

DePIN Tokenomics Under Stress

A Comparative Stress Evaluation Using the Onocoy Network as
an Anchor Case

Devinson Peña
Florian Baumann
Cristian Carvalho

CAS Blockchain
Lucerne University of Applied Sciences and Arts (HSLU)

Thesis period: Autumn Semester 2025 / Spring Semester 2026

Program management: Prof. Dr. Georges Grivas and Jan Carlos Janke

Supervisor: Alexander Denzler

Classification: Public

Submission date: March 8, 2026

Declaration of Originality

We hereby declare that we have completed this thesis independently and have used only the sources and aids indicated. All passages taken directly or indirectly from other works have been identified as quotations or paraphrases and referenced accordingly. We further confirm that this thesis does not contain unauthorized third-party material, that confidential information has been handled in accordance with the applicable submission requirements, and that applicable copyright and intellectual-property rules have been respected throughout.

Place, date:

Signatures:

Use of Generative AI

Generative AI tools were used during the drafting, editing, and review process of this thesis, including support for structural planning, prose revision, consistency checks, and code- and LaTeX-oriented implementation assistance. These tools were not treated as primary sources or as substitutes for academic judgment. All argument structure, source selection, methodological choices, interpretation, and final verification remain the responsibility of the authors. Every substantive claim included in the thesis was reviewed against the cited literature, protocol documentation, empirical material, or explicit modeling assumptions before inclusion in the final manuscript.

Abstract

We examine how DePIN tokenomic mechanisms behave under physical and economic stress, and how their robustness can be evaluated before failure becomes operationally visible. The thesis focuses on the contrast between Burn-and-Mint-Equilibrium-oriented and capped-supply designs, while using the Onocoy network as the primary anchor case for a bounded comparative analysis. Our argument begins from a cyber-physical premise: unlike software-only crypto systems, DePINs rely on hardware deployment, maintenance effort, and site-specific operating constraints, so stress propagates not only through token prices but also through provider retention, service continuity, and incentive–usage alignment.

We combine three layers of analysis. First, we develop a comparative tokenomic framework that distinguishes major mechanism classes and the stress channels most relevant to DePIN systems. Second, we examine historical stress windows to identify recurring empirical signatures and to clarify the limits of what can be observed directly for an early-stage network such as Onocoy. Third, we introduce DTSE, the DePIN Tokenomic Stress Evaluator, as a bounded rule-based simulation framework that applies matched stress inputs across fixed mechanism profiles under explicit assumptions.

The results show that DePIN robustness is better understood through stress-transmission paths than through binary stable/unstable labels. In the accepted DTSE experiment set, different stress channels register first in different subsystems: demand contraction weakens usage-linked fields before participation visibly falls, liquidity shock compresses the market-facing layer before provider economics and churn deteriorate, competitive-yield pressure shifts participation before service utilization becomes the first-moving field, and provider-cost inflation compresses margins before visible participation loss follows. Across those cases, top-line participation metrics can lag underlying economic deterioration, indicating that DePIN networks may enter phases of latent structural stress before participation visibly declines.

We therefore argue for a narrower but stronger evaluative stance. The thesis does not claim to forecast live-network outcomes. Instead, it contributes a comparative framework, an empirical observability boundary, and a bounded simulation-based evaluator for diagnosing how DePIN tokenomic designs absorb, transmit, and expose stress under explicit assumptions.

Contents

Use of Generative AI	i
Abstract	ii
List of Figures	vi
List of Tables	vi
1 Introduction	1
1.1 Why This Problem Matters	1
1.2 Problem Setting	1
1.3 Research Question and Contribution	2
1.4 Scope and Thesis Roadmap	3
2 Foundations: DePIN as a Cyber-Physical System	3
2.1 DePIN as a Cyber-Physical Coordination Problem	3
2.2 Physical Constraints and Economic Stickiness	4
2.3 Path Dependence and Recovery Asymmetry	5
2.4 Theoretical Stress Constructs	6
3 Comparative Tokenomic Framework	7
3.1 Comparator Frame and Time-Stamped Market Snapshot	7
3.2 Emission Logic and Supply Regimes	9
3.3 Reward Logic, Work Verification, and Quality Adjustment	11
3.4 Monetization, Sinks, and Value Accrual	13
3.5 Demand Regimes and Utility–Speculation Profiles	13
3.6 Framework Synthesis	14
4 Onocoy as the Anchor Case	15
4.1 System Function and Network Participants	15
4.2 Token and Payment Architecture	16
4.3 Settlement-Layer Scope and Decision Boundary	18

4.4	Case Relevance, Documentation Gaps, and Evidence Boundary	19
5	Empirical Stress Layer	21
5.1	Empirical Analysis Orientation	21
5.2	Stress Windows and Comparator Scope	22
5.3	Metric Applicability and Observability Limits	23
5.4	Observed Stress Signatures in Historical Windows	25
5.5	Empirical Boundary and Need for DTSE	27
6	DTSE Methodology	28
6.1	DTSE Purpose, Scope, and Interpretation Boundary	29
6.2	Architecture, Evidence Layers, and Core Assumptions	30
6.3	Scenario Grid and Experimental Setup	33
6.4	Metric Selection, Metric Families, and Stage 1 Detection Logic	35
6.5	Stage 2 Failure-Signature Logic	38
6.6	Reproducibility, Calibration Boundary, and Method Limits	39
7	Simulation Results	40
7.1	Reporting Conventions and Interpretation Boundaries	40
7.2	Baseline Trajectories (The Neutral Reference)	41
7.3	Cross-Scenario Detection Pattern Summary	42
7.4	Onocoy (ONO) Baseline-Relative Signal Sequencing	42
7.5	Scenario-by-Scenario Deviations	43
7.6	Cross-Profile and Cross-Scenario Synthesis	47
8	Discussion and Conclusion	50
8.1	Discussion Orientation	50
8.2	What the Results Mean for DePIN Robustness	50
8.3	Comparative Implications for DePIN Design	52
8.4	Governance and Intervention Archetypes	53
8.5	Limitations and Claim Boundaries	54
8.6	Answer to the Research Question and Contributions	55

8.7	Future Research Agenda	56
8.8	Closing Statement	57
A	Appendix	58
A.1	DTSE Scenario-Formula Map	58
A.2	DTSE Threshold Definitions	59
A.3	Override Ledger	59
A.4	DTSE Run Manifest	60
A.5	Semi-Structured Interview Context Summary	62

List of Figures

- 1 DePIN cyber-physical coordination loop 4
- 2 Onocoy tokenomics flow 17
- 3 Empirical narrowing funnel 23
- 4 DTSE evaluation pipeline 30
- 5 DTSE ONO baseline reference trajectories 42
- 6 DTSE ONO demand-contraction deviations from baseline 44
- 7 DTSE ONO liquidity-shock transmission 45
- 8 DTSE ONO competitive-yield pressure response 46
- 9 DTSE ONO provider-cost inflation response 46
- 10 DTSE ONO cross-scenario synthesis 48
- 11 DTSE cross-profile final-week comparison by scenario 49

List of Tables

- 1 Core stress constructs 6
- 2 Comparator snapshot (18 February 2026) 8
- 3 Emission-regime families 10
- 4 Reward and verification families 12
- 5 Value-accrual classes 13
- 6 Onocoy case summary 18
- 7 Onocoy Dune observability anchor 20
- 8 Comparator lineage funnel 22
- 9 Metric applicability matrix 24
- 10 Historical stress windows 27
- 11 DTSE architecture overview 31
- 12 DTSE scenario setup 34
- 13 DTSE metric selection 36
- 14 Descriptive metric definitions 37

15	Stage 2 failure-signature logic	38
16	Cross-scenario detection pattern summary	43
17	Onocoy (ONO) baseline-relative signal sequencing	44
18	Observed Stage 2 summary for ONO	47
19	Governance and intervention archetypes	54
20	Limitation-to-claim map	55
A.1	Scenario logic and key fixed settings for the DTSE experiment set.	58
A.2	Excerpt of the fixed thresholds used in the DTSE reporting logic.	59
A.3	Main profile-specific and global overrides used in the reported DTSE runs.	60
A.4	Run-manifest excerpt for the frozen DTSE reporting set.	61
A.5	Deterministic seed schedule with fixed-rule vs variable-draw separation.	61
A.6	Baseline excerpt (no stress event): same fixed DTSE rules, different seed path.	61
A.7	Liquidity-shock excerpt (event at week 20): same fixed DTSE rules, different seed path.	62

1 Introduction

1.1 Why This Problem Matters

A token holder can exit a position in seconds. An operator who has bought, installed, and calibrated site-specific hardware cannot. In DePIN, a rooftop Global Navigation Satellite System (GNSS) station, wireless device, or other physical asset is not rebalanced with the same ease as software-mediated capital. Once infrastructure is deployed, exit and re-entry become asymmetric: participation can deteriorate quickly when profitability disappears, while restoring capacity is slower because it requires renewed capital, logistics, and operating effort [1, 2, 3].

This speed mismatch is what makes DePIN analytically different from software-only crypto systems. DePIN networks coordinate real-world infrastructure through token-based incentives, but the token is not merely a peripheral exchange layer. It is the coordination mechanism that determines whether infrastructure is deployed, retained, and kept serviceable under changing market conditions [4, 5]. When incentives weaken, providers do not adjust only by moving tokens; they may defer maintenance, reduce operational commitment, or eventually power down hardware altogether.

Much of the public discussion around DePIN still emphasizes adoption, network growth, and token appreciation. Growth can be visible long before underlying robustness is tested. Those themes matter, but they do not answer the harder operational question: what happens when demand weakens, liquidity tightens, provider costs rise, or adjacent networks offer better yields for similar hardware [6, 7]? For hardware-dependent systems, robustness is not mainly an equilibrium story. It is a stress story.

1.2 Problem Setting

We therefore treat DePIN tokenomics as a cyber-physical coordination problem rather than as a purely financial mechanism. Under adverse conditions, the incentive system must continue to support three coupled functions:

1. **Provider retention**, so participation does not erode faster than the network can absorb;
2. **Service continuity**, so usable capacity and output remain available; and
3. **Incentive–usage alignment**, so reward outflows do not remain detached from usage-linked value capture for too long.

Those functions do not necessarily fail at the same time. A demand contraction can weaken usage-linked sinks before participation visibly falls. A liquidity shock can compress provider

margins before any service metric moves. Competitive yield pressure can pull supply away from a network even when internal demand has not yet collapsed. For that reason, DePIN stress evaluation is less about assigning static mechanism labels and more about tracing transmission paths: where deviations appear first, how they propagate, and when they become operationally meaningful [8, 9].

1.3 Research Question and Contribution

The preceding argument narrows the problem in a specific way. If DePIN robustness depends on how stress propagates through provider retention, service continuity, and incentive–usage alignment, then our central task is not simply to label tokenomic designs as stable or unstable. We instead ask how different mechanisms transmit stress, where those pressures register first, and how robustness can be evaluated before breakdown becomes clearly visible.

We organize the analysis around one primary research question and one subordinate evaluative question:

How do different DePIN tokenomic architectures, specifically Burn-and-Mint Equilibrium (BME)-oriented and capped-supply designs, transmit physical and economic stress through provider participation, token markets, and service utilization?

How can robustness be evaluated before failure becomes operationally visible?

Our answer is comparative rather than predictive. We do not aim to forecast token prices or identify a universal winner among mechanism classes. Instead, we examine where stress concentrates first, how deviations propagate across provider retention, service continuity, and incentive–usage alignment, and how those sequences can be evaluated before breakdown becomes clearly visible under standardized stress conditions.

This framing yields three contributions. First, we shift the evaluation of DePIN tokenomics away from growth-phase narrative and toward stress transmission under physical constraints [1, 2]. Second, we use the Onocoy network as an empirical anchor for translating documented mechanism rules into a tractable evaluative setup. Third, we develop and apply DTSE, the DePIN Tokenomic Stress Evaluator, as a reproducible comparative framework for evaluating how alternative mechanism profiles react to matched stress channels under explicit assumptions [8, 9].

1.4 Scope and Thesis Roadmap

Our scope is deliberately bounded. We do not claim to measure live-network behavior directly, and we do not treat simulation outputs as forecasts. Instead, we combine documented mechanism facts, empirical stress context, explicit model assumptions, and DTSE outputs to support comparative statements about sensitivity, failure signatures, and robustness within a defined evaluative boundary.

The remainder of the thesis follows that same sequence. We first establish why DePIN is analytically distinct from software-only crypto systems, then develop the comparative tokenomic vocabulary needed for stress evaluation. We next ground the analysis in the Onocoy case, clarify what can and cannot be observed empirically, formalize DTSE, and finally use that framework to evaluate stress behavior and interpret its implications.

2 Foundations: DePIN as a Cyber-Physical System

The pressure points in DePIN tokenomics do not begin with token design alone. They begin with the fact that the network is built on physical infrastructure that must be purchased, placed, operated, and maintained in the real world. That difference changes how participation forms, how capacity erodes, and how quickly a network can recover once adverse conditions appear. In this chapter, we establish the physical and economic foundations that make DePIN stress analytically distinct from software-only crypto systems.

2.1 DePIN as a Cyber-Physical Coordination Problem

Decentralized Physical Infrastructure Networks (DePIN) coordinate the deployment and operation of real-world infrastructure through token-based incentives. Rather than organizing only capital positions or software-mediated state transitions, they coordinate tangible service delivery through assets such as wireless hotspots, GNSS reference stations, compute nodes, storage nodes, and sensor devices [5, 1, 10]. In that setting, tokenomics is not a detachable financial wrapper around an otherwise independent service. It is part of the coordination mechanism that shapes whether infrastructure is deployed, retained, and kept serviceable over time [4, 5].

That makes DePIN closer to a cyber-physical coordination problem than to a purely digital token economy. In software-only systems, participation is usually liquid, reversible, and relatively easy to reallocate. In DePIN, participation is tied to physical assets with specific operating requirements and location constraints. A minimal coordination loop therefore links four elements: infrastructure providers who deploy hardware, end users who consume the

resulting service, protocol rules that determine what counts as valid work, and a token layer that connects service contribution to compensation and governance [11, 7]. The important point is not the taxonomy itself, but the dependence it creates. If one part of this loop weakens, the effects do not remain financial; they propagate into physical capacity and service quality.

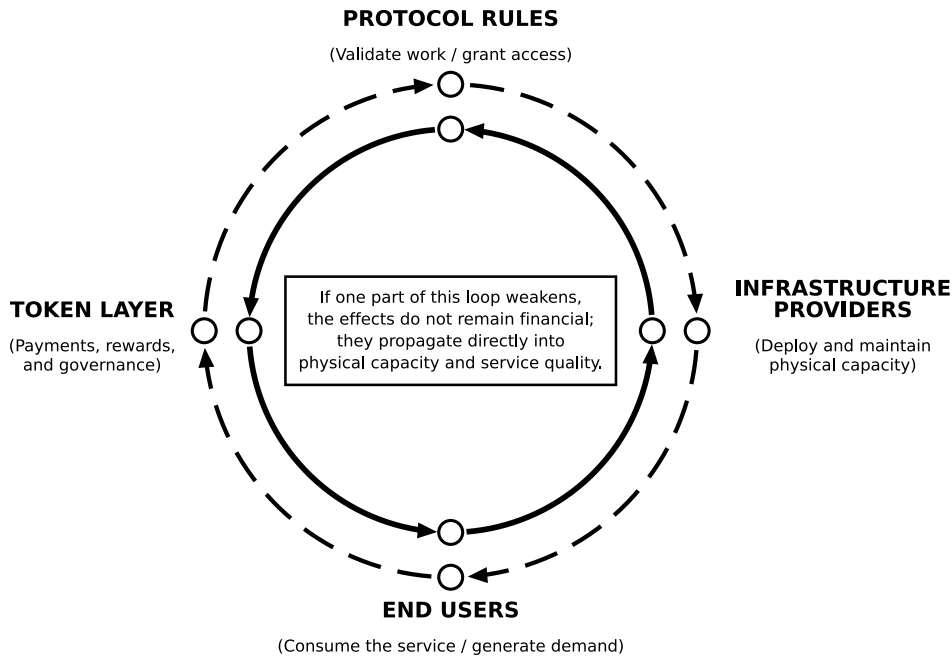


Figure 1: The cyber-physical coordination loop in a DePIN, adapted from Chiu et al. (2024). The outer solid ring shows service and data flows; the inner dashed ring shows token, payment, reward, and governance flows. A weakness in any one segment propagates structurally through the others.

2.2 Physical Constraints and Economic Stickiness

The hardware layer introduces a cost structure that software-only systems do not face. Providers incur capital expenditure through device purchase, installation, setup, and site preparation, and they continue to face operational expenditure through electricity, connectivity, maintenance, and uptime management [3, 2]. In GNSS-style networks, these requirements are not incidental. Equipment quality, placement, calibration, and continuity of operation all affect the usefulness of the service delivered downstream.

Once capital is committed to site-specific infrastructure, participation becomes economically sticky. A token position can often be exited by selling or transferring an asset. Software-mediated capital can often be reallocated very quickly; physical capacity cannot. Uninstalling, transporting, and redeploying a rooftop GNSS station, hotspot installation, or other site-specific hardware requires time, logistics, and renewed operating effort. This is the practical significance of sunk costs in DePIN: they slow immediate exit, but they also make the network

less flexible when conditions deteriorate [3, 1]. Geographic friction deepens that effect. A well-placed station or node does not merely represent one interchangeable unit of capacity; it occupies a location that may matter for redundancy, coverage, or service quality. As a result, participation is not only less liquid, but also less substitutable.

