



Prifysgol Cymru
Y Drindod Dewi Sant
University of Wales
Trinity Saint David

Programme	Master of Business Administration (BMMB7003D)
Student	Atapattu Wasala Mudalige Lasindu Sharlin Ratnayake (2031417)
Supervisor	Dr Laura James, PhD
Dissertation	Assessing Internal Audit's Contribution to ERM Effectiveness in Financial Institutions
Submission date	16 May 2026

DECLARATION:

This work has not previously been accepted in substance for any degree and is not being concurrently submitted in candidature for any degree.

Signed: **Atapattu Ratnayake**

Date: **16 May 2026**

STATEMENT 1

This work is the result of my own investigations, except where otherwise stated. Other sources are acknowledged by footnotes giving explicit references. A bibliography is appended.

Signed: **Atapattu Ratnayake**

Date: **16 May 2026**

STATEMENT 2

I hereby give consent for my work, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed: **Atapattu Ratnayake**

Date: **16 May 2026**

Acknowledgements

To my dissertation supervisor, Dr Laura James, whose patience, detailed guidance, and unwavering support throughout this study made the journey both possible and rewarding. Her constructive feedback and commitment to this research are deeply appreciated.

To the eleven participants who generously spared their time and shared a wealth of professional knowledge and experience, this study would not exist without your candour and willingness to contribute. Your insights form the foundation of everything that follows; I am truly grateful for your support.

To Jon, a thought leader and pillar of guidance in the risk and audit profession, whose knowledge, integrity, and professional approach represent the standard to which I aspire. Thank you for the inspiration.

To Almighty God, and His Son Jesus, by whom I am saved and blessed.

To my beloved wife Betsy and baby daughter Anayah: God's greatest gifts to me. You light up my life. Your love and encouragement through many gruelling days and nights of work and study, your steadfast support at every turn, without which none of this would have been possible. This dissertation is as much yours as it is mine.

Abstract

Within the contested terrain of contemporary Australian financial-services governance, this study evaluates Internal Audit's (IA) contribution to Enterprise Risk Management (ERM) effectiveness through practitioner narratives and artefacts, addressing a gap in the existing literature, which has tended to treat IA's contribution at the level of structural attribute, namely independence, reporting lines, and charter codification, without adequately specifying the operational mechanisms that translate structural arrangement into governance impact. Four research questions are addressed, namely how IA's role and boundaries are described by practitioners, which IA practices are most impactful, what organisational conditions enable or constrain IA's influence, and how sector-specific factors shape its contribution.

Adopted for this study is a qualitative multi-case design, with eleven semi-structured interviews conducted across Australian banks, non-bank lenders (NBLs), and the prudential regulator. Participants were drawn from senior practitioner roles including Chief Risk Officers, Chief Audit Executives, audit committee chairs, prudential supervisors, and senior compliance and finance leaders. The analysis followed the Gioia methodology, with first-order codes aggregated into second-order themes and three theoretical dimensions.

The central analytical contribution is a three-dimension framework. Anchored Authority captures the anchoring of audit work to the Risk Appetite Statement and the systemic conception of independence required to protect that anchoring. Disciplined Engagement captures the conditional advisory boundary, end-to-end information quality work, and evidence-grounded governance challenge. Enabling Context captures institutional capability investment and sector-specific regulatory conditions. Specified by the framework are mechanisms which the existing literature has tended to leave implicit, including the conditional advisory principle, trust as the operational medium through which IA's findings convert into governance impact, and the systemic, continuously-contestable conception of independence.

For institutions, regulators, and IA practitioners, the unifying implication is that the effectiveness of IA's contribution to ERM is a system property requiring active maintenance across multiple conditions rather than a stable institutional feature established by structural arrangement.

Keywords: Internal Audit; Enterprise Risk Management; financial services governance; Risk Appetite Statement; Gioia methodology; Three Lines Model; information quality; Australian financial services; non-bank lenders; independence; qualitative research

Contents

Abstract	1
1. Introduction	10
1.1 Research context	12
1.2 Research aim, objectives, and questions	13
1.3 Significance	13
2. Literature Review.....	15
2.1 Background	15
2.2 The Evolution of ERM: From Compliance to Strategic Capability	15
2.3 Measuring ERM Effectiveness: A Persistent Challenge	17
2.4. Internal Audit's Role in ERM: Mandate, Boundaries, and Influence.....	18
2.4.1 The Assurance–Advisory Tension.....	18
2.4.2 Organisational Drivers of Effectiveness.....	19
2.4.3 Capability as a Constraint	19
2.4.4 Information Quality as a Mechanism of ERM Effectiveness	20
2.4.5 Governance Challenge and Risk-Informed Action	21
2.4.6 Sector-Specific Dynamics: Banks and Non-Bank Lenders	22
2.4.7 Methodological Gaps and the Case for Qualitative Inquiry.....	23
2.4.8 Synthesis and Research Justification	24
3. Methodology	25
3.1 Research Philosophy	25
3.2 Qualitative Multi-Case (Organization) Design	26
3.3 Organisation Selection	27
3.4 Participants and Sampling	28
3.5 Data Collection	29
3.6 Data Management and Security.....	30

3.7 Data Analysis	31
3.8 Researcher Positionality	32
3.9 Limitations of the Research.....	33
3.10 Ethical Considerations.....	35
4. Findings.....	37
4.1 Introduction	37
4.2 Theme 1: Anchoring IA in the Risk Appetite Statement.....	37
4.3 Theme 2: Independence as structural, behavioural, and continuously contestable	38
4.4 Theme 3: Advisory has legitimate scope but a sharp ownership boundary	41
4.5 Theme 4: Information quality is the operational core of IA's ERM contribution....	43
4.6 Theme 5: Effective governance challenge is precise, evidence-grounded, and tone- aware.....	44
4.7 Theme 6: Capability is the binding constraint on contemporary IA contribution..	47
4.8 Theme 7: Sector and regulatory context shape the substance and rhythm of IA work	49
5. Discussion.....	53
5.1 Introduction	53
5.2 Aggregate Dimension 1: Anchored Authority.....	54
5.3 Aggregate Dimension 2: Disciplined Engagement	55
5.4 Aggregate Dimension 3: Enabling Context	57
5.5 Longitudinal contrast and comparative case	59
5.6 Implications.....	61
5.6.1 Institutions and boards	61
5.6.2 Regulators	62
5.6.3 IA practitioners.....	62

6. Conclusion	64
6.1 Research aim, objectives, and questions revisited	64
6.2 Key findings and contribution	64
6.3 Limitations.....	66
6.4 Recommendations for future research.....	66
6.5 Concluding reflection	67
7. References	69
Appendix A – Ethics and Consent	74
A1. Participant Information Sheet	74
A2. Consent Form	77
A3. Ethics Approval	79
Appendix B — Sample and Participants	80
B1. Participant Table	80
B2. Organisation Selection Rationale	83
Appendix C — Data Collection	86
C1. Interview Guide.....	86
Appendix D — Data Analysis Artefacts.....	93
D1. Codebook	93
D2. Data Structure Diagram.....	99
D3. Sample Coded Transcript Excerpt	103
D4. Saturation Analysis	106
D5. Analytical Decision Log	109
D6. Reflexive Memos.....	115
Appendix E — Researcher Positionality.....	118
E1. Researcher Positionality Statement.....	118

List of Abbreviations

The following abbreviations are used throughout this dissertation. Each is introduced in full at first use in the relevant chapter and thereafter used in shortened form.

Abbreviation	Full Term
AI	Artificial Intelligence
AI/ML	Artificial Intelligence / Machine Learning
ALCO	Asset and Liability Committee
AML	Anti-Money Laundering
APRA	Australian Prudential Regulation Authority
ASIC	Australian Securities and Investments Commission
BCBS	Basel Committee on Banking Supervision
BCBS 239	Basel Committee on Banking Supervision Principles for Effective Risk Data Aggregation and Risk Reporting
CAE	Chief Audit Executive
CAMELS	Capital Adequacy, Asset Quality, Management, Earnings, Liquidity, Sensitivity
CFC	Chief Financial Controller
COSO	Committee of Sponsoring Organisations of the Treadway Commission
CPS	Prudential Standard (APRA)
CPS 220	APRA Prudential Standard: Risk Management
CPS 230	APRA Prudential Standard: Operational Risk Management
CPS 234	APRA Prudential Standard: Information Security
CRO	Chief Risk Officer
ECB	European Central Bank
ERM	Enterprise Risk Management
EU	European Union
IA	Internal Audit / Internal Auditor
IIA	Institute of Internal Auditors
ISO	International Organization for Standardization
MBA	Master of Business Administration
MI	Management Information
NBL	Non-Bank Lender
NHMRC	National Health and Medical Research Council
QA	Quality Assurance
QDA	Qualitative Data Analysis

RAG	Red, Amber, Green (status reporting)
RAS	Risk Appetite Statement
RG	Regulatory Guide (ASIC)
RG 209	ASIC Regulatory Guide: Credit Licensing — Responsible Lending Conduct
RTO	Recovery Time Objective
SQL	Structured Query Language
UK	United Kingdom
US	United States
UWTSD	University of Wales Trinity Saint David

1. Introduction

1.1 Research context

Within the conventional three-lines model of financial-services governance, Internal Audit (IA) occupies a structurally distinctive position as the third line of defence, charged with providing independent assurance over Enterprise Risk Management (ERM), the management control systems, and the broader governance framework. The relationship between IA and ERM is, in practice, more contested than the structural conception suggests, with practitioners disagreeing about how much advisory engagement is appropriate before assurance is compromised, regulators expecting IA to anticipate emerging risks rather than catch up to them after the fact, and the technical complexity of contemporary financial-services risk having progressively outstripped the traditional accounting orientation of the IA workforce. At the level of structural attribute, namely independence, reporting lines, and charter codification, the academic literature has tended to treat IA's contribution well, however, the operational mechanisms that translate structural arrangement into governance impact have been less adequately specified.

Particularly informative is the Australian financial-services environment as a setting for examining this contested relationship. Expectations have been progressively tightened by APRA, the prudential regulator, through CPS 220, CPS 230, and CPS 234, while consumer-protection expectations have been reinforced by ASIC, the conduct regulator, through RG 209. For banks, the pressures include stress testing, BCBS 239 data aggregation, operational resilience, and model risk governance. Operating outside APRA's prudential perimeter, non-bank lenders (NBLs) nevertheless face investor scrutiny that frequently exceeds regulatory minimums. The sectoral diversity arising from these conditions creates the basis for a multi-case study capable of surfacing both the common framework of effective IA contribution and the sector-specific variations within it.

1.2 Research aim, objectives, and questions

The aim of this study is to evaluate IA's contribution to ERM effectiveness through practitioner narratives and artefacts. Four objectives are addressed: to explain how IA defines and operationalises ERM assurance; to identify the enablers and barriers (mandate, sponsorship, access, capability) that shape that contribution; to examine IA's influence on information quality and governance challenge; and to develop a grounded conceptual model of IA contribution across banks and NBLs.

To form a definitive outcome, four research questions will therefore guide this study:

RQ1: How do IA professionals and risk executives describe IA's role and boundaries in ERM governance?

RQ2: Which IA practices are perceived as most impactful for ERM effectiveness in retail banking and finance?

RQ3: What organisational conditions enable or constrain IA's influence on the quality of risk information and governance challenge?

RQ4: How do sector-specific factors (e.g. stress testing in banks; securitisation and warehouse reporting and covenants in finance companies) shape IA's ERM contribution?

1.3 Significance

The study contributes in related, yet distinct ways, for academic and practitioner discourse. Academically, the existing IA-ERM literature is extended through the specification of operational mechanisms, namely the conditional advisory boundary, trust as the medium of IA's contribution, and the systemic conception of independence, mechanisms which the existing accounts have tended to leave implicit. From a practitioner perspective, a framework is provided to institutions, regulators, and IA practitioners against which the effectiveness of IA's contribution to ERM can be evaluated and improved. Drawing on senior practitioners across banks, NBLs, and the

prudential regulator, the empirical material produces a multi-perspective view of the IA-ERM relationship in contemporary Australian financial services.

Across the following six chapters covering the literature review, methodology, findings, discussion, and conclusion, the above study is hereby presented.

2. Literature Review

2.1 Background

Enterprise Risk Management (ERM) has matured considerably as a discipline over the past two decades, moving from a compliance exercise built solely around risk registers and policies, and toward a strategic capability that should inform how financial institutions price, lend, allocate capital, and sustain operations under stress. Pivoting away from simply testing and treating for compliance, modern industry standard frameworks such as COSO ERM (2017) and ISO 31000 (2018) standardise governance, culture, and the quality of risk information as foundations for effective ERM. Internal Audit (IA) is positioned within the Three Lines Model (IIA, 2020) as an independent assurance function. IA is widely acknowledged as a contributor to ERM, but in practice, the nature, extent, and measurement of that contribution vary across institutions and sectors, with a notable gap in existing research to understand these nuances.

This review examines the key themes bearing on IA's role in ERM effectiveness in financial institutions, and more specifically within the Australian financial regulatory environment: the evolution of ERM from compliance to strategy, the persistent difficulty of measuring ERM outcomes, the factors that enable or constrain Internal Audit's influence, the centrality of information quality and governance challenge as mechanisms of effectiveness, and the sector-specific dynamics which distinguish Australian retail banks from non-bank lenders (NBLs). The research paper builds the case for a qualitative, multi-case investigation into the mechanisms through which Internal Audit contributes to ERM in practice, and the conditions impacting these mechanisms' effectiveness.

2.2 The Evolution of ERM: From Compliance to Strategic Capability

Until the 1980s (and in some emerging markets still), risk management in financial institutions was organised in silos, with credit risk, market risk, and operational risk sitting in separate teams, with separate tools and separate reporting lines. This arrangement persisted until being proven dangerously insufficient by the 2007-2008

global financial crisis, which exposed the inability of major banks to aggregate and report their risk exposures with any speed or accuracy (BCBS, 2013). Forced by the crisis into a fundamental rethink of risk management design, the industry adopted ERM as the successor model, promising an integrated and organisation-wide approach to identifying, assessing, treating, and managing the full range of risks a firm faces (COSO, 2017; ISO, 2018).

The updated COSO ERM framework (2017), titled “Enterprise Risk Management—Integrating with Strategy and Performance”, was a notable tangent from the preceding controls-and-compliance-focused version, entwining ERM in the organisation’s strategic drive by being built around five aspects: governance and culture, strategy and objective-setting, performance, review and revision, and information, communication, and reporting. Complementing this design methodology, ISO 31000 (2018) positioned risk management as a value-creating activity grounded in integration, synergy, and continual improvement, shifting the focus from whether an organisation has ERM structures to whether those structures actually improve decisions and outcomes.

Despite this design evolution, existing research argues that many organisations have struggled to move beyond checkbox compliance. Power (2007; 2009) has argued that the institutionalisation of risk management frequently produces rituals of verification such as auditable artefacts and process documentation that substitute for genuine engagement with uncertainty, describing this phenomenon as the “risk management of nothing,” where frameworks exist in form but do not materially influence behaviour or decisions. The critique points to a gap between the structural existence of ERM and its functional effectiveness, a gap that several researchers have since documented (Mikes, 2009; Paape and Speklé, 2012; Adam, Soliman and Mahtab, 2023).

2.3 Measuring ERM Effectiveness: A Persistent Challenge

One of the more challenging aspects of ERM effectiveness is measuring its usefulness and material contribution. Prior empirical work indicates that ERM adoption has often been treated as a binary event, with proxies such as the appointment of a Chief Risk Officer or credit-rating agency recognition used to indicate whether ERM is “in place” (Hoyt and Liebenberg, 2011; Beasley, Clune and Hermanson, 2005). This approach oversimplifies an instrumental, detailed process that unfolds across a complex space, as Lundqvist (2014) notes when advocating for continuous metrics that capture the scale of ERM integration rather than its mere presence. Building on this, Adam, Soliman and Mahtab (2023) developed a multifaceted measurement model for banking that combines ERM components with the CAMELS performance framework to better reflect the sector-specific nature of risk management outcomes. Significant variation in ERM adoption and design across institutions, as Paape and Speklé (2012) noted, further reinforces the argument that a one-size-fits-all assessment cannot adequately measure ERM effectiveness.

Empirical research has also been observed to rely on maturity scores or self-reported survey data to gauge ERM effectiveness, instruments which tend to capture perceptions of design completeness rather than evidence of material impact on decisions made. Kaplan and Mikes (2012) argue that risk management practices need to be matched to the nature of the preventable, strategic, and external risks involved, and that compliance-oriented approaches are unsuitable for effectively managing strategic and external risks. The implication is that conventional maturity assessments may not be measuring the necessary aspects of the ERM framework, scoring the architecture of controls rather than the quality of risk-informed action. This study attempts to address that gap by defining ERM effectiveness in terms of decision-useful information and observable, risk-informed actions, and by carrying out a qualitative assessment of evidence that can examine material outcomes for a more complete view on ERM effectiveness and, by design, the contribution that IA makes to it.

2.4. Internal Audit's Role in ERM: Mandate, Boundaries, and Influence

2.4.1 The Assurance–Advisory Tension

The IIA's Three Lines Model (IIA, 2020) positions IA as an independent assurance provider, distinguishing first-line functions (which own and manage risk) and second-line functions (which oversee it). Explicit on this point, the IIA's paper on IA's role in ERM (IIA, 2009) states that the function's fundamental objective is to provide independent, objective, and unbiased assurance on the effectiveness of risk management processes, with any extension into advisory work treated as a consulting engagement with safeguards to protect independence.

In practice, though, there is evidence of the boundary between assurance and advisory blurring. Drawing on qualitative interviews with Chief Audit Executives in the United States and Belgium, Sarens and De Beelde (2006) found that internal auditors consistently contributed to the development and maturation of risk management systems through advisory interactions, helping risk owners build awareness, formalise processes, and identify control gaps. This pioneering role was especially pronounced in organisations where risk management was still relatively immature. The finding highlights that advisory engagement can accelerate organisational capability but can also compromise the independence that makes assurance credible. How the IA function navigates this conflict, and how boards and committees respond to it, also remains under-explored in historical research.

2.4.2 Organisational Drivers of Effectiveness

By surveying 153 Italian companies, Arena and Azzone (2009) identified three categories of organisational drivers that determine IA effectiveness: the characteristics of the audit team (size, professional affiliation, competence), the processes and activities employed (such as control risk self-assessment), and the organisational links between IA and other governance actors, particularly the audit committee. Their work underscores that effectiveness is not a property of the audit function in isolation but a product of its organisational positioning and the relationships it maintains. When the audit committee is actively involved in directing IA's work, effectiveness increases; when internal auditors hold professional affiliations and apply structured methodologies, effectiveness increases. These are relational and contextual factors rather than technical ones, and difficult to capture through surveys alone.

Previously, Beasley, Clune and Hermanson (2005) had also found that the extent of ERM implementation was associated with board and senior management support, the presence of a CRO, and industry type, with financial institutions and regulated industries tending toward more developed ERM programmes. The implication for IA is that its ability to contribute depends partly on the organisational environment in which it operates, specifically the degree of sponsorship it receives and the access it is granted to risk data, systems, and governance forums. Where sponsorship is weak or access restricted, even a technically capable audit function may struggle to influence ERM outcomes. Reinforcing this, a recent review of 77 studies published between 2004 and 2024 classified IA's ERM involvement into core assurance roles, legitimate consulting roles with safeguards, and activities that cross into management responsibility, and highlighted the need for more behavioural and qualitative research to understand how internal auditors exercise judgment in risk assessment, technology use, and reporting (Nkansa, Barr-Pulliam and Walker, 2025).

