}{H_{\text{net}}}\right)\cdot(\text{subsidy}+\text{fees})\cdot P_t \;-\; C_{\text{energy}}(t) \;-\; C_{\text{capex/opex}}(t). \tag{5.4}$$

Retirements create e-waste from obsolete ASICs/GPUs and ancillary gear (PSUs, cooling, networking).

LCA decomposition. Differentiate *operational* emissions (electricity) from *embodied* emissions (semiconductor fabrication, logistics). Poor end-of-life treatment exacerbates embodied-carbon burdens.

Stock–flow representation. With vintage stock $S_v(0)$ and lifetime L_v (years), pulse disposals at horizon T can be represented as

$$W_{\text{e-waste}}(T) \;=\; \sum_v \frac{S_v(0)}{L_v}\,\mathbb{1}\{T \bmod L_v = 0\}. \tag{5.5}$$

For staggered retirements, replace the indicator with a lifetime density (e.g., Weibull) and integrate.

Policy and disclosure.

- Require WEEE-compliant recycling attestations from miner counterparties.
- Report portfolio-level e-waste estimates using stock–flow models, including uncertainty bands.
- Recognise jurisdictional differences in recycling capacity and Basel Convention enforcement (Wendl et al., 2023).

5.2.3 Rebound and Carbon-Leakage Dynamics

Emissions can *increase* globally even as they fall locally if mining relocates to more carbon-intensive grids. Let H_k be regional hash share, E_k energy per TH/s, and EF_k the grid emissions factor. The relocation effect is

$$\Delta \mathrm{tCO}_2 \;=\; \sum_j H_j^{\mathrm{post}}\, E_j\, EF_j \;-\; \sum_i H_i^{\mathrm{pre}}\, E_i\, EF_i. \tag{5.6}$$

If $EF_{\mathrm{dest}} > EF_{\mathrm{orig}}$, total emissions rise despite domestic reductions.

Mechanism. Policy or price shocks $\Rightarrow$ relocation to cheaper or less regulated grids $\Rightarrow$ average network intensity can worsen $\Rightarrow$ secondary restrictions may trigger additional moves. This is especially salient for mobile PoW operations.

Investor actions.

- Scenario-test relocation pathways and update H_r accordingly.
- Apply *regionalised* carbon accounting; network averages can be misleading.
- Demand electricity-sourcing attestations, backed by verifiable RECs.

A relocation-adjusted portfolio measure is

$$\mathrm{tCO}_{2,\mathrm{adj}} = \sum_r H_r \cdot \mathrm{tCO}_{2,r} \cdot w_r, \tag{5.7}$$

with w_r capturing scenario probabilities.

5.2.4 Platform Snapshots and the "Green Blockchain" Segment

Sustainable portfolio design requires *protocol-level* benchmarking. Key indicators include:

- **Consensus type** (PoW vs. PoS/FBA/NPoS) as the primary energy driver.
- **Indicative energy intensity** (Wh/tx; with caveats on interpretability).
- **Annual energy** (MWh/yr) and **emissions** (tCO_2/yr), both location- and market-based.
- **Method notes**: functional unit, boundary (operational vs. LCA), and observation window.

Platforms emphasising low energy or renewable sourcing include *Cardano, Stellar, XRP Ledger, Ethereum (post-Merge), Polkadot, Avalanche, Nano, IOTA*. Independent estimates can diverge materially due to boundary and sampling differences; reporting *ranges* with footnoted assumptions is recommended (Alzoubi & Mishra, 2023; Mulligan et al., 2024).

Table 5.1: Indicative Platform Energy/Carbon Benchmarks

Protocol	**Consensus**	**Wh/tx (range)**	**MWh/yr (range)**	**tCO_2/yr (range)**	**Methodological notes**
ADA	PoS	X–Y	X–Y	X–Y	Validator set size; renewable sourcing disclosures.
XLM	FBA-like	X–Y	X–Y	X–Y	Low validator count; estimates from independent studies.
XRP	RPCA	X–Y	X–Y	X–Y	Curated UNLs; partial decentralisation.
ETH (post-Merge)	PoS	X–Y	X–Y	X–Y	∼99% energy drop post-Merge (Wendl et al., 2023).
DOT	NPoS	X–Y	X–Y	X–Y	Nomination pools; validator curation.
AVAX	PoS	X–Y	X–Y	X–Y	Subnet heterogeneity.
XNO	PoS (block-lattice)	X–Y	X–Y	X–Y	Minimal validator load.
MIOTA	DAG	X–Y	X–Y	X–Y	Coordinator roadmap impacts intensity.

Note: Populate with ranges from Alzoubi and Mishra (2023) and protocol disclosures; document functional unit and boundary for comparability.

5.3 Social Pillar (S)

The S pillar addresses human and community outcomes associated with blockchain technologies. Although E metrics dominate headlines, social factors increasingly determine product eligibility and stewardship priorities (Mulligan et al., 2024). We formalise the main channels, propose measurable indicators, and link them to recognised frameworks.

5.3.1 Financial Inclusion vs. Consumer Protection

Inclusion channel. Crypto payment rails—especially *stablecoins*—can lower remittance costs, reduce settlement latency, and bypass weak financial infrastructures in underbanked regions. Humanitarian payouts and migrant corridors are salient examples (Mulligan et al., 2024).

Protection channel. These benefits coexist with risks: high volatility, scams, opaque fees, and inadequate disclosures can harm vulnerable users. Stewardship requires transparency, plain-language risk labels, and ex-ante cost disclosure.

Quantification. With corridor transactions TX_t and inclusion-attributed TX_t^{incl},

$$\text{InclusionShare}_t = \frac{TX_t^{\text{incl}}}{TX_t}, \qquad \text{ComplaintRate}_t = \frac{N_t^{\text{complaints}}}{TX_t} \times 10^5, \tag{5.8}$$

$$\text{LossSeverity}_t = \frac{\sum_{i=1}^{N_t^{\text{loss}}} \text{Loss}_{i,t}}{N_t^{\text{loss}}}. \tag{5.9}$$

5.3.2 Labour, Community, and Local Externalities

Mining/hosting can create jobs and ancillary demand but also strain grids, generate noise/heat, and alter local conditions. Due diligence for infrastructure exposures should document: community consultation results; employment metrics (FTEs, wage premia); grid-impact assessments (peak load, curtailment protocols); nuisance controls; and grievance redress (Wendl et al., 2023).

5.3.3 Data Rights and Privacy

Public ledgers enable transaction-graph analysis and potential re-identification. While useful for AML/CFT, this raises privacy concerns. Responsible use entails: privacy-preserving tech (ZKPs, selective disclosure credentials), strong data-governance policies (retention, access, incident response), and user-facing disclosures on traceability (Mulligan et al., 2024).

5.3.4 Candidate S-Metrics and Framework Links

5.3.5 Stewardship and Product Design

- **Clarity by design:** plain-language risk labels and full cost ladders (spread, fees, gas) disclosed upfront.
- **Safeguards:** caps, escrow/insurance for retail-facing products in high-complaint corridors.

Table 5.2: Candidate Social Metrics for Crypto Portfolios

Metric	Definition	Unit	Verification	Framework
Inclusion share	Share of activity in inclusion use cases	% tx/-value	Chain analytics + partner data	GRI 203; ESRS S1–S4
Complaint rate	Complaints per 100k tx	Count/100k tx	Helpdesk + regulator logs	GRI 418; ESRS S1
Loss severity	Average loss per incident	Currency/incident	Case files; insurer evidence	GRI 418
Local jobs	Direct full-time equivalents (FTEs) from infrastructure	FTE	Payroll; audit	ESRS S1
Grid-impact MOU	Demand-response/curtailment agreement	Y/N	Signed utility MOUs	GRI 203
Privacy controls	ZK/SD deployed; DPIA conducted	Binary/maturity	Policy & tech review	GRI 418; ESRS S6
Grievance resolution	Median closure time for complaints	Days	Registry evidence	GRI 413

- **Infrastructure engagement:** require grid-impact MOUs and demand-response participation.
- **Privacy-by-default:** prefer auditable privacy controls; conduct DPIAs where appropriate.

Assurance. Compute metrics on harmonised windows, report uncertainty, and reconcile to transactions. Independent verification (ISAE 3000 or credible NGO/academic partner) strengthens assurance (Mulligan et al., 2024).

5.4 Governance Pillar (G)

Governance spans *on-chain* protocol governance and *off-chain* market infrastructure. Strong governance reduces operational risk, limits capture by dominant actors, and aligns networks with financial and sustainability norms (Mulligan et al., 2024).

5.4.1 Decentralisation and Capture-Resilience

Let validator (or miner) shares $\mathbf{s} = (s_1, \dots, s_n)$, $\sum s_i = 1$, sorted $s_{(1)} \geq \dots \geq s_{(n)}$. The *Nakamoto coefficient* is

$$N^\star = \min \left\{ k : \sum_{i=1}^{k} s_{(i)} \geq 0.5 \right\}, \tag{5.10}$$

and dispersion can be summarised by HHI, Gini, and geographic entropy. Reports should include time-series of $\{N^\star, \mathrm{HHI}, \mathrm{Gini}, \mathrm{GeoEntropy}\}$ with trend tests to detect recentralisation.

Quorum/supermajority. Disclose quorum q, supermajority threshold τ, turnout, pass rates, and proposer concentration to evaluate veto and capture risk.

5.4.2 Processes, Accountability, and Emergency Powers

High-quality governance features transparent roadmaps, open RFCs (e.g., EIPs), staged upgrades (proposal → audit → testnet → mainnet), and COI registries. Where emergency powers exist (halts, safe-modes, rollbacks), disclose scope, triggers, authorised actors, and test cadence. A simple risk proxy is

$$\mathrm{EPR} = \alpha \cdot \mathrm{Scope} + \beta \cdot (1 - \mathrm{TriggerClarity}) + \gamma \cdot (1 - \mathrm{TestCoverage}), \tag{5.11}$$

with inputs normalised to $[0, 1]$.

5.4.3 Tokenomics, Treasury, and MEV

Tokenomics/treasury. Disclose emissions schedules, burns/lockups, treasury governance, and transparency. A treasury transparency score (TTS) can aggregate audit status, reporting cadence, on-chain dashboards, and independent signer requirements.

