

ESCALATION FRICTION, FEEDBACK SUPPRESSION, AND VOICE
LEGITIMACY: A HEURISTIC INDICATOR FRAMEWORK FOR RISK ENTROPY
IN RISK CULTURE

by

Marc Pharoah

DISSERTATION
Presented to the Swiss School of Business and Management Geneva
In Partial Fulfilment
Of the Requirements
For the Degree

DOCTOR OF BUSINESS ADMINISTRATION

SWISS SCHOOL OF BUSINESS AND MANAGEMENT GENEVA

APRIL, 2026

ESCALATION FRICTION, FEEDBACK SUPPRESSION, AND VOICE
LEGITIMACY: A HEURISTIC INDICATOR FRAMEWORK FOR RISK ENTROPY
IN RISK CULTURE

by

Marc Pharoah

APPROVED BY



dr. Ljiljana Kukec, PhD
Dissertation chair

RECEIVED/APPROVED BY:

Admissions Director

Dedication

For Denis Buterin, for his mentorship.

For my case study sponsor, for their trust.

For Christoph Smiela, for his guidance and friendship.

For Alicia, Archer, and Matilda, the heart of everything I do.

Acknowledgements

I would like to express my gratitude to Denis Buterin for his guidance, constructive challenge, and steady support throughout this research journey. I am also very appreciative of my case study sponsor, whose trust and collaboration enabled this work, and of the participants within the organisation who contributed their time, insight, and openness. Their willingness to engage meaningfully provided the foundation for this study.

My thanks go to the academics who responded to my enquiries and generously shared their expertise, helping to sharpen the conceptual and methodological dimensions of this thesis. I am equally grateful to the representatives from regulatory bodies who offered their perspectives and enriched the broader context of the study.

I would also like to acknowledge Christoph Smiela for his encouragement, thoughtful conversations, and steady support throughout.

Finally, my heartfelt appreciation goes to Alicia, Archer, and Matilda. Their patience, understanding, and unwavering belief sustained me through the most demanding phases of this work. I am extremely grateful for their love and support.

ABSTRACT
ESCALATION FRICTION, FEEDBACK SUPPRESSION, AND VOICE
LEGITIMACY: A HEURISTIC INDICATOR FRAMEWORK FOR RISK ENTROPY
IN RISK CULTURE

Marc Pharoah
2026

Dissertation Chair: Dr. Ljiljana
Co-Chair: Dr. Vasiliki

This study develops and applies a heuristic indicator framework centred on escalation friction, feedback suppression, and voice legitimacy to diagnose “risk entropy” within the risk culture of a regulated financial institution. Adopting an interpretivist, embedded single-case design, the research examines one financial institution that is formally compliant and locally perceived as open. The purpose is to understand how these behavioural indicators manifest across local relationships and formal governance channels, how they combine into unit-level configurations, and how they contribute to system-level decoupling between espoused mechanisms and enacted practice.

Data were generated through 25 semi-structured interviews across three lines model, functions, job bands, and locations, complemented by organisational artefacts such as committee minutes, risk governance documents, and engagement survey outputs, and supported by reflexive journaling. Indicators were operationalised through reflexive thematic analysis, unit-level profiling, and explanation-building that linked observed patterns of communication and challenge to underlying structural and relational

mechanisms. This enabled the construction of unit-level “entropy signatures” and a two-level pattern contrasting locally trusted voice with systemic escalation failure.

Findings show that within trusted dyads and small teams, formal mechanisms for openness, raising concerns, and reinforcement can translate into legitimate and psychologically safe risk voice. However, as concerns move into formal channels, procedural complexity, relational gatekeeping, and cognitive normalisation generate escalation friction and feedback suppression, widening the discourse–practice gap, in which governance mechanisms are discussed far more than their behavioural effects are evidenced. Four distinct unit configurations, symbolic governance, control without voice, silenced executors, and boundary negotiators, combine to form a systemic pattern of governance blindness, control-function silencing, and operational normalisation that characterises risk entropy in the case institution.

The institution's primary vulnerability lies not in its inventory of mechanisms but in a behavioural decoupling between espoused governance and enacted practice. On this basis the thesis advances four contributions: the Two-Level Pattern as a diagnostic reframing of risk culture; the Sequential Friction-to-Entropy Mechanism as a process account linking procedural barriers to inferred risk entropy; Unit-Level Entropy Signatures as a configurational typology of subcultural variation; and the Discourse–Practice Ratio (3.7:1) as a documentary heuristic marker. These are offered as candidate diagnostic tools for boards, supervisors, and future comparative research.

TABLE OF CONTENTS

List of Tables xiii

List of Figures xiv

CHAPTER I: INTRODUCTION.....	1
1.1 Statement of the Problem.....	1
1.2 Significance of the Study.....	1
1.2.1 Contributions of this thesis.....	2
1.3 Research Purpose and Questions	3
1.4 Objectives and Propositions.....	4
1.5 Limitations, delimitations, and assumptions.....	5
1.6 Definition of terms.....	6
1.7 Background.....	8
CHAPTER II: REVIEW OF LITERATURE	10
2.1 Introduction to Risk Culture in Governance Literature	10
2.1.1 Inclusion criteria.....	11
2.2 Theoretical Framework.....	11
2.2.1 Definitions and Conceptualisations of Risk Culture.....	11
2.2.2 Significance for Regulatory Effectiveness.....	13
2.2.3 Formal vs. Emergent Culture	15
2.2.4 Frameworks Emphasising Compliance, Challenge, and Resilience	17
2.2.5 Psychological Safety as an Enabling Condition.....	19
2.2.6 Limitations of Existing Models.....	21
2.3 Communication and Challenge in Risk Culture: Underexplored Essentials.....	22
2.3.1 Rationale for Focusing on Communication and Challenge	22
2.3.2 Empirical Insights and Limitations	24
2.3.3 Toward a Deeper Understanding: Significance for Risk Culture Enhancement	25
2.4 Box-Ticking and Superficial Compliance	26
2.4.1 Critique of Regulatory Instrumentalism and Checkbox Mentality	26
2.4.2 Empirical Evidence and Case Studies of Superficial Risk Culture.....	27
2.4.3 Implications: Risk Entropy and Organisational Decay	29
2.5 The Problem of Risk Entropy in Organisations.....	30
2.5.1 Conceptualising Risk Entropy.....	30
2.5.2 Cultural Drift and Systemic Vulnerabilities.....	32
2.5.3 Evidence of Increasing Risk Entropy and Systemic Failures	33

2.6 Gaps in the Literature.....	34
2.6.1 Limited Empirical Research on Actual Practices of Communication and Challenge	34
2.6.2 Underdeveloped Indicators for Deep Cultural Dynamics	35
2.6.3 The Need for Deeper Theoretical and Empirical Foundations	37
2.6.4 The Critical Gap: Linking Communication and Challenge to Measurable Outcomes	38
2.7 Chapter Summary	39
CHAPTER III: METHODOLOGY	42
3.1 Overview of the Research Problem	42
3.1.1 Philosophical Positioning	42
3.1.2 Interpretivist Focus on Communication and Challenge	43
3.1.3 Researcher Positionality and Reflexive Stance	44
3.2 From Philosophy to Research Design	45
3.2.1 Embedded Single-Case Design	45
3.2.2 Case Selection and Justification	47
3.2.3 Access Negotiation and Regulatory Context	48
3.2.4 Case Boundaries: Temporal, Spatial, and Functional Definition	49
3.2.5 Embedded Units: Structure, Definition, and Inclusion Criteria	49
3.2.6 Design Implications for Sampling and Analysis	49
3.3 Construction and Operationalisation of Heuristic Behavioural Indicators	50
3.3.1 Heuristic Behavioural Indicator Framework	51
3.3.2 Escalation Friction (EF)	53
3.3.3 Feedback Suppression (FS)	53
3.3.4 Voice Legitimacy (VL)	54
3.3.5 Risk Entropy (RE)	55
3.3.6 Synthesis and Use of Indicators	57
3.4 Sampling Strategy and Participant Recruitment	58
3.4.1 Sampling Approach	59
3.4.2 Sample Targets and Achieved Sample	60
3.4.3 Saturation Monitoring and Interpretive Completeness	62
3.4.4 Non-Response Considerations and Self-Selection	63
3.5 Data Collection and Fieldwork Operations	63
3.5.1 Overview of Data Collection Strategy	64
3.5.2 Interview Guide Development	65
3.5.3 Semi-Structured Interviews: Conduct and Quality Assurance	66
3.5.4 Organisational Artefact Analysis	68
3.5.5 Reflexive Journaling	70
3.5.6 Integration Across Methods	70

3.6 Data Analysis and Indicator Construction	72
3.6.1 Reflexive thematic analysis.....	73
3.6.2 Code-to-theme development	75
3.6.3 Unit-level indicator profiling	75
3.6.4 Indicator specification, validation and intensity rubric.....	76
3.6.5 Explanation building and risk entropy synthesis	77
3.6.6 Core-Mechanism Symbiosis as Interpretive Framework.....	78
3.6.7 Findings Chapter Organisation Logic	78
3.7 Trustworthiness, Quality Assurance and Ethical Considerations	80
3.7.1 Rigour Strategy Overview.....	80
3.7.2 Credibility	81
3.7.3 Dependability and Audit Trail	82
3.7.4 Quality Assurance Mechanisms	82
3.7.5 Confirmability and Reflexivity	83
3.7.6 Transferability	84
3.7.7 Ethical Considerations and Data Security	85
3.7.8 Link to Heuristic Indicator Rigour.....	87
3.7.9 Methodological Limitations	87
3.8 Timeline and Resource Plan	88
3.8.1 Fieldwork Timeline.....	88
3.8.2 Detailed Timeline.....	89
3.8.3 Dependencies and Critical Path.....	89
3.8.4 Feasibility in Relation to Sampling and Methods	89
3.9 Data Management and Security	90
3.9.1 Data Management Principles	90
3.9.2 Storage, Access and Encryption.....	91
3.9.3 File Naming, Backup and Retention	91
3.9.4 Anonymisation and De-identification	92
3.9.5 Integration with Trustworthiness and Ethics.....	92
3.10 Chapter Summary	93
CHAPTER IV: RESULTS.....	95
4.1 Overview of the Two-Level Pattern	95
4.2 Data Sources and Triangulation Approach.....	102
4.3 Local Authenticity: Where Mechanisms Work	106
4.3.1 Openness Enables Voice.....	107
4.3.2 Mechanisms Work in Trusted Contexts	109
4.3.3 Reinforcement Enables Voice.....	110
4.3.4 Configuration Summary: Trusted Relationships Create Legitimate Voice	112
4.4 Systemic Failure: Where Mechanisms Cannot Scale	114
4.4.1 Friction Causes Suppression	115
4.4.2 Reinforcement Does Not Prevent Entropy.....	118

4.4.3 Openness Does Not Prevent Entropy	121
4.4.4 Configuration Summary: Structural Barriers Prevent Scaling	124
4.5 Unit-Level Entropy Signatures	126
4.5.1 Oversight, Towards Symbolic Governance	127
4.5.2 Risk and Compliance (2nd Line), Towards Control Without Voice	131
4.5.3 1st Line Operations and Support, Towards Silenced Executors	133
4.5.4 1st Line Business, Towards Boundary Negotiators	135
4.5.5 Configuration Synthesis: Four Signatures, One Structural Pattern	138
4.6 Organisational Decoupling: How Unit Configurations Aggregate.....	139
4.6.1 Diagnostic Cuts: Hierarchical and Geographic Perspectives	140
4.6.2 Towards Governance Blindness, Distance and Invisibility in the Oversight Configuration.....	142
4.6.3 Towards Control Function Silencing, Absorption and Authority in the 2nd Line Configuration	143
4.6.4 Towards Operational Normalisation, Workarounds in the Operations Configuration.....	145
4.6.5 The 3.7:1 Discourse-Practice Gap: Quantifying Decoupling.....	146
4.7 Artefact Triangulation: Validating the Two-Level Pattern.....	150
4.7.1 ERCC Committee Minutes and Packs: Persistent Reds and Deferral Patterns.....	151
4.7.2 Peakon Engagement Survey: Freedom of Opinion vs. Escalation Friction	154
4.7.3 Convergence, Divergence and Novel Insights	156
4.8 Summary of Findings.....	160
4.9 Chapter Summary	163
CHAPTER V: DISCUSSION.....	167
5.1 Introduction.....	167
5.2 Summary of Study and Findings	168
5.2.1 Overview of the Study	168
5.2.2 Headline Findings	169
5.3 Local Authenticity: Where Mechanisms Succeed	170
5.4 Systemic Failure: Where Mechanisms Cannot Scale	173
5.4.1 Structural Friction as a Driver of Suppression.....	173
5.4.2 The Reinforcement Paradox: Espoused but Not Enacted	175
5.4.3 Openness Without Effect: When Invitation Becomes Ritual.....	176
5.5 Unit-Level Entropy Signatures	178
5.5.1 Symbolic Governance: The Problem of Structural Distance	179
5.5.2 Control Without Voice: The Paradox of Formal Authority	180
5.5.3 Silenced Executors: Proximity Without Influence.....	182

5.5.4 Boundary Negotiators: Revenue-Legitimated Voice (Provisional).....	184
5.6 Organisational Decoupling: From Unit Configurations to Systemic Pattern	185
5.7 Artefact Triangulation: Methodological and Substantive Implications.....	186
5.8 Conclusions Aligned to Research Questions	188
5.8.1 Overarching Conclusion.....	191
5.9 Theoretical Contributions	192
5.10 Practical Implications for Risk Governance	196
5.11 Methodological Limitations.....	199
5.12 Directions for Future Research	201
5.12.1 Speculative and Longer-Horizon Directions.....	203
5.13 Chapter Summary	205
REFERENCES	207
APPENDIX A: INTERVIEW PROTOCOL AND GUIDE	216
A.1 Interview Guide (v3.5).....	216
A.2 Research Question-to-Indicator Mapping.....	219
A.3 Interview Guide Development Rationale.....	222
A.4 Participant Information Sheet and Consent Process	223
APPENDIX B: CODEBOOK EVOLUTION LOG	225
B.1 Codebook Version History (V0.1–V0.7, Nine Snapshots)	226
B.2 Post-Saturation Refinement Record	228
B.3 Reflexive Memo Extracts Supporting Code Revisions.....	229
APPENDIX C: SATURATION TRACKING MATRIX.....	231
C.1 Parent-Theme Emergence by Interview Cluster	231
C.2 Three-Consecutive-Interview Saturation Confirmation.....	232
C.3 Reference Count Progression (Saturation Signal).....	232
C.4 Reference Count Asymptotic Flattening	233
C.5 Saturation Monitoring Protocol.....	233
C.6 Per-Interview Reference Count Detail	233
C.7 Non-Response and Self-Selection Analysis.....	234
APPENDIX D: EVIDENCE BASE AND DOCUMENTATION.....	235
D.1 Triangulation Instances by Indicator, Unit, and Source.....	235
D.2 Triangulation Summary Statistics.....	238
D.3 Indicator Distribution.....	238
D.4 Method-to-Indicator Support Matrix	239

D.5 NVivo Query and Node Management	239
APPENDIX E: PHILOSOPHICAL POSITIONING AND RIGOUR DOCUMENTATION	240
E.1 Ontological and Epistemological Positioning	240
E.2 Researcher Positionality and Independence Statement	241
E.3 Validity Framework.....	242
APPENDIX F: CASE BOUNDARIES AND EMBEDDED UNIT DEFINITIONS	243
F.1 Case Boundaries: Temporal, Spatial, and Functional Definition	243
F.2 Embedded Unit Inclusion Criteria.....	244
F.3 Embedded Unit Definitions and Targets	245
APPENDIX G: INDICATOR OPERATIONAL DEFINITIONS AND DISCRIMINATION RULES.....	246
G.1 Full Indicator Framework Overview.....	246
G.2 Escalation Friction (EF): Operational Definition and Observable Cues.....	247
G.3 Feedback Suppression (FS): Operational Definition and Observable Cues.....	248
G.4 Voice Legitimacy (VL): Operational Definition and Observable Cues	249
G.5 Risk Entropy (RE): Operational Definition and Convergent Patterns.....	250
G.6 Indicator Intensity Rubric and Cross-Unit Categories.....	256
G.7 Explanation-Building Framework (Core-Mechanism Symbiosis)	257
G.8 Boundary-Case Discrimination Rules.....	258
APPENDIX H: FIELDWORK TIMELINE AND DATA MANAGEMENT.....	259
H.1 Research Workflow Timeline	259
H.2 Dependencies and Critical Path	260
H.3 Data Storage, Access and Encryption	261
H.4 Project Chronology Table	262
APPENDIX I: PARTICIPANT REFERENCE	262
I.1: Participant Reference Counts by Theme	263

LIST OF TABLES

Table 3.1	Heuristic Indicators and Context Mechanisms	52
Table 3.2	Response Distribution Across Recruitment Stages.....	61
Table 3.3	Sample Structure by Governance Line and Business Function	61
Table 3.4	Summary of Participant Characteristics.....	62
Table 3.5	Indicator Support Across Methods	71
Table 3.6	Indicator Intensity Levels and Typical Evidence Basis.....	77
Table 3.7	Analytical Template Elements and Functions.	79
Table 4.1	Indicator reference table	97
Table 4.2	Chapter organisation	103
Table 4.3	Overview of Four Unit Entropy Signatures	127
Table 4.4	Oversight Indicator profile.....	128
Table 4.5	2nd Line Unit, EF, FS, VL Intensity and RE Inference.....	131
Table 4.6	1 st Line Operations and Support Indicator profile	133
Table 4.7	1st Line Business Indicator profile	136
Table 4.8	Indicator Intensity by Job Band (Hierarchical Cut).....	141
Table 4.9	Geographic Indicator Profile.....	142
Table 4.10	Decoupling Index by Embedded Unit.....	146
Table 4.11	Triangulation Evidence Summary	157
Table 4.12	Internal Corroborating Evidence Summary	159

LIST OF FIGURES

Figure 3.1 Two-Level Indicator Architecture.....	58
Figure 3.2 Iterative Process.....	73
Figure 4.1: The Emergent RE Inference.....	98
Figure 4.0: Chapter IV Analytical Figure Sequence.....	101
Figure 4.2: Local Authenticity Configuration	114
Figure 4.3: The Sequential Barrier Model	118
Figure 4.4: Scaling Failure Model, Six Compounding Barriers	125
Figure 4.5: Unit Entropy Signature Matrix.....	139
Figure 4.6: Four Decoupling Mechanisms as a Self-Reinforcing Cycle	149
Figure 4.7: Triangulation Evidence Map	151

CHAPTER I: INTRODUCTION

1.1 Statement of the Problem

Financial regulators have invested heavily in strengthening risk governance since the 2008 financial crisis, yet repeated supervisory findings and post-mortem reviews continue to show that organisations may appear structurally compliant while cultural and communicative failures accumulate unseen. Boards and regulators typically rely on structural proxies (policies, committees, documented processes) and aggregate survey scores to infer the health of “risk culture”, but these instruments cannot reliably detect where escalation friction, feedback suppression, and distorted risk voice are producing a decoupling between espoused governance mechanisms and their enactment in day-to-day practice. In particular, existing tools rarely distinguish between local experiences of openness and psychological safety and the systemic efficacy of formal escalation channels, leaving supervisors with limited means to identify institutions where concerns are voiced in trusted dyads but do not reach formal decision-makers in time to prevent harm. The problem this study addresses is the absence of behaviourally grounded, multi-level diagnostic indicators that reveal how communication and challenge operate within a regulated financial institution and how their failure contributes to what the author has termed “risk entropy” (Pharoah, 2025). However, this is derived from the term “information entropy”, defined as the “measure of our confusion regarding the state of the system” (Vranken *et al.*, 2015).

1.2 Significance of the Study

This study is significant for both professional practice and academic understanding of risk culture. For practitioners, it offers a set of operational indicators, including

escalation friction, feedback suppression, voice legitimacy, and the discourse–practice ratio. The first three constitute the core heuristic indicator framework developed in Chapter II (Section 2.6), while the discourse–practice ratio is introduced as a derived diagnostic tool in Chapter III. These indicators can be traced through interviews, governance artefacts, and survey data to diagnose where risk voice succeeds and where it breaks down. For supervisors and boards, the findings demonstrate that an institution that appears structurally sound and locally open to speaking up can still experience systemic escalation failure, highlighting the need for diagnostics that go beyond structural proxies and aggregate culture scores. For the scholarly community, the study reframes risk culture through the lens of risk voice, develops a process account linking structural barriers to suppressed challenge, and introduces a typology of unit-level entropy signatures, thereby contributing to theory on organisational voice, decoupling, and cultural drift in financial governance.

1.2.1 Contributions of this thesis

The thesis advances four contributions to the study and practice of risk culture, each developed in depth in Chapter V (Section 5.9) and synthesised here for the reader's orientation:

The Two-Level Pattern — a diagnostic reframing. The study evidences local authenticity and systemic failure coexisting within the same institution under the same formal governance arrangements, reframing risk culture from a unitary organisational attribute into a construct that must be read at two levels simultaneously. This shifts the diagnostic question from "does this organisation have a healthy risk culture?" to "at which level does risk culture function, and where does it break down?".

The Sequential Friction-to-Entropy Mechanism — a process account. The study specifies an empirically grounded chain through which procedural barriers produce relational hesitancy, hesitancy triggers cognitive normalisation, normalisation generates suppression, and suppression accumulates into conditions from which risk entropy is inferred. The contribution is not any single link in the chain but the mechanism-level account of how the links compound.

Unit-Level Entropy Signatures — a configurational typology. The study identifies four analytically distinct configurations, Symbolic Governance, Control Without Voice, Silenced Executors, and Boundary Negotiators (provisional), showing that different parts of the same organisation contribute to risk entropy in qualitatively different ways, and providing a typology usable as a diagnostic instrument.

The Discourse–Practice Ratio (3.7:1) — a heuristic diagnostic marker. The study operationalises simulated alignment as a decomposable, documentary-record-based ratio of contextual governance discourse to proximate behavioural friction indicators. Its diagnostic value lies in direction and internal distribution rather than absolute value, and it is offered as a replicable procedure for future comparative work.

Together, these four contributions reframe risk culture as a two-level, configurationally diagnosable, and procedurally mechanistic phenomenon, rather than a unitary cultural attribute.

1.3 Research Purpose and Questions

The overall purpose of this research is to develop and apply a heuristic behavioural indicator framework for diagnosing risk entropy in a regulated financial institution, by examining how communication and challenge behaviours operate across local relationships and formal governance channels. Specifically, the study seeks to understand how escalation friction, feedback suppression, and voice legitimacy manifest in a live organisational setting, how these mechanisms combine into unit-level configurations, and how they contribute to system-level decoupling between espoused governance mechanisms and enacted practice.

The study is guided by the following research questions:

ID	Research Question
RQ1	How are alternative views about risk encouraged and expressed within the case institution, and under what conditions does risk voice appear legitimate and safe at the local level?
RQ2:	How do formal risk governance mechanisms for communication and challenge operate in practice, and where do escalation barriers and feedback suppression arise as concerns move through formal channels?
RQ3:	How can escalation friction, feedback suppression, voice legitimacy, and risk entropy be operationalised as diagnostic indicators to support boards and supervisors in assessing risk culture in structurally compliant institutions?
RQ4:	How do the observed unit-level entropy configurations and indicator patterns aggregate into a systemic governance pattern, and what does this reveal about organisational decoupling and the conditions under which risk entropy emerges?

Appendix A.2 provides a detailed research-question-to-indicator mapping, showing how each question is linked to specific elements of the heuristic indicator framework.

1.4 Objectives and Propositions

As this is an interpretivist, qualitative case study, the research does not test statistical hypotheses. It does, in fact, pursue the following research objectives and propositions:

O1: To construct and refine a heuristic behavioural indicator framework—escalation friction (EF), feedback suppression (FS), voice legitimacy (VL), and risk entropy (RE)—that captures observable patterns in communication and challenge behaviour within a regulated financial institution.

O2: To identify and explain the relational and structural conditions under which risk voice is locally legitimate and effective, and the conditions under which it is inhibited or silenced as concerns move into formal governance channels.

O3: To develop unit-level “entropy signatures” that configure EF, FS, VL, and RE at the level of embedded governance units, and to examine how these configurations aggregate into a systemic pattern of organisational decoupling.

O4: To derive candidate diagnostic tools, including a discourse–practice ratio and a concern-thread-derived integrity score, that can inform future quantitative and comparative studies of risk culture and guide practical interventions in risk governance.

These objectives operationalise the research questions and provide propositions about how the indicator framework will reveal patterned relationships between risk voice behaviours, unit-level configurations, and system-level risk entropy.

1.5 Limitations, delimitations, and assumptions

This study has several limitations that shape the interpretation and generalisability of its findings. First, it is based on an embedded single-case design within one financial

institution, which provides contextual depth but limits the extent to which findings can be statistically generalised; the configurations and mechanisms identified are transferable analytical constructs whose applicability to other settings must be assessed through replication rather than assumed. Second, fieldwork was conducted over a six-week period (21 October to 12 November 2025), capturing a coherent organisational moment rather than a longitudinal trajectory, so the temporal stability of the observed configurations and the scrutiny-dependency of cultural dynamics remain open questions. Third, the study adopts a relativist ontology and subjectivist epistemology, meaning that findings represent co-constructed interpretations of participants' situated sense-making rather than objective measurements of organisational attributes.

The study is deliberately delimited to one regulated financial institution operating under FSB-influenced governance frameworks, with case boundaries defined temporally, spatially, and functionally in Section 3.2.4 and Appendix F. It focuses on communication and challenge in risk governance rather than on other dimensions of risk culture such as incentives or capability, and it treats escalation failure primarily as a function of internal architectural factors rather than attempting to model wider societal trust dynamics in full. The analysis assumes that participants are broadly honest and reflective within the constraints of organisational role and confidentiality, that governance artefacts (minutes, board packs, policy documents) are sufficiently complete to support triangulation, and that the heuristic indicators, while interpretive, can be applied consistently across data sources using the operational definitions in Section 3.3 and Appendix G.

1.6 Definition of terms

For the purposes of this study, the following terms are defined as follows; detailed operational definitions and observable cues are provided in Section 3.3 and Appendix G.

Risk voice: The expression, escalation, and challenge of concerns, questions, or dissenting views relating to risk, whether informally within trusted relationships or through formal governance channels.

Escalation friction (EF): Structural and procedural barriers that increase the effort, uncertainty, or perceived cost of raising a concern through formal mechanisms, including unclear routes, complex templates, and multi-stage approvals.

Feedback suppression (FS): Patterns in which responses to raised concerns are delayed, diluted, or absent, leading participants to infer that speaking up does not lead to meaningful action or visible outcomes.

Voice legitimacy (VL): The perceived appropriateness, credibility, and likely impact of speaking up about risk, shaped by relational trust, role expectations, and the way challenge is received by those in authority.

The rationale for selecting these three constructs as the core diagnostic indicators, rather than other dimensions of risk culture, is set out in Chapter II, Section 2.6, with reference to the empirical gaps identified in the literature and to the FSB's (2014) definition of effective communication and challenge.

Risk entropy (RE): An interpretive synthesis construct denoting a cumulative, system-level governance disorder in which escalation friction, feedback suppression, and distorted voice accumulate across units, producing a decoupling between espoused governance mechanisms and their enactment in practice. The term was introduced in (Pharoah, 2025) and is operationalised here as a heuristic indicator.

Discourse–practice ratio: A heuristic measure comparing the volume of governance discourse (e.g., statements about risk appetite, cultural aspiration, and challenge) to the volume of documented actions addressing raised concerns, used to quantify the espoused/enacted gap.

Concern-thread integrity score: A proposed diagnostic indicator assessing the extent to which individual raised concerns can be traced from initial articulation through formal response, disposition, and evidenced outcome.

Local Voice Safety (LVS) and Systemic Escalation Efficacy (SEE): Two dimensions of risk voice distinguishing between perceived safety of speaking up within immediate teams (LVS) and confidence that concerns raised through formal mechanisms will reach appropriate decision-makers and lead to action (SEE).

1.7 Background

Since the global financial crisis of 2008, regulators and industry bodies have identified “risk culture” as a critical determinant of financial stability, repeatedly highlighting failures of communication and challenge as central to major supervisory interventions and enforcement actions. Supervisory frameworks such as the Financial

Stability Board’s risk culture guidance and the subsequent approaches of APRA, the ECB, the PRA, DNB, and MAS place behavioural expectations around openness, escalation, and challenge at the heart of prudential oversight, yet provide limited tools for assessing how these behaviours actually play out in practice. At the same time, empirical studies of risk culture and psychological safety show that local experiences of openness can coexist with systemic escalation failure, and that structural investment in governance mechanisms can produce “culture veneers” where performance of compliance masks underlying drift and decoupling. Against this backdrop, this thesis focuses on a single regulated institution that appears structurally compliant and locally open, and uses it as a lens to examine how risk voice operates across levels, how escalation friction and feedback suppression accumulate, and how their interaction can be rendered visible through a heuristic indicator framework.

CHAPTER II:

REVIEW OF LITERATURE

2.1 Introduction to Risk Culture in Governance Literature

Since the 2008 financial crisis, regulators and academics have made sustained efforts to bring structure and clarity to financial institutions through risk management governance. A further niche has been carved out within the wider risk framework to examine the behaviours and habits underpinning organisations, specifically to address one of the driving forces behind the crisis: the cultural and communicative failures that allowed risk to accumulate undetected. A key cornerstone of effective governance and financial stability has been the desire to manage organisations' risk culture. The concepts of norms, values, and behaviours remain crucial, though they are often described as elusive by researchers and regulators, who rely on the high-level guidance from the Financial Stability Board (FSB, 2014). This guidance was adapted from an earlier definition provided by the Institute of International Finance (IIF) (Kunz and Heitz, 2021).

Regulatory authorities, including APRA, the ECB, the FCA, De Nederlandsche Bank (DNB), and the PRA, have developed definitions and supervisory guidance aimed at embedding sound risk culture within financial institutions. However, despite this concentrated effort, not all the risk culture dimensions identified by the FSB have received equal research consideration. In particular, the dimension of communication and challenge lacks meaningful empirical evaluation, and it is this cultural lever that this chapter develops further. The review proceeds from definitional and theoretical foundations, through the underexplored territory of communication and challenge, to the phenomenon of box-ticking and superficial compliance, before introducing the concept of risk entropy as an organising framework for diagnosing cultural fragility.

2.1.1 Inclusion criteria

The literature reviewed for this chapter was identified primarily through searches of the EBSCO database and major regulatory publications (FSB, APRA, ECB, PRA, BIS, DNB, MAS), supplemented by targeted backwards and forward citation tracing. This was supported by contacting highly cited authors and representatives from some of the regulatory bodies, who also provided potential directions for literature reviews. The review focused on sources published between 2008 and 2025, reflecting the post-crisis regulatory emphasis on risk culture, with selective inclusion of earlier foundational works where they remain central to current debate. Priority was given to peer-reviewed academic articles, prudential supervisory reports, and high-quality practitioner analyses directly addressing risk culture, communication and challenge, and organisational voice; sector-specific studies outside financial services were included only where they offered transferable theoretical insights. Searches continued throughout the study, ensuring that any new material, post the literature chapter creation, was also considered.

2.2 Theoretical Framework

2.2.1 Definitions and Conceptualisations of Risk Culture

Despite its growing significance (Pan, Siegel and Wang, 2017; Osman and Lew, 2021; Sheedy and Canestrari-Soh, 2023), clear definitions of risk culture remain scarce across regulators, industry bodies, and academic research. As a result, the concept is often vague, superficial, and inconsistently addressed (Anton and Nucu, 2020). The Financial Stability Board (2014) defines risk culture as "norms, attitudes and behaviours related to risk awareness, risk-taking and risk management," a definition that has served as a foundation for subsequent regulatory and academic descriptions. The European Banking

Authority (EBA, 2021) emphasises observable and governable behaviour; the Australian Prudential Regulatory Authority (APRA, 2016, 2019) has extended this framing to include "the norms of behaviour that determine the way risks are identified, understood, discussed and acted upon." The ECB (2024) similarly stresses the "collective mindset" and "embedded behavioural norms," while the Monetary Authority of Singapore (MAS, 2020) distinguishes between structural "hardware" and behavioural "software."

The De Nederlandsche Bank strongly develops upon the work of Schein (Schein, 1984; DNB, 2015), adopting a multilayered view of behaviour and culture and framing risk culture as "a function of mindset, group dynamics, leadership behaviour, and communication patterns" that shape both formal and informal governance. This iceberg model is central: only those behaviours above the waterline are directly observable, while group dynamics and mindset remain hidden beneath. The FCA (2018) adopts a more holistic view, defining culture as "the habitual behaviours and mindsets that characterise an organisation" and explicitly acknowledging that no one culture fits all firms. The PRA aligns with this position, noting that "the PRA does not set a 'right culture' but assesses whether a firm's culture is likely to support prudent management" (PRA, 2014, 2021).

An early empirical perspective defines risk culture as the organisation's propensity to take risks as perceived by its managers, finding that primary drivers are not sector affiliation but trust in employees, clarity of mission, and levels of bureaucratic red tape, with considerable variance observed across otherwise comparable organisations (Bozeman and Kingsley, 1998). This has been refined further to shared perceptions among employees of the relative priority given to risk management, with expectations, behaviours, and psychological safety identified as key attributes (Sheedy and Griffin, 2018). Drawing on

Schein, risk culture has been positioned as embedded in artefacts, assumptions, and enacted norms, what people do when no one is watching (McConnell, 2013). Expanding beyond behaviour, Gorton, Grennan and Zentefis (2022) view culture as a coordination infrastructure, an epistemic system that governs decision-making under ambiguity. Kaplan and Mikes (2012) distinguish between calculative and visionary cultures, showing how managerial logics shape how risk is framed and internalised, while Weick (1995) and Aldag and Fuller (1993) locate culture in shared narratives and sensemaking, emphasising interpretive practice over codification. Building directly on the FSB framework, Dürst and Kunz (2025) operationalise these definitional commitments into a validated psychometric instrument, the Risk Culture Scale, structured around the FSB's four dimensions (tone from the top, accountability, effective communication and challenge, and incentives) and 25 empirically confirmed factors. By translating regulatory language into measurable survey items, their scale contributes a rare bridge between conceptual definitions of risk culture and its empirical assessment in banking contexts.

2.2.2 Significance for Regulatory Effectiveness

Regulators increasingly emphasise risk culture because they recognise that structural controls and formal policies offer little protection without aligned behaviours. The 2008 financial crisis revealed that many governance failures stemmed not from the absence of rules but from suppressed challenges, muted communication, and cultural drift. As Power (2009) argued, rules alone are inadequate unless supported by cultural systems that promote transparency, accountability, and dissent. Embedded risk culture has been shown to enhance resilience to systemic shocks (Palermo, Power and Ashby, 2017), while shared meaning and challenge behaviours are crucial to navigating complex risks such as financial crises and climate disruption (Klinke and Renn, 2021, 2021).

In response, the FSB (2014) defined four supervisory pillars: tone from the top, accountability, communication and challenge, and incentives, placing behaviour at the centre of prudential oversight. These have since been adopted across jurisdictions, APRA's CPS 220 (2017), MAS's (2020) conduct dashboards, and the ECB's revised SREP (2024) embed behavioural cues into governance evaluation. The PRA's Senior Managers and Certification Regime (SMCR) (2017) goes further, linking cultural expectations to individual accountability and enabling behavioural tracing through roles and consequences. APRA's Governance, Culture, Remuneration and Accountability (GCRA) strategy (2019) reflects a parallel shift, from procedural oversight to proactive cultural engagement, expanding supervisory tools to include thematic reviews, behavioural diagnostics, and board-level risk ownership. Yet APRA itself has cautioned against complacency, noting that 'now is not the time to slow momentum' in embedding robust risk cultures (APRA, 2022), a warning that points to how easily structural progress can stall without sustained behavioural reinforcement."

Yet significant blind spots remain. A disconnect exists between board perception and employee reality, particularly around "avoidance," where staff feel concerns are ignored (Sheedy and Griffin, 2018). Supervisors often misinterpret silence as alignment (Ashby, Palermo and Power, 2012), and Power (2009) cautions against rituals of verification, where visible compliance eclipses meaningful behavioural scrutiny. Without direct access to day-to-day cultural behaviour, regulators have turned to simple proxies, such as whether a Chief Risk Officer has been appointed, to infer whether robust risk management is in place (Horvey and Odei-Mensah, 2023). Final Notices have similarly served as post hoc cultural proxies, illuminating failures in tone, challenge, or escalation

only after harm had occurred (Ring *et al.*, 2016). Effective supervision of risk culture demands more than risk appetite statements and board minutes; it requires behaviourally grounded diagnostic tools.

2.2.3 Formal vs. Emergent Culture

Theoretical and regulatory discourse increasingly recognises that effective risk governance cannot be reduced to formal structures alone. Emergent risk culture, the informal, evolving constellation of shared meanings, relational behaviours, and interpretive routines, plays an equally, if not more, decisive role in shaping how risk is assessed, discussed, and responded to. Empirical work has documented that even within organisations with strong, uniform risk governance structures, considerable variation exists between business units (Sheedy and Griffin, 2018). This intra-organisational drift suggests that colleague interactions and interpretive context generate distinct cultural expressions, independent of centralised policies: may dictate what should happen, whereas emergent culture determines how it does happen, creating the possibility that local experiences of openness coexist with systemic escalation failure.

Drawing on Schein's iceberg model (1984), culture is framed not as statements or codes, but as a lived system of meaning, embedded in unconscious assumptions and daily behaviour, particularly how individuals act when unobserved (McConnell, 2013). Culture emerges through collective sensemaking, how actors interpret ambiguous risk environments and coordinate meaning through situated practice (Weick, 1995; Sandberg and Tsoukas, 2015). This interpretive infrastructure is what Gorton, Grennan and Zentefis (2022) identify as the firm's true "coordination architecture." In environments marked by

uncertainty, formal contracts and risk controls are insufficient; it is culture, anchored in shared beliefs, that fills the void when the path ahead is unclear.

APRA's (2016) supervisory guidance recognises that shared expectations and informal norms are often the primary levers by which culture sustains or undermines formal governance. Where behaviours are misaligned, even robust frameworks collapse into hollow performance. Pan, Siegel and Wang (2017) demonstrate that risk preferences persist across leadership regimes, showing that culture is not imposed episodically but evolves organically through institutional routines, selection criteria, and embedded logics. Power, Ashby and Palermo (2013) crystallise this tension into two distinct institutional responses: 'engineered' cultures, shaped through formal regulatory mandates, centralisation, and measurable proof of change; and 'organic' cultures, sustained through informal networks, internal ownership, and relational trust, cautioning that the post-2008 drift toward engineered approaches risks crowding out the organic dynamics upon which genuine behavioural change depends (Palermo, Power and Ashby, 2017). The institutional resistance to cultural supervision is itself revealing, as when De Nederlandsche Bank pushed into behavioural oversight, critics claimed it 'treaded on the territory of consultants', to which Frank Elderson responded that anything threatening financial stability falls squarely within its remit (DNB, 2015). A follow-on review of banks between 2015 and 2023 found that while boards showed improved awareness of group dynamics, roughly half of those examined held an optimistic perception of their own risk culture that was unsupported by evidence, with persistent short-term action reflexes crowding out the sustained reflection that genuine behavioural change requires (DNB, 2023). This defensiveness illustrates how the boundary between formal governance and emergent culture remains contested not only within organisations, but between regulators and the

institutions they supervise. To supervise and sustain effective risk cultures, both formal structure and emergent meaning must be rendered visible, measurable, and actionable. This aligns with classic institutional theory's distinction between formal organisational structures and the values, meanings, and adaptive patterns through which organisations become institutionalised (Selznick, 1996). This pattern of structural investment without behavioural embedding has repeated across jurisdictions since 2008, characterised first by cultural experimentation, then initiative fatigue, and ultimately the veneer problem McConnell (2013) describes as frameworks that perform reform without producing it. The gap between the governance mechanisms organisations espouse and the behaviours they enact has a theoretical precedent in Argyris and Schön's distinction, as discussed by Edmondson (1999), between espoused theory, the stated beliefs and values individuals use to explain their actions, and theory-in-use, the implicit action programmes that actually govern behaviour (Argyris and Schön, 1978, cited in Edmondson, 1999). At the organisational level, individuals who use what Argyris terms 'Model I' theories-in-use create O-I learning systems characterised by defensiveness, self-fulfilling errors, and the suppression of the feedback loops needed to correct them (Edmondson and Moingeon, 1998). Single-loop learning corrects errors within existing frameworks without questioning the governing assumptions that produce them; double-loop learning requires examining and altering those assumptions, a process that the present study argues is structurally inhibited within the observed institution's escalation architecture (Tosey, Visser and Saunders, 2012).

2.2.4 Frameworks Emphasising Compliance, Challenge, and Resilience

Risk culture frameworks have traditionally prioritised compliance-based approaches, emphasising adherence to formal rules, controls, and regulatory expectations.

While these systems provide procedural clarity, scholars and supervisors increasingly criticise them for failing to foster deep behavioural alignment, constructive dissent, and adaptive learning. This challenge is underscored by adapting the Treating Customers Fairly framework into a behavioural diagnostic for risk culture, focusing on open challenge, empowerment, and accountability, elements often excluded from traditional compliance checklists (McConnell, 2013). Within the control typology, interactive controls, which facilitate dialogue and dissent, are essential for managing risk in ambiguous contexts, whereas diagnostic controls may inadvertently reinforce silence or box-ticking behaviours (Kunz and Heitz, 2021).

A distinction has been drawn between calculative cultures, which rely on quantification and strict control, and risk environments that promote scenario thinking, dialogue, and reflexivity (Mikes, 2009). The subsequent work on boundary-work (Mikes, 2011) reveals why this distinction matters in practice, as risk managers must actively contest organisational jurisdictions and earn legitimacy for their challenge function, making the difference between structurally authorised oversight and cultural marginalisation a relational achievement rather than a structural given. The Macquarie Risk Culture Scale, structured around four dimensions, Valued, Proactive, Avoidance, and Manager (Sheedy and Griffin, 2018), offers a rare psychometric tool for measuring staff perceptions of whether risk is genuinely prioritised, whether issues are anticipated, and whether concerns are silenced. The "Avoidance" dimension captures the degree of suppressed dissent, serving as an early indicator of cultural fragility. Dürst and Kunz (2025) advance the psychometric landscape further by developing a comprehensive, FSB-aligned Risk Culture Scale comprising 78 items across 25 validated factors, covering all four regulatory dimensions in substantially greater granularity than prior instruments. Where

Sheedy and Griffin's scale provide a four-construct summary measure derived bottom-up from employee data, Dürst and Kunz construct their scale deductively from the FSB framework itself, enabling direct regulatory benchmarking. Critically for the present study, their Effective Communication and Challenge dimension disaggregates into two validated sub-constructs: Open to Alternate Views (whether managers actively seek and value alternative opinions) and Stature of Control Functions (whether control functions are respected and integrated as independent). This sub-dimensional structure, grounded in FSB indicators 3.3.1 through 3.3.5, makes Dürst and Kunz's instrument the most granular quantitative tool currently available for operationalising the communication and challenge dimension of risk culture.

Challenge is not only a governance issue but also a psychological and sociological one. Jetten and Hornsey (2014) show that the willingness to dissent is deeply tied to identity safety: employees speak up when they feel included, not when challenge is perceived as betrayal. Adler, Forbes and Willmott (2007) argue that governance systems must enable value-based authority, in which behaviour is guided by internalised norms rather than external rules. Stulz (2016) reinforces that challenge must be structurally empowered, noting that "despite their title, risk managers, for the most part, do not manage risk. They primarily measure it, monitor it, and help those who do manage risk." When formal governance fails to grant real authority to challengers, cultures of silence and complicity follow. Resilient risk cultures must therefore be designed to surface challenges, support dissent, and embed dialogue into everyday decision-making.

2.2.5 Psychological Safety as an Enabling Condition

Despite growing attention from regulators and scholars, current models of risk culture remain conceptually constrained and diagnostically limited. Dominant frameworks overemphasise structural elements, such as formal policies, governance processes, and compliance procedures, while overlooking the more elusive dimensions of behavioural dynamics, cultural drift, and communicative silencing. Sheedy and Griffin (2018) highlight this misalignment: structural indicators, such as the presence of CRO's or board committees, often create a false sense of assurance, and their fieldwork reveals gaps between institutional declarations and employee perceptions, especially regarding escalation norms, incentives, and voice. Without behavioural metrics, routine feedback, and real accountability, tone at the top becomes a symbolic gesture rather than a transformational lever, a critique of the proliferation of superficial culture programmes (McConnell, 2013).

Power (2009) warns against the transformation of culture audits into "rituals of verification," where documentation substitutes for insight and cultural performance is staged for supervisory reassurance. Kaplan and Mikes (2012) underscore that top-down systems work poorly in judgment-based environments where risk perception, interpretation, and informal norms override formal control. Many diagnostic tools fail to account for normative drift, culture is not static but unfolds through social interactions and interpretive feedback, yet most diagnostics are snapshot-based and lack the resolution to detect early signs of misalignment (Kunz and Heitz, 2021).

The most recent addition to this measurement landscape, Dürst and Kunz (2025), represents a significant advance, yet also clarifies the boundaries of what quantitative, single-level instruments can achieve. Their Risk Culture Scale was validated within a

single organisation using self-reported survey data, and the authors themselves acknowledge that socially desirable responding may systematically distort results, particularly around communication and challenge, where norms of compliance exert the strongest social pressure. Because the scale captures perceptions at one organisational level at a point in time, it cannot detect the two-level divergence between formal escalation channels and informal behavioural norms that this thesis identifies as the central diagnostic problem. The present qualitative, multi-level design is therefore not a substitute for instruments such as Dürst and Kunz's scale but an explicit complement to it: where survey instruments map the landscape of stated perceptions, this study examines the behavioural terrain that perceptions conceal.

2.2.6 Limitations of Existing Models

The preceding analysis of psychological safety and communicative silencing foregrounds a broader set of structural deficiencies in how risk culture is currently modelled and measured. Existing frameworks, despite their proliferation across regulatory guidance and academic literature, share a common tendency to privilege observable, codifiable outputs, policies, procedures, training records, over the informal, relational, and often concealed dynamics through which risk norms are actually reproduced. This Section draws together the principal limitations identified across the literature, arguing that the field's diagnostic toolkit remains inadequate to the complexity of the phenomenon it purports to assess.

Hardy and Maguire (2016) demonstrate how dominant narratives structure discourse boundaries, excluding dissent, deviance, or ambiguity from being acknowledged. McGoey (2012) builds on this, theorising "strategic ignorance" as an institutional practice:

actors may deliberately avoid troubling knowledge, thereby reinforcing dominant risk narratives even amid deteriorating conditions. Bromiley et al. (2015) argue that ERM models frequently treat culture as an exogenous variable rather than an internal, self-reinforcing force. Taken together, these critiques expose how many risk culture models remain conceptually elegant but empirically fragile: reforms must move beyond checklists and dashboards to engage with how people interpret, enact, and challenge risk in real time.

2.3 Communication and Challenge in Risk Culture: Underexplored Essentials

2.3.1 Rationale for Focusing on Communication and Challenge

Despite its centrality to governance rhetoric, communication and challenge remain underexplored levers in the operationalisation of risk culture. Institutions routinely assert that they promote openness, speak-up cultures, and robust escalation frameworks, but few examine the conditions under which *risk voice*, the raising and escalation of risk-relevant concerns is generated, legitimised, or suppressed. Although the FSB has not formally defined the terms "communication" and "challenge" within its risk culture guidance, this review draws upon institutional practice and behavioural research to offer working definitions. Communication within a risk culture is understood as the transparent, inclusive, and continuous flow of information across all hierarchical levels of the institution, enabling the early and unimpeded identification, discussion, and escalation of risk concerns. Challenge refers to the organisational capacity to support and value alternative viewpoints, critical questioning, and dissenting perspectives, sustained through both formal mechanisms (governance structures, decision-making protocols) and informal norms (psychological safety, modelling of desired behaviours).

Research on risk culture converges on a central claim: challenge behaviours are produced by local conditions, not by declared policies. Sheedy and Griffin (2018) operationalise this through "Avoidance", the tendency to ignore or downplay risks, which they identify as a key marker of cultural failure and link to local modelling and psychological safety rather than to formal codes. Governance-oriented critiques sharpen the picture: McConnell (2013) shows that firms espousing openness to dissent frequently lack the structures, incentives, or relational trust to embed it, producing what he terms "ritualised alignment" in which staff learn which truths are safe to surface; Stulz (2016) adds that risk managers, despite formal authority, typically hold limited power to challenge business decisions. Behavioural and social-psychological work reinforces the same mechanism from a different angle, with willingness to speak up shown to be socially contingent on inclusion, identity safety, and normative support (Jetten and Hornsey, 2014), and psychologically safe dissent shown to improve integrative thinking and risk calibration in high-stakes settings (De Dreu, Nijstad and Van Knippenberg, 2008), positioning challenge as strategically necessary for decision quality rather than merely tolerable. Silence, in turn, is theorised as an active, motivated response to perceived relational risk, governed by managerial receptivity, psychological safety, and fear of retaliation (Morrison, 2014); a decade-on review confirms that hundreds of subsequent studies have replicated these findings while exposing persistent gaps, notably the interaction of formal voice channels with informal challenge behaviour and the continued neglect of prohibitive voice on ethical and risk-related issues as a distinct construct (Morrison, 2023). The present study responds to these gaps by treating risk voice as a behavioural process that can be distorted before articulation, through escalation friction and feedback suppression, and by developing indicators (EF, FS and VL) that track how risk-relevant concerns are generated, legitimised, or silenced across formal and informal channels.