2.4.3 Capability as a Constraint

With ERM growing more technically sophisticated, IA teams are consulted to evaluate complex models, stress-testing methodologies, data pipelines, and algorithm-driven decision processes. Arena and Azzone (2009) noted that the shift from compliance to risk

management demands auditors who can understand business drivers, engage with diverse sources of risk, and work collaboratively with line management. Where these capabilities are lacking, IA risks being marginalised or confined to routine compliance testing, regardless of the formal mandate it holds: particularly relevant in financial institutions where attributes such as model risk, data lineage, and aggregation integrity are central to risk governance. This identified capability gap, and its consequences for Internal Audit's influence on ERM, is a thread this study intends to follow through its analysis.

2.4.4 Information Quality as a Mechanism of ERM Effectiveness

A central argument of this review is that information quality serves as a critical mechanism linking ERM design to ERM effectiveness. The BCBS 239 principles (BCBS, 2013), issued after the global financial crisis, set explicit expectations for banks regarding the accuracy, completeness, timeliness, and adaptability of risk data aggregation and reporting. The rationale is straightforward: decision-makers cannot govern risk effectively if the information they receive is inaccurate, late, or difficult to interpret. Data quality is not a technical concern to be delegated to IT; it is a governance imperative that shapes the quality of strategic and operational decisions.

Compliance with BCBS 239 has proven remarkably difficult. Progress reports published by the Basel Committee have consistently found that the majority of globally systemically important banks have not achieved full compliance, with weaknesses concentrated in governance, data architecture, timeliness, and the clarity of risk reports (BCBS, 2013; Forvis Mazars, 2024). The European Central Bank's 2018 thematic review concluded that none of the 25 significant institutions it examined had fully implemented the principles, identifying unclear accountability for data quality as a recurring weakness (ECB, 2018). These findings are striking, in that, even in heavily regulated banking environments with well-resourced second and third lines, the translation of data quality standards into consistent organisational practice seemingly remains incomplete.

Theoretically, IA is well placed to address these shortcomings. By testing data lineage, validating reconciliations, challenging exception-handling processes, and assessing the explainability of model outputs, IA can improve the reliability of risk information reaching

committees and decision-makers. Yet, literature offers little empirical evidence on how these practices improve governance behaviour. Studies that invoke information quality do so at the level of framework prescription rather than lived organisational experience. There is a gap between knowing that information quality matters and understanding how IA practices improve it in a way that changes what committees do; a gap this study addresses directly.

2.4.5 Governance Challenge and Risk-Informed Action

Governance effectiveness depends not only on the availability of quality information but on the willingness and capacity of committee members and executives to challenge that information, ask probing questions, and convert deliberation into concrete action. Power (2004; 2007) has written extensively on the tendency of risk management to become symbolic—a performance of accountability that substitutes for substantive engagement with uncertainty. When committees receive dashboards and risk reports but do not interrogate the assumptions behind the numbers or follow through on identified weaknesses, ERM exists in form but not in function. The quality of challenge in these settings is what separates a governance process that drives action from one that merely documents activity.

Mikes (2009), in the study of two large banking organisations, distinguished between different “calculative cultures” that shape how risk information is used. In one institution, risk models were treated as accurate representations of economic reality and used diagnostically to drive capital allocation decisions—a logic of “quantitative enthusiasm.” In the other, risk figures were regarded as trend indicators that needed to be complemented and sometimes overridden by senior managerial judgment—a logic of “quantitative scepticism.” In the sceptical culture, qualitative discussion and expert judgment played a more prominent role in committee deliberation, whereas in contrast, the enthusiastic culture sees that the numbers themselves were expected to drive action.

The point is that ERM does not operate in a cultural vacuum, and Internal Audit’s ability to strengthen governance challenge is shaped by the prevailing attitudes toward risk information within the institution. IA can contribute to governance challenge by ensuring

that the questions raised in committee settings are grounded in evidence, that decisions and their rationale are documented, and that action items are tracked to closure; contribution which sits between information quality and risk-informed action in the conceptual pathway this study proposes.

2.4.6 Sector-Specific Dynamics: Banks and Non-Bank Lenders

The way IA contributes to the ERM in the financial services landscape varies between banks and non-bank lenders, given the defined nuances within most mature financial services industry global regions. Banks operating under prudential supervision face extensive regulatory obligations that shape both ERM design and audit priorities. In the Australian context, APRA administers standards including CPS 220 (Risk Management), CPS 234 (Information Security), and CPS 230 (Operational Risk Management, effective 1 July 2025), which establish expectations for governance structures, risk appetite articulation, data security, and business continuity (APRA, 2019; 2023; 2025). BCBS 239 adds a further layer of expectation around risk data aggregation and reporting. IA assurances in banks therefore tend to focus on stress-testing governance, expected credit loss model validation, risk data lineage, and capital adequacy reporting; all areas where the quality of information directly affects prudential outcomes.

NBLs operate in a different risk environment. Reliant on securitisation, warehouse funding, and investor commitments rather than deposits, these institutions face risks that centre on the integrity of trustee reporting, pool eligibility determinations, waterfall calculations, and covenant monitoring. A failure in any of these areas can trigger funding instability, investor withdrawal, or breach of warehouse facility terms. Internal Audit in non-bank lenders therefore plays a critical role in verifying the accuracy and completeness of information flows that underpin investor and creditor confidence. The regulatory footprint is lighter (ASIC's RG 209 on responsible lending applies, along with general obligations under Australian financial services licensing), but the commercial stakes around data and reporting integrity are no less significant (ASIC, 2019).

Despite these differences, the literature has tended to study IA's ERM role within single institutional settings, with comparative studies that examine how regulatory intensity, funding models, and risk profiles moderate IA's contribution remaining rare.

2.4.7 Methodological Gaps and the Case for Qualitative Inquiry

The preceding review reveals a body of literature characterised by methodological tendencies that limit the depth of understanding available to scholars and practitioners. First, survey-based designs dominate; surveys can identify statistical associations between key attributes (such as that between audit committee engagement and IA effectiveness) but cannot illuminate the social processes through which those associations operate (Arena and Azzone, 2009). How a supportive audit committee actually influences the way auditors frame their findings, or how sponsorship translates into access and access into impact, are interpretive questions that require qualitative methods. Second, the literature has concentrated on the existence and design of ERM structures rather than their functional effects on decision-making; studies that assess whether an institution has adopted ERM or score its framework maturity tell us little about whether risk information actually influences pricing decisions, limit adjustments, portfolio steering, or remediation prioritisation. Third, the role of context has been acknowledged conceptually but under-explored empirically. Paape and Speklé (2012) demonstrated significant cross-institutional variation in ERM adoption and design; Mikes (2009) showed how calculative cultures shaped the use of risk information within individual banks. Yet few studies have systematically compared how different organisational and sectoral contexts moderate IA's influence on ERM outcomes, and a multi-case design guided by the methodological principles of Yin (2018) and Gioia, Corley and Hamilton (2013) is specifically constructed to address this gap.

2.4.8 Synthesis and Research Justification

The case for this research rests on several converging observations. ERM has evolved from a compliance-driven checklist into a strategic capability, but evidence suggests many financial institutions have not completed this transition in practice. The measurement of ERM effectiveness remains contested, with most empirical studies relying on structural or perceptual proxies rather than examining the decision-making outcomes that COSO (2017) and BCBS 239 (BCBS, 2013) prioritise. IA is widely recognised as a key governance actor, but its influence is contingent on organisational factors (mandate clarity, sponsorship, access, capability) that vary markedly across institutions and sectors. Information quality and governance challenge represent critical but under-researched mechanisms linking ERM design to ERM effectiveness, and the literature lacks empirical work that traces how IA practices improve these mechanisms and under what conditions. The literature also reveals a methodological need: survey-based and maturity-assessment studies are poorly suited to capturing the social dynamics of governance such as the framing of questions, the quality of debate, the documentation of decisions, and the follow-through on actions. This study therefore responds to a clear gap at the intersection of internal audit research, ERM scholarship, and financial services governance, aiming to move the conversation beyond whether IA matters for ERM to how it matters, for whom, and under what conditions.

3. Methodology

3.1 Research Philosophy

For a study on how Internal Audit (IA) influences Enterprise Risk Management (ERM) effectiveness, a philosophical approach is useful, because central constructs of the field such as information quality, governance challenge, and risk-informed action are not objective and measurable indicators that can be rated outside of the people who produce and interpret them.

This study therefore adopts a relativist ontology, where, although concepts such as ERM effectiveness, governance challenge, and the quality of risk information do not exist as fixed, externally verifiable states, they are still maintained by perception and commentary aspects: through the way an industry practitioner would interpret risk data, frame questions in committee settings, negotiate mandate boundaries, and decide what counts as a satisfactory response to an identified weakness. Institutions with identical ERM frameworks on paper may produce differing governance outcomes in ways that a structural assessment cannot capture, caused by the people, cultures, and power dynamics that drive the frameworks. Understanding ERM effectiveness therefore requires understanding how practitioners experience and interpret it within their own institutional contexts, not only measuring whether certain structures are present.

An interpretivist epistemology holds that meaning is not discovered but co-created through the interaction between researcher and participant, and that the researcher's role is to surface patterns, mechanisms, and explanations that make sense of how governance processes work in particular settings (Creswell and Poth, 2018). This does not mean the findings are purely subjective; rather, they are contextualised, transparent about the interpretive process, and open to scrutiny through a documented audit trail of analytical decisions (see Appendix D5).

A quantitative methodology was also considered. The existing literature on IA and ERM is well supplied with survey-based studies that measure statistical associations between variables such as audit committee engagement, CAE tenure, ERM maturity scores, and organisational performance indicators (Arena and Azzone, 2009; Beasley, Clune and

Hermanson, 2005; Paape and Speklé, 2012). These studies have been valuable in establishing that certain structural and relational factors are associated with IA effectiveness and ERM adoption. However, a common limitation is evident: they state that associations exist, but neither how nor why they operate in practice. A survey can report that audit committee involvement correlates to IA effectiveness, but cannot explain what that involvement looks like day to day, how it changes the way auditors frame their findings, or what happens in the committee room when the involvement is absent.

The research questions driving this study are open-ended, “how” and “under what conditions” questions about meaning, process, and context rather than frequency or magnitude. Quantitative design would require operationalising constructs like governance challenge and information quality into Likert-scale items, which would flatten exactly the complexity that makes these aspects consequential, and would require a sample size large enough for statistical inference that is difficult to achieve when the target population is senior governance practitioners in a specific sector and geography. A mixed-methods design was also considered, but as the analytical weight of the study rests on the qualitative case work, a qualitative-only approach was adopted. The overall design is consistent with the “research onion” framework of Saunders, Lewis and Thornhill (2019), which conceptualises the design process as nested decisions moving outward from underlying philosophical assumptions through to the specific data collection techniques used.

3.2 Qualitative Multi-Case (Organization) Design

Given the interpretivist stance and the nature of the research questions, a qualitative multi-case methodology is adopted. Yin (2018) states that a qualitative case study, research is appropriate when the research questions seek to explain how or why a phenomenon occurs within a real-world context, and when the boundaries between phenomenon and context are not clearly delineated. For this research paper, the methodological concept of a case study is applied to an institution, or organization.

Both conditions apply here. Internal Audit’s influence on ERM effectiveness is not a discrete, bounded event that can be isolated from its organisational setting. It is

embedded in relationships, cultures, governance routines, and institutional histories. Separating the mechanism from the context would strip away the very factors that determine whether it works.

Thereafter, a multi-organisation design is preferred over the single-organisation approach, because the research questions explicitly seek to compare how Internal Audit contributes to ERM across different institutional settings in banks and non-bank lenders. Yin (2018) describes this as replication logic: each case (or for the purpose of this paper, organisation) is selected not for statistical representativeness, but for its ability to confirm, extend, or challenge the explanatory propositions that are seen within the sampled cohort. Eisenhardt and Graebner (2007) similarly argue that multiple-case designs enable theory-building through pattern recognition across settings. If noticeable patterns appear across organisations with different regulatory and funding contexts, the finding is considerably more robust than if it appeared in only one. If there appear to be marked differences, those differences are themselves analytically valuable because they reveal the contextual conditions under which Internal Audit's contribution changes. The target is twelve to fifteen interviews with industry practitioners – a range that balances depth and comparability; fewer than ten would limit the potential for cross-organisation synthesis, while more than fifteen would affect the time available for analysis, given the time constraints of the research paper, and the practical difficulty of securing access to a wider participation base of industry practitioners.

3.3 Organisation Selection

Organisations used for research were selected through criterion-based sampling with maximum variation across sector and operational scale, moderated by participant availability (Patton, 2015). Selection was attempted to best represent institution types (bank versus NBL), size and complexity (regional versus major), and, where possible, comparable institutions across other geographic locations (one multinational financial services organisation was also selected, as it operates in both Australian and global markets). The rationale is that if IA's contribution to ERM operates through consistent mechanisms across varied contexts, those mechanisms are likely robust and transferable; if they differ, the variation helps explain why.

Confidentiality is a practical constraint in governance research, as institutions are understandably cautious about sharing committee minutes and audit reports even when redacted. The research design accommodates this by focusing on non-confidential interview evidence, audit reports, committee action logs, and ERM dashboard templates, and by committing to strict anonymisation of institutions and participants in all outputs.

One participant was included as a deliberate comparative case. This individual holds a financial-services audit background but has since transitioned into a different industry, and was selected to provide a contrasting lens on how IA contributes to organisational objectives in a non-financial-services setting. This participant's account is treated separately from the primary thematic analysis and used in the discussion chapter to contextualise the sector-specific dynamics of the core sample, consistent with the comparative analytical orientation of the multi-case design (Stake, 2006). The comparative case was coded separately against its own analytical lens, with the resulting artefact retained in the analytical audit trail (see Appendix D5, Decision 001).

This sampling logic is consistent with Miles, Huberman and Saldaña's (2014) guidance on purposive selection in qualitative research, which emphasises that the analytical quality of the sampling is proportional to the participants' information richness rather than their statistical representativeness. Stake (2006) similarly argues that multiple-case studies should select organisations that, together, offer the strongest opportunity to learn.

3.4 Participants and Sampling

Participants were selected through purposive sampling, targeting individuals with direct involvement in ERM governance and IA activities for their respective business areas. The target roles include Chief Audit Executives (CAEs), Internal Audit Managers, Chief Risk Officers (CROs), and other senior risk and business leaders such as heads of credit, model risk, or operational resilience. In NBLs, finance and business operations leads may also be included given the central role that finance and funding-related reporting plays in those institutions' risk governance.

The selection of these specific roles is deliberate, as each provides a different vantage point on the same process: the CAE sees the mandate, capability constraints, and the choices IA makes about where to direct its effort; the CRO sees whether audit findings actually change behaviour through adjustments in limits, pricing, model overlays, or remediation priorities; senior risk leaders and committee-adjacent roles see whether the questions asked in governance forums are evidence-based, whether decisions are documented, and whether action items are followed through to closure. Interviewing across these roles is a form of triangulation that guards against relying on any single participant's account.

A total of eleven interviews were conducted. Saunders and Townsend (2016), reviewing 798 qualitative studies in management and organisational research, argued that participant numbers should be justified by their analytical adequacy rather than meeting a statistical target. Guest, Bunce and Johnson (2006), in an empirical examination of when thematic saturation is reached in qualitative interview research, similarly observed that the majority of high-frequency themes are typically identifiable within the first six to twelve interviews. The eleven interviews were assessed for thematic saturation during analysis, with the final interviews predominantly confirming and refining themes evident in earlier transcripts (see Appendix D4). The sample is therefore considered sufficient to support the analytical claims advanced in the findings, while the limits of any qualitative sample of this size are acknowledged in Section 3.9. A full participant table is provided at Appendix B1.

3.5 Data Collection

Semi-structured interviews of 30-45 minutes each were used to collect the data, conducted in person or via secure video and recorded and transcribed via Microsoft Teams onto a single secure storage device with the participants' informed consent.

The interview guide was structured around six sections that map directly to the research questions and conceptual framework: role, mandate, and scope; access and capabilities; information quality and management information; governance challenge and decisions; sector-specific dynamics; and future outlook (see Appendix C1). Each section used an open-ended questioning approach with encouragement to share actual

examples. For instance, rather than asking whether IA improves information quality (a question that invites an affirmative answer and little else) the questioning structure asked participants to describe specific instances where audit practices changed the way risk data was collected, reconciled, or reported, and what happened as a result.

A shared working definition of ERM effectiveness is offered at the start of each interview as a reference point, though it is open to participant interpretation. Participants may hold different assumptions about what “effective ERM” means (some default to maturity frameworks, others to regulatory compliance, others to the quality of committee discussion), and the shared definition helps guide conversation without imposing a single framing. Without such a reference point, the interviews would risk producing data that is internally inconsistent and difficult to compare across organisations.

Social desirability bias is a known risk in governance research, with participants potentially presenting their practices in an overly positive light when discussing their own function’s contribution. The interview design addresses this by asking for specific examples and observable outcomes rather than self-assessments, making it harder to offer generic positive narratives without grounding them in something verifiable. Interviewing multiple roles within an organisation allows the study to compare individual accounts against each other, though this control is less readily available where the participant base does not consist of multiple roles within the same institution.

The semi-structured interview format is well established in qualitative research as a means of balancing thematic consistency across participants with the flexibility to follow diverse lines of enquiry (Kvale and Brinkmann, 2009; Galletta, 2013).

3.6 Data Management and Security

All recordings and transcripts were stored in encrypted formats with access restricted to the researcher, and all recordings are deleted after a transcript is deidentified in accordance with university ethics requirements, as explained below. Participant identities and institution names were de-identified using pseudonyms in all transcripts, coding files, and outputs. These data management practices align with the principles set out in the National Statement on Ethical Conduct in Human Research (NHMRC, ARC and

UA, 2023) and the Australian Code for the Responsible Conduct of Research (NHMRC, ARC and UA, 2018).

3.7 Data Analysis

The analysis followed a hybrid inductive-deductive approach, combining the structure of the conceptual framework with openness to themes that the framework does not anticipate. A purely deductive approach would have imposed the framework on the data and risk missing insights that did not fit neatly into predefined categories. A purely inductive approach would have ignored the substantial body of prior research that informs the study's design. The hybrid approach starts with the framework as a set of sensitising categories (IA practices, information quality, governance challenge, contextual conditions), allowing the data to reshape, extend, or challenge those categories as coding progresses.

The analytical process completes in stages. First-cycle coding works through each transcript, tagging segments with descriptive codes that capture what is being said; some of these codes correspond to framework constructs, others emerge from the data itself. Second-cycle coding groups first-cycle codes into broader themes and examines relationships between them. The Gioia methodology (Gioia, Corley and Hamilton, 2013) provides a structure for this progression: first-order concepts drawn from participant language, second-order themes that represent the researcher's interpretive aggregation of those concepts, and aggregate dimensions that connect themes to the theoretical constructs. Saldaña's (2021) coding manual provides further methodological grounding for the two-cycle approach.

Within-organisation analysis is conducted first, mapping the mechanisms and conditions that operate within each institution: what IA does in relation to ERM, how that affects information quality, how it shapes governance challenge, and what contextual factors enable or constrain those effects. Cross-organisation analysis then compares the within-organisation findings using Yin's (2018) approach of identifying patterns, similarities, and differences. The comparison between banks and NBLs is a key axis, but other dimensions of variation such as size, ERM maturity, and geographic operations are

also examined. Braun and Clarke's (2006) guidance on thematic analysis provides additional methodological grounding for the coding and theme development process.