These frictions can temporarily act as a retention buffer. Providers may tolerate periods of weaker profitability because hardware is already installed and operating costs are lower than full redeployment elsewhere. But that same rigidity has a second edge. When adverse conditions persist, the network does not simply lose abstract token participation. It risks losing physical coverage, operational discipline, and location-specific capacity that is costly to rebuild.

2.3 Path Dependence and Recovery Asymmetry

Because physical assets cannot be reallocated as easily as software-mediated capital, DePIN networks exhibit path dependence. Expansion, contraction, and recovery do not follow the same timeline. A network can lose momentum relatively quickly when token rewards weaken or provider economics deteriorate, but restoring that capacity is slower because re-entry requires renewed capital commitment, logistical effort, and confidence that conditions will remain viable long enough to justify redeployment [3, 2].

This asymmetry matters because visible network deterioration often lags the underlying stress. Providers may remain nominally present while reducing maintenance effort, delaying upgrades, or tolerating declining service quality. In networks where geographic density contributes directly to output quality, the loss of a few strategically important nodes can matter more than raw node counts suggest. The analytical problem is therefore not only whether participation declines, but how physical rigidity changes the sequence by which economic pressure becomes operational degradation.

That is why we treat robustness as a question of stress transmission rather than static equilibrium. A mechanism may appear workable when demand is stable, token prices are supportive, and participation remains nominally intact. Those conditions matter, but they do not show how quickly pressure moves through the network once provider economics weaken, coverage starts to thin, or the link between rewards and service demand begins to loosen. Our central issue is therefore not whether a mechanism looks sustainable under favorable average conditions. It is whether the system can absorb adverse conditions without losing provider retention, service continuity, or incentive–usage alignment faster than it can recover. In DePIN, the decisive question is often one of sequence: where strain appears first, how long physical frictions delay visible failure, and when that delay stops acting as a buffer and starts becoming a liability.

This dynamic also implies that DePIN networks may enter periods of latent stress, in which provider economics deteriorate before observable participation metrics such as node counts or network coverage visibly adjust.

2.4 Theoretical Stress Constructs

Stress, in this thesis, refers to adverse economic or structural conditions that threaten the continued operation of a DePIN network. In practice, those conditions can emerge through lower demand, weaker liquidity, rising provider costs, or broader market dislocation. What makes them analytically important is that they do not remain confined to token price. They propagate through provider economics and can eventually reach service delivery itself.

To organize that problem, we use three theoretical stress constructs. They are not presented as an exhaustive industry taxonomy. They are the interpretive anchors that make the later empirical and DTSE analysis legible.

Concept	Meaning in this thesis	Operational pressure
Reward Addiction	A condition in which participation can only be maintained through continued or rising token incentives, even when underlying service demand is not keeping pace.	If incentives flatten before demand-side support strengthens, provider retention can deteriorate quickly.
Subsidy Gap	The distance between providers' real-world operating burden and the fiat-equivalent value entering the system through actual service usage.	The wider the gap, the more the network depends on speculative token value to sustain supply.
Speculative Fragility	The degree to which participation and security become sensitive to token-price volatility rather than to service demand.	Price dislocations can transmit into provider economics before service metrics visibly move.

Table 1: Core stress constructs used in the thesis.

Taken together, these constructs explain why DePIN sustainability cannot be evaluated through growth narratives alone. Reward Addiction highlights the danger of participation supported mainly by emissions. The Subsidy Gap clarifies why real provider economics matter even when token incentives appear attractive in nominal terms. Speculative Fragility explains why price and liquidity shocks may affect infrastructure provision even when end-user demand

has not yet visibly collapsed. In the next chapter, we turn these pressures into a comparative tokenomic vocabulary. We ask how different design choices attempt to manage them, where those attempts remain structurally exposed under stress, and why superficially similar DePIN models can still fail through different transmission paths.

3 Comparative Tokenomic Framework

In this chapter, we define the comparative vocabulary used to distinguish DePIN tokenomic designs before the thesis narrows to the Onocoy case and, later, to DTSE. Our goal is not to rank projects by short-horizon token performance. Instead, we identify the design features that determine how infrastructure demand, incentive issuance, work verification, monetization, and value capture interact once a hardware-dependent network comes under stress. In DePIN, these elements do not operate only in digital markets; they also shape whether physical infrastructure is deployed, maintained, and kept usable under changing conditions. We therefore combine a dated market snapshot, used only to situate the comparator set, with a mechanism-level comparison of emission logic, reward architecture, monetization, and value accrual [1, 2, 8, 9].

3.1 Comparator Frame and Time-Stamped Market Snapshot

We use a bounded sample of large Solana-native or Solana-bridged DePIN protocols that provide verifiable infrastructure services, such as wireless coverage, GNSS correction, mapping, compute, or data infrastructure. The purpose of the sample is not exhaustive market representation. It is to create a comparable set of mechanism profiles that sit close enough to one another to make structural tokenomic differences visible, while still spanning several infrastructure verticals [4, 12, 13].

Solana is used here as a pragmatic sampling frame rather than as a claim of universal architectural superiority. For the period in which many of these projects matured, it provided a combination of low transaction costs, throughput suited to high-frequency infrastructure interactions, and a sufficiently dense DePIN ecosystem to make cross-project comparison analytically coherent [14, 13, 15]. That framing matters because the chapter is comparing tokenomic structures inside a shared execution environment, not arguing that later DePINs could not make different chain-level choices.

The market layer is included here as dated scene-setting context rather than as durable evidence of robustness. A market snapshot can help explain why these projects matter in the DePIN landscape and why certain designs were visible comparators at the time of writing, but it cannot by itself establish which tokenomic structures are more resilient. For that reason, the

figures in Table 2 are explicitly tied to a single observation window and are used descriptively rather than inferentially.

Project	Infrastructure focus	Snapshot market cap	Reason for inclusion
Helium	Wireless connectivity	\$282.99M	Canonical DePIN reference point for burn-credit demand coupling and infrastructure bootstrapping.
Grass	Bandwidth / data collection	\$202.91M	Useful contrast case for weakly specified value accrual and strong participation narrative.
GEODNET	GNSS correction	\$57.35M	Close vertical comparator for location-based sensing and correction services.
io.net	Distributed compute	\$39.03M	Important example of revenue-linked tokenomics in compute-oriented infrastructure.
Hivemapper	Mapping	\$21.64M	Strong reference case for burn-and-mint credit logic tied to service consumption.
XNET	Wireless offload	\$15.10M	Illustrates revenue-funded buyback logic in connectivity infrastructure.
Nosana	Compute marketplace	\$8.43M	Useful example of pay-in-token utility with weaker explicit sink logic.
Aleph.im	Compute / storage / cloud	\$5.41M	Helps capture utility-token infrastructure with partially opaque accrual mechanics.
Onoco	GNSS correction	\$1.53M	Included as the primary empirical anchor because mechanism access and case-specific evidence are available.

Table 2: Comparator snapshot on 18 February 2026 (dated market context). Values are included only to situate the comparator universe at the observation date and are not used as robustness evidence or ranking criteria.

Two points matter immediately. First, the sample is heterogeneous by infrastructure

category: wireless, mapping, GNSS, compute, bandwidth, and hybrid cloud layers do not face identical adjustment dynamics. Second, superficially similar scarcity narratives can still conceal very different tokenomic regimes. A capped supply, a burn mechanism, or a tokenized payment rail says little on its own unless the chapter also asks how tokens enter circulation, how work is verified, how usage is monetized, and whether demand creates a measurable path back into token value.

3.2 Emission Logic and Supply Regimes

Emission design is the first comparative fault line. In this chapter, emission logic refers to the rule set that governs how tokens enter circulation over time, how distribution adjusts or fails to adjust under changing conditions, and whether issuance remains linked to service demand, activity, or a fixed release path. Across the sampled projects, nominally capped supply is common, but the path by which supply reaches circulation differs sharply. That difference matters because long-run scarcity and short-run dilution are not the same thing. A token can have a hard upper bound and still place considerable pressure on provider economics if distribution remains aggressive, weakly demand-coupled, or slow to adjust once market conditions deteriorate.

Table 3 groups the comparator set into regime families rather than treating each project as a fully unique case. The point is not to erase project-specific details. It is to make visible which supply logics are likely to behave differently once stress begins to propagate through demand, price, and provider retention.

Regime family	Representative projects	Mechanism logic	Stress relevance
Step-down issuance	Helium, GEODNET, Onocoy	Distribution falls according to a scheduled reduction path, usually independent of short-run demand conditions.	Mitigates long-run dilution pressure, but can still leave acute mismatch windows between reward decline and realized usage.
Fixed-rate or long-run scheduled issuance	XNET, io.net, Grass	Issuance follows a predefined release path or fixed distribution logic rather than a hard halving cycle.	Predictable distribution can simplify planning, but may sustain subsidy pressure under weak demand if the path is slow to adjust.
Activity-mediated net issuance	Hivemapper	Net supply pressure depends partly on service consumption, burn intensity, or remind constraints.	Can strengthen demand coupling, but also introduces sensitivity to measurement, governance parameters, and procyclical feedback.
Effectively fixed circulating supply	Nosana, Aleph.im	Ongoing protocol issuance is limited or no longer central relative to the already circulating stock.	Reduces dilution pressure, but shifts the robustness question toward utility strength, sell pressure, and accrual clarity.

Table 3: Emission-regime families used for comparative analysis.

The comparative implication is straightforward but important: a capped token is not automatically a low-pressure token. What matters under stress is whether issuance remains

rigid when demand weakens, whether supply reduction is delayed or state-sensitive, and whether the design leaves room for prolonged subsidy dependence even while preserving a scarcity narrative. For this thesis, we therefore treat emission logic less as a branding distinction and more as a first indicator of how quickly token-side pressure can build when real usage no longer keeps pace [16, 17, 18, 19, 20].

3.3 Reward Logic, Work Verification, and Quality Adjustment

Emission rules alone do not explain how a DePIN tokenomics design behaves. The second fault line concerns what is actually being rewarded and how the protocol determines whether a contribution is valid. A network that rewards raw participation, one that rewards validated coverage, and one that rewards outcome-quality data are all distributing tokens, but they are not funding the same thing. The reward architecture therefore has to be read together with the verification primitive.

This matters because DePINs do not just need supply; they need usable supply. Wireless networks must verify that coverage is real. Compute networks must verify that jobs were completed correctly. Sensing and mapping networks must verify that data is not merely submitted, but reliable enough to matter operationally.

Across the sample, this produces a small but important family of reward logics: proof-of-coverage style verification for connectivity networks, outcome-quality or threshold-gated rewards for sensing and GNSS systems, and job-completion or validator-backed checks for compute services. In practice, the difference is not only technical. It determines whether the protocol is paying for nominal participation, for service availability, or for output that clears a quality threshold. Where those verification layers are stronger, rewards are more defensibly tied to infrastructure value creation rather than to nominal participation alone [1, 7, 21].

Project family	Primary verification logic	What the reward architecture is trying to align	Stress relevance
Wireless coverage	Helium, XNET	Coverage attestations, authenticated offload, or measurable traffic events stand in for usable network service.	Weak demand can expose whether coverage rewards remain tied to actual traffic or persist mainly as subsidy.
Sensing / mapping / GNSS	GEODNET, Hivemapper, Onocoy	Availability, data quality, and location relevance are used to distinguish usable output from low-value supply.	Quality-adjusted rewards can defend scarce high-value nodes, but they also reveal where maintenance erosion becomes operationally costly.
Compute / cloud	io.net, Nosana, Aleph.im	Job completion, validator checks, node scoring, or stake-backed reliability mechanisms proxy usable compute service.	Under stress, reward credibility depends on whether verification remains tied to reliable service rather than residual token incentives.

Table 4: Reward and verification families across the comparator set.

The main comparative point is that verification and quality adjustment are not secondary implementation details. They are part of the tokenomic design. A project with stronger verification can target rewards more selectively, reduce leakage to low-value supply, and make provider incentives more legible under stress. Conversely, where quality weighting is weak, underdocumented, or difficult to audit, the network is more exposed to the classic subsidy problem: rewards may continue to flow even when the protocol has only limited assurance that it is purchasing resilient infrastructure.

3.4 Monetization, Sinks, and Value Accrual

The third fault line concerns how real usage enters the system and whether that usage creates a credible path back to token value. This is where superficially similar narratives often diverge most sharply. A burn mechanism, a buyback rule, a token payment rail, and a fiat-priced credit system can all be described as “value accrual,” but they do not produce the same type of linkage between demand and token-side support.

For comparative purposes, the sampled projects can be grouped into three main accrual classes.

Accrual class	Representative projects	Comparative meaning
Mechanistic usage-burn loops	Helium, Hivemapper, Onocoy	Service usage is converted into credits or equivalent units through a token-linked burn step, creating the clearest direct pathway from demand to token removal.
Revenue-mediated buyback regimes	GEODNET, XNET, io.net	Demand produces revenue that may later support buybacks or burns, but the link between usage and token support depends on revenue attribution, rule enforcement, and governance credibility.
Utility/payment models without systematic burn	Nosana, Grass, Aleph.im	Demand may still require the token or support token utility, but the pathway back to token value is weaker, more indirect, or less publicly specified.

Table 5: Value-accrual classes across the comparator set.

This distinction is central for the rest of the thesis. Mechanistic usage-burn loops are easiest to interpret because the demand path is visible and comparatively close to protocol logic. Revenue-mediated regimes can be powerful, but they are institutionally more fragile because the token’s support depends on accounting transparency, revenue routing, and credible commitment to buyback execution. Utility/payment models without systematic burn may still matter economically, yet they make it harder to infer whether rising usage should translate into persistent token support or only into short-lived transactional demand [16, 17, 18, 22, 20].

3.5 Demand Regimes and Utility–Speculation Profiles

Comparative tokenomics cannot stop at supply and sinks. The same mechanism can behave very differently depending on who the demand-side user is and how directly token demand is tied to service consumption. For that reason, two additional bounded lenses are useful here:

demand regime and utility–speculation profile. Both are best treated as comparative heuristics rather than as hard empirical rankings.

The first lens distinguishes whether a project primarily targets enterprise buyers, consumer participation, or developer and API-facing infrastructure use. This matters because enterprise and developer demand often arrives through integration workflows, contracts, or pricing expectations that are very different from consumer or narrative-driven demand [2, 6]. The second lens asks whether token demand appears to be driven mainly by service utility, by speculative holding and trading, or by a mixed combination of the two.

These lenses should remain bounded. They do not measure realized adoption, and they do not directly prove robustness. What they do provide is a way to read later case material more carefully. A GNSS network oriented toward professional demand and fiat-stable usage credits is exposed differently from a network whose token utility is weaker, more optional, or more entangled with narrative demand. Likewise, a project with strong on-chain usage signals and explicit sink mechanics can be interpreted differently from one whose accrual story remains opaque or underdocumented.

For Chapter 3, the value of these lenses is therefore comparative and diagnostic. They help us explain why projects that all fall under the DePIN label may nonetheless differ materially in monetization structure, sensitivity to speculative pressure, and the plausibility of demand-linked value support. In that sense, they do not replace the harder tokenomic comparison above. They sharpen it.

3.6 Framework Synthesis

Taken together, the comparative picture is more structured than the category label “DePIN” suggests. Projects differ not only in infrastructure vertical, but also in how tokens reach circulation, what counts as rewardable work, how usage enters the protocol, and whether value accrual is mechanistic, revenue-mediated, or only weakly specified. Those differences matter because they determine where stress is likely to register first: in dilution pressure, in verification leakage, in usage-to-token linkage, or in the gap between speculative demand and operational demand.

This chapter therefore establishes four comparative questions that remain active for the rest of the thesis. First, how rigid is the supply path once adverse conditions appear? Second, how selectively does the protocol target rewards toward usable work? Third, how strong is the path from real service demand to token-side support? Fourth, how much of the token’s observed demand appears to be utility-linked rather than narrative-driven? These questions do not yet answer which design is more robust. They do, however, give us the mechanism vocabulary needed to interpret the Onocoy case in the next chapter and to understand why

later DTSE comparisons are meaningfully more than token-price simulations.

4 Onocoy as the Anchor Case

Onocoy functions as the empirical anchor for this thesis because it combines three features that are analytically useful together: a hardware-dependent DePIN service, a capped-supply token design, and a dual-layer payment architecture that separates user-facing service pricing from direct token volatility. That combination gives us a strong case for asking how infrastructure quality, user demand, and token-side incentives interact when stress enters a network that cannot rely on purely digital capital mobility.

We do not treat Onocoy as a proxy for the whole DePIN landscape. We use the case more narrowly. Onocoy provides a documented mechanism design that can be described with enough clarity to support the later empirical and methodological chapters, while still exposing the kinds of documentation gaps and design tradeoffs that matter for a bounded stress-evaluation thesis.

4.1 System Function and Network Participants

Onocoy operates a decentralized GNSS correction network based on Real-Time Kinematic (RTK) positioning. RTK combines correction data from reference stations with rover-side measurements in order to improve positioning accuracy from meter-scale error bands toward centimeter-level precision, which matters for surveying, precision agriculture, robotics, and industrial workflows where positional error is operationally costly [22].

The network depends on geographically distributed reference stations rather than on a single centralized infrastructure operator. In practical terms, this means that the quality of the system depends not only on whether stations exist, but on where they are installed, how reliably they remain online, and whether their correction output is usable within the network's service layer [22]. That makes Onocoy analytically useful for this thesis because it sits inside an already established GNSS ecosystem: instead of inventing demand from nothing, it inserts a decentralized correction layer into positioning workflows that already care about continuity, precision, and coverage quality.