MEV controls. Track MEV rate (MEV/fees), PBS adoption, builder HHI, and user-harm proxies (failed-tx share, inclusion latency). Absence of MEV mitigation harms both G (integrity) and S (user welfare).

5.4.4 Market Infrastructure: Custody, Exchanges, Compliance

Custody. Key elements: asset segregation, proof-of-reserves (PoR) coverage/cadence, key management (HSM/MPC; rotation; cold/hot split), and incident SLAs. Define

$$\mathrm{PoR\,Coverage} = \frac{\text{attested assets}}{\mathrm{AUM}}, \qquad \mathrm{PoR\,Cadence} = \frac{\text{attestations in year}}{12}. \tag{5.12}$$

Compliance and segmentation. AML/CFT, sanctions, and Travel Rule adherence are table stakes. Cross-border segmentation creates basis risk:

$$\pi_t^{(k,\ell)} = 100 \cdot \frac{P_t^{(k)} - P_t^{(\ell)}}{P_t^{(\ell)}} \ (\%), \tag{5.13}$$

which can impair hedging when capital controls or fiat rails bind (Makarov & Schoar, 2020).

5.4.5 Candidate G-Metrics and Framework Mapping

5.4.6 Stewardship and Policy

- **Protocol engagement:** set guardrails for $N^\star$, HHI, and MEV controls; escalate on adverse trends.
- **Voting:** support proposals improving decentralisation, transparency, PBS/fair-ordering, and emergency-power accountability.
- **Infrastructure covenants:** require PoR with independent assurance, robust key management, incident SLAs, and segregation attestations.
- **Cross-border risk:** maintain venue whitelists and basis monitors; prefer same-regime hedges when $\pi^{(k,\ell)}$ widens.

5.5 Data Integrity and Digital MRV

A pervasive challenge in crypto ESG is *measurement dispersion*—large spreads in reported Wh/tx, MWh/yr, tCO_2/yr, or e-waste estimates due to divergent boundaries, assumptions, and windows (Mulligan et al., 2024). Digital MRV (dMRV) uses blockchains to secure sustainability data from capture to disclosure (Figueiredo et al., 2022).

5.5.1 Role of dMRV

Trusted inputs. Energy meters, REC registries, IoT for hardware utilisation, and privacy-preserving geolocation attestations.

Oracles. Attested data are hashed on-chain. With payload d_t and nonce ν_t, publish $c_t = H(\text{serialize}(d_t)\|\nu_t)$; later disclosures reveal ν_t to verify.

Table 5.3: Governance Metrics for Digital-Asset Portfolios

Metric	**Definition**	**Unit**	**Target/Flag**	**Framework**
Nakamoto $N^{\star}$	Entities to control > 50%	Count	Higher better	ESRS G1
Validator HHI/Gini	Power concentration	Index	Lower better	ESRS G1; GRI 205
Geo-entropy	Regional dispersion	Index	Higher better	ESRS G1
Turnout/quorum	Governance participation	%	Higher better	ESRS G1
Proposer concentration	Top-k proposer share	%/HHI	Lower better	ESRS G1
TTS	Treasury transparency score	0–4	Higher better	ESRS G1; IFRS S1
MEV rate/PBS/builder HHI	MEV exposure/mitigation	%, %, Index	↓ MEV, ↑ PBS	ESRS G1
PoR coverage/cadence	Attested assets; frequency	%, /12	Higher better	ESRS G1
Key posture	MPC/HSM; rotation; cold share	Binary/%	Stronger better	ESRS G1
Incident MTTR/MTTD	Response/containment	Hours	Lower better	ESRS G1
Jurisdictional compliance	Licences; AML/CFT	Y/N	Full	ESRS G1; IFRS S1
Cross-venue premium π	Fragmentation risk	%	Lower better	ESRS G1

Note: Candidate governance KPIs and their desired direction. Acronyms: $N^{\star}$ = Nakamoto coefficient; HHI = Herfindahl–Hirschman Index; TTS = Treasury Transparency Score; MEV = Maximal Extractable Value; PBS = Proposer–Builder Separation; PoR = Proof of Reserves; MPC = Multi-Party Computation; HSM = Hardware Security Module; MTTR/MTTD = Mean Time To Repair/Detect; AML/CFT = Anti-Money Laundering/Combating the Financing of Terrorism; ESRS = European Sustainability Reporting Standards; GRI = Global Reporting Initiative; IFRS S1 = ISSB General Requirements. "Target/Flag" indicates the preferred direction for risk management (e.g., ↓ MEV means lower is better). *Source: The author.*

On-chain series. For metric $m \in \{\mathrm{Wh/tx}, \mathrm{MWh/yr}, \mathrm{tCO_2/yr}\}$, store $(t, m_t, u_t, \mathrm{fu}, \mathrm{sb}, \mathrm{win})$, with uncertainty u_t, functional unit, system boundary, and window.

Benefits. Tamper resistance, versioning, and near real-time transparency.

5.5.2 Assurance Triad

A credible programme combines: (i) on-chain attestations; (ii) independent assurance of provenance/calibration; (iii) replication by academic/NGO teams (Mulligan et al., 2024). Reports must specify functional unit, system boundary, time horizon, and geographic allocation rules.

Uncertainty. Publish ranges or $\hat{e} \pm z_{\alpha/2}\widehat{\sigma}$ with notes on hardware, block fill, and sampling.

5.5.3 Portfolio dMRV Implementations

Validator Carbon Dashboard. Location-based: $\mathrm{tCO}_{2,r} = E_r \cdot EF_r$. Market-based: $\mathrm{tCO}_{2,r}^{\mathrm{MB}} = \max(E_r EF_r - \theta_r \mathrm{REC}_r, 0)$, with REC serials on-chain.

E-Waste Ledger. Vintage-level retirements $W_{\mathrm{retire}}(t) = \sum_v S_v(0) f_v(t)$, with recycler certificates hashed on-chain.

Green PoS Registry. Validator energy sourcing with REC IDs; detect double counting through registry cross-checks.

5.5.4 Framework Mapping

Disclosure schema. For each metric: `metric_id, value, unit, lower, upper, fu, sb, win, oracle_tx, audit_hash`.

5.6 Alignment with ESRS/CSRD, ISSB, and GRI

Regulatory alignment is essential. ESRS (under CSRD), ISSB IFRS S1/S2, and GRI provide interoperable disclosure scaffolds. In crypto—where data quality and methodological divergence are persistent—framework mapping enhances credibility and assurance-readiness (Mulligan et al., 2024).

Table 5.4: dMRV Data Objects and External Frameworks

On-chain Field	Definition / Unit	Boundary	Framework Link
energy_mwh	Network energy use (MWh/yr; Wh/tx)	Operational vs. LCA	GRI 302; ESRS E1
ci_lb / ci_mb	tCO_2/yr (location-/market-based)	Grid mix; REC-adjusted	GRI 305; ESRS E1
weee_t	E-waste (t/yr) by vintage	Disposal path; recycler ID	ESRS E5
rec_id	Serialised REC linkage	Eligibility/retirement	ESRS E1; GRI 305
fu/sb/win	Functional unit, boundary, window	Versioned	All (assurance context)

Note: dMRV = digital Monitoring, Reporting, and Verification. `energy_mwh` records operational energy; `ci_lb`/`ci_mb` are carbon intensities under location-/market-based methods; `weee_t` tracks e-waste tonnage and recycler provenance; `rec_id` links Renewable Energy Certificates (RECs); `fu`/`sb`/`win` denote functional unit, system boundary, and time window. “Boundary” distinguishes operational-only vs. life-cycle assessment (LCA) scopes. “Framework Link” maps each object to disclosure standards (GRI, ESRS) for assurance-ready reporting. *Source: The author.*

5.6.1 ESRS/CSRD Integration

ESRS E1 (Climate): Scope 1/2/3 GHG, transition plans, targets, energy consumption/intensity, and renewable share. For crypto: protocol-level MWh/yr; tCO_2/yr (location- and market-based) with regional splits; REC evidence. **ESRS E5 (Resource Use)**: hardware lifecycles, e-waste tonnage, recycling/disposal. **ESRS S1/S4 (Social)**: workforce, community impact, data privacy. **ESRS G1 (Governance)**: structures, stakeholder engagement, anti-corruption, capture resilience.

5.6.2 ISSB S1/S2 Mapping

IFRS S1 requires governance, strategy, risk management, and metrics for material sustainability risks/opportunities (e.g., validator concentration, PoR cadence, MEV mitigation). **IFRS S2** parallels ESRS E1, emphasising climate resilience, scenarios, and transition planning—critical under relocation risk (§5.2.3).

5.6.3 GRI Relevance

GRI 302 (Energy), **GRI 305** (Emissions), **GRI 418** (Customer Privacy), and **GRI 205/206** (anti-corruption/anti-competitive) anchor disclosures for crypto portfolios, including exchange and DeFi governance.

5.6.4 ESG Reporting Alignment Matrix

Table 5.5: ESG Reporting Alignment Matrix (Illustrative)

Internal Metric	**Framework Reference**	**Disclosure Item**	**Gaps / Remedies**
Network energy (MWh/yr)	GRI 302; ESRS E1; IFRS S2	Method, boundary, window	High dispersion ⇒ ranges; document FU (Mulligan et al., 2024)
tCO_2/yr (regional)	GRI 305; ESRS E1; IFRS S2	LB vs. MB scopes	Dual reporting; REC attestation (Wendl et al., 2023)
E-waste (t/yr)	ESRS E5	Cohorts; disposal paths	Data gaps ⇒ WEEE requirements; sampling (Wendl et al., 2023)
Validator concentration	ESRS G1; IFRS S1	Capture resilience	No universal threshold ⇒ trend disclosure
PoR cadence	ESRS G1; IFRS S1	Custody integrity	Standardise independent assurance
Inclusion outcomes	ESRS S1; GRI 203	Beneficiary metrics	Attribution ⇒ third-party surveys (Mulligan et al., 2024)
Privacy controls	GRI 418; ESRS S4	Data governance; incidents	Clarify analytics/PII policies

Note: This matrix maps internal portfolio metrics to external disclosure frameworks and flags common assurance gaps with suggested remedies. Acronyms—GRI: Global Reporting Initiative; ESRS: European Sustainability Reporting Standards (under CSRD); IFRS S1/S2: ISSB standards (General Requirements / Climate); LB/MB: location-based / market-based emissions accounting; REC: Renewable Energy Certificate; WEEE: Waste Electrical and Electronic Equipment directive; PoR: Proof of Reserves; PII: Personally Identifiable Information; FU: functional unit. *Source: The author.*

5.6.5 Governance for Reporting

Operationalise the mapping by: (i) making framework alignment an onboarding requirement; (ii) updating annually or on major protocol changes (e.g., Ethereum's

Merge); and (iii) reproducing the mapping in portfolio ESG reports for transparency and auditability.