The options to use qualitative data analysis software or AI tools require explicit consideration, because the choice carries trade-offs that are not always acknowledged. Dedicated QDA software offers systematic organisation of large volumes of transcript data, consistent code application, and a documented audit trail of coding decisions. However, it can also encourage treating coding as a mechanical process of tagging and sorting rather than as an interpretive act, fragmenting narratives into decontextualised segments and producing a tidier but less rich analysis than careful, repeated human reading of whole transcripts. This study therefore did not adopt dedicated QDA software, instead organising the coding work through structured documentation: a versioned codebook (Appendix D1), per-participant coded transcript summaries (Appendix D3), decision log (Appendix D5), and reflexive memos (Appendix D6), supplemented by AI-assisted analytical tools.

The study used ChatGPT 5.1 and Microsoft Copilot as cross-working AI assistance tools, employed to provide timely and accurate analysis of the collected and de-identified interview responses. The interpretive work, deciding what each segment means, how segments relate to each other, and what patterns do and do not say, was done by the researcher through close, repeated reading of AI-assisted pattern identification and analysis. The two AI tools cross-check each other's work to ensure the results remain factual and to minimise the risk of hallucinations. Recent literature suggests that large language models can support pattern recognition and code suggestion in qualitative work, but that their use raises concerns about transparency, reproducibility, and the potential for the model's training distribution to shape interpretation (Christou, 2023; Morgan, 2023). The researcher's qualitative research judgment therefore remains central, with AI tools used as research aids rather than substitutes for analytical reasoning.

3.8 Researcher Positionality

A further consideration is the researcher's own professional experience in this domain. Familiarity with IA and ERM governance can sharpen insight, making it easier to follow

technical explanations, understand industry jargon, and ask informed follow-up questions that a researcher without industry experience might miss. However, familiarity may also introduce preconceptions which the researcher needs to consciously neutralise during the interview process. A researcher who has worked in or alongside IA may unconsciously interpret participant accounts through the lens of their own experience, privileging certain explanations or taking for granted assumptions that ought to be examined.

The interview design addresses this through several mechanisms, none of which eliminates the influence of prior experience but all of which make that influence visible and manageable. Crucially, the interview guide uses open-ended prompts that do not presuppose particular answers, asking participants to describe, illustrate, and explain rather than confirm or deny. The approach of sensitising categories drawn from the literature rather than from personal experience provides a more defensible starting point. These safeguards provide a defensible positionality management strategy rather than a false claim of neutrality (see Appendix E1). Reflexivity in qualitative research is not merely a technique for managing bias but a core methodological commitment to making the researcher's role visible and accountable (Berger, 2015; Finlay, 2002), with Bourke (2014) arguing that positionality must be treated as an ongoing aspect rather than a one-time disclosure.

3.9 Limitations of the Research

The research design proposed in this study is intentionally scoped to maximise the depth of analysis whilst working within practical limitations – of which a number of limitations should be acknowledged at the outset.

First, the sample is purposive and non-probabilistic. While this is acceptable in the qualitative orientation of the study (Patton, 2015; Miles, Huberman and Saldaña, 2014), it also means findings cannot be generalised in a statistical sense. As with Lincoln and Guba (1985), the study targets analytical or theoretical generalisability rather than statistical generalisability, with detailed patterns and reasoning provided to enable readers to assess transferability to their own environments.

Second, the geographic and sectoral scope is bounded. The study focuses on retail banking and non-bank lending within the Australian financial services context, with limited extension to comparable institutions operating in other markets. Findings may therefore not transfer cleanly to insurance, superannuation, and other financial industry participants where the regulatory environment, risk profile, business structure, and IA mandate likely differ materially.

Third, access constraints inherent to governance research limit the depth of documentary triangulation possible. Reporting entities are understandably averse to sharing internal classified documents such as board meeting minutes, audit reports, and risk dashboards even in redacted form. The analysis therefore relies primarily on participant accounts. While interviewing across multiple roles within an organisation provides one form of triangulation, the absence of comprehensive documentary evidence is a material limitation.

Fourth, the researcher's prior professional experience in IA and ERM constitutes a source of interpretive influence. This is managed through reflexivity, sensitising categories drawn from the literature, and structured questioning, but cannot be eliminated. Readers should treat the analysis as situated rather than detached, in line with the broader interpretivist tradition (Berger, 2015; Bourke, 2014).

Fifth, the use of AI-assisted analysis tools introduces a methodological dependency that requires caution. These tools accelerate pattern identification but may surface interpretations shaped by their training data rather than by the participant accounts themselves (Christou, 2023). The researcher's direct, repeated engagement with transcripts is therefore essential, with AI-identified patterns treated only as inputs to interpretive judgment rather than as research conclusions.

Finally, the time and resource limitations of an MBA research program constrain both the number of interviews achievable and the duration over which IA's influence on ERM can be observed. The design captures a single point in time and cannot, by itself, demonstrate longitudinal change in governance behaviour, a limitation that future longitudinal research could address.

3.10 Ethical Considerations

The research involves senior practitioners discussing governance practices that may be commercially sensitive, professionally consequential, or both. The design follows the principles set out in the National Statement on Ethical Conduct in Human Research (NHMRC, ARC and UA, 2023) and the Australian Code for the Responsible Conduct of Research (NHMRC, ARC and UA, 2018).

Interview participation is voluntary, with informed consent obtained in writing from each participant before the interview. The consent process explains the purpose of the research, the topics to be discussed, the use made of the data, the storage and retention arrangements, and the participant's right to withdraw at any point without consequence. Recording only occurs with explicit, documented consent, and participants are informed that they may decline to answer any specific question and may request that any portion of their account be excluded from the analysis.

Confidentiality and anonymity are protected through the use of pseudonyms for individuals and institutions in all transcripts, coding files, draft outputs, and final documents. Identifiable information such as company and individual names, major public events, and unique organisational structures is generalised or removed where it could enable identification. Reporting focuses on aggregated themes and cross-organisation patterns rather than organisation-specific evaluations, protecting both individual participants and their institutions.

Data security is maintained through encryption of recordings and transcripts, restricted access controlled by the researcher, and the secure deletion of recordings once de-identified transcripts have been verified. Transcripts and analysis files are retained for the period of this research and thereafter destroyed in accordance with the research authoriser's records management requirements.

Conflicts of interest are relevant given the researcher's prior professional experience in the sector. Any current or recent professional relationship between the researcher and a prospective participant or institution is declared in advance, and where such a relationship could compromise the credibility or independence of the data, the relevant

case or participant is excluded. The dissemination of findings respects the integrity of the data and the legitimate interests of all participants, with member checking undertaken where the conditions of the engagement permit. Direct quotations are reviewed in all cases to ensure they cannot be traced to a specific individual or institution. Where findings touch on regulatory or supervisory matters, particular care is taken to present them as analytical observations rather than as evaluative judgments of any specific organisation.

4. Findings

4.1 Introduction

This chapter presents the findings from ten interviews with senior practitioners across the Australian financial services sector, supplemented by a comparative case (P11) reserved for the discussion chapter and a historical-perspective interview (P07) used selectively to anchor longitudinal observations. Full participant details are provided at Appendix B1. The participants were drawn from a deliberately varied set of vantage points: a Chief Risk Officer with extensive experience across multiple Australian banks (P05); a Chairman of Risk, Audit and Compliance Committees across multiple regulated financial-service entities (P08); a Head of Compliance in a multinational banking group (P03); two senior supervisors from a prudential regulator (P01, P09); a Chief Financial Controller in a non-bank lender (P02); a Head of Enterprise and Credit Risk in a multinational bank (P10); and two senior internal auditors at non-bank lenders (P04, P06). The role triangulation captured in this sample, combined with coverage of banks, non-bank lenders, and the prudential regulator itself, supports the analytical claim that the findings reflect a multi-perspective view of the IA-ERM relationship in contemporary Australian financial services rather than the perspective of a single role or institutional type.

4.2 Theme 1: Anchoring IA in the Risk Appetite Statement

Most consistently surfaced across the sample is that Internal Audit's contribution to ERM is constituted not by generic process review, but by its focus on measuring whether the institution operates within the risk appetite set by its board. The Risk Appetite Statement (RAS) emerges in participants' accounts as the analytical anchor that gives IA's work its governance significance: without the anchor, IA findings remain process observations; with it, the findings become enterprise risk signals that drive board engagement.

The senior risk leader P05 described IA's referencing of the RAS as an operational standard testing for adherence:

Internal audit's responsibility is to pick up those instances of where good and bad things are happening in accordance with the risk appetite statement. (P05, L93-95)

The senior supervisor articulated this from their regulatory perspective, observing that effective ERM requires the appetite statement to be operationally applied at the point of decision:

[The risk appetite statement should be] cascaded all the way down and... actually applied. Not just sitting in a board pack, but referenced when somebody is setting a [credit limit. (P09, L50-53)]

A particularly important practice surfaced in the regulator's account is the deliberate reframing of process lapses as risk appetite breaches. P09 described an instance where this anchoring move transformed the governance significance of a finding:

[IA] escalated this not as a process lapse but as a risk appetite breach... once you frame it as a risk appetite breach, the board has to engage. (P09, L328-331)

Several participants spoke similarly about ERM's application in practice. P02 characterised effective ERM by what he called the "single narrative" test, namely whether the same risk story is told consistently in the boardroom, the management committee, and the front line. P09 used the related distinction between ERM as a "living system" and ERM as a SharePoint page, observing that the difference is not in the documentation but in whether the documents are actually referenced when decisions are made.

4.3 Theme 2: Independence as structural, behavioural, and continuously contestable

IA's independence cannot be adequately characterised as a single attribute. Participants described independence as operating across multiple dimensions (structural, behavioural, goal-oriented, board-aligned, and contestable), each necessary but none alone sufficient. These accounts came particularly from participants who had observed independence failing in structurally well-configured functions, and who therefore framed independence as a continuously applied condition rather than simply an institutional concept.

The structural dimension is the most consistently mentioned. Participants described the same attributes: direct reporting from the Head of Internal Audit to the chair of the audit committee or the board, separation from operational management lines, dedicated audit committee sessions held without management present, and the codification of these arrangements in an Internal Audit Charter. The senior risk leader (P05) described a clean implementation where the IA function reported directly to the chair of the board, and at every quarterly meeting:

[the room was emptied] “of all management... with the exception of the CEO who was also an executive director. So basically they emptied the room out of all management and they got the internal audit management team in” (P05, L175-178).

The Chairman of Risk, Audit and Compliance Committees (P08) described an analogous arrangement:

[They] “report direct to me, so I know what’s going on. So when they go in, they have an unfettered access” (P08, L240-242).

The behavioural dimension is articulated less frequently but more pointedly. Several participants distinguished between an IA function that is structurally independent on paper and an IA function whose practitioners actually exercise that independence in their work. The senior supervisor observed that the independence of an IA function is tested at the moment a finding is drafted:

“When IA issues a finding, does it go to the board essentially as drafted, or does it get filtered through management?” (P09, L190-192)

The Head of Enterprise and Credit Risk in a multinational bank (P10) reinforced this distinction, framing structural independence as a precondition rather than a guarantee:

“Structurally, IA reports directly to the regional Audit Committee Chair, not to management. That separation is crucial” (P10, L129-131).

Independence as contestable, being the third dimension, is the most interesting feature of the data. The Head of Internal Audit at a non-bank lender (P04) framed independence as something that must be operationally protected by structural arrangements that are themselves continuously enforced:

[IA] owns its budget. Reports directly to the board audit committee. Has access to every data source without negotiation. (P04, L113-114)

The phrase "without negotiation" is significant: it positions independence as a condition that is liable to be negotiated away if not actively protected.

The Quality Assurance / Internal Audit distinction is seen across the data as the conceptual move that operationalises independence. P05 mentioned this in three-lines-of-defence terms:

“Quality assurance happens in line one and line two... internal audits are a critical component of that whole risk governance framework.” (P05, L86-89)

P08 also described the same separation from the governance vantage point:

“Compliance sits in there as our second line of defence. Our third line of defence is the internal audit function.” (P08, L221-222)

The Head of Compliance observed that the lines blur in practice and identified the maintenance of the distinction as one of the recurring governance tasks rather than a settled architectural feature:

Compliance, technically, is a second-line function. Audit is third line. But in big banks, those lines blur sometimes. (P03, L64-65)

The cumulative observation across the theme is that independence in IA functions is a system property rather than an attribute of any single arrangement; if any one element is weak, the entire system is vulnerable regardless of the collective strength of the others.

4.4 Theme 3: Advisory has legitimate scope but a sharp ownership boundary

The advisory dimension of Internal Audit's role produced more variation across the sample than any other theme in this chapter, but is not a random occurrence, as the data resolves into three distinguishable positions on the advisory question.

The most permissive position came from the senior risk leader (P05), who treated advisory as a natural and additive extension of IA's anchoring work:

"[the function should be able to] suggest or recommend or provide an advisory service on what controls may or may not be available within the broader industry, there is something we can we can automate or go and buy something off the shelf or something we can create ourselves, or even a manual solution" (P05, L133-137).

The advisory work was bounded, in that it offers options but does not direct decisions. The strictest separation position came from the Chairman of Risk, Audit and Compliance Committees (P08), who treated any blending of advisory with assurance as a structural conflict to be avoided rather than managed. Asked about a scenario in which the business wanted external advice on a software implementation, P08 was unequivocal that the institution's internal auditor would not be permitted to take on that work:

"they wouldn't be allowed to use [IA1]... we would say no because they're conflicted. I need IA1 to be able to go in over the top and give me the information I need" (P08, L562-572).

The argument is structural rather than ideological: the chair needs IA to provide assurance over the implemented system in due course; if IA has previously advised on its design, that future assurance is compromised.

The principle that IA should not own the design of what it will subsequently audit is widely shared across the sample, though participants articulate it differently. P09, P10, and P04 articulate it explicitly as a conditional rule in terms such as "never own", "guide but cannot design", "early consulting without crossing into design work", converging on essentially the same operational principle. P05 operates within the principle implicitly, describing IA's advisory contribution in terms of suggesting and recommending rather than designing. P08 draws an even stricter conclusion from the same underlying concern about future assurance, treating any advisory engagement as a structural conflict to be

avoided. P02's account of IA's involvement in a new funding-channel project demonstrates the principle operationally:

"IA asked questions early on" (P02)

..., without taking ownership of the project design.

The cumulative analytical claim is that the advisory question is not really a debate about whether IA should provide advisory work, but about where ownership begins. The conditional position maps the middle precisely. Advisory engagement is acceptable when IA's contribution stops short of design ownership. Where the boundary blurs in instances such as IA does take ownership of design work, write any procedures or policies, or configure the very systems it will subsequently audit, the IA function loses both its independence and its capacity to provide credible assurance over its own work.

4.5 Theme 4: Information quality is the operational core of IA's ERM contribution

If anchoring to the RAS is what gives IA its governance authority, and independence is what protects it, then information quality is what IA materially contributes to ERM. The convergence on this point was unanimous in the entire dataset, with every contemporary interview producing material that fell into this theme.

The senior risk leader (P05) provided the most context. Following a multi-million unexpected provisioning movement at one of his earlier institutions, IA was brought in to conduct an end-to-end review of the consumer credit collective provisioning process:

“[The review covered] from a data governance point of view. Where are you getting your data from? Is it documented? Is it... Is it right? What transformations have occurred from that data, from the source of truth into my modelling process and then exactly what the model was doing to spit out the number at the end of the day that I needed to book against the P&L” (P05, L346-352).

The IA review in this case was a traceability audit, testing whether the chain from raw data to reported number was sound at every step. P05's reflection on why IA was brought in despite his own confidence in his controls captures the analytical principle:

“having that third or fourth set of eyes over something that I ran and I built that model... provided a lot of comfort and support” (P05, L355-362).

This was especially interesting as it is evident that even an experienced and capable leader may be perceived as subject to confirmation bias.

The senior supervisor (P09) provided the most analytically structured account, articulating an explicit four-component framework for assessing information quality:

“...I look at [these], broadly. Lineage. Timeliness and completeness. Explainability. And aggregation with drill-down” (P09, L293-294).

Lineage was framed as:

“[the capacity to] trace a number on the board back all the way back to the source system, through every transformation, every reconciliation, every override, every overlay” (P09, L296-298).

Two specific failure modes recurred across multiple participants. The first is the self-confirming metric, a measurement that appears robust but actually only validates the limits of its own detection. The Head of Compliance (P03) provided the cleanest example: the bank's transaction reporting team produced green RAG indicators based on low error rates, but IA pointed out that "the error rate was based on detected errors. Not total errors. So if the detection mechanism was incomplete, the metric was meaningless. It was measuring its own confidence in itself" (P03, L94-97). The IA team then traced the lineage from source to regulator and identified silent overrides, logic inconsistencies, and an undocumented vendor patch. The second failure mode is the metric that aggregates cleanly upward but cannot be decomposed downward. P10 described how three different functions produced different numbers for the same underlying credit exposures:

"[IA traced] the transformation logic across four countries. Identified a misalignment in the consolidation methodology that had been there for years and had survived multiple system migrations because nobody had ever tested it end-to-end" (P10, L262-265).

The regulator (P09) had made the same observation in general terms:

"the number of times we've seen a board metric that nobody can really [follow]" (P09, L312-313).

Aggregated metrics that cannot be drilled down are governance illusions: authoritative on the board pack, but offering no analytical handle when the board wants to understand the underlying drivers. Across participants, IA's contribution to information quality is described in three operational moves: testing end-to-end traceability, challenging the substance of metrics rather than accepting the numbers, and forcing structural remediation rather than accepting workarounds.

4.6 Theme 5: Effective governance challenge is precise, evidence-grounded, and tone-aware

The capacity of IA to provide governance challenge, that is, to ask the questions that compel substantive engagement and, ultimately, change - is the moment where all the preceding themes either translate into ERM impact or fail to do so. Across the sample,

participants stated a consistent account of what distinguishes effective challenge from ineffective challenge.

The senior supervisor (P09) provided the most analytically structured account. Asked what good challenge looks like, P09 was unequivocal:

“What good challenge looks like - it’s evidence-based. It’s framed against the institution’s own appetite or policy, not against some abstract benchmark. And it’s specific enough that management can engage with it, agree or disagree, and the conversation can move forward” (P09, L377-381).

The contrast with weak challenge was equally precise:

“Bad challenge is vague. ‘We have concerns around the control environment.’ That’s almost meaningless. Management can’t action it. The board can’t follow up on it. It just becomes wallpaper” (P09, L382-385).

The metaphor implies that wallpaper is just decorative – does not load-bear and therefore not substantive.

The Head of Internal Audit at a non-bank lender (P04) described the same standard from the practitioner’s seat:

“‘Here’s what we found. Here’s why it matters to the appetite. Here’s what we recommend.’ Concise, specific, actionable” (P04, L282-284).

The regulator’s standard and the Head of IA’s own standard converging here is significant.

An important refinement emerged across the sample set: effective challenge requires precision about the underlying business reality and is not just based on analytical thoroughness. The senior risk leader (P05) articulated this point with what is perhaps the most pointed warning in the dataset. P05 argued:

“[Where IA produces a report that is] accurate but doesn’t reflect what’s happening on the ground... if you tell the board something that might be nearly factually correct but isn’t the way that the operation actually operates because you’ve missed a manual control or you don’t quite understand how the calculation works, that erodes that erodes the level of trust that the board will put in your report” (P05, L443-457).”