That creates two primary participant groups. First, reference-station operators deploy and maintain GNSS hardware, connectivity, and quality-relevant uptime. Their role is not passive. Installation quality, calibration, and maintenance discipline affect the usefulness of the correction data they contribute [22, 23]. Second, rover-side users purchase correction access for high-accuracy positioning tasks. In the Onocoy case, the expected demand profile

is closer to professional or industrial usage than to purely narrative retail usage, which makes continuity, measured quality, and coverage density central to the service proposition [22].

Onocoy’s public Explorer adds an operational transparency layer to that service model. The explorer does not only display that stations exist; it visualizes a coverage heatmap and exposes station-level details such as approximate validated location, validation status, uptime, signals, quality scale, and estimated rewards [24]. For this thesis, that matters because it makes the network spatially tangible: users and operators can observe where coverage is already dense, where overlap exists, and where additional stations may improve useful correction availability.

Separate from the thesis interview process, the Onocoy ecosystem also appears to include an active community-support layer around operators. Public documentation describes ambassador and community channels that support ecosystem participation, while bounded practitioner context suggests these spaces are also used to exchange installation practices, reward information, and deployment feedback [25]. This thesis treats that layer as ecosystem context rather than as a mechanism fact, but it still helps explain why participation in a GNSS DePIN is operationally richer than simply switching a device on and waiting for rewards.

4.2 Token and Payment Architecture

Onocoy uses a dual-layer design with ONO as the transferable utility and governance token, while Data Credits function as the non-transferable unit for service consumption. Official tokenomics documentation describes ONO as a capped utility token with a maximum supply of 810 million units, while Data Credits are purchased against fiat-denominated value and then consumed during network use [22, 20, 26]. That separation matters because it allows the immediate user-side service price to remain more stable than the market price of ONO itself.

This architecture matters because it separates two questions that are often blurred in simpler token systems. The first is whether users face a predictable service price. The second is whether the protocol still creates a credible path from service demand back into token-side support. In the documented flow, users prepay for service through Data Credits, those credits are burned when GNSS data services are consumed, and the resulting fiat revenue finances operations and may also be used for ONO buybacks. The tokenomics documentation further states that bought-back ONO can be split across the reward pool, ecosystem pool, and token burn [26]. The key point is not only that ONO and Data Credits are different units. Pricing, service consumption, and token-side settlement are separated enough to shield users from direct token volatility while still keeping demand economically relevant to the incentive system.

Public Onocoy documentation therefore supports three mechanism facts particularly

clearly: a capped ONO supply, a release path described as a 16% annual reduction in newly distributed units, and a Data Credit layer whose use is tied to burn and conditional buyback routing [22, 26]. That makes the payment architecture unusually important to the case. If the user-facing credit layer creates stability but the routing back into token-side support is too weak, the network can still face stress on the supply side even when users see a smoother service price.

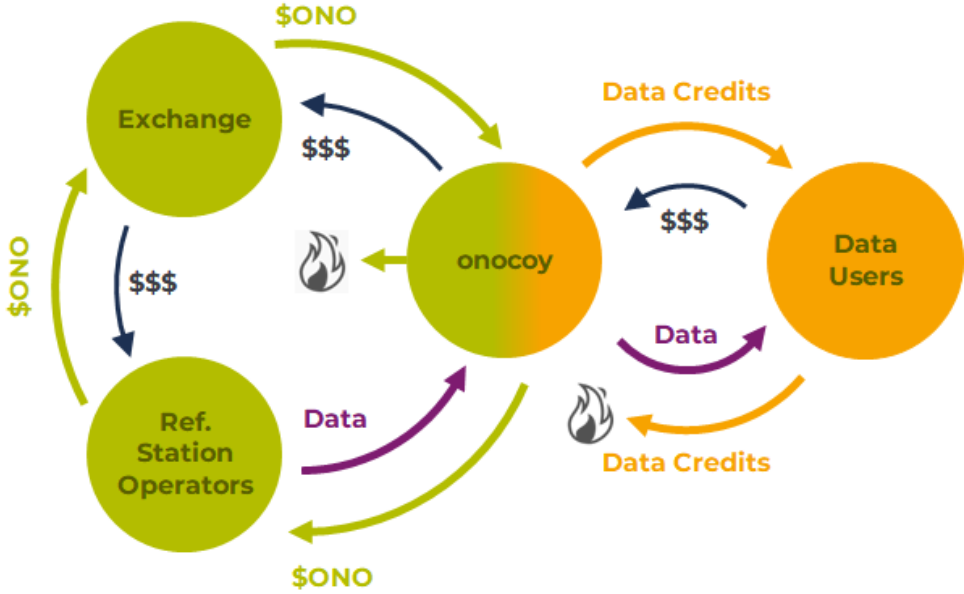


Figure 2: Official Onocoy tokenomics flow, showing how ONO, Data Credits, data users, exchange liquidity, and reference-station operators interact in the public design description [26].

The public-doc picture is stronger for ONO issuance and Data Credit mechanics than for every operational participation detail. For this thesis, that distinction matters. Emission and payment architecture can be described as mechanism facts. At the same time, the current public design does not center participation on an explicit staking- or collateral-based miner model. Where later chapters require assumptions, they arise from other operational unknowns rather than from a missing stake-based participation layer [20, 23].

Dimension	Onocoy case summary
Infrastructure domain	Decentralized GNSS correction network using RTK positioning services.
Provider role	Reference-station operators deploy and maintain GNSS hardware, connectivity, and quality-relevant uptime.
Demand side	Rover-side users purchase correction access for professional and industrial positioning workflows.
Token / credit architecture	Transferable ONO token plus non-transferable Data Credits used for service consumption.
User-side pricing logic	Data Credits are priced against fiat-denominated value rather than exposing users directly to ONO market volatility.
Supply-side release path	Public documentation describes capped ONO supply with a decaying distribution path.
Documentation boundary	Public materials are explicit about ONO and Data Credit mechanics and do not present a staking- or collateral-based participation model; the thinner areas lie in other operational participation details.

Table 6: Onocoy case summary synthesized from public protocol and tokenomics documentation cited in this section; descriptive only.

4.3 Settlement-Layer Scope and Decision Boundary

We treat Onocoy’s blockchain choice as a settlement-layer design decision tied to timing, throughput and cost requirements, and off-chain service delivery constraints rather than as a general claim that one chain is universally superior. The practical requirement in scope is straightforward: the network must support frequent, low-value accounting events while keeping latency-sensitive GNSS correction delivery off-chain for operational reasons [22, 20].

That decision is best understood in the project window in which many DePIN teams were choosing between operating sovereign chain infrastructure and deploying into an existing settlement ecosystem. Published chain-selection material for Onocoy describes a structured screening funnel and links the final choice to throughput, cost, block-time expectations, ecosystem readiness, and engineering execution constraints during that decision period [14]. In this chapter, we use those factors as project-primary evidence for a time-bounded architecture choice, not as a timeless benchmark for all future DePINs.

For Onocoy specifically, the practical advantage is not abstract chain branding. It is operational fit. A settlement layer with relatively low transaction cost and high-frequency processing makes it easier to meter small-value service usage, support fine-grained accounting,

and avoid turning the economic layer into a bottleneck for either providers or users.

On the supply side, that reduces friction around recurring reward accounting and token-side settlement. On the demand side, it supports a service model in which pricing and access can be broken down more granularly than a slower or more expensive chain would comfortably permit. In that sense, settlement-layer choice becomes part of DePIN mechanism viability rather than a purely backend engineering preference.

The broader industry context reinforces that bounded reading. A DePIN project can either build and maintain dedicated chain infrastructure or use an existing settlement environment while concentrating product effort on service delivery and mechanism operations. Helium's migration record is relevant here as contextual evidence that chain-maintenance burden can compete directly with product-layer execution capacity, but it is not treated as proof that Onocoy's own choice is universally optimal [15]. More generally, DePIN projects are forced to make execution-layer tradeoffs under the constraints of their own era, product requirements, and ecosystem options. That is why this thesis treats Onocoy's chain choice as a historically situated design decision rather than as a universal model for every later DePIN.

4.4 Case Relevance, Documentation Gaps, and Evidence Boundary

Onocoy matters to this thesis because it narrows the comparative framework developed in Chapter 3 into a concrete capped-supply design that still depends on physical infrastructure quality, measurable service demand, and token-side coordination. The case is especially useful because the project does not only present a token architecture. It also exposes how difficult it is to document every relevant participation rule with equal clarity once a DePIN begins to combine off-chain service delivery, quality-sensitive hardware, and on-chain economic settlement.

We have to handle that asymmetry in documentation explicitly. Public sources are sufficiently clear to support claims about RTK service function, the dual-layer ONO/Data Credit design, and the general supply-decay path. They also indicate that the current public design is not framed around staking or collateral participation. The thinner areas lie elsewhere. Some operational participation details remain less explicit than the token and payment architecture itself. Where those items matter later, they must be treated either as unknown mechanism facts or as declared modeled assumptions rather than as silently completed facts.

The same discipline applies to the Explorer, to Dune, and to the interview material used in this thesis, but for different reasons. In this chapter, we treat the Explorer as an operational and spatial transparency layer: it shows coverage, station-level quality signals, and deployment geography. Dune, by contrast, is used as a public on-chain analytics layer that can support dated, auditable snapshots of selected token-side and activity-side indicators [24, 27]. The

point of using a dated Dune snapshot here is not to prove growth, robustness, or long-run trend. It is to show which layers of the Onocoy case are publicly visible before the thesis later asks what can be observed empirically and what must instead be handled through bounded evaluator logic.

Observable Layer	Snapshot Values (25 February 2026)	What the Snapshot Makes Visible
Token-holder and supply visibility	1,568 holders; 75,713,648 circulating ONO	Shows that token-side participation and circulating supply are at least partly visible through public accounting surfaces rather than entirely hidden from the analysis.
Buyback and burn accounting	2,928,777 total ONO buybacks; 2,024,536 total ONO burned	Shows that some sink- and treasury-related token flows can be observed directly enough to support dated descriptive claims about the settlement layer.
Participation-side network visibility	7,668 validated stations; 5,638 online stations	Shows that network participation is visible not only as a token system, but also as a live infrastructure layer with measurable validated and online station counts.

Table 7: Dune observability anchor for Onocoy on 25 February 2026. Values are used as a dated visibility check of publicly observable fields, not as trend or performance evidence.

This snapshot anchors observability rather than performance. It shows that some token-side, accounting-side, and participation-side fields are publicly visible at a defined point in time, while other economically important dimensions remain much less settled. On its own, it cannot establish trend direction or resolve questions such as revenue depth, churn history, or client concentration. That asymmetry is part of what makes Onocoy analytically useful as an anchor case.

The thesis also includes bounded interview-based practitioner context, while the public ambassador and community material is treated separately as ecosystem context rather than as a mechanism description. We use the interview layer to identify documentation gaps, clarify how operators and the project team describe practical conditions, and surface where public materials are thinner than the mechanism would ideally require. It does not carry the same evidentiary weight as protocol documentation or directly auditable public artifacts, and it is not treated here as a substitute for them. A compact summary of the interview themes used as bounded practitioner context is retained in Appendix A.5.

Taken together, the public documentation, official tokenomics description, Explorer visibility layer, dated Dune snapshots, bounded interview context, and ecosystem-community layer make Onocoy especially relevant to the thesis. The case is documented enough to describe clearly, operational enough to visualize concretely, and incomplete enough to make evidence discipline visible rather than optional. That combination is precisely what makes it a useful anchor before the thesis turns to empirical observability limits and then to the narrower logic of DTSE.

5 Empirical Stress Layer

5.1 Empirical Analysis Orientation

Chapters 2 and 3 established the stress vocabulary used in this thesis. Stress was defined there as pressure on the DePIN system when subsidy support weakens, market volatility compresses operator economics, or provider incentives become misaligned with continued infrastructure participation. Chapter 4 then translated that vocabulary into the Onocoy case. Our next step is not yet to model new scenarios. It is to ask what the historical record can already show about these pressures when they appear in live networks.

This chapter therefore works with *historical stress windows*. A historical stress window is a bounded real-world episode in which a pressure becomes visible in public market, activity, or participation data. The term matters because the chapter is retrospective. We look at episodes that actually happened, not hypothetical episodes created for a model. That makes the empirical layer useful for observing which signals became visible in practice, but it also means the evidence is constrained by whatever the public record happened to capture.

Historical stress windows can show where live DePIN systems first register pressure. They can show price dislocation, participation stickiness, weakening reward–demand coupling, or rising cost sensitivity after the fact. They cannot, by themselves, show where a newer network with thinner history would cross an unobserved failure threshold. That boundary matters for Onocoy, where several economically relevant fields remain only partially observable in the public record.

In this chapter, we add a retrospective empirical layer to the thesis. We use event-study logic to read documented stress windows in Helium, Geodnet, Hivemapper, Render, io.net, and Onocoy without treating those windows as proof of future outcomes [28, 5]. The task here is narrower: identify what historical data can reveal, make the limits of cross-network comparison explicit, and show why the move to DTSE becomes necessary once empirical observability runs out.

5.2 Stress Windows and Comparator Scope

Retrospective comparison is only useful if scope and evidence roles stay clear. We use a Solana-centered DePIN sample as a practical frame because these networks share enough execution-era context to make mechanism and window comparisons interpretable, even though their hardware profiles and operating costs remain different [5, 4]. Mechanism facts remain anchored to protocol and project-primary documentation. Ecosystem reports are used more narrowly to situate the observed window and explain why certain comparators matter.

The comparison set in this chapter is broader than the smaller frozen set used later for controlled evaluation. The immediate point is simple: we first look widely across public historical cases, then narrow to the cases and metrics that remain interpretable, and only later move to a fixed evaluation set for controlled comparison under explicit assumptions.

Comparison Scope	Projects Included	What This Layer Does	Why the Scope Narrows
Broad empirical universe	Helium, Geodnet, Onocoy, Render, Hivemapper, io.net	Retrospective empirical stress benchmarking under observed windows	Starts with the widest public comparison set, but not every project-window supports the same degree of inference.
Empirically emphasized windows	Helium, Geodnet, Onocoy, Hivemapper (plus contextual references to Render/io.net where relevant)	Channel-specific interpretation of observed stress signatures	Narrows to the cases that best illustrate observable liquidity, yield, and cost-pressure signatures under the chapter’s metric-applicability rules.
Later frozen evaluation set	Onocoy, Helium, Hivemapper, Grass, Geodnet	Provides the smaller fixed set used later for controlled evaluation	The later evaluator can only use profiles that can be mapped cleanly into a common rule set; the broader empirical universe remains context, not the final controlled test set.

Table 8: How the comparison set narrows from a broad empirical universe to a smaller later evaluation set.

Read from top to bottom, the table shows why later chapters do not simply reuse the full empirical universe unchanged. This chapter still does empirical work. We look backward across public cases and ask which signals were visible. The funnel matters because only part of that broader universe can later be carried into a controlled evaluator without mixing unlike cases or forcing weak assumptions.

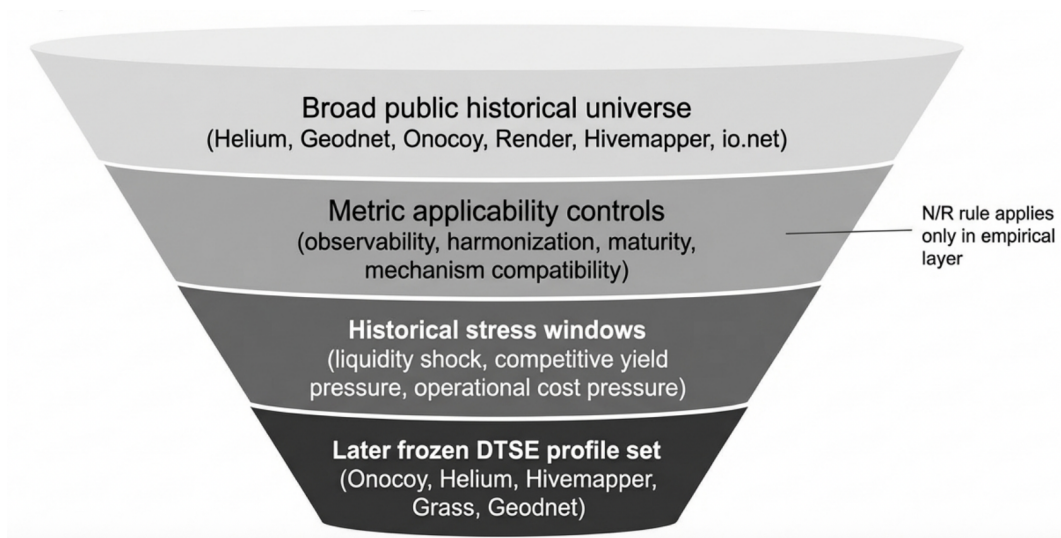


Figure 3: Empirical narrowing from broad historical comparison toward the later frozen DTSE profile set.

Figure 3 provides an at-a-glance visual of the same narrowing logic summarized in Table 8.

5.3 Metric Applicability and Observability Limits

A shared metric vocabulary helps comparison, but it also creates a risk: not every metric is observable, comparable, or interpretable across projects at different maturity stages. Some inputs exist for mature networks and disappear in early-stage ones. Some ratios look comparable on paper but rely on thin liquidity, unstable revenue denominators, or protocol-specific definitions that do not harmonize cleanly.