5.7 Blockchain and the Clean-Energy Transition

From a systems perspective, blockchain can advance the clean-energy transition (CET) by serving as decentralised *data integrity* infrastructure for energy and carbon markets, mitigating information asymmetry, verification lag, and fraud risk (Figueiredo et al., 2022; Popkova et al., 2023). Classical CET bottlenecks include: incomplete provenance information; high transaction/settlement costs; and trust deficits in certificates and emissions reporting.

5.7.1 Programmable RECs and Carbon Assets

RECs/carbon credits become tokenised objects with metadata linked to metered generation/emission reductions. Smart contracts enforce uniqueness, retirement finality, and eligibility filters, reducing double issuance/counting and compressing settlement latency relative to siloed registries (Popkova et al., 2023).

5.7.2 Grid Flexibility and DER Coordination

With rising VRE penetration, flexible demand L_t must adapt to $G_t - VRE_t$. Blockchain-based flexibility markets lower C_{tx} and support small-scale DER participation by self-bidding capacity, oracle-fed telemetry, and conditional payments on verified delivery ΔL_t.

5.7.3 dMRV for Construction and Heavy Industry

IoT sensors capture E_{input}, M_{use}, and W_{out}; hashes anchor records on-chain; material passports link batches to projects. Lifecycle embodied carbon is

$$C_{\text{emb}} = f\Big(M_{\text{use}}, E_{\text{input}}, EF_{\text{mat}}, EF_{\text{energy}}\Big), \tag{5.14}$$

with each input auditable and timestamped (Figueiredo et al., 2022).

5.7.4 Portfolio Implications

CET-enabling assets can contribute to positive scope-3 outcomes by facilitating decarbonisation elsewhere. A CET-adjusted score for asset a:

$$\mathrm{ESG}_a^{\mathrm{CET}} = w_E E_a + w_S S_a + w_G G_a + w_{\mathrm{CET}} \mathrm{CET}_a, \tag{5.15}$$

where CET_a is the verifiable share of activity from REC/carbon/DER markets and w_{CET} is mandate-specific.

5.8 Translating E/S/G into Portfolio Policy

Integrating ESG into allocation requires transforming intent into *quantitative* rules. Define the investable set

$$\mathcal{C}(P) = \{a_i \in P \mid E(a_i) \geq E_{\min},\ S(a_i) \geq S_{\min},\ G(a_i) \geq G_{\min}\}. \tag{5.16}$$

Environmental eligibility and carbon limits. Consensus type is a decisive filter: PoS (and similar) generally passes; PoW may be capped and accompanied by regionalised carbon accounting. Portfolio footprint

$$CF_P = \sum_i \frac{NAV_i}{NAV_P} CF_i \ \leq\ CF_{\max}, \tag{5.17}$$

aligns with ESRS E1 targets.

Embodied impacts and e-waste. Screen on hardware turnover δ_{HW} and recycling rates ρ_{recycle}; exclude or engage where circular-economy thresholds are not met.

Social safeguards. Monitor inclusion share

$$FI\_\mathrm{share}^{(i)} = \frac{\mathrm{Transactions}_{FI}}{\mathrm{Transactions}_{\mathrm{total}}}, \tag{5.18}$$

and consumer-protection indicators (complaints per thousand tx). Infrastructure exposures require community evidence and grid-impact documentation.

Governance thresholds. Set bounds for Nakamoto coefficient NC_i (e.g., $NC_i \geq 7$) and stake concentration (e.g., top share $\leq 33\%$). Custody requires periodic PoR with third-party audits; MEV mitigation must be disclosed.

ESG-constrained optimisation. Solve

$$\max_{\mathbf{w}} \boldsymbol{\mu}^\top \mathbf{w} - \frac{\lambda}{2}\mathbf{w}^\top \Sigma \mathbf{w} \quad \text{s.t.} \quad \mathbf{1}^\top \mathbf{w} = 1,\ \mathbf{w} \geq 0,\ \mathbf{w} \in \mathcal{C}(P),\ CF_P \leq CF_{\max},\ \text{G-thresholds}, \tag{5.19}$$

so allocation jointly satisfies financial and ESG constraints.

Telemetry and audit trail. Continuous monitors ingest on/off-chain data, triggering reviews on breaches (e.g., validator concentration rises, inclusion share falls). All checks are logged to create a verifiable audit trail for regulators and investors. This rules-based approach aligns with ISSB S1/S2 and ESRS expectations and ensures ESG performance is *integral* to investment decisions, not merely narrative.

Chapter 6

ESG-Integrated Portfolio Policy for Digital Assets

6.1 Purpose and Positioning

The purpose of this chapter is to convert the environmental, social, and governance pillars from high-level investment doctrine into a *codified* portfolio policy: a set of rules that are operational, quantifiable, and enforceable on a continuous basis. Within sustainable crypto investing, declarative commitments to "invest sustainably" are insufficient unless they are embedded directly in allocation and rebalancing mechanics. Accordingly, this section develops a framework in which ESG criteria are formalised as *eligibility screens*, *hard limits* on risk and exposure, *objective-function terms*, and *breach-handling* playbooks. Crucially, these constraints are designed as *ex-ante* requirements that shape the optimiser's feasible set *before* trades are placed, not as ex-post overlays.

The design aggregates empirical and theoretical insights from the literature. Environmental heterogeneity across consensus mechanisms and geographies (Alzoubi & Mishra, 2023; Wendl et al., 2023) informs the environmental screens. The prevalence of measurement dispersion and the imperative for harmonised disclosures (Mulligan et al., 2024) motivate explicit uncertainty bands and machine-auditable trails in score construction. The time-varying, regime-dependent nature of crypto–ESG connectedness (Alharbi et al., 2025; Ali et al., 2024; Bibi et al., 2025) guides dynamic risk overlays, while systemic and tail-risk behaviour (Mzoughi et al., 2024) shapes diversification and hedging rules. Macro-financial asymmetries (Lin et al., 2025) are incorporated through conditional exposure adjustments, and market-microstructure frictions (Makarov & Schoar, 2020) impose feasibility constraints on hedging and rebalancing.

6.2 Policy Architecture at a Glance

The policy is organised as a layered stack, progressing from the investable universe to real-time controls and disclosure alignment:

1. **Universe construction.** Assets must meet minimum E/S/G thresholds based on normalised scores derived from protocol characteristics and market-infrastructure evidence.

2. **Quantitative limits.** Binding constraints on environmental footprint, social outcomes, governance decentralisation, and market-structure vulnerability restrict the allocation domain.

3. **ESG-integrated optimisation.** ESG enters the objective alongside return and variance, and a connectedness penalty captures systemic fragility beyond pairwise correlation.

4. **Monitoring and breach management.** On-chain and off-chain telemetry feed continuous compliance checks; breaches trigger engagement or automatic de-risking. Reporting is mapped to ESRS, ISSB, and GRI.

The architecture ensures ESG requirements function as *pre-trade* constraints—binding at construction time—rather than narrative rationales after positions exist. Each rule is traceable to evidence in the literature, anchored to a precise metric, and paired with a specified remediation action.

6.3 Universe Construction and Hard Screens

Let P denote the raw universe spanning (i) native cryptoassets, (ii) crypto-adjacent equities and ETPs, and (iii) "green diversifiers" such as green bonds or clean-energy equities. The ESG-eligible investable set $\mathcal{C}(P)$ is:

$$\mathcal{C}(P) = \{a_i \in P \mid E(a_i) \geq E_{\min},\ S(a_i) \geq S_{\min},\ G(a_i) \geq G_{\min}\}, \tag{6.1}$$

where $E(\cdot)$, $S(\cdot)$, and $G(\cdot)$ are normalised $[0, 100]$ scores constructed from protocol-level features (e.g., consensus, validator dispersion) and market-infrastructure evidence (e.g., custody set-up, proof-of-reserves). Each score is accompanied by provenance and an uncertainty interval, explicitly acknowledging measurement dispersion (Mulligan et al., 2024). The screen is *conjunctive*: failure on any pillar disqualifies the asset, thereby preventing cross-pillar "netting out." Thresholds $E_{\min}, S_{\min}, G_{\min}$ are set by the investment committee, typically 50–70 depending on mandate strictness.

6.3.1 Environmental Eligibility: Consensus and Geography

The environmental filter has three axes: protocol preference, PoW caps, and region-sensitive carbon accounting.

PoS and other low-energy designs (e.g., FBA) exhibit orders-of-magnitude lower operational energy than PoW (Wendl et al., 2023); they therefore qualify by default, subject to other constraints. PoW assets may only be included for liquidity/diversification and face a portfolio cap:

$$\sum_{i \in \mathrm{PoW}} w_i \;\leq\; \bar{w}_{\mathrm{PoW}}, \tag{6.2}$$

with $\bar{w}_{\mathrm{PoW}}$ commonly 5–10%.

Each asset a_i receives a carbon factor CF_i (tCO_2/yr) that is *geographically adjusted* to miners'/validators' grid mix, avoiding misleading global averages. The portfolio footprint is:

$$CF_P \;=\; \sum_i \frac{NAV_i}{NAV_P} CF_i \;\leq\; CF_{\max}, \tag{6.3}$$

and is disclosed in both location- and market-based formats, aligned with ESRS E1 (Mulligan et al., 2024).