Technically accurate findings that misread operational reality do more damage than no challenge at all because they compound; each technically-correct-but-operationally-naive report progressively erodes the standing of the function.

Tone, the third dimension of effective challenge, surfaced across participants consistently. P09 framed it explicitly:

“There’s a difference between holding management to account and being adversarial. The best CAEs hold the line on substance but they bring it in a way that management can engage with rather than retreat from” (P09, L425-429).

P04 made essentially the same point from the CAE’s perspective:

“The CAE’s job isn’t to win arguments. It’s to drive change. That requires a tone that brings management along, not one that puts them on the defensive. Sometimes I have to remind my team we’re not trying to be right. We’re trying to be effective” (P04, L291-296).

A recurring failure mode identified across the sample is paper closure the situation in which an audit action is recorded as closed without the underlying issue having been substantively resolved. P09 articulated this directly:

“what we sometimes call paper closure. The action gets ticked off. There’s an email saying it’s done. But the underlying issue hasn’t actually been fixed. Closure evidence has to be substantive. Independent verification. Test results. Walkthrough by IA. That’s the standard” (P09, L405-410).

Paper closure is invisible to the audit committee unless the closure-evidence standard is actively enforced.

A second failure mode is challenge fatigue, where management progressively disengages from a function that is perceived to keep raising the same concerns without acknowledging progress or differentiating new issues from recurring ones. P09 acknowledged that challenge fatigue is a genuine phenomenon, particularly in institutions that have been through extensive remediation, and identified the practitioner moves that prevent it from becoming destructive:

“The strong CAEs manage it by varying the lens, by being thoughtful about how they bring findings, by celebrating closure when it actually happens, and by being clear when something is genuinely a new issue versus a recurrence” (P09, L420-424).

The Chairman (P08) described how the governance forum itself manages this dynamic, with the risk compliance and audit committee absorbing routine red-flagging and only escalating to the board what genuinely requires board-level engagement:

“A lot of the challenging happens at the risk compliance and audit committee level... where workplace health and safety reports, business continuity planning, compliance reports, if there’s any red flags in there, this committee will raise it. And go back to management and say, right, you need to... bring IA1 in, go back and redo that branch next quarter” (P08, L728-737).

The cumulative analytical claim is that effective governance challenge is a composite of three distinct properties: substantive precision (evidence-grounded, specific, actionable), operational accuracy (reflective of business reality), and tonal calibration (firm on substance, constructive in delivery). Where all three are present, IA’s challenge translates into decisions changing. Where any of the three is absent, challenge degenerates into either wallpaper, eroded trust, or adversarial standoff.

4.7 Theme 6: Capability is the binding constraint on contemporary IA contribution

A theme that emerged with notable consistency across the contemporary participants is that the binding constraint on IA’s contribution to ERM is the technical capability of the auditors themselves. The risks that financial institutions now manage — model risk, AI/ML governance, cloud and cyber resilience, vendor ecosystems, data lineage at scale — are technical in a way that traditional accounting-trained audit teams are not equipped to test.

The senior risk leader (P05) articulated the issue from his own historical experience.

“[When the audit work required someone] who knew internal systems, bank systems, data transformation, statistics, I needed an all-singing all-dancing auditor, couldn’t find one so we had actually three auditors in that process. One to look at the financial aspects, one to do the stats and the modelling and the data and one to look at the process” (P05, L404-409).

The implication is that the scope of what IA needs to test in a contemporary financial institution exceeds the skill set of any single auditor.

The senior supervisor (P09) provided the most synoptic statement of the same problem, framing the technical gap directly:

“Modern risks are technical. Models.. Automation... Vendor ecosystems... Data lineage.. Cloud architecture. Audit teams that are still dominated by chartered

accountants [...] they struggle with these areas. They can audit a balance sheet, but they can't audit a machine learning model" (P09, L260-265).

The Head of IA at a non-bank lender (P04) described actively reweighting recruitment:

"Data literacy. My analysts should be able to read SQL, understand data lineage, test transformation logic, challenge models. I'm hiring for this skill more aggressively than for traditional accounting backgrounds" (P04, L319-323).

P04 also articulated the rhetorical dimension:

"IA must talk in board language. Not 'Control 12 is ineffective'. But 'this undermines your risk appetite in this specific area, with the following implications for the strategy you've signed off.' That reframing changes how the report lands" (P04, L335-338).

The Head of Enterprise and Credit Risk (P10) provided what may be the sharpest statement of the strategic stakes:

"If IA cannot challenge future-looking risks, it becomes historical commentary. The board doesn't need[] that" (P10, L316-318).

The phrase "historical commentary" captures the failure mode precisely. An IA function that can describe what happened but cannot interrogate the systems that determine what will happen is not contributing to ERM in any meaningful sense.

A particular sub-theme is the implication of AI for IA capability. P05 articulated it most extensively, describing what he saw as a fundamental shift in the accessibility of model-building. The mathematics behind contemporary AI models, P05 observed, is the same mathematics he had used to build credit scorecards in the 1990s; what has changed is the speed and accessibility of the modelling process.

"It just takes 45 minutes to build a model, not full half months" (P05, L631-632).

The democratisation of model-building changes the scope of what IA needs to audit, both because the volume of models in production has increased and because the practitioners building them no longer have the training that earlier generations of model-builders had. Where previous-generation modellers had spent the bulk of their careers learning to clean data and understand its rules:

"75% or 80% of the time spent on building models getting the data right" (P05, L645).

...the current generation can build models on already-clean data without that learning.

P05 further says that the result is a population of:

“potentially, [...] uneducated people who are unaware of the potential implications of the models they’re building because they haven’t grown up taking 4 1/2 months to clean this data and actually understand all of the rules and register behind it” (P05, L648-654).

The capability burden on IA is to be able to interrogate these models and to challenge the assumptions of practitioners who may not themselves understand the implications of their work. The senior supervisor (P09) identified data analytics as the single capability most likely to compound across all the others:

“Data analytics. Without question. Because analytics underpins everything else. If you can interrogate the data yourself, you can challenge models, you can validate metrics, you can do continuous monitoring, you can identify outliers without waiting for management to flag them. It’s the multiplier capability” (P09, L280-285).

P10 added a complementary observation on the timing of capability investment, noting that institutions whose IA functions are most respected are “investing six to twelve months ahead” (P10, L330-331). The cumulative claim is that capability has become the binding constraint on what IA can credibly contribute to ERM. The structural arrangements, the disciplined advisory boundary, and the rigour of information-quality work all presuppose that the IA function has the technical capability to do the work.

4.8 Theme 7: Sector and regulatory context shape the substance and rhythm of IA work

The seventh theme to emerge is that the sector in which an IA function operates, and the regulatory pressures bearing on that sector at any particular moment, materially shape what IA actually does. The framework of effective IA contribution described in Themes 1 to 6 does not vary by sector in its underlying logic; what varies is the substantive content of the audits, the regulatory pressures driving the audit plan, and the rhythm at which the function pivots to address emergent risks.

The clearest statement of how regulation steers the audit programme was provided by the senior supervisor (P09):

“Massively. When major reforms emerge, [...] IA needs to pivot quickly. We expect them to be testing readiness, identifying blind spots, raising concerns before the regulator does” (P09, L437-441).

The regulator's expectation is not that IA will catch up to regulatory changes after the fact but that the function will anticipate them and surface readiness gaps before the regulator finds them in a thematic review. P09 was direct about what happens when this expectation is not met:

"Institutions that treat IA as a lagging function — they get caught out" (P09, L442-443).

The substantive content of the audit work, however, varies by sector. For banks, P09 identified the recurring pressure points as:

"...stress testing, BCBS 239 data aggregation, operational resilience, model risk governance" (P09, L447-448).

For non-bank lenders, the pressures reorganise around different fundamentals:

"...securitisation reporting integrity, covenants and triggers, funding stability, consumer credit obligations" (P09, L448-451).

The framing P09 offered captures the structural point precisely:

"Different shapes, same fundamental requirement — IA has to be ahead of the regulator, not behind" (P09, L451-453).

The Chief Financial Controller (P02) provided the operational corollary from inside an NBL. The NBL audit programme is structured around four substantive drivers:

"responsible lending outcomes that affect provisioning and reporting. Funding compliance, especially the securitisation trust obligations. Hardship reforms. Basel-like expectations in treasury practices, even though we're not a bank. And data governance trends generally" (P02, L256-261).

NBLs are not subject to APRA's prudential regime in the same way banks are, but the regulatory expectation in adjacent areas has been progressively imported into the sector. A particular feature of the NBL context that recurred across P02 and P04 is the role of investor scrutiny in shaping IA's audit work:

"When we explored a new funding channel with a global investor, IA asked questions early on. Could our data lineage withstand external validation? Could the originations system produce loan-level audit logs at the granularity the investor expected? Did our hardship processes meet the global investor expectations, which can be more stringent than our own?" (P02, L172-178).

NBL IA functions effectively answer to two governance audiences (the board, and the investors funding the securitisation programme) with the higher of the two standards typically setting the audit bar.

The Head of Compliance in a multinational bank (P03) described an analogous but distinct dynamic for institutions operating across jurisdictions. Each major jurisdiction introduces thematic reviews such as:

“consumer protection, AML uplift, operational resilience, complaints handling, vulnerability frameworks” and “these themes cascade into the audit plan” (P03, L228-230).

The examples P03 cited capture the cross-jurisdictional dynamic with operational specificity:

“UK conduct risk rules influenced our global complaints framework. APAC AML reforms forced a redesign of our sanctions filter governance. US model risk expectations impacted our transaction monitoring upgrades” (P03, L232-236).

The Head of Enterprise and Credit Risk (P10) described the same multi-jurisdictional dynamic from inside the institution, articulating the methodological implication for IA:

“the US regulator recently tightened expectations around AI model governance. The EU is pushing for heightened third-party risk visibility. Locally, Australia is emphasising operational resilience and CPS-like themes even for foreign ADIs” (P10, L237-242).

The audit programme, as P10 articulated it, must be designed against unforgiving terms:

“IA must be able to test readiness across all those overlapping frameworks. And that’s a real challenge because the frameworks aren’t perfectly consistent with each other. IA has to develop a methodology that works for the most stringent jurisdiction and then scale appropriately for others” (P10, L243-247).

P10 also provided the most consequential single example of sector-driven IA impact in the entire dataset. He stated that two years prior, the bank’s IA function had completed a review of operational resilience. Internally, the bank had treated its operational resilience programme as mature. IA went deeper and identified that:

“...[core systems] had recovery time objectives that exceeded our risk appetite. The RTOs documented were aspirational, [did not pass testing or assign limitations]. When IA actually compared them against the appetite statement, the gap was material” (P10, L156-160).

The finding triggered a board-level escalation, immediate remediation funding from the global CFO, and a complete rebuild of the resilience playbooks. The relevance of this example is captured in P10's punchline:

“six months later, we had a regional outage. Significant. If IA hadn’t forced that work, we’d have been exposed. Instead, we recovered within tolerance” (P10, L165-167).

5. Discussion

5.1 Introduction

This chapter develops the central analytical contribution of the study by aggregating the seven empirical themes from Chapter 4 into three theoretical dimensions, namely Anchored Authority, Disciplined Engagement, and Enabling Context. Each dimension is engaged with the existing literature at the point where the analytical claim is being advanced, rather than as a separate engagement section. The chapter then turns to the longitudinal contrast and comparative case, and to the implications for institutions, regulators, and IA practitioners.

5.2 Aggregate Dimension 1: Anchored Authority

The first aggregate dimension describes the conditions under which IA's voice has both the substantive grounding and the institutional standing to contribute meaningfully to ERM. Anchoring without independence produces compliant audit work that lacks governance traction; independence without anchoring produces unfettered work that lacks analytical substance. The two are mutually constitutive, and together specify the mechanism by which IA's authority is actually constituted, a question the existing literature has tended to address at the level of structural arrangements without adequately accounting for the analytical anchoring that gives those arrangements force.

The IIA's position paper on IA's role in ERM (IIA, 2009) and the more recent Three Lines Model (IIA, 2020) treat IA's contribution as defined primarily by its structural separation from operational management and by its position as the third line of defence, and COSO's enterprise risk management framework (COSO, 2017) similarly treats IA structurally. Both treatments are confirmed by the findings (Theme 2): participants described the line-of-defence architecture as the normative baseline, and institutions adhering to it produced IA functions with greater ERM contribution than those in which the lines were blurred. None of these contributions, however, fully addresses what makes any specific audit finding count as a governance signal rather than a process observation. The data answer, developed in Theme 1, is the anchoring move: IA's testing has to be conducted against the institution's stated risk appetite if findings are to compel board engagement, since without anchoring the same finding is management-handleable, whereas with it the finding becomes governance-level. This claim extends rather than contradicts the broader risk-management literature. Power's work on the risk management of nothing (Power, 2009) raises the concern that risk management can become ritualised performance disconnected from substantive work, however the data suggests the anchoring move is precisely what distinguishes substantive ERM from ritualised ERM: where IA tests against actual risk appetite, the function operates as a genuine assurance mechanism, and where it tests against generic process standards, it produces the compliant-but-empty work Power critiques. Mikes' work on calculative cultures (Mikes, 2009) describes how risk-management practice differs across

institutions according to dominant cultural orientation; the data suggest anchoring to the RAS is itself a specific cultural orientation, whose presence or absence largely determines whether IA's contribution is substantive or performative.

The independence component of the dimension extends the literature differently. Structural independence is treated extensively in the Three Lines Model (IIA, 2020), but the data contribute two further dimensions which the literature has tended to leave implicit. The first is behavioural independence, the willingness of practitioners to actually exercise the independence that structural arrangements provide. The senior supervisor's observation that "when IA issues a finding, does it go to the board essentially as drafted, or does it get filtered through management?" (P09, L190-192) reframes independence as a behavioural test rather than an attribute. The second is contestable independence, the recognition that independence is a position to be defended rather than a stable condition: the Head of Internal Audit's framing of access "without negotiation" (P04, L113-114) captures the contestable point. The systemic conception that emerges has particular resonance with Adam, Soliman and Mahtab's (2023) recent work on measuring ERM implementation, which treats ERM effectiveness as depending on the active operation of multiple components rather than the establishment of any single arrangement. The same logic applies to independence: Charter codification, reporting line, QA/IA distinction, behavioural exercise, and audit committee chair sponsorship operate concurrently, and weakness in any single one compromises the system as a whole.

5.3 Aggregate Dimension 2: Disciplined Engagement

The second aggregate dimension describes the substantive practice through which IA contributes to ERM on a day-to-day basis. Where Anchored Authority describes the conditions under which IA can speak with weight, Disciplined Engagement describes what the function should be doing with that weight, and the dimension comprises three interdependent practices: the conditional advisory boundary, end-to-end information quality work, and evidence-grounded governance challenge. Advisory drift erodes the standing required for credible challenge; information quality work without effective challenge produces findings that do not move; challenge without information quality

grounding becomes wallpaper. The three practices operate together rather than in sequence.

The advisory boundary is the most contested component (Theme 3), with the data resolving into three positions, namely permissive, strict-separation, and conditional, the conditional being the most articulated. Three participants articulated the same operational principle from different vantage points, converging on the view that IA may guide but never own, may consult early but not cross into design work. Sarens and De Beelde (2006) examine the advisory question through perception studies, and the IIA's Three Lines Model (IIA, 2020) treats advisory as a permissible but boundaried activity, however neither articulates the operational principle with the precision the data surface, the specific contribution being the identification of design ownership as the relevant boundary line. This permits the analytically valuable forms of advisory engagement, such as early consulting, pre-implementation review, and industry-practice translation, while ruling out the forms that erode the function's later capacity to test, namely writing policy, configuring systems, and taking the pen on design decisions.

Information quality work is what IA actually does on a day-to-day basis to contribute to ERM, with Theme 4 surfacing the strongest convergence in the dataset. The regulator's four-component framework (lineage, timeliness, explainability, and aggregation with drill-down) provides an analytically structured account of what good ERM information looks like, while the Chief Financial Controller's parallel articulation of end-to-end traceability and the senior risk leader's extended provisioning case provide operational corroboration. This finding has substantial resonance with BCBS 239's principles for risk data aggregation and reporting (BCBS, 2013), which articulate similar principles at the regulatory level, however the data contribute the empirical observation that BCBS 239's principles describe what IA actually tests in practice, not just what regulators think IA should test. The European Central Bank's findings on persistent implementation difficulties (ECB, 2018) and Forvis Mazars' more recent assessment (Forvis Mazars, 2024) both report that institutions continue to struggle with the principles in practice, and the data here suggest IA is the operational mechanism through which compliance is tested and forced.

Effective governance challenge is the moment at which IA's testing translates into decisions changing, with Theme 5 surfacing three components, namely substantive precision, operational accuracy, and tonal calibration. The data converged with consistency across all three: the regulator and the senior practitioner articulated identical standards from different vantage points, while the senior risk leader's contrasting account of how technically-correct-but-operationally-naive findings progressively erode board trust (P05, L443-457) added a particular refinement, namely that substantive rigour is necessary but not sufficient. This refinement reframes the existing IA literature on impact and credibility, particularly the work of Beasley, Clune and Hermanson (2005), Hoyt and Liebenberg (2011), and Lundqvist (2014), none of which substantially addresses how a single audit finding actually translates into a decision changing. Credibility, therefore, is not an output of the function but the operational medium through which the function works. The trust mechanism is prior to substantive work rather than after, with the implication that a function with intact trust converts relatively modest analytical work into substantial governance impact, whereas a function whose trust has eroded may not compel action even where its findings are well-grounded.

5.4 Aggregate Dimension 3: Enabling Context

The third aggregate dimension describes the structural conditions external to the IA function itself that determine whether the practices captured in the first two dimensions can actually be sustained in any given institution. The dimension is analytically distinct because its determinants sit largely outside the function's own control, being what institutions provide to IA and what the regulatory environment imposes, rather than what IA itself constructs. The dimension comprises capability investment (Theme 6) and the sectoral and regulatory conditions that shape the substantive content of audit work (Theme 7), with the data showing both operating as preconditions for the framework's other dimensions.

Capability emerged in the data as the binding constraint on contemporary IA contribution, with the findings showing that the risks financial institutions now manage (model risk, AI/ML governance, cloud and cyber resilience, vendor ecosystems, data

lineage at scale) are technical in a way the traditional audit-trained workforce is not equipped to test. The regulator's observation that audit teams dominated by chartered accountants struggle to audit machine learning models captures the structural problem with particular precision, while the senior risk leader's three-auditor solution to a single audit, the practitioner's reweighting of recruitment toward data-literate analysts, and the multinational risk leader's framing of capability investment ahead of need together describe the same structural shift. This finding has resonance with the broader literature on the changing nature of audit work, particularly Power's earlier observations on the audit explosion and organised uncertainty (Power, 2007), although the data add a more specific empirical claim, that the capability gap is not a problem of audit-quality variability across firms but a structural insufficiency of the traditional audit workforce relative to the technical demands of contemporary risk. The regulator's identification of data analytics as the multiplier capability, the single skill area whose presence improves the quality of every audit rather than just analytics audits, provides an additional refinement. The implication is that capability-building in IA is not a continuous-improvement question but a categorical one: institutions which have made the transition can audit contemporary risks; those which have not, cannot.