2.3.2 Empirical Insights and Limitations

The theoretical salience of communication and challenge in risk culture is widely affirmed (Renn, Klinke and Van Asselt, 2011a; Kaplan and Mikes, 2012). However, empirical research into how these behaviours operate in organisational practice remains limited, particularly in high-pressure, regulated environments. Sheedy and Griffin (2018) demonstrate that employees who score low on "Avoidance", indicating comfort with voicing concerns, correlate with fewer instances of regulatory breach and misconduct, positioning challenge capacity as a predictive proxy for risk maturity. However, their study acknowledges a methodological limitation: cross-Sectional surveys cannot reveal team-level silences, discursive suppression, or moment-to-moment communication breakdowns. Hardy and Maguire (2010) argue that dominant organisational discourses predetermine what is "speakable," filtering out dissent through strategic framing, so that tools focused solely on documented procedures underrepresent cultural pressure, tone, and social legitimacy. Jurisdictional evidence corroborates this, such as the Central Bank of Ireland's (2018) review of five retail banks found persistent behavioural barriers, including senior managers deferring routine decisions upward, producing 'congestion in decision-making' that diluted accountability and obscured risk signals from frontline staff.

McGoey (2012) introduces the concept of strategic ignorance, describing how organisations selectively omit or deflect uncomfortable knowledge, silence becomes not a lapse but a practised routine that frequently escapes both empirical measurement and supervisory detection. Knowledge sharing studies offer complementary insights, as Wasko and Faraj (2005) find that communication relies less on incentives and more on reciprocity norms and perceived alignment of values, undermining assumptions that challenge can be

invoked via policy mandates alone. Supervisory bodies have begun to note the same limitations, as the FSB (2014) concedes that assessing culture demands multiple lenses and warns that frameworks should not be treated as checklists, as cultural indicators often fail to detect erosion until it manifests in external failure. Together, the empirical terrain surrounding communication and challenge remains underdeveloped, technically constrained, and behaviourally incomplete.

2.3.3 Toward a Deeper Understanding: Significance for Risk Culture Enhancement

This study fills a critical gap, moving from conceptual endorsement to behavioural clarity, by diagnosing what prompts voice, sustains challenge, and triggers corrective action across organisational contexts. Sheedy and Griffin (2018) demonstrate that "Avoidance" scores, reflecting suppressed voice and weak challenge, are the strongest predictors of rule-bending and observed misconduct, silence, in this sense, is not merely passive but an active signal of vulnerability. Kunz and Heitz (2021) argue that challenge behaviours must be embedded in interactive control systems, frameworks that legitimise dissent and shape its social meaning, while Adler, Forbes and Willmott (2007) hold that normative cohesion, not policy alone, enables challenge to become habitual rather than exceptional.

Regulatory momentum supports this pivot. The FSB (2014, 2023, 2024) calls for behavioural data and judgment-sensitive oversight to match the pace of strategic shifts in business models, while the DNB (2015) champions "credible and sceptical conversation" in board dialogue. These shifts mark a departure from symbolic compliance toward behavioural learning, where tone, reflexivity, and communicative permeability define supervisory relevance. By embedding challenge and communication as strategic assets,

organisations move from procedural compliance to cultural governance, where risk is not just expressed but actively interpreted. This study, therefore, charts the frontier from challenge imagined to challenge enacted, offering both theoretical insight and operational indicators to enhance cultural resilience.

2.4 Box-Ticking and Superficial Compliance

2.4.1 Critique of Regulatory Instrumentalism and Checkbox Mentality

Despite growing regulatory and scholarly interest in embedding risk culture, many institutions still default to procedural compliance over behavioural substance. This phenomenon, often labelled regulatory instrumentalism, manifests through checkbox mentalities that prioritise documented adherence over cultural embodiment. Power (1994) describes this tendency as "rituals of verification," where audit logic, polished dashboards, and policy artefacts offer symbolic reassurance while leaving informal behaviours and internal norms untouched. This tendency toward ritualistic compliance is reinforced at the systems level: a systematic review of 69 empirical ERM studies finds that risk culture 'seems hard to change' and that a compliance-focused tone from the top, combined with checklist-based risk reporting, can create a substitutive relationship that impedes employees' understanding of strategic risks and reduces open risk communication, the very behaviours that challenge-based voice depends upon (Braumann, Hiebl and Posch, 2024). McConnell (2012, 2013) similarly critiques the rise of "culture veneers," noting that firms often adopt frameworks that serve as compliance tokens rather than drivers of genuine change. Morrison and Milliken (2000) establishes the theoretical foundation for this dynamic, demonstrating that organisational silence becomes institutionalised through three reinforcing mechanisms, employee fear of negative consequences, managerial belief that speaking up is futile, and implicit norms that treat dissent as disloyal, creating a structural

barrier to change that compliance frameworks alone cannot dissolve. A systematic review of 74 case studies of institutional failure, spanning financial scandals, nuclear disasters, and corporate fraud, finds that organisational silence (employees not speaking up) and management's failure to listen are among the eight most frequently cited 'corrective culture' failures that represent missed opportunities to avert disaster before it develops (Hald, Gillespie and Reader, 2021). Ashby, Palermo and Power (2012) show that organisations simulate cultural alignment through artefacts and statements, but sidestep the contested, messy behavioural terrain underneath.

The gap between surface governance and lived experience is also quantifiable. Sheedy and Griffin (2018) reveal that even firms with robust governance elements exhibit high "Avoidance" scores, indicating suppressed challenge and muted escalation. Regulators are increasingly acknowledging these limitations: the FSB (2014, 2018, 2023, 2024) cautions that indicators "should not be treated as a checklist" and calls for supervision that engages culture holistically, not just procedurally. Similarly, APRA (2019, 2022) voices concern that institutions may fulfil cultural mandates superficially, focusing on dashboard generation and formal statements without addressing behavioural misalignment. Even technologically advanced oversight tools carry risk: Culture Conduct Dashboards (CCDs) and Management Information Systems (MIS) can become performance theatre if metrics overshadow meaning, offering governance imagery rather than governance insight.

2.4.2 Empirical Evidence and Case Studies of Superficial Risk Culture

Empirical studies and supervisory reviews increasingly affirm that many financial institutions engage in culture signalling, projecting compliance and alignment through

polished governance structures, while critical behaviours like dissent, escalation, and challenge remain suppressed or performative. Sheedy and Griffin (2017; 2018) found significant cultural divergence among seven large Australian and Canadian banks, despite comparable governance frameworks, with Canadian institutions exhibiting notably more favourable risk cultures, lower Avoidance scores, and higher rates of professional risk certification than their Australian counterparts, suggesting that national and institutional context shapes risk culture independently of structural governance design. Sheedy and Griffin (2018) found that even across seven major banks with comparable governance frameworks, 'senior leaders tended to have a rosier perception of risk culture than staff generally', with senior leaders in five of seven institutions scoring significantly more favourably on Avoidance than frontline employees, indicating that leadership-level optimism about challenge and escalation can be structurally disconnected from lived employee reality. This misalignment compromises supervisory tone and weakens feedback loops, fostering blind spots at the leadership level. Sheedy and Canestrari-Soh (2023) offer a contemporary test case in Australia's Banking Executive Accountability Regime (BEAR), finding that enhanced individual accountability does improve risk management practices and cultural tone, but that the regime's most significant unintended consequence is administrative burden that reproduces the very box-ticking behaviour it seeks to displace, illustrating how even well-designed accountability mechanisms can generate compliance performance rather than cultural transformation.

This dynamic is visible in recent crisis events, captured in case studies of the Silicon Valley Bank collapse (Federico, Napoli and Notte, 2024), Credit Suisse's governance breakdown (FINMA, 2023), and the DNB's supervisory reports (DNB, 2015, 2023). All reveal the same pattern of formal obligations, committees, codes, with management

information systems nominally fulfilled, yet escalation and dissent remained dormant or discouraged. The challenge failures in these institutions were not the result of regulatory absence, but of internal misalignment, in which governance structures were adopted symbolically and never fully embedded. A parallel literature documents how the routine instruments of assurance can themselves reinforce this symbolic layer: audit rituals, ethics rollouts, and sentiment surveys have been shown to mask institutional silence and normalise passive obedience and avoidance (Magee, Schilling and Sheedy, 2019; Liu, Liu and Sathye, 2021), while regulators, on the evidence of public Final Notices, typically infer cultural misalignment only after the damage has materialised (Ring *et al.*, 2016). McGoey (2012, 2016) theorises this wider pattern as strategic ignorance, a form of organisational blindness in which inconvenient cultural knowledge is avoided to preserve strategic narrative, letting plausible deniability operate as an adaptive coping strategy.

2.4.3 Implications: Risk Entropy and Organisational Decay

The cumulative effect of symbolic governance contributes directly to a phenomenon first introduced in the practitioner literature by the author (Pharoah, 2025) and further developed in this thesis as **risk entropy**, the gradual unravelling of organisational attentiveness, behavioural vigilance, and cultural integrity. As Power (1994, 2004) conceptualises, this state emerges when organisations are flooded with documentation and assurance rituals that simulate oversight but suppress challenge, delay action, and dilute ownership. These "rituals of compliance" create a false sense of stability while masking cultural deterioration. Over time, systems affected by risk entropy begin to appear governed while losing their capacity to interpret, escalate, or contest risks meaningfully. Challenge suppression becomes normalised, dissent becomes discomforting, and governance loses its reflexivity.

Risk entropy arises not from missing controls but from failed behavioural activation. Empirical research by Klinke and Renn (2021) and Wiedemann et al. (2024) shows that systemic risk accumulation stems from organisations that favour harmony over inquiry and preserve consensus by discouraging challenge. Even promising culture efforts may fall prey to entropy: McConnell (2012, 2013) warns that in the absence of behavioural feedback and reinforcement, cultural initiatives often regress into symbolic gestures, rendering institutions vulnerable to failure without visible warning signs. Supervisory bodies, including APRA (2022), the BIS (2024), and OSFI (2024), now emphasise that behavioural embedding, not structural form, is essential to counter systemic fragility. Challenge suppression and communication decay are early-stage cultural fault lines: if made measurable, through indicators tracking escalation latency, feedback distortion, and voice legitimacy, they could serve as high-resolution diagnostics for risk entropy.

Paul Lockhart, in 'A Mathematician's Lament' (2009), argues that mathematics education fails when students are trained to arrive at correct answers without experiencing the discovery process that gives those answers meaning. The parallel for risk culture is exact, such as organisations that produce compliant outputs, signed-off reports, closed findings, and green RAG statuses, without the underlying process of genuine deliberation, challenge, and sense-making, are not managing risk. They are performing it, and it is precisely this substitution of process for outcome that risk entropy describes, and that makes its early detection so diagnostically urgent.

2.5 The Problem of Risk Entropy in Organisations

2.5.1 Conceptualising Risk Entropy

Risk entropy refers to the gradual deterioration of organisational attentiveness and behavioural resilience, a phenomenon in which institutions drift from meaningful risk engagement toward symbolic oversight and superficial calm. The term was introduced by the author in practitioner-oriented work published by Starling Trust (Pharoah, 2025) and is developed here as a theoretically grounded, empirically operationalised construct, building on the foundational work of Power (1994, 2004). The concept captures how routine practices and formal compliance rituals can create a false veneer of control, masking a deeper erosion of challenge, vigilance, and accountability. Subsequent chapters operationalise this erosion behaviourally, using escalation friction, feedback suppression and voice legitimacy to trace how risk-relevant concerns are delayed, diverted or muted within such ritualised environments. As Sheedy and Griffin (2018) show, high "Avoidance" scores, indicating suppressed escalation, correlate with rule bending and overconfidence, even in environments with robust governance structures. Entropy does not manifest through overt failure but through superficial harmony, where cultural decay is hidden beneath procedural success.

The impact of entropy deepens when cultural initiatives lack reinforcement or behavioural anchoring. McConnell (2013) cautions that when tone and messaging are symbolic rather than embedded, organisations decay quietly, losing their capacity for challenge, transparency, and ethical reflexivity despite outward signs of alignment. Power (2012, 2004) notes that rituals of oversight, particularly dashboard-style reporting and static indicators, may obscure deterioration, acting as "mechanisms of comfort" that simulate vigilance while discouraging uncomfortable reflection. More insidiously, entropy may be cultivated deliberately: McGoey (2012, 2016) introduces strategic ignorance as a form of institutional unknowing, not a failure of competence, but a defence against

cognitive dissonance, allowing risk concerns to drift into normalised silence. Hardy and Maguire (2010) further argue that as risk discourse is shaped by power structures, risk narratives outside accepted logic disappear from formal governance altogether through discursive silencing.

2.5.2 Cultural Drift and Systemic Vulnerabilities

Cultural drift is the silent architecture of organisational decline, describing the slow divergence of actual risk behaviours from espoused principles. Unlike sudden breakdowns, drift operates gradually, fuelled by misalignment, passivity, and ritualised governance, forming the behavioural bedrock of risk entropy. Klinke and Renn (2021; 2011) emphasise that systemic risks often arise not from governance voids, but from cultures that grow inattentive to complexity and hostile to dissent. Pan, Siegel and Wang (2017) observe that organisational risk profiles tend to persist across leadership changes, revealing deep-seated cultural inertia: while stability may appear beneficial, it becomes brittle when contextual demands evolve, leaving legacy cultures misaligned with emerging realities.

Kunz and Heitz (2021) caution that under pressure, risk cultures can regress into compliance formalism, disengaging staff and reinforcing echo chambers, so that control systems ossify and firms become blind to surprise risks. Ashby, Palermo and Power (2012) document how risk culture programmes lose traction once internal champions leave or initiatives lack structural reinforcement, culture fades not through confrontation, but through abandonment. Sheedy and Griffin (2018) reveal that even within a single institution, subcultures may evolve in Divergent directions; without active measurement and reinforcement, entropy takes hold at the organisational periphery, manifesting in functional misalignment, inconsistent escalation, and disparate tone perceptions. Bromiley

et al. (2015) further argue that formal ERM systems may stabilise around outdated threat profiles, encouraging staff to disengage when routines feel irrelevant, resulting in strategic blind spots where risks mutate silently. Systemic vulnerability is not just a failure of control; it is a failure of cultural continuity.

2.5.3 Evidence of Increasing Risk Entropy and Systemic Failures

Evidence across financial and supervisory domains increasingly reveals that risk entropy is not a theoretical abstraction but a recurrent precursor to systemic failure. During phases of apparent stability, both institutions and regulators tend to rely on "comfort mechanisms" and governance proxies, failing to interrogate whether internal behaviours reflect meaningful risk intent (Palermo, Power and Ashby, 2017). Sheedy and Griffin (2018) show that business units exhibiting high Avoidance scores were more likely to report noncompliance and unjustified risk optimism, following formal governance but tolerating cultural misalignment, thereby reinforcing the thesis that entropy can lurk beneath structural conformity. McConnell's (2013) review of UBS documents how risk escalation failures and challenge avoidance, though structurally anticipated, were repeatedly ignored, representing not purely financial losses but a breakdown in institutional reflexivity, tone transmission, and ethical durability. Wiedemann et al. (2024) confirm that critical anomalies are routinely overlooked not due to the absence of control, but due to erosion of challenge, suppression of discomfort, and organisational overconfidence.

This fragility has led regulators to adopt more behaviourally attuned frameworks. APRA (2022) emphasises that employees require psychological safety to escalate uncomfortable issues; without it, voice diminishes, challenge recedes, and entropy deepens. The DNB (2023) affirms that open dialogue across organisational layers is

essential to maintain interpretive clarity and prevent drift. These sources reveal a repeating pattern: institutions do not collapse because they lack rules, but when ritualised governance replaces behavioural responsiveness and when challenge is symbolised rather than enacted. Recognising entropy as a behavioural failure, not just a structural oversight, reshapes how regulators and institutions should respond.

2.6 Gaps in the Literature

2.6.1 Limited Empirical Research on Actual Practices of Communication and Challenge

Although theoretical models and supervisory frameworks consistently emphasise open communication and constructive challenge as hallmarks of resilient risk culture (Renn, Klinke and Van Asselt, 2011b; Sheedy and Griffin, 2018), empirical research examining how these behaviours unfold within organisations remains fragmentary and underdeveloped. Most existing studies rely on structural proxies, whistleblower channels, escalation procedures, or board committee mandates as markers of communicative and challenge capacity (Li *et al.*, 2021). These formal indicators tell us little about the relational nuances, psychological safety, and behavioural context within which challenges either thrive or falter. Sheedy and Griffin (2018), despite conducting a large-scale survey of 30,000+ employees, caution that cross-Sectional instruments cannot adequately capture dynamic, situated risk communication, especially regarding escalation hesitancy, selective voice, and silencing cues.

Hardy and Maguire (2016) demonstrate how organisational discourse structures define what can be spoken, and what remains unthinkable; risk topics outside accepted frames are often sidelined through discursive suppression, rendering genuine challenge

behaviour invisible to traditional survey or reporting tools. McGoe (2012) introduces "strategic ignorance," in which firms institutionalise unknowing through subtle routines and avoidance strategies, which enable silence to be not merely absence but behaviour. The challenge for empirical researchers lies in detecting non-articulated dissent and interpreting deflection behaviours that evade visible speech acts or formal escalation paths. Existing survey-based diagnostics, while useful for gauging the general climate, are structurally limited in detecting pre-articulation suppression and escalation friction; they measure whether staff feel able to speak in local contexts, not whether risk voice travels intact into formal governance channels. Considering these limitations, this thesis proposes to address this empirical gap through three behaviourally grounded indicators, escalation friction, feedback suppression, and voice legitimacy, selected precisely because they target the pre-articulation, relational, and discursive dimensions of challenge that structural proxies and cross-sectional surveys cannot reach. Together, they aim to illuminate the invisible scaffolding of organisational learning, dissent, and adaptation.

2.6.2 Underdeveloped Indicators for Deep Cultural Dynamics

Although the importance of risk culture is widely acknowledged, the tools used to assess cultural depth remain underdeveloped and predominantly structural or statistical. Current frameworks lean heavily on quantitative proxies, such as the number of reported incidents, frequency of risk committee meetings, or compliance test scores (de Araújo Lima, Crema and Verbano, 2020), that are useful for tracking procedural consistency but fail to capture the qualitative, behavioural dynamics that underpin resilient risk cultures, such as psychological safety, escalation trust, interpretive clarity, and openness to dissent. McConnell (2013) critiques most assessment models for lacking behavioural granularity, focusing on artefacts rather than experiences, and calls for indicators that reflect how

employees interpret risk-related interactions, whether policies are formally present matters less than how staff understand and experience their safety and legitimacy.

Many risk culture indicators fail to connect control systems to dynamic behaviour, underestimating how real teams adapt, interpret, and respond in live conditions (Kunz and Heitz, 2021). Without feedback loops and behavioural reflexivity, assessments risk creating an illusion of control while masking cultural drift. Few frameworks distinguish between categories of risk, controllable versus judgmental, which require diverse cultural responses: some demand procedural assurance, others require escalation comfort, discursive flexibility, and psychological fortitude (Kaplan and Mikes, 2012). APRA (2022) emphasises that entities must "understand the level of understanding" across teams; without insight into cultural comprehension, institutions may miss how risk is internally interpreted, tolerated, or misread. The literature reveals a gap between formal structures and behavioural realities. This thesis responds by treating voice legitimacy as a heuristic operationalisation of two distinct theoretical conditions, at the psychological level, Edmondson's (1999) psychological safety model, extended by Morrison's (2014) dual safety-efficacy calculus, which together explain why individuals judge speaking up as safe and likely to be effective and at the discursive level, Hardy and Maguire's (2016) account of how dominant organisational narratives predetermine whose voice is institutionally credible, a condition psychological safety alone cannot explain. Escalation friction is treated as a measure of challenge viability under real organisational conditions, and feedback suppression as a marker of deterioration in adaptive tone. Three indicators are selected because they map directly onto the qualitative, relational dynamics that quantitative proxies consistently miss.

2.6.3 The Need for Deeper Theoretical and Empirical Foundations

The literature on risk culture is rich in aspiration but remains methodologically fragmented and theoretically constrained. While concepts such as psychological safety, leadership tone, and social norms are widely recognised as enablers of challenge behaviour (Edmondson, 1999), few studies have systematically tested these mechanisms in live organisational settings, especially across industry types, governance contexts, and regulatory environments. Organisational theory must prioritise practical rationality, how people make sense of institutional life, over externally imposed categories, a gap particularly striking in risk culture research (Sandberg and Tsoukas, 2015).

Power asymmetries, trust gradients, and local norms strongly influence escalation choices, which vary significantly across geographies, organisational types, and maturity levels, underscoring the need for context-sensitive, longitudinal inquiry (Li *et al.*, 2021; Wiedemann *et al.*, 2024). They rarely distinguish between the voice conditions experienced by senior or tenured staff and those of junior or newer employees, despite consistent evidence that hierarchical position shapes both perceived safety and the reception of dissent. Power, Ashby and Palermo (2013) state that culture evaluation is routinely "shoehorned into audit categories," ignoring its embedded, informal, and multilevel nature, and call for a more reflexive foundation based on behavioural observation, interpretive analysis, and sensemaking.

What is missing is a comprehensive framework linking micro-level enablers to macro-level governance outcomes. This thesis addresses that gap by proposing a heuristic indicator framework in which escalation friction, feedback suppression, and voice legitimacy are treated as measurable distortions of risk voice, with risk entropy inferred as

the emergent condition when these distortions co-occur. The selection of these three indicators is grounded in both regulatory definition and academic evidence. The FSB (2014) defines effective communication and challenge as requiring that decision-making processes allow for testing of current practices, that alternate views be expressed across all layers, and that open and constructive engagement be actively maintained, with criteria that map directly onto escalation friction, voice legitimacy, and feedback suppression, respectively. Each indicator is further supported by the academic literature: escalation friction responds to the discursive suppression identified by Hardy and Maguire (2016) and the avoidance patterns evidenced by Sheedy and Griffin (2018); feedback suppression operationalises the strategic ignorance theorised by McGoe (2012) and the absence of feedback loops critiqued by Kunz and Heitz (2021); voice legitimacy addresses the identity safety and power asymmetry conditions that Jetten and Hornsey (2014) and Wiedemann et al. (2024) identify as determinants of whether challenge is enacted at all, with these conditions theorised here as a social-identity sub-dimension of Edmondson's (1999) psychological safety framework rather than as an independent theoretical anchor. Future research must embrace multi-method designs, behavioural observation, longitudinal tracking, and multilevel modelling to capture how challenging behaviour propagates across institutional layers.

2.6.4 The Critical Gap: Linking Communication and Challenge to Measurable Outcomes

Although communication and challenge are widely recognised as central tenets of risk culture, empirical tools that connect these behaviours to tangible governance outcomes remain underdeveloped. Most existing studies rely on correlational surveys or descriptive assessments, offering conceptual affirmation but lacking causal validity or operational

applicability. Sheedy and Griffin (2018) provide empirical traction, showing that business units with high "Avoidance" scores report more unreported breaches and unjustified optimism; yet few diagnostics exist to translate these behavioural pathologies into actionable insights or early warning indicators. De Dreu and van Knippenberg (2008) demonstrate that functional dissent enhances decision quality in groups equipped to tolerate ambiguity and engage constructively, but governance models rarely incorporate such findings, leaving challenge acknowledged but unmeasured in supervisory design.

Without validated, behaviourally grounded indicators, such as challenge frequency, escalation responsiveness, or perceived safety scores, regulators and institutions cannot monitor, benchmark, or evolve their governance practices with confidence. Voice is shaped by social identity and norm signals, yet current frameworks fail to capture these antecedents, limiting the predictive utility of challenge diagnostics (Jetten and Hornsey, 2014). This thesis aims to bridge that critical gap: converting conceptual insights into measurable indicators of communicative and challenge dynamics and testing their validity in live organisational contexts. This thesis operationalises that ambition through three specific indicators, escalation friction, feedback suppression, and voice legitimacy, whose theoretical grounding and selection rationale are set out in Sections 2.6.1–2.6.3 above. These tools promise to support targeted interventions, supervisory precision, and cultural resilience, enabling firms to learn, adapt, and lead with greater integrity amid complexity.

2.7 Chapter Summary

This chapter has mapped the conceptual terrain of risk culture, from its definitional multiplicity across regulatory and academic discourse, through the theoretical models that frame formal and emergent culture, to the critically underexplored dimensions of

communication and challenge, and finally to the phenomenon of risk entropy as an organising diagnosis for cultural fragility. Several clear themes emerge.

First, definitions of risk culture remain contested and inconsistently applied, yet converge on the centrality of behavioural norms, shared perceptions, and the conditions under which challenge is voiced or withheld. Second, existing frameworks, while increasingly sophisticated, tend to overemphasise structural compliance, creating the conditions for the "culture veneer problem" McConnell identifies: governance that is declared rather than demonstrated. Third, communication and challenge, despite being formally endorsed by the FSB and major supervisory bodies, receive scant empirical attention, and the tools available to diagnose their presence or absence remain inadequate. Fourth, the concept of risk entropy provides a theoretically grounded and operationally useful construct for capturing the gradual, cumulative erosion of organisational attentiveness that these diagnostic gaps allow to persist undetected. It is on the basis of these four themes that this thesis selects escalation friction, feedback suppression, and voice legitimacy as its core diagnostic indicators, with each addressing a gap identified in the literature, each is theoretically traceable to the mechanisms through which risk entropy emerges, and together they constitute a framework for making the invisible visible.

These gaps in the literature directly motivate the research questions and methodological choices developed in Chapter III.

CHAPTER III: METHODOLOGY

3.1 Overview of the Research Problem

Financial institutions have sophisticated governance structures, but risk failures often stem from muted communication and suppressed challenge rather than flawed policies. This study explores how everyday practices of voicing, withholding, and escalating risk concerns reveal cultural vulnerabilities underlying formal compliance; these practices are conceptualised through three proximate indicators (EF, FS, VL) and an emergent synthesis condition (RE), each defined in Section 3.3. The selection of these indicators is grounded in the literature analysis in Chapter II (Sections 2.6.1–2.6.4), which identifies escalation friction, feedback suppression, and voice legitimacy as the most salient behavioural distortions of risk voice and conceptualises risk entropy as the emergent condition when these distortions co-occur.

Regulators emphasise communication and challenge as key pillars of risk culture, but diagnostics remain limited to surveys and disclosures that often overlook local differences in how staff experience escalation pathways or perceive speaking-up legitimacy. An embedded single-case study within a Channel Islands deposit-taker examines patterned differences across governance units (Board/SMT, Risk/Compliance, First-Line Business, Operations/Support), despite uniform structures. Interviews (n=25), artefact analysis, and reflexive memos map how EF, FS, and VL cluster, identifying where RE emerges as vigilance declines.

3.1.1 Philosophical Positioning

The study adopts an interpretivist-constructivist paradigm because the research aim is to explain how communication and challenge operate within a specific financial institution, how they are experienced, made meaningful, and enacted differently across governance units. This stance directly shapes the choice of semi-structured interviews, organisational artefact analysis, and reflexive explanation-building described in Sections 3.5 and 3.6, each of which enables contextualised, applied insights into governance practice (Yin, 2009), with Eisenhardt's (1989b) logic of theory building from case studies informing the embedded single-case design described in Section 3.2.1.

Ontologically, the study adopts a relativist stance in which risk culture has no single fixed reality but is shaped by local histories, power relations, and role expectations; epistemologically, knowledge of EF, FS, and VL is co-constructed by researcher and participants rather than observed from a neutral standpoint, with reflexive journaling used to document how interpretations are reached (Guba and Lincoln, 1994). Full exposition of the ontological and epistemological positioning is provided in Appendix E.1.

3.1.2 Interpretivist Focus on Communication and Challenge

Within this philosophical framework, risk culture is regarded as the interpretive infrastructure through which communication and discussion about risk occur. This infrastructure encompasses the stories, categories, metaphors, and conversational routines that influence whether concerns are voiced at all (FS), how they evolve or remain static within the system (EF), and how legitimate the speaker and their concern are perceived to be (VL), each detailed in Section 3.3.

Sense-making and narrative processes are therefore essential. Employees continually evaluate how dissent will be interpreted, whether as constructive engagement, drawing on past responses, or as a matter of identity safety (Jetten and Hornsey, 2014, an identity dimension of Edmondson's (1999) psychological safety condition). At the discursive level, dominant organisational narratives structure which voices are regarded as institutionally credible: Hardy and Maguire (2016) show that what can be spoken, and by whom, is pre-determined by power-structured discourse boundaries, with the primary theoretical anchor for the VL indicator's second dimension. Leaders develop accounts that may encourage participation or implicitly suppress it (Palermo, Power and Ashby, 2017; Hardy *et al.*, 2020), while groups collectively invoke strategic ignorance to avoid troubling knowledge (McGoey, 2012, 2016), each understood here as an effect of discursive power rather than a separate theoretical framework.

3.1.3 Researcher Positionality and Reflexive Stance

Reflexive thematic analysis, as Byrne (2022) demonstrates in a worked example of Braun and Clarke's approach, treats the researcher's subjectivity not as a contaminant to be eliminated but as a resource that shapes how data are engaged with, coded, and interpreted. This study explicitly adopts that stance, with the researcher an external risk governance practitioner with no employment, advisory, or reporting relationship with the case organisation (the full independence statement appears in Appendix E.2). Professional familiarity with regulatory vocabulary and the three lines model of governance brought interpretive sensitivity to escalation dynamics and challenge norms, but also risked normalising certain institutional practices as unremarkable. A reflexive journal and

structured post-interview memos (Section 3.5.5) served as the primary mechanisms for surfacing and interrogating these pre-understandings throughout fieldwork and analysis.

A concrete example of this corrective process, the reclassification of 'appropriate gatekeeping' to 'informal gatekeeping (EF)' following MemoINTMS09, and the full independence statement are documented in Appendix E.2.

3.2 From Philosophy to Research Design

The study uses an exploratory, inductive qualitative approach to uncover how communication and challenge function within a specific financial institution. Instead of starting with fixed hypotheses about the frequency or impact of EF and FS, the study begins with open, theory-informed questions about when, how, and why individuals raise or withhold risk-related concerns, how they attempt to challenge decisions, and how escalation pathways are experienced in practice. The goal is to develop and refine a heuristic indicator framework in which EF, FS, and VL are proximate mechanisms and RE is inferred as an emergent synthesis condition from their cross-unit configuration.

3.2.1 Embedded Single-Case Design

Relativism promotes choosing a single financial institution with multiple embedded units, maintaining contextual depth while enabling comparison of how communication and challenge are locally enacted across divisions, jurisdictions, and roles (Guba and Lincoln, 1994; Yin, 2009). The four embedded units (Board/Oversight/SMT, Risk and Compliance, First Line Business, and First Line Operations) facilitate examination of patterned variation in how risk-related voice, silence, and dissent are constructed across different

organisational contexts, while avoiding the decontextualisation that would occur from aggregating experience into a single organisation-wide score.

This study adopts an embedded single-case study approach to examine how EF, FS, and VL manifest within a complex financial institution, with RE inferred as an emergent condition from their cross-unit configuration rather than coded directly at the event level. Case-study methodology is particularly suited to "how" and "why" questions where phenomena cannot be meaningfully separated from their institutional context, such as governance frameworks, regulatory requirements, and psychological safety structures (Yin, 2009). Following Eisenhardt's (1989b) inductive logic, the four heuristic indicators (the three proximate behavioural indicators and the emergent RE condition) are developed through iterative engagement with the data rather than imposed beforehand, with theoretical sampling, replication logic across embedded units, and multi-source triangulation underpinning the within-case comparison carried out here.

The embedded configuration distinguishes governance units that operate under formally uniform arrangements but may experience escalation and challenge differently. This allows for detailed, context-aware comparison of local risk-culture dynamics that a solely holistic organisational profile would obscure (Kallio *et al.*, 2016; Adeoye-Olatunde and Olenik, 2021; Ruslin *et al.*, 2022). Case-boundary decisions and unit-level inclusion criteria, informed by Baxter and Jack (2008) and Baskarada (2014), are detailed in Appendix F.2 and F.3.

This design is intentionally distinct from a formal multi-case study, in which multiple organisations would each constitute a separate case with independent sampling

frames and cross-case replication logic (Eisenhardt, 1989b; Yin, 2009). The four embedded units here share a common regulatory framework, governance architecture, and organisational history; they are treated as analytically distinct sub-contexts within one case, not as separate cases. The embedded single-case design is therefore a deliberate methodological choice, offering contextual depth and analytical richness that a multi-organisation design would risk sacrificing.

Following the logic of theory building from case studies (Eisenhardt, 1989b), this embedded structure provides a framework for systematic within-case comparison whilst maintaining contextual richness, illustrating how patterns of EF, FS, and VL across units contribute to emerging risk entropy within a formally well-governed organisation.

3.2.2 Case Selection and Justification

The institution was chosen as the single case for three connected reasons. First, it publicly aligns with the Financial Stability Board (FSB, 2014) risk culture principles and related supervisory frameworks, maintains mature governance arrangements, and makes explicit commitments to psychological safety. These features indicate organisational intent to promote communication and challenge, creating an ideal context for exploring how these goals are realised in daily practice.

Second, internal diagnostics and prior survey work reveal significant variation in escalation viability and voice legitimacy across business divisions, geographies, and management levels, despite uniform formal structures. This documented divergence supports broader evidence that risk culture can differ substantially within the same firm, with some units demonstrating high levels of challenge and others showing avoidance or

deference. The institution is therefore well positioned to identify early signs of RE (Section 3.3) across its embedded governance units.

Third, the institution provided extensive research access, allowing interviews across the three lines model, direct interaction with governance teams, and review of organisational artefacts, including committee minutes and internal frameworks. Such multi-source access is uncommon in sensitive risk-culture research but is strongly advised in qualitative governance work to improve credibility and triangulation. This access was facilitated in part by a pre-existing professional network relationship with the Chief Risk Officer, whose prior familiarity with the researcher supported the trust required for institutional approval. This relationship is acknowledged as a positionality dimension and is addressed in the independence statement at Appendix E.2.

3.2.3 Access Negotiation and Regulatory Context

Access was negotiated directly with the Chief Risk Officer over a four-week period (August–September 2025). During this time, the proposal was examined by senior governance bodies and legal and compliance functions to ensure it aligned with confidentiality obligations, data protection requirements, and ethical protocols. These negotiations determined the final scope of the fieldwork and informed the safeguards outlined in Section 3.9.

The institution's regulatory positioning strengthens its suitability as a theoretically purposeful case. As a licensed deposit-taker operating in the Channel Islands (Jersey and the Isle of Man), it is subject to FSB-aligned expectations on risk culture and governance (FSB, 2014). This combination of stringent regulatory oversight and explicit cultural

aspirations makes the organisation a particularly revealing setting for studying how regulation-prescribed cultural mechanisms operate in practice.

3.2.4 Case Boundaries: Temporal, Spatial, and Functional Definition

Fieldwork was bounded temporally (21 October–12 November 2025, aligned with the governance meeting cycle), geographically (Jersey and the Isle of Man, covering headquarters, retail banking, private banking, compliance, and operations), and conceptually (communication and challenge behaviours within the formal governance framework, excluding outsourced functions such as external audit). These boundaries were designed to capture a consistent organisational moment across the three lines model without triggering year-end resource constraints. Full boundary definitions and exclusion criteria are documented in Appendix F.1.

3.2.5 Embedded Units: Structure, Definition, and Inclusion Criteria

To facilitate systematic within-case comparison while maintaining fidelity to the three lines model, the case is organised into four embedded units: Board/Oversight/SMT; Risk and Compliance (Second Line); First Line Business Divisions; and First Line Operations and Support Functions. Each unit is defined by formal responsibility for risk identification, oversight, and escalation, and was staffed to ensure analytical adequacy rather than proportional headcount. Inclusion criteria and the full unit-level target table are provided in Appendix F.2 and F.3; Figure 1 is reproduced there.

3.2.6 Design Implications for Sampling and Analysis

The embedded-unit structure guides two subsequent methodological choices. First, it directs a stratified purposive sampling procedure where the sampling strata reflect the

embedded units. Within-unit quotas guarantee representation across different seniority levels and functions. Theoretical sampling extensions are initiated if certain units continue to display thematic diversity or if cross-unit comparisons uncover unexpected patterns; unit definitions and sampling targets for each embedded unit are provided in Appendix F.3.

Second, the design promotes a within-case explanation-building approach that concentrates on how EF, FS, and VL cluster within and across units, with RE inferred as an emergent condition from these cross-unit configurations, rather than on pinpointing a single organisation-wide risk-culture profile. This unit-based perspective enhances causal inference by uncovering localised cultural vulnerabilities that endure despite uniform formal governance.

Each unit's saturation criteria, defined as thematic redundancy across its final two interviews, are operationalised in Section 3.4.3, which details the dual-level monitoring protocol and documents how interpretive completeness was assessed.

3.3 Construction and Operationalisation of Heuristic Behavioural Indicators

The three heuristic behavioural indicators of EF, FS, and VL function as proximate behavioural mechanisms that shape how risk-related concerns are raised, heard, and acted upon, whereas the fourth indicator, RE, is an emergent synthesis condition that arises from particular configurations of EF, FS, and VL rather than a fourth independent mechanism. Together, all four are diagnostic constructs adapted from FSB (2014) supervisory expectations on risk culture and from the literature-derived rationale developed in Chapter II (Section 2.6), where EF, FS, and VL are justified as core distortions of effective communication and challenge. Developing them directly responds to calls for culture

diagnostics that go beyond symbolic disclosures and risk appetite statements towards observable patterns of communication and challenge. None are fixed theoretical variables; all were refined iteratively through engagement with data and are grounded in participants' sense-making and organisational artefacts rather than numerical thresholds, supporting interpretivist explanations consistent with the study's relativist ontology and subjectivist epistemology (Guba and Lincoln, 1994; Braun *et al.*, 2019a).

3.3.1 Heuristic Behavioural Indicator Framework

The framework maps the four constructs onto observable communication and challenge behaviours, typical triggers, and documentary traces, with a full overview provided in Appendix G.1. Each indicator is defined conceptually and operationalised through a set of discriminating behavioural cues, with evidence sources and triangulation logic documented in Appendix D, and the full research question-to-indicator mapping set out in Appendix A.2. Indicators are synthesised into unit-level interpretive profiles that reveal how local practices enable or inhibit communication and challenge, and how these configurations give rise to RE (see Section 3.3.5).

In addition to the three proximate indicators and the emergent synthesis construct (EF, FS, VL, RE), the framework includes a set of supporting mechanism labels (BRE, ICF, MRC, OAV, RC, RCF) that operate as child-level codes and contextual descriptors during analysis, helping to differentiate how the core indicators manifest in practice (for example, through tensions between risk and commercial pressures, the visibility of control functions, or the reinforcement of challenge in governance forums). Table 3.1 summarises the heuristic indicators and associated context mechanisms, linking each to its methods section, observable indicators, and primary data sources. The three proximate indicators

(EF, FS, VL) and the synthesis construct RE each receive a dedicated subsection (Sections 3.3.2–3.3.5). The six contextual mechanisms are described in full in Appendix G.1 and are introduced here as analytical context rather than as independently profiled constructs.

Table 3.1 Heuristic Indicators and Context Mechanisms
Overview linking context mechanisms (BRE/ICF/MRC/OAV/RC/RCF) — derived from the FSB's (2014) effective communication and challenge pillar, to behavioural indicators (EF/FS/VL/RE). Guides coding/triangulation.

Code	Short label	Full label	Methods Section	Observable indicators (brief)	Primary data sources
A	BRE	Balancing risk, efficiency and innovation	Section 3.3.1 and Appendix G.1	Tension between control, efficiency, and commercial innovation	Interviews, artefacts
B	EF	Escalation Friction	3.3.2	Procedural, relational, and cognitive barriers to escalation	Interviews, artefacts, memos
C	FS	Feedback Suppression	3.3.3	Ignored input, absent follow-up, tokenism, and voice fatigue	Interviews, artefacts, memos
D	ICF	Independence of Control Functions	Section 3.3.1 and Appendix G.1	Formal distance, structural authority, and separation from business	Interviews, artefacts
E	MRC	Mechanisms for Raising Concerns	Section 3.3.1 and Appendix G.1	Availability and usability of formal channels for concerns	Interviews, artefacts
F	OAV	Open to Alternate Views	Section 3.3.1 and Appendix G.1	Managerial openness to dissenting or alternate perspectives	Interviews, artefacts
G	RC	Reinforcement of Challenge	Section 3.3.1 and Appendix G.1	Active encouragement, follow-through, and learning from challenge	Interviews, artefacts
H	RCF	Representation of Control Functions	Section 3.3.1 and Appendix G.1	How control roles are portrayed and perceived in the organisation	Interviews, artefacts
I	RE	Risk Entropy Synthesis		Evidence basis: cross-unit comparison of EF, FS, and VL configurations; cross-unit explanation-building table (Appendix G.5.1); RE is inferred, not directly coded as an event-level indicator.	All sources

J	VL	Voice Legitimacy	3.3.4	Who is seen as authorised and credible to speak up	Interviews, artefacts, memos
---	----	------------------	-------	--	------------------------------

3.3.2 Escalation Friction (EF)

Escalation friction includes procedural, relational, and cognitive barriers that are interpreted here as blocking, distorting, or preventing the upward flow of risk-related concerns. It involves formal obstacles (e.g., unclear procedures, excessive paperwork), informal gatekeeping (e.g., managers filtering or reframing issues), and cognitive biases (e.g., normalisation of deviance, minimisation of problems).

- Informal gatekeeping: intermediaries (managers, supervisors) reframe or minimise issues before they reach control functions
- Selective documentation: omitting or sanitising dissenting views in minutes and escalation records
- Escalation silence: lack of concerns from formal channels despite evidence of informal raising

Evidence is drawn from interview narratives, committee minutes compared with frontline accounts, and reflexive memos; full triangulation logic is documented in Appendix D. Full operational definitions and boundary-case discrimination rules, including the distinction between EF and FS at the point of channel entry, are provided in Appendix G.2.

3.3.3 Feedback Suppression (FS)

Feedback suppression describes observed patterns interpreted as blocking the upward or sideways flow of evaluative information about decisions, processes, or

behaviours. It includes active discouragement (e.g., punitive responses, reputational costs), passive discouragement (e.g., lack of acknowledgement, absent follow-up), and structural disincentives (e.g., reward systems that favour conformity over challenge).

- Absent follow-up: raised issues are not tracked, assigned, revisited, or closed at subsequent meetings.
- Punitive signalling: dissenters face reputational costs, reassignment, or stalled career progression.
- Voice fatigue: evidence that repeated, ignored feedback leads individuals to stop attempting to escalate.

The study draws on interview accounts, action logs compared with interview claims, and reflexive memos; full triangulation logic is documented in Appendix D. Full operational definitions and boundary-case discrimination rules, including the distinction between FS and EF at the point of receipt, are provided in Appendix G.3.

3.3.4 Voice Legitimacy (VL)

Perceived safety is a necessary but insufficient condition for risk voice: an individual may feel safe to speak yet still judge that their voice carries no institutional weight. This indicator therefore operationalises two theoretical conditions, Edmondson's (1999) psychological safety model, extended by Morrison's (2014) dual safety-efficacy calculus, and Hardy and Maguire's (2016) account of discursive legitimacy, to capture both the individual's willingness to speak and the organisation's readiness to recognise that voice as credible.

Voice legitimacy concerns the perceived appropriateness and credibility of speaking up within a specific role, unit, or context. It defines who is considered authorised to raise concerns, which forms of expression are regarded as acceptable, and how legitimacy is negotiated through status, expertise, relational ties, and institutional narratives.

- Differential reception: identical concerns are treated differently based on the speaker's role, seniority, or background.
- Legitimacy narratives: explanations that frame speaking up as constructive engagement rather than disloyalty or troublemaking.
- Role-legitimacy boundaries: views that voice is appropriate in some roles or issues but not in others.

VL is evidenced through interview narratives and meeting minutes that record who speaks and how contributions are documented; full triangulation logic is documented in Appendix D. Full operational definitions and boundary-case discrimination rules, including the distinction between VL and FS where silence reflects delegitimised voice rather than suppressed feedback, are provided in Appendix G.4.

3.3.5 Risk Entropy (RE)

Risk Entropy (RE) occupies a different ontological position from the three proximate indicators. EF, FS, and VL are coded directly at the level of individual transcript passages, they describe observable behaviours or perceptions. RE is not coded at this level. It is an inferred systemic condition, attributed at the unit and case level once EF, FS, and VL profiles are complete, based on the convergent weight and cross-unit configuration of

proximate evidence. The 'RE synthesis segments' counted in Appendix I and referenced in Section 4.9 are passages where this higher-order inference is most analytically legible; they are not direct observations of entropy but passages that collectively support the inference. RE is therefore best understood as a diagnostic conclusion drawn from the EF/FS/VL evidence base, analogous to a clinical diagnosis inferred from multiple symptoms, rather than as a fourth parallel symptom observed alongside them.

Risk entropy refers to increasing disorder in organisational risk communication and governance when hazard information becomes fragmented, distorted, or suppressed across local practices. Unlike EF, FS, and VL, which are coded at the level of observable events, RE is not a direct coding destination but an interpretive condition inferred post-analysis from cross-unit configurations of the three proximate indicators. It is seen as a cumulative state rather than a single event, marked by uneven information flows, inconsistent interpretations, and weakened feedback loops that can correct misunderstandings.

RE is inferred when cross-unit analysis of EF, FS, and VL profiles reveals the following convergent patterns across units; these are not coded as a separate RE node at the event level.

- Information fragmentation: different reports of the same event, hazard, or decision across units or communication channels
- Normalisation of ambiguity: routine acceptance of unclear, incomplete, or contradictory risk information without seeking clarification
- Disconnected escalation: evidence of concerns reaching governance forums but not resulting in a coordinated organisational response

Evidence for RE is constituted through cross-unit explanation-building, comparing EF, FS, and VL configurations across embedded units, supported by cross-unit interview comparison and artefact triangulation, with the full triangulation logic is documented in Appendix D. The inferential warrant criteria, specifying under which configurations of EF, FS, and VL elevated RE is inferred, together with worked unit-level applications and disconfirming evidence considered, are provided in Appendix G.5 and G.5.1.

3.3.6 Synthesis and Use of Indicators

Together, EF, FS, VL, and RE provide a heuristic lens for diagnosing how communication and challenge operate across embedded units, rather than offering fixed numerical scores. EF, FS, and VL were used to build unit-level profiles of escalation dynamics, feedback cycles, and perceived legitimacy of voice; cross-unit comparison of these profiles informed the inference of RE as an emergent condition of cultural fragility. Detailed synthesis procedures, warrant criteria, and the explanation-building framework used to translate indicator configurations into findings are documented in Section 3.6 and Appendix G.6–G.7, with evidence base and documentation in Appendix D and codebook evolution in Appendix B.

The heuristic indicator framework operates across two analytical levels, and the distinction shapes how evidence is treated throughout this study. Escalation friction, feedback suppression, and voice legitimacy are coded at the event level, each tied to an identifiable moment in the data. Risk entropy is not coded directly; it is inferred at the cross-unit level from the configurational pattern of the three proximate indicators across embedded units, visible only when read across units. Figure 3.1 sets out the architecture,

which informs the sampling, collection, coding, and synthesis choices that follow (Section 3.4–3.7) and is returned to in Section 5.5.




Figure 3.1 Two-Level Indicator Architecture.
EF, FS, and VL are coded at the event level; RE is inferred at the cross-unit level from their configurational pattern across embedded units.

3.4 Sampling Strategy and Participant Recruitment

Stratified purposive sampling aimed to maximise interpretive completeness across all four embedded units, prioritising governance, risk, business, and operational coverage criteria for selecting interview participants. Sampling was designed to support interpretivist

explanation-building rather than statistical generalisation, emphasising interpretive completeness and capturing the full range of contextual variations in how EF, FS, and VL are experienced across embedded units, so that RE can be inferred as an emergent condition from their cross-unit configuration, rather than early stopping at thematic saturation (Guba and Lincoln, 1994; Sandberg, 2005; Braun *et al.*, 2019a).

3.4.1 Sampling Approach

A stratified purposive sampling design was adopted, aligned with the embedded single-case design and the organisation's three-line governance model (Etikan, Musa and Alkassim, 2016). Strata were defined by line of responsibility (Board/SMT, First Line Business and Operations, Second Line Risk and Compliance), business unit (Risk, Compliance, Operations, Retail and Private Banking, Finance, Marketing, External Assets, and Independent Non-Executive Directors), and geographical location (Jersey and Isle of Man). Within each stratum, participants were selected on the basis of direct involvement in risk communication and challenge, including those who raise, escalate, or receive risk-related concerns, ensuring information-rich cases for the phenomena under investigation and a diversity of perspectives on EF, FS, and VL across roles, functions, and jurisdictions (Kallio *et al.*, 2016; Guba and Lincoln, 1994).

Campbell *et al.* (2020) demonstrate through three worked case studies how purposive sampling strategies can be explicitly linked to Lincoln and Guba's four trustworthiness criteria. In the present study, credibility is served by selecting participants across all embedded units and seniority levels; transferability is supported by documenting the rationale for each inclusion decision, enabling readers to judge whether similar dynamics might exist in comparable institutions; dependability is strengthened by

consistent application of unit-level sampling quotas; and confirmability is enhanced by the transparency of the sampling framework, which is fully documented in Table 3.2 and Appendix F, applied embedded-unit quotas aligned with the indicator and context-mechanism structure introduced in Section 3.3.1. Since risk culture encompasses multiple, context-dependent meanings rather than a single measurable trait, the overarching goal was to achieve interpretive completeness across different parts of the organisation rather than stopping at the initial point of thematic saturation (Guba and Lincoln, 1994; Sandberg, 2005).

The Participant Information Sheet (Appendix A.4.) distributed to invitees described the selection process as 'stratified random sampling' to convey proportionality of coverage across governance units and seniority levels in accessible, non-technical language; the operative research design throughout this study is stratified purposive sampling as defined in this section. The PIS wording reflects a lay simplification made for participant communication purposes and cannot be amended as it forms part of the ethics consent record.