6.3.2 E-Waste and Embodied Impacts

Hardware churn—notably ASIC replacement under PoW—creates e-waste and embodied emissions (Wendl et al., 2023). For each asset, estimate hardware turnover $\delta_{\mathrm{HW},i}$ and unrecycled fraction $1 - \rho_{\mathrm{recycle}}(i)$. The portfolio e-waste load is:

$$EW_P \;=\; \sum_i w_i \, \delta_{\mathrm{HW},i} \Big(1 - \rho_{\mathrm{recycle}}(i)\Big) \;\leq\; EW_{\max}, \tag{6.4}$$

with $EW_{\max}$ aligned to ESRS E5 (resource use and circularity). Persistently low recycling performance triggers exclusion or engagement.

6.3.3 Social Eligibility

The social screen evaluates inclusion and protection. Inclusion is proxied by the share of transactions with characteristics consistent with inclusion use cases (e.g., low-cost remittances):

$$FI_{\text{share},i} \;=\; \frac{\text{Transactions}_{\text{inclusion}}(i)}{\text{Transactions}_{\text{total}}(i)}. \tag{6.5}$$

A portfolio floor enforces:

$$\sum_i w_i \cdot FI_{\text{share},i} \;\geq\; \overline{FI}, \tag{6.6}$$

with $\overline{FI}$ typically 10%. Uncertainty is reported via ranges/confidence intervals (Mulligan et al., 2024). Consumer protection is monitored through fraud/loss incidents per 1,000 transactions, complaint ratios, and dispute resolution times; assets above a policy percentile are ineligible or put on watch.

6.3.4 Governance Eligibility

Governance screens cover decentralisation, upgrade process, emergency powers, MEV mitigation, and custody transparency. Require a minimum Nakamoto coefficient $NC_i \geq NC_{\min}$ and a maximum stake concentration $SC_i \leq SC_{\max}$ for the top validator cohort (e.g., $\leq 33\%$). Protocols must disclose governance pathways, any centralised emergency powers, and MEV controls; non-disclosure is disqualifying (Mulligan et al., 2024). Listed securities/ETPs require at least quarterly audited proof-of-reserves, including scope and liabilities coverage.

6.4 Risk-Aware Limits Informed by the Literature

The static screens of §6.3 define eligibility, but portfolio risk management must adapt to the time-varying transmission documented empirically (Bibi et al., 2025; Lin et al., 2025; Mzoughi et al., 2024). We overlay three limit classes: connectedness/regime overlays, macro-quantile overlays, and market-structure guardrails.

6.4.1 Connectedness and Regime Overlays

During system-wide stress (e.g., COVID-19 onset; Russia–Ukraine escalation), TCI readings for ESG–crypto systems surged into the 70–85% range, with negative-return spillovers dominant and large-cap crypto often acting as net transmitters (Bibi et al., 2025). Let TCI_t denote the connectedness level. If

$$\text{TCI}_t \;\geq\; \tau_{\text{TCI}}, \tag{6.7}$$

with τ_{TCI} calibrated to a high-percentile threshold, then:

- **Tighten concentration** on known transmitters by lowering their caps $\bar{w}_i$.
- **Reweight diversifiers** toward low-connectedness green instruments (e.g., green bonds, certain clean-energy equities).
- **Shift hedging paradigm** from MVP to connectedness-sensitive constructions (MCoP/MCP), prioritising spillover reduction over pure variance minimisation (Bibi et al., 2025; Mzoughi et al., 2024; Ul Haq et al., 2023).

6.4.2 Macro-Quantile Overlays

Crypto betas to macro variables are quantile-dependent. Lin et al. (2025) show that BTC's PPI beta turns more negative in upper return quantiles, implying greater reversal risk for energy-intensive PoW assets during inflationary shocks. Let Q_t^{return} denote the return-quantile state. When

$$Q_t^{\mathrm{return}} \in \text{upper tail} \quad \text{and} \quad \mathrm{PPI}_t > \mathrm{PPI}_{\mathrm{thr}}, \tag{6.8}$$

then:

- Cut PoW weights,
- Add hedges (e.g., short BTC futures, long green-bond sleeves),
- Reverse tilts as inflation pressure abates and quantiles revert.

6.4.3 Market-Structure (Segmentation) Guardrails

Persistent cross-border segmentation (e.g., the "Kimchi premium" at 20–60%) impairs hedge efficacy and introduces execution risk (Makarov & Schoar, 2020). Define a segmentation index S_t as average cross-venue dispersion relative to a reference. If $S_t > S_{\mathrm{thr}}$:

- Constrain routing to venues aligned with the exposure's jurisdiction,
- Inflate slippage and transaction-cost inputs in optimisation,
- Temporarily reduce leverage and hedge sizes.

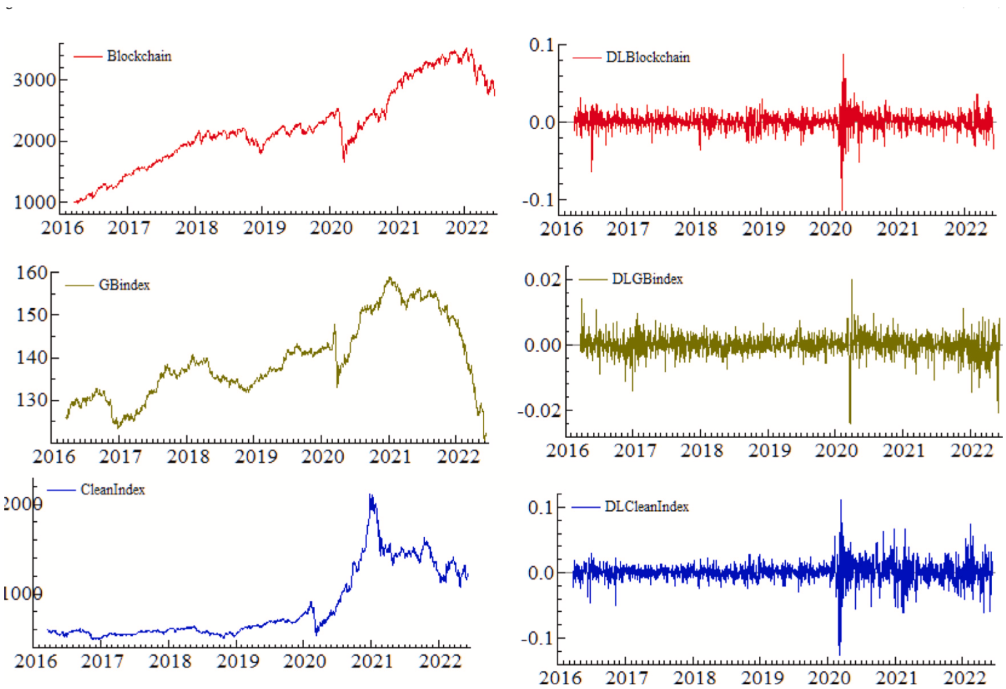

Figure 6.1: Daily prices and returns for green bonds, equities, and blockchain markets.

Note. This figure displays daily price series and corresponding returns for green bond indices, selected equity markets, and blockchain-related assets. The return series highlight volatility clusters and abrupt changes linked to market shocks, consistent with typical financial time-series characteristics.

Table 6.1: Risk-Aware Limit Triggers and Actions

Trigger	**Threshold**	**Action**
Total Connectedness Index	$\mathrm{TCI}_t \geq \tau_{\mathrm{TCI}}$	Tighten caps; tilt to diversifiers; adopt MCoP/MCP hedging.
Macro + Quantile	Upper-tail returns & $\mathrm{PPI} > \mathrm{PPI}_{\mathrm{thr}}$	Reduce PoW; add hedges; reverse on normalisation.
Segmentation Index	$S_t > S_{\mathrm{thr}}$	Jurisdictional routing; higher slippage; curb leverage/hedge size.

Note: Triggers implement regime-aware controls. TCI = Total Connectedness Index; τ_{TCI} = policy threshold (e.g., high-percentile level); PPI = Producer Price Index; PoW = Proof of Work; MCoP/MCP = Minimum Connectedness / Minimum Correlation Portfolio; S_t = segmentation index (cross-venue price dispersion). "Tighten caps" = reduce issuer/asset-class limits; "tilt to diversifiers" = increase weights in low-connectedness assets (e.g., green bonds); "reverse on normalisation" = unwind overlays when indicators fall back below thresholds. *Source: The author.*

6.5 Optimisation with ESG and Connectedness

Let $\mathbf{w} \in \mathbb{R}^n$ be portfolio weights, $\boldsymbol{\mu}$ expected returns, Σ the covariance matrix, and C_t a normalised connectedness matrix (PCI-based, scaled to $[0, 1]$) (Bibi et al., 2025). The objective blends return, variance, and systemic-linkage aversion:

Objective

$$\max_{\mathbf{w} \in \mathcal{C}(P)} \quad \boldsymbol{\mu}^\top \mathbf{w} \; - \; \frac{\lambda}{2} \mathbf{w}^\top \Sigma \, \mathbf{w} \; - \; \frac{\gamma}{2} \mathbf{w}^\top C_t \mathbf{w}, \tag{6.9}$$

with $\lambda > 0$ (variance aversion) and $\gamma > 0$ (connectedness aversion). The C_t term provides a continuous analogue of MCoP, discouraging allocations that co-load systemic spillovers.

Constraints

$$\mathbf{1}^\top \mathbf{w} = 1 \qquad \text{(budget)} \tag{6.10}$$

$$0 \le w_i \le \bar{w}_i \qquad \text{(box)} \tag{6.11}$$

$$\sum_{i \in \text{PoW}} w_i \le \bar{w}_{\text{PoW}} \qquad \text{(PoW cap)} \tag{6.12}$$

$$\sum_i w_i \, CF_i \le CF_{\max} \qquad \text{(carbon limit; ESRS E1)} \tag{6.13}$$

$$\sum_i w_i \, \delta_{\text{HW},i} \Big(1 - \rho_{\text{recycle}}(i)\Big) \le EW_{\max} \qquad \text{(e-waste cap; ESRS E5)} \tag{6.14}$$

$$\sum_i w_i \, FI_{\text{share},i} \ge \overline{FI} \qquad \text{(inclusion floor)} \tag{6.15}$$

$$w_i > 0 \Rightarrow (NC_i \ge NC_{\min} \;\wedge\; SC_i \le SC_{\max}) \qquad \text{(governance gates)} \tag{6.16}$$

Segmentation-aware trading costs When $S_t > S_{\text{thr}}$, augment (6.9) by

$$-\,\kappa(S_t) \cdot \text{Turnover}(\mathbf{w}),$$

with $\kappa(\cdot)$ increasing in segmentation severity.