For that reason, this chapter does not treat missing comparability as a nuisance to smooth away. A metric is reported only when its inputs are observable at the necessary resolution, can be harmonized across projects without strong additional assumptions, and remain interpretable in the event window at hand. Where those conditions are not met, the metric is marked *N/R* (not reliably observable) and excluded from comparative ranking. The *N/R* rule is a control, not a missing-data problem to be patched through imputation. It applies only to this empirical layer and does not alter the later DTSE metric framework.

N/R reason codes

- **NR1** Required inputs are not publicly observable at the necessary resolution for the event window.
- **NR2** Native metric definitions differ materially across protocols and cannot be harmonized without strong additional assumptions.

- **NR3** Market or revenue observability is too thin or distorted in the window for stable liquidity-, valuation-, or revenue-dependent ratios.

The empirical layer uses four gates before any metric is allowed to do inferential work:

1. **Mechanism-compatibility gate:** the metric must have a valid interpretation under the underlying mechanism class.
2. **Observability gate:** required inputs must be publicly visible at sufficient resolution without imputation.
3. **Harmonization gate:** definitions must be alignable across protocols without importing strong extra assumptions.
4. **Maturity gate:** market- and revenue-dependent metrics require enough depth to avoid unstable denominators and microstructure-dominated noise.

Metrics marked *N/R* may still appear as bounded context, but they do not carry cross-project ordering in this chapter. Table 9 applies that rule across the main comparator set.

Metric	He	Ge	On	Re	Hi	Primary Constraint
Burn-to-Mint Ratio	A	A	P	A	A	Onocoy uses an adapted capped-supply burn-to-distribution analogue rather than a native BME measure.
30-Day Node Retention	A	A	P	A	A	Onocoy stress-window history remains short, so inference is directional rather than mature.
Revenue per Node	A	A	N/R	A	P	Onocoy revenue depth is immature (NR3); Hivemapper carries labor-intensity comparability caveats.
Token Turnover	A	A	N/R	A	P	Onocoy post-TGE liquidity depth is insufficient for stable inference (NR3).
FDV / Annualized Revenue	A	A	N/R	A	P	Onocoy denominator instability under early-stage revenue conditions (NR3).

Table 9: Metric applicability matrix for empirical comparison windows (He = Helium, Ge = Geodnet, On = Onocoy, Re = Render, Hi = Hivemapper; A = applicable; P = partially applicable; N/R = not reliably observable).

The Onocoy case is where this control matters most. In the current empirical layer:

- adapted burn-to-mint remains a directional burn-to-distribution proxy rather than a native BME magnitude,

- retention and churn are usable as primary participation signals, but only over a short public history,
- revenue per node, token turnover, and FDV / annualized revenue remain N/R under current revenue depth and liquidity conditions.

When these observability limits are read together with the evidence boundary already established in Chapter 4, a clear empirical ceiling appears. Historical windows can reveal direction, sensitivity, and comparability limits, but they cannot independently identify the threshold at which a newer capped-supply network would fail under matched stress.

5.4 Observed Stress Signatures in Historical Windows

The historical record is most useful when it is read as a set of realized stress signatures rather than as a scoreboard. The point is not to prove that one network class is robust and another is fragile in the abstract. The point is to see which signals became visible, which ones stayed hidden, and which kinds of comparison remain defensible once the available data are thinned by project-specific measurement conditions.

Liquidity Shock and Crypto Winter (2022–2023)

Liquidity shock refers here to an abrupt deterioration in market depth and token price support. In DePIN, that matters because provider rewards may be token-denominated while many operator costs behave more like fiat costs. A sharp market dislocation can therefore compress real reward value before the physical network visibly adjusts. The 2022–2023 crypto winter is the broad market downturn used here as a historical reference window for that kind of pressure.

Helium is a wireless DePIN and one of the best-documented public stress cases in the sector. Exchange data for the trading pair HNT/USDT, meaning Helium’s token priced against the dollar-linked stablecoin Tether, show severe price compression in this period [29, 6]. Reported hotspot counts did not collapse one-for-one with the drawdown, which is consistent with physical exit friction and sunk-cost stickiness in deployed hardware systems [1]. The useful empirical lesson is not that Helium was simply “robust.” It is that market-side dislocation can arrive well before the installed base visibly unwinds.

For Onocoy, that window matters as a bounded analog rather than as a forecast. GNSS infrastructure may also remain sticky during a drawdown even if growth, maintenance intensity, or quality buffers weaken earlier than the raw installed-base count suggests. Historical observation can reveal that kind of lag between price stress and visible network response. It cannot show where the threshold to broad failure sits for a newer network with thinner public history.

Competitive Yield Pressure (2024–2025)

Competitive yield pressure refers to the risk that operators compare rewards across adjacent networks and redirect attention, hardware, or deployment effort toward the option that offers better returns. In DePIN, this matters because the relevant comparison is not only token price. It is also whether a provider believes similar work can earn more elsewhere. That makes competitive pressure especially important where hardware profiles, geographies, or participation models overlap.

Geodnet is useful here because it is a neighboring GNSS correction network rather than a distant vertical. Its location-governance logic illustrates how a protocol can shape local supply density and yield conditions when operator competition matters [30]. Where hardware adjacency or multi-homing is feasible, provider reallocation becomes a plausible empirical stress signature even without a single market-wide crash [2, 7].

That matters for Onocoy because a provider deciding where to place or maintain GNSS infrastructure is not making that decision in isolation. If an adjacent network offers better expected returns for similar geographic coverage, similar hardware effort, or a more favorable local reward environment, the outside option starts to matter even before a large visible exit appears in Onocoy’s own public metrics. Competitive pressure can therefore affect growth quality, coverage density, and participation commitment before it shows up as a simple collapse in station counts.

The empirical comparison cannot prove the scale of realized migration inside Onocoy from public data alone, so the relevance here remains directional rather than final. What the historical window does show is why an early-stage GNSS network should not be evaluated as if its providers were insulated from neighboring yield comparisons. It identifies competitive yield as a plausible stress channel that can shape participation behavior even when the public record is still too thin to measure the full magnitude of that response.

Operational Cost Pressure (Hivemapper 2024)

Operational cost pressure appears when the cost of continued participation rises relative to the value of expected rewards. In DePIN, this matters because infrastructure does not remain online for free. Even when hardware is already deployed, operators still face recurring burdens such as electricity, connectivity, maintenance, calibration effort, or ongoing labor. If those burdens start to dominate expected reward value, participation can weaken before the network looks visibly smaller from the outside.

Hivemapper provides a useful comparator because it is a mapping network in which contribution depends on continuing active labor rather than on a relatively passive fixed

installation. Contributors gather data through ongoing driving and device use, which means time and fuel behave like a moving cost floor [7, 5]. When token-denominated reward value weakens, contribution can plateau or retreat faster than in more passive sensor networks.

The Onocoy relevance is again analogical and bounded. Onocoy does not share Hivemapper’s labor model, and this chapter should not imply that it does. The lesson is narrower: provider economics can deteriorate before broad network discontinuity becomes obvious in top-line counts. For Onocoy, whose operators still bear recurring infrastructure and quality-maintenance obligations, that makes cost pressure an important stress channel even if the exact burden differs from more labor-intensive DePIN models.

Stress Window	What Became Visible	Why It Matters for Later Evaluation
Liquidity shock and crypto winter (2022–2023)	Price compression appeared quickly, while installed-base persistence weakened more slowly.	Shows why market dislocation and visible participation decline may not occur at the same time.
Competitive yield pressure (2024–2025)	Provider reallocation pressure became plausible where adjacent networks competed for similar hardware or geography.	Shows why outside opportunity sets matter when later evaluation asks how participation responds to competing yields.
Operational cost pressure (Hivemapper 2024)	Participation became sensitive to variable effort and operating costs when reward value weakened.	Shows why provider economics can deteriorate before broad top-line network contraction becomes obvious.

Table 10: Historical stress windows and why they matter for the later controlled evaluation stage.

Across these windows, the repeated value of historical observation is diagnostic rather than predictive. Price shocks show up first in market and provider-economics fields. Competitive pressure shows up through participation elasticity rather than immediate service collapse. Cost pressure can register in provider economics before broad top-line discontinuity becomes visible. Those recurring signatures are informative, but they remain retrospective and project-specific.

5.5 Empirical Boundary and Need for DTSE

Historical windows are valuable, but they do not answer every question the thesis needs to ask. They show what happened in a specific episode, under the specific market structure, governance response, and maturity stage that happened to exist at that time. They do not let us rerun the same shock across projects on equal terms. That is why they cannot, on their own, test thresholds under matched conditions.

This point is easier to see if the terms are made explicit. A *counterfactual* question asks what would happen under a condition that did not actually occur in the historical record. A *matched condition* means applying the same stress input across different mechanism profiles on an equal basis. Historical windows rarely provide either. Helium’s 2022–2023 market shock, for example, is a real episode, but it is not the same thing as asking how a capped-supply GNSS network and a BME-oriented wireless network would respond if both faced the same shock under comparable rules and observability.

That limit becomes especially constraining for early-stage cases such as Onocoy. Several market- and revenue-dependent ratios remain N/R in the empirical layer, and the public record is still too thin to support stronger cross-project inference on those fields. In practice, that means the empirical layer can identify recurring stress signatures and clarify where comparison becomes unreliable, but it cannot yet show where Onocoy would cross from tolerable pressure into a failure-relevant threshold under a controlled shared shock.

That is why the next chapter turns to DTSE. The role of DTSE is narrow and methodological, not predictive. It provides a rule-based comparative evaluator that can apply the same stress inputs across a fixed set of profiles under explicit assumptions. In that controlled setting, we can ask threshold-style questions that the historical record cannot answer cleanly on its own, such as where a metric first departs materially from baseline or which subsystem registers stress first. The empirical layer therefore does not get replaced. It determines what should be tested, which metrics deserve caution, and how later controlled outputs should be interpreted.

6 DTSE Methodology

Historical observation cannot reveal failure thresholds before they occur. Chapter 5 showed that empirical windows remain useful for identifying how stress signatures appear, whether as price compression, cost inflation, or shifts in outside-option pressure. But once the question shifts from what happened in the past to how a different mechanism profile responds under the same adverse condition, the real-world record becomes insufficient for counterfactual comparison.

This chapter addresses that limitation by moving from empirical observation to controlled evaluation. To compare how mechanisms respond under matched pressure, the thesis uses a rule-based comparative framework: the DePIN Tokenomic Stress Evaluator (DTSE). The emphasis remains on transparent, bounded evaluation rather than on exhaustive market replication.

6.1 DTSE Purpose, Scope, and Interpretation Boundary

The purpose of DTSE is diagnostic, not predictive. It asks where stress appears first, how it propagates through incentive and participation fields, and which subsystems remain most sensitive under the stated assumptions. By applying matched inputs across a fixed set of mechanism profiles, it reports baseline-relative deviations rather than absolute performance scores or market forecasts [8, 9].

Several non-goals are equally important. DTSE does not forecast live token prices, predict the future of a given network, or rank live protocols as universally superior or inferior. It does not simulate full market microstructure, governance adaptation inside a run, or rich behavioral psychology. DTSE instead serves a narrower evaluative role for DePIN tokenomics: not as a general market replica, but as a controlled environment for comparing how different mechanism profiles respond to the same adverse conditions. Outputs are therefore interpreted as scenario-conditional patterns inside a bounded experimental design rather than as empirical proof about what a real network will do next [8, 31].

This chapter fixes that methodological setup in a fixed sequence. It first defines the DTSE architecture and the evidence layers used to interpret it. It then states the frozen scenario grid, the selected metric families, the Stage 1 signal-detection rules, the Stage 2 failure-signature logic, and the reproducibility and limit conditions that bound later claims. Settlement-layer choice remains fixed environmental context for a given profile; DTSE varies stress inputs and modeled response rules, but it does not model chain migration decisions within a run.

In plain terms, DTSE is run as a controlled comparison. For each profile, the thesis first runs a neutral 52-week baseline with stress channels turned off. It then runs one stress channel at a time under the same setup and compares that stressed path to the profile's own baseline path. Each condition is repeated 60 times under a fixed random-start schedule so the readout reflects stable patterns rather than one lucky or unlucky run. Interpretation then follows sequence: which field moves first, when that movement becomes material, and how deterioration propagates afterward.

End-to-end DTSE workflow

1. Build each protocol profile from its documented mechanism rules and the bounded modeled assumptions required for simulation control.
2. Freeze the shared experiment configuration: reporting horizon, scenario set, threshold rules, and run-repetition policy.
3. Run the neutral baseline for each profile over 52 weeks.
4. Repeat that baseline 60 times using the same fixed random-start schedule to capture distributional behavior.

5. Run each stress scenario for that same profile under the same fixed configuration, again with 60 repetitions.
6. Compare each stressed trajectory to that profile’s own baseline trajectory, week by week, for the selected metric families.
7. Apply Stage 1 logic to detect the first material baseline-relative deviation and its timing.
8. Apply Stage 2 logic to classify the resulting deterioration pattern as a failure signature, then compare those patterns across profiles.

Within this workflow, random starts vary only the stochastic draws in each replication (for example demand noise, provider heterogeneity, and stochastic exit realizations). They do not change profile rules, scenario definitions, threshold settings, or the reporting horizon.

Figure 4 summarizes the full DTSE flow from bounded inputs, through baseline and stress runs, into Stage 1 and Stage 2 readouts and cross-profile comparison.

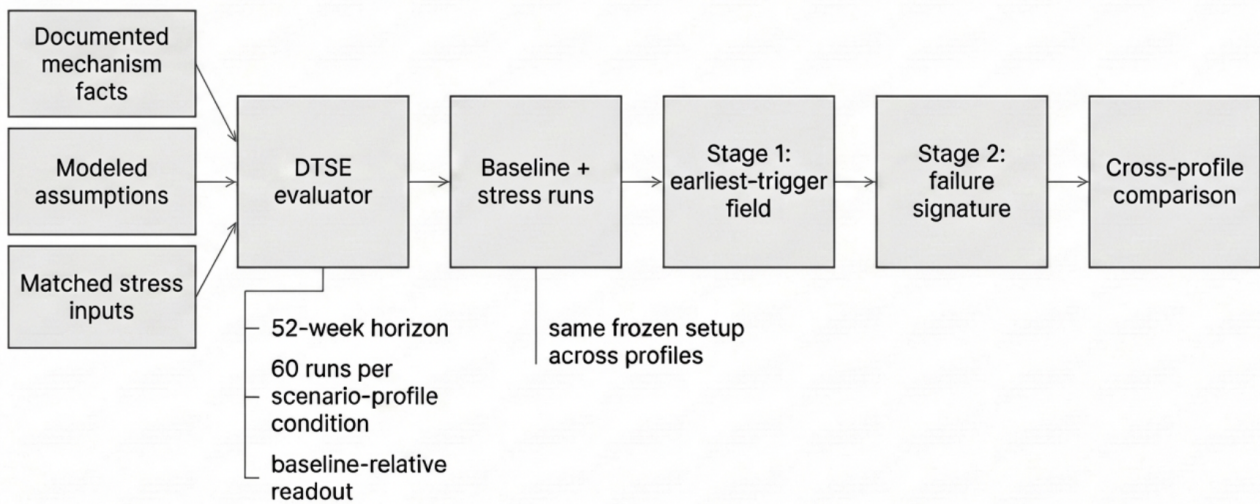


Figure 4: DTSE evaluation pipeline: bounded inputs feed a fixed evaluator, results are read baseline-relative, then interpreted through Stage 1 earliest-trigger detection and Stage 2 failure-signature classification before cross-profile comparison.

6.2 Architecture, Evidence Layers, and Core Assumptions

DTSE is implemented as a rule-based agent simulation in which heterogeneous infrastructure providers interact with fixed protocol rules under exogenous stress inputs [31, 32]. The methodology distinguishes between what is anchored in documented mechanism design, what is introduced as an explicit modeling assumption, and what is generated later as a DTSE output.

Component	What It Contains	Methodological Role
Mechanism facts	Documented rule sets, reward logic, supply path, sink structure, participation conditions	Anchor each profile in public protocol evidence rather than in ad hoc simulation behavior.
Modeled assumptions	Demand regimes, reduced-form price signal, provider-cost distributions, continuation thresholds, switching sensitivity	Supply the controlled abstractions needed where live-network evidence is incomplete or not directly observable.
Matched stress inputs	Baseline neutral, demand contraction, liquidity shock, competitive-yield pressure, provider-cost inflation	Expose every profile to the same classes of adverse pressure so later comparison remains interpretable.
DTSE evaluator	Weekly update loop with fixed profile rules, heterogeneous providers, and repeated runs	Generates the baseline and stress trajectories that later chapters compare.
Stage 1 readout	Baseline-relative signal detection	Identifies which metric family moves first and when a deviation becomes materially visible.
Stage 2 readout	Failure-signature classification	Classifies the broader pattern of deterioration across metric families after the first signal appears.

Table 11: DTSE architecture summarized as a sequence from inputs and assumptions to diagnostic readout.

Table 11 makes the separation logic explicit. Mechanism facts come from protocol documentation and anchor what a profile is allowed to do. Modeled assumptions provide the controlled abstractions needed because live DePIN systems are only partially observable. DTSE outputs are then interpreted only within that stated methodological boundary.