Sector and regulatory context, the second component, was articulated most directly by the regulator (Theme 7), whose framing of how major reforms (operational resilience, CPS 230 themes, cyber uplift, consumer protection reforms) cascade into audit programmes captures the rhythm at which contemporary IA functions must pivot. The sector-shaping observation has particular resonance with Australian-specific governance literature on prudential regulation, including APRA's CPS 220 (APRA, 2019), the more recent CPS 230 (APRA, 2025), and ASIC's RG 209 (ASIC, 2019), which together shape the institutional context within which Australian financial-services IA functions operate. What the data add is the empirical observation that the framework of effective IA contribution to ERM is sector-invariant in its underlying logic but sector-specific in its substance, with banks facing stress testing, BCBS 239 data aggregation, operational resilience, and model risk governance pressures, while NBLs face securitisation reporting integrity, covenants and triggers, funding stability, and consumer credit obligations.

This sector-specificity has particular force when applied to NBLs, where the Chief Financial Controller's account of investor scrutiny as a substantive driver of audit work reframes a question the existing literature has tended to treat structurally. NBLs are not subject to APRA's prudential regime in the same way banks are, however the data show their IA functions effectively answer to two governance audiences (the board and the investors funding the securitisation programme), with the higher of the two standards typically setting the audit bar. For multinational institutions, jurisdictional heterogeneity functions as a forcing mechanism driving audit standards upward to the most stringent applicable jurisdiction. The most consequential single example of sector-driven IA impact in the dataset, the multinational risk leader's account of the operational resilience review that prevented a regional outage, operationalises the sector dimension into a specific impact pattern.

A broader observation that emerges across all three dimensions is that, while the IIA's frameworks (IIA, 2009; IIA, 2020) treat structural independence as a primary marker of IA effectiveness, the data invert this hierarchy: structural arrangements were necessary but ubiquitous in the sample, with variation in IA effectiveness tracking variation in the behavioural, practice-based, and contextual factors developed across the three dimensions rather than in the structural arrangements themselves. The three dimensions developed in this and the preceding sections together constitute the central analytical contribution of this study.

5.5 Longitudinal contrast and comparative case

Two interviews, the historical-perspective interview with P07 and the deliberate comparative case with P11, were excluded from the primary thematic analysis but retained for testing the scope of the framework. P07 provides a longitudinal contrast that locates the framework in time; P11 provides a sectoral contrast that locates it in institutional context.

P07's career history captures a financial-services era in which the contemporary architecture of IA-ERM relationships did not exist. The participant's 22-year career at a major Australian bank, "18 of them was in head office departments, group financial controllers, different financial or finance departments within various divisions of the

bank” (P07, L36-38), describes an institutional context in which finance, audit, accounting, and compliance were not yet specialised functions. The analytical use of this contrast is to locate the framework in time. The structural arrangements at the heart of Anchored Authority did not exist in P07’s professional era in the form contemporary participants describe; the substantive practices at the heart of Disciplined Engagement (model risk audit, BCBS 239 lineage testing, CPS 230 readiness reviews) did not yet exist as auditable subjects; and the capability requirements at the heart of Enabling Context were configured around accounting and process competencies that contemporary participants describe as inadequate. The framework describes a recent state of governance practice rather than a timeless feature of organisational life, historically situated, configured around specific developments in regulatory environment, technical environment, and post-financial-crisis governance emphasis.

P11 was included as a deliberate comparative case (Decision 001 of the analytical decision log — see Appendix D5 — drawing on Stake’s (2006) multiple-case methodology), the participant carrying a financial-services audit background but having since transitioned into a non-financial-services industry context. The most analytically interesting contribution is P11’s structured/unstructured distinction, framing the financial-services audit environment as “structured” (regulator-driven, documentation-heavy) and contrasting it with the non-financial-services environment as “unstructured” in the documentary mechanism but structured “as a practise” (P11, L380-384). The contrast supports the analytical claim that the three-dimension framework is sector-shaped rather than universal: Anchored Authority depends on a Risk Appetite Statement developed under regulatory expectation, and P11’s industry context does not have an analogous artefact; Enabling Context takes regulator-driven sectoral pressures as a primary determinant of audit content, and P11’s industry is differently shaped. A second contribution is comparative validation of the IA-Tickbox-Trap failure mode identified in the primary analysis (P05, P06), with P11 articulating this from outside the financial-services sector and providing evidence that the failure mode is a more general pathology of contemporary IA practice rather than a financial-services artefact. A third contribution concerns the insider/outsider distinction: P11’s observation that “we feel the pinch. As an outsider, you will not feel the pinch” (P11, L355-360) frames a structural advantage of

the in-house IA function over external audit or consultancy engagement, where the “pinch” is the organisational stake an employed IA practitioner carries in remediation outcomes. The cumulative contribution of P07 and P11 is to bound the framework rather than to extend it: the three-dimension framework is not advanced as a universal theory of IA-ERM relationships but as an empirically-grounded account of a specific governance configuration in a specific sector at a specific moment.

5.6 Implications

The three-dimension framework has implications for three audiences.

5.6.1 Institutions and boards

Assessment across all three aggregate dimensions, rather than against any single feature, is required to evaluating IA effectiveness. Boards evaluating IA performance against single-attribute criteria (independence, audit-plan completion, finding counts) miss the systemic conception that the data show actually distinguishes effective from ineffective functions. The audit committee chair occupies a more analytically central position than the existing structural literature has tended to indicate. Across the data, the chair was the institutional actor whose willingness to defend the function’s prerogatives sustained independence, whose appetite to fund the function determined capability, and whose handling of management reactions to challenge shaped the tonal calibration through which findings were received.

A third implication concerns the trust mechanism. Where institutions have allowed IA’s standing to be eroded (through reports that misread operational reality, paper closure, or tonal misjudgement) the recovery of effectiveness is structurally slower than its loss. Institutions cannot rebuild board trust simply by improving audit-output quality; the rebuilding process requires sustained demonstration of operationally accurate, substantively precise, and tonally calibrated work over time. A fourth implication concerns the timing of capability investment: data showed that institutions whose IA functions are most effective invest in capability “ahead of need rather than after the fact” (P10, L330-331), with leading functions building skills six to twelve months ahead of risks

materialising. Capability investment is not appropriately funded as a residual category in the audit-plan budget; it has to be planned ahead of the risk horizon.

5.6.2 Regulators

The data confirm the value of the regulator-driven anchoring that distinguishes financial-services governance from less-regulated sectors, with the comparative-case analysis making this point at the level of framework scope. Regulators considering whether prudential expectations should be extended to currently-unregulated entities can take the data as evidence that the extension produces substantive governance changes rather than merely additional compliance costs.

A second implication concerns the operational specificity with which regulator expectations are communicated to institutions. The data showed that institutions whose IA functions are most responsive to regulatory pressure are those whose senior auditors interpret regulatory communications operationally rather than literally, with the Head of Compliance's account of cross-jurisdictional cascade (P03, L232-236) describing translation work that is itself a marker of mature audit function. Regulators communicating expectations in operationally-translatable terms enable this work; less translatable communications require more interpretive labour from institutions and produce more variable compliance outcomes. A third implication concerns the regulator's role in distinguishing performative from substantive ERM compliance, with the senior supervisor's identification of paper closure as a recurring failure mode (P09, L405-407) describing a regulatory practice that materially reduces the risk of performative compliance.

5.6.3 IA practitioners

Capability investment is no longer a discretionary professional development activity but a structural requirement for analytical relevance. Practitioners considering long-term careers in IA need to be actively building data, model risk, cyber, and behavioural risk competencies, since the IA practitioner of the next decade will need a substantially different skill profile than that of the previous decade.

The cultivation of effective challenge as a craft is a second implication. The CAE's job, in P04's framing, is "not to win arguments. It's to drive change" (P04, L292-294); tone is observable, learnable, and consequential for whether substantive work translates into governance impact. The active maintenance of conditions is a third implication: practitioners who treat independence as established through structural arrangements, the advisory boundary as a one-time agreement, or trust as a stable attribute will find their effectiveness eroding even when formal arrangements remain intact. A fourth implication, surfaced through P11's comparative case, concerns the structural advantage of the in-house IA function relative to external audit or consultancy engagement: the "pinch" of organisational stake in remediation outcomes that P11 describes (P11, L355-360) is part of what gives in-house IA its distinctive standing in ERM, an advantage realised through the substantive ongoing accountability that the in-house role makes possible rather than automatic.

6. Conclusion

6.1 Research aim, objectives, and questions revisited

The study aimed to evaluate Internal Audit's (IA) contribution to Enterprise Risk Management (ERM) effectiveness through practitioner narratives and artefacts. For this, four objectives were set out: to explain how IA defines and operationalises ERM assurance; to identify the enablers and barriers, namely mandate, sponsorship, access, and capability, that shape that contribution; to examine IA's influence on information quality and governance challenge; and to develop a grounded conceptual model of IA contribution across banks and non-bank lenders (NBLs). To steer this analysis, the study addressed four research questions:

- **RQ1:** How do IA professionals and risk executives describe IA's role and boundaries in ERM governance?
- **RQ2:** Which IA practices are perceived as most impactful for ERM effectiveness in retail banking and finance?
- **RQ3:** What organisational conditions enable or constrain IA's influence on the quality of risk information and governance challenge?
- **RQ4:** How do sector-specific factors (e.g. stress testing in banks; securitisation and warehouse reporting and covenants in finance companies) shape IA's ERM contribution?

The conclusion recaps how the dissertation answered these questions, identifies the resulting contribution, acknowledges limitations, and proposes directions for future research that extend the answers developed here.

6.2 Key findings and contribution

The central analytical contribution is a three-dimension framework describing the conditions under which IA's contribution to ERM is effective: Anchored Authority, Disciplined Engagement, and Enabling Context. The framework was developed inductively through Gioia-style coding of eleven semi-structured interviews with senior

practitioners across Australian banks, NBLs, and the prudential regulator. The four research questions are answered as follows.

RQ1 — IA's role and boundaries in ERM governance. Participants described IA's role as constituted by the anchoring of audit work to the Risk Appetite Statement (Theme 1) and the systemic conception of independence required to protect that anchoring (Theme 2). The boundary of IA's role is most contested at the advisory-assurance interface, where the data resolved into three positions, the most articulated being the conditional principle that IA may guide but never own (Theme 3).

RQ2 — IA practices perceived as most impactful for ERM effectiveness. The three practices identified as most impactful, together constituting the dimension of Disciplined Engagement, were the conditional advisory boundary (Theme 3), end-to-end information quality work (Theme 4), and evidence-grounded governance challenge (Theme 5). Information quality work surfaced with the strongest convergence in the dataset and was identified by participants as IA's most substantive day-to-day contribution. The trust mechanism emerged as the operational medium through which IA's findings convert into governance impact.

RQ3 — Organisational conditions enabling or constraining IA's influence. The data identified capability as the binding constraint on contemporary IA contribution (Theme 6), with traditional accounting-trained audit teams progressively unable to interrogate technical risks such as model risk, AI/ML governance, cyber resilience, and data lineage at scale. The enabling context further depends on institutional investment in data analytics as a multiplier capability, on capability investment ahead of need, and on the active sponsorship of independence and challenge by the audit committee chair.

RQ4 — Sector-specific factors shaping IA's ERM contribution. The framework was shown to be sector-invariant in its underlying logic but sector-specific in its substance (Theme 7). Banks face stress testing, BCBS 239 data aggregation, and operational resilience pressures; NBLs face securitisation reporting and investor scrutiny that frequently exceeds regulatory minimums; multinational institutions face overlapping jurisdictional frameworks that drive audit standards toward the most stringent applicable jurisdiction.

The framework is bounded historically and sectorally. The longitudinal contrast provided by the historical-perspective interview locates the framework as a recent specialisation in financial-services governance, configured around regulatory, technical, and governance developments of the past two decades. The comparative case locates the framework as sector-specific, dependent on the regulator-driven scaffolding that distinguishes financial services from less-regulated industries.

6.3 Limitations

The study's limitations are detailed in Section 3.9 and warrant brief restatement. The sample of eleven interviews, while consistent with established qualitative standards on saturation, remains a small purposive sample drawn from a single national context; findings are generalisable in the analytical sense rather than the statistical. Participating organisations were not directly accessible for documentary triangulation due to the confidentiality constraints of governance research, with analysis relying on participants' accounts. The researcher's own positionality in the financial services sector creates both opportunities and risks that have been managed through reflexive practice but cannot be fully eliminated. The use of AI-assisted analytical tools, while consistent with recent qualitative methodology literature (Christou, 2023; Morgan, 2023), introduces dependencies that the audit trail attempts to make transparent (see Appendix D5 and D6). The time-horizon constraint of an MBA dissertation has shaped both the recruitment process and the depth of within-case analysis possible per participant.

6.4 Recommendations for future research

Three directions for future research arise from the present study, each connecting to extending the answers to the research questions developed above.

The first recommendation is for empirical investigation of the framework's applicability outside contemporary Australian financial services, addressing the scope-limitation surfaced through RQ4. Corresponding work in other national contexts, in other sectors (such as insurance, superannuation, asset management, healthcare governance), or at other points in time would test whether the three-dimension framework holds, requires modification, or fails altogether. Future researchers undertaking this work should retain

the Gioia analytical structure (see Appendix D2) and the saturation-based sampling logic adopted here (see Appendix D4), while substituting sector-appropriate regulatory and institutional anchors. Such research would clarify the framework's analytical scope and identify the conditions under which alternative frameworks may be more appropriate.

The second is for investigation of the trust mechanism as a discrete research object, extending the answer to RQ2 on impactful IA practices. The present data identify trust as the operational medium through which IA's findings convert into governance impact, but the data do not directly address how trust is built, how it degrades, or how it can be recovered once lost. A longitudinal study tracking IA-board relationships across the stages of trust formation, erosion, and rebuilding would substantially extend the empirical understanding of this mechanism. Future researchers should consider methodologies that permit observation across time, such as embedded ethnography or repeated interviews across multiple audit cycles, since cross-sectional designs of the kind adopted here cannot capture the asymmetry between trust loss and trust recovery that the present data surface.

Thirdly, investigation of the capability transition currently underway in IA practice, extending the answer to RQ3 on enabling conditions. The data showed that traditional accounting-trained audit teams are progressively unable to interrogate the technical risks that increasingly dominate financial-services governance, and that institutions whose IA functions are most effective are investing in capability six to twelve months ahead of need. A longitudinal study of IA hiring, training, and team composition across institutions of different sizes and sectoral positions would surface the dynamics of this transition more precisely and could inform institutional decision-making about how to fund and structure the transition. Particular attention to data analytics capability must be paid by future researchers, identified in the present study as the multiplier skill that compounds across all other audit competencies.

6.5 Concluding reflection

The research aim was to evaluate IA's contribution to ERM effectiveness through practitioner narratives and artefacts. The empirical work shows that this contribution is more analytically demanding than the structural literature has tended to indicate,

depending on the concurrent operation of three interlocking conditions, namely Anchored Authority, Disciplined Engagement, and Enabling Context, with the absence of any single condition degrading the others. The four research questions are answered by reference to these three dimensions, with RQ1 and RQ2 substantially addressed by the Anchored Authority and Disciplined Engagement dimensions, RQ3 by the Enabling Context dimension and the systemic conception of independence, and RQ4 by the sector-specific application of the dimension framework across banks, NBLs, and multinational institutions. The study contributes a grounded conceptual model of IA contribution that extends the existing literature, satisfying the fourth objective. For institutions, regulators, and IA practitioners, the unifying implication is that the effectiveness of IA's contribution to ERM is a system property requiring active maintenance across multiple conditions rather than a stable institutional feature established by structural arrangement.

7. References

- Adam, M., Soliman, A.M. and Mahtab, N., 2023. Measuring enterprise risk management implementation: A multifaceted approach for the banking sector. *The Quarterly Review of Economics and Finance*, 87, pp.244–256.
- Arena, M. and Azzone, G., 2009. Identifying organisational drivers of internal audit effectiveness. *International Journal of Auditing*, 13(1), pp.43–60.
- Australian Prudential Regulation Authority (APRA), 2019a. *CPS 220 Risk Management*. Canberra: APRA.
- Australian Prudential Regulation Authority (APRA), 2019b. *CPS 234 Information Security*. Canberra: APRA.
- Australian Prudential Regulation Authority (APRA), 2023. *APS 220 Credit Risk Management*. Canberra: APRA.
- Australian Prudential Regulation Authority (APRA), 2025. *CPS 230 Operational Risk Management*. Canberra: APRA.
- Australian Securities and Investments Commission (ASIC), 2019. *RG 209 Credit Licensing: Responsible Lending Conduct*. Sydney: ASIC.
- Basel Committee on Banking Supervision (BCBS), 2013. *Principles for Effective Risk Data Aggregation and Risk Reporting (BCBS 239)*. Basel: Bank for International Settlements.
- Beasley, M.S., Clune, R. and Hermanson, D.R., 2005. Enterprise risk management: An empirical analysis of factors associated with the extent of implementation. *Journal of Accounting and Public Policy*, 24(6), pp.521–531.
- Berger, R., 2015. Now I see it, now I don't: Researcher's position and reflexivity in qualitative research. *Qualitative Research*, 15(2), pp.219–234.
- Bourke, B., 2014. Positionality: Reflecting on the research process. *The Qualitative Report*, 19(33), pp.1–9.

Braun, V. and Clarke, V., 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), pp.77–101.

Christou, P.A., 2023. The use of artificial intelligence (AI) in qualitative research for theory development. *The Qualitative Report*, 28(9), pp.2739–2755.

Committee of Sponsoring Organizations of the Treadway Commission (COSO), 2017. *Enterprise Risk Management—Integrating with Strategy and Performance*. New York: AICPA.

Creswell, J.W. and Poth, C.N., 2018. *Qualitative Inquiry and Research Design: Choosing Among Five Approaches*. 4th ed. Thousand Oaks, CA: Sage.

Eisenhardt, K.M. and Graebner, M.E., 2007. Theory building from cases: Opportunities and challenges. *Academy of Management Journal*, 50(1), pp.25–32.

European Central Bank (ECB), 2018. *Report on the Thematic Review on Effective Risk Data Aggregation and Risk Reporting*. Frankfurt: ECB.

Finlay, L., 2002. Negotiating the swamp: The opportunity and challenge of reflexivity in research practice. *Qualitative Research*, 2(2), pp.209–230.

Forvis Mazars, 2024. *Risk Data Aggregation & Risk Reporting: Navigating BCBS 239 Implementation Progress*. [online] Available at: <https://www.forvismazars.us/forsights/2024/02/risk-data-aggregation-risk-reporting-navigating-bcbs-239-requirements> [Accessed 28 April 2026].

Galletta, A., 2013. *Mastering the Semi-Structured Interview and Beyond: From Research Design to Analysis and Publication*. New York: NYU Press.

Gioia, D.A., Corley, K.G. and Hamilton, A.L., 2013. Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational Research Methods*, 16(1), pp.15–31.

Guest, G., Bunce, A. and Johnson, L., 2006. How many interviews are enough? An experiment with data saturation and variability. *Field Methods*, 18(1), pp.59–82.

Hoyt, R.E. and Liebenberg, A.P., 2011. The value of enterprise risk management. *Journal of Risk and Insurance*, 78(4), pp.795–822.

Institute of Internal Auditors (IIA), 2009. *The Role of Internal Auditing in Enterprise-wide Risk Management*. IIA Position Paper. Altamonte Springs, FL: IIA.

Institute of Internal Auditors (IIA), 2020. *The IIA's Three Lines Model: An Update of the Three Lines of Defense*. Lake Mary, FL: IIA.