Regime adaptivity

- If TCI_t is high, *raise* γ to penalise systemic linkages more strongly.
- Under inflationary, upper-quantile states, *tighten* $\bar{w}_{\mathrm{PoW}}$ and $CF_{\max}$.
- With elevated S_t, *inflate* trading-cost penalties and *limit* leverage.

Interpretation ESG is embedded at construction—not after the fact—and the connectedness penalty hard-wires crisis-resilient diversification. Regime-sensitive parameters operationalise findings in (Bibi et al., 2025; Lin et al., 2025; Makarov & Schoar, 2020).

Table 6.2: Optimisation Constraints and Parameters

Item	Definition	Source	Default	Stress Rule
Budget	$\mathbf{1}^\top \mathbf{w} = 1$	Portfolio	1	N/A
Box	$0 \le w_i \le \bar{w}_i$	Policy	Policy	Tighten if TCI high
PoW cap	$\sum_{i \in PoW} w_i \le \bar{w}_{PoW}$	Protocol map	5–10%	Lower in inflationary stress
Carbon	$\sum_i w_i CF_i \le CF_{\max}$	dMRV/ESRS	Policy	Tighten if inflationary stress
E-waste	$\sum_i w_i \delta_{HW,i}(1-\rho)$	Supplier/WEEE	Policy	N/A
Inclusion	$\sum_i w_i FI_{\mathrm{share},i} \ge \overline{FI}$	Analytics	10%	N/A
Gov gates	NC_i, SC_i thresholds	On-chain	$NC_{\min}, SC_{\max}$	N/A
λ, γ	Variance / connectedness	Estimation	Policy	$\gamma \uparrow$ if TCI high
Segmentation	$-\kappa(S_t)$Turnover	Market data	0	$\kappa \uparrow$ if S_t high

Note: Summary of binding constraints and regime-sensitive parameters used in the optimiser. Symbols: w_i = asset weight; $\bar{w}_i$ = asset cap; PoW = Proof of Work; CF_i = asset carbon factor; $CF_{\max}$ = portfolio carbon limit; $\delta_{\mathrm{HW},i}$ = hardware turnover rate; ρ = recycling rate; $FI_{\mathrm{share},i}$ = inclusion share for asset i; $\overline{FI}$ = portfolio inclusion floor; NC_i = Nakamoto coefficient (decentralisation); SC_i = stake/validator concentration; λ = variance aversion; γ = connectedness aversion; S_t = segmentation index (cross-venue dispersion); $\kappa(S_t)$ = turnover penalty increasing in S_t; TCI = Total Connectedness Index; dMRV = digital Monitoring, Reporting & Verification; ESRS = European Sustainability Reporting Standards; WEEE = Waste Electrical and Electronic Equipment directive. "Default" lists baseline policy values; "Stress Rule" describes automatic tightening in high-TCI or inflationary states. *Source: The author.*

6.6 Hedging Overlays and Systemic-Risk Controls

Empirical evidence shows hedge ratios and hedging effectiveness tend to increase during crises (Alharbi et al., 2025; Ali et al., 2024; Bibi et al., 2025). But higher effectiveness arrives with higher financing and execution costs; hence hedging must scale with systemic indicators.

6.6.1 Hedge-Ratio Calibration

For asset i and hedge j, the variance-minimising hedge ratio is

$$h_t \;=\; \frac{\mathrm{Cov}(r_{i,t}, r_{j,t})}{\mathrm{Var}(r_{j,t})}. \tag{6.17}$$

To accommodate tail-dependence shifts, estimate h_t using rolling *quantile* regressions, adjusting faster in stress (Ul Haq et al., 2023).

6.6.2 Minimum Coverage in High-TCI Regimes

When $\mathrm{TCI}_t \geq \tau_{\mathrm{TCI}}$, enforce

$$\text{Hedge Coverage} \;\geq\; \overline{HC}, \tag{6.18}$$

with coverage proxied by hedging effectiveness,

$$\mathrm{HE}_t \;=\; 1 - \frac{(R^{(H)}_{P,t})}{(R^{(U)}_{P,t})}, \tag{6.19}$$

calibrated to historical peak transmission (Bibi et al., 2025).

6.6.3 Instrument Selection and ESG Consistency

Hedge instruments must clear a minimum ESG score: preferred choices include green bonds and sustainability-linked bonds for macro hedges, ESG-screened equity futures, and derivatives on PoS/low-carbon crypto for sector-specific coverage.

6.6.4 Systemic Risk via CoVaR and Tail Dependence

Systemic exposure is tracked via CoVaR:

$$\mathrm{CoVaR}_q^{j|i} = \mathrm{VaR}_q(R_j \mid R_i = \mathrm{VaR}_q(R_i)), \quad \Delta\mathrm{CoVaR}_q^{j|i} = \mathrm{CoVaR}_q^{j|i} - \mathrm{VaR}_q(R_j \mid R_i = \mathrm{Median}(R_i)). \tag{6.20}$$

Large ΔCoVaR flags transmitters for immediate overlays or deweighting. Copula-based tail-dependence coefficients (λ_U, λ_L) complement CoVaR for regime recognition (Mzoughi et al., 2024).

6.6.5 Shock Taxonomy and Hedge Activation

Following Fry and Cheah (2016), distinguish:

- **Exogenous** shocks (policy bans, geopolitics): use broad ESG-compliant index hedges.
- **Endogenous** shocks (negative bubbles, microstructure failures): use asset-specific shorts or liquidity-matched pairs.

6.6.6 Carry and Cost Budgeting

Because costs rise in stress, set a regime-varying hedge budget:

$$\text{Hedge Budget}_t = \theta_t \cdot \text{NAV}_t, \qquad \theta_t \in [0.5\%, 2\%], \tag{6.21}$$

with monitoring of realised carry, funding, slippage, and variance reduction per cost.

Table 6.3: Hedging Policy Parameters

Parameter/Trigger	**Definition**	**Policy Action**
High TCI	$\text{TCI}_t \geq \tau_{\text{TCI}}$	Enforce HE $\geq \overline{HC}$; tighten caps.
ΔCoVaR	Above cut-off	Overlays/deweight high transmitters.
ESG filter	Min ESG score	Use green bonds, ESG futures, PoS derivatives.
Budget θ_t	$0.5\% \to 2\%$ by regime	Scale coverage; rank by HE per cost.
Quantile OHR	Rolling QR h_t	Faster updates in stress; slower in calm.

Note: Parameter definitions and actions for the hedging overlay. TCI = Total Connectedness Index; τ_{TCI} = policy threshold for TCI; HE = hedging effectiveness; $\overline{HC}$ = minimum hedge coverage target; ΔCoVaR = change in conditional Value-at-Risk (systemic downside when the conditioning asset is stressed); ESG = Environmental, Social, and Governance; "ESG futures" = futures on ESG-screened indices; PoS = Proof of Stake (low-energy consensus); θ_t = hedge budget as a share of NAV; "rank by HE per cost" = prioritise instruments delivering the largest variance reduction per unit of carry/slippage; OHR = optimal hedge ratio; QR h_t = hedge ratio estimated via rolling quantile regression (faster adaptation in stress regimes). "Transmitters" denotes net sources of spillovers in the connectedness network. *Source: The author.*

6.7 Monitoring, Breaches, and Rebalancing

Policy effectiveness hinges on real-time telemetry and disciplined breach management. The monitoring stack blends on-chain, off-chain, and market feeds, producing a compliance vector and a tiered response.

6.7.1 Telemetry Stack

On-chain Validator stake concentration (top-N share), governance participation and outcomes, protocol incidents (governance attacks, exploits, emergency forks).

Off-chain Jurisdictional carbon factors (location-/market-based), protocol e-waste/recycling metrics, PoR attestations with scope and liabilities treatment.

Market TCI/PCI matrices, CoVaR and tail-dependence estimates, segmentation index S_t, and macro-quantile indicators for PPI, FX, and yields (Lin et al., 2025).

6.7.2 Compliance Vector and Breach Detection

Define $\mathbf{g}_t(\mathbf{w})$ such that

$$g_{t,j}(\mathbf{w}) \leq 0 \quad \forall j, \tag{6.22}$$

with one component per constraint (carbon, e-waste, inclusion, PoW cap, governance, etc.). A breach occurs if any $g_{t,j}(\mathbf{w}) > 0$. Classification:

- **Soft:** Marginal, transient (e.g., $< 2\%$ over a limit), often due to drift.
- **Hard:** Material, persistent (e.g., $> 5\%$ over), or governance minimums violated.

6.7.3 Breach Response Protocol

Soft Watch-list and issuer/protocol engagement; shorten review cadence (e.g., monthly $\rightarrow$ weekly).

Hard Immediate de-risking (partial/full exit). If remediation fails within a grace window, remove from $\mathcal{C}(P)$.

6.7.4 Rebalancing Mechanics

Rebalancing is both scheduled (e.g., monthly/quarterly) and trigger-based (breaches; regime switches such as high TCI; macro tail events like PPI spikes in upper quantiles). Turnover penalties and jurisdiction-aware routing mitigate costs, especially under elevated segmentation (Makarov & Schoar, 2020).

6.7.5 Disclosure Mapping

Every breach and action is logged with timestamp, constraint, data sources, uncertainty, and rationale. Logs map to ESRS/ISSB/GRI items for audit readiness (Mulligan et al., 2024).