3.4.2 Sample Targets and Achieved Sample

Initial sampling targets ranged from 15 to 30 semi-structured interviews, with an indicative potential of 4 to 5 participants per embedded unit contacted, where feasible (Hennink and Kaiser, 2022). These targets were regarded as guidance rather than strict stopping rules (Braun *et al.*, 2019a). The study carried out a full enumeration of the accessible population ($n = 62$ eligible staff (24+13+25)), completed 25 interviews (40% response rate), and continued recruitment until all consenting participants had been interviewed. Table 3.2 summarises the response distribution across recruitment stages and

business unit, showing how many eligible staff were contacted, declined, or participated at each stage.

Table 3.2 *Response Distribution Across Recruitment Stages*

	No Response	Declined	Pilot	Not contacted	Accepted	Back-up	Grand Total
External Asset	1			1	1		3
Finance		1	1	2	2	1	7
Risk & Compliance	1	1	1	3	5	2	13
Marketing		1		1	1		3
Private & Personal Banking	2	1		6	2	3	14
Operations	12	8	1	32	9		62
Retail Banking	4			7	1	2	14
Human Resources		1			1		2
INED/CEO	3				3		6
Audit	1						1
Grand Total	24	13	3	52	25	8	125

Table 3.3 presents the realised sample structure by governance line and business function, highlighting how the 25 participants are distributed across oversight, second line, and first line units.

Table 3.3 *Sample Structure by Governance Line and Business Function*

Realized sample (n=25)	% of sample	Three Lines classification	Business departments included	Geographic split (Jersey/Isle of Man)
5	20%	Oversight	Board (CEO, CRO, COO INED×2)	3 Jersey, 1 Isle of Man
5	20%	2nd Line	Risk and Compliance (5)	2 Jersey, 3 Isle of
15	60%	1st Line	Operations (8), Finance (2), Human Resource (1), Marketing (1), Private and Personal Banking (1), Retail Banking (1), External Asset (1)	8 Jersey, 7 Isle of Man

25	100%	Total	9 business areas	14 Jersey (56%), 11 Isle of Man (44%)
----	------	-------	------------------	---------------------------------------

Reference counts by participant and theme, confirming coding coverage across all 25 participants, are provided in Appendix I.1. Table 3.4 provides a summary of participant characteristics by job band, function, and location, illustrating the breadth of coverage across roles and the two geographic jurisdictions.

Table 3.4 Summary of Participant Characteristics

Participant	Location	Job Band	Business Department	3 Lines Model	Senior Management Team Member	Board Member
Cases\\MS01	Jersey	SL	Risk & Compliance	Oversight	Yes	Yes
Cases\\MS02	Jersey	S4	Marketing	1st Line	Yes	No
Cases\\MS03	Jersey	SL	Board/CEO	Oversight	Yes	Yes
Cases\\MS05	Jersey	S4	Human Resource	1st Line	Yes	No
Cases\\MS06	Jersey	S5	Operations	1st Line	Yes	Yes
Cases\\MS08	Jersey	S2	Operations	1st Line	No	No
Cases\\MS09	Jersey	S2	Risk & Compliance	2nd Line	No	No
Cases\\MS10	Jersey	S3	Private & Personal Bankin	1st Line	No	No
Cases\\MS12	Jersey	S3	External Asset	1st Line	No	No
Cases\\MS15	Jersey	S2	Finance	1st Line	No	No
Cases\\MS16	Jersey	S3	Operations	1st Line	No	No
Cases\\MS20	Isle of Man	S3	Finance	1st Line	No	No
Cases\\MS21	Isle of Man	S2	Risk & Compliance	2nd Line	No	No
Cases\\MS23	Isle of Man	S4	Risk & Compliance	2nd Line	No	No
Cases\\MS29	Isle of Man	S2	Operations	1st Line	No	No
Cases\\MS37	Isle of Man	S3	Operations	1st Line	No	No
Cases\\MS39	Isle of Man	S1	Retail Banking	1st Line	No	No
Cases\\MS41	Isle of Man	S5	Operations	1st Line	Yes	No
Cases\\MS44	Isle of Man	S2	Operations	1st Line	No	No
Cases\\MS45	Jersey	S4	Risk & Compliance	2nd Line	Yes	No
Cases\\MS46	Isle of Man	S2	Risk & Compliance	2nd Line	No	No
Cases\\MS47	Jersey	SL	Board/INED	Oversight	No	Yes
Cases\\MS48	Isle of Man	SL	Board/INED	Oversight	No	Yes
Cases\\MS49	Isle of Man	S2	Operations	1st Line	No	No
Cases\\MS50	Jersey	S3	Operations	1st Line	No	No

3.4.3 Saturation Monitoring and Interpretive Completeness

Saturation was monitored at two levels: parent-theme saturation across the whole sample, reached at interview 18 and confirmed over three consecutive interviews; and unit-

level saturation within each embedded unit, assessed as thematic redundancy across its final two interviews. Because a census of all consenting participants was conducted, interpretive completeness across all four embedded units was confirmed rather than relying on early stopping. Full saturation tracking, reference-count progression, and the dual-level monitoring protocol are documented in Appendix C.1–C.6, with three-consecutive-interview confirmation checks detailed in Appendix C.2 and per-interview reference counts tabulated in Appendix C.6.

3.4.4 Non-Response Considerations and Self-Selection

Of the 62 eligible staff contacted, 37 did not respond or declined to participate. Because this study investigates voice and suppression, non-participants may disproportionately include those whose experiences of FS and low VL are most acute, potentially resulting in a conservative picture of EF, FS, and VL in lower-participation units (Morrison and Milliken, 2000). Four design features partially offset this risk: embedded-unit coverage ensured no governance function was unrepresented; artefact triangulation provided suppression evidence independent of self-report; reflexive memos treated absence as analytically meaningful; and minimum per-unit thresholds preserved cross-unit comparison. Full analysis of non-response patterns and self-selection effects is provided in Appendix C.7; the caveat is revisited in the limitations discussion in Chapter 5.

3.5 Data Collection and Fieldwork Operations

Fieldwork rested on three overlapping data streams, gathered concurrently and read against one another throughout the study. Semi-structured interviews revealed how participants experienced escalation, challenge, and voice; organisational artefacts anchored

those accounts in the formal governance record; and reflexive memos tracked the researcher's evolving interpretations as patterns took shape across units. The three proximate indicators and the emergent synthesis condition were handled differently within this design: EF, FS, and VL were coded at the event level, each tied to an identifiable moment of escalation, challenge, or voice; RE was not coded directly but inferred as an emergent condition from cross-unit configurations of those three proximate indicators, becoming visible only when patterns were read across the embedded units.

The subsections that follow set out the collection strategy, participant selection, sampling and saturation logic, operational conduct, and the fieldwork challenges encountered. A short behavioural survey had been scoped at the outset as a supplementary triangulation device but was discontinued once the depth available through interviews and artefacts made an additional instrument disproportionate to the marginal evidence it could add.

3.5.1 Overview of Data Collection Strategy

Interviews explored lived experiences and participants' interpretations of escalation, challenge, and voice in daily practice. Organisational artefacts offered insights into formal governance structures and recorded communication, allowing comparison between documented processes and frontline accounts. Reflexive memos recorded the researcher's positionality, interpretive growth, and analytic choices, increasing transparency and trustworthiness.

Fieldwork procedures were implemented iteratively, beginning with early interviews and artefact readings that informed initial working definitions of EF, FS, and

VL, and progressing through reflexive thematic development, explanation building, and triangulation (Guba and Lincoln, 1994; Tobin and Begley, 2004; Yin, 2009). RE was reserved as a synthesis construct, inferred from cross-unit configurations rather than refined as an event-level code. Organisational access and case boundary definition followed the negotiation described in Section 3.2.3 and 3.2.4; embedded units were predefined using the three-line model's logic, and case classification attributes (line of responsibility, function, location, job band) were recorded in NVivo to support within-case comparison (Eisenhardt, 1989a; Yin, 2009).

3.5.2 Interview Guide Development

The interview guide was developed using Kallio et al.'s (2016) five-step framework for constructing theory-based semi-structured interview guides, progressing from research question mapping through indicator-linked probe design to pilot testing. To support participant-level clarity, each formal research question in Section 1.3 was translated into a short interview theme (Appendix A.2), which in turn anchored the semi-structured interview guide (Appendix A.1); this translation preserves conceptual fidelity while keeping questions accessible and non-leading in the field. This ensured that every question block was traceable to a specific RQ and heuristic indicator (EF, FS, VL), preventing conceptual drift during data collection. The full guide (v3.5) is provided in Appendix A.1, and the formal RQ-to-interview-theme-to-indicator mapping is provided in Appendix A.2; indicator discrimination rules for boundary-case coding are provided in Appendix G.8. The development rationale applying the Kallio framework is documented in Appendix A.3, and participant information and consent materials are provided in Appendix A.4.

The guide underwent two rounds of piloting with three senior governance practitioners from within the case organisation. Their data were excluded from the main analysis dataset to maintain sample integrity, resulting in eight question revisions and the addition of four clarifying probes. Piloting confirmed that questions about escalation silence required reframing from hypothetical to experience-based phrasing to reduce social desirability bias. Full details of the revisions and their procedural triggers are documented in Appendix A.3.

3.5.3 Semi-Structured Interviews: Conduct and Quality Assurance

Semi-structured interviews served as the main data source, allowing participants to describe how communication, challenge, and escalation functioned on their own terms while exploring tensions between formal policies and informal practices. This method aligned with the study's subjectivist epistemology: knowledge of EF, FS, and VL was co-created through dialogue; RE is inferred later as an emergent condition from cross-unit patterns in these indicators, in times when speaking up felt possible, risky, or futile (Kallio *et al.*, 2016; Adeoye-Olatunde and Olenik, 2021; Ruslin *et al.*, 2022).

Recruitment and consent. Recruitment followed the stratified purposive approach described in Section 3.4.1, using a standardised CRO-approved invitation process. Invitations were issued in phases: first to the Board and SMT, then to Risk/Compliance and business units, to minimise perceived coercion. The invitation letter clarified that the researcher had no employment or advisory relationship with the organisation and that participation would not affect performance evaluations or career progression. Reminder emails were sent after one week if unit coverage was incomplete. All participants received the Participant Information Sheet 48 hours prior to interview and were given the

opportunity to ask questions before consenting (Appendix A.4). Written consent was provided by responding to the initial request for participation; verbal consent was recorded at the start of each interview and covered audio recording, anonymised use of verbatim quotes, and the right to withdraw at any time (Guba and Lincoln, 1994). Board and SMT interviews were conducted in person in private meeting rooms away from usual workspaces to support candour; all other interviews were conducted via Microsoft Teams.

Interview conduct. Twenty-five interviews were conducted between 21 October and 12 November 2025; full participant breakdown is in Tables 3.3 and 3.4 (for the full fieldwork chronology, see Appendix H.4). Interview duration ranged from 45 to 90 minutes (mean: 62 minutes). Questions balanced consistency with flexibility to follow participants' sense-making paths, in line with the subjectivist epistemology. The achieved sample (n=25; see Section 3.4.2) was interviewed using the finalised guide (v3.5), with saturation monitoring detailed in Section 3.4.3 (Hennink and Kaiser, 2022).

Recording and transcription quality assurance. Interviews were recorded via Microsoft Teams or in private rooms with consent. Teams' auto-transcripts were checked against the audio; unclear terms were clarified where necessary, or the participant was asked for clarification. Personally identifiable information was removed through a line-by-line audit, and pseudonymous MS codes (MS01 to MS25) were applied; a separate attribute file links these codes to non-identifying role descriptors. Each participant received a brief summary of key themes from their interview for confirmation before coding began, supporting credibility and co-construction of meaning.

Interview data provided the core empirical foundation for all four heuristic indicators (Section 3.3), surfacing evidence of EF, FS, and VL across embedded units; this evidence base later supports RE inference through cross-unit explanation-building.

3.5.4 Organisational Artefact Analysis

Organisational artefacts such as committee minutes, governance documents (policies, escalation frameworks), organisational charts, and role descriptions were analysed not only for their content but also for framing, omissions, and silences that reveal how communication and challenge are managed within the institution (Yin, 2009; Power and Gendron, 2015). All artefact-to-indicator linkages and triangulation instances are recorded in Appendix D.1.

Artefact corpus. Access to organisational artefacts was negotiated under confidentiality and ethical protocols that excluded certain highly sensitive sources. Artefacts drawn upon include board and committee minutes (paraphrased to avoid re-identification), agenda packs and presentation materials, governance policies and escalation frameworks, organisational charts and role descriptions, and aggregated internal survey summaries. Sources explicitly excluded, on formal legal and compliance grounds and to mitigate re-identification risk, include detailed risk registers, incident and breach reports, email or messaging threads, and individual-level survey data. Artefacts specifically referenced by participants (for example, committee minutes, risk and declined-business registers, survey dashboards, and SharePoint guidance) were examined under existing governance arrangements, with any potentially identifying details retained only in anonymised form.

Analytic approach. Instead of detailed line-by-line coding, structured analytic memos were created for each artefact or group of artefacts, using headings aligned with the codebook for EF, FS, and VL; artefact memos also informed the later RE synthesis through pattern-level observations. These memos documented important patterns, recurring language, and notable omissions, serving as additional data sources for systematic comparison with coded interview material. Analytically, the artefact set and related memos were used to verify formal records against frontline narratives, identify omitted examples, trace routinised agreement or sanitised dissent, and follow pathways by which concerns were documented, escalated, and closed through governance channels. This selective, memo-based triangulation enhanced credibility and reliability by testing key participant claims about EF, FS, and VL against contemporaneous organisational records, while recognising that artefact review did not constitute a comprehensive survey of governance documentation.

Confidentiality safeguards. Raw artefacts such as emails, instant messages, detailed risk registers, and commercially sensitive figures are not included in this thesis. Instead, artefact material is paraphrased, anonymised, and linked to indicators through a documented audit trail. Appendix D provides the triangulation tracking matrices showing how interview evidence, artefact readings, and reflexive observations were compared, including triangulation instances by indicator and source (D.1), summary statistics (D.2), indicator distribution (D.3), and the method-to-indicator support matrix (D.4); Appendix B records the codebook version history (B.1), post-saturation refinement (B.2), and reflexive memo extracts supporting code revisions (B.3); and Appendix C summarises interview-by-interview saturation tracking.

3.5.5 Reflexive Journaling

Reflexive journaling was maintained as an ongoing record of the researcher's impressions, positionality, and analytical choices, making clear how assumptions and interpretations developed during fieldwork and analysis. Reflexivity was central to the subjectivist epistemology: the researcher influenced the identification and interpretation of EF, FS, and VL at event level, and the later cross-unit synthesis through which RE is inferred as an emergent condition, expectations, and emotional responses influenced the identification and interpretation of EF, FS, VL, and RE.

Structured reflexive memos (e.g., Memo_INT_MS16, Memo_INT_MS49, Memo_INT_MS50; sample extracts in Appendix B.3) were prepared for each interview using a consistent template: date/context, researcher reflections, interpretive notes, assumptions/biases, and action points. These memos documented participant tone, hesitancy, perceived power dynamics, and moments that challenged prior assumptions, along with analytic decisions such as codebook refinements and sampling adjustments. For instance, reflexive observations from early interviews led to the addition of child codes distinguishing passive exhaustion (EF "Fatigue") from active structural barriers (EF "burden and inefficiency") and helped clarify the boundary between performative openness (RE "Ritualised Openness") and genuine challenge. Memos were coded alongside interviews and artefacts to triangulate participant accounts with documentary evidence and to identify and reduce potential bias, contributing to both the transparency of the analytical process and the robustness of emerging interpretations.

3.5.6 Integration Across Methods

Data sources were integrated through systematic triangulation to strengthen indicator validity and support explanatory development. In NVivo, each source type was classified, and the shared codebook for EF, FS, and VL was applied across interviews, artefacts, and reflexive memos; RE was developed through cross-unit explanation building (Section 3.6.5 and Appendix G.5), inferring an emergent condition from cross-unit configurations, enabling both convergent and divergent queries at the indicator and unit levels.

Triangulation proceeded by comparing narratives, artefactual traces, and reflexive observations for each indicator, testing where patterns converged or diverged, and using explanation-building to link local mechanisms (such as blocked escalation routes or ignored feedback) to broader manifestations of risk entropy. Table 3.5 summarises indicator support across methods, showing how semi-structured interviews, organisational artefacts, and reflexive journaling contribute to EF, FS, VL, and RE and their respective roles in the study.

Table 3.5 Indicator Support Across Methods

Method	EF	FS	VL	RE	Role in study
Semi-structured interviews	✓	✓	✓	✓	Primary source of lived experience
Organisational artefacts	✓	✓	✓	Δ	Documentary evidence of structures and recorded challenge
Reflexive journaling	✓	✓	✓	✓	Positionality tracking and interpretive evolution

Note: ✓ = primary coding destination; Δ = indirect/partial support (artefacts inform RE synthesis but RE is not coded at event level from artefacts).

Summary statistics on triangulation coverage, indicator distribution across sources, and the method-to-indicator support matrix are provided in Appendix D.2, D.3, and D.4 respectively.

3.6 Data Analysis and Indicator Construction

The analysis sought to translate coded interview and artefact segments into validated heuristic Behavioural Indicators of communication and challenge, grounded in participants lived experience. This was achieved through four successive stages: reflexive thematic coding of transcripts and artefacts into parent themes and child codes; unit-level profiling of how EF, FS, and VL were configured across roles and jurisdictions; indicator specification and triangulated validation through structured analytical memos combining artefacts and reflexive notes; and explanation-building that linked unit configurations to causal mechanisms and inferred RE as an emergent cross-unit condition rather than coding it at the event level. Figure 3.2 depicts this four-stage analytic pipeline, showing how reflexive thematic coding, unit-level indicator profiling, indicator specification and triangulated validation, and explanation-building cumulate through iterative refinement cycles.




Figure 3.2 Iterative Process

The process was iterative, with coding refinements guiding unit profiling, validation findings leading to re-profiling, and explanation-building initiating further validation checks, all documented through structured analytical memos on ERCC minutes, packs and Terms of Reference, reflexive memos, and supervisor debriefs. This pipeline ensures claims about cultural fragility are traceable to coded material and are systematically triangulated across methods.

3.6.1 Reflexive thematic analysis

Braun, Clarke, Hayfield and Terry (2019a) distinguish coding-reliability TA, codebook TA, and reflexive TA, each characterised by different assumptions about coding procedures and the status of themes. This study adopts reflexive TA because it is consistent with the subjectivist epistemology (Section 3.1.1) and because it accommodates the iterative, heuristic development of EF, FS, and VL, and the post-saturation synthesis through which RE is inferred as an emergent cross-unit condition.

In reflexive thematic analysis, the researcher serves as the primary interpretive instrument; dependability depends on transparent coding rules and reflexive accountability. Inter-rater reliability statistics are unsuitable for reflexive TA because they assume that themes exist objectively in the data and can be identified equally by any trained coder (Braun *et al.*, 2019b). Dependability is demonstrated through transparent codebook development documented in Appendix B.1, post-saturation refinements recorded in Appendix B.2, and reflexive memo extracts in Appendix B.3, alongside the saturation tracking matrices in Appendix C.1–C.5.

Reflexive memos documented analytical decisions, researcher positionality, and observed divergences between espoused governance and lived practice, offering an additional source for triangulation and tracking code refinement over time. Coding extended beyond parent-theme saturation to capture cross-unit variation, as documented in Appendix C and 3.4.3.

While this study adopts the interpretive orientation and reflexive analyst positioning of Braun and Clarke's reflexive thematic analysis, it incorporates structured elements, codebook versioning, saturation monitoring, and reference-count tracking, that are more characteristic of codebook-based or abductive thematic approaches (see Appendix B and C). These elements were added to support auditability and unit-level comparability across four embedded governance units, not to shift the analytic paradigm toward positivism. The resulting design is best described as a reflexively oriented, codebook-supported case analysis, as the researcher's interpretive judgement remains primary, and the counting and tracking procedures serve as transparency tools rather than

validity criteria. This hybrid reflects Braun and Clarke's (2019a) acknowledgement that different analytic purposes may warrant structured documentation without abandoning reflexive positioning.

3.6.2 Code-to-theme development

RE was not a direct coding destination but was reserved for later cross-unit synthesis (Section 3.6.5) to ensure consistent handling of boundary cases; the full boundary-case discrimination rules are documented in Appendix G.8. Child codes were added to capture context-specific manifestations such as “voice fatigue” or “governance burden”. Themes were then triangulated with organisational artefacts (minutes, logs, policies) and reflexive memos to confirm, refine, or challenge emerging interpretations. Post-saturation child-code refinements, including the addition of sub-codes distinguishing passive exhaustion (EF Fatigue) from active structural barriers, are documented in Appendix B.2. The full codebook evolution is documented in Appendix B, and NVivo query and node management procedures are detailed in Appendix D.5.

3.6.3 Unit-level indicator profiling

Unit-level profiling followed within-case logic, examining each embedded unit individually before cross-unit comparison. For each unit, interviews and artefacts were synthesised to identify prominent child codes for EF, FS and VL and to map how these clustered into qualitatively distinct configurations, labelled with concise mechanism descriptors such as “symbolic governance” or “control without voice”.

Within-unit variation by role or geography was documented using calibrated claims language, treating very small subgroups as illustrative rather than definitive evidence. Profiles were then triangulated with organisational artefacts to corroborate patterns of filtering, silence or selective acknowledgement, while respecting restrictions on email content and granular registers. This triangulation was operationalised through NVivo-ready analytical memos for ERCC minutes, packs, action logs, CRO reports and Terms of Reference, applying the EF/FS/VL/RE framework and the intensity rubric defined in Sections 3.3 and 3.6.4. The deliverable comprised four-unit profiles showing how formally similar governance structures were experienced differently, providing the evidentiary basis for locating risk entropy as a distributed, context-dependent phenomenon. Indicator reference counts by participant, supporting the intensity assessments reported in the unit profiles, are tabulated in Appendix I.1.

3.6.4 Indicator specification, validation and intensity rubric

Heuristic Behavioural Indicators are defined as patterned configurations of lived experience and interpretive sense-making that signal how risk-related communication and challenges operate in practice. EF, FS, and VL are distinguished by mechanism-level specificity; RE is distinguished as a synthesis-level construct inferred from cross-unit configurations of these indicators, contextual grounding at the unit, role and jurisdictional levels, and diagnostic value for identifying cultural vulnerabilities.

An indicator is considered present within a unit when multiple participants describe the pattern, triangulated artefacts corroborate it, the pattern recurs across sub-contexts, and structured analytical memos and reflexive notes document its persistence. Subgroup claims follow a calibrated hierarchy, and intensity is communicated through qualitative categories

defined by combinations of frequency, corroboration and reflexive observation rather than numerical scores.

Table 3.6 Indicator Intensity Levels and Typical Evidence Basis

Intensity Level	Definition	Typical Evidence Basis
Illustrative	Pattern present but limited to isolated instances; not generalisable across the unit	1-2 independent participant accounts; no artefact corroboration
Common	Pattern recurrent across multiple participants or sources within a unit; not universal	3+ accounts or 1 account with artefact corroboration; cross-role consistency
Ubiquitous	Pattern saturation across the unit; present in all or nearly all accounts and corroborated by artefacts	Convergence across participant roles, seniority levels, and independent artefact sources

Cross-unit intensity patterns fall into three interpretive categories:

- Stratified by role, for example VL MINIMAL-STRATIFIED at Board where control functions are presumed legitimate and operations are treated as conditional.
- Evidence of structural barriers, for example, FS UBIQUITOUS at Board and ILLUSTRATIVE at Operations, indicating that governance structures do not experience the suppression as others do.
- Decoupling signals, for example, RC HIGHEST discourse combined with ICF COMMON/UBIQUITOUS, where reinforcement is espoused but not enacted.

3.6.5 Explanation building and risk entropy synthesis

Unit-level configurations of EF, FS, and VL were linked to potential causal mechanisms and to the broader idea that distributed cultural fragility generates risk entropy. Pattern matching identified recurring indicator clusters within units, compared them with

emerging theoretical propositions, and outlined specific mechanisms, such as informal gatekeeping or ritualised governance, supported by interview, artefactual, and reflexive evidence.

Mechanisms were tested for coherence and alternative explanations, including disconfirming cases, and then combined to illustrate how local practices produced organisation-level manifestations of risk entropy through processes such as information fragmentation, delayed corrective action, normalised ambiguity, and cross-unit inconsistency; the full explanation-building framework and Core-Mechanism Symbiosis procedure are documented in Appendix G.7.

3.6.6 Core-Mechanism Symbiosis as Interpretive Framework

Explanation-building distinguished between contextual mechanisms (such as governance structures, policies, and escalation channels that formally exist in the organisation) and core heuristic Behavioural Indicators (such as friction, suppression, legitimacy constraints, and emergent entropy that staff report experiencing). This two-tier perspective was employed to interpret unit-level configurations, preventing the misinterpretation of the absence of an indicator in a unit, such as Board-level Escalation Friction, as a data gap rather than a real governance-level condition. Detailed patterns of co-occurrence and examples from this framework are presented in the findings chapter. The full two-tier framework structure, including the distinction between contextual mechanism codes and proximate behavioural indicators, is documented in Appendix G.7.

3.6.7 Findings Chapter Organisation Logic

The four analytical stages described in Sections 3.6.1–3.6.5 (reflexive thematic coding, unit-level profiling, indicator specification, and explanation-building) combined to produce findings presented in Chapter 4 using a structured analytical template. The template is the output format of the explanation-building process: it operationalises the requirement for traceable, evidence-grounded interpretive findings (Yin, 2009) by ensuring that every analytical claim is explicitly anchored to its heuristic indicator, stated as a testable heuristic pattern, supported by at least two independent source types, interpreted through visible inferential reasoning, explained mechanistically, and bounded by the conditions under which it holds.

Each analytical Subsection in Chapter 4, Sections 4.3–4.4, uses a seven-element template comprising two reader-navigation elements and five evidential elements. The navigation elements orient the reader before evidence is introduced, combined with the evidential elements constitute the analytical architecture of each claim. Table 3.7 presents the analytical template elements and their functions.

Table 3.7 Analytical Template Elements and Functions.

Element	Type	Function
Purpose	Navigation	Orients the reader to the specific analytical question the Subsection addresses
Indicator link	Navigation	Anchors the Subsection to its specific heuristic indicator (EF, FS, or VL), preventing conflation between indicators with overlapping surface content
Pattern statement	Evidential	States the heuristic claim before evidence is presented, making it testable against the evidence that follows
Evidence (dual-source)	Evidential	Meets the triangulation requirement (Section 3.5.6) by requiring at least two independent source types before a pattern claim is accepted
Analysis	Evidential	Provides the interpretive link from raw evidence to the stated claim, making the inferential move visible and open to challenge

Mechanism explanation	Evidential	Moves from description to explanation by specifying the causal or relational process through which the pattern operates
Boundary conditions	Evidential	Prevents over-generalisation by specifying the contextual limits within which the pattern holds, consistent with the sampling-explicit interpretation principles in Section 4.1

Section 4.6 uses a compressed three-part variant, Pattern, Mechanism, Consequence, rather than the full seven-element template. This is a deliberate design choice: the evidential base has been established through the full-template Subsections in Sections 4.3–4.4 and the unit profiles in Section 4.5, and the synthesis in Section 4.6 draws on that base rather than re-presenting it. The compressed variant is not a reduction of the analytical standard; it reflects the different function of synthetic cross-unit analysis relative to the evidential analysis that precedes it.

3.7 Trustworthiness, Quality Assurance and Ethical Considerations

Rigour, trustworthiness, and ethics were incorporated into the design from the beginning, organised around Lincoln and Guba's criteria of credibility, dependability, confirmability, and transferability. These principles guided the design, conduct, and reporting of the study, ensuring that EF, FS, and VL are based on a transparent, traceable evidence base, with RE inferences grounded in the cross-unit warrant documented in Appendix G.5 and G.5.1.

3.7.1 Rigour Strategy Overview

The quality framework adopts Baskarada's (2014) integration of Yin's (2009) six-stage case-study process, combining construct, internal, and external validity criteria; the full validity specification is tabulated in Appendix E.3. Construct validity is underpinned by triangulation across interviews, artefacts, and reflexive memos; internal validity through

explanation-building and pattern-matching across embedded units; and external validity via analytical generalisation of indicator configurations as transferable diagnostic heuristics.

3.7.2 Credibility

Prolonged engagement through 25 interviews spanning job bands, functions and locations provided repeated opportunities to probe how communication, challenge and escalation were experienced in practice, supporting a nuanced understanding of EF, FS, VL and emergent manifestations of RE.

Method triangulation involved systematically comparing interview narratives with organisational artefacts (committee minutes, aggregate risk dashboards and registers, escalation policies and related governance documents) and reflexive memos to check consistency, surface discrepancies and refine emerging mechanisms. The interview guide was piloted with three senior participants; feedback informed minor refinements to question ordering and wording, and pilot data were excluded from the analytic dataset. Credibility was further reinforced by convergent evidence across methods and sustained immersion, rather than reliance on participant consensus alone. Saturation evidence supporting credibility is documented in Section 3.4.3 and Appendix C.

Member feedback was offered as an optional interpretive check rather than a formal validation procedure. A summary of emerging thematic findings was shared with all participants, with the majority of whom provided responses. Respondents confirmed broad recognition of the patterns described and raised no substantive objections to the two-level interpretation. No changes to the analytic conclusions were made as a result of this process,

though some passages were amended where participants indicated that specific terms required clarification or were interpreted differently than intended. Given the optional nature of participation and the small response base, this process is characterised here as interpretive feedback rather than systematic member validation, as it contributes to credibility through transparency rather than through replication.

3.7.3 Dependability and Audit Trail

The solo-coder design aligns with reflexive thematic analysis, in which the researcher serves as the primary interpretive instrument; dependability in this paradigm is demonstrated through transparency, consistent application of coding rules, and reflexive accountability rather than inter-rater reliability statistics. Post-saturation code refinements are documented in Appendix B.2; the full codebook evolution log constituting the primary audit trail is in Appendix B.1, and triangulation tracking across all data sources is recorded in Appendix D.

3.7.4 Quality Assurance Mechanisms

Analytic quality was further supported by convergent checks that made interpretive decisions explicit and open to challenge. Systematic codebook application was ensured by recording a brief rationale for each code assignment in NVivo annotations, allowing later scrutiny of why particular segments were interpreted as EF, FS, or VL, and how these unit-level patterns contributed to the synthesis of RE as an emergent condition.

Codebook stability was assessed by formally comparing codebook states at interview 18 (V0.5) and interview 25 (V0.6); only two inductive refinements were added after saturation, indicating that later coding primarily elaborated existing concepts rather

than introducing new ones. Deviant case analysis involved active searches for disconfirming evidence, such as cases where high EF co-occurred with strong VL, to refine configuration-based explanations and avoid overly homogeneous narratives.

Reflexive documentation through systematic memos (MemoINTMS01–MemoINTMS50) recorded evolving interpretations, positionality considerations and analytic uncertainties, further strengthening dependability and credibility. Together, these mechanisms create a transparent, challengeable record of analytic choices and outcomes; the full validity framework underpinning these criteria is documented in Appendix E.3.

3.7.5 Confirmability and Reflexivity

Confirmability was maintained by separating raw data from interpretive products and by systematically examining researcher positioning. Each interview was accompanied by a structured reflexive memo recording contextual details, initial interpretations, assumptions, and perceived biases, while a digital reflexive journal chronicled longitudinal changes in understanding of EF, FS, and VL, and the subsequent cross-unit synthesis through which RE is inferred.

Alternative interpretations were explicitly documented and reconsidered when later data indicated tension or contradiction, ensuring that competing explanations were actively evaluated rather than implicitly dismissed. These practices make the researcher's influence transparent and provide a structured way to verify and refine interpretive claims about communication, challenge, and the emergence of RE.

Specific instances where practitioner vocabulary was identified and corrected during coding are documented with full traceability in Appendix E.2; the reflexive memo extracts that triggered those code revisions are reproduced in Appendix B.3, and the full codebook version history is in Appendix B.1.

3.7.6 Transferability

Transferability was supported through thick description and systematic case classification, enabling readers to judge the applicability of findings to other settings. The organisational setting, governance architecture and embedded units are described elsewhere in the chapter, providing contextual information necessary to assess the relevance of EF, FS, and VL configurations beyond this case, and whether analogous cross-unit patterns support similar inferences about RE as an emergent condition.

NVivo attributes captured the three lines model, function, location, job band and governance role for each participant, enabling construction of unit-level indicator profiles (for example, control functions versus frontline; Jersey versus Isle of Man). Vignettes and configuration tables link indicator patterns to local structures and narratives, enabling practitioners and researchers to assess whether similar communication and challenge dynamics may exist in their own institutions, thereby supporting analytical rather than statistical generalisation.

Additionally, the self-selection dynamics described in Section 3.4.4 mean that transferability claims should be understood as bounded by the perspectives of those who consented to participate; the study's ability to capture the most suppressed voices remains a structural constraint of voluntary interview-based designs (Kristensen and Ravn, 2015).

3.7.7 Ethical Considerations and Data Security

Ethical safeguards and data security measures were implemented throughout, consistent with institutional requirements and qualitative research guidance. The study was conducted under prior institutional ethics approval, with permissible data sources, recruitment boundaries and confidentiality safeguards defined in the proposal. Participants received an information sheet outlining the purpose, voluntary participation, withdrawal rights and the intended use of anonymised quotations, and verbal consent was recorded at the start of each interview as the primary consent mechanism, given the organisational sensitivity of the setting.

Transcripts and obtainable artefacts were de-identified, with pseudonymous MS codes (for example, MS01) applied. Data storage, encryption, access control, backup and retention followed organisational information-security procedures, and sensitive material was paraphrased in reporting to avoid verbatim extracts that could risk deductive disclosure.

Artefact access was constrained by organisational confidentiality requirements and ethical protocols: analysis was limited to governance documents, committee minutes, policies, and aggregated survey summaries, explicitly excluding individual-level operational data. This ensured that triangulation respected both participant anonymity and institutional sensitivity while still supporting robust indicator development.

Given the insider–outsider dynamics of research within a single organisation, participants were reminded that their contributions would not affect employment, thereby balancing the protection of participants with the preservation of a robust analytic record.

The materials used to inform participants are reproduced in Appendix A.4 (participant information sheet and consent process). Written confirmation was invited by return email, though not always received. Where this occurred, verbal consent recorded at interview commencement served as the operative consent record, consistent with ethical guidance permitting verbal consent in organisationally sensitive qualitative research settings.

A structural tension exists between the access arrangement and claims of full researcher independence. Invitations were dispatched through a CRO-approved process, meaning that potential participants were aware, or could reasonably infer, that participation had senior-level endorsement. This creates a risk that employees with concerns about escalation failure, the study's core subject, may have been less candid than they would have been under fully independent recruitment. Three mitigations apply. First, the researcher had no employment, advisory, or reporting relationship with the institution (Appendix E.2). Second, interviews were conducted individually, with recordings and transcripts accessible only to the researcher, and with explicit assurance against deductive identification (Section 3.9.4). Third, the pattern of findings, with multiple participants describing dysfunction with specificity and candour, including senior leaders' self-critical reflections on the bonus-pool measurement problem, suggests that, if coercion effects were present, they did not uniformly suppress critical disclosure. The access arrangement is nonetheless acknowledged as a limitation that future independent replications, using unconditional blind recruitment, should seek to address.

3.7.8 Link to Heuristic Indicator Rigour

Collectively, the rigour and ethical practices described above strengthen the credibility and stability of EF, FS, and VL as proximate diagnostic indicators grounded in participants' lived experiences, and RE as an emergent synthesis construct inferred from their cross-unit configurations of communication and challenge. Prolonged, stratified engagement and multi-source triangulation increase confidence that observed configurations reflect durable patterns within the case rather than artefacts of a single method or moment.

The documented audit trail, reflexive journal, and explicit search for disconfirming evidence support the claim that indicator profiles and risk-entropy narratives are analytically defensible, transparent, and open to scrutiny.

3.7.9 Methodological Limitations

Every methodological design entail trade-offs, and transparency about limitations strengthens rather than undermines the study's contribution.

Single-case design limits statistical generalisability; the findings describe how EF, FS, VL, and RE manifest within one Channel Islands deposit-taker, and their applicability to other jurisdictions, institution types, or regulatory regimes is a matter of analytical transferability rather than replication (Yin, 2009; Baskarada, 2014). Thick description and unit-level configuration tables in Chapter 4 enable readers to assess transferability on their own terms. The indicators, EF, FS, VL, and the inferred RE condition, are interpretive constructs developed through the researcher's engagement with participant accounts; their boundaries reflect analytical judgements documented in the codebook version history

(Appendix B) and the reflexive journal (Appendix E.2), enabling scrutiny and replication by future researchers working in comparable governance contexts. Interview data are retrospective; artefact triangulation, comparing participant accounts against committee minutes and governance documents collected during the same fieldwork window, provided contemporaneous evidence for cross-checking narrative claims. Non-participation may introduce systematic bias, as discussed in Section 3.4.4, with the most acute experiences of suppression or delegitimised voice potentially underrepresented in the achieved sample. Finally, the exclusion of certain artefact types, such as emails, messaging threads, individual-level survey data, and detailed risk registers are constrained documentary evidence to the "committee view" of governance, a boundary imposed by organisational confidentiality conditions and acknowledged as a limit on the depth of artefact-based triangulation.

3.8 Timeline and Resource Plan

A twelve-week schedule structured the study from participant recruitment in Week 1 through validation and write-up by Week 12, with a three-week interview phase (21 October–12 November 2025) within a broader six-month project at its centre (Yin, 2009; Braun *et al.*, 2019a). The complete project chronology, covering all phases from access negotiation through write-up, is documented in Appendix H.4 and reflects the iterative, reflexive approach of the study.

3.8.1 Fieldwork Timeline

Interviews ran from 21 October to 12 November 2025 (three weeks). Coding commenced on 24 November 2025, following transcript quality assurance; parent-theme saturation was confirmed at interview 18 (see Section 3.4.3). Post-fieldwork code

refinement (1–18 December 2025) produced child-level elaborations without altering the parent-theme structure (see Appendix B for codebook evolution).

3.8.2 Detailed Timeline

Main-study interviews were conducted between 21 October and 12 November 2025 (three weeks), following three pilot interviews in the week commencing 13 October 2025. Coding commenced on 24 November 2025; post-fieldwork code refinement (1–18 December 2025) produced child-level elaborations without altering the parent-theme structure (see Appendix B). See Appendix H.1 for the complete project timeline.

3.8.3 Dependencies and Critical Path

Main fieldwork could not commence until the three-participant pilot had verified the interview guide (Kallio et al., 2016). Concurrent coding required sufficient transcripts by Week 6 to enable saturation monitoring alongside data collection (Braun et al., 2019a). Tracking parent-theme saturation during transcription confirmed the rationale for census completion detailed in Section 3.4.3; dependencies and the critical path for these activities are documented in Appendix H.2.

3.8.4 Feasibility in Relation to Sampling and Methods

The 12-week schedule was realistic given the study’s embedded design, rapid on-boarding of the researcher and negotiated access. The target of 15–30 interviews was comfortably achieved (25 completed), reflecting strong organisational support via sponsor liaison and senior-level sponsorship. The embedded case design, spanning multiple business units and locations, allowed interviews to be scheduled efficiently without external recruitment delays (Eisenhardt, 1989b).

NVivo coding, reflexive memoing, and unit-level explanation-building ran concurrently across the nine-week analysis period, with artefact review embedded in each interview cycle rather than deferred to a separate phase (Yin, 2009; Braun et al., 2019a). Organisational artefacts were accessed alongside interviews, enabling real-time triangulation during coding and avoiding bottlenecks (Power and Gendron, 2015).

3.9 Data Management and Security

All data handling complied with institutional ethics approval and organisational confidentiality requirements. This Section outlines the procedures used to protect participants and maintain an auditable link from raw materials (audio files, transcripts, memos, coding tables, NVivo projects and artefacts) to the heuristic Behavioural Indicators EF, FS, and VL, and to the synthesis of RE as an emergent condition from their cross-unit configuration. These procedures support the trustworthiness criteria of credibility, dependability and confirmability within an interpretivist design in which participants' accounts of communication and challenge are treated as context-bound and sensitive (Guba and Lincoln, 1994; Tobin and Begley, 2004).

3.9.1 Data Management Principles

Data management was designed to safeguard confidentiality, integrity, and secure storage throughout the full data lifecycle, collection, transfer, storage, analysis, sharing, and destruction, ensuring that no individual participant or organisational unit was directly identifiable in transcripts, coding outputs, or reported findings.

These principles serve the study's subjectivist epistemology: participants were more likely to describe escalation barriers, discouraging feedback responses and legitimacy judgements candidly when confident their contributions would be protected and could not be traced back to them or their units (Sandberg, 2005). An auditable link from raw data through to indicator construction strengthens dependability by demonstrating how analytical claims are traceable to source material (Tobin and Begley, 2004).

3.9.2 Storage, Access and Encryption

Primary data were stored on institutionally approved, encrypted drives and in secure cloud storage. Access was restricted to the researcher, and viewing was limited to de-identified materials. Data handling followed the organisation's standard data and information security procedures, including the automatic deletion of Teams recordings after 90 days and their transfer to approved encrypted drives, ensuring the confidentiality, integrity, and secure retention of all research materials. These arrangements create the infrastructural conditions for a dependable audit trail and protect the evidentiary base for EF, FS, VL and RE (Guba and Lincoln, 1994; Tobin and Begley, 2004) with the full storage, access, and encryption arrangements detailed in Appendix H.3.

3.9.3 File Naming, Backup and Retention

Standardised ID-based filenames were used for all audio files, transcripts, memos and NVivo exports, enabling consistent linkage between raw segments, coded nodes and unit-level indicator profiles. Daily encrypted backups were run during fieldwork and analysis, with periodic verification checks. Raw audio recordings were securely deleted once transcription accuracy had been confirmed; de-identified versions were safely stored in line with dependability and confirmability requirements (Guba and Lincoln, 1994).

3.9.4 Anonymisation and De-identification

Participants were assigned pseudonymous codes (MS01–MS25). The key linking identities to codes was stored separately, encrypted, and excluded from NVivo projects and shared folders. Transcript cleaning replaced or generalised personal names, specific roles, locations and organisational units while preserving analytic meaning, so that patterns in EF, FS, VL and RE could be interpreted without risking deductive disclosure (Guba and Lincoln, 1994).

Artefact excerpts were paraphrased or aggregated so that neither individual employees nor sensitive transactions could be re-identified. All quotations used in the thesis were checked against anonymisation criteria, with additional masking applied in small-N contexts or highly specific episodes. This approach balances participant protection with sufficient contextual detail to support analytic credibility and transferability (Tobin and Begley, 2004), with the full anonymisation and de-identification procedures documented in Appendix

.

3.9.5 Integration with Trustworthiness and Ethics

These data-management procedures underpin dependability by maintaining a coherent audit trail from raw recordings through cleaned transcripts, coding decisions and indicator construction. They support confirmability by separating raw data from reflexive memos, analytic notes and decision logs (Guba and Lincoln, 1994; Tobin and Begley, 2004).

By minimising identifiability and ensuring storage and transfer comply with organisational standards, the protocol protects participants while enabling rigorous, traceable analysis of communication and challenge dynamics. This alignment of data management with ethical commitments strengthens the trustworthiness of EF, FS, VL and RE as interpretive indicators within the approved research design.

3.10 Chapter Summary

Four iteratively refined heuristic behavioural indicators reveal how communication and challenge function within the governance structure of a formally regulated deposit-taker. The embedded single-case design, implemented through refining heuristic indicators (EF, FS, VL, RE), captures patterns of escalation friction and voice suppression that uniform governance arrangements conceal. Interview data (n=25), read against committee artefacts and reflexive memos (Appendices B–D), underpin each unit-level indicator profile with overlapping evidence.

The resulting behavioural framework provides governance practitioners with a practical diagnostic tool to identify early signs of risk entropy before it leads to governance failure. Chapter 4 illustrates where EF and FS cluster within the four embedded units and what these concentrations reveal about the organisation's governance culture.

CHAPTER IV:

RESULTS

4.1 Overview of the Two-Level Pattern

The analysis reported in this chapter addresses the four research questions established in Section 1.3 through the heuristic behavioural indicators operationalised in Chapter III, using the indicator mappings detailed in Appendix A.2 to structure each subsection. Three data sources, covering 25 interviews, organisational artefacts, and reflexive memos, are consistent with the same two-level pattern in how communication and challenge function within the case institution. Escalation Friction (EF), Feedback Suppression (FS), and Voice Legitimacy (VL) are used as proximate behavioural indicators, applying the operational definitions set out in Section 3.3 (indicator specifications) and the analytical pipeline described in Section 3.6. RE is inferred as an emergent synthesis construct from their cross-unit configuration rather than being applied as a fourth parallel indicator.

The two-level architecture adopted here is not an arbitrary editorial choice. It is grounded in a theoretical distinction, developed in Chapter II (Sections 2.2.5 and 2.6.3), between two qualitatively different conditions under which risk voice operates. At the local level, Edmondson's (1999) psychological safety model and Morrison's (2014) dual safety-efficacy calculus together explain why individuals judge speaking up as safe and likely to be effective, with conditions that are principally relational and dyadic. At the systemic level, Hardy and Maguire's (2016) account of discursive power boundaries, and Power's (2009) rituals of verification, explain why concerns that feel safe to raise locally nonetheless fail to translate into corrective governance outcomes, conditions that are principally structural and architectural. The organisational-level theoretical warrant for

treating these as two analytically separate levels comes from Rousseau et al. (1998), who distinguish relational trust, arising from repeated interaction, reciprocity, and interpersonal attachment, from institution-based trust, which operates through roles, procedural safeguards, and reputational mechanisms rather than personal familiarity. Mayer et al. (1995) supply the dyadic mechanism: willingness to be vulnerable is a function of perceived ability, benevolence, and integrity in a specific trustee, and they separate this willingness from the act of risk-taking itself, which is also conditioned by the perceived risk of the situation. An employee may therefore regard their direct manager as trustworthy in Mayer et al.'s terms while judging formal escalation too structurally risky to attempt, the safety-efficacy split Morrison (2014) operationalises. The data are consistent with precisely this two-level structure, not a gradient of voice safety, but a boundary condition at which the conditions governing voice change qualitatively.

The main finding is that contextual mechanisms of Openness to Alternate Views (OAV), Mechanisms for Raising Concerns (MRC), and Reinforcement of Challenge (RC), the three behavioural expressions of the FSB's (2014) effective communication and challenge pillar, are widely discussed, and in the contexts where staff genuinely feel safe, translate into Voice Legitimacy (VL). OAV, MRC, and RC are used as the structural organising device for both the local-authenticity analysis (Section 4.3) and the systemic-failure analysis (Section 4.4), allowing the same mechanisms that enable local voice to be examined as the sites of its failure at formal organisational levels, a structural parallel that gives the chapter its two-level architecture. Across the coded data, codes for these mechanisms account for 1,127 references (full source coverage) and 31 positive co-occurrence passages where openness, mechanisms, and reinforcement lead to perceived voice legitimacy. However, when voice attempts extend beyond these local boundaries into

formal organisational settings, the same mechanisms fail to prevent widespread Escalation Friction, contested Voice Legitimacy, and the conditions from which Risk Entropy is inferred at the systemic level. Across the coded data, EF, FS, and VL together account for 301 proximate behavioural references (EF 95, FS 43, VL 163), with 7 proximate negative co-occurrences (EF–FS) and 9 cross-indicator configurations, comprising 5 RC–RE co-occurrence passages and 4 OAV–RE co-occurrence passages, from which risk entropy is inferred as an emergent condition. The FS distribution is structurally concentrated rather than evenly spread, with two participants, one in Oversight (MS01, 15 references) and one in 2nd Line (MS41, 9 references), together accounting for 56% of the 43 FS references, consistent with the finding in Section 4.4 that suppression concentrates where formal feedback loops exist rather than where voice is pre-filtered before reaching them. Separately, 160 analytical segments are marked as RE synthesis passages, these are not a fourth set of coded events but rather the passages where the emergent entropy condition, inferred from the cross-unit EF–FS–VL configuration, is most analytically legible (Section 3.6.5).

For ease of reference throughout this chapter, the four indicators are summarised in Table 4.1 below.

Table 4.1 Indicator reference table

Indicator	Full name	One-sentence operational definition
EF	Escalation Friction	Procedural, relational, and cognitive barriers that impede the upward movement of risk concerns through formal channels (Section 3.3.2).
FS	Feedback Suppression	Patterns of ignored input, absent follow-up, tokenistic acknowledgement, and voice fatigue that reduce the likelihood of future challenge (Section 3.3.3).
VL	Voice Legitimacy	The perceived appropriateness and credibility of speaking up, who is seen as authorised to raise concerns, and in what contexts (Section 3.3.4).

RE	Risk Entropy (emergent)	An emergent synthesis condition, inferred from cross-unit EF-FS-VL configurations, indicating fragmented, delayed, or inconsistent risk information flow (Section 3.3.5).
-----------	-------------------------	---

The two levels map to the research questions as follows: the local authenticity configuration (Section 4.3) principally addresses RQ1; the systemic failure analysis (Section 4.4) principally addresses RQ2; Sections 4.5 and 4.6 together address RQ3 and RQ4 by operationalising the indicators as diagnostic tools and showing how unit configurations aggregate into a systemic pattern.




Figure 4.1: The Emergent RE Inference
How EF, FS, and VL Configurations Give Rise to Risk Entropy. The local expression of this architecture, where EF, FS, and VL co-occur favourably, is shown in Figure 4.2.