International Organization for Standardization (ISO), 2018. *ISO 31000:2018 Risk Management—Guidelines*. Geneva: ISO.

Kaplan, R.S. and Mikes, A., 2012. Managing risks: A new framework. *Harvard Business Review*, 90(6), pp.48–60.

Kvale, S. and Brinkmann, S., 2009. *InterViews: Learning the Craft of Qualitative Research Interviewing*. 2nd ed. Thousand Oaks, CA: Sage.

Lincoln, Y.S. and Guba, E.G., 1985. *Naturalistic Inquiry*. Beverly Hills, CA: Sage.

Lundqvist, S.A., 2014. An exploratory study of enterprise risk management: Pillars of ERM. *Journal of Accounting, Auditing and Finance*, 29(3), pp.393–429.

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

Miles, M.B., Huberman, A.M. and Saldaña, J., 2014. *Qualitative Data Analysis: A Methods Sourcebook*. 3rd ed. Thousand Oaks, CA: Sage.

Morgan, D.L., 2023. Exploring the use of artificial intelligence for qualitative data analysis: The case of ChatGPT. *International Journal of Qualitative Methods*, 22, pp.1–10.

National Health and Medical Research Council (NHMRC), Australian Research Council (ARC) and Universities Australia (UA), 2018. *Australian Code for the Responsible Conduct of Research*. Canberra: NHMRC.

National Health and Medical Research Council (NHMRC), Australian Research Council (ARC) and Universities Australia (UA), 2023. *National Statement on Ethical Conduct in Human Research*. Canberra: NHMRC.

Nkansa, P., Barr-Pulliam, D. and Walker, K., 2025. From compliance to strategic partnerships: The role of internal audit in enterprise risk management and opportunities for future research. *Journal of Risk and Financial Management*, 18(12), p.707.

Paape, L. and Speklé, R.F., 2012. The adoption and design of enterprise risk management practices: An empirical study. *European Accounting Review*, 21(3), pp.533–564.

Patton, M.Q., 2015. *Qualitative Research and Evaluation Methods*. 4th ed. Thousand Oaks, CA: Sage.

Power, M., 2004. *The Risk Management of Everything: Rethinking the Politics of Uncertainty*. London: Demos.

Power, M., 2007. *Organized Uncertainty: Designing a World of Risk Management*. Oxford: Oxford University Press.

Power, M., 2009. The risk management of nothing. *Accounting, Organizations and Society*, 34(6–7), pp.849–855.

Saldaña, J., 2021. *The Coding Manual for Qualitative Researchers*. 4th ed. Thousand Oaks, CA: Sage.

Sarens, G. and De Beelde, I., 2006. Internal auditors' perception about their role in risk management: A comparison between US and Belgian companies. *Managerial Auditing Journal*, 21(1), pp.63–80.

Saunders, M.N.K., Lewis, P. and Thornhill, A., 2019. *Research Methods for Business Students*. 8th ed. Harlow: Pearson.

Saunders, M.N.K. and Townsend, K., 2016. Reporting and justifying the number of interview participants in organization and workplace research. *British Journal of Management*, 27(4), pp.836–852.

Stake, R.E., 2006. *Multiple Case Study Analysis*. New York: Guilford Press.

Yin, R.K., 2018. *Case Study Research and Applications: Design and Methods*. 6th ed. Thousand Oaks, CA: Sage.

Appendix A – Ethics and Consent

A1. Participant Information Sheet

Study Title: Assessing Internal Audit's Contribution to ERM Effectiveness in Financial Institutions

Researcher: Atapattu Wasala Mudalige Lasindu Sharlin Ratnayake
Candidate for Master of Business Administration
University of Wales Trinity Saint David
Email: 2031417@student.uwtsd.ac.uk
Tel: +61 426 870 074

Supervisor: Dr Laura James, University of Wales Trinity Saint David

Purpose of the Study

This research aims to explore how Internal Audit (IA) contributes to the effectiveness of Enterprise Risk Management (ERM) frameworks within financial institutions. The study seeks to identify enablers and barriers to IA's role in strengthening ERM and to develop conceptual insights that inform governance best practices.

Why You Are Being Invited

You have been invited because your professional role involves Internal Audit, Risk Management, or Governance oversight within a financial institution. Your insights are critical to understanding how IA interacts with ERM processes in practice.

What Participation Involves

- A semi-structured interview (approximately 30–45 minutes) conducted via a secure online platform or, where feasible, in a professional setting.
- The interview will focus on IA's role in ERM assurance and advisory activities, governance mechanisms, and decision-useful risk information.

- Participation is entirely voluntary. You may decline to answer any question, pause, or withdraw at any time before anonymisation of your data.

Risks and Benefits

- **Risks:** Minimal. The discussion will focus on professional practices, not personal matters. Confidentiality safeguards will prevent identification of individuals or organisations.
- **Benefits:** Your contribution will help advance understanding of IA's role in ERM and may inform improved governance practices across the sector.

Confidentiality and Data Protection

- All data will be anonymised during transcription; identifiers such as names and organisation details will be removed.
- Audio recordings will be deleted after transcription and quality checks.
- Data will be stored securely in encrypted folders, accessible only to the researcher and supervisor.
- Findings will be reported in aggregate form or using generic descriptors (e.g., "Audit Manager, mid-sized lender").

Withdrawal Rights

You may withdraw at any time before anonymisation of your data. After anonymisation and integration into thematic analysis, withdrawal will not be possible because individual contributions cannot be isolated.

Ethical Approval

This study has been reviewed and approved by the University of Wales Trinity Saint David Research Ethics Committee in accordance with the University's Research Ethics and Integrity Policy.

Contact for Questions or Concerns

For any questions about the study, please contact:
Shar Ratnayake (Lead Researcher) – 2031417@student.uwtsd.ac.uk

For concerns about ethical conduct, contact: pgresearch@uwtsd.ac.uk
(Subject: Ethical Approval – Shar Ratnayake)

Next Steps

Please read this sheet carefully before signing the consent form. Your signature will confirm that you understand the study and agree to participate voluntarily.

Name:

Position:

Employer:

Signed date:

A2. Consent Form

Study Title: Assessing Internal Audit's Contribution to ERM Effectiveness in Financial Institutions

Lead Researcher: Atapattu Wasala Mudalige Lasindu Sharlin Ratnayake
Master of Business Administration candidate, University of Wales Trinity Saint David

Supervisor: Dr Laura James, University of Wales Trinity Saint David

Please tick each box to confirm your agreement before signing.

	Statement
☐	I confirm that I have read and understood the Participant Information Sheet (<i>Appendix A1</i>) for this study, dated [insert date], and have had the opportunity to ask questions.
☐	I understand that my participation is entirely voluntary and that I may withdraw at any time before anonymisation of my data, without giving a reason and without consequence to me.
☐	I understand that the interview will be audio recorded for transcription purposes, that audio recordings will be deleted after transcription is verified, and that transcripts will be de-identified using pseudonyms.
☐	I understand that direct quotations from my interview may be used in the dissertation and any subsequent publications, but only in a form that does not identify me or my organisation.
☐	I understand that the de-identified data will be stored securely in encrypted folders, accessible only to the researcher and supervisor, and retained for the period of the research and thereafter destroyed in accordance with University of Wales Trinity Saint David data management requirements.
☐	I understand that withdrawal after the point of anonymisation and integration into thematic analysis is not possible, as individual contributions cannot be isolated.
☐	I understand that this study has been reviewed and approved by the University of Wales Trinity Saint David Research Ethics Committee.
☐	I agree to participate in the study described above.

Participant Details

Name (printed): _____

Position: _____

Employer: _____

Signature: _____

Date: _____

Researcher Confirmation


I confirm that I have explained the nature and purpose of the study to the participant named above, and that I have answered any questions to the best of my ability.

Researcher signature: _____

Date: _____

A3. Ethics Approval

The following email confirms approval of the ethics application by the supervisor on behalf of the University of Wales Trinity Saint David Research Ethics Committee, granting authorisation to commence primary research.



The sender is external to the organisation. Ensure caution is taken with any links or attachments included in the email below.

Thanks for this, Shar. I am happy to approve this Ethics form so you are now able to undertake your primary research when best. Please do let me know if you would like a meeting to discuss next steps

Kind Regards

Laura

From: Shar Ratnayake <shar.ratnayake@uwtsd.ac.uk>

Sent: Tuesday, January 20, 2026 05:47

To: Laura James <laura.james@uwtsd.ac.uk>

Cc: Shar Ratnayake (2031417) <2031417@student.uwtsd.ac.uk>; sharlin.ratnayake@uwtsd.ac.uk

Subject: RE: Ethical considerations form and Papers

RHYBUDD: Deilliodd yr e-bost hwn o du allan i system E-bost Prifysgol Cymru Y Drindod Dewi Sant. Peidiwch ag ateb, na chlicio ar ddolenni nac agor atodiadau oni bai eich bod yn adnabod cyfeiriad e-bost yr anfonwr ac yn gwybod bod y cynnwys yn ddiogel.
WARNING: This email originated from outside University of Wales Trinity Saint David's Email System. Do not reply, click links or open attachments unless you recognise the sender's email address and know the content is safe.

Good morning Laura,

THE ORIGINAL ETHICS APPLICATION FORM, SUBMITTED ON 20 JANUARY 2026 AND APPROVED ON 22 JANUARY 2026, IS AVAILABLE ON REQUEST.

Appendix B — Sample and Participants

B1. Participant Table

The following table summarises the eleven participants who contributed to this study. All identifying information (names, employers, specific locations) has been removed in accordance with the consent and confidentiality arrangements set out in Appendices A1 and A2. Participants are identified by pseudonymous codes (P01 to P11) used throughout the findings and discussion chapters.

Code	Role	Sector / Institutional Context	Vantage Point	Use in Analysis
<i>P01</i>	Senior Supervisor	Prudential regulator (Australia)	Regulatory oversight across multiple regulated institutions	Primary thematic analysis
<i>P02</i>	Chief Financial Controller	Non-bank lender (Australia)	Senior finance and treasury, NBL governance	Primary thematic analysis
<i>P03</i>	Head of Compliance	Multinational banking group	Cross-jurisdictional compliance and risk	Primary thematic analysis
<i>P04</i>	Head of Internal Audit	Non-bank lender (Australia)	Day-to-day practitioner leadership, NBL IA	Primary thematic analysis
<i>P05</i>	Chief Risk Officer	Multiple Australian banks (historical and current)	Senior risk leadership, cross-institutional experience	Primary thematic analysis
<i>P06</i>	Senior Internal Auditor	Non-bank lender (Australia)	Operational practitioner, NBL IA	Primary thematic analysis

P07	Senior Finance Leader (retired)	Major Australian bank (historical)	22-year career across head office finance and control functions	Historical-perspective interview; longitudinal contrast
P08	Chairman, Risk, Audit and Compliance Committees	Multiple regulated financial-service entities	Audit committee chair vantage; governance forum perspective	Primary thematic analysis
P09	Senior Supervisor	Prudential regulator (Australia)	Regulatory oversight, governance challenge expectations	Primary thematic analysis
P10	Head of Enterprise and Credit Risk	Multinational bank	Cross-jurisdictional risk and capital, group-level vantage	Primary thematic analysis
P11	Senior practitioner (financial-services audit background)	Non-financial-services industry	Comparative case (financial-services background, currently outside financial services)	Comparative case; sectoral scope contrast

Composition summary

Total participants: 11

Sectoral coverage:

- Australian banks (including multinational): 4 participants (P03, P05, P08, P10)
- Non-bank lenders (NBLs): 3 participants (P02, P04, P06)
- Prudential regulator: 2 participants (P01, P09)
- Historical and comparative perspectives: 2 participants (P07, P11)

Role coverage:

- Chief-level (CRO, CAE, CFC, Head of Compliance, Head of Enterprise Risk): 5 participants
- Audit committee chair and historical finance leadership: 2 participants

- Prudential supervisors: 2 participants
- Operational IA practitioners: 2 participants

Analytical treatment:

- Primary thematic analysis (P01-P06, P08, P09, P10): 9 contemporary financial-services interviews forming the core dataset for the Gioia-style coding and the seven empirical themes developed in Chapter 4
- Comparative case (P11): treated separately, deployed in the discussion chapter to test the sectoral scope of the framework, consistent with Stake's (2006) multiple-case methodology and Decision 001 of the analytical decision log (Appendix D5)
- Longitudinal contrast (P07): treated as a historical-perspective interview deployed in the discussion chapter to locate the framework in time, consistent with Decision 002 of the analytical decision log (Appendix D5)

B2. Organisation Selection Rationale

This appendix sets out the rationale for organisational sampling in this study, consistent with the methodological principles set out in Section 3.3 of the dissertation.

Sampling logic

Selection of organisational contexts followed criterion-based sampling with maximum variation across sector and operational scale, moderated by participant availability (Patton, 2015; Stake, 2006). Three distinguishing criteria informed the selection.

First, institution type. Australian retail banks and non-bank lenders (NBLs) were both targeted to enable cross-sectoral analysis. Banks operate within the APRA prudential regime (CPS 220, CPS 230, CPS 234), face stress-testing and BCBS 239 obligations, and have ERM frameworks shaped by extensive supervisory expectation. NBLs operate outside APRA's prudential perimeter, rely on securitisation and warehouse funding rather than deposits, and face investor scrutiny that frequently exceeds regulatory minimums. The sectoral contrast was selected deliberately to surface whether IA's contribution to ERM operates through consistent mechanisms across the two contexts or through sector-specific mechanisms.

Second, size and complexity. Within each sector, attempts were made to include institutions of varied scale, from regional to major. Larger institutions tend to have more developed ERM frameworks, dedicated risk and audit functions of meaningful size, and direct supervisory engagement; smaller institutions face the same governance expectations with leaner resources and tighter capability constraints. Including both surfaces how scale moderates IA's substantive contribution.

Third, geographic and jurisdictional reach. One multinational financial services organisation was included alongside Australia-only institutions, on the basis that it operates in both Australian and global markets. Multinational institutions face overlapping jurisdictional frameworks (US model risk expectations, EU third-party visibility requirements, APRA prudential standards), and including such an institution allowed the data to surface how jurisdictional heterogeneity functions as a forcing mechanism on audit methodology.

Confidentiality constraint and its accommodation

Confidentiality is a structural feature of governance research, with institutions understandably cautious about sharing committee minutes and audit reports even in redacted form. The research design accommodated this by focusing on interview evidence (rather than documentary triangulation), by committing to strict anonymisation of institutions and participants in all outputs, and by referring to organisations only by generic descriptor (e.g., "non-bank lender", "multinational bank") in the dissertation. No institution is identifiable by name in this dissertation or in any analytical artefact.

Comparative case selection (P11)

One participant was selected as a deliberate comparative case (P11), drawing on Stake's (2006) multiple-case methodology. This participant carries a financial-services audit background but has since transitioned into a non-financial-services industry context. The analytical purpose of this inclusion was to test the sectoral scope of the framework developed from the contemporary financial-services sample, with the participant's account treated separately from the primary thematic analysis and deployed in the discussion chapter to contextualise the sector-specific dynamics of the core dataset. The rationale for this treatment is documented in Decision 001 of the analytical decision log (Appendix D5).

Historical-perspective interview (P07)

A second participant (P07) was included with a deliberately historical orientation. P07's 22-year career at a major Australian bank spanned an era in which the contemporary IA-ERM architecture did not yet exist as a specialised configuration, with finance, audit, accounting, and compliance functions less differentiated than in contemporary practice. The analytical purpose was longitudinal contrast, allowing the framework to be located in time rather than presented as a timeless feature of organisational life. P07 was treated separately from the primary thematic analysis and used selectively in the discussion chapter, consistent with Decision 002 of the analytical decision log (Appendix D5).

Reflection on sample composition

The final sample of eleven participants offers role triangulation (regulator, chair, CRO, CAE, CFC, senior risk and audit practitioners), sectoral diversity (banks, NBLs, regulator, comparative-case industry), and temporal range (historical and contemporary). The composition is consistent with Miles, Huberman and Saldaña's (2014) principle that the analytical quality of qualitative sampling is proportional to the participants' information richness rather than their statistical representativeness, and with Saunders and Townsend's (2016) review establishing that participant numbers in qualitative organisational research should be justified by analytical adequacy rather than meeting a statistical target. Saturation analysis (Appendix D4) confirms that the sample was sufficient for the analytical claims advanced in the dissertation.

Appendix C — Data Collection

C1. Interview Guide

This appendix reproduces the semi-structured interview guide used during fieldwork. The guide is organised in six sections aligned with the conceptual framework set out in the literature review and methodology chapters, with Section 5 (sector-specific dynamics) incorporating enablers and barriers as a sub-component. Each section includes core themes for exploration and suggested probes. The semi-structured format permitted the interviewer to adapt the order and depth of questioning to each participant's role and vantage point, while ensuring thematic consistency across interviews.

Pre-interview confirmation

Before commencing each interview, the interviewer:

- Confirmed receipt and understanding of the Participant Information Sheet (Appendix A1)
- Confirmed signed consent (Appendix A2) and explicit consent to audio recording
- Confirmed the participant's right to decline any question, pause the interview, or withdraw before anonymisation
- Confirmed the working definition of ERM effectiveness for shared reference, while noting that participant interpretations were welcomed

Section 1 — Role and Mandate

Core themes:

- Definition of IA's ERM assurance scope: how ERM assurance is framed within the organisation, and which ERM components IA covers (risk appetite, stress testing, operational resilience).
- Boundaries between assurance and advisory: where advisory begins and ends, and tensions between independence and influence.

- Independence maintenance: mechanisms for safeguarding independence.

Suggested probes:

- How is IA's ERM assurance scope documented and communicated internally?
- Does the scope include strategic risk areas or only compliance-driven elements?
- What advisory activities does IA undertake, and how do you manage independence?
- Have there been situations where advisory blurred boundaries? How were these resolved?
- What protocols or safeguards ensure IA remains independent when providing advice?
- How does the Audit Committee monitor independence in ERM-related work?

Section 2 — Access and Capabilities

Core themes:

- Access to risk data and systems: level of visibility into risk data pipelines.
- Access to governance forums: participation in ALCO, Credit Committee, Board Risk Committee.
- Skills and resources: technical capabilities (analytics, model risk expertise).
- Capability gaps and development needs: resource constraints and future needs.

Suggested probes:

- What systems and data sources does IA have access to for ERM assurance?
- Are there any restrictions that limit IA's ability to validate risk information?
- Does IA attend key governance forums? How does this influence its ability to challenge risk information?

- What impact does forum access have on IA's credibility and effectiveness?
- What specialist skills does IA currently have for ERM assurance?
- Where do you see gaps in technical capability (data analytics, stress testing, model governance)?
- What capability gaps most limit IA's ERM contribution?
- If you could strengthen one area of IA capability, what would it be and why?

Section 3 — Information Quality (Decision-Useful Management Information)

Core themes:

- Data lineage and integrity across source systems, transformations, reconciliations, and reporting layers.
- Timeliness and completeness of risk metrics for committee decision cycles.
- Explainability of models (assumptions, overlays, governance evidence) used in credit, capital, liquidity, and operational risk.
- Aggregation and drill-down capability that preserves traceability without re-identifying institutions or individuals in outputs.

Suggested probes:

- Walk me through how a key ERM metric (e.g., arrears rate, liquidity buffer) is produced end-to-end. Where has IA intervened to improve reliability?
- What reconciliations or exception logs are most informative, and how does IA validate them?
- When management uses model overlays, how does IA assess the documentation and rationale?
- How quickly can decision-makers obtain corrected MI after an identified issue?