Evidence layers and interpretation boundary

- *Mechanism facts* are documented protocol rules such as issuance paths, sink structures, reward allocation logic, and participation conditions.
- *Modeled assumptions* are abstractions required for experimental control, including exogenous demand regimes, reduced-form price dynamics, provider-cost distributions, and decision thresholds.
- *DTSE outputs* are descriptive metric patterns produced by the frozen experiment set. They are comparative indicators under stated assumptions, not empirical measurements of live networks.

Within this architecture, *exogenous* inputs are imposed from outside the model so that they remain matched across profiles. *Endogenous* states are generated internally as outcomes of

profile rules, stress inputs, and provider decisions. Users are not modeled as individual agents. Service demand is specified as an exogenous input so that differences in outcomes can be attributed to mechanism rules and provider response rather than to different modeled user behavior [8, 9].

Providers are represented as heterogeneous agents with different cost structures, effective capacity contributions, and continuation thresholds. Their behavior is deliberately reduced-form. DTSE allows economically motivated continuation, deterioration, or exit responses, but it does not attempt to reproduce full strategic coordination, social signaling, or off-model governance negotiation [31, 32]. The modeled token-price series plays a similarly bounded role: it is an internal signal used to translate token-denominated rewards into in-model revenue and margin conditions, not an attempt to reproduce market microstructure or forecast real prices [8, 31].

The theory-to-model translation follows the thesis stress constructs established in Chapter 2. Reward Addiction is operationalized through burn-to-mint and distribution-pressure parameters. The Subsidy Gap is operationalized through provider-cost assumptions and fiat-demand baselines. Speculative Fragility is operationalized through adverse price conditions and churn sensitivity rules. For capped-supply profiles such as Onocoy, the model retains cross-profile metric names such as “emissions” for comparability, but those fields refer to modeled reward distributions from a fixed inventory path rather than to uncapped minting.

Parameterization follows a range-and-regime discipline rather than a point-estimate discipline. Public documentation and the empirical layer help determine which parameters can be anchored directly and which remain modeled assumptions. Practitioner interviews may narrow plausible ranges where documentation is incomplete, but they do not create mechanism facts and they do not substitute for empirical validation. Calibration in this chapter therefore means disciplined bounding of assumptions for comparative experiments, not recovery of true real-world parameter values.

Rather than adapting a more open-ended agent-based modeling environment, we implement DTSE as a purpose-built comparative evaluator. That narrower architecture fits the task of this thesis: holding stress inputs constant across profiles, preserving a clear distinction between documented mechanism rules and modeled assumptions, and reading outputs under the same reporting logic. The aim is not to build a general market simulator, but a disciplined instrument for comparative stress evaluation under explicit bounds.

Weekly update sequence

One DTSE time step corresponds to one week, matching the frozen reporting horizon used later in the results chapter. Each week follows the same high-level order:

1. Exogenous scenario inputs update the background conditions for that week, such as demand path, event-driven liquidity shock status, outside-yield pressure, or provider-cost multipliers.
2. Fixed protocol rules translate those conditions into reward distributions, usage-linked sink activity, and accounting-side state updates.
3. Provider economics are recalculated from the internal price signal, expected reward value, and modeled operating conditions.
4. Continuation decisions are then evaluated. Providers may remain active, continue under worsening economic conditions, or exit the active pool if modeled continuation thresholds are no longer met.
5. Aggregate outcomes such as capacity, utilization, demand satisfaction, profitability, retention, and churn are recorded for that week.

Operationally, deterioration and exit are not the same event. Deterioration means that a provider remains inside the modeled system while its economics or effective contribution weaken. Exit means that the provider no longer satisfies the continuation rule and leaves the active pool, which then becomes visible in churn and active-provider fields. Under competitive-yield pressure, switching sensitivity determines how strongly an outside opportunity changes that continuation decision, so different provider types can respond differently to the same relative-yield shock.

6.3 Scenario Grid and Experimental Setup

If stress were defined as an outcome, such as price decline, provider exit, or service deterioration, evaluation would become circular. To keep the experiment interpretable, stress in DTSE is defined as an externally imposed adverse condition applied as a model input. The question is not whether a network suffered historically, but how a fixed rule set responds when the same type of pressure is applied under matched conditions.

The experiment set contains one neutral baseline and four distinct stress channels. Table 12 summarizes the primary pressure logic without moving the full scalar configuration into the chapter body.

Scenario	Primary Pressure	Main Modeled Lever	Why It Is Included
Baseline neutral	Neutral reference path	Stress channels disabled under fixed background conditions	Provides the reference trajectory against which all later deviations are measured.
Demand contraction	Reward–demand decoupling	Exogenous service-usage deterioration with stylized volatility	Tests whether usage-linked value capture weakens before other fields recover or adjust.
Liquidity shock	Token-liquidity dislocation	Discrete unlock-style sell event under finite market depth	Tests how quickly a modeled market dislocation transmits into provider-economics fields.
Competitive yield pressure	Outside-option pressure	External-yield scalar applied through switching sensitivity	Tests whether participation reacts to adjacent-network opportunity before service demand is the first-moving field.
Provider cost inflation	Subsidy-Gap pressure	Exogenous upward shift in provider operating costs	Tests whether margin compression appears before visible contraction in participation or service capacity.

Table 12: Frozen DTSE scenario setup used for the controlled experiment set. Full scalar values are summarized in the appendix rather than repeated in the chapter body.

The scenario grid is frozen before results interpretation. Stress channels are therefore not tuned in response to later outputs. Background macro settings and demand regimes remain stylized by design, because the objective is not to reconstruct one historical market path exactly. The objective is to expose different mechanism profiles to the same classes of adverse pressure and read their baseline-relative responses under a shared experimental setup [8].

This also clarifies the role of calibration anchors. Historical windows from Chapter 5 motivate which stress channels deserve inclusion, but they do not validate the specific DTSE scenario formulas. Public protocol material anchors mechanism rules where those rules are documented clearly. Interview-derived practitioner context is used only to bound plausible ranges for operating frictions or continuation thresholds when those quantities are not directly

specified in public evidence. Full scalar grids, override values, and run identifiers are kept in the appendix rather than repeated in the main narrative.

6.4 Metric Selection, Metric Families, and Stage 1 Detection Logic

DTSE uses a deliberately selective metric suite rather than every output the simulator can generate. The main reporting families were chosen because they speak directly to the three robustness dimensions used throughout the thesis: provider retention, service continuity, and incentive–usage alignment. They also remain comparatively legible across different mechanism profiles and can be explained without forcing more precision than the model actually supports.

Not every available output is promoted into the main methodology spine. Some derived indicators are better treated as secondary or appendix-level detail because they depend heavily on case-specific commercial assumptions, unstable denominators, or project-specific business logic. Payback periods are one example. Additional summary fields may still appear in supporting exports, but the thesis centers the harmonized metric families above because they travel more cleanly across profiles and support the actual research question more directly.

Profile-specific rule sets are not forced to be structurally identical. A capped-supply GNSS profile and a BME-oriented wireless profile may use different underlying rule mappings and profile overrides. What remains fixed is the reporting horizon, the scenario grid, the metric vocabulary, and the baseline-relative comparison rule. Outputs are therefore read as deviations from each profile’s own neutral baseline under the same stress input, not as magnitude-equivalent live-network rankings.

Chapter 5 introduced an empirical *N/R* policy for historical windows. That control does not carry over into DTSE generation. In the methodology chapter, the selected metric suite is produced under controlled assumptions because the model is generating the required fields directly. The key requirement here is different: the reported metrics must be transparent enough to explain, reproducible enough to audit, and restrained enough to interpret without overclaiming.

Metric Family	Primary Robustness Dimension	What It Captures	Why It Is Kept in the Main Spine
Burn-to-Mint Ratio / Net Emissions	Incentive–usage alignment	Balance between modeled reward outflow and usage-linked value capture	Keeps subsidy pressure visible in a way that remains comparable across mechanism classes.
Provider Profitability	Provider retention	Whether token-denominated reward value still covers modeled provider economics	Shows margin compression before participation visibly contracts.
Provider Retention / Churn	Provider retention	Persistence of active supply and intensity of operator exit	Makes participation stress legible without reducing it to price alone.
Capacity Utilization / Demand Satisfaction	Service continuity	Extent to which deployed capacity continues to serve modeled demand	Shows whether service degradation appears before or after participation stress.
Volatility Proxy	Incentive translation under market stress	Dispersion in the modeled internal price signal	Keeps market-side transmission visible without claiming to model full market microstructure.
Incentive Efficiency	Incentive–usage alignment / service continuity	Modeled reward cost per unit of validated capacity	Shows whether maintaining service becomes materially more incentive-intensive under stress.
Velocity Proxy	Secondary market-side context	Turnover-style activity signal under the reduced-form price layer	Retained only as bounded support because it is less central and more sensitive to abstraction choices.

Table 13: Selected DTSE metric families and the robustness dimensions they are meant to illuminate.

The main computations are introduced descriptively before they are stated symbolically. Here, t denotes a week and \mathcal{T} the reporting horizon.

Metric / Rule	Metric Description	Compact Definition
Burn-to-Mint Ratio	Tokens burned in week t divided by tokens distributed in week t ; it asks whether usage-linked value capture is keeping pace with reward outflow.	$Burn\text{-}to\text{-}Mint_t \approx Burns_t / Emissions_t$
Net Emissions	Tokens distributed in week t minus tokens burned in week t ; it shows whether the system is adding or absorbing net token pressure in that period.	$Net\ Emissions_t = Emissions_t - Burns_t$
Provider Profitability	Average reward value minus modeled operating cost for active providers in week t ; it shows whether continued participation remains economically viable.	weekly provider margin signal
Retention / Churn	Share of providers still active and the flow of providers exiting relative to baseline; these show how economic stress becomes participation stress.	active-provider share and exit events over time
Capacity Utilization / Demand Satisfaction	Used capacity relative to available capacity, and served demand relative to modeled demand; these show whether service continuity is weakening.	used capacity / available capacity; served demand / total demand
First-Signal Timing	First week in which a selected metric departs materially from its baseline path under the fixed threshold for that metric family.	first t where deviation from baseline exceeds θ_m
Incentive Efficiency	Cumulative reward outlay over \mathcal{T} divided by cumulative validated capacity over \mathcal{T} ; it shows how incentive-intensive service becomes under a scenario.	cumulative reward outlay / cumulative validated capacity

Table 14: Descriptive definitions of the main DTSE metric families and the detection rule used in Chapter 7.

Stage 1 detection uses these equations to identify the first material departure from baseline for each metric family. The threshold rule is fixed ex ante by metric class and does not change across profiles inside the frozen experiment set. Stage 1 therefore answers a narrow question: *which field moved first, and when did that departure become material under the*

stated threshold rule? It does not yet classify the broader pattern of deterioration. That second task belongs to Stage 2.

6.5 Stage 2 Failure-Signature Logic

Stage 2 converts baseline-relative deviations into a diagnostic vocabulary. The point is not to declare a network simply stable or unstable. The point is to classify how stress is propagating across metric families once Stage 1 has identified where material deviation begins. Stage 2 therefore uses patterned deterioration across multiple fields and their time order, rather than a single scalar threshold, to define a failure signature.

Failure Signature	Operational Rule	What It Diagnoses
Reward–Demand Decoupling	Usage-linked service fields weaken versus baseline while modeled reward distribution remains active and the incentive-solvency proxy deteriorates.	A breakdown in incentive–usage alignment before the mechanism has stopped distributing rewards.
Profitability-Induced Churn	Provider-profit fields remain negative for a sustained window and are followed by rising churn or active-provider deterioration versus baseline.	Margin compression severe enough to turn economic pressure into participation stress.
Liquidity-Driven Compression	A discrete price dislocation triggers material price deviation and near-window deterioration in provider-economics or churn fields.	Fast transmission from market-side dislocation into the real incentive environment facing providers.
Elastic Provider Exit	Under outside-opportunity pressure, churn and active-provider deterioration trigger before utilization becomes the first-moving field.	Supply-side mobility under adjacent-network opportunity rather than immediate demand collapse.
Latent Capacity Degradation	Participation erodes first while utilization or demand-satisfaction fields deteriorate later as redundancy is consumed.	Hidden resilience loss beneath still-functional top-line service conditions.

Table 15: Stage 2 failure-signature vocabulary used to classify patterned DTSE deterioration. Exact scalar cutoffs are listed in the appendix rather than in the chapter body.

These signatures are formal methodological categories, not results in themselves. The results chapter later instantiates them under the frozen scenario grid. Here the chapter fixes only the rules of use: Stage 1 identifies earliest-trigger fields and timing; Stage 2 classifies the larger deterioration pattern that follows. The two stages are related but not interchangeable. A field may trigger early in Stage 1 without by itself determining the full Stage 2 signature, and

some trajectories may satisfy more than one diagnostic condition before a dominant pattern is assigned.

The benefit of Stage 2 is interpretive discipline. Instead of forcing every trajectory into a single robustness score, DTSE can describe whether a scenario primarily expresses reward–demand misalignment, margin-driven churn, liquidity-driven compression, elastic exit under outside opportunity, or delayed service loss after participation has already eroded. That vocabulary is what later allows the thesis to compare mechanism response paths without overstating numerical precision or pretending that one scalar captures the whole stress process.

6.6 Reproducibility, Calibration Boundary, and Method Limits

Reproducibility is treated as a core requirement rather than an afterthought. Each scenario–profile condition is defined by explicit parameterization, a fixed scenario identity, and a deterministic seed schedule. Because DTSE includes stochastic elements, such as demand variation, provider heterogeneity sampling, decision noise, and price-related stochasticity, a single illustrative trajectory would be misleading.

To address this, the methodology relies on repetition. Each scenario–profile condition is evaluated over repeated runs so the thesis can report medians and dispersion bands while keeping the experimental setup fixed. The run-count and seed policy are documented in the appendix materials so that the resulting distributions can be independently reproduced.

Verification in this chapter means internal validity under stated assumptions, not empirical validation against live-network history. Baseline runs are checked for boundedness and invariants, isolated perturbations are used to confirm directional response, and sensitivity passes are used on key parameter families such as provider-cost ranges, churn thresholds, and distribution controls [31, 9]. Those checks are designed to show that the evaluator behaves coherently under its own rules. They do not prove that the model reproduces a real network.

That calibration boundary is essential. Chapter 5 showed which historical windows motivate the chosen stress channels, and Chapter 4 established which Onocoy mechanism facts are documented well enough to anchor profile rules. Neither step turns DTSE into an empirically validated replica of a live protocol. Public documentation anchors rule mappings. Historical observation motivates which pressures deserve testing. Interview material narrows some plausible assumption bounds. The simulator remains a comparative evaluator built on explicit abstractions.

The main limits follow directly from those abstractions:

1. **Demand abstraction:** demand regimes are exogenous stylizations rather than endogenous adoption models.

2. **Price-process abstraction:** the internal price signal transmits stress into provider economics but does not reproduce detailed market microstructure or strategic trading.
3. **Behavioral simplification:** provider decisions are reduced-form and do not model rich coordination or adaptive governance behavior inside runs.
4. **Governance and verification scope:** governance intervention, adversarial behavior, and detailed cyber-physical verification layers remain out of scope unless explicitly modeled.

Accordingly, DTSE outputs support comparative statements about sensitivity, signal timing, and failure-signature patterning within the frozen experiment set. They do not support causal proof, prevalence claims about live networks, or universal mechanism rankings. The next chapter therefore reports outputs under this fixed methodological setup: baseline-relative, scenario-conditional, and bounded by the assumptions made explicit here.

7 Simulation Results

Chapter 7 reports the comparative DTSE results under the frozen scenario grid defined in Chapter 6. The focus is on how the Onocoy anchor case and alternative mechanism profiles respond to matched stress inputs, and on the order in which those responses become visible across price, incentive, participation, and service fields.

7.1 Reporting Conventions and Interpretation Boundaries

This chapter reports the DePIN Tokenomic Stress Evaluator (DTSE) outputs for the stress scenarios defined in Section 6.3. Results are descriptive and comparative: each scenario evaluates how the Onocoy (ONO) anchor case and comparator DePIN mechanism profiles respond to matched stress inputs. As established in the methodology, these outputs represent model-conditional behavior under frozen assumptions; they are not empirical measurements, market forecasts, or absolute claims of real-world protocol superiority.

Chapter 5 also established that several ONO-relevant market and revenue fields remain not reliably observable in historical documentation alone. This results chapter therefore uses DTSE outputs to examine those fields under matched conditions rather than treating retrospective empirical observation as sufficient on its own.

Because different DePINs operate on different numeric scales, raw magnitude can mislead. The more informative question is not only what breaks, but when a field breaks relative to its own baseline path. All results in this chapter are therefore governed by a strict *Sequence Over Magnitude* reading rule. The analytical focus is on timing: which metric family registers stress

first, how quickly that deviation propagates across subsystems, and whether service continuity degrades before or after participation erodes.

To operationalize this, scenario outcomes are evaluated entirely as deviations from their corresponding neutral baseline trajectories. Crucially, baseline trajectories can naturally drift over the 52-week reporting horizon due to deterministic mechanism rules. Stress is therefore defined and measured as a material acceleration, reversal, or threshold crossing relative to that drifting reference path, rather than against a static equilibrium state.

Signal timing follows the two-stage detection logic. Stage 1 identifies the earliest-trigger metric family, and Stage 2 classifies the subsequent pattern of deterioration into predefined operational failure-mode signatures. Governance adaptation, discretionary intervention, and other unmodeled human responses remain outside the evidentiary scope of this chapter and are taken up later in the discussion.