The data indicate that, for every proximate behavioural reference, mechanism discourse accounts for 3.7 references, indicating that mechanisms are discussed far more extensively than the friction and suppression they are designed to prevent. This disparity is referred to throughout this chapter as the discourse-practice gap. Throughout this chapter, the code references for contextual mechanisms are collectively referred to as Mechanism Discourse (MD), comprising six codes: BRE, ICF, MRC, OAV, RC, and RCF. Within this set, OAV, MRC, and RC function as the primary analytical organising device throughout this chapter, they are used to structure Sections 4.3 and 4.4 as standalone subsections, as they most directly capture the behavioural mechanisms participants experience in practice: openness to dissent, channels for raising concerns, and reinforcement of challenge. ICF and RCF, grounded in the FSB's (2014) "Stature of control functions" sub-heading within Section 3.3, surface within the unit-level profiles where control-function authority and perception are analytically relevant, principally the Control Without Voice configuration (Section 4.5.2) and the Control Function Silencing mechanism (Section 4.6.3), but are not used as independent subsection-level organising devices. BRE reflects FSB risk-reward balance themes distributed across the framework and surfaces in commercial-tension contexts rather than as a recurring communication and challenge theme. Proximate behavioural indicator references are referred to collectively as the Proximate Behavioural count (PB), comprising EF, FS, and VL. It is operationalised as the ratio of MD (six codes, totalling 1,127 references) to PB (three indicators, totalling 301 references), yielding a ratio of 3.71. Risk Entropy (RE) synthesis segments (n=160) mark passages where the emergent entropy condition is most analytically legible from the EF–FS–VL configuration; they are shown separately rather than added to PB, as RE is an emergent construct rather than a fourth direct code. The full derivation and unit-level breakdown are provided in Section 4.6.5, and the two-level configuration is visualised in


Figure 4.1. The pattern was not pre-existing; it surfaced through the four-stage analysis process outlined in Chapter III (Section 3.6).

Throughout this chapter, intensity labels (Ubiquitous, Common, Illustrative, and Minimal) describe the prevalence of a given indicator among participants sampled within each embedded unit, not across the organisation as a whole. Cross-unit comparisons are accordingly configurational, illuminating how indicators manifest differently across governance units rather than ranking units on a single scale. Full intensity definitions and the rubric governing their application are provided in Section 3.6.4 and Appendix G.6.

The seven figures in this chapter form the connected analytical sequence mapped in Figure 4.0, progressing from local authenticity (Figure 4.2) through systemic failure (Figures 4.3–4.4), unit-level configuration (Figure 4.5), system-level synthesis (Figure 4.6), and evidential corroboration (Figure 4.7).

Chapter IV Analytical Logic Map

Indicator reading logic and analytical progression across Figures 4.1-4.7



Revised visual logic map based on supplied Chapter IV sequence.

Figure 4.0: Chapter IV Analytical Figure Sequence

Figure 4.1 provides the indicator inference logic governing all sections. Figures 4.2–4.7 present the analytical progression from local authenticity through systemic failure, unit-

level configuration, synthesis, and evidential corroboration. Each figure is referenced in the section where its evidence is developed.

4.2 Data Sources and Triangulation Approach

The following findings integrate three data sources as specified in Section 3.6:

- 25 semi-structured interviews conducted 21 October–12 November 2025 (Section 3.6.2), coded for proximate behavioural indicators (EF, FS, VL); RE was inferred as an emergent condition from cross-unit configurations of EF, FS, and VL rather than coded directly at segment level and contextual mechanisms (BRE, ICF, MRC, OAV, RC, RCF) using the hierarchical codebook (Appendix B, Section 3.6).
- Organisational artefacts, including ERCC committee minutes and packs (October, December 2025; May 2025), action logs, forum Terms of Reference and Peakon engagement survey data (December 2023–April 2024), analysed through structured analytic memos (Section 3.6.3).
- Reflexive memos documenting analytical insights and researcher positionality (Section 3.6.4; exemplars in Appendix B.3).

Evidence is presented using convergence/divergence/novel insight tags as specified in the Triangulation Tracking Matrix (Appendix D); a summary of triangulation outcomes across all 18 artefact instances is provided in Table 4.11 (Section 4.7.3). This three-source design, interviews, artefacts, and reflexive memos, reflects the triangulation logic established in Chapter III (Section 3.5.6), selected as appropriate for an interpretivist embedded case study in which no single source can fully represent how communication and challenge operate across governance levels (Yin, 2009). Here, 'convergent' indicates interview and artefact evidence support the same pattern, 'Divergent' highlights policy–

practice gaps where they conflict, and 'novel insight' marks patterns visible only in artefacts. Citations employ MS-xx for Main Study interview excerpts, ART-xx for artefact evidence, and Memo_INT_MS-xx for reflexive memo observations. For each major finding, evidence from at least two sources is included.

Table 4.2 summarises the analytical function of each Section and the findings it contributes to the two-level pattern.

Table 4.2 Chapter organisation

Section	Finding
Section 4.3	Identifies where and how contextual mechanisms OAV, MRC, RC appear to translate into legitimate voice, identifying trusted relationships as the boundary condition under which this authenticity holds, and establishing the local configuration against which systemic failure is measured in Section 4.5.
Section 4.4	Examines how the same mechanisms that enable local voice appear to become the sites of failure at formal organisational levels. Examines how the three EF barrier types (procedural, relational, cognitive) produce escalation friction, how reinforcement remains scrutiny-dependent rather than internalised, and how openness is subject to procedural conversion and confidence-induced closure, producing the discourse-practice gap quantified in Section 4.6.5. OAV, MRC, and RC are used as the structuring device for Sections 4.4.1–4.4.3.
Section 4.5	Identifies unit-specific entropy signatures through indicator profiling, producing four provisional labels: Symbolic Governance (Oversight:

	governance appears complete but is structurally disconnected from the friction it mandates); Control Without Voice (2nd Line: formal independence is procedurally absorbed so that authority does not reliably translate into influence); Silenced Executors (Operations: those closest to risk have the least governance access, resolving dysfunction through workarounds); and Boundary Negotiators (Business, provisional: voice legitimacy depends on commercial framing rather than operational substance).
Section 4.6	Synthesises how the four-unit configurations, when read together, produce three aggregate decoupling patterns, governance blindness, control function silencing, and operational normalisation, each examined through a compressed three-part structure (Pattern, Mechanism, Consequence) rather than the full five-component template used in Sections 4.3 and 4.4. This abbreviated form is used because the evidence base has already been established in the unit profiles; Section 4.6 draws on it synthetically. The Section closes with the 3.7:1 discourse-practice gap as the quantitative expression of aggregate decoupling.
Section 4.7	Triangulates interview findings against three independent artefact streams, ERCC committee minutes and packs, Peakon engagement survey data, and forum Terms of Reference, across 18 artefact instances. Reveals convergence, divergence, and Novel Insights that independently corroborate the two-level pattern: governance documents confirm persistent reds and deferral patterns; survey data confirm the freedom-of-opinion / escalation-friction split; and novel artefact signals surface patterns not accessible through interview data alone.

Section 4.8	Summarises the two-level pattern, four-unit entropy signatures, four decoupling mechanisms, and the 3.7:1 discourse-practice ratio as the integrated analytical outputs of the chapter and previews the theoretical and practical implications developed in Chapter 5.
--------------------	--

Each finding Subsection in this chapter is structured by a named analytical template, visible to the reader as a set of formatted italic headings. These headings, Purpose, Indicator link, Pattern statement, Evidence, Analysis, Mechanism explanation, and Boundary conditions, are not arbitrary formatting choices; they correspond directly to the analytical pipeline described in Chapter 3, Section 3.6.7, and each carries a specific evidential function. They are also intended as navigational anchors: a reader wishing to examine only the evidential basis of a claim can go directly to the Evidence element of any Subsection; a reader wishing to assess the scope or limits of a claim can go directly to Boundary conditions. Section 4.7 uses a compressed three-part version, Pattern, Mechanism, Consequence, because its synthesis function differs from the evidential function of Sections 4.3 and 4.4; this is explained in the template description immediately below.

Each analytical Subsection in Sections 4.3–4.4 follows the seven-element analytical template established in Chapter 3, Section 3.6.7. The first two elements are reader-navigation devices: a Purpose statement that names the analytical question the Subsection addresses, and an Indicator link that identifies the specific heuristic indicator (EF, FS, or VL) under examination, ensuring that evidence relevant to, for example, procedural barriers (EF) is not read as evidence about feedback reception (FS), even where participant language overlaps. The remaining five elements constitute the evidential

architecture: a Pattern statement (the heuristic claim, stated before evidence so it is testable), Evidence from at least two independent sources (meeting the triangulation requirement established in Section 3.5.6 and operationalised in the analytical template in Section 3.6.7), Analysis (the interpretive link from evidence to claim, making the inferential move visible and open to challenge), Mechanism explanation (the causal or relational process through which the pattern operates), and Boundary conditions (the contextual limits within which the pattern holds, consistent with the sampling-explicit interpretation principles in Section 4.1). Section 4.7 uses a compressed three-part variant, Pattern, Mechanism, Consequence, because the evidence base has been established in full across the unit profiles in Section 4.5, and the synthesis draws on that base rather than representing it; this is a deliberate design choice, not a departure from the analytical standard.

4.3 Local Authenticity: Where Mechanisms Work

Purpose: Within trusted relationships, OAV, MRC, and RC, the three behavioural expressions of the FSB's (2014) effective communication and challenge pillar (Table 3.1), reliably translate into Voice Legitimacy (VL), the local boundary within which RQ1 (how alternative views are encouraged and expressed) and RQ2 (how formal mechanisms operate in practice) are first answered.

This purpose applies across all three Subsections: 4.3.1 examines openness as the invitation condition, 4.3.2 examines formal mechanisms as the channel condition, and 4.3.3 examines reinforcement as the sustaining condition, together constituting the local configuration against which systemic failure is measured in Section 4.4.

Throughout this section, "trusted relationships" refers to dyadic or small-team relationships characterised by relational proximity (direct managers and familiar colleagues), interpersonal reciprocity (mutual attentive listening), operational scope (concerns framed around process rather than strategic governance), and low reputational risk (challenge framed as constructive rather than disloyal). These are the boundary conditions under which OAV, MRC, and RC reliably translate into Voice Legitimacy, the local configuration visualised in Figure 4.2.

4.3.1 Openness Enables Voice

Indicator link: Voice Legitimacy (VL)—perceived appropriateness and credibility of speaking up (Section 3.3.4).

Purpose: To examine whether genuine openness in trusted relationships reliably produces voice legitimacy, and under what relational conditions that link holds.

Pattern statement: Seventeen passages, read together, reveal a recurring co-occurrence pattern in which OAV and VL appear together, offering a heuristic indication that where staff experience genuine openness within trusted relationships, speaking up is regarded as legitimate and worthwhile rather than risky or futile.

Evidence (dual-source requirement met):

MS15: "We want your opinion... I've never seen a situation where they said we're not interested."

Convergent, Aligns with Peakon Freedom of Opinion score 8.9 (ART-Peakon/Apr2024).

Analysis: The participant presents openness as unconditional ("never seen") and invitation-based ("we want"), rather than tolerance-based ("you're allowed"). This distinction indicates psychological safety, where voice is encouraged, not merely permitted.

Evidence (dual-source requirement met):

MS02: "Everyone is enabled and empowered to give their opinion... there's no limitations on that."

Convergent, Aligns with Ideas Bank mechanism described in artefacts (ART-ERCC/Oct2025).

Analysis: The dual emphasis on "enabled" (structural permission) and "empowered" (felt authority) indicates OAV translates into VL through both formal mechanisms and cultural norms.

Mechanism explanation: Where OAV and VL co-occur positively, openness functions through three interconnected elements: (1) direct invitation ("we want your opinion"), (2) attentive listening (ideas acknowledged rather than dismissed), and (3) low reputational risk (challenging is framed as constructive rather than disloyal). These elements point towards a localised feedback loop where initial contribution is encouraged through positive response, fostering further involvement.

Boundary conditions: All 17 OAV→VL positive passages share two contextual constraints: (1) voice occurs within existing relationships, direct managers and familiar colleagues, not formal escalation pathways, and (2) voice concerns operational process improvements, not strategic governance issues. When participants attempt to challenge beyond these boundaries, OAV→VL breaks down. The boundary condition is the relationship itself, with openness operating through familiarity rather than through formal design, a distinction that sets the stage for the scaling failure examined in Section 4.4.

4.3.2 Mechanisms Work in Trusted Contexts

Indicator link: Voice Legitimacy (VL) and Escalation Friction (EF), perceived channel usability within relational contexts (Sections 3.3.4 and 3.3.2).

Purpose: To examine whether formal mechanisms for raising concerns function as usable channels independently of relational trust, or only when embedded within existing trusting relationships.

Pattern statement: Eight passages, read together, reveal a recurring relational pattern in which MRC and VL co-occur, offering a heuristic indication that formal channels are perceived as usable only when embedded within existing, trusting relationships.

Evidence (dual-source requirement met):

MS01: "Very good relationship... I feel I can talk to them."

Novel insight, MEM-INT01: Researcher noted participant's relaxed posture when describing direct manager relationship, shifting to visible tension when discussing cross-

unit escalation, adding researcher-positionality corroboration to the relational dependency reading not fully explicit in the transcript itself.

MS05: "We have access to the SFS board at any time."

Novel insight, MS05 raises the board-access claim from a non-Jersey jurisdiction, indicating that perceived channel usability is not Jersey-specific.

Analysis: These accounts collectively suggest that access to channels is perceived as safe only when part of a trusted dyadic relationship. The MEM-INT01 observation corroborates that this relational dependency is reflected not only in what participants say but in how they say it, tone and posture shift when the referent moves from local to cross-unit escalation.

Mechanism explanation: Formal channels (Ideas Bank, direct manager access, board access) are perceived as usable when trust is already established. The mechanism is not the channel itself but the relational embedding that makes the channel feel secure.

Boundary conditions: When the same mechanisms are deployed in formal escalation, for example, bypassing a manager to contact Oversight, they create relational friction rather than promoting voice (Section 4.4). This relational dependency means that formal channel design, on its own, cannot substitute for the trust that makes those channels feel safe to use, a structural finding that Section 4.4.1 examines through the lens of procedural burden.

4.3.3 Reinforcement Enables Voice

Indicator link: Voice Legitimacy (VL), interpersonal reinforcement as perceived appropriateness of speaking up (Section 3.3.4).

Purpose: To examine whether interpersonal reinforcement sustains voice legitimacy over time, and whether that effect is institutionally or relationally anchored.

Pattern statement: Six passages, read together, reveal a recurring interpersonal pattern in which RC and VL co-occur, offering a heuristic indication that, in this material, reinforcement supports legitimate voice primarily through reciprocity and mutual regard rather than through any formal reward system.

Evidence (dual-source requirement met):

MS03: "Both sides listen to each other."

MS01: "Very good relationship... I feel I can talk to them."

Novel insight, MEM-INT01: The posture shift recorded here, relaxed when describing the direct manager relationship, visibly tense when discussing cross-unit escalation, reflects a coding decision: what initially appeared as individual affect was reinterpreted, through positionality review, as learned behavioural adaptation, consistent with reinforcement being relationally bounded rather than institutionally generalised.

Analysis: These passages illustrate reinforcement as a relational, two-way process, mutual listening, felt permission to speak, rather than as a formal reward system. The MEM-INT01 posture observation, read as behavioural adaptation rather than personal affect, emphasises that this reinforcement is local and fragile.

Mechanism explanation: Reinforcement appears to operate on an interpersonal level, through mutual listening and reciprocal respect, rather than via formal reward systems. Participants describe reinforcement as relational ("my manager encourages me") rather

than institutional ("the organisation rewards challenge"). This interpersonal emphasis is independently visible in the Peakon engagement artefact, where a high Freedom of Opinion score indicates staff feel their views are valued, suggesting comfort in expressing opinions without reliance on formal reward structures (VL, Illustrative).

Boundary conditions: Reinforcement is specific to relationships, not organisation-wide. Staff experience reinforcement through particular managers but not within formal governance forums. Reinforcement is interpersonally strong but institutionally absent, participants can identify specific relationships that promote challenge, but cannot point to a structural mechanism that extends it beyond those relationships. This interpersonal ceiling is the condition under which systemic erosion becomes possible, examined in Section 4.5.

4.3.4 Configuration Summary: Trusted Relationships Create Legitimate Voice

Across 31 positive co-occurrence passages (OAV→VL: 17; MRC→VL: 8; RC→VL: 6), a consistent pattern emerges: openness, mechanisms, and reinforcement translate into voice legitimacy when embedded in trusted relationships. The configuration works through relational proximity (voice occurs within dyads of familiarity, not distant governance bodies) and interpersonal reciprocity ('both sides listen'), bounded by scope (concerns are operationally framed) and low reputational risk (challenge is framed as problem-solving, not dissent).

OAV provides the invitation, MRC supplies a usable channel, and RC maintains reciprocity over time, but only within trusted relationships, what this analysis provisionally

labels authentic behavioural effectiveness: mechanisms appear to operate as intended within these relational boundaries, though this is an interpretive characterisation rather than a measured state. Participants describe genuine felt permission ("I feel I can talk"), rather than reluctant compliance ("I'm technically allowed"). Local reinforcement is effective but limited, which identifies the conditions under which systemic erosion becomes possible, examined in Section 4.5.

However, this authenticity is locally situated and contingent. The same mechanisms (MRC, OAV, RC) that enable voice in trusted dyads fail when deployed beyond those boundaries, in formal escalation, cross-unit challenge, or governance-level dissent. The boundary condition, relationships, not formal design, is the analytical pivot: the same mechanisms that make voice authentic within dyads make its absence structurally invisible at scale. This three-part configuration is visualised in Figure 4.2. Section 4.5 examines what happens when those same mechanisms are deployed beyond relational boundaries and why they fail.

Local Authenticity Configuration

How contextual mechanisms converge through trusted relationships to generate Voice Legitimacy

Configuration logic:
OAV + MRC + RC → [Trust] → VL




Figure 4.2: Local Authenticity Configuration

How OAV, MRC, and RC Translate into Voice Legitimacy within Trusted Relationships. The breakdown of these conditions under formal escalation is shown in Figures 4.3 and 4.4.

4.4 Systemic Failure: Where Mechanisms Cannot Scale

The following three sub-sections examine where the theoretical conditions for authentic voice, Edmondson's (1999) relational safety and Morrison's (2014) efficacy calculus, are met in practice, establishing the local configuration against which Section 4.5 measures failure.

The three barrier types examined here, procedural, relational, and cognitive, correspond analytically to the distinction Argyris developed between single-loop and double-loop learning (Edmondson and Moingeon, 1998; Tosey, Visser and Saunders,

2012): procedural and relational barriers prevent concerns from being corrected at the surface level; cognitive barriers prevent the governing variables of the escalation architecture itself from being examined, the condition to which Chapter V (Section 5.9) returns.

Purpose: When voice attempts to move beyond trusted relationships into formal organisational settings, the mechanisms that enabled it locally to become the sites of its failure.

4.4.1 Friction Causes Suppression

Indicator link: Escalation Friction (EF) and Feedback Suppression (FS), procedural, relational, and cognitive barriers impeding upward movement of risk concerns, and the suppression of future challenge that follows (Sections 3.3.2 and 3.3.3).

Purpose: To examine how three types of escalation barrier, procedural, relational, and cognitive, operate sequentially to convert initial willingness to raise concerns into withdrawal or silence.

Pattern statement: Seven passages, read together, reveal a recurring sequential pattern in which procedural, relational and cognitive barriers operate in sequence, such that EF and FS co-occur sequentially in a pattern consistent with cumulative barrier logic, offering a heuristic indication of how cumulative barriers transform initial willingness to raise concerns into withdrawal or silence.

Evidence (dual-source requirement met, procedural barrier, Section 3.3.2: formal obstacles, ambiguous pathways, onerous documentation):

MS46: "[Issue] dropped beforehand because of time, resources."

Convergent, ART-ERCC/Dec2025: RSCA uplift deferred to 2026 due to R23 resourcing.

MS01: "Agreed it with my head of compliance and financial crime as well before it goes out."

Convergent, ART-ERCC/Oct2025: Action 192 (Target Country) deferred 2023–2025, repeated rescoping, Group approval dependencies.

Analysis: Escalation requires sequential approval across multiple functions, creating structured opportunities for delay and dropout. The artefact evidence corroborates that deferral is not an isolated perception, the same pattern appears in formal governance records spanning years.

Evidence (dual-source requirement met, relational barrier, Section 3.3.2: informal gatekeeping, managers filtering or reframing issues):

MS20: "I highlighted something from our accounts payable process... but then Fin Controller and Head of Fin Planning do not necessarily have the full context of it."

MS41: "not documented or minuted."

Novel insight, MEM-INT03: Participant's voice dropped noticeably when describing upward escalation, contrasting with animated tone when discussing peer-level challenge. Positionality review led to reinterpretation of this shift as learned behavioural adaptation, consistent with relational EF rather than individual reluctance.

Analysis: As concerns move up the hierarchy, operational context is abstracted into summary form, losing urgency and specificity. Relational hesitancy reflects learned experience that escalation transforms the concern itself, not fear alone.

Evidence (dual-source requirement met, cognitive barrier, Section 3.3.2: normalisation of deviance, problem minimisation):

MS46: "Escalation pathway exists but [I] dropped [it] beforehand."

Divergent, ART-Peakon/Apr2024: "Would have to find whistleblowing policy", mechanism exists but accessibility and awareness are weak.

Analysis: Informal norms discourage use of formal channels ("don't bother them, work it out yourself"), creating cognitive friction. Staff have internalised which concerns are worth raising, a learned heuristic that pre-filters issues before they reach any formal channel.

Mechanism explanation: The three barriers operate sequentially, not in isolation. An initial willingness to raise a concern is first filtered by cognitive doubt, then by an internal assessment of procedural costs, and finally by relational hesitancy; at each stage, the concern either diminishes or is discarded entirely. The same staff members (most clearly in Operations and 1st Line Business accounts, 4.6.3–4.6.4) who describe genuine openness within their teams are also the ones who pre-filter concerns before they reach any formal channel, suggesting, in this material, that the friction is structural rather than personal. This sequential dynamic is visualised in Figure 4.3.

Boundary conditions: The sequential barrier model applies specifically to formal escalation contexts, the three barrier types operate in combination when concerns move beyond trusted dyadic relationships into cross-unit or governance-level channels. Within trusted dyads, the same staff members navigate all three barriers successfully, which is the local authenticity condition described in Section 4.3.

The Sequential Barrier Model: From Willingness to Silence

How sequential barrier types convert willingness toward voice into withdrawal




Figure 4.3: The Sequential Barrier Model

Procedural, relational, and cognitive barriers operating in sequence during formal escalation. The dashed feedback arrow indicates that withdrawal validates learned futility, reducing the probability of future attempts. The six compounding barriers are detailed in Figure 4.4; evidential basis developed across 4.4.1–4.4.3. Source: 25 semi-structured interviews; artefact corroboration per 4.7.

4.4.2 Reinforcement Does Not Prevent Entropy

Indicator link: Reinforcement of Challenge (RC) and Risk Entropy (RE, emergent), the highest-frequency contextual mechanism code co-existing with the cross-unit EF–FS–VL configuration from which entropy is inferred (Sections 3.3.2 and 3.3.5).

Purpose: To examine how Reinforcement of Challenge, the most extensively documented mechanism in the corpus, co-exists with the conditions from which Risk Entropy is inferred, and what this paradox reveals about the scrutiny-dependency of challenge norms.

Pattern statement: Five passages, read together, reveal a recurring configurational pattern in which Reinforcement of Challenge (RC), the highest-frequency code in the codebook (408 references), co-exists with the RC–EF–FS–VL configuration associated with the emergent entropy condition. This offers a heuristic account of how governance investment in challenge norms does not reliably prevent behavioural withdrawal under scrutiny removal.

Evidence (dual-source requirement met, complete absence of reinforcement):

MS02: "Never been really in existence of something impacting marketing as an operational risk."

Convergent, ART-ERCC/Oct2025: 1st Line Business Operational Risk Forum: "No items to escalate", despite 2nd Line dashboard showing TPRM red, liquidity concentration red.

Analysis: Where reinforcement is entirely absent, operational risks are neither tracked, documented, nor escalated. The artefact convergence corroborates that this is a recognised governance gap, not merely a participant perception.

Evidence (dual-source requirement met, scrutiny-dependent reinforcement):

One participant, in a position with direct oversight of risk governance, articulates the temporal character of this pattern without prompting:

MS03: "Challenge is enhanced when we're under that level of scrutiny... and practice's a little bit quieter when we're not."

Divergent, from universal RC discourse: despite 408 references to reinforcement, actual practice is scrutiny-dependent, not continuous.

Analysis: Reinforcement functions as an external response to regulatory review or audit pressure rather than as an internalised cultural norm. When scrutiny diminishes, challenge becomes quieter, suggesting reinforcement is not self-sustaining.

Evidence (dual-source requirement met, compartmentalised reinforcement):

MS10: "They may well have to afford some post-mortems, but I don't see them... not something that they published."

Convergent, ART-ERCC/May2025: 1st Line Operational Risk Forum, no formal minutes, only action tracker. Learning not documented or shared.

Analysis: Reinforcement takes place within silos, the risk team conducts post-mortems but does not share insights across units. Lessons stay compartmentalised, restricting organisational learning.

Mechanism explanation: The governance paradox in this material is structural rather than individual: Reinforcement of Challenge carries 408 event-level references, the highest count in the codebook, yet the RE synthesis segments (n=160), marking where the emergent entropy condition is most analytically legible from the EF–FS–VL configuration, cluster most densely precisely where RC discourse is most extensive.

The gap between what RC references record (extensive discussion of challenge, training, and encouragement) and what RE segments infer (conditions consistent with entropy) is not a contradiction but an explanation: reinforcement is enacted conditionally, intensifying under regulatory scrutiny and receding when scrutiny is removed.

RC therefore functions as a signal of governance intent rather than evidence of behavioural embedding, which is why its frequency is the highest in the corpus while its behavioural correlates remain scrutiny-dependent.

Boundary conditions: Scrutiny-dependent reinforcement is one of the primary drivers of the discourse-practice gap developed in Section 4.6.5, mechanism discourse is most extensive precisely where behavioural enactment is most conditional. This temporal pattern addresses RQ4 directly: rather than a sustained cultural embedding, reinforcement functions reactively, intensifying under regulatory or audit pressure then receding once that pressure lifts.

4.4.3 Openness Does Not Prevent Entropy

Indicator link: Openness to Alternate Views (OAV) and Escalation Friction (EF, procedural, relational, and cognitive sub-types), the co-existence of professed openness with the structural constraints that prevent it from generating genuine escalation conditions (Sections 3.3.1 and 3.3.2).

Purpose: To examine how Openness to Alternate Views co-exists with the OAV–EF–FS–VL entropy configuration, and how procedural conversion, performative substitution, and confidence-induced closure prevent discourse-level openness from generating systemic-level voice effectiveness.

Pattern statement: Four passages, read together, reveal a recurring configurational pattern in which OAV co-exists with the OAV–EF–FS–VL entropy configuration, such that even where openness is professed, structural gates, ritualised practices, and confidence-based

closure can still allow entropy to build, offering a heuristic indication that discourse-level openness does not guarantee systemic-level voice effectiveness.

Evidence (dual-source requirement met, procedural EF, Section 3.3.2: formal obstacles converting invitation into compliance):

MS07: "The Ideas Bank is good but you have to submit three ideas a quarter, it becomes about hitting the quota, not the ideas."

Convergent, ART-ERCC/Oct2025: Ideas Bank referenced in governance artefacts as active channel; no record of submissions escalating to corrective action.

Analysis: Where openness is enacted through formal mechanisms, the procedural requirements of those mechanisms transform invitation into compliance. The quota obligation reframes contributing from a voluntary expression of genuine concern into a performance target. The artefact evidence is convergent with a reading in which Ideas Bank submissions are tracked as participation metrics rather than reviewed as risk signals, the channel is documented as functioning while the content it captures is not escalating into governance.

Evidence (dual-source requirement met, relational EF, Section 3.3.2: performative enactment substituting ritual for genuine inquiry):

MS14: "There are consultation sessions, but the outcome is usually decided before you walk in. You can say whatever you like."

Novel insight, MEM-INT04: Researcher noted that several participants distinguished between being asked and being heard, using language suggesting ritual rather than genuine inquiry. Positionality review confirmed this distinction is felt rather than inferred, participants modulate their language when describing formal consultation relative to informal conversation, suggesting they have learned to read the difference.

Analysis: Where openness is enacted relationally, it can present as consultation while functioning as confirmation. Participants distinguish between being invited to speak and experiencing that speech as consequential.

Evidence (dual-source requirement met, cognitive EF, Section 3.3.2: confidence closure removing the perceived need for escalation):

MS01: "I'm fairly confident that if there was something that the organisation needed to know, it would get to me."

Divergent, Operations and 2nd Line data in Sections 4.6.1 and 4.6.2–4.6.3 describe concerns being pre-filtered or dropped before reaching governance level.

Analysis: Where openness is genuinely experienced at senior levels, confidence in the system's functioning can itself close down inquiry. A senior participant's confidence that concerns would surface is consistent with having received filtered, polished escalations, which is precisely the governance blindness pattern developed in Section 4.7.1. Confidence is not misplaced at the individual level; it is structurally produced.

Analysis (consolidated): Across the three evidence blocks, openness in formal settings operates through a single logic, the form of invitation is preserved while its substance is withdrawn, whether by procedural gating, performative substitution, or confidence-based closure.

Mechanism explanation: The three constraints operate differently from the sequential EF barriers in Section 4.5.1: procedural gating (procedural EF) converts invitation into compliance before a concern is formed; performative enactment (relational EF) substitutes

the ritual of consultation for its substance; and confidence closure (cognitive EF) removes the perceived need for escalation at the point where it would otherwise be generated. Together, all three locally functioning mechanisms, OAV, RC, and MRC, are subject to systemic failure, establishing the analytical ground for the unit-level entropy profiling in Section 4.6.

Boundary conditions: Openness is authentic within trusted relationships (Section 4.4), but at the systemic level it is subject to procedural conversion, performative substitution, and confidence-induced closure. Where Section 4.5.1 shows concerns being filtered after they are raised, and Section 4.5.2 shows reinforcement failing to sustain challenge once raised, this Section is consistent with a pattern, visible across Oversight, 2nd Line, and 1st Line Business accounts (MS01, MS14, MS07), in which openness fails to generate the conditions under which raising concerns feels worth attempting in the first place.


4.4.4 Configuration Summary: Structural Barriers Prevent Scaling

Three types of escalation barrier, procedural (procedural gates, resource constraints), relational (accountability fragmentation, manager gatekeeping), and cognitive (confidence closure, scrutiny dependency), operate in sequence and in combination, their six discrete forms together converting local authenticity into systemic failure, as visualised in Figure 4.4.

Scaling Failure Model — Six Compounding Barriers

How three barrier types, expressed as six discrete forms, convert local authenticity into systemic voice failure

Scaling logic:
Local Context → Escalation → Barriers → Systemic Failure



Critical finding:

Structural design problem, not isolated individual failures. For the temporal sequence by which these barriers convert willingness into withdrawal, see Figure 4.3.

Source: 25 semi-structured interviews; artefact corroboration per Section 4.7.

Figure 4.4: Scaling Failure Model, Six Compounding Barriers
Converting Local Authenticity into Systemic Voice Failure. The unit-level expression of these barriers is shown in Figure 4.5.

Across 301 proximate behavioural references (EF: 95; FS: 43; VL: 163), with 7 negative co-occurrences between EF and FS at the proximate level and 9 cross-indicator configurations (RC→RE: 5; OAV→RE: 4) from which entropic conditions are inferred; RE is inferred as an emergent synthesis condition from these configurations, with 160 synthesis segments marking passages where that inference is most analytically legible. Mechanisms that work locally cannot scale to formal organisational contexts: procedural gates and resource constraints together produce the most frequent escalation dropout (most visible in Operations and 2nd Line accounts, 4.6.2–4.6.3); accountability fragmentation allows issues to persist through repeated rescoping (observed across 2nd Line and Oversight data, 4.6.1–4.6.2); and confidence closure and scrutiny dependency then sustain

the cycle (concentrated in Oversight accounts, 4.6.1), removing both the perceived need and the external pressure for genuine challenge.

Within this single-case design, this configuration provides the closest available answer to the 'why' aspect of RQ2 by identifying procedural and relational mechanisms that normalise delay and silence. These barriers appear to function as structural features rather than individual faults. The data do not point toward deliberate concealment: participants across all units describe genuine commitment to openness and reinforcement, and organisational artefacts record sustained investment in these mechanisms, the gap between discourse and practice appears structural in character rather than individually motivated. Section 4.6 examines how this aggregate failure manifests differently across the four embedded units, producing distinct entropy signatures from a common structural logic.

4.5 Unit-Level Entropy Signatures

Unit-level profiling is necessary because aggregate assessments, including Peakon data and organisation-wide mechanism discourse counts, cannot detect the subcultural variation the embedded case design was selected to reveal (Ashby, Palermo and Power, 2012). The four units map directly onto the three lines model (Section 3.2.5), Oversight as mandate layer, 2nd Line as independent authority, 1st Line Operations as executor, 1st Line Business as commercial interface and treating them as a single culture would obscure the structural inversions that Sections 4.3 and 4.4 have established as the analytical question. The four-unit labels are inductively derived from indicator profiles rather than imported from existing typologies, which operate at the organisational level and cannot differentiate governance function within a three lines model.

Purpose: Each embedded unit produces a distinct entropy signature, a characteristic configuration of proximate indicator intensities (EF, FS, VL) from which Risk Entropy (RE) is inferred as an emergent condition, revealing how and where cultural vigilance erodes within that unit's structural position. Comparing these signatures across the four units diagnoses the distribution of entropy rather than assuming it is evenly spread. Each unit profile follows a three-part structure: an indicator profile table summarising EF, FS, VL, and RE intensities; evidence blocks linking coded data to artefact evidence; and a configuration summary assigning the provisional entropy label. Table 4.3 below provides an orientation to the four signatures; each is examined in detail in Sections 4.5.1–4.5.4.

Table 4.3 Overview of Four Unit Entropy Signatures
Each label is a provisional characterisation derived from the unit's indicator profile; the derivational logic is made explicit in 4.6.5.

Entropy Signature	Embedded Unit	EF Type Dominant	Key Mechanism Failure	Confidence
Symbolic Governance	Oversight / Board / SMT	Procedural (architectural)	Non-reception of filtered escalations	HIGH
Control Without Voice	2nd Line Risk and Compliance	Procedural + Relational	Procedural absorption of formal authority	HIGH
Silenced Executors	1st Line Operations	All three (procedural, relational, cognitive)	Learned futility + workaround normalisation	HIGH
Boundary Negotiators	1st Line Business	Cognitive (revenue gating)	Self-regulation; commercial framing filters	PROVISIONAL

4.5.1 Oversight, Towards Symbolic Governance

The Oversight data point toward what this analysis provisionally labels Symbolic Governance, a configuration in which formal governance appears complete, yet the data suggest it remains largely disconnected from the friction and silencing described in

practice. Census of accessible governance participants (n = 5), covering Board/Oversight and SMT members with direct responsibility for risk oversight across Jersey and the Isle of Man. All consenting senior leaders in scope were interviewed, providing 100% accessible coverage for this unit. Saturation was reached at n = 3, with no new governance escalation patterns emerging in later interviews; subsequent cases reinforced rather than extended the pattern. Confidence is therefore high that these findings reflect governance-level realities rather than sampling artefacts. The indicator intensity profile for this unit, including RE inference, is set out in Table 4.4.

Table 4.4 Oversight Indicator profile

Indicator	Intensity	Key finding
EF	Common (4/5)	Escalation processes shaped by consensus gates, multi-function approval requirements and Group dependencies that they mandate but experience only as procedural thoroughness (procedural delay, informal gatekeeping). Governance appears to generate friction architecturally, while receiving only polished, filtered escalations, obscuring the friction those structures may produce.
FS	Ubiquitous (5/5)	Feedback is procedurally absorbed rather than substantively actioned; committee procedures ensure items are tabled, minuted and formally closed, but formal recording substitutes for corrective response (tokenistic recognition, absent follow-up). Governance interprets procedural completeness as robust challenge, obscuring the same "noting only" pattern documented in 2nd Line experience.
VL	Common (4/5)	Control-function voices are treated as inherently authoritative; operational voices are heard selectively and via managers. Institutionalised stratification of voice: "second line by right, operations by permission."
RE	Ubiquitous	The EF/FS/VL configuration for this unit is consistent with an RE inference of Ubiquitous; variations in practice across units are acknowledged but framed as inevitable contextual difference, consistent with the normalisation of ambiguity and inconsistency (Section 3.3.5).

RE (inferred from EF/FS/VL configuration; not directly coded)

Evidence:

MS21 : "couldn't understand why [their] CPD was a higher requirement."

MS23 : "it doesn't always mean that everybody can give that same response"

DIVERGENT, Operations staff describe this same inconsistency as systemic dysfunction requiring workarounds (Section 4.5.3).

The Oversight data are consistent with a pattern in which policy-practice gaps are framed as acceptable variation, the normalisation of ambiguity observable indicator (Section 3.3.5), rather than read as signals of the EF and FS patterns evident at lower levels.

MS12 : "If a member of the compliance, risk or financial crime team have significant concerns about that piece of business, they will speak up... They may hold slightly more of a higher weighting than the front office because obviously they've got that complete independent view of it and unbiased view. Well, the frontline might be pushed towards seeing the pound signs."

MS06 : "I generally don't escalate them to [CEO] until I've had discussions with CRO to consider whether it's appropriate"

CONTRAST, 2nd Line describes having to argue for credibility (Section 4.4.2).

All sampled Oversight members regard control-function voices as inherently legitimate ("that's their job") and operational voices as conditionally legitimate, reflecting institutionalised stratification, a visible indicator of VL (such as attribution of authority and differential reception; Section 3.3.4).

Configuration Summary: Towards Symbolic Governance

The Oversight data are consistent with Symbolic Governance, formal governance appears to function as intended (papers received, minutes approved, policies mandated) yet the data suggest it remains largely disconnected from the friction and silencing described in practice. The census sample is consistent with this being a role-specific setup: senior leaders consistently state they receive polished escalations and aggregated dashboards rather than raw operational struggles.

What is structurally distinctive about Oversight is not active concealment, but what the data are consistent with labelling architectural non-reception: friction reaches governance only in filtered form, because the procedures governance mandates are also what filter it. The data suggest no deliberate concealment: governance receives filtered evidence precisely because the architecture it mandates produces the filtering, a mechanism the analysis develops further in Section 4.6.2. *Convergent, ERCC committee minutes and packs record persistent dashboard reds and multi-year action deferrals without corrective closure (ART-ERCC/Oct2025, Dec2025, May2025), independently consistent with a governance architecture that acknowledges risk while remaining structurally removed from the friction that produces it.* Geographic location subtly influences how Oversight perceives these dynamics: Jersey-based leaders tend to view cross-jurisdictional inconsistency as normal contextual variation, while staff on the Isle of Man more often see such variation as a source of friction that necessitates workarounds. What distinguishes this configuration is not that governance is uninformed, but that the architecture it mandates is precisely what prevents it from receiving unfiltered evidence, friction is not hidden from governance, it is filtered by it.

4.5.2 Risk and Compliance (2nd Line), Towards Control Without Voice

The 2nd Line data point toward what this analysis provisionally labels Control Without Voice, a configuration in which control functions hold formal independence on paper, but the data suggest their challenge is procedurally absorbed or deferred, so that structural authority may not reliably translate into behavioural influence. Target-achieved sample (n=5), including Risk and Compliance officers across Jersey and the Isle of Man. Saturation was reached at n=5, with consistent challenge-credibility patterns observed across all participants. Confidence remains high that these findings reflect structural features of the 2nd Line role rather than individual differences. The indicator intensity profile for this unit, including RE inference, is summarised in Table 4.5.

Table 4.5 2nd Line Unit, EF, FS, VL Intensity and RE Inference

Indicator	Intensity	Key finding
EF	Ubiquitous (5/5)	Escalation occurs but requires navigating Group dependencies and multi-function consensus. Challenge reaches governance but often after dilution.
FS	Illustrative (2/5)	Despite formal independence, 2nd Line experiences challenge being absorbed, deferred, or reframed. "Noted" but not actioned is a recurring pattern.
VL	Ubiquitous (5/5)	Authority granted structurally ("that's their job") but contested in practice. Must argue for credibility rather than assume it.
RE	Common	Information provided to Oversight/Board but concerns whether all concerns surface. Aware of gaps but unable to close them systemically.

RE intensity is assessed at unit-configuration level, not individual participant level; the UBIQUITOUS rating reflects the strength of the cross-unit EF–FS–VL inference pattern, not participant presence.

Evidence:

MS46 : "are actually been dropped beforehand because of time resources, the person just doesn't have capacity or they don't see the benefits in it versus BAU."

MS16 : "sometimes when you just know that you're going to get the same response... you kind of know that it's not going to go anywhere... sometimes actually it's not worth it."

CONVERGENT, ART-ERCC/Dec2025: Multiple "noting only" items on ERCC agenda across three consecutive packs.

These records show actions being documented but then deferred or closed with minimal corrective change, so formal documentation replaces substantive response, illustrating FS (absent follow-up, Section 3.3.3); the delayed corrective action that follows is among the EF/FS/VL configurations from which risk entropy is inferred as an emergent condition (Section 3.6.5).

Configuration summary: towards control without voice

The 2nd Line data suggest a pattern consistent with Control Without Voice, structurally supported by formal independence charters and escalation rights, yet appearing experientially limited by procedural absorption, deferral, and credibility contests. The structural inversion here, highest formal authority, most attenuated felt influence, is the 2nd Line's distinctive diagnostic signal: independence and absorption operate simultaneously in the same role.

The distinctive tension here is the inversion of mandate and experience: structural authority is highest in 2nd Line, yet felt influence is most attenuated, formal independence and procedural absorption pull in opposite directions. These patterns are more prominently expressed in the Isle of Man, where second-line roles are concentrated. IoM participants report more frequent encounters with escalation bottlenecks and deferred actions, while

Jersey participants more often mention proximity to senior decision-makers as a partial safeguard.

4.5.3 1st Line Operations and Support, Towards Silenced Executors

The Operations data point toward what this analysis provisionally labels Silenced Executors, a configuration in which those closest to operational risk appear to develop workarounds but, the data suggest, seldom escalate concerns into governance forums. An extended sample ($n = 12$) was deliberately increased beyond the initial saturation at $n = 7$ to capture thematic diversity across IT, Operations, and Change Management. Saturation was confirmed at $n = 8$, with subsequent interviews ($n = 9\text{--}12$) reinforcing rather than expanding core patterns. One participant within this unit (MS10) held an internal system administration role with limited cross-unit escalation responsibility; this account is treated as a sampling boundary consistent with bounded operational scope rather than a counter-instance to the unit pattern and does not affect the indicator intensity designations in Table 4.6. Confidence is high that these findings reflect operational realities rather than sampling artefacts.

Table 4.6 1st Line Operations and Support Indicator profile

Indicator	Intensity	Key finding
EF	Common (11/12)	All three barrier types present: procedural (time/resources), relational (manager filtering), cognitive (learned helplessness). Escalation drops before reaching formal channels.
FS	Illustrative (4/12)	Concerns raised locally but not escalated further. Feedback loops stop at immediate manager. Voice fatigue evident—"What's the point?"
VL	Ubiquitous (12/12)	Operational voice seen as situational input, not strategic. Legitimacy bounded by role—"not my place to question" is recurring phrase.
RE	Ubiquitous	Highest entropy concentration. Cross-unit inconsistency, workarounds normalised, information fragmentation across silos.

RE (inferred from EF/FS/VL configuration; not directly coded)

MS46 (S2, Operations, IoM): "I know the issue, I've flagged it to my manager, but taking it further? That would need time I don't have and approvals I can't navigate."

Convergent, ART-ERCC/Dec2025: RSCA uplift deferred due to resourcing.

EF seems to operate across all three barrier types simultaneously, procedural (time/resources), relational (manager as gatekeeper), cognitive (learned futility). The artefact convergence indicates that governance is aware of the gap but lacks the resources to address it, supporting the operational view of structural constraints rather than individual reluctance.

MS21 : " "We have had one case last year when we were launching a new product and it was assumed because the Jersey regulator almost just ticked a box and said yes, that's fine, that the Isle of Man regulator would be the same, and it wasn't."

MS23 : "they were all thinking it was plain sailing... but it doesn't always mean that everybody can give that same response."

DIVERGENT, Oversight/Board describes this as "contextual adaptation" (MS01), not dysfunction.

The RE inference signal is strongest in Operations because staff experience cross-unit inconsistency as daily friction, requiring workarounds, while governance perceives it as acceptable variation, the same reality yielding opposite interpretations based on structural position. The strength of the RE inference signal aligns with the Risk Entropy indicators specified in Methods 3.3.5: normalisation of ambiguity, cross-unit inconsistency, and disconnected escalation.

Configuration summary: towards silenced executors

The Operations data are consistent with Silenced Executors, those closest to operational risk and process dysfunction appear, from the data, to be structurally isolated from governance forums where these insights could inform policy. This unit records the highest EF and VL reference counts and the strongest RE inference signal among all units, along with illustrative FS, creating the widest overall indicator spread, although FS intensity is lower than in the 2nd Line and Oversight.

The diagnostic irony in Operations is quantitative as well as qualitative: the unit with the highest EF reference count and the strongest RE inference signal is also the unit with the least governance access, knowledge most relevant to governance is least accessible to it, not through concealment but through accumulated learned futility. The conditions from which entropy is inferred are particularly concentrated in the Isle of Man operations sample, where participants report both heavy reliance on informal workarounds and limited insight into governance forums. Jersey operations accounts, by contrast, refer to slightly more direct routes to senior managers but show similar experiences of escalation fatigue.

4.5.4 1st Line Business, Towards Boundary Negotiators

The Business unit data offer an indicative pattern, provisionally labelled Boundary Negotiators, in which the data suggest business units operate at the interface between customers and organisational risk appetite, using informal influence networks and revenue-framed arguments so that only commercially framed concerns gain voice legitimacy and visibility in governance. This label carries lower confidence than the three saturated

configurations and is treated accordingly. Extracted sample (n = 3), representing roles in relationship management and business development across Jersey; indicator intensities, including RE inference, are set out in Table 4.7.

Table 4.7 1st Line Business Indicator profile

Indicator	Intensity	Key finding
EF	Common (2/3)	Escalation occurs informally through relationship networks. Formal escalation rare - business units navigate through influence rather than channels.
FS	Minimal (1/3)	Suppression less through blocking and more through self-regulation - business units pre-filter what reaches governance to maintain credibility.
VL	Ubiquitous (3/3)	Business voice is revenue-legitimated. When tied to customer/revenue impact, voice is heard; operational concerns without revenue framing are marginalised.
RE	Ubiquitous	Lower inferred entropy than Operations; direct governance access partially mediates the disconnect between revenue priorities and operational risk.

RE (inferred from EF/FS/VL configuration; not directly coded)

Evidence:

MS16 : "we're most affected because we hear first-hand from the client... we see the impact of a certain process just not being fit for purpose"

MS12 : "Our job is to try and make money for the business, but also protect the business as well"

CONVERGENT - ART-ERCC/Oct2025: 1st Line Business Operational Risk Forum reports "no items to escalate" despite 2nd Line dashboard showing reds.

VL in this unit appears to rely on revenue. Concerns raised in business terms seem more credible, while those in operational terms are dismissed. The artefact convergence illustrates this filtering, with business forums reporting no escalations even when risk dashboards display red items, because those risks are not framed as revenue concerns.

Configuration summary: towards boundary negotiators (provisional)

The data from business units are consistent with what this analysis provisionally labels Boundary Negotiators, roles operating at the interface between customer demands and organisational risk appetite, with voice legitimacy appearing to depend on revenue framing. This unit tends to manage escalation through influence networks rather than formal channels, experiencing less friction than Operations but more strategic filtering, only concerns framed in business impact terms seem to gain traction.


The Business configuration points toward a pattern provisionally described here as revenue-gated voice, elaborated as a component of the Control Function Silencing mechanism in Section 4.6.3, where its lower-confidence status relative to the three saturated configurations is explicitly maintained. Commercially framed concerns gain traction while operational risks without a revenue dimension go unescalated; formal channels are bypassed in favour of influence networks; and the recurring 'no items to escalate' pattern alongside dashboard reds suggests a filtering mechanism that may prevent certain risk categories from surfacing. Despite its provisional status, this configuration provides early evidence relevant to RQ1 (voice legitimacy seems contingent on revenue rather than role) and RQ2 (formal escalation mechanisms seem to be bypassed in favour of influence networks). If the study were expanded, prioritising business-unit sampling would be necessary to achieve full saturation and increase confidence from provisional to high. The provisional status of this configuration reflects genuine evidential limits, three participants cannot saturate a commercial role, but the revenue-gated voice pattern, visible across all three accounts, is sufficiently consistent to carry forward as an indicative contribution rather than a confirmed finding.

4.5.5 Configuration Synthesis: Four Signatures, One Structural Pattern

The indicator profiling process conducted in Sections 4.5.1–4.5.4 has produced four distinct entropy signatures, Symbolic Governance, Control Without Voice, Silenced Executors, and Boundary Negotiators (provisional), each reflecting a different structural position in relation to the two-level pattern identified in Sections 4.2 and 4.3. Read together, they reveal a shared logic: in every unit, formal mechanism discourse is extensive, yet behavioural evidence of those mechanisms functioning is consistently attenuated. Oversight mandates the architecture but does not encounter the friction it creates; 2nd Line raises concerns that are procedurally absorbed rather than actioned; Operations develops workarounds that prevent entropy from surfacing; and Business units gate voice through revenue framing. The inversions visible in Table 4.3, high EF with low FS in Operations, high FS with moderate RE in 2nd Line, are the diagnostic signals that distinguish how each unit's entropy is produced, not merely that entropy exists. What no individual unit can see is the pattern they collectively produce: each failure is locally coherent, locally explained, and structurally invisible from within. The four signatures and their inversion patterns are synthesised in Figure 4.5.