Table 6.4: Breach Classes and Actions

Class	**Definition**	**Actions**	**Disclosure**
Soft	$<$ 2% over threshold; transient	Watch-list; engage; tighter monitoring	Internal log; periodic ESG report
Hard	$>$ 5% and/or persistent; min G breach	De-risk; grace remediation; exclude if unresolved	Incident note; assurance pack; audit trail

Note: Breach taxonomy and the corresponding response protocol. "Soft" = minor, short-lived breach typically caused by drift; "Hard" = material and/or persistent breach, or violation of the governance minimum ("min G breach"). *Watch-list* = flag the asset for heightened review; *engage* = contact issuer/protocol or service provider with remediation requests; *de-risk* = reduce or exit exposure; *grace remediation* = time-bounded window to fix issues before exclusion. *Assurance pack* = evidence bundle (data sources, calculations, and third-party attestations) supporting disclosure; *audit trail* = timestamped record of breach, analysis, and actions taken. ESG = Environmental, Social, and Governance. *Source: The author.*

6.8 Illustrative Parameterisation (Committee Menu)

Thresholds should be literature-grounded, regulator-consistent, and reviewed annually or upon structural change.

6.8.1 ESG Eligibility Thresholds

Pillar minimums (0–100 scale)

$$E_{\min} = 60 \quad \text{(favouring low-energy protocols (Wendl et al., 2023))}, \qquad S_{\min} = 50 \quad \text{(baseline inclusio}$$

Scores are composite, with documented sources and uncertainty bands.

6.8.2 Environmental Constraints

PoW cap

$$\sum_{i \in \text{PoW}} w_i \le 5\text{–}10\%, \quad \text{tighten to 5\% in inflationary or high-TCI regimes (Bibi et al., 2025; Lin et al.,}$$

Carbon limit

$$\sum_i w_i \, CF_i \leq CF_{\max},$$

aligned to transition plans and disclosed in LB/MB forms.

E-waste cap

$$\sum_i w_i \, \delta_{\mathrm{HW},i} \Big(1 - \rho_{\mathrm{recycle}}(i)\Big) \leq EW_{\max},$$

consistent with ESRS E5.

6.8.3 Social Constraints

Inclusion floor

$$\sum_i w_i \, FI\_\mathrm{share}^{(i)} \geq 10\%.$$

Consumer-protection outliers (e.g., fraud/complaints $> 1.5\times$ median) require remediation to remain eligible.

6.8.4 Governance Constraints

Decentralisation

$$NC_i \geq 7, \qquad SC_i \leq 33\%.$$

Proof-of-reserves at least quarterly, with liabilities coverage disclosed.

6.8.5 Optimisation Weights

Risk aversion λ calibrated by stress tests; γ (connectedness aversion) raised in high-TCI_t to approximate MCoP discipline (Bibi et al., 2025).

6.8.6 Summary Table

Table 6.5: Constraint Menu and Default Thresholds

Type	Metric	Threshold	Trigger	Action	Reference
PoW cap	$\sum_{i \in PoW} w_i$	5–10%	High TCI / high PPI	Reduce PoW	Wendl et al. (2023)
Carbon	$\sum_i w_i CF_i$	$CF_{\max}$	Exceedance	De-risk	Mulligan et al. (2024)
E-waste	$\sum_i w_i \delta_{HW,i}(1-\rho)$	$EW_{\max}$	Exceedance	Engage / limit	Mulligan et al. (2024)
Inclusion	$\sum_i w_i FI\_\text{share}^{(i)}$	$\geq 10\%$	Shortfall	Rebalance	Mulligan et al. (2024)
Governance	NC_i, SC_i	≥ 7, $\leq 33\%$	Deterioration	Remove asset	Mulligan et al. (2024)

Note: Definitions and acronyms—PoW: Proof of Work; w_i: portfolio weight of asset i; TCI: Total Connectedness Index; PPI: Producer Price Index. *Carbon*: CF_i is the asset-level carbon factor (e.g., tCO_2/yr), $CF_{\max}$ is the portfolio carbon cap. *E-waste*: $\delta_{HW,i}$ is the hardware turnover rate; ρ is the recycling rate; $EW_{\max}$ is the portfolio e-waste cap. *Inclusion*: $FI\_\text{share}^{(i)}$ is the inclusion-share metric for asset i (e.g., proportion of transactions/use consistent with financial inclusion). *Governance*: NC_i is the Nakamoto coefficient (decentralisation); SC_i is stake/validator concentration (e.g., top cohort share). *Actions*: "Reduce PoW" lowers aggregate PoW exposure; "De-risk" trims or exits positions exceeding limits; "Engage / limit" combines issuer/protocol engagement with temporary caps; "Rebalance" reallocates to restore compliance; "Remove asset" excludes on persistent governance deterioration. The *Reference* column lists indicative sources used to motivate thresholds; align with your bibliography. *Source: The author.*

Chapter 7

Conclusion and Policy Outlook

7.1 Synthesis of Findings

This thesis set out to examine the economic and sustainability interdependencies between ESG-aligned assets and digital-asset markets and to translate those insights into a rules-based portfolio policy that is measurable, enforceable, and aligned with leading disclosure frameworks. Across empirical evidence, methodological architectures, and policy design, four core conclusions emerge.

(i) Interconnectedness is high, state-dependent, and asymmetric. Time-varying parameter connectedness analyses reveal structurally elevated integration between ESG equities (and related factors) and cryptocurrencies, with sharp intensifications in systemic regimes. During COVID-19 onset and the Russia–Ukraine escalation the Total Connectedness Index (TCI) reached the 80–85% range, implying that in such windows the majority of forecast error variance is attributable to cross-series spillovers. Decomposition by return sign shows persistent dominance of downside transmission: negative shocks propagate more strongly than positive shocks. Directionality analyses attribute transmitter roles to large-cap crypto (e.g., BTC, ETH, LTC) and receiver roles to "green" cryptocurrencies (ADA, MIOTA, XLM, XNO, XRP), particularly when ESG components and G7 equities are themselves under stress. These findings rationalise crisis-specific policy switches, including concentration caps on transmitters and protective tilts toward low-connectedness diversifiers.

(ii) Diversification is horizon- and regime-contingent. Time–frequency (wavelet) methods show that integration is scale dependent. Green bonds and ESG equity indices co-move strongly at lower frequencies, compressing long-horizon diversification. By contrast, long-run anti-phase behaviour between BTC and broad ESG equity benchmarks can deliver structural diversification in strategic sleeves. The portfolio takeaway is that hedging and allocation rules should be horizon-aware: instruments that diversify tactically at short horizons may not diversify strategically, and vice versa.

(iii) Systemic risk concentrates in equity-like sustainability exposures; bonds diversify. Tail-dependence and ΔCoVaR diagnostics indicate that clean-energy equities and blockchain-related equities exhibit pronounced crisis co-crash patterns, while green bonds retain diversification benefits under stress. Consequently, sustainability alignment at the theme level does not guarantee risk mitigation: pro-cyclical clean-energy equities can amplify blockchain-equity drawdowns, whereas green fixed income tends to cushion them. This asymmetry informs the hedging palette and the selection of "green" diversifiers in crisis overlays.

(iv) Environmental externalities are protocol- and geography-specific; governance and social factors are first-order risk controls. Proof-of-Work (PoW) systems entail materially higher operational energy use, e-waste, and relocation-sensitive carbon intensity than Proof-of-Stake (PoS) and related low-energy consensus mechanisms. Grid-mix heterogeneity and miner migration generate carbon-leakage risks that defeat simple network-average accounting. On the social and governance pillars, financial inclusion benefits exist alongside non-trivial consumer-protection risks, while decentralisation, MEV governance, and custody integrity are central to operational resilience. The thesis therefore embeds environment, social, and governance metrics not as narrative addenda but as *binding constraints* in allocation.

7.2 Methodological Contributions

The thesis contributes a unified decision architecture that fuses econometric evidence with ESG policy design:

1. **Connectedness-aware optimisation.** The objective function augments mean–variance with a connectedness penalty $\mathbf{w}^\top C_t \mathbf{w}$, explicitly discouraging portfolios that co-load systemic spillover channels. This continuous analogue of the Minimum Connectedness Portfolio (MCoP) internalises systemic fragility rather than relying solely on pairwise covariances.

2. **ESG as pre-trade feasibility.** ESG criteria are encoded as eligibility screens and hard constraints: PoW exposure caps; carbon and e-waste budgets

aligned to ESRS E1/E5; inclusion floors; decentralisation gates via a minimum Nakamoto coefficient and stake concentration limits; and custody governance requirements (e.g., proof-of-reserves cadence and scope). By imposing these at construction time, the optimiser's feasible set is ESG-compliant by design.

3. **Regime-adaptive overlays.** Empirical stylised facts are operationalised as triggers: (i) TCI thresholds initiate transmitter caps, diversifier reweights, and hedging-mode shifts (MCoP/MCP); (ii) macro-quantile overlays translate inflation and rate shocks into time-varying caps (e.g., tighter PoW bounds in high-PPI right-tail states); (iii) segmentation guardrails curb leverage and re-route execution when cross-venue dispersion widens.

4. **Digital MRV for auditability.** A data integrity layer specifies on-chain attestations, oracle commits, and versioned metric objects (functional unit, system boundary, observation window, uncertainty bands) to address measurement dispersion and to map directly into ESRS/ISSB/GRI disclosures. This closes the loop between portfolio controls and external reporting.

Together, these elements create a *rules-based* framework that is both implementable in production and traceable for assurance.

7.3 Implications for Investors and Policymakers

For asset owners and managers. The results counsel against viewing "ESG" as a monolithic hedge to "crypto." Integration is uneven and asymmetric, with co-movements intensifying where diversification is needed most. A connectedness-aware optimiser combined with ex-ante ESG feasibility constraints yields portfolios that are more resilient to contagion and easier to certify against evolving disclosure regimes. In practice this means: prefer PoS exposure for core allocations; cap PoW tactically with relocation-aware carbon accounting; use green bonds as first-line diversifiers in crises; elevate hedging coverage when TCI signals cross regime thresholds; and recognise that governance (decentralisation, MEV controls, custody integrity) is not cosmetic—it changes tail risk.

For regulators and standard setters. Two policy gaps emerge. First, *methodological dispersion* in energy and emissions accounting invites selective reporting; the chapter on digital MRV proposes a credible template that couples on-chain audit trails with third-party assurance. Adoption of versioned, machine-readable disclosures would improve comparability and reduce compliance friction. Second, cross-border segmentation creates execution risk that can undermine investor protection; harmonised guidance on proof-of-reserves scope, key management standards, and fair-ordering/relay decentralisation would lower systemic vulnerabilities without prescribing specific technologies.