7.2 Baseline Trajectories (The Neutral Reference)

Baseline runs establish the neutral reference trajectories for core DTSE state variables (e.g., modeled price signal, token supply and emissions/burns, provider participation, and utilization) in the absence of scenario-driven stress inputs. They serve as comparison anchors for this chapter and are not interpreted as static equilibrium states.

All baseline experiments use the standard DTSE configuration (as documented in Section 6.2). In this neutral setup, we disable all stress scenario logic and hold exogenous inputs steady. This ensures that any baseline dynamics arise purely from the organic interaction between each protocol's tokenomic rule set and the modeled provider economics.

Baseline drift matters for interpretation. In DTSE, the baseline is the neutral reference path generated by the frozen rule set under neutral inputs; it is not a claim of long-term stability. Stress results are therefore strictly compared against that drifting reference path rather than against a flat line.

In the ONO baseline, for example, modeled price and active provider counts naturally decline over the 52-week horizon due to expected token emission decay. Under stress, the diagnostic question is not whether these metrics fall, but whether they fall faster, reverse earlier, or cross critical thresholds sooner than they would have under neutral conditions. This is how the chapter reads every subsequent scenario: by measuring when and where a stress input pulls the system away from its own baseline trajectory, and which metric family registers the departure first.

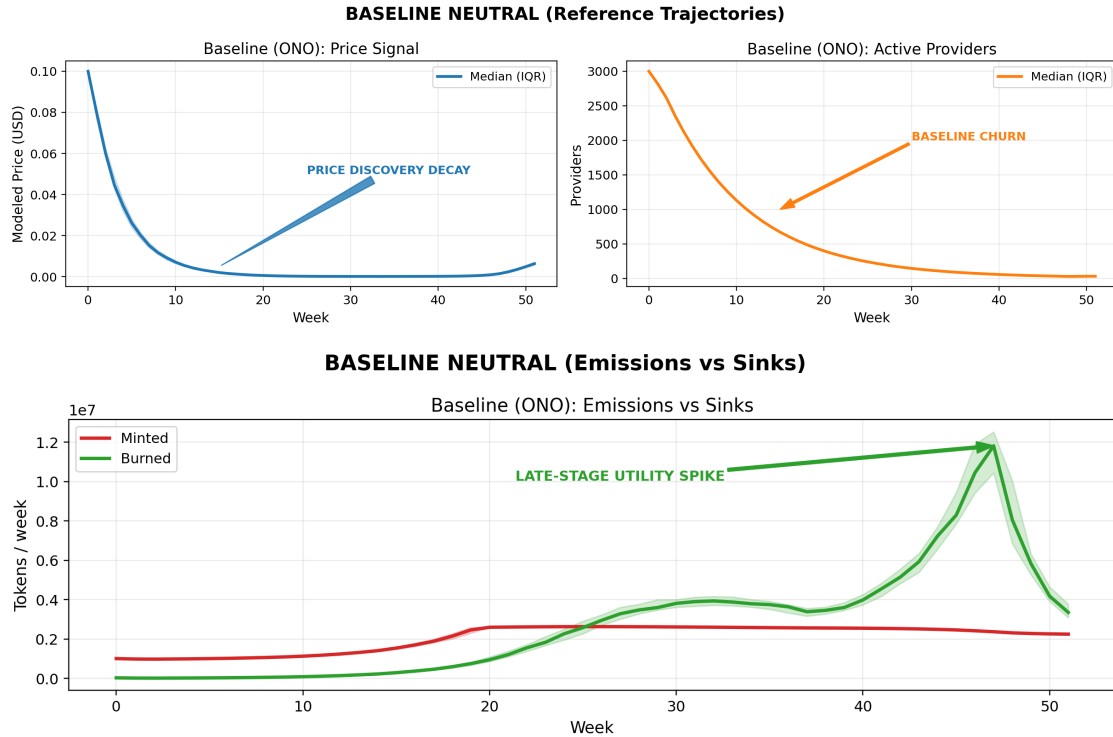


Figure 5: DTSE ONO baseline reference trajectories under neutral inputs: modeled price and active provider count (top) and emissions versus sinks (bottom). These curves provide the drifting reference path against which later stress deviations are read.

7.3 Cross-Scenario Detection Pattern Summary

Across the DTSE experiment set, each stress channel produces a distinct sequence of metric changes when read as baseline deviations. The purpose of this section is to summarize those recurring response patterns before the chapter turns to the ONO-specific scenario figures. Table 16 maps each modeled stress channel to its earliest signal, its participation-side follow-on field, and the Stage 2 signature observed under that response pattern.

7.4 Onocoy (ONO) Baseline-Relative Signal Sequencing

Table 17 makes the same sequencing logic concrete for the Onocoy (ONO) profile under the frozen stress scenarios. The table is read strictly against the ONO baseline: it shows which subsystems move first, which fields follow, and which Stage 2 signature is observed under each channel without claiming exact live-network thresholds or token-price outcomes.

Stress Channel (Modeled Input)	Earliest Signal	Participation-Side Follow-On	Observed Stage 2 Signature
Demand Contraction:			
Exogenous service-usage reduction under a specified demand regime.	Utilization; usage-linked sinks.	Active provider count deteriorates materially after usage-linked weakening.	Reward–Demand Decoupling
Liquidity Shock:			
Discrete price dislocation via modeled finite-depth liquidity channel.	Modeled price signal.	Provider profitability and provider churn deteriorate around the event window.	Liquidity-Driven Compression
Competitive Yield Pressure:			
Exogenous increase in outside-opportunity attractiveness.	Provider churn rate.	Active provider count weakens before utilization becomes the first-moving field.	Elastic Provider Exit
Provider Cost Inflation:			
Exogenous upward shift in operator hardware/OpEx costs.	Provider profitability.	Active provider count and provider churn weaken after sustained margin compression.	Profitability-Induced Churn & Latent Capacity Degradation

Table 16: Cross-scenario detection pattern summary under the frozen DTSE setup. The table reports how each modeled stress channel first appears and which Stage 2 signature is observed once deterioration propagates across subsystems.

7.5 Scenario-by-Scenario Deviations

The preceding tables summarized each stress channel’s detection sequence in compact form. This section provides the detailed simulation evidence behind those summaries, walking through the observed scenario deviations for the ONO baseline one channel at a time. The focus remains on which metric families register first and how the initial deviation propagates through later fields.

Demand Contraction (Reward–Demand Decoupling)

Demand contraction isolates a scenario where real-world usage weakens while reward logic remains active. The diagnostic question is which subsystem registers first under baseline-relative detection, and whether participation effects appear only after profitability and provider churn thresholds begin to bind.

Under this channel, utilization and usage-linked sink fields deviate from the baseline first. However, the active provider count does not meaningfully break from the baseline trajectory

Stress Channel	First-Moving Field (Stage 1 Detection)	Secondary Detection Field (Stage 2)	Diagnosed Signature for ONO
Demand Contraction	Utilization and usage-linked sinks deviate first.	Active provider count deviates materially later.	Reward–Demand Decoupling.
Liquidity Shock	Modeled price drawdown triggers immediately.	Provider profitability and provider churn break following the price event.	Liquidity-Driven Compression.
Competitive Yield Pressure	Provider churn rate and active provider count deviate immediately.	Utilization remains stable initially, lagging behind participation drop.	Elastic Provider Exit.
Provider Cost Inflation	Provider profitability turns structurally negative.	Active provider count and provider churn degrade following margin collapse.	Profitability-Induced Churn & Latent Capacity Degradation.

Table 17: Onocoy (ONO) baseline-relative signal sequencing under the frozen stress scenarios. The table summarizes which fields move first, which fields follow, and which Stage 2 signature is observed in each case.

until detectably later in the simulation horizon. The network continues its scheduled reward distribution while users are no longer buying sufficient Data Credits to offset the outflow.

Taken together, the ONO demand-contraction run shows a clear *Reward–Demand Decoupling* pattern. Usage-linked fields weaken first, while provider participation remains comparatively sticky and only deteriorates materially much later. The result is not immediate network discontinuity, but a widening gap between ongoing reward distribution and weakening service-side demand.




Figure 6: DTSE ONO demand-contraction deviations from baseline: utilization, sinks, and provider profit. The channel breaks first in usage-linked fields, with participation effects arriving materially later.

Liquidity Shock (Liquidity-Driven Compression)

The liquidity shock scenario isolates the transmission of a discrete token-liquidity dislocation into provider-side stress fields. The simulation applies a discrete token sell-off directly against a finite-depth liquidity channel, creating an immediate price drop without an accompanying initial collapse in real-world service demand.

When the configured token unlock event hits the market, the abrupt price dislocation immediately compresses the fiat-equivalent reward value for providers. Within a short window, provider profitability falls below zero and provider churn rises sharply. These signals appear together even though the scenario does not begin with a simultaneous collapse in service demand.

In the ONO liquidity-shock run, the decisive feature is the short transmission distance between the price event and provider-side deterioration. The modeled price dislocation compresses fiat-equivalent reward value around the event window, and profitability and churn respond almost immediately thereafter. The observed Stage 2 pattern is therefore *Liquidity-Driven Compression* rather than a demand-led contraction.




Figure 7: DTSE ONO liquidity-shock transmission: modeled price dislocation and provider churn response. The figure highlights how a financial shock is transmitted into participation-side stress around the event window.

Competitive Yield Pressure (Elastic Provider Exit)

This channel isolates provider sensitivity to outside opportunities without requiring internal demand collapse or a discrete internal liquidity event. In this scenario, operators suddenly gain access to a much more attractive alternative for their hardware or capital, parameterized to reflect competing yield networks like adjacent hardware protocols or re-purposed GNSS arrays.

As soon as the outside yield pressure increases, participation fields register immediate deviations. Provider churn rises, and the active provider count drops sharply in the earliest stages of the simulation. During this initial period, the internal network utilization metrics remain comparatively stable, so the first visible change is on the supply side rather than the service side.

For ONO, competitive-yield pressure first appears as a participation-side response to a more attractive outside option. Churn and active-provider counts deviate before utilization becomes the first-moving field, indicating that operator withdrawal can begin while service-

side conditions still appear comparatively stable. The observed Stage 2 pattern is therefore *Elastic Provider Exit*.




Figure 8: DTSE ONO competitive-yield pressure response: active provider count and provider churn deviations versus baseline. Participation moves before visible service-side deterioration when outside yield becomes more attractive.

Provider Cost Inflation (Latent Capacity Degradation)

Provider cost inflation tests exposure to real-world operating environments (e.g., electricity prices, connectivity costs, facility rents). The scenario applies an exogenous upward shift to the physical operating-cost conditions of the providers.

This channel immediately triggers the *Profitability-Induced Churn* signature: rising fiat-equivalent costs compress provider margins into negative territory immediately. However, the simulation also shows that the active provider count does not synchronously collapse. Active nodes begin to deteriorate only after sustained windows of unprofitability.

For ONO, provider-cost inflation first appears as margin compression and only later as visible participation loss. The run therefore shows a two-step deterioration pattern: profitability weakens immediately, while active-provider deterioration follows after sustained unprofitability. In Stage 2 terms, the scenario expresses *Profitability-Induced Churn* together with *Latent Capacity Degradation*.




Figure 9: DTSE ONO provider-cost inflation response: margin compression and subsequent active provider count dynamics. The figure shows profitability deteriorating before participation metrics visibly break from baseline.

Stress Channel	Earliest Material Deviation	Observed Propagation Pattern	Observed Stage 2 Signature
Demand Contraction	Utilization and usage-linked sinks weaken before provider participation does.	Reward distribution remains active while service-side demand weakens; participation only deteriorates later.	Reward–Demand Decoupling
Liquidity Shock	Modeled price dislocation appears around the event window.	Price compression is followed quickly by weaker provider economics and higher churn.	Liquidity-Driven Compression
Competitive Yield Pressure	Participation fields break before utilization becomes the first-moving field.	Outside-option pressure appears first as churn and active-provider deterioration while service-side measures remain comparatively stable.	Elastic Provider Exit
Provider Cost Inflation	Provider profitability turns negative before participation breaks materially.	Margin compression persists first, then propagates into provider churn and active-provider deterioration.	Profitability-Induced Churn + Latent Capacity Degradation

Table 18: Observed Stage 2 summary for the ONO profile under the frozen DTSE scenarios. The table closes the scenario readouts by showing how earliest deviations propagate into diagnosed failure signatures.

7.6 Cross-Profile and Cross-Scenario Synthesis

Having diagnosed the operational failure signatures for each stress channel individually, this section synthesizes the results across all four scenarios and compares them against alternative mechanism profiles.

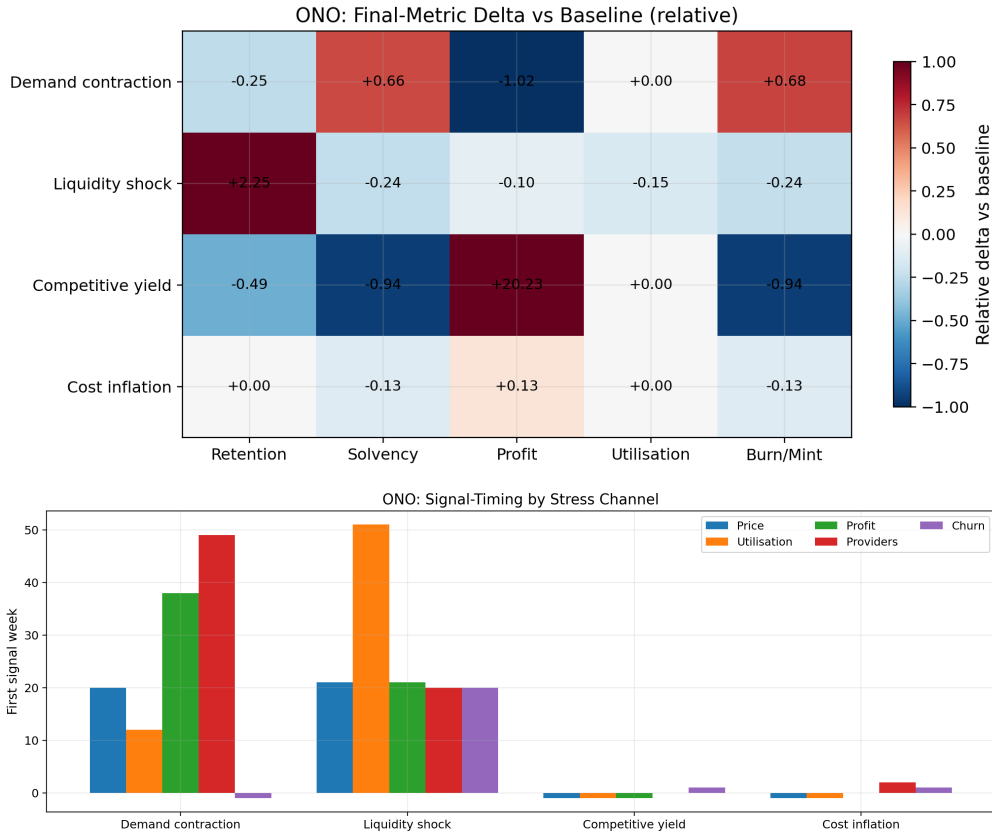


Figure 10: DTSE ONO cross-scenario synthesis: final-metric delta heatmap (top) and first-signal timing by stress channel (bottom). Together, the panels demonstrate both where stress concentrates operationally and how quickly each channel becomes visible in the pre-specified metric suite.

Cross-Scenario Observation: Figure 10 makes the sequence-over-magnitude pattern visible across the ONO stress set. The heatmap and timing charts show that stress rarely strikes the physical participation layer first. Across nearly all channels, financial indicators (profitability) or usage indicators (utilization) act as the early-warning layer, while active provider counts tend to move later.

Across the experiment set, stress channels manifest in distinct metric families and timing windows (Figures 11 and 10). Demand-side stress is detected earlier in utilization-linked channels. Liquidity-shock signals appear in price, provider-economics, and provider-churn fields around the event window. Cost and yield channels are first detected in provider-economics and provider-churn fields.




Figure 11: DTSE cross-profile final-week comparison by scenario: provider retention (top) and incentive-solvency proxy (bottom). These outputs must be read directionally under matched stress inputs rather than as magnitude-equivalent live-protocol rankings.

Cross-Profile Observation: Figure 11 shows that mechanism profiles channel matched stress differently. Under identical liquidity and yield shocks, BME-oriented profiles exhibit larger swings in their incentive-solvency proxies. In contrast, the capped-supply ONO profile absorbs the same pressures through a more bounded proxy path, with stress appearing more clearly in provider margin compression.

Incentive-Solvency Proxy Caveat: The proxy used in Figure 11 serves as a comparative stress signal rather than a standalone economic health score. For capped-supply profiles like ONO, the “mint” denominator represents scheduled reward distributions from fixed inventory rather than uncapped token creation. A larger proxy swing (such as those seen in BME comparators) therefore indicates a different structural transmission path under stress, not a categorically “worse” absolute outcome.

Because the proxy behaves like a ratio, it can move sharply when one side deteriorates faster than the other. Accordingly, the proxy bars in Figure 11 are interpreted jointly with profitability, provider churn, provider retention, and utilization rather than in isolation.

In summary, the cross-scenario synthesis supports a sequential reading of DePIN stress under the frozen DTSE setup. Financial, usage-linked, and provider-economics fields absorb the initial shock first, while physical participation metrics often move later. These observed sequences provide the basis for the discussion that follows.