Unit Entropy Signature Matrix

Indicator intensity profiles across four organisational units, revealing structurally significant inversions



Source: 25 semi-structured interviews.

Figure 4.5: Unit Entropy Signature Matrix

Indicator intensity profiles across the four embedded units reveal diagnostic inversions among EF, FS, VL, and RE. Each unit's profile draws on the barrier types identified in Figures 4.3-4.4. Source: 25 semi-structured interviews; unit profiles in Sections 4.5.1–4.5.4. The aggregate system-level pattern produced by these signatures is shown in Figure 4.6.

4.6 Organisational Decoupling: How Unit Configurations Aggregate

Having established the cross-level pattern in Sections 4.3 and 4.4, this section profiles how EF, FS, and VL manifest differently within each embedded unit, with the subcultural variation that Power's (2009) rituals of verification and the embedded case design were selected precisely to reveal.

Purpose: Four structurally distinct units converge on the same overall condition: formal mechanism discourse is extensive, while behavioural enactment remains consistently

reduced. The question is not whether decoupling exists, but how it is maintained across three separate mechanisms.

Three analytically distinct decoupling mechanisms, governance blindness (Oversight, 4.5.1), control function silencing (2nd Line, 4.5.2), and operational normalisation (Operations, 4.5.3), each sustain the conditions for the others, and no unit perceives the system-level pattern it collectively produces. The Business unit's indicative revenue-gated voice pattern is incorporated into Section 4.5.4 at lower confidence. The three mechanisms address distinct dimensions: voice failure (RQ1–RQ2), diminished control authority (RQ3), and gradual erosion of organisational vigilance (RQ4). Before examining the three mechanisms individually, Table 4.8 and Table 4.9 provide complementary diagnostic perspectives on the same coded material, by seniority band and by geographic jurisdiction respectively. Each mechanism is examined through a compressed three-part structure, Pattern, Mechanism, and Consequence, rather than the full five-component template applied in Sections 4.3 and 4.4. This is intentional: the evidence underpinning each pattern has already been presented in the unit profiles (Section 4.5) and systemic failure analysis (Section 4.4); the purpose here is synthesis, not re-presentation.

4.6.1 Diagnostic Cuts: Hierarchical and Geographic Perspectives

Before examining the four decoupling mechanisms (Sections 4.6.2–4.6.4), Table 4.8 provides a complementary analytical perspective. While Section 4.6.5 profiles the data by embedded unit, Table 4.10 reanalyses the same coded material by seniority band, offering a second diagnostic view of the same underlying patterns rather than a separate quantitative test. The hierarchical gradient is clear: EF, VL, and RE steadily increase from Oversight/Board/SMT (Common) to Operational staff (Common – though 11/12), corroborating that the two-level pattern remains consistent across two independent

analytical dimensions, unit and band. This supports the governance blindness diagnosis (Section 4.5.1): seniority levels responsible for designing escalation architecture experience the least friction, while those operating within it encounter the most. FS suggests a distinct pattern, Common at the Oversight/Board/SMT level, decreasing to Illustrative at Operational levels, aligning with the finding that suppression concentrates where formal feedback loops exist (governance, 2nd Line), rather than where voice is pre-filtered before reaching those loops (Operations, Section 4.5.3). This inversion indicates that suppression is not merely a top-down force but a structural by-product of proximity to formal governance processes.

Table 4.8 Indicator Intensity by Job Band (Hierarchical Cut)

Job Band	EF	FS	VL	RE
Oversight/Board/SMT (S5/SL)	Common (5/9)	Common (6/9)	Common (5/9)	Common
Senior Mgmt (S4)	Ubiquitous (4/4)	Common (2/4)	Ubiquitous (4/4)	Ubiquitous
Mid-Mgmt (S3)	Ubiquitous (5/6)	Illustrative (1/6)	Ubiquitous (6/6)	Ubiquitous
Operational (S2/S1)	Ubiquitous (9/9)	Illustrative (3/9)	Ubiquitous (9/9)	Ubiquitous

RE (inferred from EF/FS/VL configuration; not directly coded)

A parallel geographic cut across the same coded corpus is summarised in Table 4.9. The data indicate that the decoupling pattern is not evenly distributed across jurisdictions, with the Isle of Man consistently recording higher indicator intensity than Jersey across EF, VL, and RE. This jurisdictional gradient is consistent with a pattern in which structural distance from governance compounds the entropy that unit proximity to risk would otherwise surface. One IoM participant (MS20) articulates the mechanism directly: *"I feel like if I was in Jersey, it would have been a conversation... that same issue took weeks going back and forth on Teams with people who didn't have the full context."* The account

is analytically significant because it isolates the geographic variable without attributing the friction to individual reluctance: the same person, the same concern, the same willingness to escalate, different outcomes as a function of physical proximity to decision-makers.

Table 4.9 Geographic Indicator Profile

Indicator	Jersey (n = 14)	Isle of Man (n = 11)
EF	13/14 (93%); 3.6 refs/pp. Widespread but experienced as procedural thoroughness.	10/11 (91%); 4.0 refs/pp. All three barrier types active; escalation drops before formal channels.
FS	8/14 (57%); 1.9 refs/pp. Procedural absorption concentrated in 2nd Line at HQ.	4/11 (36%); 1.5 refs/pp. Feedback loops stop at immediate manager; voice fatigue evident.
VL	14/14 (100%); 7.1 refs/pp. Relationally embedded in governance networks.	10/11 (91%); 5.8 refs/pp. Conditional, bounded by role and proximity to decision-makers.
RE	13/14 (93%); 5.5 refs/pp. Inconsistency framed as acceptable "variation."	11/11 (100%); 7.5 refs/pp. Inconsistency experienced as daily friction requiring workarounds.

RE (inferred from EF/FS/VL configuration; not directly coded)

4.6.2 Towards Governance Blindness, Distance and Invisibility in the Oversight Configuration

Pattern: Although EF and FS are common in Oversight discourse (Section 4.5.1), they appear as structurally displaced phenomena, with governance participants describing friction and suppression as features of lower-level operations rather than barriers they personally encounter. This displacement does not suggest absence but rather structural distance. In the Oversight data (MS01, MS06; ART-ERCC convergences), governance participants describe friction and suppression as features of lower-level operations rather than barriers they personally encounter, suggesting that what governance does not directly experience may be difficult to address.

Mechanism: Three compounding factors produce blindness. By the time concerns reach Oversight level, procedural gates, manager gatekeeping, and cognitive pre-screening have stripped away operational texture... Cross-unit inconsistency is simultaneously framed as contextual adaptation by governance and experienced as systemic dysfunction below it... Finally, structural checks build confidence that mechanisms function, masking behavioural failures.

Consequence: Oversight participants in this sample describe what they experience as a well-functioning risk culture, with formal structures in place, and escalations are refined, while friction is described as occurring elsewhere. However, Operations and the 2nd Line report systemic barriers and workarounds. The data suggest this is not a façade, rather, it is that governance actors do not directly experience what organisational architecture prevents them from receiving. What is analytically legible here is structural non-reception: governance may be architecturally unable to receive what its own procedures prevent from surfacing, the friction is not hidden from governance, it is filtered by governance.

What distinguishes this configuration from the others, if the inference holds, is that the data are consistent with entropy building not through active blocking but through structural non-reception, and that governance may not directly encounter the friction the architecture it mandates generates below it.

4.6.3 Towards Control Function Silencing, Absorption and Authority in the 2nd Line Configuration

Pattern: The second line's ubiquitous FS despite structural independence (Section 4.5.2) is consistent with the interpretation that formal authority does not ensure behavioural influence. Mechanisms are present but are absorbed through procedural acknowledgement. Business-unit experience (Boundary Negotiators, provisional) is used in this synthesis to illustrate how commercial framing shapes escalation decisions, but the weight of the decoupling argument rests on the three fully saturated configurations. The 2nd Line data, read together, are consistent with three absorption patterns, though whether these constitute stable mechanisms or context-specific responses is a question addressed in Chapter 5.

Mechanism: Three absorption patterns operate. Items are formally recorded but not actioned, procedural acknowledgement substitutes for corrective response. Despite structural legitimacy at governance level, 2nd Line staff must argue for uptake case by case. Reinforcement is scrutiny-dependent rather than internalised: challenge intensifies under audit or regulatory review and recedes when scrutiny lifts.

Consequence: For the 2nd Line, structural independence and formal escalation rights result in frequent agenda items but limited corrective impact—items are noted, minuted, and closed without meaningful change (one senior 2nd Line participant notes an emerging norm of avoiding written records for sensitive communications, consistent with this absorption dynamic: MS41) or are deferred through repeated rescoping. The data suggest a divergence between formal representation and behavioural experience: at the governance level, minutes record items as raised and discussed, while 2nd Line participants describe absorption, deferment, and repeated credibility contests; what is formally visible and what is experienced appear, in this material, to be tracking different things.

Where Governance Blindness appears to operate through distance, this configuration suggests a different dynamic: concerns are formally visible, raised, minuted, closed, yet the data point toward a pattern in which procedural recording may substitute for corrective response. Procedural acknowledgement, in this configuration, appears to substitute for corrective response; formal visibility and behavioural impact appear to track different things.

4.6.4 Towards Operational Normalisation, Workarounds in the Operations Configuration

Pattern: Operations' ubiquitous EF (Section 4.6.3), and the correspondingly strong RE inference signal, do not translate into escalation volume because staff develop localised workarounds rather than formal challenges. Adaptation appears to prevent entropy from surfacing. The Operations data point toward three normalisation processes, their systemic character is examined further in Chapter 5.

Mechanism: Three normalisation processes operate. Learned futility converts initial willingness into pre-emptive self-censorship. Each workaround represents an unescalated signal of dysfunction, resolving inconsistency locally without surfacing it to governance. Voice fatigue depletes motivation, turning repeated experience of concerns noted but not resolved into durable silence. A fourth process operates at the level of documentation: at least one senior participant (MS41) identifies an active cultural shift away from written records, linked to the risk of discoverability in employment disputes, "there is now a culture of don't write things down, if you can help it." Where the first three processes describe how concerns fail to reach governance, this fourth describes how concerns that do circulate informally are prevented from entering the institutional record, producing information

fragmentation through omission rather than blockage, a pattern consistent with the absence of forum minutes noted in Section 4.7.3.

Consequence: Over time, workarounds and learned futility become normalised; concerns are managed locally rather than escalated, and the absence of formal escalation is misinterpreted as the absence of problems. This creates a governance blind spot—the unit closest to operational risk has the least influence on policy, and governance perceives silence as reassurance rather than structural suppression.

The data from this unit suggest that entropy may build not through structural distance alone but through learned response, repeated experience of concerns being noted but not resolved appears, in this sample, to reduce motivation for future escalation. The unit with the highest proximity to risk appears, in this material, to have the least access to the governance processes where that knowledge could act.

4.6.5 The 3.7:1 Discourse-Practice Gap: Quantifying Decoupling

To summarise the pattern across units, a simple decoupling gap indicator was created by comparing the relative prominence of contextual mechanisms to core heuristic behavioural indicators within each unit's coded material; the results are presented in Table 4.10. This index is a qualitative diagnostic tool derived from code frequencies, and it is not intended as a generalisable or statistically precise measure.

Table 4.10 Decoupling Index by Embedded Unit

Unit	Mechanism Discourse (MD)	Proximate Behavioural (PB)	RE Synthesis Segments ‡	Decoupling Gap (MD-PB)	Ratio (MD:PB)	Configuration Label
------	--------------------------	----------------------------	-------------------------	------------------------	---------------	---------------------

Oversight (Board/SMT)	307	77	45	230	4.0:1	Symbolic Governance
2nd Line Risk and Compliance	235	59	27	176	4.0:1	Control Without Voice
1st Line Ops and Support	444	132	67	312	3.4:1	Silenced Executors
1st Line Business (prov.)	141	33	21	108	4.3:1	Boundary Negotiators (Provisional)
Aggregate	1,127	301	160	826	3.7:1	

‡ RE synthesis segments mark passages where the emergent entropy condition is most analytically legible from the EF/FS/VL configuration; they are not added to the PB count as a fourth equivalent variable and are shown here for interpretive context only.

MD = mechanism discourse references across all six contextual mechanism codes (BRE: Balancing Risk-Efficiency; ICF: Independence of Control Functions; MRC: Mechanisms for Raising Concerns; OAV: Open to Alternate Views; RC: Reinforcement of Challenge; RCF: Representation of Control Functions). PB = proximate behavioural references (EF: Escalation Friction; FS: Feedback Suppression; VL: Voice Legitimacy). Ratio = MD ÷ PB.

Table 4.10 decomposes the aggregate 3.7:1 ratio by embedded unit; the distribution is uneven, with Operations carrying the largest absolute gap (312), consistent with the Silenced Executors configuration. Operations carry the largest absolute gap (312) consistent with the Silenced Executors configuration: the unit with the highest proximity to risk also has the widest distance between what is discussed and what is behaviourally experienced. Oversight's gap (230) is notable because it reflects high contextual discourse (governance extensively discusses mechanisms) alongside substantial but structurally displaced behavioural friction, friction that governance mandates architecturally but does not encounter directly (Section 4.6.2).

Although the empirical material is cross-Sectional, several artefact sequences (for example, multi-year deferral of Action 192 and repeated ‘noting only’ items across three ERCC packs) provide limited temporal depth. These sequences do not constitute a full longitudinal study, but they do show how the same friction and suppression patterns recur across decision cycles, offering a minimal evidential basis for inferring erosion of cultural vigilance over time.

The 3.7:1 ratio does not suggest hypocrisy; participants across all units describe genuine commitment to openness, mechanisms for raising concerns, and reinforcement of challenge, and organisational artefacts show sustained investment in these structures. Used as a heuristic diagnostic marker rather than a precise statistical measure, the ratio's significance lies not in the number itself, but in what it consistently indicates across every unit, mechanisms are discussed far more extensively than they are behaviourally experienced.

Figure 4.6 synthesises the four decoupling mechanisms as a self-reinforcing cycle, with the 3.7:1 discourse–practice gap as its quantified output: each mechanism sustains the conditions for the others, and no unit perceives the system-level pattern it collectively produces.




Figure 4.6: Four Decoupling Mechanisms as a Self-Reinforcing Cycle Governance Blindness, Control Function Silencing, and Operational Normalisation, with the 3.7:1 Discourse-Practice Gap as Quantified Output. Artefact corroboration of this pattern across three independent evidence streams is shown in Figure 4.7.

The ratio operationalises what Ashby, Palermo and Power (2012) term simulated alignment, the condition in which governance mechanisms are visibly active while their behavioural effects remain unverified and responds directly to Sheedy and Griffin's (2018) observation that existing risk-culture instruments lack decomposable behavioural metrics capable of distinguishing governance discourse from enacted practice. Its diagnostic value lies not in the absolute value of 3.71 but in the internal distribution across units (Table 4.10): Operations carries the largest absolute gap (312), confirming that the unit with the greatest proximity to operational risk has the widest distance between what is espoused and

what is behaviourally experienced; the theoretical implications of this distribution are developed in Chapter V (Section 5.9).

4.7 Artefact Triangulation: Validating the Two-Level Pattern

The 3.7:1 ratio is derived from coded interview data, but the pattern it captures is also legible in the institution's own records.

Purpose: Three streams of formal artefact evidence, (1) ERCC committee minutes and packs (October, December 2025; retrospective May 2025), analysed in Section 4.7.1; (2) Peakon engagement survey data (December 2023–April 2024), analysed in Section 4.7.2; and (3) action logs and forum Terms of Reference, which appear across both streams, independently validate the two-level pattern, revealing policy-practice divergence that interview accounts alone cannot establish. Artefacts are presented through anonymised paraphrases and descriptive summaries; each reference is linked to the Triangulation Tracking Matrix (Appendix D); raw documents are withheld to preserve participant and institutional confidentiality, with anonymised extracts and analytic memos retained in the audit trail (Appendix D and Appendix B.3). Across 18 triangulation instances (summarised in Table 4.11), convergent evidence validates interview claims through independent corroboration, Divergent evidence diagnoses where mechanisms exist formally but fail behaviourally, and Novel Insights surface patterns not visible in interviews alone.

The three streams were selected because each occupies a different position in the governance architecture: ERCC committee packs are the primary formal record of how risk concerns are received at governance level; Peakon engagement survey data were encountered within those packs as an instrument the institution was itself using to assess voice culture, providing an artefact-embedded self-assessment independent of the research

design; and forum Terms of Reference are the structural documentation of the escalation channels examined in Sections 4.3 and 4.4, access constraints meant board-level minutes and internal audit reports were not available (Section 3.2.3)."




Figure 4.7: Triangulation Evidence Map
Cross-referencing interview findings against three independent artefact streams. This figure provides evidential corroboration for the analytical architecture shown in Figures 4.1–4.6.

This synthesis section draws on the evidential base established in Section 4.6; it presents no new primary evidence but shows how the four-unit configurations, read together, produce a self-reinforcing system-level pattern that no single unit perceives.

4.7.1 ERCC Committee Minutes and Packs: Persistent Reds and Deferral Patterns

The ERCC record provides the most direct independent test of the two-level pattern. If the interview accounts are accurate, governance minutes should simultaneously document formal process adherence and record unresolved risks, procedural activity and substantive inaction coexisting in the same document. The data are consistent with precisely this configuration across three consecutive packs (October and December 2025; retrospective May 2025). Three patterns emerge from the record that interview data alone cannot establish: procedural deferral as a documented governance reality rather than participant perception; noting-without-action as the precise mechanism through which feedback suppression operates at formal governance level; and cross-unit inconsistency as a documented condition from which the emergent entropy inference gains independent evidential weight.

Procedural deferral

The most persistent item in the action log is Action 192, Target Country regulatory framework implementation, deferred from 2023 to 2025, repeatedly rescoped, and cited against Group approval dependencies and resource constraints (ART-ERCC/Oct2025). *Convergent, consistent with MS01 and MS46 descriptions of procedural barriers producing escalation dropout.* What the record adds is not the fact of deferral but its duration: the same action recurs across multiple governance cycles without corrective closure, indicating that procedural friction is embedded in the governance architecture rather than attributable to individual circumstances.

The December 2025 pack records the RSCA (Risk and Control Self-Assessment) uplift as formally deferred to 2026, citing R23 resourcing as the specific constraint (ART-ERCC/Dec2025). *Convergent, this is consistent with MS46: "dropped beforehand because*

of time, resources." What the artefact adds is precision: the resourcing constraint is not a perception but a named, minuted decision. The interview account describes the experience of procedural friction; the governance record documents its formal institutionalisation.

Noting-without-action

The reference to TPRM (Third Party Risk Management) appears red across three consecutive ERCC packs, October 2025, December 2025, and the retrospective May 2025 pack, without any change in status across quarterly reviews (ART-ERCC/Oct2025, Dec2025, May2025). *Convergent, this is consistent with MS03's description of items being "noted" but action delayed.* The artefact record is more diagnostically precise than the interview account: it is not simply that action is delayed, but that the same risk status is formally recorded, reviewed, and closed as "noted" across three consecutive cycles, suggesting that procedural acknowledgement has become the governance response rather than a precursor to it.

A recurring classification across both the October and December 2025 ERCC agendas is 'for noting only', no action required, no follow-up recorded (ART-ERCC/Oct2025, Dec2025). *Divergent, artefacts record formal acknowledgement while interviews describe an absence of substantive response (MS03: 'noted, then sits there').* The analytical significance is not that acknowledgement replaces action in any single instance, but that it appears as a recurrent governance category, a formal designation that closes the item without requiring a response.

Cross-unit inconsistency

The starkest cross-unit discrepancy in the October 2025 pack is the gap between two governance artefacts produced in the same cycle: the 1st Line Business Operational

Risk Forum records 'no items to escalate' while the 2nd Line dashboard simultaneously shows TPRM and liquidity concentration as red (ART-ERCC/Oct2025). *Divergent, this is not a perception difference between participants but a documented discrepancy between two records.* It is consistent with the cross-unit inconsistency from which the emergent entropy condition is inferred (Section 3.3.5): fragmented risk information flow visible in the institution's own records, not only in interview accounts.

What the ERCC record makes possible is a shift from inference to independent corroboration. The discourse-practice gap identified through interview coding and quantified in the 3.7:1 ratio (Section 4.6.5) is not an artefact of participant perception or researcher interpretation, it is legible in the documentary record the institution itself generates. Procedural deferral, noting-without-action, and cross-unit inconsistency are each recorded in committee minutes and action logs. The artefact evidence does not simply validate the interview findings; it anchors them in the institution's own account of what occurred.

4.7.2 Peakon Engagement Survey: Freedom of Opinion vs. Escalation Friction

The Peakon survey data introduce a different kind of triangulation problem from the ERCC record, not a gap between formal procedure and formal outcome, but a gap between what a well-regarded workforce sentiment instrument captures and what the interview data reveal. The April 2024 Freedom of Opinion score of 8.9 on a ten-point scale (ART-Peakon/Apr2024) is, on its face, strongly positive. Read in isolation, it would suggest an organisation in which staff feel genuinely free to express views. The analytical interest of this score lies not in what it confirms but in what it cannot see.

The 8.9 score is consistent with the 17 OAV→VL positive co-occurrences documented in Section 4.4.1, where participants describe genuine openness within trusted relationships, direct manager access, and team-level voice as authentic rather than performative. *Convergent, the survey independently corroborates what the interview data indicate about local authenticity.* This convergence matters methodologically: it establishes that the Section 4.4 findings are not a researcher artefact or an optimistic reading of ambiguous data. Staff do experience genuine freedom to speak, and a standardised, anonymised workforce instrument independently records the same thing.

The most significant feature of the 8.9 score is that it coexists with interview accounts of pervasive escalation friction, self-censorship, concerns being "dropped beforehand", and challenge becoming "quieter when not under scrutiny" (MS03; MS46). Divergent, novel insight. The score does not contradict those accounts; it illuminates why they are structurally invisible to standard survey instruments. The Freedom of Opinion question asks whether staff feel able to speak up but does not distinguish between speaking up within a team and escalating concerns beyond a direct manager into formal governance channels. Participants answer based on the experience most salient to them, which, as the Section 4.3 data indicate, is local and relational. The systemic failure documented in Section 4.4 operates at a different level, one that the survey instrument, by design, does not reach. This measurement bias is not confined to a single account: five independent participants across Oversight (MS01), Operations–IoM (MS21), 2nd Line (MS41), Board (MS47), and Operations–Jersey (MS29), none of whom were in contact with each other during the fieldwork window, each describe the same mechanism, namely that low Peakon scores are linked to reduced bonus pools, creating an incentive to inflate ratings. The cross-unit consistency of this account, spanning all three lines and both jurisdictions, adds

independent triangulation weight beyond the survey artefact alone: the 8.9 score is not simply a methodological curiosity but a structurally produced outcome.

The result is that the high score and the widespread friction are not in tension: they are both accurate readings of the same organisation at different levels of resolution. What the two-level pattern predicts, genuine local openness coexisting with formal escalation failure, is precisely what the convergent/divergent pairing here captures in a single instrument. The survey score is high because mechanisms work locally; escalation friction is widespread because the same mechanisms fail when voice attempts to scale. The 8.9 score is therefore not evidence against the two-level pattern, it is independently consistent with it.

This has a methodological implication the data make visible in this case, though its generalisability is a question for Chapter 5: standard engagement surveys measuring team-level psychological safety may systematically miss cross-level escalation friction, recording the local condition as representative of the whole. Within this case, the survey score appears to capture the Section 4.4 reality accurately while remaining structurally unable to register the Section 4.5 condition, not because the instrument is poorly designed, but because the two levels of the pattern operate below the threshold of resolution that team-level survey items are built to detect.

4.7.3 Convergence, Divergence and Novel Insights

The 18 artefact instances analysed across Sections 4.7.1 and 4.7.2 do not tell a uniform story, and that is their analytical value. Eleven instances (61%) converge with interview-derived patterns, five (28%) diverge from them, and two (11%) surface insights

that interview data alone could not have produced. Table 4.11 summarises these outcomes across the three artefact streams, categorised by convergence, divergence, and novel insight.

Table 4.11 Triangulation Evidence Summary

Category	Count	%	Examples
Convergent	11	61%	Persistent reds (TPRM x 3 packs); "Noting only" items; multi-year deferrals; Freedom of Opinion score validates local authenticity
Divergent	5	28%	1st Line "No escalations" vs. 2nd Line dashboard reds; Freedom of Opinion vs. friction; Forum terms of Reference (ToRs) vs. no minutes; Policy exists vs. "would have to find it"
Novel Insights	2	11%	Forum consolidation (monthly to quarterly); AI minutes trials revealing process vs. outcome focus

The 61% convergent rate means that the two-level pattern, genuine local voice legitimacy coexisting with formal escalation failure, is independently legible in the institution's own records. Persistent dashboard reds, multi-year action deferrals, and noting-only committee patterns are not inferences drawn from participant accounts; they are documented in the minutes, packs, and action logs the institution generates for its own governance purposes. What this convergence establishes is not simply that the interview findings are credible, but that they are corroborated by a source that has no stake in the researcher's interpretation.

The 28% divergence rate is, however, the most analytically productive dimension of the triangulation, not a weakness in the evidence base, but the precise location where the discourse-practice gap is most structurally visible. The Divergent instances share a common architecture: formal mechanisms are documented and apparently active, Terms of Reference for escalation forums, whistleblowing policies, risk dashboards showing live

red statuses, yet behavioural engagement with those mechanisms is weak. Staff report that they "would have to find" the whistleblowing policy. The 1st Line Business forum records "no items to escalate" against a backdrop of 2nd Line dashboard reds. The gap between what artefacts record and what participants experience is not a measurement inconsistency; it is the structural condition the two-level pattern is designed to describe. Formal completeness and behavioural enactment appear to track different things.

The two Novel Insights (11%) are worth distinguishing from each other. Forum consolidation, the shift from monthly to quarterly escalation frequency, is a structural change that reduces the available pathways for voice to reach governance, compounding the friction documented in Section 4.5. This is not visible in interview data, where participants describe the current forum cadence as normal rather than as a recent reduction. The second novel insight, the AI-generated minutes trial, is analytically significant for a different reason: minutes focused on process adherence rather than substantive challenge content reveal that the recording of governance activity and the capture of its quality are not the same operation. The artefact analysis itself exposed a methodological limitation, that formal documentation of meeting occurrence cannot be treated as evidence of genuine challenge behaviour within those meetings.

Taken together, the three categories of artefact evidence contribute differently to the overall argument. The convergent instances indicate what fails, specific risks persist unresolved, actions accumulate beyond their target dates, and feedback loops close without corrective action. The Divergent instances indicate how failure is sustained, formal structures provide the appearance of functioning governance while operating below the threshold at which substantive challenge would need to occur. The Novel Insights indicate

what accelerates the failure, structural changes to forum frequency and documentation practice that reduce the density of genuine challenge opportunity over time. This three-part contribution is captured in Table 4.12, which consolidates the quantitative artefact evidence into a single reference point.

Table 4.12 Internal Corroborating Evidence Summary

Claim	Evidence	Source
Escalation lag	Average exception resolution time increased from 12 to 31 days	Action log analysis (ERCC Dec 2025)
Action log accumulation	47 of 192 open actions; 25% open >6 months	ERCC action logs (Oct/Dec 2025)
Dashboard noise	Standard reviews persistently red; high-risk periodic reviews green	ERCC dashboards (May/Oct 2025)
Geographic quality divergence	IoM operations 22% higher error rates than Jersey HQ	Operational quality metrics
Peakon survey divergence	Freedom of Opinion 8.9 vs. interview accounts of friction	Peakon surveys (Dec 2023–Apr 2024)
Measurement bias	Staff soften scores because low ratings reduce bonus pools	Interview evidence (INT-02, INT-03)

The operational metrics in Table 4.12 anchor the qualitative patterns in concrete terms. The escalation lag increases from 12 to 31 days (ERCC Dec 2025, action log analysis) and the accumulation of 47 open actions, 25% of which have been open for more than six months (ERCC Oct/Dec 2025), provide a measurable dimension to the procedural friction described throughout Section 4.4. The geographic quality divergence, IoM operations recording 22% higher error rates than Jersey, independently corroborates the interview pattern in which the Isle of Man unit describes the highest intensity of EF and learned futility. The location experiencing the most pronounced escalation barriers in interview data is also the location where artefact-level operational quality metrics are most degraded. This cross-source consistency, across qualitative and quantitative artefact evidence, is the strongest triangulation finding in the Section: not that single instances

agree, but that the overall configuration of evidence converges on the same structural account.

4.8 Summary of Findings

This chapter applied three proximate behavioural indicators, Escalation Friction (EF), Feedback Suppression (FS), and Voice Legitimacy (VL), and inferred Risk Entropy (RE) as an emergent synthesis construct from their cross-unit configuration, across data drawn from 25 semi-structured interviews, organisational artefacts, and reflexive memos. The analysis is consistent with a two-level pattern in how communication and challenge function within the case institution. Across the coded corpus, EF, FS, and VL account for 301 proximate behavioural references (EF: 95, FS: 43, VL: 163), with 160 passages where the cross-indicator configuration from which RE is inferred is most analytically concentrated (see Section 3.3.5 for the ontological distinction between RE as inference and EF/FS/VL as proximate coded indicators). Contextual mechanism codes account for 1,127 references across the same data, producing a 3.7:1 discourse-practice ratio, an exploratory heuristic indicator of the gap between espoused and enacted risk culture. This ratio is derived from two non-equivalent code families (contextual mechanism codes and proximate behavioural indicator codes) and should be interpreted as a descriptive, directional signal rather than a validated metric. Its value lies in comparative illustration, highlighting the relative weight of governance discourse against behavioural evidence, not in precision measurement. Sensitivity to code-family definition, interview prompt density, and transcript length normalisation is acknowledged, with further testing against those variables is identified as a priority for future comparative work (Section 5.12).

At the local level, contextual mechanisms, Openness to Alternate Views (OAV), Mechanisms for Raising Concerns (MRC), and Reinforcement of Challenge (RC), translate into genuine voice legitimacy within trusted relationships characterised by relational proximity, interpersonal reciprocity, operational framing, and low reputational risk (Section 4.3). At the systemic level, the same mechanisms fail to prevent escalation friction and feedback suppression when voice attempts to scale into formal organisational settings. Three sequential EF barrier types, procedural, relational, and cognitive, convert initial willingness to raise concerns into withdrawal or silence (Section 4.4.1). Reinforcement remains scrutiny-dependent rather than culturally internalised, intensifying under audit or regulatory pressure and receding when that pressure lifts, a temporal inference drawn from repeated patterns across artefact sequences, not from direct longitudinal observation of the same actors (Section 4.4.2). Openness is subject to procedural conversion, performative substitution, and confidence-induced closure, each of which prevents discourse-level openness from generating systemic-level voice effectiveness (Section 4.4.3). Together, six compounding barriers, procedural gates and resource constraints (procedural); accountability fragmentation and manager gatekeeping (relational); and confidence closure and scrutiny dependency (cognitive), together convert local authenticity into systemic failure (Figure 4.4; Section 4.4.4).

Indicator profiling across the four embedded units identifies four distinct unit-level entropy signatures (Section 4.6): Symbolic Governance (Oversight), where governance mandates the escalation architecture but does not encounter the friction it creates, resulting in ubiquitous FS through procedural absorption; Control Without Voice (2nd Line), where formal structural independence is behaviourally challenged by procedural absorption and credibility contests; Silenced Executors (1st Line Operations), where the unit with the

highest proximity to operational risk and the widest indicator spread has the least governance influence, with entropy particularly pronounced in the Isle of Man subsample; and Boundary Negotiators (1st Line Business, provisional), where voice legitimacy is contingent on revenue framing rather than operational substance, this signature carries lower confidence than the three saturated configurations and is carried forward as an indicative contribution.

Read together, the four unit configurations produce three aggregate decoupling mechanisms (Section 4.6): governance blindness, in which structural distance prevents Oversight from encountering the friction its own architecture produces; control function silencing, in which 2nd Line challenge is formally visible but procedurally absorbed without corrective impact; and operational normalisation, in which Operations resolves dysfunction through workarounds, preventing entropy from surfacing to governance. These three mechanisms are self-reinforcing, each sustains the conditions for the others, and no unit perceives the system-level pattern it collectively produces. Diagnostic cuts confirm that EF, VL, and RE intensities increase consistently from Oversight/Board/SMT to Operational staff (Table 4.8), and that the Isle of Man records higher indicator intensity than Jersey across EF, VL, and RE (Table 4.9), two independent analytical dimensions converging on the same structural account. The 3.7:1 ratio is not a measure of individual intent: participants across all units describe genuine commitment to openness and challenge, and organisational artefacts record sustained investment in these structures; the gap between discourse and practice appears structural rather than individually motivated.

Artefact triangulation across 18 instances (Section 4.7) independently corroborates the two-level pattern: 61% of instances converge with interview-derived findings, 28%

diverge in ways that make the discourse-practice gap structurally visible through formal governance records, and 11% surface Novel Insights not accessible through interview data alone. Quantitative artefact evidence anchors these patterns concretely: escalation lag increased from 12 to 31 days; 47 open actions accumulated across the review period, 25 of which had been open for more than six months; and the Isle of Man recorded 22% higher operational error rates than Jersey, cross-source consistency that extends the two-level pattern from qualitative interpretation into measurable governance terms (Table 4.12). Chapter 5 develops the theoretical and practical implications of these findings, situating the two-level pattern within existing risk-culture literature and examining what the structural character of the discourse-practice gap signifies for governance design, each research question's contribution, and the diagnostic potential of the heuristic indicator framework.

4.9 Chapter Summary

Three data sources, 25 participants, and twelve evidence tables converge on the same structural account: in this institution, mechanisms for challenge are extensively documented and genuinely experienced within trusted relationships, yet they fail to scale into the formal governance channels where their effects would matter most. Escalation Friction, Feedback Suppression, and Voice Legitimacy, the three proximate behavioural indicators applied across Sections 4.3 and 4.4, each indicate the same boundary condition: the relational dyad, not the governance architecture, is the unit within which challenge operates authentically.

The unit-level entropy signatures developed in Section 4.5 make that boundary visible at different structural locations. Oversight is characterised by architectural non-

reception, filtered escalations are received in polished form, and friction-generating structures are also what filter the friction they produce. The 2nd Line holds formal independence but experiences procedural absorption: challenge authority does not reliably convert into behavioural influence. Operations records the highest indicator spread and the strongest RE inference signal, those closest to operational risk have the least governance access, resolving dysfunction through workarounds rather than escalation. The Business unit, at lower confidence, adds a fourth variant: commercial framing functions as a legitimacy gate, pre-filtering which concerns acquire voice. Each configuration is distinct; all four share the same underlying mechanism, visible in the 3.7:1 discourse-practice gap quantified in Section 4.6.5: mechanisms are discussed and documented at 3.7 times the rate at which proximate behavioural evidence suggests they are enacted.

Risk Entropy, inferred as an emergent synthesis condition from the cross-unit EF–FS–VL configuration rather than coded directly, is the term applied to this gap between mechanism discourse and behavioural enactment. It is not a finding about individual conduct: participants across all units describe genuine commitment to openness and challenge, and artefact evidence confirms sustained governance investment in these mechanisms. The 3.7:1 ratio, and the unit configurations it summarises, point instead toward a structural condition, governance architecture that mandates the mechanisms that simultaneously produce and obscure the friction those mechanisms are designed to prevent.

Artefact triangulation in Section 4.7 provides independent corroboration across three streams: ERCC committee records confirm that persistent dashboard reds and multi-year action deferrals coexist with procedural completeness; the Peakon engagement data capture the Freedom of Opinion/escalation-friction split that interview accounts describe;

and novel artefact signals, forum frequency reduction, AI-generated minutes focused on process adherence rather than challenge content, surface accelerants of the decoupling pattern not visible in interview data alone. The convergence across sources is the chapter's strongest evidential finding: not that individual instances agree, but that a qualitatively consistent and artefact-corroborated structural account emerges from data collected through four independent analytical channels.

The four research questions are addressed at the evidential level throughout. How alternative views are invited and expressed (RQ1) and how formal mechanisms operate in practice (RQ2) are answered locally by the authenticity configuration of Section 4.3 and systemically by the three-barrier model of Section 4.4. How escalation pathways function in practice (RQ3) is addressed through the 2nd Line control-without-voice configuration and the governance-blindness mechanism of Section 4.6.2. How challenge, monitoring, and learning mechanisms interact over time (RQ4) is addressed by the scrutiny-dependent reinforcement pattern of Section 4.4.2 and the 3.7:1 discourse-practice gap of Section 4.6.5. What the data cannot resolve, and what the heuristic framing of EF, FS, VL, and RE is designed to preserve, is whether these configurations are specific to this institution, this regulatory context, or this analytical moment. That question, and the theoretical work of positioning these findings against the risk-culture and organisational-silence literature developed in Chapter II, is the task of Chapter V.

CHAPTER V:

DISCUSSION

5.1 Introduction

The findings presented in Chapter 4 are consistent with a two-level pattern in how communication and challenge operate within the case institution. Contextual mechanisms, Openness to Alternate Views (OAV), Mechanisms for Raising Concerns (MRC), and Reinforcement of Challenge (RC), appear to operate with genuine behavioural effectiveness within trusted local relationships, yet appear unable to prevent Escalation Friction (EF), contested Voice Legitimacy (VL), Feedback Suppression (FS), and the emergent Risk Entropy (RE) inferred from their systemic configuration when voice attempts to scale into formal governance structures. This two-level pattern is heuristically illustrated through a 3.7:1 discourse-practice ratio, which is an exploratory descriptive signal rather than a validated measure, derived from 1,127 contextual mechanism references against 301 proximate behavioural indicator references, and surfaced inductively through the reflexive thematic analysis specified in Chapter 3.

This chapter interprets those findings in dialogue with the theoretical frameworks established in Chapter 2, moving from local authenticity and systemic failure mechanisms through unit-level entropy signatures and the decoupling synthesis, before drawing research-question-level conclusions and setting out theoretical, practical, methodological, and future-research implications.

The overarching research question provides the interpretive thread throughout: How do communication and challenge behaviours function as early indicators of risk entropy within financial institutions, and what mechanisms either suppress or enable them?

EF, FS, and VL are treated as proximate heuristic diagnostic indicators of *risk voice*, the act of raising, escalating, or challenging a risk-relevant concern together with its narrative provenance, iteratively refined during fieldwork, and RE is treated as an emergent synthesis construct inferred from their cross-unit configuration, not a directly coded fourth parallel indicator.

5.2 Summary of Study and Findings

5.2.1 Overview of the Study

This study investigates the fourth and least-researched pillar of the Financial Stability Board's (FSB, 2014) supervisory framework: communication and challenge. Preceding scholarship has concentrated predominantly on the other three pillars, tone from the top, accountability, and remuneration, leaving the communication and challenge dimension analytically underdeveloped. The study addresses this gap through an embedded single-case design, using semi-structured interviews and interpretive analysis to examine how communication and challenge behaviours function as early indicators of risk entropy in a multi-jurisdiction Channel Islands financial services institution.

The analytic sample comprised twenty-five main-study participants, selected to provide cross-unit coverage across all four embedded governance layers (Oversight/Board/SMT, Second Line Risk and Compliance, First Line Operations, and First Line Business) and across both jurisdictions (Jersey n=14; Isle of Man n=11). Three pilot interviews were conducted in advance to refine the interview guide and were excluded from the analytic sample, consistent with the protocol set out in Section 3.5.2. Selection was deliberately proportional by unit rather than by headcount to prevent overrepresentation by larger functions. The case institution was selected on access grounds:

the Chief Risk Officer's participation in the researcher's professional network was essential to establishing the trust required for a sensitive topic. The institution is multi-jurisdictional, UK-parent-regulated, and had Board-acknowledged communication as an identified development area, making it a particularly relevant setting for the study's diagnostic focus.

Three theoretical frameworks from Chapter 2 run throughout the discussion. The FSB's (2014) four-pillar supervisory framework provides the regulatory anchor against which findings are assessed. Edmondson's (1999) psychological safety model establishes the baseline for understanding local authenticity and its limits. Morrison's (2014) dual safety-efficacy calculus provides the bridging mechanism for understanding why local safety does not prevent systemic failure. Power's (1994, 2009) rituals of verification provide the theoretical lens through which decoupling is interpreted. Together, these frameworks support the move from structural proxies to heuristic behavioural indicators of voice, escalation friction, and challenge.

5.2.2 Headline Findings

The analysis produced four interconnected findings, presented in the same sequence as Chapter 4.

Finding 1, Local Authenticity. Contextual mechanisms translate into genuine Voice Legitimacy within trusted relationships, bounded by four conditions: relational proximity, interpersonal reciprocity, operational scope, and low reputational risk. Thirty-one positive co-occurrence passages and a Peakon Freedom of Opinion score of 8.9 (ART-Peakon/Apr2024) independently corroborate this authenticity.

Finding 2, Systemic Failure. The same mechanisms fail when voice attempts to scale. Three barrier types, procedural, relational, and cognitive, operate cumulatively and sequentially, converting initial willingness to raise concerns into withdrawal. Seven EF--FS negative co-occurrences and nine cross-indicator configurations provide the evidential base.

Finding 3, Unit-Level Entropy Signatures. Four configurations were identified: Symbolic Governance (Oversight), Control Without Voice (Second Line), Silenced Executors (First Line Operations), and Boundary Negotiators (First Line Business, provisional).

Finding 4, Organisational Decoupling. The four configurations aggregate into a systemic pattern captured by the 3.7:1 discourse-practice ratio. The total decoupling gap of 826 is unevenly distributed, with Operations carrying the largest absolute gap (312), the unit with the highest proximity to risk has the widest distance between what is discussed and what is behaviourally experienced.

In plain terms, people can speak up and address issues within their own teams, but once concerns enter formal governance, they encounter delays, filtering, and confusion, so that many real risks never produce timely action at the top. Findings suggest this is less a design failure than a structural feature of this institution, and the data point toward potentially FSB-compliant governance architectures more broadly, in which the architecture of challenge culture appears to have been built without consistently closing the gap between espousing it and enacting it.

5.3 Local Authenticity: Where Mechanisms Succeed

Chapter 4 established that communication and challenge mechanisms function genuinely within trusted local relationships. Thirty-one positive co-occurrence passages, OAV to VL (17); MRC to VL (8); RC to VL (6), indicate that Openness, Mechanisms for Raising Concerns, and Reinforcement of Challenge translate into perceived Voice Legitimacy when four boundary conditions are present: relational proximity, interpersonal reciprocity, operational scope, and low reputational risk. This authenticity does not appear performative; participants describe felt permission, 'I feel I can talk' (MS01) and 'we want your opinion' (MS15), rather than reluctant compliance.

A concrete illustration of what happens when one condition is absent clarifies the analytical importance of this boundary framework. When relational proximity is removed, specifically, in the relationship between the local entity and group-level dotted-line functions, MS06 describes how challenge is replaced by passive acceptance: 'it's sometimes easier to take the super-compliant route... rather than leverage a local understanding.' The capacity for genuine challenge does not disappear; it is suppressed by the absence of the relational foundation that local authenticity requires. The same dynamic surfaces in mirror form: MS47 describes absorbing significant social cost to challenge the group legal team over an inherited tax liability, because without interpersonal reciprocity the full burden of the challenge falls on the individual challenger alone.

The findings are consistent with Edmondson's (1999) foundational work on psychological safety, which establishes that team-level voice depends on interpersonal trust and the perceived consequences of speaking up. The seventeen OAV-to-VL co-occurrences support Edmondson's proposition that safety is a group-level construct mediated by leader behaviour and relational norms. The present study extends this

framework: by identifying the four specific boundary conditions, the findings suggest that psychological safety operates as a local phenomenon that does not automatically transfer to formal organisational escalation, a scaling limitation that Edmondson's original team-level model did not directly address.

At the institutional level, this scaling failure carries documented consequences: the workarounds and informal resolutions documented in the Operations configuration remove the escalation signal that governance would otherwise receive, sustaining the causal conditions for systemic failure even as local problems are resolved.

Fukuyama's (1995) distinction between interpersonal and institutional trust provides the broader theoretical grounding. The Peakon Freedom of Opinion score of 8.9 accurately captures interpersonal trust; the EF, FS, and VL data reveal that institutional trust, the extent to which formal escalation channels are believed to produce consequential action, operates independently, and at a substantially lower level of confidence.

This two-level distinction is not merely a conceptual inference from the present data: Ghafoori et al. (2023), in developing a comprehensive validated measure of risk culture in financial institutions, explicitly operationalise this separation between felt voice safety and systemic escalation efficacy, providing empirical evidence that the two dimensions are independently measurable and do not co-vary predictably within the same institution.

Morrison's (2014, 2023) dual-calculus model provides the bridging architecture: even employees who perceive interpersonal safety may withhold voice when they judge

escalation to be futile. Within local contexts, both safety and efficacy perceptions are positive, producing the authentic voice behaviour captured in the 31 co-occurrence passages. When voice attempts to scale, the safety calculus may remain intact while the efficacy calculus collapses under the cumulative weight of the barriers examined in Section 5.4.

What the local authenticity finding establishes is a critical interpretive baseline: the mechanisms designed to enable voice are not inherently flawed. They appear to succeed where relational conditions permit. The question Section 5.4 addresses are why the same mechanisms fail when voice moves beyond those boundaries.

5.4 Systemic Failure: Where Mechanisms Cannot Scale

5.4.1 Structural Friction as a Driver of Suppression

The clearest single illustration of the three-barrier sequential model comes from the Ideas Bank case. MS16, a Relationship Director, identified an anomaly in product-process alignment, formulated a solution, and submitted it to the Ideas Bank. The concern cleared the cognitive barrier, she understood the problem and had a solution but stalled at the procedural barrier: the routing through a line manager who lacked authority to champion it upwards created a dead end. The concern never reached the relational barrier, because it dissolved silently at the procedural level, invisible to senior leadership. MS01 independently observed the downstream pattern without knowing it was happening: 'I do wonder how some of these ideas have never been escalated before... Do people go to my Head of Compliance, get told no and then just accept that as an answer and go away without realising they could come to me?' The case demonstrates that the three-barrier model is not merely theoretical, it is a real failure chain observable in the data.

Chapter 4 pointed toward a sequential, cumulative process in which EF produces FS. Seven co-occurrence passages indicate that concerns initially intended for escalation encounter procedural barriers, relational hesitancy, and cognitive constraints that convert willingness to raise them into non-escalation or silent withdrawal. The concept of rituals of verification (Power, 2009) provides the directly relevant theoretical lens: procedural gates create the appearance of deliberation while producing dropout. Each procedural step functions as a verification point that governance interprets as due diligence but that operational staff experience as an obstacle course. This extends the verification thesis from an organisational-level critique to a process-level explanation of how governance architecture may produce the very silence it is designed to prevent.

The cognitive constraint dimension represents the behavioural enactment of discourse boundaries. Dominant organisational narratives predetermine what is 'speaking' (Hardy and Maguire, 2016), and staff appear to construct internal barriers through experience: they have learned which concerns are unlikely to survive the escalation process. This cognitive pre-filtering finds grounding in Morrison's (2014, 2023) identification of implicit voice theories, taken-for-granted schemas about the riskiness of challenging authority, which the present study points toward extending from a personality-level schema to an organisationally produced one. The three-barrier model reframes escalation failure as architectural rather than motivational, with the important implication that targeting only one barrier type may simply shift the dropout point rather than eliminating it.

Formal control additions without corresponding redesign of the behavioural architecture simply relocate rather than resolve the escalation barrier, a finding consistent with the three-barrier sequential logic documented here (Grieser and Pedell, 2022).

5.4.2 The Reinforcement Paradox: Espoused but Not Enacted

The paradox that animates this Section is crystallised in a single exchange. MS01 observed: 'We have a minimum score threshold linked to bonus performance, so people also know that if the scores are low, the bonus pool gets lowered. So, are people just going in, hitting nine out of ten, knowing that if I hit a three then the bonus pool might come down? I think the link to bonus pool is poor if I'm honest.' Against this, the April 2024 Peakon artefact records a Freedom of Opinion score of 8.9, 0.8 points above benchmark. The artefact declares an above-average voice culture; the individual architecturally responsible for that culture says the instrument producing the declaration is structurally compromised. This is the paradox at maximum compression: the highest-frequency governance discourse in the dataset, 408 coded references to Reinforcement of Challenge, coexists with a primary measurement instrument that its own custodian does not fully trust.

Three dysfunctional reinforcement forms were identified. *Absent* reinforcement, in which operational risks are neither tracked nor escalated, is evidenced in artefact data: the First Line Business Operational Risk Forum reported no items to escalate while the Second Line dashboard simultaneously showed TPRM red and liquidity concentration red (ART-ERCCOct2025). *Reactive* reinforcement, in which challenge intensifies under regulatory scrutiny then recedes, is captured by MS03: 'challenge is enhanced when we're under that level of scrutiny... and practices a little bit quieter when we're not.' This oscillation is consistent with Power's (2009) thesis on governance activity that intensifies in response to

external stimuli but lacks self-sustaining momentum. *Compartmentalised* reinforcement, in which post-mortems occur within silos without cross-unit propagation, is documented in the First Line Operational Risk Forum (ART-ERCCMay2025: action tracker only, no formal minutes, no dissemination record).

Critically, the finding that contradicts the existing literature is not the existence of scrutiny-dependency, but its durability. The literature largely treats regulatory-pressure-driven improvement as a transitional state on a path toward internalised practice. The data point toward something more structural, such as scrutiny-dependency, may not be a transitional phase but rather the stable equilibrium this institution appears to have reached, at least across the fieldwork window. Baumeister et al.'s (2001, Starling 2025) demonstration that bad events exert disproportionate psychological force on behaviour provides an explanatory mechanism: isolated noting-only episodes may degrade the efficacy calculus at a rate incommensurate with the volume of positive reinforcement discourse surrounding them, a ratchet effect rather than a reversible one.