Evidence sought:

- Examples of closed audit actions that improved lineage and reconciliations.
- Artefacts such as data dictionaries, lineage maps, model governance packs, and committee MI change logs before/after IA input.

Pitfalls explored:

- MI that is timely but not explainable; comprehensive but not decision-useful.
- Over-reliance on manual fixes with no sustainable control improvement.

Section 4 — Governance Challenge and Decisions

Core themes:

- Quality of challenge: how evidence-based questions are framed by IA and responded to by management and committees.
- Actionability: clarity of ownership, due dates, and closure evidence for risk actions.
- Influence pathways: how IA insights translate into changes in limits, pricing, portfolio mix, funding actions, or remediation priorities.

Suggested probes:

- Describe a recent instance where IA's work changed a decision or prompted a re-think. What enabled that impact?
- How are challenges and responses recorded in minutes, and who tracks through to closure?
- Where does challenge fatigue occur, and how do you keep it constructive?

Evidence sought:

- Committee minutes with documented challenge, action logs and closure proofs; thematic summaries showing trend reduction in repeat issues.

- Examples where IA assurance led to limit/overlay adjustments or decommissioning of unreliable MI.

Pitfalls explored:

- Advisory drift that dilutes independence.
- Actions logged but not meaningfully closed (paper compliance).

Section 5 — Sector-Specific Dynamics (including enablers and barriers)

Core themes (banks, APRA-regulated):

- Stress testing workflow: scenario design, overlays, model risk governance, use of results in setting limits and capital appetites.
- Risk data aggregation (BCBS 239): accuracy, completeness, timeliness, and adaptability; IA's role in validating aggregation and reporting controls.
- Operational resilience and CPS 230 readiness: mapping of critical operations, impact tolerances, and evidence that MI supports executive and board decisions.

Core themes (non-bank lenders and finance companies):

- Securitisation and warehouse reporting integrity: trustee reports, pool eligibility and stratifications, waterfall calculations, cash-reconciliation controls.
- Covenants and triggers: monitoring accuracy, exception management, and governance of investor communications.
- Funding stability: how IA assurance influences confidence of funders and board oversight.

Enablers and barriers (cross-sectoral):

- Mandate clarity and sponsorship from Audit Committee and executive.
- Access to systems, forums, and capability (analytics, model risk, data engineering literacy).

- Culture of openness to challenge versus defensive posture.
- Resourcing and prioritisation within the audit plan.

Suggested probes:

- Which sector pressures most shape IA's ERM priorities here?
- Where has sector regulation directly steered the audit programme?
- Which enablers most consistently correlate with IA impact here?
- What barriers keep recurring despite attempts to address them?
- How do you negotiate access when it is initially constrained?

Evidence sought:

- Clear ERM mandate statements in the IA Charter; standing invitations to key committees; defined data access protocols.
- Training logs or recruitment plans that address capability gaps.

Pitfalls explored:

- Over-scoping without sponsorship; ad-hoc access that jeopardises consistency; reliance on one or two "heroes".

Section 6 — Future Outlook

Core themes:

- Regulatory horizon and expected shifts in prudential expectations affecting ERM information and resilience.
- Skills evolution: analytics, model risk, AI literacy for explainability and control testing.
- Operating model: earlier engagement (e.g., pre-implementation reviews) while preserving independence.

Suggested probes:

- Looking 12 to 24 months ahead, what changes in regulation or funding conditions will most affect ERM MI and IA's work?
- Which skills will be most critical for IA to deepen its ERM contribution?
- Where can IA engage earlier without crossing advisory or management boundaries?

Interview flow guidance

The interview was structured as follows:

- Open with context-setting: confirmation of participant's role and involvement in ERM.
- Open-ended prompts under each theme to encourage narrative responses.
- Sector-specific probes applied where relevant (banks vs non-bank lenders).
- Closing reflection: "What would make IA's ERM contribution more impactful in your organisation?"

Compliance and governance notes

- All interviews were conducted within the ethical scope set out in the approved ethics application: voluntary participation, right to withdraw before anonymisation, confidentiality, and anonymised reporting.
- Special category data was deliberately not collected; conversations were steered toward professional practice and organisational processes.
- Consent forms were stored separately from research data; the coding key linking pseudonyms to participants was maintained offline until analysis was complete; audio recordings were deleted after verified transcription.

Appendix D — Data Analysis Artefacts

D1. Codebook

This codebook documents the analytical coding structure developed during qualitative analysis of the eleven interview transcripts. The structure follows the Gioia methodology (Gioia, Corley and Hamilton, 2013): first-order codes drawn from participant language, second-order themes representing the researcher's interpretive aggregation, and aggregate dimensions connecting themes to the theoretical constructs.

The codebook was developed iteratively across eight versions. Earlier versions (v1 to v4) captured first-cycle descriptive coding of individual transcripts; v5 introduced the first-order code clusters; v6 consolidated the second-order themes; v7 applied the three aggregate dimensions; v8 (this final version) reflects the structure as it appears in the dissertation findings and discussion chapters. Decisions affecting codebook structure are documented in the analytical decision log (Appendix D5).

Structure summary

- **47 first-order codes** grouped into **12 code clusters**
- **7 second-order themes** (each corresponding to a section in the findings chapter)
- **3 aggregate dimensions** (each corresponding to a section in the discussion chapter)

Aggregate dimensions

Code	Dimension	Themes consolidated
AD1	Anchored Authority	Theme 1 (RAS Anchoring), Theme 2 (Independence)
AD2	Disciplined Engagement	Theme 3 (Advisory Boundary), Theme 4 (Information Quality), Theme 5 (Governance Challenge)
AD3	Enabling Context	Theme 6 (Capability), Theme 7 (Sector and Regulatory Context)

Theme-level structure

Theme 1 — Anchoring IA in the Risk Appetite Statement

Cluster 1.1 — RAS as operational standard

Code	Label	Definition	Exemplar participants
1.1.1	RAS-as-anchor	IA tests adherence to the Risk Appetite Statement as the operational standard for audit findings	P05, P09
1.1.2	RAS-cascade	Articulation of the RAS at the point of decision rather than as a board-pack document	P09
1.1.3	Appetite-breach-reframing	Deliberate reframing of process lapses as risk appetite breaches to compel board engagement	P09
1.1.4	Living-system-vs-SharePoint	Distinction between ERM as embedded practice and ERM as documented framework	P09

Cluster 1.2 — ERM coherence

Code	Label	Definition	Exemplar participants
1.2.1	Single-narrative-test	Consistency of risk story across boardroom, management committee, and front line	P02
1.2.2	ERM-as-living-system	ERM as continuously referenced rather than periodically documented	P02, P09
1.2.3	Risk-signal-vs-process-observation	Distinction between findings that compel governance engagement and findings that remain management-handleable	P05, P09

Theme 2 — Independence as structural, behavioural, and continuously contestable

Cluster 2.1 — Structural independence

Code	Label	Definition	Exemplar participants
2.1.1	Reporting-line	Direct reporting from CAE to audit committee chair or board	P05, P08, P10
2.1.2	Charter-codification	Formalised arrangements in the Internal Audit Charter	P05, P10
2.1.3	Management-exclusion-from-AC	Dedicated audit committee sessions held without management present	P05, P08
2.1.4	Unfettered-access	CAE access to all data sources without negotiation	P04, P08

Cluster 2.2 — Behavioural and contestable independence

Code	Label	Definition	Exemplar participants
2.2.1	Finding-filtering-test	Whether findings reach the board as drafted or are softened by management filtering	P09
2.2.2	Continuous-exercise	Independence as a behavioural practice that must be exercised, not just structured	P04, P09
2.2.3	Defensible-budget-control	IA owning its budget as a protection against operational interference	P04

Code	Label	Definition	Exemplar participants
2.2.4	QA-IA-distinction	Conceptual separation between line one or two quality assurance and line three internal audit	P03, P05, P08
2.2.5	Audit-committee-sponsorship	Active defence of IA's prerogatives by the audit committee chair	P08

Theme 3 — Advisory has legitimate scope but a sharp ownership boundary

Cluster 3.1 — Three positions on the advisory question

Code	Label	Definition	Exemplar participants
3.1.1	Permissive-advisory	Position that advisory is a natural and additive extension of IA's anchoring work	P05
3.1.2	Strict-separation	Position that any blending of advisory with assurance is a structural conflict to be avoided	P08
3.1.3	Conditional-advisory	Position that IA may guide but never own; may consult early but not cross into design work	P09, P10, P04

Cluster 3.2 — Design-ownership boundary

Code	Label	Definition	Exemplar participants
3.2.1	Guide-not-own	Operational principle articulated as conditional rule across multiple participants	P09, P10, P04
3.2.2	Early-consulting	Advisory engagement at pre-implementation stage that does not cross into design ownership	P04
3.2.3	Advisory-drift	Failure mode in which IA becomes so embedded in management's risk-management work that it loses standing to provide independent assurance	P04, P08
3.2.4	Pre-implementation-review	Specific form of advisory engagement that surfaces design weaknesses before they are committed	P02, P04

Theme 4 — Information quality is the operational core of IA's ERM contribution

Cluster 4.1 — The four-component framework

Code	Label	Definition	Exemplar participants
4.1.1	Lineage	Capacity to trace a number from the board pack back to source system through every transformation	P09, P05
4.1.2	Timeliness-completeness	Risk metrics arriving on the committee decision cycle in usable form	P09
4.1.3	Explainability	Assumptions, overlays, and governance evidence supporting model outputs	P09

Code	Label	Definition	Exemplar participants
4.1.4	Aggregation-with-drill-down	Capacity to drill from aggregated metrics to underlying drivers	P09, P10

Cluster 4.2 — Information quality failure modes

Code	Label	Definition	Exemplar participants
4.2.1	Self-confirming-metric	Measurement that validates the limits of its own detection rather than the underlying reality (the "detected errors not total errors" failure mode)	P03
4.2.2	Aggregation-decomposition-failure	Metric that aggregates cleanly upward but cannot be decomposed downward	P10, P09
4.2.3	Manual-workaround-as-fix	Reliance on manual fixes rather than structural remediation	P09
4.2.4	Provisioning-data-integrity	Specific information-quality failure mode surfaced in the unexpected provisioning movement case	P05

Theme 5 — Effective governance challenge is precise, evidence-grounded, and tone-aware

Cluster 5.1 — Three components of effective challenge

Code	Label	Definition	Exemplar participants
5.1.1	Substantive-precision	Challenge that is evidence-based, framed against appetite, and specific enough that management can engage with it	P09, P04
5.1.2	Operational-accuracy	Challenge that reflects business reality and is not just analytically thorough	P05
5.1.3	Tonal-calibration	Firm on substance, constructive in delivery, brings management along rather than putting them on the defensive	P09, P04

Cluster 5.2 — Failure modes of governance challenge

Code	Label	Definition	Exemplar participants
5.2.1	Wallpaper	Vague challenge that management cannot action and the board cannot follow up on	P09
5.2.2	Trust-erosion-from-naivety	Technically-correct-but-operationally-naive findings that progressively erode board trust	P05
5.2.3	Paper-closure	Audit actions recorded as closed without underlying issue being substantively resolved	P09

Code	Label	Definition	Exemplar participants
5.2.4	Challenge-fatigue	Management progressive disengagement from a function perceived to raise the same concerns	P09, P08
5.2.5	Adversarial-standoff	Tone failure where holding the line on substance tips into adversarial posture that management retreats from	P09, P04

Cluster 5.3 — Trust as operational medium

Code	Label	Definition	Exemplar participants
5.3.1	Trust-as-medium	Credibility framed not as output of the function but as operational medium through which findings convert into governance impact	P05
5.3.2	Trust-asymmetry	Trust as easily lost and slowly rebuilt	P05

Theme 6 — Capability is the binding constraint on contemporary IA contribution

Cluster 6.1 — Technical capability gap

Code	Label	Definition	Exemplar participants
6.1.1	Chartered-accountant-limit	Traditional accounting-trained audit teams unable to interrogate technical risks	P09
6.1.2	Multi-disciplinary-team-need	Single audit requiring multiple specialists (finance, stats, process) because no single auditor has the full skill set	P05
6.1.3	Data-literacy-recruitment	Active reweighting of recruitment toward analytics-capable practitioners	P04
6.1.4	Historical-commentary-failure	IA function that describes what happened but cannot interrogate forward-looking systems	P10
6.1.5	AI-tooling-democratisation	Implication of widely-accessible AI for scope of models IA must audit and capability of practitioners building them	P05

Cluster 6.2 — Capability investment

Code	Label	Definition	Exemplar participants
6.2.1	Multiplier-capability	Data analytics as the single skill area whose presence improves the quality of every audit, not just analytics audits	P09
6.2.2	Ahead-of-need-investment	Institutions investing in capability six to twelve months ahead of risks materialising	P10
6.2.3	Board-language-rhetoric	Translating audit findings into appetite-framed terms that compel governance engagement	P04

Code	Label	Definition	Exemplar participants
6.2.4	IA-Tickbox-Trap	Failure mode in which IA becomes mechanical rather than judgement-led (also articulated by P11 in the comparative case)	P05, P06, P11

Theme 7 — Sector and regulatory context shape the substance and rhythm of IA work

Cluster 7.1 — Bank-specific drivers

Code	Label	Definition	Exemplar participants
7.1.1	Stress-testing	Stress-testing workflow as a primary audit focus for banks	P09
7.1.2	BCBS-239-data-aggregation	Risk data aggregation against BCBS 239 principles as a substantive audit focus	P09, P10
7.1.3	Operational-resilience-CPS230	Operational resilience and CPS 230 readiness reviews as emergent audit focus	P09, P10
7.1.4	Model-risk-governance	Model risk governance, particularly for credit and capital models	P09, P10, P05

Cluster 7.2 — NBL-specific drivers

Code	Label	Definition	Exemplar participants
7.2.1	Securitisation-reporting-integrity	Audit focus on trustee reports, pool eligibility, and waterfall calculations	P02, P09
7.2.2	Covenants-triggers	Covenants and triggers monitoring as a substantive audit focus for NBLs	P02
7.2.3	Investor-scrutiny-as-driver	Investor scrutiny operating as a substantive driver of audit work that frequently exceeds regulatory minimums	P02
7.2.4	Two-audience-governance	NBL IA effectively answering to both board and investors, with the higher of two standards setting the audit bar	P02, P09
7.2.5	Hardship-reforms	Consumer credit hardship obligations as an emergent NBL audit driver	P02

Cluster 7.3 — Multinational and jurisdictional drivers

Code	Label	Definition	Exemplar participants
7.3.1	Cross-jurisdictional-cascade	Thematic reviews in one jurisdiction cascading into audit plans globally	P03
7.3.2	Most-stringent-jurisdiction-design	Audit methodology designed for the most stringent applicable jurisdiction and scaled appropriately for others	P10

Code	Label	Definition	Exemplar participants
7.3.3	Regulator-anticipation	Expectation that IA will anticipate regulatory shifts rather than catch up after thematic reviews	P09
7.3.4	Operational-resilience-RTO-case	Specific impact pattern documented in P10's account of operational resilience review preventing regional outage	P10

Codebook evolution

The codebook went through eight iterations during analysis. Significant transitions:

- **v1 to v3:** First-cycle descriptive coding of P01 through P05 transcripts, generating an initial population of 78 descriptive codes
- **v4:** Consolidation pass following P06, P08, P09 transcripts; codes reduced to 51 through merging of conceptually overlapping codes
- **v5:** Introduction of the 12 code clusters as second-cycle aggregation; codes refined to 47 first-order codes
- **v6:** Second-order theme structure (the 7 themes) formalised; theme boundaries tested against the dataset
- **v7:** Three aggregate dimensions (Anchored Authority, Disciplined Engagement, Enabling Context) introduced; theme-to-dimension mapping documented
- **v8 (this version):** Final consolidation following inclusion of P10 and the deliberate comparative case (P11); minor refinements to code labels for consistency with the findings and discussion chapters

The analytical decision log (D5) records the substantive decisions that shaped the codebook's evolution, including Decision 001 (treatment of P11 as comparative case) and Decision 002 (treatment of P07 as historical-perspective interview).

D2. Data Structure Diagram

This appendix presents the Gioia three-level data structure that summarises the analytical relationships between first-order codes, second-order themes, and aggregate dimensions, as developed during qualitative analysis of the eleven interview transcripts.

The data structure is the standard visual artefact required by the Gioia methodology (Gioia, Corley and Hamilton, 2013) and demonstrates how participant language was aggregated through successive interpretive moves into the theoretical framework advanced in the discussion chapter.

The full underlying codebook (47 first-order codes consolidated into 12 clusters) is documented in Appendix D1. This appendix presents the structure at the level of cluster-to-theme-to-dimension relationships, consistent with the Gioia convention.

How to read the data structure

The data structure reads left-to-right as a movement from participant language to theoretical abstraction. Each column represents a level of interpretive aggregation:

First-order code clusters capture the language and practices participants themselves used, grouped into conceptually-related sets. These clusters consolidate 47 individual first-order codes documented in Appendix D1.

Second-order themes represent the researcher's interpretive move from descriptive coding to thematic aggregation, identifying patterns and analytical claims that emerge from the clustered codes.

Aggregate dimensions represent the highest level of abstraction, organising the themes into three theoretical constructs that together constitute the central analytical contribution of this study.

The three aggregate dimensions are not independent of one another. Anchored Authority creates the conditions under which Disciplined Engagement is possible; Enabling Context determines whether the engagement can actually be sustained in any given institution. The relationships between dimensions are developed in the discussion chapter (Chapter 5).


Theme-to-research-question mapping

For reference, the mapping between the seven themes and the four research questions is as follows:

Research Question	Themes Primarily Addressed
RQ1: How do IA professionals and risk executives describe IA's role and boundaries in ERM governance?	Themes 1, 2, 3
RQ2: Which IA practices are perceived as most impactful for ERM effectiveness in retail banking and finance?	Themes 3, 4, 5
RQ3: What organisational conditions enable or constrain IA's influence on the quality of risk information and governance challenge?	Themes 6, plus the independence component of Theme 2
RQ4: How do sector-specific factors shape IA's ERM contribution?	Theme 7

This mapping is consistent with the answers to the research questions developed in Section 6.2 of the conclusion chapter.

Data Structure Diagram (Gioia Methodology)



Adapted from the Gioia methodology (Gioia, Corley and Hamilton, 2013). Full codebook in Appendix D1.

D3. Sample Coded Transcript Excerpt

Presented in this appendix is a sample of the coding applied during qualitative analysis, drawn from the transcript of Participant P08 (Chairman of Risk, Audit and Compliance Committees across multiple regulated financial-service entities). Selected as an analytical exemplar because of the governance-forum vantage point it represents, the excerpt is the only contribution in the dataset drawn from a board chair cum audit committee chair, and as such it provides a perspective on IA's contribution to ERM that sits above the management layer and above the IA function itself. All identifying information has been removed from the transcript in accordance with the de-identification protocol described in Section 3.6 of the methodology chapter, with the participant referred to throughout by pseudonymous code (P08) and no institutional identifiers, personal names, or details that could enable re-identification appearing in this or any other analytical artefact.