8 Discussion and Conclusion

In this chapter, we interpret the DTSE results reported in Chapter 7 without reopening the methodology or adding new outputs. Our central question is not whether one mechanism class can be declared simply robust or fragile. We ask how stress moves through DePIN systems, which subsystems register deterioration first, and what those sequences imply for the design and governance of hardware-dependent tokenized networks.

That interpretive move matters because the results chapter was intentionally descriptive. It established earliest-trigger fields, baseline-relative deviations, and Stage 2 failure signatures under a frozen scenario grid. Here, we ask what those observed sequences mean for DePIN robustness, how they can be translated into response categories without overstating what DTSE can claim, and where our conclusions must remain bounded by evidence and abstraction.

8.1 Discussion Orientation

The results support a comparative reading of robustness rather than a binary one. Stress did not appear in a single universal order across the experiment set, and no single metric captured the full deterioration process. Instead, different channels expressed themselves first in different places: usage-linked fields under demand contraction, price and provider-economics fields under liquidity shock, participation fields under competitive-yield pressure, and margin fields under provider-cost inflation. This is why we have treated robustness as a question of transmission path rather than as a one-shot stability label.

The same point also clarifies why the Stage 1 and Stage 2 structure was worth preserving. Stage 1 timing identifies where deterioration first becomes materially visible. Stage 2 shows what broader pattern that early movement turns into once stress propagates across the rest of the system. Taken together, those two readings let us discuss DePIN robustness as a sequence: first signal, propagation path, and resulting failure signature. That is a more informative basis for interpretation than asking only whether a protocol still appears operational at the end of a run.

8.2 What the Results Mean for DePIN Robustness

Across the DTSE results, we understand DePIN robustness as the management of coupled pressures across provider retention, service continuity, and incentive–usage alignment. Those three dimensions were introduced in the foundations chapters as the thesis lens, and the results chapter showed that stress rarely enters all three at once. Some channels first strike usage-linked fields, others market-facing or provider-facing fields, and only later do the effects

spread outward into the rest of the system. Figures 10 and 11, together with Tables 16 and 18, therefore support an interpretation centered on sequencing rather than on end-state labels.

The DePIN Illusion

A central interpretive finding of this thesis is the deceptive persistence of installed capacity under stress. In Chapter 2, we argued that DePIN participation is shaped by sunk costs, deployment friction, and path dependence in a way that software-only crypto participation is not. Once operators have acquired, installed, and calibrated site-specific hardware, participation becomes economically sticky. That does not make the network immune to stress. It means that deterioration can accumulate inside provider economics and incentive alignment before the physical participation layer visibly contracts.

In Chapter 5, we saw the empirical version of the same problem. In the Helium stress window, severe price compression did not produce a one-for-one collapse in the installed base, even though the underlying economic environment had clearly weakened. That historical persistence matters because it can be read too quickly as proof of resilience. A more cautious interpretation is that physical exit friction can delay visible participation decline, allowing top-line counts to remain comparatively stable while the underlying tokenomic environment is already under strain.

The DTSE results reproduce that same sequence under controlled assumptions. In the ONO demand-contraction run, utilization and usage-linked sinks weaken before active-provider counts deteriorate materially. Under provider-cost inflation, provider margins compress first and participation loss follows later. Even under liquidity shock, the initial dislocation appears first in the modeled price layer and then propagates into provider-economics and churn fields. Taken together, these results suggest that active node or provider counts can function as a lagging indicator of DePIN stress rather than as an early proof of underlying robustness.

The cross-profile comparison also sharpens the thesis answer. BME-oriented and capped-supply profiles do not escape DePIN's physical constraints; they distribute stress differently. BME-oriented profiles can show sharper movement in incentive-solvency proxies under matched liquidity and yield shocks, while the capped-supply ONO profile keeps that proxy more bounded but shifts the stress burden more visibly into provider margin compression and usage-linked alignment conditions. In other words, the relevant difference is not that one class is free of vulnerability. It is that the location and timing of vulnerability change with the mechanism design.

For the ONO anchor case, that trade-off is especially clear. The dual ONO/Data-Credit architecture reduces direct user-price exposure to token volatility, which helps explain why user-facing service demand is not the first field to move under every channel. At the same time,

the same design places more interpretive weight on whether usage-linked sinks remain strong enough and whether provider economics remain viable as stress accumulates. ONO therefore illustrates a broader point developed across the thesis: insulating one interface of a DePIN economy can still leave other parts of the system, especially provider margins and participation persistence, as the main route through which stress becomes operationally significant.

For the Onocoy anchor case, these results also imply a specific interpretive caution. Apparent stability in top-line participation should not be read too quickly as evidence of underlying robustness, because provider margins, usage-linked sinks, and competitive-yield pressures can weaken before node or provider counts visibly contract. The ONO/Data-Credit design reduces one channel of direct user-price volatility, but it does not remove the need to preserve provider viability and demand-linked support under stress. In that sense, the central issue is not whether the network can avoid pressure altogether, but whether stress is detected early enough in the fields that move first.

8.3 Comparative Implications for DePIN Design

Within the bounded comparator set used in this thesis, the DTSE results also provide design implications beyond the Onocoy anchor case. Because the evaluator applies matched stress inputs across several mechanism profiles, the final comparison yields design implications for DePIN systems that differ in issuance logic, reward routing, and provider-cost structure.

For BME-oriented profiles, the main implication is that market-facing stress can transmit rapidly into incentive-solvency conditions when token issuance remains closely tied to reward support under adverse price conditions. In the DTSE comparison, these profiles showed larger proxy swings under matched liquidity and yield shocks than the capped-supply ONO profile. That pattern does not establish that BME-oriented systems are categorically weaker; within this experiment set, it suggests that vulnerability can concentrate more directly in the market-facing layer, which makes state-sensitive reward adjustment and issuance logic especially important under stress.

A different implication follows for networks whose participation depends on higher recurring effort or variable operating burden. Where provider economics are strongly exposed to fuel, time, electricity, maintenance, or other continuing costs, the Subsidy Gap can widen quickly when token-denominated reward value weakens. In those settings, provider profitability becomes a more immediate early-warning field than top-line participation. The central lesson is not that all such networks fail in the same way, but that recurring real-world operating burden can compress participation more quickly than installed hardware counts alone would suggest.

The comparative results also matter for networks exposed to adjacent-yield competition.

When hardware or operating setups can be repurposed across nearby protocols, outside-option pressure can pull participation away before internal demand becomes the first-moving field. In the DTSE comparison, this appeared as Elastic Provider Exit: participation-sensitive metrics moved early under competitive-yield pressure even when service utilization was not yet the first field to deteriorate. For DePIN systems of that kind, robustness depends not only on internal reward logic, but also on how exposed providers remain to a more attractive external yield environment.

Taken together, these implications support a broader reading of DePIN robustness. In hardware-dependent networks, end-state node counts and static market snapshots are not sufficient as primary diagnostics. A more useful perspective asks which fields move first, how quickly stress crosses from token-side pressure into provider economics and service quality, and whether physical participation is masking deterioration that has already begun elsewhere in the system.

8.4 Governance and Intervention Archetypes

The accepted results also make it possible for us to discuss governance and intervention without pretending that DTSE modeled governance directly. The archetypes below are interpretive response categories derived from the observed failure signatures. They are not DTSE outputs, and they are not claims about how frequently real networks adopt any particular response. Their value is narrower: they help organize what kinds of intervention become thinkable once a given stress pattern is visible.

Two clarifications matter here. First, the archetypes are anchored in the accepted results, not in speculative intent labels. When the ONO runs show margin compression preceding visible participation loss, the relevant discussion category is not a claim about what governance *will* do. It is the narrower observation that any response would have to confront provider economics before service continuity visibly collapses. Second, not every response class is equally representable inside DTSE. Rule changes and exogenous-input changes can be modeled when translated into explicit parameters, whereas informal coordination, signaling, or governance latency remain interpretive context unless they are turned into modeled changes.

That boundary keeps the discussion useful without letting it outrun the evidence. The archetypes do not claim that real DePIN governance is uniform, nor do they claim that one response class is universally preferable. They instead offer a disciplined vocabulary for moving from observed stress signatures to response categories: subsidy continuation where usage weakens first, incentive expansion or re-targeting where provider retention becomes the first pressure point, coordinated operational adjustment where exogenous cost or demand conditions dominate, and non-structural response where visible deterioration outpaces modeled

Archetype	Observed Result Pattern That Motivates It	Typical Response Focus	Main Boundary or Trade-Off
Subsidy continuation	Reward–Demand Decoupling: usage-linked fields weaken while reward distribution remains active.	Maintain existing incentive rules while hoping usage or demand recovers.	Can preserve participation temporarily while worsening incentive–usage misalignment if sink strength does not recover.
Broad incentive increase	Profitability-Induced Churn or Elastic Provider Exit: provider economics and participation fields weaken first.	Raise aggregate incentive intensity to stabilize provider retention.	Can soften participation loss in the short run, but may intensify subsidy pressure and weaken budget discipline.
Incentive re-targeting	Participation stress appears before service failure, especially under outside-option pressure.	Redirect rewards toward higher-commitment or higher-value capacity rather than raising total reward magnitude.	Depends on rule changes, verification quality, and implementation capacity that DTSE does not model directly [21, 7].
Coordinated operational intervention	Demand-side or cost-side channels register first in exogenous operating conditions.	Change operating conditions outside the token rule set, such as cost support, partner coordination, or demand-side support.	Becomes representable in DTSE only when operational changes can be translated into exogenous-input adjustments.
Non-structural response	Stress is visible, but no modeled rule or input change accompanies it.	Communication, signaling, or informal coordination without parameter change.	May matter in practice, but remains outside DTSE unless it alters modeled rules, thresholds, or exogenous inputs.

Table 19: Governance and intervention archetypes derived from the observed DTSE failure signatures. These are discussion-level response categories, not DTSE outputs or empirical prevalence claims.

intervention.

8.5 Limitations and Claim Boundaries

The interpretive strength of Chapter 7 depends on staying clear about what DTSE can and cannot establish. External validity is bounded by design. We use exogenous demand paths, a reduced-form internal price signal, simplified provider decision rules, and fixed scenario identities in the experiment set. Those choices are appropriate for a comparative evaluator, but they do not reproduce full market microstructure, endogenous adoption, rich governance behavior, or the entire cyber-physical verification layer of a live DePIN network [8, 9, 31].

The Onocoy anchor introduces an additional evidence boundary. Chapter 4 established mechanism facts around ONO, Data Credits, reward logic, and public observability, while Chapters 5 and 7 showed why several market- and revenue-dependent fields remain only partially observable in historical documentation. DTSE therefore helps interpret bounded

stress behavior under explicit assumptions, but it does not turn those assumptions into empirical proof. That distinction is especially important for claims about thresholds, governance efficacy, and generalization beyond the documented experiment set.

Limitation	What It Constrains	Interpretive Consequence
Exogenous demand stylization	Claims about endogenous adoption, user formation, and real-world demand elasticity.	Results can compare how profiles react to the same demand shock, but they cannot establish how demand would form or recover outside the modeled scenarios.
Reduced-form price process	Claims about detailed market microstructure, strategic trading, and institutionally mediated liquidity.	Liquidity-shock results show transmission into provider economics under the modeled price layer, not a full reconstruction of real token markets.
Simplified provider decision rules	Claims about coordination, multi-homing, and richer strategic adaptation.	Participation-side findings should be read as directional sensitivity under stated rules rather than as exhaustive behavioral prediction.
Governance and verification out of scope	Claims about intervention efficacy, adversarial resilience, and cyber-physical validation layers.	Governance discussion can interpret response categories, but DTSE alone cannot prove that a given intervention would succeed in practice.

Table 20: Limitation-to-claim map for the discussion chapter. The table shows which parts of the interpretation remain bounded by DTSE abstraction rather than by live-network validation.

Two further reading rules help keep those limits explicit. First, baseline trajectories drift even under neutral inputs, so the thesis does not classify profiles as simply stable or unstable. It asks whether a stress channel accelerates, reverses, or materially worsens a profile relative to its own neutral path. Second, cross-profile metrics do not share one universal numeric scale. They therefore support directional comparison under matched stress inputs, not absolute rankings across live protocols. Those two rules are what keep the closing claims comparative, conditional, and proportionate to the available evidence.

8.6 Answer to the Research Question and Contributions

We answer the research question in comparative rather than predictive terms. DePIN tokenomic mechanisms behave differently under physical and economic stress because their reward rules, sink pathways, and participation conditions create different transmission paths across utilization, provider economics, participation, and incentive–usage alignment. The relevant question is therefore not whether a mechanism is universally stable, but which subsystem weakens first under a given channel and how that weakness propagates through the rest of the system.

Within the DTSE experiment set used in this thesis, the answer appears in earliest-trigger metric families, baseline-relative timing, and Stage 2 failure signatures. Chapter 7 showed

four recurring patterns: demand contraction weakens usage-linked fields before participation visibly contracts; liquidity shock first compresses the internal price layer and then provider economics and churn; competitive-yield pressure pulls participation away before utilization becomes the first-moving field; and provider-cost inflation compresses margins before visible participation loss follows. Taken together, those patterns show that robustness in DePIN is not a single property. It is a question of where sensitivity concentrates and how quickly deterioration crosses from one subsystem into another.

For the capped-supply ONO anchor, the main trade-off is now clear. The ONO/Data-Credit design dampens one route of user-facing token-volatility transmission, but it does not remove the deeper requirement that usage-linked sinks and provider economics remain aligned over time. More generally, neither capped-supply nor BME-oriented designs escape DePIN's physical constraints. They differ in where stress appears first, how much of it is absorbed by market-facing layers versus provider-facing layers, and how long apparently stable service conditions can persist before accumulated pressure becomes operationally visible.

We make three main contributions that speak to distinct audiences. First, for academic researchers, we build a coherent line from DePIN stress theory and comparative mechanism vocabulary to a bounded evaluator that makes its assumptions explicit. Second, for protocol designers, we establish a comparative diagnostic language based on earliest-trigger fields and failure-signature sequences rather than binary stable/unstable labels. Third, for governance teams and protocol stewards, we provide an interpretive layer that remains clearly separated from DTSE output itself, allowing discussion of response categories without confusing modeled results with live intervention efficacy.

8.7 Future Research Agenda

The next extensions follow directly from the limits established above. One path is broader mechanism coverage: additional documented DePIN profiles could be brought into the same stress grid while preserving the same metric vocabulary and reporting logic. A second path is stronger empirical anchoring where data quality permits it, especially through more systematic benchmarking of selected DTSE outputs against historical stress windows without turning the evaluator into a forecasting system [32, 31]. A third path is explicit intervention modeling. If governance-response latency, verification failures, or targeted re-allocation rules are to be studied rigorously, they should enter DTSE as explicit rule or input changes rather than as interpretive commentary.

These extensions would deepen the thesis in different ways. Broader profile coverage would test how portable the diagnostic vocabulary remains across DePIN sectors. Stronger empirical anchoring would help refine assumption bounds where public observability improves.

Explicit intervention modeling would make it possible for us to compare response classes under the same failure-signature vocabulary used here. Each path would strengthen the evaluator, but each also depends on preserving the same evidence discipline that shaped the current thesis.

8.8 Closing Statement

We present DTSE in this thesis as a bounded comparative evaluator for reasoning about DePIN tokenomics under stress when closed-form analysis is insufficient and empirical observability is incomplete. By making assumptions explicit, reporting outputs as baseline-relative comparisons, and interpreting deterioration through failure-signature sequences, the framework supports a more disciplined discussion of robustness than either growth narratives or binary collapse labels can provide on their own.

We therefore close the thesis with a narrower but stronger claim. Our framework does not prove how any live DePIN network will behave under future stress. Instead, we show how different tokenomic designs can be compared before catastrophic failure by asking where stress appears first, how it propagates, and which response categories become salient once those sequences are visible.

A Appendix

This appendix retains only compact audit material that supports the accepted empirical and DTSE chapters. Core chapter argumentation, scenario framing, and methodology exposition remain in the main body.

A.1 DTSE Scenario-Formula Map

Table A.1 gives a compact summary of why each scenario is included, how it enters the evaluator, and which fixed settings matter most for interpretation.

Scenario	Why Included	Rule Summary	Key Fixed Settings
Baseline Neutral	Establish the reference path for timing and magnitude comparisons.	Stable demand and background settings generate the neutral baseline against which all later deviations are measured.	No adverse event trigger; no signal threshold needs to be crossed for the baseline to exist as a reference path.
Demand Contraction	Test reward–demand decoupling under weakening service use.	Demand decays through the scenario-specific contraction path under higher volatility.	Utilization becomes a material signal once it departs ten percent from the baseline path.
Liquidity Shock	Test fast transmission from token-liquidity stress into provider-economics fields.	A discrete unlock-style sell event is introduced under limited market depth, producing a sudden price dislocation.	Event week: 20. Sell-pressure scalar: 35%. Material price deviation threshold: 10%.
Competitive Yield Pressure	Isolate provider reallocation pressure without requiring demand collapse.	External yield pressure raises switching sensitivity, with provider types responding at different intensities under the same outside opportunity.	External-yield scalar: 1.5×. Provider-count threshold: 5%. Churn-increase threshold: 20%.
Provider Cost Inflation	Test Subsidy-Gap exposure under higher operating costs.	Provider cost baselines are shifted upward across the active provider set while all other scenario channels remain neutral.	Cost multiplier: 1.35×. Profitability threshold: 10%. Provider-count threshold: 5%.

Table A.1: Scenario logic and key fixed settings for the DTSE experiment set.

A.2 DTSE Threshold Definitions

Table A.2 summarizes the main fixed thresholds referenced in Sections 6 and 7.