The performance consequences of this equilibrium are visible in the case data: TPRM remained red across three consecutive ERCC packs without change in status, 47 of 192 open actions remained unresolved, and the 22% IoM-Jersey error differential maps precisely onto the unit carrying the largest absolute decoupling gap. The implication is that scrutiny-dependent reinforcement, which by definition is not structurally embedded, cannot sustain the efficacy perceptions that voice behaviour requires (Kunz and Heitz, 2021).

5.4.3 Openness Without Effect: When Invitation Becomes Ritual

Three dysfunctional forms of openness were identified. The *ritualised* form is most directly evidenced by MS16, who provides a natural articulation of the promotive-prohibitive voice divide that Morrison (2014) theorises. On contributing an idea through the Ideas Bank, she is enthusiastic: 'It's only because we're in this time at the moment where we've got this ideas bank that it might potentially... be something that I think they probably would take on board.' But when asked about raising a risk concern outside that structure, she shifts register: 'I'm not aware of a way that I could report that formally or informally, apart from mentioning it to my line manager, who might think, might be really busy that day... there's not so...' She then explicitly proposes that the bank adopt something analogous to the Ideas Bank for risk concerns, treating the two as categorically separate things. The asymmetry appears structural rather than simply cultural in this case: the organisation has created visible, incentivised infrastructure for promotive voice while leaving prohibitive voice dependent on a line manager who might be busy, with MS16 arriving at this distinction without being prompted toward it.

The conduct implications of this structural asymmetry are visible in the data: the absence of an independent prohibitive voice channel means that concerns dependent on line manager willingness are, in practice, concerns that can silently disappear, a structural condition that creates the precise conditions under which conduct risk accumulates undetected.

The *structurally constrained* form operates through consensus gates that convert openness from a communicative invitation into a bureaucratic process, extending the argument on discourse boundaries (Hardy and Maguire, 2016) to suggest that procedural architecture enacts the same filtering function as cultural norms. The *confidence-based*

closure form operates through a different mechanism: senior participants' genuine belief that the system is functioning effectively is itself structurally produced by receiving filtered, polished escalations. Governance cannot identify what is not being raised when those closest to the issue do not perceive a problem requiring escalation, a form of inadvertent unknowing (McGoey, 2012) sustained through institutional confidence rather than deliberate avoidance.

This structural production of governance confidence is not idiosyncratic to the case institution: Sheedy and Griffin (2018) demonstrate empirically that senior leaders in financial institutions consistently hold rosier perceptions of risk culture than staff, a top-down perception bias arising because assessment instruments capture only the visible, espoused layer, leaving the behavioural gap, and the confidence it generates, structurally undetected.

5.5 Unit-Level Entropy Signatures

The four configurations together map onto the four research questions. Symbolic Governance (Section 5.5.1) addresses RQ2's question about how formal mechanisms operate in practice at governance level; Control Without Voice (Section 5.5.2) addresses RQ3's question about how EF, FS, VL, and RE can be operationalised as diagnostic indicators of the espoused–enacted gap in risk culture; Silenced Executors (Section 5.5.3) speaks to RQ1 and RQ4; and Boundary Negotiators (Section 5.5.4) provides provisional evidence relevant to RQ1. This mapping is consistent with the chapter-level structure shown in Figure 4.1 and Table 4.2, and with the unit-level summaries in Table 4.3 and Appendix A.2, which link each configuration and indicator set back to the four research questions.

5.5.1 Symbolic Governance: The Problem of Structural Distance

The Oversight configuration formally validated (n=5, census, HIGH confidence), describes a governance layer in which formal structures operate as designed yet remain largely disconnected from the friction experienced in practice. Three features are analytically distinctive. First, Oversight participants consistently report receiving polished escalations and aggregated dashboards. The board pack curation mechanism (MS47) makes this structural: the Board Paper Template requires authors to distil concerns to the 'top three points' as an 'elevator pitch' and explicitly instructs that the Management Governance Section, containing operational detail, early-stage debate, and dissenting views, be removed 'before onward provision to the Board.' The polishing of concerns before they reach Oversight is not incidental behaviour by individual gatekeepers; it is an institutionally mandated process. Second, FS is Ubiquitous (5/5) in Oversight, indicating that feedback suppression is not confined to frontline layers but operates structurally at governance level itself. Third, the hierarchical band analysis confirms that EF, FS, and VL are Ubiquitous at every seniority level from S1/S2 through to Board/SMT, confirming that friction and suppression are architecturally pervasive.

Power's work on verification rituals (1994, 2009) provides the foundational lens: organisations substitute visible compliance, documented policies, signed charters, and completed audits for meaningful behavioural scrutiny. Symbolic Governance appears, in this case, less a failure of attention than a product of architectural design, with the escalation system that governance mandates being the same system that may prevent governance from seeing the friction those mandates create. The empirically validated perception gap, Sheedy and Griffin (2018) found that 'senior leaders tended to have a rosier

perception of risk culture than staff generally' is mechanistically explained here: the gap is not produced by governance complacency or frontline exaggeration, but by the three barrier types filtering concerns before they reach governance level. The present study points toward extending this finding further, in which the gap appears not merely hierarchical but structural, embedded in the four entropy signatures and the 3.7:1 discourse-practice ratio. Sheedy and Canestrari-Soh (2023) further demonstrate that formal accountability structures do not reliably close this gap in the absence of behavioural embedding.

Sheedy and Canestrari-Soh (2023) further demonstrate that formal executive accountability structures, including Senior Managers and Certification Regime obligations, do not reliably close this perception gap in the absence of behavioural embedding: accountability mandates establish formal responsibility without generating the feedback loops through which governance would receive unfiltered operational intelligence, leaving the Symbolic Governance configuration intact even under strengthened regulatory accountability frameworks.

5.5.2 Control Without Voice: The Paradox of Formal Authority

The Second Line configuration, fully saturated (n=5, HIGH confidence), names the paradox of structural empowerment coexisting with experiential constraint. The noting-only pattern documented across three consecutive ERCC packs (ART-ERCCOct2025, Dec2025, May2025) provides the artefact-level evidence: concerns are formally recorded in committee minutes but not actioned, creating the impression of governance engagement while corrective action remains absent.

Grieser and Pedell (2022) identify this noting-only dynamic as the central failure mode of formal risk culture controls: in their study of how risk culture development can be controlled, they find that formal recording mechanisms systematically diverge from the behavioural responses they are designed to produce, with governance bodies treating documentation completeness as a proxy for cultural enactment, the same substitution the present artefact data evidence at the level of committee practice.

The motivational consequence of this pattern is captured most vividly by MS06: 'there's almost an exhaustion of trying to meet unclear compliance guidelines... you end up with that exhaustion of... there's been no event, there's been no material event... you get that sort of sigh, why are we doing this.' This is the psychological arc of a concern that has been raised, received, noted, tracked as Amber on dashboards for months, and never fully resolved, to the point where the raising itself begins to feel futile. MS01 independently describes the downstream consequence at his level: staff submitting ideas to the Ideas Bank that he has never seen, suggesting they were turned down lower in the hierarchy and the challengers accepted that answer as final without realising they could escalate further.

The structural critique that risk managers primarily measure and monitor rather than directly manage risk (Stulz, 2016) finds support in this configuration. Three absorption tactics were identified: noting-only absorption; credibility contests, in which the Second Line must argue for uptake despite structural legitimacy; and scrutiny-dependent reinforcement. The regulatory assumption embedded in frameworks such as APRA's CPS 220, that structural independence translates into effective challenge, is not confirmed by this evidence. Structural independence appears necessary but insufficient; the diagnostic

question is not whether the Second Line can challenge, but whether challenge, once exercised, produces observable governance response.

The DNB's direct observational supervision programme found that 34 of 54 Dutch financial institutions reviewed had fundamental behavioural risks invisible in their financial data (Scott, 2024); corroborated by Ashby, Palermo and Power (2012) and Anton and Nucu (2020), both of whom demonstrate that risk culture assessments and disclosures systematically capture the visible, artefact-level layer of cultural practice while leaving behavioural gaps invisible in conventional governance data. This population-level reference point suggests the Control Without Voice pattern may represent the institutional mechanism through which behavioural risk becomes structurally invisible to conventional data systems, a pattern normative rather than exceptional across regulated institutions.

5.5.3 Silenced Executors: Proximity Without Influence

The Operations configuration, fully saturated with high confidence (n=12, saturation at n=7, confirmed at n=8), carries all three EF barrier types simultaneously: almost ubiquitous across every sampled participant (1st Line Operations was 11/12). The 22% higher error rate in Isle of Man operations compared to Jersey HQ (artefact data, Table 4.12) provides measurable performance consequences for what the interview data describe qualitatively. The 22% differential rests on artefact evidence alone; no participant directly connects their escalation experience to a confirmed output error, a gap that is analytically significant: the causal link between escalation failure and quality outcomes is not legible to those experiencing it, visible only in the downstream data.

Morrison (2014, 2023) identifies three motivational pathways to silence: quiescent (fear-based), acquiescent (resignation-based), and prosocial (protecting others). The Silenced Executors configuration appears to exhibit all three simultaneously. Quiescent silence appears in the role-bounded legitimacy pattern ('not my place to question'). Acquiescent silence maps onto the voice fatigue endpoint: staff who have repeatedly experienced escalation producing no corrective action have, in this material, ceased perceiving escalation as viable. Prosocial silence appears in the workaround accumulation phase, where operational staff resolve problems locally rather than escalate through channels perceived as unresponsive. The low FS intensity in Operations (Illustrative: 4/12) does not indicate satisfaction, it may reflect the absence of formal feedback attempts: staff have moved from quiescent silence, where fear inhibits voice, to acquiescent silence, where learned helplessness renders voice irrelevant. This distinction matters for intervention design: quiescent silence responds to safety interventions, whereas acquiescent silence requires efficacy interventions that demonstrate voice produces observable change.

Morrison's (2023) decade-later review further distinguishes acquiescent silence as a distinctly post-experience phenomenon, produced not by fear of immediate reprisal but by repeated observation that escalation produces no consequential response, a distinction that matters for intervention design, since quiescent silence responds to safety-enhancing interventions while acquiescent silence requires efficacy-restoring interventions that demonstrate voice produces observable change.

The paradoxical implication is that the more adaptive and capable the operational workforce, the less governance sees of the friction they experience, because their

competence converts systemic problems into local solutions that never surface as formal escalation.

5.5.4 Boundary Negotiators: Revenue-Legitimated Voice (Provisional)

The First Line Business unit data point toward what this analysis provisionally labels Boundary Negotiators, a configuration in which voice legitimacy appears to be revenue-contingent, with concerns framed in customer or revenue-impact terms receiving governance uptake while operationally framed concerns are deflected. This label carries lower confidence than the three saturated configurations (n=3, partial saturation, PROVISIONAL throughout).

While interview evidence is limited by partial saturation, the convergence of interview and artefact evidence, specifically, the systematic absence of escalation records in Business forums (ART-ERCCOct2025: no items to escalate) despite the Second Line dashboard simultaneously showing TPRM and liquidity concentration red, strengthens the provisional inference beyond what interview data alone would support.

Were future sampling to reveal a marked difference in Voice Legitimacy scores between operational first-line participants and client-facing revenue-generating first-line participants, this would provide evidence that PandL accountability confers voice credibility within financial institutions, a finding with significant implications for how risk culture is managed in commercially-oriented business units. Conversely, if no such differential emerged, the provisional label would require revision toward a simpler seniority-based explanation.

The findings are presented as hypotheses for future research rather than established contributions, consistent with the confidence calibration framework established in Chapter 3.

5.6 Organisational Decoupling: From Unit Configurations to Systemic Pattern

The four unit-level configurations, read together, are consistent with a systematic divergence between the cultural mechanisms the institution promotes, and the behavioural experiences reported by its participants. This pattern is quantified through the 3.7:1 discourse-practice ratio, as derived in Table 4.10: for every proximate behavioural reference (EF, FS, or VL), there are 3.7 references to mechanisms, openness, and reinforcement. The decoupling gap is distributed as follows: Operations 312, Oversight 230, Risk and Compliance 176, and Business (provisional) 108. This distribution is consistent with the proximity-influence inversion, the unit with the highest proximity to risk (Operations) carries the largest absolute decoupling gap, while Oversight, structurally furthest from operational friction, carries the smallest. This is labelled proximity-influence inversion as an interpretive observation drawn from the unit-level data, rather than as a separate theoretical contribution.

The concept of simulated alignment (Ashby, Palermo and Power, 2012) describes the aggregate pattern precisely. The organisation produces 3.7 units of cultural discourse for every unit of proximate cultural friction, creating a documentary landscape in which governance appears culturally healthy while behavioural reality points toward systemic barriers. Evidence that awareness of the gap does exist at EXCO level, both the senior leader participants sense it, reinforces this reading rather than undermining it. The instinct (MS03) is to add a new channel, with the other (MS01), to encourage more informal access.

Neither frames the problem as one requiring a recalibration of the governance architecture itself. This is precisely why decoupling remains emergent: the actors closest to the evidence interpret it as a channel design deficit, not as evidence of a systemic gap between formal risk infrastructure and operative behavioural reality.

This channel-addition instinct is a predictable misdiagnosis, where formal control and channel additions without architectural redesign of the behavioural processes through which they operate relocate rather than resolve the barriers to effective escalation (Grieser and Pedell, 2022).

The ritual thesis (Power, 1994, 2009) argues that the audit explosion produces organisations extensively documented, verified, and assessed yet no more transparent in their actual functioning. The present study is consistent with this thesis at the level of risk culture, pointing toward a recursive reinforcing mechanism: governance mandates escalation structures; those structures generate friction that filters concerns; governance receives filtered evidence and interprets it as evidence of functioning structures; confidence in the system grows; and less interrogation follows. Each mechanism, governance blindness, control function silencing, and operational normalisation, sustains the conditions for the others. Within this case, no unit appears to perceive the system-level pattern it collectively produces.

5.7 Artefact Triangulation: Methodological and Substantive Implications

Chapter 4 documented 18 triangulation instances across ERCC committee minutes and packs, action logs, forum Terms of Reference, and Peakon engagement survey data, yielding 11 convergent (61%), 5 divergent (28%), and 2 novel insight (11%) outcomes.

The action log evidence provides direct artefact-level corroboration of the decoupling thesis. Forty-seven of 192 open actions remained outstanding and 25 had been open for more than six months (ERCC action logs, Oct--Dec 2025). Persistent reds, TPRM appearing red across three consecutive ERCC packs without change in status, provide convergent evidence of exactly the noting-only pattern identified qualitatively. Governance bodies should systematically audit not merely whether items appear on dashboards, but whether dashboard entries produce changes in operational practice over time.

The most analytically powerful single triangulation finding is the Peakon divergence. A Freedom of Opinion score of 8.9 coexists with widespread EF, self-censorship, and scrutiny-dependent challenge in interview accounts. The Peakon instrument measures perceived freedom to voice, a safety perception, but cannot capture deliberate withholding, a silence behaviour. The high score therefore validates local authenticity while potentially missing systemic failure, creating a survey instrument that confirms the organisation's self-image of cultural health while leaving the governance blind spot undiagnosed. Morrison (2014) anticipates this precisely: scales measuring the frequency of voice behaviour cannot be used to infer silence.

The most recent validated instrument development work corroborates the diagnostic gap this creates: Ghafoori et al. (2023) developed a comprehensive risk culture measure for financial institutions specifically designed to separately capture voice safety and escalation efficacy, while Dürst and Kunz (2025) propose effective risk culture assessment protocols that distinguish between safety-level and efficacy-level dimensions as a precondition for diagnostic validity, both confirming that single-dimension

instruments of the Peakon type will systematically misread institutions exhibiting the two-level pattern identified here.

A further distortion layer is documented in the interview data: staff soften Peakon scores because low ratings reduce bonus pools (MS01, MS21), a perverse incentive that suppresses the very signal the instrument is designed to capture, and a clear illustration of a measurement instrument functioning as a ritual of verification.

Sheedy and Canestrari-Soh (2023) demonstrate that this type of incentive-instrument corruption is not idiosyncratic: in their study of executive accountability and risk culture, they find that remuneration linkages to cultural metrics systematically distort the signals those metrics produce, with staff and managers calibrating their cultural reporting behaviour to the incentive architecture rather than to underlying cultural reality.

5.8 Conclusions Aligned to Research Questions

This Section presents the main conclusions of the study in direct alignment with the four formal research questions stated in Section 1.3. For each research question, the conclusion synthesises evidence from the relevant findings Subsections in Chapter 4, structured through the interview themes and indicator mappings outlined in Appendix A.2.

RQ1: How are alternative views about risk encouraged and expressed within the case institution, and under what conditions does risk voice appear legitimate and safe at the local level?

The findings suggest the institution may have conflated the presence of mechanisms with a functioning culture of challenge. Alternative views are genuinely expressed, but

only within trusted, proximate relationships. The Peakon Freedom of Opinion score of 8.9 and 31 positive co-occurrence passages confirm that local authenticity is real rather than performative. However, when voice moves beyond a familiar dyad into formal escalation or governance settings, the same mechanisms designed to enable it, such as open-door policies, the Ideas Bank, and reinforcement frameworks, become barriers: procedural gates pre-filter concerns, reinforcement activates only under regulatory scrutiny and recedes when it lifts, and consultation forums arrive with outcomes already decided. The 3.7:1 discourse-to-practice ratio is the clearest expression of this: the institution has not failed to create space for alternative views; it appears not to have consistently made that space available beyond where personal trust already does.

RQ2: How do formal risk governance mechanisms for communication and challenge operate in practice, and where do escalation barriers and feedback suppression arise as concerns move through formal channels?

Formal mechanisms are designed as open, accessible channels for challenge, but in practice, they appear to function closer to compliance theatre than active challenge infrastructure, documented, tracked, and reported as active, while the concerns they capture rarely escalate into consequential action. The Ideas Bank converts genuine invitations into quota performance; ERCC committee records show persistent, red-rated risks deferred across years without resolution; and the whistleblowing policy exists, but staff report they would have to search to find it. The mechanisms are not broken in design; they are operationally decoupled: the institution can demonstrate that channels exist, while the information flowing through them never reliably reaches the governance level where it would change decisions.

RQ3: How can escalation friction, feedback suppression, voice legitimacy, and risk entropy be operationalised as diagnostic indicators to support boards and supervisors in assessing risk culture in structurally compliant institutions?

The control function's structural role, formal independence, escalation rights, governance access, does not reliably translate into behavioural influence. The Second Line data reveal what the findings label Control Without Voice: concerns are raised, minuted, and formally closed, but procedural acknowledgement substitutes for corrective response. Credibility must be argued for case by case rather than assumed from the mandate, and reinforcement intensifies only under regulatory scrutiny before receding once it lifts. Structural authority and felt influence are inverted, the function with the highest formal standing experiences the most procedural absorption. The dominant literature on three lines models largely treats structural independence as sufficient for effective challenge; the data expose a gap the literature has not adequately theorised: the distinction between formal authority and behavioural uptake. Procedural completeness and behavioural enactment can diverge indefinitely, and formal governance records, the very evidence base most research relies on, are precisely the artefacts that make the divergence invisible.

RQ4: How do the observed unit-level entropy configurations and indicator patterns aggregate into a systemic governance pattern, and what does this reveal about organisational decoupling and the conditions under which risk entropy emerges?

The data suggest reinforcement of challenge was not so much eroded as never consistently institutionally embedded, at least within the governance architecture observed here. The data show it is sustained locally through interpersonal reciprocity within trusted relationships and activated externally by regulatory or audit pressure; in the absence of

both conditions, it recedes. The telling signal is that RC carries the single highest code frequency in the entire corpus, 408 references, yet coexists with the configurations from which entropy is inferred. The institution does not lack commitment to reinforcement in discourse; it lacks a mechanism that makes reinforcement self-sustaining, independent of scrutiny or personal relationships. The most genuinely surprising finding is the inversion of volume and effect: the literature largely treats reinforcement as a cumulative process in which more extensive articulation of challenge norms produces deeper embedding over time. The data flatly contradict this. More reinforcement talk does not produce more reinforcement behaviour, it may actually obscure its absence, because extensive mechanism discourse gives governance the appearance of a functioning challenge culture without requiring the behavioural evidence that would confirm it.

5.8.1 Overarching Conclusion

Financial institutions have spent two decades building the architecture of challenge culture (FSB, 2014), codes of conduct, speak-up lines, risk forums, three lines mandates, board-level attestations of openness. The implicit theory underlying all of it is that if you design enough channels, articulate norms loudly enough, and assign formal responsibility for independent challenge, the behaviour will follow. This study's most important finding is that this theory is materially incomplete.

What the data appear to show is a decoupling problem rather than a design problem. The early signal of risk entropy is not the absence of mechanisms, it is the gap between how extensively those mechanisms are discussed and how rarely they produce consequential action. The mechanisms that suppress challenge are not dramatic: nobody is silencing dissent, nobody is punishing whistleblowers. What the data show instead is

quieter and structurally harder to fix: concerns are pre-filtered before they reach any channel; procedural gates absorb energy that would otherwise escalate; feedback is received, minuted, and formally closed without a corrective response; reinforcement intensifies under scrutiny and recedes without it. Cumulatively, they produce a system in which the people closest to operational risk have the least influence on governance, while those with the most governance access receives only filtered signals that confirm the system is working.

The practical implication is that institutions need diagnostic tools that measure the behavioural gap, not the mechanism inventory. The question is not 'do we have a speak-up line?' but 'when someone used it, what happened next?' Not 'is our Freedom of Opinion score high?' but 'does that freedom extend to formal escalation, or only to conversations with people you already trust?' The distance between those two questions is where risk entropy lives, and where it is measurable, before anything goes wrong.

5.9 Theoretical Contributions

The findings yield four contributions that collectively reframe how risk culture is diagnosed, how its failure propagates, and how governance might respond. In summary form:

- The Two-Level Pattern reframes risk culture from a unitary attribute into a construct read at both relational and institutional levels.
- The Sequential Friction-to-Entropy Mechanism specifies the chain through which procedural barriers compound into inferred risk entropy.
- Unit-Level Entropy Signatures provide a configurational typology, Symbolic Governance, Control Without Voice, Silenced Executors, and

Boundary Negotiators (provisional), for diagnosing subcultural variation within a single institution.

- The Discourse–Practice Ratio (3.7:1) operationalises simulated alignment as a decomposable, documentary-record-based diagnostic marker.

Each is developed in turn below.

The Two-Level Pattern as a Diagnostic Reframing. The identification of local authenticity coexisting with systemic failure within the same institution under the same formal governance arrangements challenges frameworks that assess risk culture as a unitary organisational attribute. Instruments including the Macquarie Risk Culture Scale (Sheedy and Griffin, 2018) and Edmondson's psychological safety model (1999) each assess culture at a single level, producing scores that do not distinguish between where mechanisms succeed and where they fail. The theoretical grounding for this reframing draws on Fukuyama's (1995) distinction between interpersonal and institutional trust, which may reflect a fundamental property of trust itself, operating through different mechanisms at relational and institutional levels. This shifts the diagnostic question from “does this organisation have a healthy risk culture?” to “at which level does risk culture appear to function, and where does it break down?”. In this sense, the two-level pattern can be read as a split between relationally robust risk voice in trusted dyads and structurally fragile risk voice once concerns move, or fail to move, into formal governance channels. Taken together, this framing is original in treating risk voice as a reframed construct, specifying escalation friction, feedback suppression and voice legitimacy as its measurable distortions, and defining risk entropy through their cross-unit configurations rather than as a generic cultural weakness.

The Sequential Mechanism from Friction to Inferred Entropy. Structural friction does not operate in isolation. Procedural barriers produce relational hesitancy, which triggers cognitive normalisation, which in turn generates suppression, which accumulates into conditions from which risk entropy is inferred. Prior literature theorises each element independently: Power's (2009) rituals of verification; the veneer concept (McConnell, 2013); discourse boundaries (Hardy and Maguire, 2016); and strategic ignorance (McGoey, 2012). The sequential chain presented here links these into a process account. The contribution is not any single construct but the empirically grounded specification of how they compound, providing a mechanism-level explanation of how governance architecture may produce the very silence it is designed to prevent. This pattern is not merely theoretical, as a contemporaneous independent assurance review of a court-enforceable bank remediation programme, found that Line 2 oversight would need to move beyond verifying activity completion to testing whether risk outcomes had materially improved in a sustainable manner. The same root cause analysis identified a 'good news culture' and a 'reluctance to challenge' as the persistent behavioural drivers of that failure (ANZ: *Promontory 1st Report*, 2026).

Unit-Level Entropy Signatures as a Configurational Diagnostic Tool. The identification of four analytically distinct configurations, Symbolic Governance, Control Without Voice, Silenced Executors, and Boundary Negotiators (Provisional), points toward different parts of the same organisation contributing to conditions consistent with risk entropy in qualitatively different ways. This extends the call for plural, subcultural approaches to risk culture (Ashby, Palermo and Power, 2012) and provides a configuration-based diagnostic framework, a typology of entropy signatures, that moves beyond

monolithic organisational profiles. Each configuration carries a distinct unit-level contribution to the overall decoupling pattern.

The Discourse-Practice Ratio as a Heuristic Diagnostic Marker. The 3.7:1 ratio operationalises simulated alignment (Ashby, Palermo and Power, 2012), translating an abstract diagnostic concept into a decomposable, documentary-record-based marker. The procedure is replicable in its logic, code interview and artefact data for contextual mechanism discourse on one side and proximate behavioural friction indicators on the other, then compare volumes, though direct numerical comparability across institutions requires standardisation of code taxonomy boundaries, the RE exclusion rule, segmentation units, data source scope, unit architecture, and FS sensitivity thresholds. The ratio's diagnostic value lies in its direction and internal distribution, not in the absolute number. Sheedy and Griffin (2018) identify the absence of decomposable, behaviourally grounded metrics as the central weakness in current assessment practice, instruments such as the Macquarie Risk Culture Scale produce useful organisational-level aggregates but cannot distinguish where mechanisms succeed from where they fail, situating the discourse-practice ratio directly within this gap.

The 3.7:1 discourse-practice ratio operationalises, in a risk governance context, the process that Edmondson (1999), drawing on Argyris and Schön, 1978, describes as learning through the detecting and correcting of error, and the interpersonal dynamics that impede it. The four unit-level entropy configurations observed in this study are consistent with what Edmondson and Moingeon (1998) describe, following Argyris, as O-I learning systems: self-reinforcing dynamics in which escalation barriers and challenge suppression prevent the double-loop learning, examining and altering the governing variables of the

governance architecture itself, that would be required to resolve them. Single-loop learning corrects surface-level errors; the EF/FS/VL indicator framework, including pre-articulation suppression where concerns do not form into articulable issues at all, provides the mechanism-level account of why organisations appear to remain structurally unable to reach the governing variables that produce those errors, consistent with the single-loop / double-loop distinction as reviewed by Tosey, Visser and Saunders (2012), who trace its origin to Argyris and Schön's early work on organisational learning.

On the status of Risk Entropy itself: within this study, RE functions as an interpretive synthesis construct, constituted through the post-saturation cross-unit configuration of EF, FS, and VL rather than through independent event-level coding. It names something real and theoretically distinctive, a cumulative, system-level governance disorder qualitatively different from any individual proximate barrier. Whether RE can be operationalised as a formal, independently measurable construct is a question this study deliberately leaves open, and which the findings position as a tractable research agenda for multi-site, mixed-methods investigation.

5.10 Practical Implications for Risk Governance

The findings carry targeted implications for practitioners, regulators, and governance designers, framed as enhancements to existing frameworks. The single most implementable action for a resource-constrained deposit-taker is a time-boxed documentary exercise, using existing ERCC minutes, board packs, and action logs to establish the discourse–practice baseline, followed by standing up structured process channels (digital or otherwise) across two or three live governance processes, which activates the entire measurement architecture simultaneously. The practical contribution

lies in proposing behavioural indicators and concern-thread integrity as concrete, traceable diagnostics for the discourse–practice gap, extending current reliance on single-level survey instruments and policy inventories.

Escalation pathway redesign. Escalation pathways should be audited for cumulative friction across all three barrier types rather than assessed on a single dimension. The FSB's (2014) communication and challenge pillar assumes that mechanisms, once established, function as designed; the present findings suggest mechanisms require ongoing behavioural audit. Practical interventions include simplified escalation templates with clear routing and expected timeframes; structured pre-formal intake channels that capture concerns before procedural gates; and transparency reports documenting concerns that stalled before formal review. This recommendation applies to any institution operating formal multi-step escalation governance where cumulative friction has not been separately audited.

Behavioural verification alongside structural verification. The Symbolic Governance configuration suggests that boards supplement structural verification with behavioural verification: mechanisms that provide unmediated access to operational experience, including direct board observation of governance forums, structured skip-level conversations, and third-party behavioural audit methodologies. This recommendation aligns with any institution operating a three-line governance model in which board-level assurance currently relies on mediated reporting.

Survey instrument reform. The Peakon divergence suggests the need for instruments that distinguish between Local Voice Safety (LVS) and Systemic Escalation

Efficacy (SEE). This distinction is empirically grounded in both Ghafoori et al.'s (2023) validated two-dimensional risk culture measure and Dürst and Kunz's (2025) assessment framework, both of which operationalise the LVS-SEE separation as a precondition for diagnostic validity. Example items operationalising this distinction include: *LVS-1: 'In my immediate team, if I raise a concern about a risk or process issue, my view is taken seriously regardless of my job level.'* And *SEE-1: 'If I raised a significant risk concern through the formal escalation process, I am confident it would reach the right decision-makers in a timely way.'* The High LVS / Low SEE combination is precisely the Silenced Executors configuration, and it is the cell that conventional single-level engagement surveys systematically misclassify as healthy. This recommendation is additive rather than substitutive, as existing surveys retain value, and the reform augments their diagnostic range. In very small institutions where relational proximity is the dominant governance mechanism, the LVS–SEE distinction may be less diagnostically actionable and should be adapted accordingly.

Unit-level cultural diagnostics. Organisation-wide cultural assessments may obscure unit-level variation that carries governance significance. Indicator profiling that maps EF, FS, VL, and RE inference across embedded governance units would enable the four configurations to serve as diagnostic reference points for institutions with comparable three lines architectures. At the level of individual concerns, a concern thread integrity score, assessing whether a raised issue can be traced from initial articulation through response, disposition, and evidenced outcome, offers a candidate operationalisation of the discourse–practice gap, though its validation lies beyond this single-case design.

Scrutiny-independent reinforcement design. Scrutiny-dependency is a predictable property of conformist social learning, requiring architectural rather than motivational remediation. Reinforcement mechanisms should be designed to operate between regulatory review cycles. Tracking whether challenge behaviours persist between regulatory engagement periods provides a leading indicator of cultural sustainability versus fragility. This recommendation transfers to any institution subject to periodic external regulatory review where challenge behaviours have not been separately tracked between cycles.

Continuous behavioural diagnostics. The Peakon divergence and the DNB 34/54 finding together point toward a diagnostic gap that periodic surveys cannot reliably close. The discourse-practice ratio and the unit-level entropy signatures represent early steps toward the continuous, process-level diagnostic capability that the regulatory direction of travel indicates. This recommendation is most directly applicable to institutions with existing digital governance data environments and should be treated as a directional aspiration rather than immediate prescription where such infrastructure is not yet in place.

5.11 Methodological Limitations

The embedded single-case design provides contextual depth at the cost of statistical generalisability, a trade-off made deliberately. The configurations, mechanisms, and ratios identified represent transferable analytical constructs whose applicability to other settings must be assessed through replication rather than assumed. The study's relativist ontology and subjectivist epistemology mean that findings represent co-constructed interpretations of participants situated sense-making rather than objective measurements of organisational attributes, and that silence or apparent acquiescence must be interpreted with explicit

barrier attribution checks, examining structural and contextual constraints before inferring motivational or character-based explanations. Methodologically, the study's main contribution is to demonstrate how a sampling-explicit, template-based, and barrier-attribution-sensitive case design can generate diagnostic indicators that are both theoretically meaningful and operationally traceable in risk governance settings.

Fieldwork was conducted over a six-week window (21 October to 12 November 2025), capturing a coherent organisational moment rather than a longitudinal trajectory. The scrutiny-dependency finding suggests cultural dynamics fluctuate with regulatory cycles, a hypothesis requiring longitudinal investigation. The case institution operates within the Channel Islands regulatory jurisdiction, a distinctive environment whose scale and labour-market dynamics may influence the manifestation of EF, FS, VL, and RE in ways that differ from larger jurisdictions.

The researcher occupied an insider-outsider position, external to the organisation but professionally familiar with financial services governance. Professional background likely sharpened sensitivity to the structural nature of the barriers identified, having previously observed 'box-ticking' exercises (Power, 1994) under the assumption that more boxes to tick were the answer. The notable feature of this study was finding the indicators present in an institution that appeared to be working well, with high internal survey scores, suggesting that FSB-compliant organisations may harbour these configurations in varying degrees, regardless of regulatory review status, a hypothesis this single case can raise but not confirm. This insight is itself a finding that deserves emphasis: problem organisations are not the only ones generating entropy. The indicators can be present even when the CRO

is engaged and the Board acknowledges communication as a development area, suggesting the problem is structural and systemic rather than symptomatic of individual failure.

Three scope limitations merit acknowledgement. First, the study's practical implications assume a stable macro-regulatory architecture; a material deregulatory turn has been documented across several major jurisdictions (Scott, 2025), and in such contexts the external incentive structures that give the recommendations their regulatory force may weaken substantially. Second, the study's frame treats escalation failure as a function of internal architectural factors, yet some proportion of the efficacy-calculus failure documented here may reflect exogenous societal trust dynamics (Fukuyama and Edelman, cited in Scott, 2025) that cannot be fully disentangled from endogenous governance design. Third, supervisory agencies themselves may be subject to comparable decoupling dynamics (Byres, 2024; Enria, 2024, cited in Scott, 2024), a relevant scope consideration for the practical implications set out in Section 5.10.

5.12 Directions for Future Research

If a follow-on study were possible, the most directly productive design would be a comparative multi-case study across three FSB-compliant institutions of different scale and regulatory maturity, or, with particular methodological advantage, across three distinct entities within the same group structure, to test whether the 3.7:1 ratio is institution-specific or whether a comparable range (± 1.5 either side) represents a normative parameter for this type of governance architecture. The within-group design would hold the parent governance architecture, regulatory regime, and cultural tone-from-the-top broadly constant, isolating unit-level and jurisdictional variation as the primary explanatory dimension; cross-institutional replication would remain the longer-term goal for claims of

broader generalisability. The specific hypothesis would be that institutions with a discourse-practice ratio above 3.5 will show measurably higher operational error rates, regulatory findings, or governance deferrals than those below 2.5, controlling for institutional size and regulatory regime. Establishing this threshold relationship would convert the ratio from a heuristic to a benchmarkable diagnostic instrument. Within such a design, a concern-thread-derived integrity score could be tested as a complementary diagnostic, examining whether institutions with systematically incomplete concern threads exhibit higher error rates or regulatory findings than those where raised concerns can be traced to evidenced outcomes.

Multi-case replication across varying jurisdictions, scales, and regulatory regimes is the immediate priority. This would enable cross-case comparison to assess whether Symbolic Governance, Control Without Voice, Silenced Executors, and Boundary Negotiators recur in other settings, substantially strengthening the argument for analytic generalisability.

Longitudinal tracking would enable investigation of the normative drift hypothesis, whether the configurations identified here represent stable states or drift trajectories, and whether scrutiny-dependent reinforcement is phase-specific or reflects the stable equilibrium suggested by the present data.

Survey instrument development. The dual-level voice measurement instrument proposed in Section 5.10, which distinguishes Local Voice Safety from Systemic Escalation Efficacy, is a directly tractable product of the present study's diagnostic framework and requires validation across multiple institutions and regulatory

environments. Dürst and Kunz's (2025) assessment protocols and Ghafoori et al.'s (2023) validated measure provide the methodological scaffolding for this development work.

Ratio quantification and benchmarking could enable the discourse-practice ratio to function as a comparative diagnostic tool across institutions, requiring standardised coding protocols for committee minutes, board packs, and escalation records. Sheedy and Griffin's (2018) identification of the senior leader perception gap as an unresolved limitation of current instruments provides the theoretical grounding for what cross-institutional ratio variation would need to demonstrate to influence governance practice.

5.12.1 Speculative and Longer-Horizon Directions

The following directions extend beyond what the current dataset can support and are offered as research conjectures rather than findings-derived propositions. They are included because the indicator framework, if validated through replication, may ultimately reach these applications.

AI-assisted and continuous behavioural diagnostics represent the emerging frontier: EF, FS, and VL constitute the indicator categories that a continuous, process-level diagnostic system would need to detect, and the discourse-practice ratio provides the theoretical input to such a system's design. As AI-mediated communication becomes embedded in governance processes, a further question arises as to whether systems trained on institutional communication patterns may inadvertently encode the same friction and suppression structures this study identifies, producing, at scale, the polishing and filtering mechanisms currently enacted through human gatekeeping.

Extension to external parties, fraud, and disinformation. A more speculative but potentially high-impact direction concerns the application of the indicator framework beyond organisational boundaries. EF, FS, and VL dynamics may operate analogously across contractual and service relationships: suppliers who learn that escalated concerns are absorbed without response may develop the same acquiescent silence observed in the Operations configuration, while customers with low perceived Voice Legitimacy tend to exit or self-censor rather than challenge. The same logic extends to AI-enabled communication, where the capacity to synthetically replicate high-legitimacy voice, mimicking the relational and cognitive conditions that lower escalation barriers, means that anomalous VL or EF patterns in real-time communication streams may become detectable pre-event signals for conduct risk and fraud. Critically, this mechanism is not limited to deliberate deception: misinformation and disinformation persist within governance structures for the same architectural reason, that the conditions under which legitimate challenge to false or misleading information does not occur, or does not escalate, are precisely those captured by elevated EF and suppressed VL, including the two-tier VL environments suggested by this case, in which senior voices retain legitimacy while junior voices do not. The framework thus points toward a diagnostic application considerably broader than its original governance context, and one that aligns with the WEF's (2025) identification of misinformation and disinformation, cyber espionage, and third-party risk management as among the leading systemic threats of the current period. Future research could test this two-tier VL hypothesis explicitly across institutions, using stratified measures of voice safety and escalation efficacy by hierarchy level rather than aggregate scores alone. These directions are generated by, and extend beyond, this single case by turning the heuristic devices of this study—the discourse–practice ratio, concern-thread

integrity, and two-tier voice legitimacy—into explicit hypotheses and candidate design specifications for future comparative and quantitative work.

5.13 Chapter Summary

This chapter set out to interpret the findings of Chapter 4 in dialogue with the theoretical frameworks of Chapter 2, and to advance a response to the overarching research question: How do communication and challenge behaviours function as early indicators of risk entropy within financial institutions, and what mechanisms either suppress or enable them?

The discussion revealed a consistent two-level pattern. Communication and challenge mechanisms appear to function genuinely within the relational proximity of trusted local relationships, a finding corroborated by 31 positive co-occurrence passages and a Peakon Freedom of Opinion score of 8.9. Yet the same mechanisms appear unable to prevent the sequential chain from friction through suppression to conditions from which risk entropy is inferred as an emergent condition once voice attempts to scale into formal governance. The reinforcement paradox, 408 coded references to challenge norms coexisting with the configurations from which entropy is inferred, points toward governance discourse about challenge potentially substituting for the challenge itself.

The four unit-level entropy signatures aggregate into organisational decoupling quantified by the 3.7:1 discourse-practice ratio, distributed as Operations gap 312, Oversight gap 230, Risk and Compliance gap 176, and Business gap 108, a distribution consistent with the proximity-influence inversion. Artefact evidence, 47 of 192 open actions outstanding, TPRM red across three consecutive ERCC packs, and a 22% IoM-

Jersey error rate differential, provides documentary corroboration that the behavioural gap carries measurable performance consequences.

Four theoretical contributions were advanced and six practical implications identified, each with an explicit transferability statement. The overarching conclusion, that this appears to be a decoupling problem rather than a design problem in this case, directs attention toward the behavioural gap rather than the mechanism inventory as the appropriate level of diagnostic and remedial intervention. The future research directions extend this logic progressively: from multi-case replication and ratio benchmarking, through AI-assisted continuous diagnostics, to the more speculative application of EF, FS, and VL beyond organisational boundaries, to suppliers, third-party providers, customers, and the conditions under which AI-enabled communication and disinformation go unchallenged. The framework, developed to diagnose risk entropy within a single institution, points toward a diagnostic application that may extend considerably further.

REFERENCES

- Adeoye-Olatunde, O.A. and Olenik, N.L. (2021) 'Research and scholarly methods: Semi-structured interviews', *Journal of the american college of clinical pharmacy*, 4(10), pp. 1358–1367.
- Adler, P.S., Forbes, L.C. and Willmott, H. (2007) '3 Critical Management Studies', *Academy of Management Annals*, 1(1), pp. 119–179. Available at: <https://doi.org/10.5465/078559808>.
- Aldag, R.J. and Fuller, S.R. (1993) 'Beyond fiasco: A reappraisal of the groupthink phenomenon and a new model of group decision processes.', *Psychological bulletin*, 113(3), p. 533.
- Anton, S.G. and Nucu, A.E.A. (2020) 'Enterprise risk management: A literature review and agenda for future research', *Journal of Risk and Financial Management*, 13(11), p. 281.
- ANZ Non-Financial Risk management: Promontory First Report (2026). Available at: <https://www.anz.com.au/newsroom/media/2026/may/anz-non-financial-risk-management-promontorys-first-report/> (Accessed: 5 May 2026).
- APRA (2016) *Risk Culture*. Information Paper. Australian Prudential Regulation Authority.
- APRA (2017) *Prudential Standard CPS 220*. Australian Prudential Regulation Authority.
- APRA (2019) *Transforming governance, culture, remuneration and accountability - APRA's approach*. Australian Prudential Regulation Authority.
- APRA (2022) *No room for complacency on bank risk culture*. Australian Prudential Regulation Authority. Available at: <https://www.apra.gov.au/news-and-publications/no-room-for-complacency-on-bank-risk-culture>.
- de Araújo Lima, P.F., Crema, M. and Verbano, C. (2020) 'Risk management in SMEs: A systematic literature review and future directions', *European management journal*, 38(1), pp. 78–94.
- Ashby, S., Palermo, T. and Power, M. (2012) 'Risk culture in financial organisations: An interim report'.
- Baskarada, S. (2014) 'Qualitative case study guidelines', *Başkarada, S. (2014). Qualitative case studies guidelines. The Qualitative Report*, 19(40), pp. 1–25.

Baxter, P. and Jack, S. (2008) 'Qualitative case study methodology: Study design and implementation for novice researchers', *The qualitative report*, 13(4), pp. 544–559.

Behaviour and Culture Report | Central Bank of Ireland (2018) *Central Bank of Ireland - English*. Available at: <https://www.centralbank.ie/publication/behaviour-and-culture-report> (Accessed: 14 April 2026).

BIS (2024) *Core Principles for effective banking supervision BCP40*. Bank of International Settlements.

Bozeman, B. and Kingsley, G. (1998) 'Risk culture in public and private organizations', *Public administration review*, pp. 109–118.

Braumann, E.C., Hiebl, M.R. and Posch, A. (2024) 'Enterprise risk management as part of the organizational control package: Review and implications for management accounting research', *Journal of Management Accounting Research*, 36(2), pp. 7–29.

Braun, V. *et al.* (2019a) 'Thematic analysis', *Handbook of research methods in health social sciences*. Springer, pp. 843–860.

Braun, V. *et al.* (2019b) 'Thematic analysis 48', *Handbook of research methods in health social sciences*, pp. 843–860.

Brinsfield, C.T. (2013) 'Employee silence motives: Investigation of dimensionality and development of measures', *Journal of organizational behavior*, 34(5), pp. 671–697.

Bromiley, P. *et al.* (2015) 'Enterprise risk management: Review, critique, and research directions', *Long range planning*, 48(4), pp. 265–276.

Byrne, D. (2022) 'A worked example of Braun and Clarke's approach to reflexive thematic analysis', *Quality & quantity*, 56(3), pp. 1391–1412.

Campbell, S. *et al.* (2020) 'Purposive sampling: complex or simple? Research case examples', *Journal of research in Nursing*, 25(8), pp. 652–661.

De Dreu, C.K.W., Nijstad, B.A. and Van Knippenberg, D. (2008) 'Motivated Information Processing in Group Judgment and Decision Making', *Personality and Social Psychology Review*, 12(1), pp. 22–49. Available at: <https://doi.org/10.1177/1088868307304092>.

DNB (2015) *Supervision of Behaviour and Culture*. De Nederlandsche Bank.

DNB (2023) *Moving from reflex to reflection*. De Nederlandsche Bank.

Dürst, N. and Kunz, J. (2025) 'How to conduct effective risk culture assessments', *Journal of Management Control*, pp. 1–46.

EBA (2021) *Final report on Guidelines on internal governance under CRD*. European Banking Authority.

ECB (2024) *Draft guide on governance and risk culture*. European Central Bank.

Edmondson, A. (1999) 'Psychological Safety and Learning Behavior in Work Teams', *Administrative Science Quarterly*, 44(2), pp. 350–383. Available at: <https://doi.org/10.2307/2666999>.

Edmondson, A. and Moingeon, B. (1998) 'From organizational learning to the learning organization', *Management learning*, 29(1), pp. 5–20.

Eisenhardt, K.M. (1989a) 'Agency theory: An assessment and review', *Academy of management review*, 14(1), pp. 57–74.

Eisenhardt, K.M. (1989b) 'Building theories from case study research', *Academy of management review*, 14(4), pp. 532–550.

Etikan, I., Musa, S.A. and Alkassim, R.S. (2016) 'Comparison of convenience sampling and purposive sampling', *American journal of theoretical and applied statistics*, 5(1), pp. 1–4.

FCA (2018) *Transforming Culture in Financial Services*. Financial Conduct Authority.

Federico, D., Napoli, F. and Notte, A. (2024) 'An Overview on Silicon Valley Bank Failure in Europe: Causes and Effects', *Contemporary Research in Business, Management and Economics*, 5, pp. 63–79.

FINMA (2023) *Lessons Learned from the CS Crisis*. FINMA Report. Bern: FINMA. Available at: <https://www.finma.ch/en/documentation/finma-publications/reports/bericht-zur-credit-suisse-krise/>.

FSB (2014) *Guidance on Supervisory Interaction with Financial Institutions on Risk Culture*. Financial Stability Board (FSB).

FSB (2018) *Strengthening Governance Frameworks to Mitigate Misconduct Risk: A Toolkit for Firms and Supervisors*. Financial Stability Board (FSB).

FSB (2023) *2023 Bank Failures*. Financial Stability Board (FSB). Available at: <https://www.fsb.org/2023/10/2023-bank-failures-preliminary-lessons-learned-for-resolution/>.

FSB (2024) *Peer Review of Switzerland*. Review report. Financial Stability Board (FSB). Available at: <https://www.fsb.org/2024/02/peer-review-of-switzerland/>.

Fukuyama, F. (1995) 'Social capital and the global economy', *Foreign Aff.*, 74, p. 89.

Ghafoori, E. *et al.* (2023) 'Measuring risk culture in finance: Development of a comprehensive measure', *Journal of Banking & Finance*, 148, p. 106720.

Global Risks Report 2025: Conflict, Environment and Disinformation Top Threats (2025) *World Economic Forum*. Available at: <https://www.weforum.org/press/2025/01/global-risks-report-2025-conflict-environment-and-disinformation-top-threats/> (Accessed: 20 April 2026).

Gorton, G.B., Grennan, J. and Zentefis, A.K. (2022) 'Corporate culture', *Annual Review of Financial Economics*, 14(1), pp. 535–561.

Grieser, F. and Pedell, B. (2022) 'Exploring risk culture controls: to what extent can the development of organizational risk culture be controlled and how?', *Journal of Accounting & Organizational Change*, 18(5), pp. 752–788.

Guba, E.G. and Lincoln, Y.S. (1994) 'Competing paradigms in qualitative research', *Handbook of qualitative research*, 2(163–194), p. 105.

Hald, E.J., Gillespie, A. and Reader, T.W. (2021) 'Causal and corrective organisational culture: A systematic review of case studies of institutional failure', *Journal of Business Ethics*, 174, pp. 457–483.

Hardy, C. *et al.* (2020) 'Organizing risk: Organization and management theory for the risk society', *Academy of management annals*, 14(2), pp. 1032–1066.

Hardy, C. and Maguire, S. (2010) 'Discourse, field-configuring events, and change in organizations and institutional fields: Narratives of DDT and the Stockholm Convention', *Academy of Management journal*, 53(6), pp. 1365–1392.

Hardy, C. and Maguire, S. (2016) 'Organizing risk: Discourse, power, and "riskification"', *Academy of management review*, 41(1), pp. 80–108.

Hennink, M. and Kaiser, B.N. (2022) 'Sample sizes for saturation in qualitative research: A systematic review of empirical tests', *Social science & medicine*, 292, p. 114523.

Horvey, S.S. and Odei-Mensah, J. (2023) 'The measurements and performance of enterprise risk management: a comprehensive literature review', *Journal of Risk Research*, 26(7), pp. 778–800.

Jetten, J. and Hornsey, M.J. (2014) 'Deviance and Dissent in Groups', *Annual Review of Psychology*, 65(1), pp. 461–485. Available at: <https://doi.org/10.1146/annurev-psych-010213-115151>.