For sustainability programme leads. Portfolio alignment with climate and inclusion targets becomes tractable when ESG metrics are instrumented and enforced at the optimiser level. The proposed policy converts high-level commitments (e.g., carbon budgets, circularity targets, inclusion objectives) into hard constraints with monitoring telemetry and breach protocols. This supports audit readiness under ESRS/CSRD and ISSB while preserving investment discipline.

7.4 Limitations

The analysis is subject to several constraints which should temper inference and guide future work.

- **Sample and regime dependence.** Connectedness, tail dependence, and macro betas are estimated over windows that include discrete structural breaks (e.g., Ethereum's Merge, regulatory shocks, exchange failures). Results are informative about *observed* regimes; their extrapolation requires caution.
- **Measurement uncertainty.** Environmental metrics depend on system boundaries (operational vs. life-cycle), functional units (Wh/tx vs. MWh/yr), and geography assignments for miners/validators. Ranges and uncertainty bands mitigate, but do not eliminate, model risk.
- **Identification challenges.** Directionality and causality are inferred from reduced-form models (TVP-VAR, copulas, quantile processes). Although robust to ordering and distributional assumptions, they remain subject to omitted-variable bias and latent common factors (e.g., global dollar liquidity).
- **Market-microstructure opacity.** Execution frictions, inventory constraints, and off-exchange flows (e.g., internalisation, OTC) can blunt the mapping from policy rules to realised hedging outcomes, particularly under segmentation or collateral stress.
- **Governance observability.** Validator identities, cartelisation risks, and MEV extraction are imperfectly observable; proxy metrics (e.g., $N^\star$, HHI) may understate hidden concentration or control.

These limitations motivate the design choice to encode *ranges*, *state-contingent parameters*, and *explicit breach protocols* rather than relying on single-point estimates or static policy settings.

7.5 Directions for Future Research

Several extensions would deepen the evidence base and enhance the policy apparatus:

1. **Causal identification of spillovers.** Combine high-frequency instruments (funding rates, order-book imbalance, stablecoin premia) with structural sign/zero restrictions to isolate shock types (liquidity vs. information vs. policy) and to test the robustness of TX/RX roles.

2. **Cross-layer and cross-chain dynamics.** Extend connectedness matrices to include L2s, app-chains, and restaking/shared-security constructs; measure how sequencing markets and order-flow auctions (PBS, MEV-boost) alter governance risk and tail dependence.

3. **Granular environmental telemetry.** Standardise validator geolocation (privacy-preserving), metered energy inputs, and REC serialisation to shrink dispersion in E metrics; develop mixed-method LCAs to quantify embodied emissions and e-waste with survival/hazard models.

4. **Inclusion and protection outcomes.** Link wallet cohorts to use-case typologies (remittances, small-business payments) and construct controlled studies of complaint/loss severity with intervention tests (e.g., escrow defaults, insurance layers).

5. **Execution under segmentation.** Build agent-based simulations calibrated to venue-specific fees, latencies, and capital controls to stress-test hedging and rebalancing under widening cross-venue bases; quantify the cost of policy-imposed venue whitelists.

6. **Regulatory shocks and scenario design.** Integrate macro-quantile overlays with formal scenario sets (rate/inflation surprises, energy-price spikes, sanction regimes) and back-test policy reactions (PoW caps, carbon-limit tightening, hedge budgets).

7. **Fairness and MEV externalities.** Empirically relate MEV mitigation (PBS adoption, builder concentration) to user-welfare metrics (failed transactions,

inclusion latency) and systemic indicators (CoVaR, λ_U) to quantify G- and S-pillar benefits.

These strands would improve out-of-sample stability and make the policy toolbox more predictive rather than reactive.

7.6 Closing Remarks: From Narrative ESG to Rules-Based Stewardship

The central claim of this thesis is that sustainability and risk–return are not separable objectives in long-horizon allocation to digital assets. Market integration is too high, crisis asymmetries too pronounced, and environmental and governance externalities too material to treat ESG as a marketing narrative or an after-the-fact filter. A credible, investable approach must instead be *rules-based* and *state-aware*: connectedness belongs in the objective; ESG belongs in the feasible set; telemetry and breach protocols belong in operations; and disclosure artefacts must be machine-verifiable.

Practically, the policy advanced here operationalises that philosophy. It starts by restricting the investable universe to assets demonstrating acceptable environmental footprints (preferencing PoS and region-specific carbon accounting), tangible social outcomes (inclusion floors and consumer-protection surveillance), and resilient governance (decentralisation thresholds, MEV controls, custody integrity). It then embeds systemic-awareness by penalising connectedness in optimisation and by activating overlays when TCI, macro quantiles, and segmentation indices cross thresholds. Finally, it closes the loop with digital MRV and standards alignment (ESRS/ISSB/-GRI), ensuring that what is measured is what is managed and what is disclosed is what is assured.

The policy is intentionally modular. As protocols evolve (e.g., energy profiles, governance mechanisms), as market structure shifts (e.g., sequencing markets, cross-border rails), and as regulation matures (e.g., standardised PoR, privacy requirements), the same scaffolding can ingest new metrics, tighten or relax thresholds, and tune regime triggers. In that sense, the contribution is not merely a snapshot of present conditions; it is a template for *evolving stewardship*—capable of absorbing new evidence without abandoning discipline.

In closing, sustainable allocation to digital assets is neither a contradiction nor a foregone conclusion. It is a *governance problem* framed by data credibility, a *risk problem* conditioned by regime dynamics, and an *engineering problem* solvable through transparent standards and programmable incentives. By elevating ESG from narrative to constraint and by elevating systemic awareness from intuition to objective, this thesis proposes a path by which crypto exposure can be held to the same scientific and fiduciary standards that govern the rest of institutional capital.

Appendices

Supplement to Introduction

Scope Delimitations and Additional Context

This appendix complements Introduction by detailing inclusion/exclusion rules for assets, time windows, and policy frameworks referenced in the thesis. Table 1 summarises scope choices mapped to the research questions.

Table 1: Scope matrix aligning assets, periods, and frameworks to research questions.

Question	**Asset set**	**Baseline horizon**	**Policy/Reporting frame**
Connectedness dynamics	ESG leaders; E/S/G sub-indices; BTC/ETH; "green" cryptos (ADA, MIOTA, XLM, XRP, XNO)	2017–2024	Economic/financial lens (no policy filter)
Environmental footprints	PoW vs. PoS/FBA networks	Rolling annual	ESRS E1; IFRS S2; GRI 302/305
Market microstructure	Top fiat venues; segmentation pairs (e.g., US/KR)	Event windows	Execution governance (venue/rail)
Diversifiers	Green bonds; clean-energy equities; blockchain equities	Crisis vs. calm	Risk policy; ESRS alignment

Notes on Key Assumptions

Energy accounting boundaries (operational vs. LCA) follow (Mulligan et al., 2024; Wendl et al., 2023). Connectedness interpretation follows (Antonakakis et al., 2020; Chatziantoniou & Gabauer, 2021).

Supplement to Literature Review

Study Inventory and Coding Schema

Table 2: Inventory of reviewed empirical studies with coding fields.

Study	Sample	Method	Regime windows	Notes
Bibi et al. (2025)	ESG leaders vs BTC/ETH/ALT	TVP–VAR; FEVD; asymmetry	2019; 2020Q1; 2022Q1	TCI 55–85%, negat dominance
Alharbi et al. (2025)	ESG E/S/G vs green cryptos	Spillovers; HE	2018–2023	ESG → green crypt HE↑ in crises
Ali et al. (2024)	G7 vs green cryptos	Spillovers; hedge ratios	2018–2023	Crisis peaks; ratios↑
Ul Haq et al. (2023)	ESG/GB vs crypto	Wavelet coherence	2010s–2022	Horizon dependence
Mzoughi et al. (2024)	Clean energy vs blockchain eq.	Copula; CoVaR	2018–2023	Tail co-crash; GB div sify
Makarov and Schoar (2020)	Cross-venue BTC	Microstructure	2017–2018	Segmentation; "Kim premium"
Alzoubi and Mishra (2023) and Wendl et al. (2023)	Protocol energy/env.	Reviews; LCA	Various	PoS $\ll$ PoW energy

Robustness Conventions

TCI scaling follows (Chatziantoniou & Gabauer, 2021). Event windows align to COVID-19 onset and RU–UA escalation.

Supplement to Methods Used in the Literature

TVP–VAR Estimation Details

Kalman Filtering Settings

Discount factors $\delta_\beta \in [0.97, 0.995]$; covariance evolution $\delta_\Sigma \in [0.97, 0.995]$. Forecast horizon $H \in \{10, 20\}$.

Generalised FEVD and Connectedness

Let $A_{h,t}$ be VMA coefficients; the H-step GFEVD share is

$$\Psi^g_{ij,t}(H) = \frac{\sigma^{-1}_{jj,t} \sum_{h=0}^{H-1} (e_i' A_{h,t} \Sigma_t e_j)^2}{\sum_{h=0}^{H-1} (e_i' A_{h,t} \Sigma_t A'_{h,t} e_i)}.$$

Row-normalised shares feed TCI, directional and pairwise indices.

Pseudocode: Connectedness Pipeline

Listing 1: Pseudocode for the connectedness pipeline in Chapter 3.

```
# Inputs: returns matrix Z (T x N), horizon H, discount
   factors delta_beta, delta_sigma
# Output: TCI_t, TO/FROM/NET_t, PCI_t
estimate_tvpvar(Z, p=1, delta_beta, delta_sigma)
for t in 1..T:
    compute_VMA_coeffs(A[h,t] for h=0..H-1)
    compute_GFEVD_shares_Psi(t, H)
    row_normalize_shares()
    compute_TCI_TO_FROM_NET_and_PCI(t)
```

```
return indices
```

Hedge Ratios and Hedging Effectiveness

Minimum-variance hedge ratio $h_t = \text{Cov}(r_i, r_j)/\text{Var}(r_j)$; $\text{HE} = 1 - (R_{P,h})/(R_{P,0})$ (Ederington, 1979).