Threshold or Setting	Value	Unit	Interpretive Role
Negative-provider-profit rule	0	weekly profit	Separates non-negative provider economics from loss-making periods in profitability-based readouts.
Sustained-loss window	3	weeks	Defines how long negative profitability must persist before profitability-induced churn is treated as a Stage 2 pattern.
Liquidity event timing	20	week index	Marks the discrete sell event used in the liquidity-shock scenario.
Liquidity sell-pressure scalar	35	percent	Defines the size of the unlock-style sell event in the liquidity-shock scenario.
Competitive-yield scalar	1.5	relative multiple	Sets the strength of the outside opportunity used in the competitive-yield scenario.
Provider-cost multiplier	1.35	multiplier	Sets the exogenous increase in operating costs used in the cost-inflation scenario.
Price-signal threshold	10	percent	Defines a material price deviation from baseline in Stage 1 detection.
Utilization-signal threshold	10	percent	Defines a material utilization deviation from baseline in Stage 1 detection.
Profitability-signal threshold	10	percent	Defines a material profitability deviation from baseline in Stage 1 detection.
Provider-count threshold	5	percent	Defines a material active-provider deviation from baseline in Stage 1 detection.
Churn-increase threshold	20	percent	Defines a material churn deviation from baseline in Stage 1 detection.

Table A.2: Excerpt of the fixed thresholds used in the DTSE reporting logic.

A.3 Override Ledger

The reported DTSE runs also include a small number of profile-specific and global overrides. Table A.3 summarizes the main ones in human-readable form.

Override Family	Scope	Status	Purpose
Profile-specific churn thresholds	ONO, Helium, Geodnet	Modeled assumption	Allow continuation sensitivity to differ across mechanism profiles without treating those values as empirical estimates of live behavior.
ONO maximum weekly churn cap	ONO	Modeled assumption	Prevents unrealistic discontinuities in weekly exit intensity while preserving participation stress under adverse conditions.
Provider-type switching adjustment	Global competitive-yield rule	Scenario logic	Allows competitive-yield pressure to affect provider types differently under the same outside opportunity.

Table A.3: Main profile-specific and global overrides used in the reported DTSE runs.

A.4 DTSE Run Manifest

Table A.4 records the key features of the frozen run manifest used for Sections 7 and 8.

Replications follow a fixed random-start sequence for reproducibility. This changes stochastic draws only, while profile rules, scenario settings, threshold definitions, and the reporting horizon remain fixed.

Run-set specification

The reported DTSE run set for Chapters 7 and 8 uses one fixed experiment configuration. That setup contains: (1) five scenarios (Baseline Neutral, Demand Contraction, Liquidity Shock, Competitive Yield Pressure, Provider Cost Inflation), (2) five mechanism profiles (Onocoy, Helium, Hivemapper, Grass, Geodnet), (3) one fixed threshold set, (4) one fixed override set, and (5) a 52-week horizon per run.

Seed policy

A seed is the numeric initializer for randomized model draws. For run index $r \in \{1, \dots, 60\}$, the seed is assigned as

$$S_r = 20260224 + (r - 1).$$

Run 1 uses seed 20260224, run 2 uses 20260225, and run 60 uses 20260283. With 5 scenarios, 5 profiles, and 60 replications, the reported set contains $5 \times 5 \times 60 = 1500$ trajectories.

Manifest Field	Recorded Value
Run timestamp (UTC)	24 February 2026, 10:21 UTC
Run-set definition	Frozen Chapter 7–8 DTSE configuration used for all reported comparisons
Deterministic seed policy	Master seed plus run index under a fixed seed schedule
Horizon and repetitions	52 weeks per run; 60 repeated runs per scenario–profile condition
Mechanism profiles	ONO calibrated profile, Helium BME profile, Hivemapper profile, Grass profile, and Geodnet profile
Scenario set	Baseline Neutral, Demand Contraction, Liquidity Shock, Competitive Yield Pressure, and Provider Cost Inflation

Table A.4: Run-manifest excerpt for the frozen DTSE reporting set.

Run Component	Seed 20260224 (Run 1)	Seed 20260225 (Run 2)	Interpretation
Scenario definitions	Identical	Identical	Fixed by the reported run-set definition above.
Mechanism profile rules	Identical	Identical	Fixed per profile mapping.
Threshold values	Identical	Identical	Fixed before result interpretation.
Horizon	52 weeks	52 weeks	Fixed reporting horizon.
Demand noise draw	Realization A	Realization B	Varies by seed.
Provider heterogeneity draw	Realization A	Realization B	Varies by seed.
Decision-noise draw	Realization A	Realization B	Varies by seed.
Price-related stochastic draw	Realization A	Realization B	Varies by seed.

Table A.5: Deterministic seed schedule with fixed-rule vs variable-draw separation.

Week	Metric	Seed 20260224	Seed 20260225	Interpretation
0	Demand	10,293.60	10,541.80	Small seed-driven demand variation.
0	Providers	3,000	3,000	Initial provider stock fixed.
0	Utilization	1.58	1.62	Minor stochastic utilization spread.
1	Providers	2,825	2,835	Same provider direction; small magnitude spread.
1	Churn	175	165	Churn magnitude differs by seed draw.
1	Profit	9.57	8.72	Profit differs by seed; rules unchanged.

Table A.6: Baseline excerpt (no stress event): same fixed DTSE rules, different seed path.

Week	Metric	Seed 20260224	Seed 20260225	Interpretation
20	Price (pre-update)	0.485846	0.484401	Near-identical pre-shock price level.
20	Providers	1,167	1,298	Provider state differs by seed draw.
20	Churn	6,782	7,586	Churn magnitude differs by seed draw.
21	Price (after propagation)	0.000107	0.000107	Same post-shock price direction.
21	Profit	-37.56	-39.38	Margin compression in both seeds.
21	Providers	1,965	1,517	Provider direction aligns; spread differs.

Table A.7: Liquidity-shock excerpt (event at week 20): same fixed DTSE rules, different seed path.

Together, Tables A.6 and A.7 show the same separation logic used throughout the thesis: scenario definitions, profile rules, thresholds, and horizon remain fixed, while stochastic draws vary by seed.

Empirical applicability note

The empirical metric-applicability boundary remains in Section 5, where the thesis defines the *N/R* rule, applies the comparator-level matrix, and states the Onocoy-specific observability limits directly in the main body rather than repeating them in appendix form.

A.5 Semi-Structured Interview Context Summary

The thesis also draws on a semi-structured interview with a representative of the Onocoy project. The purpose of the interview was not to establish mechanism facts in place of public documentation, but to clarify underspecified operational issues and identify stress-relevant sensitivities where public materials remain thinner than the token and payment architecture itself. Accordingly, the interview is treated here as bounded practitioner context. It informs interpretation and evidence-gap identification, but it does not override protocol documentation, public tokenomics materials, or directly auditable public data.

Emission structure

Question focus: Is issuance time-based, demand-based, governance-adjustable, or hybrid?
Is there a formal taper plan?

Interview clarification: Issuance was described as time-based rather than demand-adjusted in the short run. Public tokenomics documentation also specifies a 16% annual reduction in newly distributed ONO.

Use in this thesis: This supports the reading that demand shocks primarily weaken buyback strength rather than automatically reducing issuance pressure; in the thesis, this remains bounded practitioner context consistent with public tokenomics documentation.

Demand measurement

Question focus: How is demand measured internally?

Interview clarification: Data Credits were described as the main internal demand proxy, with service usage observed through credit purchase and burn activity. Public documentation also makes clear that Data Credits are the non-transferable service unit priced against fiat-denominated value.

Use in this thesis: This reinforces the thesis treatment of Data Credits as the clearest available demand-linked field in the Onocoy design; it is used as bounded practitioner context broadly aligned with public documentation.

Buyback logic

Question focus: What is the current burn-to-emission relationship? How does demand affect net supply in the short run?

Interview clarification: Buybacks were described as revenue-mediated rather than as a strict mechanical burn-and-mint equilibrium. Public tokenomics documentation likewise ties Data Credit usage to burn and conditional ONO buyback routing rather than to a native BME rule.

Use in this thesis: This supports the thesis distinction between ONO and native BME systems and clarifies why demand depth matters for token-side support; it is treated as bounded practitioner context consistent with public tokenomics routing language.

Provider economics

Question focus: What kinds of provider-cost differences and economic horizons matter in practice?

Interview clarification: Provider economics were described as heterogeneous rather than uniform, with hardware choice, reward levels, and token price all affecting economic viability. Public retail listings likewise suggest that Onocoy-compatible hardware can span a few-hundred-dollar entry level, while more complete setups are materially higher in cost.

Use in this thesis: This is used only to justify why provider-cost sensitivity should remain explicit in interpretation and DTSE assumption design; it combines bounded interview clarification with dated public market-price context rather than a fixed mechanism fact.

Retention sensitivity

Question focus: Which mechanism most strongly influences provider retention today?

Interview clarification: Retention was described as more sensitive to profitability, macro conditions, and local saturation than to governance participation.

Use in this thesis: This supports the thesis view that participation elasticity may emerge through provider economics before broader service deterioration becomes obvious; it remains bounded practitioner context, not a documented mechanism rule.

Governance response

Question focus: What happens operationally if token price drops sharply, and who can change emissions or related rules?

Interview clarification: Stress adaptation was described as discretionary rather than automatic, with rule changes depending on explicit governance action rather than built-in feedback loops.

Use in this thesis: This supports the thesis distinction between algorithmic stabilization and governance-mediated intervention; it is treated as bounded practitioner context consistent with the current public evidence boundary.

Utility versus speculation

Question focus: Do participants appear mainly utility-driven, speculative, or mixed in their token behavior?

Interview clarification: Participant behavior was described as mixed: some operators appear long-horizon, while others sell rewards to offset operating costs.

Use in this thesis: This supports the thesis reading that price volatility can amplify margin stress without reducing the network to a purely speculative system; it remains bounded practitioner context only.

References

- [1] Mark C. Ballandies et al. *A Taxonomy for Blockchain-based Decentralized Physical Infrastructure Networks (DePIN)*. Center for Law & Economics Working Paper Series 03/2023, ETH Zurich. Working paper. 2023. DOI: 10.3929/ethz-b-000640008. URL: <https://doi.org/10.3929/ethz-b-000640008>.
- [2] Zhibin Lin et al. *Decentralized Physical Infrastructure Network (DePIN): Challenges and Opportunities*. Preprint. 2024. arXiv: 2406.02239 [cs.NI]. URL: <https://arxiv.org/abs/2406.02239>.
- [3] W. Brian Arthur. “Competing Technologies, Increasing Returns, and Lock-In by Historical Events”. In: *The Economic Journal* 99.394 (1989), pp. 116–131. DOI: 10.2307/2234208. URL: <https://doi.org/10.2307/2234208>.
- [4] Dylan Bane and Salvador Gala. *State of DePIN 2024*. Messari, 2024. URL: <https://messari.io/report/state-of-depin-2024>.
- [5] Muneer Maher Alshater. “Decentralized Physical Infrastructure Networks (DePIN) Tokenomics”. In: *Frontiers in Blockchain* 8 (2025). Accessed: 2026-01-19. DOI: 10.3389/fbloc.2025.1644115. URL: <https://www.frontiersin.org/journals/blockchain/articles/10.3389/fbloc.2025.1644115/full>.
- [6] C.-F. Wu, C.-M. Ho, and S.-J. Wu. “Cryptocurrency Policy Uncertainty and Bitcoin Performance”. In: *Finance Research Letters* (2022). Discusses macro conditions and structural weakness.
- [7] Michael T. C. Chiu et al. *DePIN: A Framework for Token-Incentivized Participatory Sensing*. Preprint. 2024. arXiv: 2405.16495 [cs.GT]. URL: <https://arxiv.org/abs/2405.16495>.
- [8] Tim P. Morris, Ian R. White, and Michael J. Crowther. “Using simulation studies to evaluate statistical methods”. In: *Statistics in Medicine* 38.11 (2019), pp. 2074–2102. URL: <https://onlinelibrary.wiley.com/doi/10.1002/sim.8086>.
- [9] Sietse Braakman, Pras Pathmanathan, and Helen Moore. “Evaluation Framework for Systems Models”. In: *CPT: Pharmacometrics & Systems Pharmacology* 11.3 (Jan. 2022), pp. 264–289. DOI: 10.1002/psp4.12755. URL: <https://doi.org/10.1002/psp4.12755>.
- [10] Michael S. Andrew and Mark C. Ballandies. *Are you a DePIN? A Decision Tree to Classify Decentralized Physical Infrastructure Networks*. Preprint. 2025. arXiv: 2501.17416 [cs.ET]. URL: <https://arxiv.org/abs/2501.17416>.

- [11] Jean-Charles Rochet and Jean Tirole. “Platform Competition in Two-Sided Markets”. In: *Journal of the European Economic Association* 1.4 (2003), pp. 990–1029. DOI: 10.1162/154247603322493212. URL: <https://doi.org/10.1162/154247603322493212>.
- [12] CoinMarketCap. *DePIN Report - Deep Dive*. 2024. URL: <https://s3.coinmarketcap.com/uploads/depin-deep-dive-bridging-to-the-real-world.pdf>.
- [13] Coincub. *Deep in DePIN: Why Solana is the Execution Layer for Real-World Networks*. Accessed: 2026-01-19. 2025. URL: <https://coincub.com/depin-solana/>.
- [14] Mark C. Ballandies. *Chain Selection for a Decentralized Physical Infrastructure Network (DePIN)*. WiHi Weather (Medium). Project-authored chain-selection memo; archived locally in thesis repository. Mar. 2023. URL: <https://medium.com/wihi-weather/chain-selection-for-a-decentralized-physical-infrastructure-network-depin-bf1c646d9330>.
- [15] Helium Core Developers. *HIP-70: Scaling the Helium Network*. Primary governance proposal for Solana migration and subDAO architecture. 2022. URL: <https://github.com/helium/HIP/pull/470>.
- [16] Render Network Knowledge Base. *Burn Mint Equilibrium*. Accessed: 2026-01-19. 2025. URL: <https://know.rendernetwork.com/basics/burn-mint-equilibrium>.
- [17] Mike Horton. *GEODNET Improvement Proposal 7*. Accessed: 2026-01-19. 2025. URL: <https://medium.com/geodnet/geodnet-improvement-proposal-7-c2313f643b9b>.
- [18] io.net. *2025: io.net Year in Review*. Accessed: 2026-01-19. 2025. URL: <https://io.net/blog/2025-io-net-year-in-review>.
- [19] CoinMarketCap. *Latest Hivemapper (HONEY) Price Analysis*. Accessed: 2026-01-19. 2025. URL: <https://coinmarketcap.com/cmca-ai/hivemapper/price-analysis/>.
- [20] Onocoy Association. *Your Guide to Understanding the \$ONO Token*. 2024. URL: <https://onocoy.com/blog/your-guide-to-understanding-the-ono-token>.
- [21] Lisa J. Y. Tan. *The Economics and Math of Token Engineering and DeFi: Fundamentals of Token Economics*. Comprehensive taxonomy of incentive mechanisms. Economics Design, 2020.
- [22] Onocoy Association. *ONOCOY: Enabling Mass Adoption of High Precision GNSS Positioning Using Web3 (White Paper Revision 3.0.1)*. White paper revision 3.0.1. July 2025. URL: <https://onocoy.com>.
- [23] Onocoy Documentation. *Mining rewards breakdown*. 2024. URL: <https://docs.onocoy.com/documentation/mining-rewards-breakdown>.
- [24] Onocoy Documentation. *Explorer*. Accessed: 2026-03-01. 2026. URL: <https://docs.onocoy.com/documentation/explorer>.

- [25] Onocoy Documentation. *Ambassador Program*. Accessed: 2026-03-01. 2026. URL: <https://docs.onocoy.com/documentation/ambassador-program>.
- [26] Onocoy Documentation. *Tokenomics*. Accessed: 2026-03-01. 2026. URL: <https://docs.onocoy.com/documentation/tokenomics>.
- [27] Dune. *Onocoy Dashboard*. Accessed: 2026-03-01. 2026. URL: <https://dune.com/onocoy/dashboard>.
- [28] A. Craig MacKinlay. “Event Studies in Economics and Finance”. In: *Journal of Economic Literature* 35.1 (1997), pp. 13–39. URL: <https://www.jstor.org/stable/2729691>.
- [29] Binance. *Binance Spot API: HNTUSDT weekly kline data (2022–2023 window)*. Exchange primary market data endpoint, accessed 2026-02-26. 2026. URL: <https://api.binance.com/api/v3/klines?symbol=HNTUSDT%5C&interval=1w%5C&startTime=1640995200000%5C&endTime=1703980800000>.
- [30] GEODNET Docs. *Location NFT*. Accessed: 2026-01-19. 2025. URL: <https://docs.geodnet.com/geod-token/location-nft>.
- [31] Andrew J. Collins, Matthew Koehler, and Christopher J. Lynch. “Methods That Support the Validation of Agent-Based Models: An Overview and Discussion”. In: *Journal of Artificial Societies and Social Simulation* 27.1 (2024), p. 11. DOI: 10.18564/jasss.5258. URL: <https://doi.org/10.18564/jasss.5258>.
- [32] Josie McCulloch et al. “Calibrating Agent-Based Models using Uncertainty Quantification Methods”. In: *Journal of Artificial Societies and Social Simulation* 25.2 (2022), p. 1. DOI: 10.18564/jasss.4791. URL: <https://doi.org/10.18564/jasss.4791>.