- Kallio, H. *et al.* (2016) 'Systematic methodological review: developing a framework for a qualitative semi-structured interview guide', *Journal of advanced nursing*, 72(12), pp. 2954–2965.
- Kaplan, R.S. and Mikes, A. (2012) 'Smart companies match their approach to the nature of the threats they face', *Harvard Business Review*, June [Preprint].
- Klinke, A. and Renn, O. (2021) 'The coming of age of risk governance', *Risk analysis*, 41(3), pp. 544–557.
- Kristensen, G.K. and Ravn, M.N. (2015) 'The voices heard and the voices silenced: Recruitment processes in qualitative interview studies', *Qualitative Research*, 15(6), pp. 722–737.
- Kunz, J. and Heitz, M. (2021) 'Banks' risk culture and management control systems: A systematic literature review', *Journal of Management Control*, 32(4), pp. 439–493.
- Li, K. *et al.* (2021) 'Measuring corporate culture using machine learning', *The Review of Financial Studies*, 34(7), pp. 3265–3315.
- Liu, L.X., Liu, S. and Sathye, M. (2021) 'Predicting bank failures: a synthesis of literature and directions for future research', *Journal of Risk and Financial Management*, 14(10), p. 474.
- Lockhart, P. (2009) *A mathematician's lament: How school cheats us out of our most fascinating and imaginative art form*. Bellevue literary press.
- Magee, S., Schilling, C. and Sheedy, E. (2019) 'Risk governance in the insurance sector—determinants and consequences in an international sample', *Journal of Risk and Insurance*, 86(2), pp. 381–413.
- MAS (2020) *Culture Conduct Practices of Financial Institutions*. Monetary Authority of Singapore.
- Mayer, R.C., Davis, J.H. and Schoorman, F.D. (1995) 'An integrative model of organizational trust', *Academy of management review*, 20(3), pp. 709–734.
- McConnell, P. (2012) 'STRATEGIC RISK MANAGEMENT-A TALE OF TWO STRATEGIES', *Journal of Risk and Governance*, 3(2), p. 83.
- McConnell, P.J. (2013) 'A risk culture framework for systemically important banks', *Journal of risk and governance*, 3(1).
- McGoey, L. (2012) 'The logic of strategic ignorance', *The British Journal of Sociology*, 63(3), pp. 533–576. Available at: <https://doi.org/10.1111/j.1468-4446.2012.01424.x>.

McGoey, L. (2016) 'Strategic unknowns: Towards a sociology of ignorance', *An introduction to the sociology of ignorance*. Routledge, pp. 1–16. Available at: <https://api.taylorfrancis.com/content/chapters/edit/download?identifierName=doi&identifierValue=10.4324/9781315771106-1&type=chapterpdf> (Accessed: 9 July 2025).

Mikes, A. (2009) 'Risk management and calculative cultures', *Management accounting research*, 20(1), pp. 18–40.

Mikes, A. (2011) 'From counting risk to making risk count: Boundary-work in risk management', *Accounting, organizations and society*, 36(4–5), pp. 226–245.

Morrison, E.W. (2014) 'Employee voice and silence', *Annu. Rev. Organ. Psychol. Organ. Behav.*, 1(1), pp. 173–197.

Morrison, E.W. (2023) 'Employee voice and silence: Taking stock a decade later', *Annual review of organizational psychology and organizational behavior*, 10(1), pp. 79–107.

Morrison, E.W. and Milliken, F.J. (2000) 'Organizational silence: A barrier to change and development in a pluralistic world', *Academy of Management review*, 25(4), pp. 706–725.

OSFI (2024) *Culture and Behaviour Risk Guideline*. OSFI: Office of the Superintendent of Financial Institutions.

Osman, A. and Lew, C.C. (2021) 'Developing a framework of institutional risk culture for strategic decision-making', *Journal of Risk Research*, 24(9), pp. 1072–1085.

Palermo, T., Power, M. and Ashby, S. (2017) 'Navigating institutional complexity: The production of risk culture in the financial sector', *Journal of Management Studies*, 54(2), pp. 154–181.

Pan, Y., Siegel, S. and Wang, T.Y. (2017) 'Corporate risk culture', *Journal of Financial and Quantitative Analysis*, 52(6), pp. 2327–2367.

Pham, L.T.M. (2018) 'Qualitative approach to research a review of advantages and disadvantages of three paradigms: Positivism, interpretivism and critical inquiry', *University of Adelaide*, 6, pp. 34–47.

Pharoah, M. (2025) 'Mapping Risk Culture Across Organisational Levels — Open questions for Risk Management', 17 August. Available at: <https://www.linkedin.com/pulse/weekend-reading-mapping-risk-culture-across-organisational-avazc/?trackingId=dHhB66ZzQWSSM3UULCddCQ%3D%3D>.

Power, M. (1994) *The audit explosion*. Demos.

- Power, M. (2004) 'The risk management of everything', *The Journal of Risk Finance*, 5(3), pp. 58–65.
- Power, M. (2009) 'The risk management of nothing', *Accounting, organizations and society*, 34(6–7), pp. 849–855.
- Power, M., Ashby, S. and Palermo, T. (2013) *Risk culture in financial organisations: A research report*. CARR-Analysis of Risk and Regulation.
- Power, M.K. and Gendron, Y. (2015) 'Qualitative research in auditing: A methodological roadmap', *Auditing: A Journal of Practice & Theory*, 34(2), pp. 147–165.
- PRA (2014) *The use of PRA powers to address serious failings in the culture of firms*. Prudential Regulatory Authority.
- PRA (2021) *Non-systemic UK banks: The Prudential Regulation Authority's approach to new and growing banks*. Prudential Regulation Authority.
- Renn, O., Klinke, A. and Van Asselt, M. (2011a) 'Coping with complexity, uncertainty and ambiguity in risk governance: a synthesis', *Ambio*, 40(2), pp. 231–246.
- Renn, O., Klinke, A. and Van Asselt, M. (2011b) 'Coping with complexity, uncertainty and ambiguity in risk governance: a synthesis', *Ambio*, 40, pp. 231–246.
- Ring, P.J. *et al.* (2016) 'Taking notice of risk culture—the regulator's approach', *Journal of Risk Research*, 19(3), pp. 364–387.
- Rousseau, D.M. *et al.* (1998) 'Not so different after all: A cross-discipline view of trust', *Academy of management review*, 23(3), pp. 393–404.
- Ruslin, R. *et al.* (2022) 'Semi-structured Interview: A methodological reflection on the development of a qualitative research instrument in educational studies', *IOSR Journal of Research & Method in Education (IOSR-JRME)*, 12(1), pp. 22–29.
- Sandberg, J. (2005) 'How do we justify knowledge produced within interpretive approaches?', *Organizational research methods*, 8(1), pp. 41–68.
- Sandberg, J. and Tsoukas, H. (2015) 'Making sense of the sensemaking perspective: Its constituents, limitations, and opportunities for further development', *Journal of organizational behavior*, 36(S1), pp. S6–S32.
- Schein, E.H. (1984) 'Coming to a New Awareness of Organizational Culture', *Sloan Management Review*, 25.

Scott, S. (2024) *Culture & Conduct Risk in the Banking Sector: Why it matters and what the industry is doing to address it*. 7th Annual Compendium. Starling Trust Sciences, LLC. Available at: insights.starlingtrust.com.

Scott, S. (2025) *Culture & Conduct Risk in the Financial Sector: Why it matters and what the industry is doing to address it*. 8th Annual Compendium. Starling Trust Sciences, LLC.

Selznick, P. (1996) 'Institutionalism "Old" and "New"', *Administrative Science Quarterly*, 41(2), p. 270. Available at: <https://doi.org/10.2307/2393719>.

Sheedy, E. and Canestrari-Soh, D.S. (2023) 'Does executive accountability enhance risk management and risk culture?', *Accounting & Finance*, 63(4), pp. 4093–4124.

Sheedy, E. and Griffin, B. (2018) 'Risk governance, structures, culture, and behavior: A view from the inside', *Corporate Governance: An International Review*, 26(1), pp. 4–22.

Sheedy, E.A., Griffin, B. and Barbour, J.P. (2017) 'A framework and measure for examining risk climate in financial institutions', *Journal of Business and Psychology*, 32, pp. 101–116.

SS21/15 – *Internal Governance* (2017). Available at: <https://www.prarulebook.co.uk/guidance/supervisory-statements/ss21-15-internal-governance> (Accessed: 14 April 2026).

Stulz, R.M. (2016) 'Risk management, governance, culture, and risk taking in banks', *Economic Policy Review*, Issue Aug, pp. 43–60.

Tobin, G.A. and Begley, C.M. (2004) 'Methodological rigour within a qualitative framework', *Journal of advanced nursing*, 48(4), pp. 388–396.

Tosey, P., Visser, M. and Saunders, M.N. (2012) 'The origins and conceptualizations of 'triple-loop' learning: A critical review', *Management learning*, 43(3), pp. 291–307.

Vranken, I. *et al.* (2015) 'A review on the use of entropy in landscape ecology: heterogeneity, unpredictability, scale dependence and their links with thermodynamics', *Landscape ecology*, 30(1), pp. 51–65.

Wasko, M.M. and Faraj, S. (2005) 'Why should I share? Examining social capital and knowledge contribution in electronic networks of practice', *MIS quarterly*, pp. 35–57.

Weick, K.E. (1995) 'Sensemaking in organizations'.

Wiedemann, A. *et al.* (2024) 'The risk-related tone from the top: evidence from German regional banks', *Journal of Risk Research*, 27(7), pp. 713–734.

Yin, R.K. (2009) *Case study research: Design and methods*. sage.

APPENDIX A: INTERVIEW PROTOCOL AND GUIDE

This appendix contains the full interview materials referenced in 3.5.2–3.5.3, including the finalised interview guide, research-question mapping, development rationale, and participant consent documentation.

A.1 Interview Guide (v3.5)

Instructions

- Target duration: 45–60 minutes
- Prioritise depth; flag core questions as Must-ask
- Note artefact/document references verbatim
- Italicise truly optional probes

1 Introduction and Consent (5 min)

Q1.1 Welcome participant; explain purpose:

“We’re exploring how open dialogue and challenge shape risk decisions here, at [Institution].”

Q1.2 Confirm consent and permission to record.

Q1.3 Reiterate confidentiality, anonymisation, and how insights will feed into a practitioner-ready toolkit.

Transition: Let’s gather some background.

2 Background and Role (5 min)

Q2.1 Please state your role/business unit, seniority level, and core risk-management responsibilities. (Must-ask)

Q2.2 Who do you collaborate with when making or challenging risk decisions?

Q2.3 How does your reporting line influence your comfort in raising concerns?

Transition: We’ll start with how alternate views are solicited and valued.

3 Encouraging Alternate Views (FSB 3.3.1) (10 min)

Q3.1 Must-ask: How are differing opinions or questions actively invited in meetings or reviews?

Probe: Can you recall specific language or behaviours that signalled “your view matters”?

Q3.2 Must-ask: Describe a time when someone’s challenge changed a risk decision or process.

Probe: How did that outcome affect your willingness to speak up again?

Probe: “Was that challenge publicly acknowledged or rewarded?”

prompt: “Was that challenge documented in any meeting minutes, follow-up email, or decision log?”

Consider a rapid-fire backup: “What’s the most visible sign that objection/dissent is respected?”

Q3.3 What happens when someone’s view is unwelcome? How is that managed?

Probe: “Have you ever held back feedback because it felt futile or risky?”
Consider a rapid-fire backup: “What’s the clearest cue that feedback isn’t welcome?”
Transition: Let’s look at formal channels and how openness is measured.

4 Mechanisms for Expression and Assessment (FSB 3.3.2) (10 min)

Q4.1 Must-ask: What formal channels exist for raising risk concerns (e.g., hotlines, skip-levels, forums)? Which do you use most?

Document/Artefact prompt: “Do you have an example email, survey, or policy capturing that channel?” “can you provide the name/link?”

Probe: “Have you ever hesitated to use a formal channel due to perceived gatekeeping or delay?”

Consider a rapid-fire backup: “What’s the longest it’s taken for a concern to reach the right level?”

Q4.2 Must-ask: How does the institution measure willingness to speak up? Which surveys or assessments run, how often, and who reviews results?

Probe: When openness scores dip, what follow-up actions occur?

Document/Artefact Prompt: “Can you recall or share a survey result, dashboard, or internal memo showing openness scores or related actions?”

Q4.3 How transparent are assessment results back to staff and the board?

Transition: Next, the stature and influence of control functions.

5 Stature and Independence of Control Functions (FSB 3.3.3–3.3.5) (10 min)

Q5.1 Must-ask: How are risk management, internal audit, and compliance represented in governance committees? Do they have equal voting or veto rights alongside business-line leaders?

Probe: Can you share a committee agenda or minutes illustrating that role?

Document/Artefact Prompt: “Can you share a committee agenda or minutes illustrating that role?”

Q5.2 Must-ask: Describe an instance where a control function overruled or halted a business proposal. What enabled that outcome?

Probe: Did they ever bypass management to report directly to the board? Was this documented?

Document/Artefact prompt: “Was that escalation documented in any board report, email, or governance log?” → Strengthens evidence for Escalation Friction and control-function independence.

Q5.3 How is independence of control functions maintained (e.g., reporting lines, board access)?

Transition: Let’s explore how these elements work together in practice.

6 Interplay and Cultural Reinforcement (5 min)

Q6.1 Must-ask: How do open challenge and strong control-function participation work together in your organisation?

In what ways do open challenge and strong control-function participation reinforce each other in your organisation?

Probe: Recall a moment when both pillars combined to improve a decision.

Q6.2 Optional: Where do you see tensions or gaps between encouraging challenge and control-function authority?

Transition: Finally, we'll cover learning, monitoring, and your final thoughts.

7 Monitoring, Learning and Wrap-Up (5 min)

Q7.1 Must-ask: How does the organisation track outcomes of risk challenges and decisions (e.g., dashboards, post-mortems)?

Probe: Have you observed sustained improvements or worrying declines over time?

Probe: "Have you noticed any patterns in how the challenge fades or intensifies over time?"

Document/Artefact prompt: "Do you have access to any dashboard, post-mortem report, or trend chart showing how challenge outcomes are tracked?" → Directly supports Indicator 5.3.4's entropy synthesis.

Consider a rapid-fire backup: "Which metric tells you the challenge is working, or failing?"

Q7.2 Is there any aspect of communication and challenge we haven't covered but you believe is critical?

Q7.3 If you could implement one change to strengthen open dialogue or control-function stature, what would it be?

Probe for a concrete example

Q7.4 Explain the member-checking timeline and thank the participant.

Rapid-Fire Backup Probes

Use when time is short:

- Alternate Views: "What's the single clearest sign your opinion is genuinely valued?"
- Mechanisms: "Which channel do you use first to raise a concern?"
- Control Functions: "Who is your go-to control lead when you question a decision?"
- Interplay: "What's the biggest friction between business lines and controls?"
- Monitoring: "Which metric tells you the challenge is working?"

Interviewer Cheat-Sheet and Tips

- Must-ask questions are core; cover them fully.
- Optional probes (italicised) only if time and rapport allow.
- Steering prompt: "That's interesting – how did that affect your willingness to speak up next time?"
- Ethical reminder: Respect boundaries when probing sensitive power dynamics.

Question-to-Code Mapping (for NVivo)

Q#	FSB Sub-Pillar
Q3.1–3.3	3.3.1 Open to Alternate Views
Q4.1–4.3	3.3.2 Mechanisms and Assessment
Q5.1–5.3	3.3.3–3.3.5 Stature of Control Functions
Q6.1–6.2	Interplay (Cross-Pillar)
Q7.1–7.4	Monitoring and Reflection (Culture Reinforcement)

Indicator Overview

Indicator	Description	Theoretical Lens	Evidence Type
Escalation Friction	Gatekeeping, delay, relational hesitancy	Power (rituals of verification), Gorton et al. (coordination architecture)	Quotes + email thread
Feedback Suppression	Self-censorship, discouragement cues, threshold effects	De Dreu et al. (psychological safety)	NVivo excerpts + survey item
Voice Legitimacy	Reinforcement vs. ridicule, follow-through disparities	Jetten and Hornsey (identity-safety), social exchange theory	Kudos email + meeting minutes
Risk Entropy Synthesis	Emergent. Cultural vigilance decay, divergence, trend patterns	Entropy metaphor	Time-series chart + coded patterns

Note: Introductory, consent, and closing questions (1.1-1, 7.4) support ethics and rapport; background questions (2.1-2.3) provide contextual classification for embedded-unit and role-level analysis; rapid-fire probes deepen core themes where time allows. They are not treated as separate research questions but as instruments that enrich responses to RQ1–RQ4 and strengthen EF, FS, VL, and RE indicator development.

A.2 Research Question-to-Indicator Mapping

This table maps each research question to the interview Sections, example probes, and heuristic indicators it primarily supports. The four research questions in Section 1.3 are operationalised here as interview-level themes written in practitioner-friendly language. Each A.2 question corresponds to one formal research question and is linked to specific indicators (EF, FS, VL, RE) and mechanisms (voice, control-function participation, monitoring/learning).

RQ	Primary focus	Main interview Sections	Example questions/probes	Key indicators supported
RQ1	How are alternate views and dissenting opinions actively invited, expressed, and responded to in routine risk discussions and formal meetings?	Encouraging Alternate Views (FSB 3.3.1); Background (role context)	Q3.1–3.3: how views are invited, examples of challenge changing outcomes, what happens when views are unwelcome; Q2.2–2.3: collaboration partners and	EF (relational hesitancy); FS (self-censorship); VL (who feels legitimate); RE: contributes data to later cross-unit synthesis; not coded directly at event level.

			reporting line influence	
RQ2	How do formal mechanisms for expressing concerns operate in practice, and how is willingness to speak up assessed and fed back to staff and the board?	Mechanisms for Expression/Assessment (FSB 3.3.2)	Q4.1–4.3: available channels, hesitancy to use them, openness surveys, transparency of results	EF (gatekeeping/delay); FS ("not worth it"); VL (mechanism credibility); RE: contributes data to later cross-unit synthesis; not coded directly at event level.
RQ3	How is the stature, independence, and influence of risk, compliance, and internal audit experienced in risk decision-making, including their ability to halt or override business proposals?	Stature and Independence of Control Functions (FSB 3.3.3–3.3.5) Background and Role (context for control-function reporting lines and governance committee representation)	Q5.1–5.3 (representation and voting rights in committees; examples of overrides or halts; bypassing management to report directly to the board; maintenance of independence) Document/Artefact prompts: committee agendas, minutes, board escalation logs	VL (primary): legitimacy of control-function voice relative to business lines EF: friction experienced when control functions attempt to escalate or override FS: subtle suppression or marginalisation of control-function challenge RE: pockets where control-function challenge is ritualised, bypassed, or absorbed, contributes to later cross-unit RE synthesis inference; not coded directly at event level for this RQ
RQ4	How do open challenge, control-function participation, and monitoring/learning	Interplay/Cultural Reinforcement ;	Q6.1–6.2: challenge/control tensions; Q7.1–7.3: tracking	RE (drift/vigilance decay); EF/FS/VL (evolution over time)*

	mechanisms interact over time to reinforce or erode cultural vigilance around risk?	Monitoring/Learning	outcomes, perceived changes	
--	---	---------------------	-----------------------------	--

** RE is not coded at event level in any RQ column; it is inferred through cross-unit explanation-building after EF, FS, and VL profiles are complete*

A.2.1 Research Question-to-Indicator and Interview Theme Mapping

This appendix links the formal research questions stated in Section 1.3 to the interview-level themes and heuristic indicators that guided data collection. Each interview theme is a practitioner-friendly translation of one formal research question, designed to keep questions accessible and non-leading while preserving the underlying conceptual focus. The table below shows, for each formal research question, the corresponding interview theme wording used to structure the semi-structured interviews (Appendix A.1) and the primary indicators and mechanisms (EF, FS, VL, RE and related constructs) that were operationalised in the coding and analysis.

Formal RQ (from 1.3)	Interview theme wording (A.2 question)	Primary indicators and mechanisms operationalised
RQ1: How are alternative views about risk encouraged and expressed within the case institution, and under what conditions does risk voice appear legitimate and safe at the local level?	Theme 1 – “How are alternate views and dissenting opinions actively invited, expressed, and responded to in routine risk discussions and formal meetings?”	Risk voice; Local Voice Safety (LVS); Voice Legitimacy (VL); openness to alternate views; relational trust; psychological safety; informal versus formal challenge behaviours.
RQ2: How do formal risk governance mechanisms for communication and challenge operate in practice, and where do escalation barriers and feedback suppression arise as concerns move through formal channels?	Theme 2 – “How do formal mechanisms for expressing concerns operate in practice, and how is willingness to speak up assessed and fed back to staff and the board?”	Escalation Friction (EF) (procedural, relational, cognitive); Feedback Suppression (FS); formal escalation routes and channels; treatment of speak-up metrics and survey results; feedback loops to staff, control functions, and the board.
RQ3: How can escalation friction, feedback suppression, voice legitimacy, and risk entropy be operationalised as diagnostic indicators to support boards and supervisors in assessing risk culture in structurally compliant institutions?	Theme 3 – “How is the stature, independence, and influence of risk, compliance, and internal audit experienced in risk decision-making, including their ability to halt or override business proposals?”	EF, FS, VL, and Risk Entropy (RE) as experienced in control-function roles; structural position and independence of risk, compliance, and internal audit; ability to challenge, halt, or reshape proposals; Control-Function Voice configurations; early formulation of diagnostic indicators (e.g., LVS/SEE distinctions).

RQ4: How do the observed unit-level entropy configurations and indicator patterns aggregate into a systemic governance pattern, and what does this reveal about organisational decoupling and the conditions under which risk entropy emerges?	Theme 4 – “How do open challenge, control-function participation, and monitoring/learning mechanisms interact over time to reinforce or erode cultural vigilance around risk?”	Unit-level entropy configurations (Symbolic Governance, Control Without Voice, Silenced Executors, Boundary Negotiators); Reinforcement of Challenge (RC); O-I learning systems and monitoring/learning mechanisms; scrutiny-dependent reinforcement; discourse-practice ratio; Concern-Thread Integrity Score (CTIS); organisational decoupling and risk entropy over time.
---	---	--

A.3 Interview Guide Development Rationale

The interview guide (final version 3.5, reproduced in Appendix A.1) was developed using Kallio et al.'s (no date) five-phase framework comprising (1) identification of prerequisites, (2) retrieval and use of prior knowledge, (3) formulation of draft themes and probes, (4) pilot testing, and (5) finalisation.

The finalised guide comprised seven thematic Sections aligned with FSB risk culture principles (FSB, 2014) and the four heuristic behavioural indicators: introduction and consent (5 min), background and role (5 min), encouraging alternative views (10 min), mechanisms for expression and assessment (10 min), the stature of control functions (10 min), interplay and cultural reinforcement (5 min), and monitoring and wrap-up (5 min). Core "must-ask" questions ensured coverage of critical themes, while optional probes allowed for deeper exploration. The guide included artefact prompts ("Can you share a committee agenda or minutes?") and reflexive probes ("How did that affect your willingness to speak up again?") to reveal interpretations of escalation barriers, feedback suppression, and the legitimacy of voice.

A.1 maps the four analytic research questions (RQs) to key interview Sections, example probes, and the heuristic behavioural indicators they mainly support. It demonstrates shows that RE enters primarily as a synthesis consideration for RQ4 rather than through event-level probing in RQ1–RQ3.

Pilot testing. The interview process was trialled with three senior-level participants (one at Board level, two S4/S5 management), representing diverse governance contexts, before the main fieldwork. Pilot interviews followed the same ethical protocols as the main study and lasted 30 to 45 minutes, allowing focused testing of the clarity and flow of core questions. The streamlined pilot guide (v1.1) included "must-ask" questions aligned to FSB risk culture principles (3.3.1 to 3.3.5) (FSB, 2014), deliberately excluding optional probes to prioritise feedback on fundamental question wording, sequencing, and participant comfort with sensitive topics.

Pilot feedback confirmed that the guide's thematic structure was suitable and that the questions encouraged honest discussion. Two improvements were made to the question design. Question 3.8.4.1 was clarified with concrete examples ("e.g., hotlines, skip-levels, forums") and prompts to reflect on actual usage; in practice, this redirected participants away from policy-level confirmation towards lived accounts of channel use, directly strengthening the sensitivity of the Voice Legitimacy and Escalation Friction indicators. Probe 3.8.5.2 was softened and reworded to include an artefact prompt ("Was this documented?") to enhance triangulation; this change was prompted by feedback that the control function stature questions required simpler wording and clearer document anchoring, and the revised probe produced more candid accounts alongside natural openings to corroborate claims through governance records. Three procedural changes also followed: digital consent materials and topic-area guidance were issued 48 hours before the interview after one participant reported initial discomfort; a sponsor communication was circulated to reduce resistance to participation; and active prompting was codified in the Interviewer Cheat-Sheet after it proved effective in keeping participants on track. The pilot also highlighted a need to expand the main guide with explicit indicator-focused sections and customised probes, which were incorporated into v3.5. Pilot data were excluded from the analysis dataset to maintain sample integrity.

A.4 Participant Information Sheet and Consent Process

Participant Information Sheet (PIS) – Main Study

Study Title:

RISK ENTROPY AND THE FRAGILITY OF ORGANISATIONAL VOICE: MAPPING COMMUNICATION AND CHALLENGE IN RISK CULTURE

Researcher: Marc Pharoah, Doctorate in Business Administration (DBA), Swiss School of Business Management (SSBM)	Supervisor: Removed due to NDA
--	--

Date:

14th October 2025

1. Introduction

You are being invited to take part in a research interview for my doctoral study. Before you decide, you must understand why the research is being conducted, what participation involves, and how your information will be used. Please read this sheet carefully and feel free to contact me with any questions.

2. Purpose of the Research

This study explores how risk is identified, communicated, and managed across different parts of [Institution]. The aim is to understand the behaviours, processes, and cultural factors that shape risk

management practices, with a focus on how organisations foster transparency, challenge, and control-function independence. The findings will inform a practical framework to support future organisational improvement.

3. What Participation Involves

You have been selected through a stratified random sampling process designed to ensure a balanced representation of seniority levels and business functions across the organisation. This approach helps the research capture a diverse and representative range of perspectives on risk culture.

You will be invited to take part in a semi-structured interview lasting approximately 45–60 minutes. The interview will be conducted in person or via Microsoft Teams, depending on your preference and location.

With your permission, the interview will be audio- and/or video-recorded and transcribed (including possible use of Teams' automatic transcription feature) to ensure an accurate record of the conversation.

The interview will explore several key themes aligned to [Institution] risk culture framework, including:

- How alternate views are encouraged and responded to
- How formal channels support risk escalation and openness
- The role and independence of control functions
- How these elements reinforce a culture of challenge and learning
- How outcomes of risk challenges are monitored and reflected upon

You may be invited to reference or describe internal artefacts (e.g., meeting minutes, dashboards, survey results) if relevant and comfortable. These help illustrate how risk culture operates in practice.

You are free to skip any question or pause the interview at any time.

4. Use of Data

Your responses will be analysed as part of the main study and may be quoted anonymously in the final thesis or related publications.

All data will be anonymised and securely stored. No identifiable information will be shared.

You will be invited to review a summary of the emerging findings (“member checking”) to ensure your views have been accurately represented.

5. Confidentiality and Data Protection

Your identity will remain strictly confidential. Any data collected will be anonymised and stored securely on encrypted, password-protected systems.

The Microsoft Teams recording will be downloaded promptly and deleted from the Teams platform once securely stored.

Only the researcher and supervisory team will have access to the data.

All data will be handled in accordance with SSBM research ethics policy and the requirements of [Institution].

6. Voluntary Participation and Right to Withdraw

Participation is entirely voluntary. You are under no obligation to take part.

You may decline to answer any question during the interview.

You may withdraw from the study at any time without penalty and request that any data collected up to that point be deleted.

Once the anonymised data has been integrated into the analysis (anticipated by November 30th 2025), it may not be possible to extract individual contributions.

7. Potential Risks and Benefits

There are no significant risks associated with participation.

Some questions may involve reflection on workplace practices or decision-making; you are free to skip any questions you do not wish to answer.

While you may not directly benefit from taking part, your insights will contribute to a practical framework aimed at strengthening organisational risk culture.

8. Contact Information

If you have any questions about the study, please contact:

Researcher: Marc Pharoah Doctorate in Business Administration (DBA) Candidate Swiss School of Business Management Email: Marc.Pharoah@[Institution].co.uk	Supervisor: [Supervisor, name withheld for anonymity] Chief Risk Officer Email: [withheld for anonymity]
---	--

9. Consent

If you decide to participate, you may provide your consent either by signing a Consent Form or by replying to the invitation email with a statement such as:

“I have read the Participant Information Sheet and agree to take part in the interview.”

APPENDIX B: CODEBOOK EVOLUTION LOG

This appendix is the audit trail for codebook evolution referenced in 3.6.1–3.6.4, documenting version history, post-saturation refinements, and the reflexive memos that triggered code revisions.

B.1 Codebook Version History (V0.1–V0.7, Nine Snapshots)

Nine codebook states are documented in the table below, spanning the full analytical lifecycle from initial open coding (V0.1) to the final archived state (V0.7). Intermediate sub-versions (V0.51, V0.52, V0.54, V0.55, V0.65) capture granular refinements during the active coding phase; all versions are accounted for in sequence.

Date	Version	Interviews Analysed	Total Codes	Parent Themes	Key Changes and Refinements	Evidence of Saturation
18 Nov 2025	V0.1	INT01–09 (approx.)	~60	EF, FS, VL, RE (flat structure)	Initial coding phase; hierarchical structure beginning to emerge from early interviews; RE treated as placeholder pending cross-unit synthesis	Early phase; parent themes emerging
24 Nov 2025	V0.5	INT01–18	~150	EF, FS, VL, RE (stable)	SATURATION POINT: All four parent themes stable and operational; hierarchical child-code structure fully established; RE operationalized with preliminary subcodes	Parent themes stable at INT18
25 Nov 2025	V0.51	INT01–19	~160	EF, FS, VL, RE (stable)	Minor clarifications to child-code boundaries following interview 19; zero new parent themes	Saturation confirmed (INT19)
27 Nov 2025	V0.52	INT01–20	~175	EF, FS, VL, RE (stable)	Elaboration of existing child codes; no new parent themes; complexity increasing within existing structure	Saturation confirmed (INT20)
28 Nov 2025	V0.54	INT01–21	~190	EF, FS, VL, RE (stable)	Final elaboration phase for interviews 19–21; three-consecutive-interview rule fully confirmed;	SATURATION FULLY CONFIRMED (INT19–21)

					zero new parent themes	
30 Nov 2025	V0.55	INT01-22	~220	EF, FS, VL, RE (stable)	Elaboration continues; post-saturation refinements begin; child-code density increasing through interviews 22-25	Post-saturation elaboration phase
9 Dec 2025	V0.6	INT01-25 + Artefacts	~438	EF, FS, VL, RE + refinements	POST-SATURATION REFINEMENTS APPLIED: (1) EF → "Fatigue" child code added (14 refs; boundary clarification distinguishing exhaustion from structural barriers); (2) FS → "visible challenge" elaboration (9 refs; performative vs. genuine engagement distinction); (3) RE → "Openness" and "Distortion" subcodes operationalized per Methods 3.3.5 (64 refs total; synthesis subcodes for cross-unit patterns)	Parent themes frozen; only child-code clarifications
16 Dec 2025	V0.65	INT01-25 + Artefacts (extended)	~470	EF, FS, VL, RE (stable)	Artefact triangulation intensifies; reference density increases +7%; no new codes added; codebook structure frozen	Structure stable; elaboration phase
14 Jan 2026	V0.7 (Final)	INT01-25 + All Artefacts	~501	EF, FS, VL, RE (archive)	FINAL CODEBOOK: All interviews and organisational artefacts (ERCC, Peakon, ToR, RE reflexive memos) fully coded; reference density +14% from V0.6; zero new parent or child codes; structure locked for analysis phase	Stable for findings and discussion

* RE is listed as a parent theme in NVivo for structural organisation purposes; analytically, it is treated as a synthesis construct inferred post-saturation from cross-unit EF, FS, and VL configurations, not coded at event level.

Saturation and Refinement Summary

Saturation Achievement

Saturation Point: Interview 18 (4–6 November 2025)

Saturation Confirmation: Interviews 19–21 (7–8 November 2025)

Evidence: Codebook v0.5 (24 Nov) documents all four parent themes (EF, FS, VL, RE) stable; zero new parent codes in any subsequent version (v0.51 through v0.7)

Duration: 9 weeks from saturation point (24 Nov) to final archive (14 Jan) with zero new parent-level emergence

B.2 Post-Saturation Refinement Record

This Section records the refinements made to codes after saturation was confirmed, documenting the rationale for each change.

The formal comparison of codebook states at V0.5 (interview 18) and V0.6 (interview 25), documenting the two inductive child-code refinements added post-saturation.

Post-Saturation Refinements (Child-Code Level)

Refinement	Indicator	Child Code	First Applied	Type	Final References	Rationale
1	EF	"Fatigue" subcode	V0.6 (9 Dec)	Boundary clarification	14	Distinguishes voice fatigue (exhaustion from ignored feedback) from procedural barriers (structural friction)
2	FS	"visible challenge" elaboration	V0.6 (9 Dec)	Boundary clarification	9	Elaborates distinction between performative openness (ritualised consultation) and genuine suppression (unheard input)
3–4	RE	"Openness" and "Distortion"	V0.6 (9 Dec)	Synthesis operationalisation	64	Per Methods 3.3.6: operationalizes Risk Entropy as synthesis construct with "Openness" and "Distortion" subcodes for cross-unit pattern mapping

Two substantive refinements (EF, FS) represent boundary clarifications within existing parent structures; two RE operationalisations are framework implementation (planned post-saturation synthesis), not new parent-level emergence.

B.3 Reflexive Memo Extracts Supporting Code Revisions

This Section reproduces selected extracts from the structured reflexive memos (Memo-INT-MS series) that directly triggered the three post-saturation child-code refinements recorded in B.2. Extracts are reproduced verbatim or near-verbatim from their original memo sections, with personal identifiers removed in line with the anonymisation protocol described in Section 3.9.4. Their purpose is to demonstrate that each refinement was induced from live memo observations rather than applied retrospectively, providing the confirmability record required under Lincoln and Guba's (1994) criterion of a traceable, open analytical audit trail.

B.3.1 Refinement 1, EF: Fatigue Subcode

Memo-INT-MS09 27 October 2025

"The willingness to challenge is framed as cultural, 'no one is scared to speak up', but there is an absence of memorable occasions where challenge changed a risk decision or process. When a challenge did arise, the process and outcome were handled informally, signalling a reliance on direct, informal communication even when escalation occurs. Reflexive memoing surfaces a potential illusion of openness in small teams, where comfort speaking within a known group may not translate to broader organisational challenge."

Analytic note (recorded at memo close): My initial code was *EF [escalation barrier, procedural rerouting]*. On reflection, the participant's language shifted across the account from confident early framing to flat resignation. The concern had reached the right person; what failed was follow-through. This is a fatigue dynamic, not a routing-friction dynamic. Coding it as *EF* risks misattributing the cause. Where withdrawal is attributed to accumulated non-response rather than procedural burden, the primary code should be *FS [voice fatigue]*. Flagged for codebook review: the *EF/FS* boundary at the point of withdrawal requires a child-code discrimination rule.

Revision applied (V0.6, 9 December 2025): *EF Fatigue* subcode added. Rule: code as *EF [fatigue]* only where withdrawal is attributed to procedural burden or routing confusion; code as *FS [voice fatigue]* where attributed to accumulated non-response or learned futility. Cross-reference: Appendix G.2, rows 3–4.

B.3.2 Refinement 2, FS: Visible Challenge Elaboration

Memo-INT-MS08 28 October 2025

"Survey feedback may capture voice in theory, but the participant does not recall ever seeing dashboard results or feedback loops, suggesting limited inclusion in broader governance processes. Formal channels exist but have not been tested; documentation of challenge is essentially absent at this operational level. Reflexive memoing surfaces the potential vulnerability that absence of formal documentation could create latent risks and invisible tensions as the organisation grows."

Analytic note (recorded at memo close): My first instinct was *MRC [mechanisms for raising concerns, underused]*. On further reflection, the survey channel is present and the participant was consulted, but the response loop is opaque. No outcomes were recalled; no changes could be described. This is not an unused channel but a channel that creates the appearance of engagement without substantive closure. A child-code is needed under FS to capture consultation that is *visible but non-substantive*, distinct from outright suppression, tokenism, and active discouragement.

Revision applied (V0.6, 9 December 2025): FS *Visible Challenge* elaboration applied. Rule: code as *FS [visible challenge / performative engagement]* where a formal consultation mechanism is available and used but the participant cannot describe substantive outcomes or follow-through. Cross-reference: Appendix G.2a, row 6.

B.3.3 Refinement 3, RE: Openness/Distortion Subcodes

Cross-memo synthesis: Memo-INT-MS06, Memo-INT-MS03, Memo-INT-MS20

[MS06]: "The [Role] feels strongly empowered to challenge and escalate issues due to seniority and established executive relationships. Assumes a high degree of psychological safety for self and direct reports, possibly overstating the openness of less senior or less confident staff."

[MS03]: "The [Role] emphasises the intentional design of an open, inclusive culture since the organisation's inception. As [Role], may underestimate the psychological distance or hesitancy that junior or non-executive staff feel, particularly outside the core governance circle."

[MS20]: "Team openness is described as present, barriers arise from geographic separation, communication via phone/Teams/email lacking face-to-face timeliness. Physical separation and leadership not visiting [Location] is recognised as a barrier to timely challenge, despite strong documentation practices at the centre."

Analytic note (cross-memo synthesis entry, post-interview 25): Reading the Oversight/SMT memos alongside the *[Location]* Operations account, a gap becomes analytically significant. At governance level, the discourse is one of authentic, designed openness. At operational and cross-jurisdictional levels, the same structures are

experienced as distant or filtered. These accounts do not describe the same organisation. The mechanism is not a discrete EF or FS event; it is a structural condition in which the *discourse* of openness at governance level co-exists with its *absence* in operational experience, a gap between espoused and enacted challenge, recurring across multiple units and both jurisdictions. RE must be operationalised with two subcodes: one marking passages where governance-level openness discourse is prominent, one where that discourse diverges from operational experience.

Revision applied (V0.6, 9 December 2025): RE *Openness* and *Distortion* subcodes operationalised. *RE, Openness*: governance or senior-actor passages describing challenge culture consistent with FSB (2014) ideals. *RE, Distortion*: operational or cross-jurisdictional passages in which those same claims are experienced as absent, tokenistic, or inaccessible. Cross-reference: Methods Section 3.6.5; Appendix D, RE-synthesis row.

Note: All extracts are reproduced from the five-Section reflexive memo template (date/context → researcher reflections → data interpretation → assumptions/biases → theoretical and methodological notes). Analytic notes represent closing annotations recorded at the time of writing and are traceable in the NVivo audit trail through the codebook version history in B.1.

APPENDIX C: SATURATION TRACKING MATRIX

This appendix documents the saturation monitoring process referenced in 3.4.3 and 3.7.3, providing theme-emergence tracking, saturation confirmation, reference-count progression, and non-response analysis.

C.1 Parent-Theme Emergence by Interview Cluster

Evidence of Parent-Theme Saturation at Interview 18 (24 November 2025)

Saturation Criterion (Methods 3.3.6): Parent-theme saturation was reached when no new parent-level themes for EF, FS, VL or RE emerged across three consecutive interviews, achieved at interview 18 and confirmed by interviews 19–21.

Interview Cluster	Date Range	Codebook Version	EF	FS	VL	RE	New Parents	Total Refs	Status
INT01-09	21 Oct–1 Nov	Early (undated)	Emerging ✓	Emerging ✓	Emerging ✓	Placeholder	4 emerging	~100	Building phase

INT10-15	2-3 Nov	Early (undated)	Stable ✓	Stable ✓	Stable ✓	Placeholder	0 new	~130	Hierarchy developing
INT16-18	4-6 Nov	24 Nov (v0.5)	✓	✓	✓	Operationalized	0 NEW	~150	SATURATION
INT19-21	7-8 Nov	25-28 Nov	Stable ✓	Stable ✓	Stable ✓	Elaborating	0 NEW	~200 +	Confirmed
INT22-25	9-12 Nov	9 Dec (v0.6)	Stable ✓	Stable ✓	Stable ✓	Stable ✓	0 NEW	438	Elaboration
Final	14 Jan	14 Jan (v0.7)	Stable ✓	Stable ✓	Stable ✓	Stable ✓	0 NEW	501	Stable

C.2 Three-Consecutive-Interview Saturation Confirmation

Three-Interview Window	Dates	New Parent Codes	Codebook	Verdict
INT16-18	4-6 Nov	0	24 Nov v0.5	Saturation reached
INT17-19	5-7 Nov	0	Post-sat. coding	Confirmed
INT18-20	6-7 Nov	0	Post-sat. coding	Confirmed
INT19-21	7-8 Nov	0	25-28 Nov	Saturation CONFIRMED

C.3 Reference Count Progression (Saturation Signal)

Codebook Version	Date	Cumulative Interviews	Total References	Growth vs. Prior	Interpretation
V0.1	18 Nov	~5	60	–	Early emergence phase
V0.5	24 Nov	18	150	+150%	Saturation inflection point
V0.52	27 Nov	~20	175	+17%	Post-sat. elaboration begins
V0.55	30 Nov	22	220	+25%	Interviews 19-25 elaboration continues
V0.6	9 Dec	25 + artefacts	438	+99%	Artefact coding + post-sat. refinements
V0.65	16 Dec	25 + extended artefacts	470	+7%	Triangulation deepens
V0.7	14 Jan	25 + all artefacts	501	+7%	Final archive; minimal change

Asymptotic Pattern: Steep growth 18–24 Nov (parent emergence), plateau 24 Nov onward (elaboration + triangulation), minimal final change (structure stable).

C.4 Reference Count Asymptotic Flattening

Codebook Date	Cumulative INT	Total Refs	Growth vs. Prior	Interpretation
18 Nov (v0.1)	~5	60	–	Early phase; steep growth expected
24 Nov (v0.5)	18	150	+150%	Saturation inflection point
9 Dec (v0.6)	25	438	+165%	Growth reflects elaboration + INT19–25, not new parents
14 Jan (v0.7)	25	501	+14%	Stable; only density increase

Interpretation: Expected asymptotic curve. Steep emergence phase (18–24 Nov) followed by elaboration plateau (24 Nov onward).

C.5 Saturation Monitoring Protocol

The dual-level monitoring procedure (parent-theme and unit-level thematic redundancy criteria) and the census-completion rationale.

Because the accessible population was small (n=62 eligible staff; 25 consenting, a 40% response rate), a census of all consenting participants was conducted rather than stopping at the aggregate saturation point (Braun et al., 2019a), eliminating concerns about premature stopping and ensuring interpretive completeness across all four embedded units. Coding continued beyond parent-theme saturation to confirm how core themes manifested across units and roles.

The following table summarises the achieved samples and saturation-based confidence levels for each embedded unit.

This protocol describes the three-consecutive-interview rule and reference-count threshold used to confirm saturation.

Embedded Unit	Sample	Saturation Status	Confidence
Oversight	n=5	Saturated at n=3	HIGH
Risk and Compliance (2nd Line)	n=5	Saturated at n=5	HIGH
1st Line Operations and Support	n=12	Saturated at n=7, confirmed n=8	HIGH
1st Line Business	n=3	Partial saturation	PROVISIONAL

C.6 Per-Interview Reference Count Detail

This table provides the per-interview reference counts that underlie the aggregate progression shown in C.3 and C.4.

Participant Data											
	Balancing risk (BRE)	Escalation Friction (EF)	Feedback Suppression (FS)	Ind. of Control Functions (ICF)	Mechs. for Raising Concerns (MRC)	Open to Alt Views (OAV)	Reinforcement of Challenge (RC)	Reps. of Control Functions (RCF)	Risk Entropy Synthesis (RE)	Voice Legitimacy (VL)	Total sample
Oversight	4	4	5	4	5	5	5	5	5	4	5
2nd	5	5	2	5	5	5	5	5	4	5	5
1st Line	5	12	4	8	12	12	12	10	12	12	12
1st Line Business	2	2	1	3	3	3	3	3	3	3	3
IOM	9	10	4	10	11	11	11	11	11	10	11
Jersey	7	13	8	10	14	14	14	12	13	14	14
S1/S2	6	9	3	6	9	9	9	8	9	9	9
S3	3	5	1	5	6	6	6	5	6	6	6
S4	2	4	2	4	4	4	4	4	3	4	4
S5	2	2	2	1	2	2	2	2	2	2	2
SL	3	3	4	4	4	4	4	4	4	3	4

RE (inferred from EF/FS/VL configuration; not directly coded)

C.7 Non-Response and Self-Selection Analysis

This Section analyses response patterns across sampled strata to assess whether non-response introduced systematic bias into the achieved sample.

The full discussion of the 37 non-participants, Morrison and Milliken (2000) framework, and the four offsetting design features.

Of the 62 eligible staff contacted through the CRO-approved invitation process, 25 consented and were interviewed, yielding a 40% response rate. The remaining 37 either did not respond to the invitation and follow-up reminder or declined to participate.

Because this study investigates voice, silence, and the perceived consequences of speaking up, the profile of non-participants warrants explicit consideration rather than being treated as a neutral sampling outcome.

Research on organisational silence shows that individuals who perceive higher personal risk in speaking up are generally less likely to volunteer for research on that topic (Morrison and Milliken, 2000; Brinsfield, 2013). In this study's context, staff experiencing increased FS may have been less inclined to participate, viewing the interview as another instance of voice with uncertain outcomes. Similarly, individuals with lower perceived VL may have self-excluded, believing they had nothing relevant to contribute—a framing that itself reflects the phenomenon being investigated.

If these patterns hold, the achieved sample may overrepresent staff who feel relatively safe or empowered to speak and underrepresent those whose experiences of suppression, friction, or delegitimised voice are most acute. The findings may therefore present a

conservative picture of EF, FS, and VL in parts of the organisation where participation was lowest.

Several design features partially offset this risk. First, the embedded-unit structure ensured that no single governance function was left unrepresented; all four units achieved the minimum threshold of 3–5 participants, preserving cross-unit comparison even where response rates varied. Second, artefact analysis, particularly the comparison of committee minutes with interview accounts, provided documentary evidence of suppression and friction that did not depend on the willingness of affected individuals to self-report. Third, reflexive memos recorded instances where non-participation itself became analytically meaningful; for example, the absence of any respondents from [specific function/location if applicable] was noted as a potential signal of disengagement or low perceived relevance, consistent with the broader FS pattern identified in adjacent units. Non-response is acknowledged as a boundary condition of the findings rather than an artefact to be explained away. Where the analysis identifies moderate or low FS/VL in a given unit, readers should consider that this assessment is based on those who chose to speak and that the experiences of non-participants may reflect a more pronounced manifestation of the very phenomena the study seeks to surface. This caveat is revisited in the discussion of limitations in Chapter 5.

APPENDIX D: EVIDENCE BASE AND DOCUMENTATION

This appendix provides the full triangulation tracking matrix referenced in 3.6.5 and 4.7. Each entry records a triangulation instance by indicator, embedded unit, and data source, documenting how interview and artefact evidence were combined to validate indicator intensity assessments.

D.1 Triangulation Instances by Indicator, Unit, and Source

The table documents each triangulation instance identified during analysis, recording the indicator activated, the embedded unit, the primary interview source, the corroborating artefact or secondary source, and the outcome classified as convergent, divergent, or novel insight. Each row represents a discrete analytical decision made during cross-source verification; the table should be read alongside Appendix B.1 (codebook) and Appendix C.1 (saturation matrix) to locate each instance within the broader audit trail.