Supplement to Market Interconnectedness: ESG Equities, Green Assets, and Crypto

Additional Connectedness Tables

Table 3: Illustrative pairwise connectedness (% FEVD share, $H = 20$).

	ESG	BTC	ETH	ADA	XLM	XRP
ESG	–	6.1	5.8	3.2	2.9	3.0
BTC	1.4	–	13.0	7.5	7.0	6.6
ETH	1.6	11.9	–	8.1	7.8	8.0
ADA	1.0	8.4	8.6	–	7.3	7.6
XLM	0.9	7.7	7.9	7.1	–	6.9
XRP	1.1	8.0	8.2	7.4	7.0	–

Wavelet and Tail-Risk Notes

Time–frequency coherence highlights short-run contagion in crises; lower-tail dependence (λ_L) and ΔCoVaR rise between clean-energy and blockchain equities (Mzoughi et al., 2024; Ul Haq et al., 2023).

Supplement to ESG Pillars for Digital Assets: Measurement, Governance, and Integration

Environmental Metric Template

Table 4: Template (report ranges and assumptions).

Metric	**Unit**	**Boundary/Notes**
Network energy E_{op}	MWh/yr	Operational, rolling 12M; method per Wendl et al. (2023).
Carbon (LB/MB)	tCO_2/yr	Location-/market-based; regional grid factors.
E-waste	t/yr	Vintage model; recycler attestations; ESRS E5.
Renewable share	%	RECs with serials; verification cadence.

dMRV: On-Chain Data Object

Record ⟨metric_id, value, unit, lower, upper, fu, sb, window, oracle_tx, audit_hash⟩ (see §5.5).

Governance Checklist

Decentralisation (Nakamoto $N^{\star}$, HHI, top-k share); MEV controls (PBS, builder HHI); Custody (PoR coverage/cadence; key posture).

Supplement to ESG-Integrated Portfolio Policy for Digital Assets

Optimisation Problem (Full Form)

Objective:

$$\max_{\mathbf{w}} \ \boldsymbol{\mu}^\top \mathbf{w} - \frac{\lambda}{2}\mathbf{w}^\top \Sigma \mathbf{w} - \frac{\gamma}{2}\mathbf{w}^\top C_t \mathbf{w}$$

subject to constraints (6.10)–(6.16).

Pseudocode: Regime-Aware ESG Rebalancing

Listing 2: Pseudocode for regime-aware ESG-constrained rebalancing.

```
# Inputs: current weights w, (mu, Sigma), connectedness C_t,
# ESG metrics (CF_i, FI_share, NC_i, ...), regime flags (TCI,
    PPI, S)
# Output: target weights w*
build_feasible_set_by_ESG_screens()
tighten_limits_if_regime(TCI, PPI, S)
solve_QP(mu, Sigma, C_t, constraints)   # objective with
   lambda, gamma
apply_turnover_and_cost_penalties()
return w_star
```

Breach Handling Playbook

Soft vs. hard breaches per Table 6.4; log incidents with data sources and uncertainty ranges for ESRS/ISSB.

Abbreviations and Notation

Abbreviations

Acronym	**Meaning**
ESG	Environmental, Social, Governance
PoW / PoS / FBA	Proof-of-Work / Proof-of-Stake / Federated Byzantine Agreement
MRV / dMRV	Monitoring, Reporting, Verification / digital MRV
ESRS / ISSB / GRI	European Sustainability Reporting Standards / International Sustainability Standards Board / Global Reporting Initiative
TVP–VAR	Time-Varying Parameter Vector Autoregression
FEVD / TCI / PCI	Forecast Error Variance Decomposition / Total Connectedness Index / Pairwise Connectedness Index
CoVaR	Conditional Value-at-Risk (systemic risk)
HE	Hedging Effectiveness
REC	Renewable Energy Certificate
HHI	Herfindahl–Hirschman Index

Notation (selected)

Symbol	Definition
$z_t,\ \mathbf{y}_t$	Return vectors.
$A_{h,t}$	VMA coefficients at horizon h.
Σ_t	Innovation covariance matrix.
$\Psi^g_{ij,t}(H)$	H-step GFEVD share from j to i.
TCI_t	Total Connectedness Index at time t.
h_t	Minimum-variance hedge ratio.
CF_i	Carbon factor for asset i (tCO_2/yr).

$N^{\star}$	Nakamoto coefficient (entities to control >50%).

References

Alharbi, S. S., Naveed, M., Ali, S., & Moussa, F. (2025). Sailing towards sustainability: connectedness between ESG stocks and green cryptocurrencies. *International Review of Economics & Finance*, *98*, 103848. https://doi.org/10.1016/j.iref.2025.103848

Ali, S., Naveed, M., Yousaf, I., & Khattak, M. S. (2024). From cryptos to consciousness: dynamics of return and volatility spillover between green cryptocurrencies and G7 markets. *Finance Research Letters*, *60*, 104899. https://doi.org/10.1016/j.frl.2023.104899

Alzoubi, Y. I., & Mishra, A. (2023). Green blockchain – a move towards sustainability. *Journal of Cleaner Production*, *430*, 139541. https://doi.org/10.1016/j.jclepro.2023.139541

Anscombe, F. J., & Glynn, W. J. (1983). Distribution of the kurtosis statistic b_2 for normal samples. *Biometrika*, *70*(1), 227–234. https://doi.org/10.1093/biomet/70.1.227

Antonakakis, N., Chatziantoniou, I., & Gabauer, D. (2020). Refined measures of dynamic connectedness based on TVP-VAR. *Energy Economics*, *90*, 104847. https://doi.org/10.1016/j.eneco.2020.104847

Bibi, R., Gulzar, S., & Shahzad, S. J. H. (2025). ESG leaders and cryptocurrency market: asymmetric TVP-VAR connectedness and investment approaches. *Research in International Business and Finance*, *76*, 102833. https://doi.org/10.1016/j.ribaf.2025.102833

Chatziantoniou, I., & Gabauer, D. (2021). TCI is not bounded between 0 and 1: a Monte Carlo investigation and adjustment. *Finance Research Letters*, *43*, 101982. https://doi.org/10.1016/j.frl.2021.101982

D'Agostino, R. B. (1970). Transformation to normality of the sample distribution of D. *Biometrika*, *57*(3), 679–681. https://doi.org/10.1093/biomet/57.3.679

Ederington, L. H. (1979). The hedging performance of the new futures markets. *The Journal of Finance*, *34*(1), 157–170. https://doi.org/10.1111/j.1540-6261.1979.tb02077.x

Elliott, G., Rothenberg, T. J., & Stock, J. H. (1996). Efficient tests for an autoregressive unit root. *Econometrica*, *64*(4), 813–836. https://doi.org/10.2307/2171846

Figueiredo, K., Hammad, A. W. A., Haddad, A., & Tam, V. W. Y. (2022). Assessing the usability of blockchain for sustainability: extending key themes to the construction industry. *Journal of Cleaner Production*, *343*, 131047. https://doi.org/10.1016/j.jclepro.2022.131047

Fry, J., & Cheah, E.-T. (2016). Negative bubbles and shocks in cryptocurrency markets. *International Review of Financial Analysis*, *47*, 343–352. https://doi.org/10.1016/j.irfa.2016.02.008

Gabauer, D. (2021). Dynamic connectedness of cryptocurrencies: a time-varying parameter VAR approach. *Finance Research Letters*, *38*, 101680. https://doi.org/10.1016/j.frl.2020.101680

Hannan, E. J., & Quinn, B. G. (1979). The determination of the order of an autoregression. *Journal of the Royal Statistical Society: Series B (Methodological)*, *41*(2), 190–195.

Jarque, C. M., & Bera, A. K. (1980). Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics Letters*, *6*(3), 255–259. https://doi.org/10.1016/0165-1765(80)90024-5

Koop, G., Pesaran, M. H., & Potter, S. M. (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, *74*(1), 119–147. https://doi.org/10.1016/0304-4076(95)01753-4

Lin, M., Liu, Y., & Ng Kim Sheng, V. (2025). Analysis of the impact of macroeconomic factors on cryptocurrency returns: a quantile regression study. *International Review of Economics & Finance*, *97*, 103757. https://doi.org/10.1016/j.iref.2024.103757

Makarov, I., & Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, *135*(2), 293–319. https://doi.org/10.1016/j.jfineco.2019.07.001

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, *7*(1), 77–91. https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

Markowitz, H. (1959). *Portfolio selection: efficient diversification of investments*. Wiley.

Mulligan, C., Morsfield, S., & Cheikosman, E. (2024). Blockchain for sustainability: a systematic literature review for policy impact. *Telecommunications Policy*, *48*, 102676. https://doi.org/10.1016/j.telpol.2023.102676

Mzoughi, H., Ben Amar, A., Guesmi, K., & Benkraiem, R. (2024). Blockchain markets, green finance investments, and environmental impacts. *Research in*

International Business and Finance, *69*, 102249. https://doi.org/10.1016/j.ribaf.2024.102249

Pesaran, M. H., & Shin, Y. (1998). Generalized impulse response analysis in linear multivariate models. *Economics Letters*, *58*(1), 17–29. https://doi.org/10.1016/S0165-1765(97)00214-0

Popkova, E. G., Bogoviz, A. V., Lobova, S. V., Vovchenko, N. G., & Sergi, B. S. (2023). Blockchain, sustainability and clean energy transition. *Global Transitions*, *5*, 64–78. https://doi.org/10.1016/j.glt.2023.04.002

Sharpe, W. F. (1994). The sharpe ratio. *The Journal of Portfolio Management*, *21*(1), 49–58. https://doi.org/10.3905/jpm.1994.409501

Ul Haq, I., Maneengam, A., Chupradit, S., & Huo, C. (2023). Are green bonds and sustainable cryptocurrencies truly sustainable? evidence from a wavelet coherence analysis. *Economic Research–Ekonomska Istraživanja*, *36*(1), 807–826. https://doi.org/10.1080/1331677X.2022.2080739

Wendl, M., Doan, M. H., & Sassen, R. (2023). The environmental impact of cryptocurrencies using proof of work and proof of stake consensus algorithms: a systematic review. *Journal of Environmental Management*, *326*, 116530. https://doi.org/10.1016/j.jenvman.2022.116530