Indicator	Embedded Unit(s)	Interview Evidence	Artefact Evidence	Triangulation Outcome	Intensity
EF	Risk and Compliance, First Line Ops	MS06: "Cyber training took forever—multiple forums"; MS49: "Training	ERCC Action Log Dec 2025: Cyber training deferred 5x (Q1 2024 → Nov 2024); 18-month delay	Divergence: Interviews claimed "timely"; artefact shows delay	Common

		kept getting pushed back"			
EF	First Line Ops, Risk and Compliance	MS21: "Interim spreadsheets while waiting"; MS49: "Dual-running painful"	ERCC Minutes Oct 2025: Holistic 3 delayed → 2026; spreadsheet workarounds Q3	Convergence: Procedural delay + workaround burden confirmed	Common
EF	Org-wide	Peakon Apr 2024: "15 forums before decision, then never communicated"	ERCC Action Log: Action 192 deferred 2023–2025, not operational	Convergence: Multi-forum drag confirmed	Common
EF	Risk and Compliance, Ops	MS49: "Suppliers slow-email-heavy, reluctant to escalate"	ERCC Minutes Oct 2025: "Cases ongoing 12 months... escalation needed"	Convergence: TPRM friction + weak escalation	Common
FS	Org-wide	MS16: "People keep quiet—no point raising it"	Peakon Apr 2024: "Never goes anywhere"; "They rolled their eyes"	Convergence: Self-censorship + discouragement cues	Common
FS	Org-wide	MS20: "Limited feedback—hardly any 1:1s"	Peakon Apr 2024: "No feedback—hardly do 1:1—no guidance"	Convergence: Absent follow-up confirmed	Common
FS	Org-wide	MS21: "We track everything—close on time"	Peakon Apr 2024: "Being looked at—not clear if acted on"	Divergence: Robust tracking claimed; ambiguous follow-through shown	Common
FS	Ops, Risk and Compliance	MS06: "Audit finding—self-assessment inaccurate"	ERCC Minutes Oct 2025: Resilience contradiction; no action owner	Convergence: Documentation gap without corrective follow-up	Illustrative
VL	Risk and Compliance	MS21: "Financial Crime team very approachable"	Peakon Feb 2024: "FC Ops team very approachable and helpful"	Convergence: Positive legitimacy—access + responsiveness	Common
VL	Ops, Risk and Compliance	MS49: "We raised callback issues—threshold adjusted";	ERCC Minutes Oct 2025: Callback approved, implemented, monitored	Convergence: Operational challenge → control adjustment + impact	Common

		MS06: "Data-led fraud optimisation"			
VL	Ops, Risk and Compliance	MS49: "SEEN escalation thresholds agreed—clear pathway"	ERCC Action Log Dec 2025: SEEN process documented, closed Dec 2021	Convergence: Formal mechanism codified	Illustrative
RE	Board/SMT, Risk, IT	MS06: "Difficult reconciling dashboards"; MS21: "Inconsistent definitions"	ERCC Minutes Oct 2025: "Mismatch between IT cyber and Group... clarification awaited"	Convergence: Reporting misalignment + definitional drift	Common
RE	Risk and Compliance, Ops	MS01: "TPRM red all year—slow progress"; MS21: "Persistent red, limited movement"	ERCC Minutes Oct 2025: "Top 5 static... TPRM only red throughout year... dependencies remain"	Convergence: Chronic red without closure—normalisation risk	Common
RE	Ops, IT	MS06: "CyberArk broken 1 year—external dependency"; MS21: "Long remediation timelines"	ERCC Minutes Oct 2025: "CyberArk broken 9 Sept 2024; resolved two weeks ago"	Convergence: 1-year delay in critical control—workaround entropy	Illustrative
RE	Ops	MS49: "Workarounds ossify into BAU—no post-change reviews"	Peakon Apr 2024: "Workarounds become BAU... need postchange reviews"	Convergence: Unresolved fixes normalised without learning	Common
RE	Ops	MS21: "R20 still causing issues—long delays"	Peakon Apr 2024: "R20 fixes taking far too long" (6+ mentions)	Convergence: Platform instability + slow stabilisation	Common
RE	Board/SMT, Risk	MS01: "We emphasise remediation—tight control, close findings"	ERCC Minutes Oct 2025: Multiple acceptances approved; one paper post-meeting	Divergence: Remediation rhetoric vs. routine acceptances, limited debate	Illustrative
RE	Ops, Jersey vs. IOM	MS15 (Jersey): "Positive culture for suggestions";	Peakon Apr 2024: "IOM not my job—terrible influence";	Novel Insight: Island-based microculture divergence—	Common

		MS20 (IOM): "Location creates challenges"	Jersey vs IOM contrasts	ownership fragmented	
--	--	--	----------------------------	-------------------------	--

D.2 Triangulation Summary Statistics

Category	Count	%	Examples
Convergent	11	61%	Persistent reds (TPRM x 3 packs); "Noting only" items; multi-year deferrals; Freedom of Opinion score validates local authenticity
Divergent	5	28%	1st Line "No escalations" vs. 2nd Line dashboard reds; Freedom of Opinion vs. friction; Forum terms of Reference (ToRs) vs. no minutes; Policy exists vs. "would have to find it"
Novel Insights	2	11%	Forum consolidation (monthly to quarterly); AI minutes trials revealing process vs. outcome focus

D.3 Indicator Distribution

Unit	EF	FS	VL	RE
Oversight (n=5)	Common (4/5)	Ubiquitous (5/5)	Common (4/5)	Ubiquitous
RISK and COMPLIANCE 2nd Line (n=5)	Ubiquitous (5/5)	Illustrative (2/5)	Ubiquitous (5/5)	Common
Operations 1st Line (n=12)	Common (11/12)	Illustrative (4/12)	Ubiquitous (12/12)	Ubiquitous
Business 1st Line (n=3) [Prov.]	Common (2/3)	Minimal (1/3)	Ubiquitous (3/3)	Ubiquitous

D.3.1 Lines Model × Heuristic Behavioural Indicators

Indicator	Intensity	Key Evidence
EF	Ubiquitous (5/5)	Relational hesitancy and gatekeeping when escalating to Board/Oversight; staff do not realise or feel able to use 'extra paths' to escalate
FS	Illustrative (2/5)	Recommendations noted, deferred, or quietly resisted; issues cycle without implementation
VL	Ubiquitous (5/5)	"Arming with regulation"; extensive evidence-gathering; reputational risk when challenging
RE	Common	Control staff recognise recurring issues and "thin" corrective action; see patterns others treat as isolated

D.3.2 Second Line "Control Without Voice" Profile

Location	EF	FS	VL	RE
Jersey (Governance)	Ubiquitous (13/14)	Common (8/14)	Ubiquitous (14/14)	Ubiquitous
Isle of Man (Operations)	Ubiquitous (10/11)	Minimal (4/11)	Ubiquitous (10/11)	Ubiquitous

D.4 Method-to-Indicator Support Matrix

The following matrix documents each analytical code's methodological grounding, observable indicators, and primary evidence sources, supporting the triangulation logic described in 3.3.1.

Code	Methods ref	Observable indicators (brief)	Primary data sources
BRE	3.3.1	Tension between control, efficiency, innovation	Interviews, artefacts
EF	3.3.2	Procedural, relational, cognitive barriers to escalation	Interviews, artefacts, memos
FS	3.3.3	Ignored input, absent follow-up, tokenism, voice fatigue	Interviews, artefacts, memos
ICF	3.3.1	Formal distance, structural authority, separation from business	Interviews, artefacts
MRC	3.3.1	Availability and usability of formal channels	Interviews, artefacts
OAV	3.3.1	Managerial openness to dissenting perspectives	Interviews, artefacts
RC	3.3.1	Active encouragement, follow-through, learning from challenge	Interviews, artefacts
RCF	3.3.1	How control roles are portrayed and perceived	Interviews, artefacts
RE	3.3.5	Fragmented, delayed, inconsistent risk information	All sources
VL	3.3.4	Who is seen as authorised/credible to speak up	Interviews, artefacts, memos

D.5 NVivo Query and Node Management

This analytical position has direct implications for how the evidence base for RE is constituted. Because RE is not coded at event level, it does not produce NVivo node references in the way EF, FS, and VL do. Instead, RE inferences are grounded in the cross-unit warrant documented in Table G.5.1 below, which records which unit-level configurations gave rise to the emergent condition, what conditions were met, and what disconfirming evidence was considered and why it did not displace the inference and why it did not displace the interpretation. Appendix D provides the triangulation tracking matrix showing how convergent and divergent evidence streams across units underpins the RE synthesis; Appendix B records how RE's operationalisation as a synthesis

construct was treated as intentionally post-saturation work within the codebook evolution.

Interviews and artefacts were loaded as separate source classifications; codes were applied via a shared codebook; matrix coding queries were run to identify indicator co-occurrences; text-search queries confirmed boundary-case consistency; segments were regrouped iteratively with rationale recorded in NVivo annotations.

APPENDIX E: PHILOSOPHICAL POSITIONING AND RIGOUR DOCUMENTATION

This appendix extends the philosophical and rigour positions summarised in 3.1 and 3.7, providing the full ontological and epistemological narrative, the researcher independence statement, and the validity framework.

E.1 Ontological and Epistemological Positioning

The detailed relativist ontology and subjectivist epistemology exposition. Ontologically, the study adopts a relativist stance, asserting that reality is multiple, context-dependent, and interpreted through human experience. There is no single "true" risk culture; instead, there are multiple, partially overlapping realities of risk discourse and challenges to viability, shaped by local histories, power relations, identities, and role expectations. The same escalation policy or speak-up channel may be experienced by one group as an invitation to raise concerns and by another as a cue to remain strategically silent.

The research design is inherently shaped by relativism, which justifies concentrating on participants' contextual accounts of how escalation and challenge operate within their part of the organisation, as well as on the local meanings linked to silence, hesitancy, and dissent. The study does not seek an overarching 'risk culture level' across the entire organisation. Instead, it identifies patterned differences in EF, FS, and VL across embedded units (Board/SMT, Risk and Compliance, First Line Business, First Line Operations and Support Functions) and examines how these local arrangements appear to reinforce or undermine the fabric of risk culture (RE, Section 3.4), within the same formally governed institution (Eisenhardt, 1989b; Yin, 2009).

The study adopts a subjectivist stance, which asserts that knowledge is co-constructed by the researcher and participants, always influenced by perspective, context, and interpretation. Understanding of EF, FS, and VL is developed from participants' accounts; RE is inferred as an emergent condition from cross-unit configurations of these indicators that participants' accounts are direct reflections of an independent reality, but by viewing them as situated sense-making related to speaking up, remaining silent, and challenging decisions in risk-relevant situations.

Within this framework, the researcher serves as an interpretive instrument rather than a neutral observer. Reflexive journaling records are used to document how assumptions, positionality, and emotional responses develop through fieldwork and analysis, clarifying how interpretations of EF, FS, and VL are reached, and how RE is inferred as an emergent condition from cross-unit configurations of these indicators, rather than imposed beforehand (Guba and Lincoln, 1994; Tobin and Begley, 2004). Reflexive memos are written after each interview in a structured format. They function as supplementary interpretive notes that enhance transparency and credibility, not as standalone "hard" evidence.

Reflexive thematic analysis, which considers themes as active analytical constructs developed through repeated engagement with the data, rather than as pre-existing entities to be "discovered." In this study, coded segments are progressively refined into EF, FS, and VL; RE is inferred later through cross-unit explanation-building, not assigned at segment level. Explanatory analysis is then employed to examine how different arrangements of these indicators correlate with governance outcomes, such as delayed escalations, muted dissent, or near misses (Eisenhardt, 1989b; Yin, 2009).

E.2 Researcher Positionality and Independence Statement

E.2.1 Independence Declaration

The researcher has no employment, directorial, advisory, or reporting relationship with the case organisation, and held no such relationship at any point during the design, fieldwork, or analysis phases of this study. The researcher is an external risk governance practitioner engaged solely for the purposes of this doctoral research.

Research access was formalised through a Non-Disclosure Agreement executed on 4 September 2025 between the researcher and the case organisation. This agreement governed the confidential treatment of all organisational data, artefacts, and participant information encountered during fieldwork. The existence of this agreement confirms the arm's-length nature of the relationship: the organisation treated the researcher as an external third party subject to formal legal confidentiality obligations, not as an employee, consultant, or affiliated party.

The NDA does not constrain the analytical independence of this study. Its terms govern data protection and confidentiality of identifying information, not the researcher's interpretive conclusions, theoretical framing, or critical findings. All findings, including those that reflect unfavourably on formal governance practices within the institution, were developed and reported without organisational review or approval.

E.2.2 Reflexive Stance: Worked Example

One concrete example illustrates this dynamic. During early coding, the label 'appropriate gatekeeping' was initially applied to several segments in which line managers filtered

concerns before escalation, a framing drawn from professional experience of risk governance practice, where some filtering is considered legitimate. On reflexive review, the memo recorded after interview MemoINTMS09 noted that this framing had suppressed the participant's own attribution of the filtering as obstructive. The code was revised to 'informal gatekeeping (EF)' and the memo observation was retained in Appendix B.3 as an audit record of how practitioner vocabulary had temporarily overridden participant meaning. This revision affected the EF profile for the Risk and Compliance unit and is an example of how the reflexive journal functioned as a corrective instrument rather than a narrative record.

The interpretivist paradigm embraced here recognises certain limitations. Pham (2018), in reviewing positivism, interpretivism, and critical inquiry, highlights subjectivity bias and limited generalisability as inherent trade-offs of interpretivist research. These limitations are seen as boundary conditions rather than flaws that disqualify the study: it prioritises contextual depth and explanatory detail over statistical breadth, aligned with the relativist ontology discussed earlier. The methodological safeguards outlined in Sections 3.7–3.9 (including triangulation, member checking, and deviant-case analysis) aim to manage, rather than eliminate, the effect of researcher subjectivity.

E.2.3 Practitioner Bias Monitoring Record

The practitioner-background risk identified in Section 3.2.3, specifically the tendency to normalise institutional filtering as legitimate governance practice, was actively monitored throughout the coding process. At least three reflexive memo entries in Appendix B.3, dated between October and November 2025, record instances where initial code assignments were revised after the researcher identified that professional vocabulary had temporarily overridden participant attribution. Each of these revisions is traceable in the codebook version history documented in Appendix B, particularly in the transition from V0.2 to V0.3, where the affected codes and the rationale for their revision are recorded. Together, these entries demonstrate that the reflexive journal functioned not merely as a record of the researcher's experience but as an active corrective mechanism within the analytical process.

E.3 Validity Framework

The Baskarada (2014) construct/internal/external validity mapping table.

Validity Criterion	Yin (2009) Test	Baskarada (2014) Application	This Study's Design Response	III Location
Construct Validity	Use multiple sources of evidence; establish chain of evidence	Triangulation across sources; member-checking	Three overlapping data streams (interviews n=25, artefact analysis, reflexive memos); participant theme summaries returned for member-checking;	3.7.2, 3.5.6

			convergent/divergent/novel tagging in D.1	
Internal Validity	Pattern-matching; explanation-building; rival explanations addressed	Cross-unit comparison; mechanism identification; disconfirming cases	Explanation-building links unit-level EF/FS/VL configurations to mechanisms; deviant case analysis actively searched for disconfirming evidence (e.g., high EF co-occurring with strong VL); codebook stability confirmed V0.5→V0.6	3.7.3, 3.7.4, 3.6.5
External Validity	Analytical generalisation (not statistical); replication logic across embedded units	Thick description enabling reader transferability judgements	Unit-level configuration tables and vignettes in Ch.4; NVivo attributes (line, function, location, band) enable profile replication; EF/FS/VL as transferable diagnostic heuristics, not statistical generalisations	3.7.6, 3.6.3
Reliability / Dependability	Use case study protocol; develop case study database	Transparent coding rules; reflexive audit trail	Codebook V0.1–V0.7 documented in Appendix B; reflexive memos MemoINTMS01–50; NVivo annotations recording rationale for each code assignment; saturation matrix in Appendix C	3.7.3, 3.7.4, 3.6.1

APPENDIX F: CASE BOUNDARIES AND EMBEDDED UNIT DEFINITIONS

This appendix documents the case and embedded-unit boundaries referenced in 3.2.5–3.2.6, including temporal, spatial, and functional boundary definitions, unit inclusion criteria, and target interview counts.

F.1 Case Boundaries: Temporal, Spatial, and Functional Definition

Fieldwork dates, geographic scope (Jersey/Isle of Man), and conceptual boundary exclusions.

Temporal Boundaries. Fieldwork spans a six-week period (21 October–7 November 2025), aligning with the organisation's governance meeting cycle and deliberately avoiding the year-end period when resource availability is limited. This time frame ensures that the data reflect a consistent organisational moment rather than a mix of shifting governance contexts.

Geographic and Functional Scope. Spatially, the case covers the institution's main operational sites: Jersey, which hosts the headquarters and retail along with private-banking activities, and the Isle of Man, which concentrates on compliance and operations. This multi-jurisdictional framework captures risk culture dynamics across


similarly regulated locations. Functionally, the case involves business divisions (Retail and Private Banking), control functions (Risk and Compliance), and enabling functions (Operations and Human Resources). These groupings follow the three lines model, which separates first-line business units, second-line risk and compliance, and third-line assurance.

Definitional Boundary. Conceptually, the case focuses on communication and challenging behaviours related to risk identification, escalation, and feedback within the formal governance framework. Other behaviours, such as commercial customer service interactions, product development, or outsourced functions (e.g., external audit), are purposely excluded. Internal Audit, delivered as a group service, meets the exclusion criteria and did not respond during the fieldwork period. The real case therefore concentrates on the local Board and Senior Management Team (Board/SMT), second-line risk and compliance functions, and first-line business and operations units.

F.2 Embedded Unit Inclusion Criteria

A unit was included in the embedded case structure if it satisfied all five of the following criteria, applied consistently across the four governance units and informed by Baxter and Jack (2008) and Baskarada (2014):

1. Maintains formal risk-communication mechanisms (e.g., escalation routes, committee reporting lines, or documented challenge processes).
2. Includes at least one senior-level participant with direct responsibility for risk identification, oversight, or escalation.
3. Operates within the fieldwork temporal boundary (21 October – 7 November 2025).
4. Contains at least three to four reachable interview participants, ensuring minimum analytical adequacy for within-unit profiling.
5. Directly participates in risk identification and escalation governance within the three lines model, rather than performing purely support or outsourced functions.



F.3 Embedded Unit Definitions and Targets

The four embedded units satisfying the inclusion criteria set out in F.2 are defined below, together with their target and achieved interview counts.

Embedded Unit	Roles Included	Geographic Span	Interview Target
Board/ Oversight/SMT	CEO, CRO, CFO, COO, INEDs	Jersey + Isle of Man	4-5
Risk and Compliance (2nd Line)	Risk managers, compliance officers, governance specialists	Jersey + Isle of Man	5-6
First Line Business (Retail and Private Banking)	Business managers, relationship directors, support staff	Jersey + Isle of Man	3-4
Operations and Support Functions	Operations managers, HR staff, process/systems leads	Jersey + Isle of Man	8-10
TOTAL CASE TARGET	–	Multiple	20-25

APPENDIX G: INDICATOR OPERATIONAL DEFINITIONS AND DISCRIMINATION RULES

This appendix consolidates the full indicator framework referenced in 3.3 and 3.6.4, covering operational definitions, observable cues, intensity rubrics, and boundary-case discrimination rules for EF, FS, VL, and RE.

G.1 Full Indicator Framework Overview

Table G.1 maps the complete analytical codebook across three layers: three proximate behavioural indicators (EF, FS, VL) coded directly at segment level; six contextual mechanism codes (BRE, ICF, MRC, OAV, RC, RCF) capturing the governance conditions in which those indicators operate; and one emergent synthesis construct (RE) inferred through cross-unit explanation-building rather than event-level coding. Each entry shows its Chapter III Section reference, observable indicators, and primary evidence sources. Boundary-case discrimination rules are in G.8; codebook version history in Appendix B.

Code	Short label	Full label	Methods Section	Observable indicators (brief)	Primary data sources
A	BRE	Balancing risk, efficiency and innovation	Section 3.3.1	Tension between control, efficiency, and commercial innovation	Interviews, artefacts
B	EF	Escalation Friction	Section 3.3.2	Procedural, relational, and cognitive barriers to escalation	Interviews, artefacts, memos
C	FS	Feedback Suppression	3.3.3	Ignored input, absent follow-up, tokenism, and voice fatigue	Interviews, artefacts, memos
D	ICF	Independence of Control Functions	Section 3.3.1	Formal distance, structural authority, and separation from business	Interviews, artefacts
E	MRC	Mechanisms for Raising Concerns	Section 3.3.1	Availability and usability of formal channels for concerns	Interviews, artefacts
F	OAV	Open to Alternate Views	Section 3.3.1	Managerial openness to dissenting or alternate perspectives	Interviews, artefacts
G	RC	Reinforcement of Challenge	Section 3.3.1	Active encouragement, follow-through, and learning from challenge	Interviews, artefacts
H	RCF	Representation of Control Functions	Section 3.3.1	How control roles are portrayed and perceived in the organisation	Interviews, artefacts
I	RE	Risk Entropy Synthesis	3.3.5	Fragmented, delayed, or inconsistent risk information	All sources

J	VL	Voice Legitimacy	3.3.4	Who is seen as authorised and credible to speak up	Interviews, artefacts, memos
---	----	------------------	-------	--	------------------------------

RE is an emergent synthesis construct derived from cross-unit configurations of EF, FS, and VL (Section 3.3); it is listed here for codebook completeness but is not applied as a direct event-level code equivalent to EF, FS, or VL.

G.2 Escalation Friction (EF): Operational Definition and Observable Cues

Escalation Friction describes procedural, relational, and cognitive barriers that block, distort, or prevent the upward flow of risk-related concerns through formal governance channels. It encompasses three distinct barrier types that were treated as child codes during reflexive thematic analysis and used to profile EF intensity at the unit level.

Procedural barriers include formal obstacles such as unclear escalation procedures, multi-function approval dependencies, and excessive documentation requirements that introduce structural delay or dropout before a concern reaches a control function.

Relational barriers involve informal gatekeeping, where intermediaries, managers, supervisors, or peers, reframe, minimise, or absorb issues before they reach governance forums. This barrier type is operationalised through evidence of selective framing, filtered escalation, and the rerouting of concerns into informal resolution.

Cognitive barriers capture internalised constraints on escalation: normalisation of deviance (concerns are not raised because dysfunction has become routine), problem minimisation (the issue is perceived as insufficiently serious to warrant escalation), and learned futility (prior experience of non-response reduces the probability of future attempts).

Observable cues, three primary signal families:

- *Informal gatekeeping*: intermediaries reframe or minimise issues before they reach control functions; concerns are noted as resolved at manager level without reaching a governance forum
- *Selective documentation*: dissenting views, concerns, or minority positions are omitted or sanitised in committee minutes and escalation records
- *Escalation silence*: an absence of formally recorded concerns from channels despite interview and artefact evidence of informal raising, the gap between what is raised informally and what appears in governance records

Evidence basis and triangulation logic:

Evidence is drawn from interview narratives coded for event sequencing and distortion, committee minutes cross-referenced with frontline accounts to identify omissions and reframings, and reflexive memos documenting researcher observations of hesitancy and learned adaptation during interviews. Full triangulation logic and instance-level tracking

are in Appendix D. Boundary-case discrimination rules, distinguishing EF from FS where overlapping surface language is present, are in Appendix G.8.

G.3 Feedback Suppression (FS): Operational Definition and Observable Cues

Feedback Suppression describes the practices and norms, active, passive, and structural, that reduce the likelihood of future challenge by blocking, discouraging, or failing to respond to evaluative input once it has been raised. Where EF concerns whether a concern reaches a governance channel, FS concerns what happens after it does: whether it is tracked, acknowledged, and acted upon, or whether the response itself signals that raising concerns is futile.

FS manifests across three child code categories used during reflexive thematic analysis:

Active discouragement involves explicit signals that challenge is unwelcome: punitive responses, reputational costs, reassignment, or visible career consequences for those who dissent.

Passive discouragement operates through absence rather than action: concerns are raised but receive no follow-up, no acknowledgement, and no visible closure. The issue is not rejected, it simply disappears from the governance record.

Structural disincentives embed suppression in organisational design: reward and recognition systems that privilege conformity, performance frameworks that do not credit challenge behaviour, and governance calendars that do not create space for iterative feedback.

Observable cues, three primary signal families:

- *Absent follow-up: raised issues are not tracked, assigned, revisited, or closed at subsequent meetings; action logs show performative entries with no demonstrable resolution*
- *Punitive signalling: dissenters face reputational costs, reassignment, or stalled career progression, whether experienced directly or observed vicariously by colleagues*
- *Voice fatigue: repeated, ignored feedback leads individuals to stop attempting to escalate, evidenced through interview narratives describing withdrawal, resignation, or deliberate self-censorship after prior non-response*

Evidence basis and triangulation logic:

Evidence is drawn from interview accounts coded for follow-up behaviour and perceived consequences of dissent, action logs compared with interview claims to identify missing or performative entries, and reflexive memos documenting researcher observations of resignation or adaptive silence during interviews. Full triangulation logic and instance-

level tracking are in Appendix D. Boundary-case discrimination rules, distinguishing FS from EF where suppression operates at the channel entry point rather than post-receipt, are in Appendix G.8.

G.4 Voice Legitimacy (VL): Operational Definition and Observable Cues

Voice Legitimacy describes the perceived appropriateness and credibility of speaking up within a specific role, unit, or context. Unlike EF, which concerns whether a concern can physically reach a governance channel, or FS, which concerns whether it is received and acted upon, VL concerns whether the speaker is perceived by themselves and others as authorised to raise it in the first place. Legitimacy is not structurally assigned; it is negotiated through status, expertise, relational ties, and institutional narratives, and can vary significantly between units even within the same formally governed organisation.

VL manifests across three child code categories used during reflexive thematic analysis:

Differential reception captures instances where the same concern, raised in equivalent terms, is treated with different weight depending on the speaker's role, seniority, or organisational background. A concern from a Board member and an identical concern from an Operations staff member may receive qualitatively different governance responses.

Legitimacy narratives are the explanatory frames that participants use to describe who is entitled to speak and on what basis. These include positive framing, speaking up as constructive engagement, professional responsibility, or valued input, and negative framing, speaking up as disloyalty, overreach, or troublemaking. Negative legitimacy narratives are particularly significant as early suppression signals.

Role-legitimacy boundaries are perceived limits on the appropriate scope of voice: views that challenge is appropriate for certain roles (Risk, Compliance) but not others (Operations, junior staff), or appropriate on certain topics (process improvement) but not others (strategic governance or senior conduct).

Observable cues, three primary signal families:

- *Differential reception: identical concerns treated differently based on the speaker's role, seniority, or background; meeting minutes record contributions selectively by seniority or function*
- *Legitimacy narratives: participants frame speaking up as requiring justification, or describe reluctance based on role boundaries rather than content of the concern*
- *Role-legitimacy boundaries: participants describe categories of concern that "aren't my place to raise" or governance forums perceived as appropriate only for certain seniority levels*

Evidence basis and triangulation logic:

Evidence is drawn from interview narratives coded for credibility attributions and role-based reasoning about who can legitimately speak, cross-checked against meeting minutes recording who speaks and how contributions are documented. Reflexive memos provided supplementary corroboration where positionality review revealed that participant hesitancy during the interview itself reflected enacted VL rather than personal affect. Full triangulation logic and instance-level tracking are in Appendix D. Boundary-case discrimination rules, distinguishing VL from FS where silence reflects delegitimised voice rather than suppressed feedback, are in Appendix G.8.

G.5 Risk Entropy (RE): Operational Definition and Convergent Patterns

Risk Entropy is a cumulative, emergent synthesis condition rather than a proximate behavioural mechanism. Unlike EF, FS, and VL, each of which describes a distinct barrier operating at the unit level and is coded directly at the event and segment level during reflexive thematic analysis, RE is not assigned to individual coded segments. It is inferred through cross-unit comparison and explanation-building once unit-level EF, FS, and VL profiles have been constructed. This distinction is analytically significant: RE has no NVivo node of its own; its evidence base is constituted through cross-unit explanation-building and the inferential warrant documented in G.5.1, rather than segment-level coding.

Inferential procedure:

Each embedded unit is first profiled for EF, FS, and VL intensity. Where multiple units independently exhibit elevated EF alongside elevated FS, the organisation's feedback loops are doubly compromised, concerns are both slowed in transmission and unreceived when they arrive. Where this coincides with low VL across units, a condition in which most staff do not perceive themselves as credible speakers, the cumulative effect is a governance system that processes risk symbolically rather than substantively. RE is the label applied to that emergent condition.

Three-configuration warrant, sufficient conditions for inferring elevated RE:

- *High EF in multiple units + FS common or ubiquitous organisation-wide: degrades both the transmission and reception of risk information simultaneously*
- *VL highly stratified + low-intensity challenge normalised across governance lines: removes the conditions under which correction is possible, even when concerns do reach a forum*
- *Artefact records and frontline accounts diverge systematically, indicates that formal governance is recording a version of risk communication that differs materially from what is experienced in practice*
- *Elevated RE is inferred when at least two of the three configurations are present concurrently; where all three are met, the inference is treated as high-confidence;*

where only two are met, the synthesis record documents why the third was absent and what alternative explanation was considered.

The inferential warrant for each unit-level RE inference is documented in G.5.1 above. Disconfirming evidence considered for each unit is recorded in the final column of that table. The triangulation tracking matrix in Appendix D provides the underlying source-level audit trail.

Convergent patterns, three primary signal families:

Information fragmentation: different accounts of the same event, hazard, or decision across units or communication channels, detectable through cross-unit interview comparison and artefact cross-referencing

Normalisation of ambiguity: routine acceptance of unclear, incomplete, or contradictory risk information without seeking clarification, evidenced through interview narratives describing tolerance of persistent ambiguity and artefact records showing unresolved items carried forward repeatedly

Disconnected escalation: concerns reaching governance forums but not resulting in coordinated organisational response, evidenced through committee minutes showing persistent reds, deferred actions, and absent closure entries

Evidence basis:

RE inferences are grounded in the cross-unit warrant documented in G.5.1, which records which unit-level configurations of EF, FS, and VL gave rise to the emergent condition, the specific cross-unit patterns that drove the inference, and the disconfirming evidence considered and why it did not displace the interpretation. Appendix D provides the triangulation tracking matrix; Appendix B records how RE's operationalisation as a post-saturation synthesis construct was treated during the codebook evolution.

G.5.1 RE inference logic and configurations

Risk entropy (RE) is a cumulative, emergent synthesis condition, not a proximate behavioural mechanism: EF, FS, and VL are each coded at the event and segment level during reflexive thematic analysis, whereas RE is not assigned to individual coded segments but is inferred from their cross-unit configuration once unit-level profiles have been constructed. Unlike EF, FS, and VL, each of which describes a distinct barrier operating at the unit level and is coded directly at the event and segment level during reflexive thematic analysis, RE is not assigned to individual coded segments. It is inferred through cross-unit comparison and explanation-building once unit-level EF, FS, and VL profiles have been constructed.

The inferential procedure works as follows. Each embedded unit is first profiled for EF, FS, and VL intensity using the coding and triangulation methods described in Sections 3.6.2 and 3.6.3. Where multiple units independently exhibit elevated EF alongside elevated FS, the organisation's feedback loops are doubly compromised: concerns are both slowed in transmission and unreceived when they arrive. Where this coincides with low VL across units, a condition in which most staff do not perceive themselves as credible speakers, the cumulative effect is a governance system that processes risk symbolically rather than substantively. RE is the label applied to that emergent condition. It is synthesised through the explanation-building stage (Section 3.6.5), drawing on cross-unit pattern matching, disconfirming case analysis, and the inferential warrant documented in G.5.1, not through an additional round of segment-level coding.

Three configurations of EF, FS, and VL cross-unit patterns are treated as sufficient warrant for inferring elevated RE: where high EF is present in multiple units and FS is common or ubiquitous organisation-wide, degrading both the transmission and reception of risk information; where VL is highly stratified and low-intensity challenge is normalised across governance lines, removing the conditions under which correction is possible; and where artefact records and frontline accounts diverge systematically, indicating that formal governance is recording a version of risk communication that differs materially from what is experienced in practice. No single pattern is decisive; RE is an interpretive judgement about their convergence and cumulative weight, made explicit in the cross-unit warrant documented in G.5.1 and opened to scrutiny through the audit trail in Appendix D.

Note on intensity labels: RE intensity descriptors (Ubiquitous, Common, etc.) are retained from the G.6 rubric to preserve cross-indicator readability. However, for RE these labels denote the strength and pervasiveness of the inferential warrant across the unit's structural configuration, not participant-presence counts as defined in G.6 for EF, FS, and VL. RE has no NVivo event-level node; the label is assigned after cross-unit EF–FS–VL profiling is complete, per the process described in G.5.

Condition	Met?	Evidence basis	Source
C1 – Escalation blocked (EF)	✓ Ubiquitous	All three EF barrier types present; procedural, relational, cognitive; highest EF reference count across all units	Interviews + ART-ERCCDec2025
C2 – Response absent (FS)	✓ Illustrative	Concerns raised locally but not escalated further; voice fatigue ("what's the point?") documented	Interviews MS46, MS16
C3 – Voice collapse (VL)	✓ Ubiquitous	"Not my place to question" recurring; operational voice treated as situational input only	Interviews across all 12 Operations participants

C4 – Artefact- account divergence	✓	Oversight describes cross-jurisdictional inconsistency as normal variation; Operations describes same pattern as daily dysfunction requiring workarounds	MS21/MS23 vs MS01; ART-ERCCDec2025
--	---	---	------------------------------------

G.5.2 Worked Examples of RE Emergence by Unit

Table G.5.2 provides a worked application of the inferential warrant criteria to all four embedded units. For each unit, the table records the proximate indicator inputs, identifies which G.5 convergent patterns are met, states the resulting RE inference, and documents the disconfirming evidence considered and why it did not displace the interpretation. This table constitutes the primary audit trail for the RE inferences reported in Tables 4.4–4.7 and Figures 4.5-4.6.

Unit	EF Input	FS Input	VL Input	G.5 Convergent Patterns Met	RE Inference	Disconfirming Evidence Considered
Oversight (<i>Symbolic Governance</i>)	Common – mandates procedural gates but experiences only polished, filtered escalations	Ubiquitous – items tabled, minuted, formally closed; formal recording substitutes for corrective action (MS01; ART-ERCC Oct/Dec 2025)	Common – control-function voices treated as inherently authoritative; operational voices heard selectively via managers	1-Information fragmentation: filtered escalations strip operational texture before reaching governance. 2-Normalisation of ambiguity: cross-jurisdictional inconsistency framed as acceptable variation (MS21, MS23) vs. experienced as dysfunction below. 3-Disconnected escalation: Action 192 deferred 2023–2025; persistent reds unresolved across three ERCC cycles	Ubiquitous – entropic condition is pervasive and self-sealing; architecture mandated by Oversight is precisely what filters friction from reaching it	MS01 expresses genuine confidence in escalation system; treated as confidence-closure (cognitive EF) rather than counter-evidence, as Operations and 2nd Line independently corroborate the filtering
2nd Line R&C (<i>Control Without Voice</i>)	Ubiquitous – Group dependencies and multi-function consensus required; challenge diluted in transit	Illustrative – challenge noted but deferred or reframed; "noted but not actioned" is recurring pattern (MS16; ART-ERCC Dec 2025 noting-only items)	Ubiquitous – authority granted structurally but contested in practice; must argue for credibility case by case	1-Information fragmentation: concerns surface at governance level but are procedurally absorbed rather than actioned. 2-Disconnected escalation: formal independence co-exists with repeated deferral and credibility contests. 3-Normalisation of ambiguity: MS16 – " <i>sometimes it's not worth it</i> " – learned futility normalised within structurally authoritative role	Common – entropy is present and consistent; partially offset by structural access to governance forums, which provides a weak corrective signal	2nd Line holds formal charter rights and direct board access – treated as structural feature that moderates but does not prevent entropy, given documented absorption pattern
1st Line Operations (<i>Silenced Executors</i>)	Ubiquitous – all three barrier types	Illustrative – concerns raised locally,	Ubiquitous – operational voice bounded by role; " <i>not my</i>	1-Information fragmentation: cross-unit inconsistency experienced as daily friction; same reality	Ubiquitous – strongest RE inference signal across all	MS10 describes compartmentalised post-mortems within risk team – treated

	simultaneously active; escalation drops before reaching formal channels (MS46; ART-ERCC Dec 2025 resourcing deferral)	feedback loops stop at immediate manager; voice fatigue evident – "what's the point?"	place to question" is recurring phrase; legitimacy situational not structural	interpreted as acceptable variation at governance level (MS21/MS23 divergence). 2-Normalisation of ambiguity: workarounds resolve dysfunction locally without surfacing it – each workaround is an unescalated signal. 3-Disconnected escalation: highest EF reference count and widest discourse-practice gap (312) yet lowest governance access – unit with most proximity to risk has least influence on policy	units; information fragmentation operates through omission and workaround normalisation, not active blockage	as partial corrective that is silo-bound and does not address cross-unit fragmentation
1st Line Business (Boundary Negotiators – Provisional)	Common – escalation occurs informally through relationship networks; formal channels bypassed in favour of influence	Minimal – suppression operates through self-regulation; business units pre-filter what reaches governance to maintain credibility	Ubiquitous – voice is revenue-legitimated; concerns tied to customer/revenue impact are heard, operational concerns without commercial framing are marginalised (MS16; ART-ERCC Oct 2025 – no items to escalate vs. 2nd Line dashboard reds)	1-Information fragmentation: 1st Line Business forum records no escalations while 2nd Line dashboard shows TPRM and liquidity concentration red simultaneously – documented in same governance cycle 2-Normalisation of ambiguity: revenue-framing acts as cognitive gate; operational risks without commercial salience are pre-filtered before any channel is used. 3-Disconnected escalation: governance receives commercially shaped concerns, not the operational substance beneath them	Ubiquitous (provisional – n=3, lower confidence) – direct governance access partially mediates disconnect; revenue-gated filtering nonetheless generates fragmentation between what is raised and what exists	Lower sample (n=3) means this inference carries provisional status; pattern is internally consistent across all three accounts but cannot be confirmed at saturation

G.6 Indicator Intensity Rubric and Cross-Unit Categories

The four intensity levels used throughout Chapter IV to characterise indicator prevalence within each embedded unit are defined below. These are interpretive thresholds applied to the sampled roles within each unit, not organisation-wide prevalence estimates.

Level	Definition	Application Threshold	Artefact Corroboration Role	Analytical Use
Ubiquitous	Every sampled participant in the unit referenced the indicator	All <i>n</i> participants in the unit	Confirms but not required to reach level	Included in unit profile and cross-unit synthesis
Common	The majority of sampled participants referenced the indicator	More than half of unit participants	Can substitute for one participant account	Included in unit profile and cross-unit synthesis
Illustrative	A minority of participants referenced it, sufficient to note as a pattern	2+ participants, non-trivial	Supports but does not elevate level	Included in unit profile; carried with lower confidence notation
Minimal	Single or isolated reference; noted but not analytically weighted	1 participant or artefact only	Does not elevate to Illustrative alone	Excluded from unit profile; noted in reflexive memos only

Cross-Unit Illustration: How Intensity Levels and Triangulation Tags Combine in Practice

The following three examples show how intensity levels and triangulation tags were applied together when reading across embedded units. These are illustrative extracts, not summaries; full unit profiles are in 4.6.

Oversight unit, EF: Common / VL: Common / RE inferred: Oversight participants produced the fewest EF and VL references in the sample. This was read as a governance-level condition, the unit that mandates escalation channels does not experience the friction those channels generate at lower levels, rather than a data gap. Triangulation tag: *Novel Insight* (the asymmetry was visible only when Oversight and Operations profiles were read against each other, not in either set of interviews alone).

2nd Line Risk and Compliance, EF: Common / VL: Illustrative / ICF: Ubiquitous: Formal independence was procedurally present but participants described their authority as contingent on business willingness to engage. Triangulation tag: *Divergent* (artefact evidence, ERCC minutes, recorded 2nd Line items as formally received, while interview evidence indicated influence was routinely absorbed rather than acted upon).

1st Line Operations, EF: Ubiquitous / FS: Ubiquitous / VL: Minimal: Every Operations participant referenced both escalation friction and feedback suppression; no participant described themselves as credibly authorised to raise governance-level concerns. Triangulation tag: *Convergent* (Peakon Freedom of Opinion score and interview accounts aligned on the same pattern of operational voice without governance reach).

Methodological Notes on Triangulation Logic

The three triangulation tags used throughout Chapter IV to classify the relationship between interview and artefact evidence are defined below. These tags are applied at the point of each dual-source evidence citation and are tracked in full across Appendix D.

Tag	Definition	Analytical function
Convergent	Interview and artefact evidence independently support the same pattern	Strengthens credibility; confirms the pattern is not a researcher artefact or participant perception only
Divergent	Interview and artefact evidence conflict, or formal records contradict frontline accounts	Identifies the discourse-practice gap; most analytically productive tag, the gap itself is the finding
Novel Insight	A pattern visible only in artefact evidence, not accessible through interview data alone	Surfaces structural dynamics that participant self-report cannot reach; supplements rather than validates

No single tag is evaluatively superior. Divergent evidence is not a weakness in the data, it is the primary location where formal governance and lived practice separate, which is central to the two-level pattern this study identifies.

G.7 Explanation-Building Framework (Core-Mechanism Symbiosis)

The Core-Mechanism Symbiosis framework distinguishes two analytical tiers that operate simultaneously throughout the findings chapter.

Tier	Codes	What they capture	How they are used
Contextual mechanisms (upper tier)	BRE, ICF, MRC, OAV, RC, RCF	Formal governance structures, policies, and escalation channels that officially exist in the organisation	Establish the governance landscape against which behavioural indicators are read; high mechanism discourse signals espoused intent

Proximate behavioural indicators (lower tier)	EF, FS, VL	What staff report actually experiencing, friction, suppression, legitimacy constraints	Reveal whether formal mechanisms are functioning as intended or have become decoupled from practice
Emergent synthesis	RE	The cumulative condition inferred from cross-unit EF/FS/VL configurations	Not coded directly; synthesised post-unit-profiling through cross-unit explanation-building (Appendix G.5)

Why the two-tier distinction matters analytically:

The absence of lived Escalation Friction within a specific unit, for example, the Oversight unit recording EF as procedural thoroughness rather than as a barrier (Table 4.4), is not a data gap. It is a governance-level condition: the unit that mandates escalation architecture experiences those same procedures as design features, not as the friction they generate at lower levels.

The 3.7:1 discourse-practice ratio (1,127 mechanism references: 301 proximate indicator references) quantifies this gap at the organisational level, formal mechanism discourse outpaces behavioural evidence of those mechanisms functioning by a factor of 3.7. This ratio is derived directly from the two-tier framework and is the primary quantitative expression of the Core-Mechanism Symbiosis pattern in Chapter IV.

G.8 Boundary-Case Discrimination Rules

The following table documents boundary-case discrimination rules applied where a single data segment exhibited features of more than one indicator, cross-reference with the codebook evolution in Appendix B.

Scenario	Primary code	Discrimination rule
A concern was raised but stalled in formal channels (delayed, rerouted, filtered by intermediary)	EF	The concern entered the system; friction describes what happened to it in transit
A concern was raised but received no acknowledgement, follow-up, or action	FS	The concern arrived but the response was absent or tokenistic
An individual stopped raising concerns after repeated non-response	FS (voice fatigue)	Code as FS when the participant attributes withdrawal to past suppression experiences; code as EF only if withdrawal is attributed to procedural burden or routing confusion
No concern was raised despite the individual recognising a risk	VL if attributed to perceived lack of authority; EF if	The participant's own attribution of why they stayed

	attributed to unclear channels; FS if attributed to anticipated non-response	silent determines the primary code
A concern reached governance forums but produced no coordinated response	EF + FS (dual code)	EF captures the transmission failure; FS captures the response failure, record co-occurrence
Dissenting views were sanitised in minutes despite being voiced in the meeting	EF (selective documentation)	The concern was in transit through the documentation process; the barrier is procedural

APPENDIX H: FIELDWORK TIMELINE AND DATA MANAGEMENT

This appendix consolidates the fieldwork timeline and data management records referenced in 3.8–3.9, providing the workflow schedule, sequencing dependencies, and storage, encryption, anonymisation, and retention protocols.

H.1 Research Workflow Timeline

This appendix documents the four-phase research workflow that structured fieldwork from October to December 2025. The schedule was designed to enable concurrent data collection and analysis, though coding commenced in Week 6 alongside ongoing transcription, rather than waiting until data collection was complete, thus reducing end-to-end compression and allowing saturation monitoring to inform both remaining interviews and emerging thematic development. Each phase is described in the body at 3.8.1–3.8.4; this table provides the consolidated operational record against which execution can be audited.

Phase	Key activities	Timeline	Key resources
Recruitment and Piloting	Coordination of access via sponsor liaison; pilot interviews with three participants; refinement of Interview Guide to V3.5 based on pilot feedback (Kallio et al., 2016).	Week 1	Sponsor liaison for participant access; pilot group (n = 3); reflexive memoing tools.
Data Collection	Twenty-five semi-structured interviews across job bands (S1-Oversight/Board), functions (Risk, Compliance, HR, Finance, Operations) and locations (Jersey, Isle of Man);	Weeks 2–5	Microsoft Teams/recorder; NVivo for preliminary coding; secure file

	collection of organisational artefacts (committee minutes and high-level risk governance documents); concurrent reflexive memoing (Kallio et al., 2016; Yin, 2009).		storage; structured memo templates.
Analysis	Transcription finalisation and quality checking; line-by-line NVivo coding; theme development via reflexive thematic analysis; indicator construction (EF, FS, VL) and cross-unit RE synthesis; saturation monitoring; explanation building (Braun and Clarke, 2006, 2019; Eisenhardt, 1989).	Weeks 6–9	NVivo software; codebook refinement; NVivo; structured reflexive memo templates (one per interview).
Write-up and Validation	Member checking with selected participants for interpretive accuracy; framework diagramming; integration of findings; supervisor review and feedback (Guba and Lincoln, 1994; Tobin and Begley, 2004).	Weeks 10–12	Participant availability for feedback; supervisor meetings; visualisation software.

The dependencies governing phase sequencing, in particular, the requirement that pilot completion precede main data collection, and that sufficient transcripts be available before NVivo coding could begin, are documented in Appendix H.2.

H.2 Dependencies and Critical Path

This appendix documents the three sequencing dependencies that governed phase transitions in the twelve-week schedule. These are not contingency notes but structural constraints built into the design: each dependency reflects a methodological requirement, guide validity, transcript sufficiency, and participant re-engagement, that determined when the next phase could begin. The narrative context for each dependency is provided in the body at 3.8.4; this appendix provides the consolidated record for audit purposes.

#	Dependency	Governed transition	Methodological basis	Resolved
---	------------	---------------------	----------------------	----------

1	Pilot completion and guide sign-off	Week 1 → Week 2 (main data collection)	Substantive revisions post-pilot would have invalidated comparability across the full 25-interview sample; guide locked at v3.5 following pilot feedback	Week 1, guide confirmed at v3.5
2	Minimum transcript base for saturation monitoring	Data collection → NVivo coding (Week 6)	Parent-theme saturation tracking required sufficient coded interviews before asymptotic patterns could be identified; five transcripts provided the baseline	Week 6, coding commenced with interviews 1–5 quality-checked
3	Participant re-engagement for member checking	Analysis → Write-up (Week 10)	Member checking required the same participants who contributed interview data to remain accessible; re-engagement was managed through the sponsor liaison relationship established in Week 1	Weeks 10–12, confirmations received from selected participants

Data storage, encryption, and file management protocols that operated across all four phases are documented in Appendix H.3.

H.3 Data Storage, Access and Encryption

This appendix documents the data storage, access, and encryption arrangements that applied across all four fieldwork phases. These arrangements were not post-hoc administrative decisions but design constraints agreed prior to fieldwork commencement, consistent with institutional ethics approval and the organisational confidentiality conditions governing artefact access. The narrative context is provided in the body at 3.9.2; this appendix provides the consolidated operational record. Storage arrangements for specific data types, transcripts, NVivo project files, artefacts, and reflexive memos, are described in the rows below.

Data type	Storage location	Encryption	Access	Retention
Audio recordings	Encrypted personal device	AES-256 device encryption	Researcher only	Deleted following transcript quality confirmation

Teams auto-transcripts	Downloaded locally; deleted from Teams	Device encryption	Researcher only	Within 90-day Teams retention window
Cleaned transcripts	Password-protected folder, encrypted device	AES-256	Researcher only	Duration of study; secure deletion post-submission
NVivo project file	Encrypted device with daily backup to encrypted external drive	AES-256	Researcher only	Duration of study; secure deletion post-submission
Organisational artefacts	Password-protected folder, separate from transcript folder	AES-256	Researcher only	Paraphrased in reporting; originals returned or deleted per organisational protocol
Reflexive memos	NVivo memo function within project file	As NVivo project file	Researcher only	Duration of study

H.4 Project Chronology Table

Phase	Activity	Dates
Access & recruitment	CRO negotiation; invitation dispatch	Aug-Sep 2025
Pilot interviews	Three pilot interviews; guide lock at v3.5	W/c 13 Oct 2025
Main fieldwork (interviews)	25 main-study semi-structured interviews	21 Oct – 12 Nov 2025
Concurrent coding	NVivo coding commenced alongside transcription	24 Nov 2025 (from interview 1–5 base)
Post-fieldwork code refinement	Child-level elaboration; codebook v0.7	1–18 Dec 2025
Member feedback	Selective participant review of emerging themes	Dec 2025 – Jan 2026
Write-up	Draft and supervisor review	Jan–Apr 2026

APPENDIX I: PARTICIPANT REFERENCE

I.1: Participant Reference Counts by Theme

Participant Data											
	Balancing risk (BRE)	Escalation Friction (EF)	Feedback Suppression (FS)	Ind. of Control Functions (ICF)	Mechs. for Raising Concerns	Open to Alt Views (OAV)	Reinforcement of Challenge (RC)	Reps. of Control Functions	Risk Entropy Synthesis (RE)	Voice Legitimacy (VL)	Totals
1 : MS01	4	11	15	19	19	16	23	23	8	11	149
2 : MS02	0	3	3	1	7	3	8	3	6	8	42
3 : MS03	0	4	2	6	14	7	22	17	4	6	82
4 : MS05	0	2	1	4	10	14	38	21	2	22	114
5 : MS06	9	8	2	0	19	21	29	15	16	8	127
6 : MS08	0	2	0	0	3	12	12	0	2	11	42
7 : MS09	1	2	1	2	7	5	6	6	4	6	40
8 : MS10	1	0	2	3	16	6	4	3	5	4	44
9 : MS12	0	4	0	3	12	8	17	7	5	5	61
10 : MS15	1	3	0	0	3	2	8	3	6	3	29
11 : MS16	0	4	0	0	9	4	8	4	6	4	39
12 : MS20	1	6	0	2	10	8	22	6	8	6	69
13 : MS21	8	3	0	7	5	5	28	5	7	5	73
14 : MS23	4	3	0	5	3	8	17	6	5	8	59
15 : MS29	3	1	0	0	4	6	17	5	5	3	44
16 : MS37	6	3	0	4	5	10	21	7	15	5	76
17 : MS39	6	5	0	3	11	10	20	11	11	13	90
18 : MS41	7	7	9	9	8	15	38	9	12	8	122
19 : MS44	0	1	1	2	5	6	10	2	1	6	34
20 : MS45	4	3	0	7	10	3	17	6	0	4	54
21 : MS46	4	10	5	7	6	8	26	9	11	9	95
22 : MS47	5	3	1	3	6	4	11	3	11	5	52
23 : MS48	2	0	1	3	2	1	3	1	6	0	19
24 : MS49	0	5	0	2	2	4	1	1	2	1	18
25 : MS50	0	2	0	1	2	3	2	0	2	2	14
Totals	66	95	43	93	198	189	408	173	160	163	