How to read this appendix

Across three columns, the excerpt is presented in three analytical layers. The first column reproduces the de-identified participant language verbatim from the verified transcript, with square brackets indicating words or phrases replaced during de-identification. The second column lists the first-cycle descriptive codes applied to each segment, drawn from the codebook (Appendix D1) and referenced by their codebook code number. The third column identifies the second-order theme to which the coded segment contributes, consistent with the data structure diagram (Appendix D2). Where a single segment attracts multiple codes, all codes are listed, and the coding decisions reflected here are consistent with v8 FINAL of the codebook.

Transcript excerpt — P08, Lines 221–222, 240–242, 562–572, 728–737

Transcript Segment	First-Order Codes Applied	Second-Order Theme
"Compliance sits in there as our second line of defence. Our third line of defence is the internal audit function." (L221–222)	2.2.4 QA-IA-distinction	Theme 2: Independence as structural, behavioural, and continuously contestable
"[They] report direct to me, so I know what's going on. So when they go in, they have an unfettered access." (L240–242)	2.1.1 Reporting-line; 2.1.4 Unfettered-access; 2.2.5 Audit-	Theme 2: Independence as structural, behavioural, and continuously contestable

	committee-sponsorship	
"They wouldn't be allowed to use [IA1]... we would say no because they're conflicted. I need [IA1] to be able to go in over the top and give me the information I need." (L562–572)	3.1.2 Strict-separation; 3.2.3 Advisory-drift (pre-emption of); 2.2.5 Audit-committee-sponsorship	Theme 3: Advisory has legitimate scope but a sharp ownership boundary
"A lot of the challenging happens at the risk compliance and audit committee level... where workplace health and safety reports, business continuity planning, compliance reports, if there's any red flags in there, this committee will raise it. And go back to management and say, right, you need to... bring [IA1] in, go back and redo that branch next quarter." (L728–737)	5.2.4 Challenge-fatigue (management of); 5.1.1 Substantive-precision; 5.1.3 Tonal-calibration; 2.2.5 Audit-committee-sponsorship	Theme 5: Effective governance challenge is precise, evidence-grounded, and tone-aware

Analytical memo — P08 excerpt

Memo date: During v6 codebook development

Participant: P08

Analytically significant in a way that distinguishes P08's account from all other participants in the dataset is the governance-forum vantage point it occupies: as Chairman of the Risk, Audit and Compliance Committees across multiple regulated entities, P08 is the institutional actor who receives IA's output rather than the actor who produces it, and whose sponsorship sustains or erodes the structural and behavioural conditions that the other participants describe as preconditions for IA effectiveness. The recurring presence of code 2.2.5 (Audit-committee-sponsorship) across all four segments extracted here was not an analytical imposition; it emerged from the data as a consistent pattern, with P08 consistently framing his role as one of active guardianship of IA's independence and access rather than passive governance oversight.

Of particular analytical interest is the advisory boundary passage at lines 562-572, in which P08 articulates what was coded as the strict-separation position (3.1.2), drawing on the same underlying concern about future assurance that motivates the conditional position held by P09, P10, and P04, but reaching a more restrictive conclusion. Rather than treating this as a contradiction of the conditional position, the analysis treats it as an

analytically informative divergence, since the strict-separation position and the conditional position both derive from the same principle (that IA must not audit what it has designed) and differ only in where they draw the practical boundary. The three-position taxonomy developed in Theme 3 (permissive, strict-separation, conditional) was formalised during v6 codebook development in direct response to this divergence, with the strict-separation position anchored in P08's account and the permissive position anchored in P05's.

Decision recorded: The L728-737 governance-challenge passage was initially coded only against 5.2.4 (Challenge-fatigue management) during first-cycle coding, on the basis that P08 is describing how the committee absorbs routine red-flagging to prevent challenge fatigue at board level. During second-cycle coding, however, the passage was recoded to include 5.1.3 (Tonal-calibration), since P08's description of the committee's handling of red flags ("go back to management and say, right, you need to... bring[IA1] in") demonstrates the tonal calibration principle in institutional practice: the escalation is firm on substance but structured to bring management back into the process rather than issuing a finding from above. This recoding is noted in the decision log (Appendix D5) under the entry for Theme 5 second-cycle refinements.

Positionality note: P08's account of the strict-separation position on advisory work was initially read by the researcher through a lens shaped by personal professional experience, in which the conditional position (guide but never own) was the more familiar convention. The reflexive journal (Appendix D6) records a specific moment of recognition during analysis of this passage where the researcher's prior assumption that the conditional position was the dominant one needed to be suspended in favour of reading P08's strict-separation position on its own terms. The three-position taxonomy that resulted is the more accurate representation of the data precisely because that suspension was made.

Note on full transcript availability

Retained in the researcher's secure encrypted storage in accordance with the data management arrangements described in Section 3.6 of the methodology chapter and the ethics approval granted by the University of Wales Trinity Saint David Research Ethics

Committee (Appendix A3), the full de-identified transcripts for all eleven participants (P01 to P11) are available for inspection by the dissertation supervisor and examiner on request.

D4. Saturation Analysis

This appendix documents the chronological saturation analysis conducted during qualitative analysis of the eleven interview transcripts, consistent with the saturation logic set out in Section 3.4 of the methodology chapter and the analytical standards described by Guest, Bunce and Johnson (2006) and Saunders and Townsend (2016). Presented here is the consolidated record of theme emergence, development, and saturation across the interview sequence, demonstrating that the sample of eleven interviews was sufficient to support the analytical claims advanced in the findings chapter.

Saturation logic

Within qualitative research, thematic saturation describes the point at which additional interviews cease to produce new analytical content, with subsequent accounts confirming and refining themes already present in the data rather than introducing new ones. Operationalised in this study across three stages of engagement:

Emergence refers to the point at which a theme or code cluster first appeared in the data, either as a participant-led concept (bottom-up emergence) or as a recognisable instance of a sensitising category drawn from the conceptual framework (top-down recognition).

Development refers to the stage during which a theme was extended, refined, or complicated by subsequent interviews, with new sub-codes, contrasting positions, or failure modes being added to the emerging cluster.

Saturation refers to the point at which no new codes were being added to a theme across successive interviews, with subsequent accounts either confirming existing codes or providing additional exemplars of patterns already present in the data.

Interview sequence and participant coding

For the purposes of the saturation analysis, interviews are referred to by sequence number (I1 to I11) rather than by participant code, since the order of interview conduct does not necessarily correspond to the P01-P11 pseudonymous sequence used in the findings chapter. The ten substantive contemporary interviews (I1 to I10) were conducted across the core sample; the historical-perspective interview (P07) and the comparative case (P11) are treated separately and their saturation is noted at the end of this appendix.

Saturation tracking table

The following table records the status of each of the seven second-order themes across the interview sequence. Key: **E** = first emergence; **D** = theme developed or extended (new codes, contrasting positions, or failure modes added); **C** = confirmed (no new analytical content; existing codes confirmed or exemplified); **S** = saturation point declared.

Theme	I1	I2	I3	I4	I5	I6	I7	I8	I9	I10	Saturation point
T1: RAS Anchoring	E	D	D	C	C	C	C	C	C	C	I4
T2: Independence	E	D	D	D	C	C	D	C	C	C	I7
T3: Advisory Boundary		E	D	D	D	C	C	C	C	C	I6
T4: Information Quality	E	D	D	D	D	C	C	C	C	C	I6
T5: Governance Challenge		E	D	D	D	D	C	C	C	C	I7
T6: Capability			E	D	D	D	D	C	C	C	I8
T7: Sector and Regulatory Context		E	D	D	D	D	C	D	C	C	I8

Narrative commentary

Themes 1 and 2 (Anchored Authority) emerged earliest in the sequence, consistent with their status as foundational conditions rather than specific practices. RAS anchoring (Theme 1) was present from Interview 1 and saturated by Interview 4, with subsequent interviews confirming rather than extending the analytical claim. Independence (Theme 2) took longer to saturate, with the strict-separation position on advisory (which is structurally a component of the independence dimension) emerging as a distinct position only during Interview 7, extending the saturation point for that theme.

Themes 3, 4, and 5 (Disciplined Engagement) emerged across the first half of the interview sequence and saturated between Interview 6 and Interview 7. The advisory boundary theme (Theme 3) was the most analytically contested, with the three-position taxonomy (permissive, strict-separation, conditional) only fully resolved once the strict-separation position emerged clearly in a later interview. The information quality theme (Theme 4) surfaced with the strongest convergence and the earliest saturation, consistent with its status as the most unanimously articulated theme in the dataset. Governance challenge (Theme 5) was developed over the longest span, with failure modes (paper closure, challenge fatigue, adversarial standoff) continuing to emerge through Interview 7 before the theme stabilised.

Themes 6 and 7 (Enabling Context) saturated latest in the sequence, consistent with their dependence on sector-specific and role-specific knowledge that was unevenly distributed across the sample. Capability (Theme 6) required input from participants with technical risk management experience before the full scope of the capability gap could be mapped. Sector and regulatory context (Theme 7) re-opened briefly at Interview 8 as multinational and cross-jurisdictional dynamics emerged from a participant with direct experience of overlapping regulatory frameworks, dynamics that the Australian-domestic participants had not articulated. Following Interview 8, both themes remained stable across Interviews 9 and 10.

Final two interviews (I9 and I10)

Across both final interviews, no new first-order codes were generated, no new clusters were identified, and no new positions on contested questions emerged. Consistent with Guest, Bunce and Johnson's (2006) observation that the majority of high-frequency themes are typically identifiable within the first six to twelve interviews in qualitative research of this kind, the final two interviews confirmed and exemplified existing analytical content rather than extending it. The saturation finding is recorded in the decision log (Appendix D5) as a formal analytical milestone.

Historical-perspective interview (P07) and comparative case (P11)

Treated separately from the primary saturation analysis and not counted in the I1 to I10 sequence, P07 and P11 were assessed against a different saturation criterion: whether each provided analytically distinct material not already present in the core dataset, rather than whether they confirmed the same themes. P07 provided the longitudinal contrast used to locate the framework in time, contributing material on pre-specialisation governance practice that was absent from the contemporary accounts. P11 provided the sectoral contrast used to test the framework's scope, contributing the structured/unstructured distinction and the comparative validation of the IA-Tickbox-Trap failure mode. Both contributed analytically distinct material in their designated roles, and neither was expected to saturate against the primary seven-theme structure. The decisions governing their analytical treatment are recorded in Appendix D5 (Decision Log, Decisions 001 and 002).

Saturation conclusion

Across the ten contemporary substantive interviews, the seven themes developed in the findings chapter reached saturation between Interview 4 (the earliest, Theme 1) and Interview 8 (the latest, Themes 6 and 7). Interviews 9 and 10 produced no new analytical content. Consistent with Saunders and Townsend's (2016) guidance that participant numbers in qualitative organisational research should be justified by analytical adequacy rather than statistical target, the sample of eleven interviews (ten substantive, one historical-perspective, one comparative case) is considered analytically sufficient for the claims advanced in this dissertation.

D5. Analytical Decision Log

Recorded in this appendix are the substantive analytical decisions made during the qualitative analysis of the eleven interview transcripts, consistent with the audit trail requirements of the Gioia methodology (Gioia, Corley and Hamilton, 2013) and the reflexive analytical practice described in Section 3.7 of the methodology chapter. Each entry documents the decision, its rationale, the alternative considered, and its implication for the analytical outputs presented in the findings and discussion chapters.

Referenced throughout the dissertation body and the preceding appendices, the decision log is the primary artefact demonstrating that the analytical choices made during this study were deliberate, reasoned, and consistently applied.

The log is presented in chronological order of analytical stage rather than participant sequence, with decisions numbered sequentially. Five decisions are recorded here as the primary decision points; minor coding decisions (individual code label changes, cluster boundary adjustments) are not individually logged but are traceable through the versioned codebook (Appendix D1, Section on codebook evolution).

Decision 001 — Treatment of P11 as comparative case

Stage of analysis: Pre-analysis, during sample design

Decision: To include P11 in the study as a deliberate comparative case rather than as a substantive participant in the primary thematic analysis, and to treat P11's account separately in the discussion chapter to contextualise the sectoral scope of the three-dimension framework.

Rationale: Carrying a financial-services audit background but having since transitioned into a non-financial-services industry, P11 occupies an analytically distinctive position in the sample. Included as a primary thematic participant, P11's account would have introduced structural noise into the codebook, since the conditions described (absence of a regulator-driven RAS, lighter documentary scaffolding, differently-shaped capability requirements) diverge materially from the conditions described by the ten contemporary financial-services participants. Treated as a comparative case, consistent with Stake's (2006) multiple-case methodology, P11's account performs a different analytical function: it tests the sectoral scope of the framework by surfacing the conditions under which the framework does not apply in its current form.

Alternative considered: Excluding P11 from the study entirely on the grounds that the non-financial-services context is beyond the stated scope of the research. This alternative was rejected because the analytical value of the comparative case, specifically its capacity to bound the framework rather than extend it, was considered to

outweigh the risk of scope creep, provided the treatment was clearly demarcated in both the methodology and the discussion chapters.

Implication for analysis: P11's transcript was coded separately against its own analytical lens, with the resulting artefact retained in the analytical audit trail but excluded from the primary seven-theme codebook structure. Deployed in Section 5.5 of the discussion chapter, P11's account is used to argue that the three-dimension framework is sector-shaped rather than universal, dependent on the regulator-driven scaffolding that distinguishes financial services from less-regulated industries.

Decision 002 — Treatment of P07 as historical-perspective interview

Stage of analysis: Pre-analysis, during sample design

Decision: To treat P07 not as a substantive participant in the primary thematic analysis but as a historical-perspective interview, used selectively in the discussion chapter to provide a longitudinal contrast that locates the framework in time.

Rationale: P07's 22-year career at a major Australian bank spans a professional era in which the contemporary architecture of IA-ERM relationships did not exist in the form the other participants describe, with finance, audit, accounting, and compliance functions less differentiated than in contemporary practice. Included in the primary thematic analysis, P07's account would have introduced historical artefacts into a framework intended to describe contemporary practice, since P07 articulates governance arrangements that contemporary participants would not recognise as analogous to their own. Treated as a longitudinal-contrast interview, P07's account performs a different function: it allows the framework to be positioned as a recent, historically-specific governance configuration rather than a timeless feature of organisational life.

Alternative considered: Treating P07 as a primary participant and adjusting the analytical framework to accommodate both historical and contemporary accounts within the same thematic structure. This alternative was rejected because the historical-contemporary boundary is analytically significant rather than merely temporal: the substantive practices at the heart of Disciplined Engagement (model risk audit, BCBS 239 lineage testing, CPS 230 readiness reviews) did not exist as auditable subjects in P07's

professional era, making direct thematic comparison between P07's account and the contemporary accounts misleading rather than informative.

Implication for analysis: P07's transcript was reviewed for analytical content consistent with the historical-perspective function and was not coded against the primary seven-theme codebook structure. Deployed selectively in Section 5.5 of the discussion chapter, P07's account is used to argue that the framework describes a recent specialisation in financial-services governance, configured around regulatory, technical, and governance developments of the past two decades, rather than as a universal model of IA-ERM relationships.

Decision 003 — Formalisation of the three-position advisory taxonomy

Stage of analysis: v6 codebook development, during second-cycle coding

Decision: To resolve the variation in participants' advisory positions into a formal three-position taxonomy (permissive, strict-separation, conditional) rather than treating it as analytical noise or collapsing it to a single dominant position.

Rationale: Encountered during second-cycle coding, the divergence between P05's permissive position, P08's strict-separation position, and the conditional position articulated by P09, P10, and P04 initially appeared to complicate the analytical claim that participants converge on an operational principle for managing the advisory-assurance boundary. On closer reading, however, the divergence was found to be analytically informative rather than analytically problematic: all three positions derive from the same underlying concern (that IA must not audit what it has designed or come to own), and they differ only in how strictly that concern is operationalised in practice. Treating the three positions as a taxonomy rather than as noise preserves this analytical structure, and allows the conditional position to be identified as most articulated without dismissing the strict-separation and permissive positions as outliers.

Alternative considered: Treating the permissive and strict-separation positions as boundary cases and presenting the conditional position as the dominant finding. This alternative was rejected as analytically dishonest: P08's strict-separation position is not a deviant case but a coherent and internally consistent stance whose rationale

(preserving future assurance) is identical to the rationale underlying the conditional position. Suppressing it would misrepresent the data.

Implication for analysis: Theme 3 in the findings chapter (Section 4.4) is structured around the three-position taxonomy, with the permissive position anchored in P05's account, the strict-separation position anchored in P08's account, and the conditional position anchored in P09, P10, and P04's accounts. The discussion chapter (Section 5.3) treats the conditional position as the most analytically developed without dismissing the other two.

Decision 004 — Exclusion of code 4.2.4 (Provisioning-data-integrity) from cross-case claims

Stage of analysis: v7 codebook development, during aggregate dimension construction

Decision: To retain code 4.2.4 (Provisioning-data-integrity) in the codebook as an exemplar code within Theme 4 but to exclude it from cross-case analytical claims, treating the provisioning case as an institution-specific exemplar rather than as a generalised pattern.

Rationale: Drawn from P05's extended account of an unexpected provisioning movement at one of his prior institutions (P05, L346-352), the provisioning case is the most operationally specific information quality example in the dataset. Including it in cross-case claims would risk overstating the generalisability of a single institutional incident, however well-documented, since the self-confirming-metric (4.2.1) and aggregation-decomposition-failure (4.2.2) failure modes were articulated independently by multiple participants and the provisioning case was not. Cross-case claims in Theme 4 therefore rest on the four-component framework codes (4.1.1 through 4.1.4) and the two convergent failure mode codes, with the provisioning case retained as an illustrative exemplar of the lineage component rather than elevated to a cross-institutional pattern.

Decision 005 — Saturation declaration following Interview 10

Stage of analysis: Post-Interview 10, during saturation analysis

Decision: To declare thematic saturation following the completion of Interview 10, on the basis that Interviews 9 and 10 produced no new first-order codes and no new positions on contested analytical questions, and to proceed to second-cycle synthesis rather than continuing to recruit additional participants.

Rationale: Consistent with the saturation logic set out in Section 3.4 of the methodology chapter and documented in the saturation analysis (Appendix D4), the decision was made that the analytical sufficiency criterion had been met following Interview 10. Across Interviews 9 and 10, all seven themes were in confirmed or saturated status, with neither interview generating new codes, new code clusters, or new positions on contested questions (particularly the advisory taxonomy and the independence dimension, which had been the last two themes to stabilise). Guest, Bunce and Johnson's (2006) empirical finding that the majority of high-frequency themes are identifiable within the first six to twelve interviews provides the methodological basis for the sufficiency claim.

Alternative considered: Continuing recruitment to twelve or fifteen participants, consistent with the upper end of the target range stated in Section 3.2 of the methodology chapter. This alternative was considered but rejected on the grounds that the saturation criterion had been met and that additional interviews were unlikely to produce new analytical content that would materially alter the seven-theme structure or the three-dimension framework. The practical access constraints described in Section 3.9 (Limitations) reinforced this decision.

Implication for analysis: The study's sample is presented throughout the dissertation as eleven interviews (ten substantive contemporary, one historical-perspective, one comparative case) rather than as a sample that fell short of a numerical target. The saturation declaration is the basis for the claim in Section 3.4 that "the sample is therefore considered sufficient to support the analytical claims advanced in the findings."

D6. Reflexive Memos

Presented in this appendix are three reflexive memos produced during the analytical process, consistent with the reflexive practice described in Section 3.8 of the methodology chapter and the commitment to researcher positionality management articulated by Berger (2015) and Bourke (2014). In qualitative research of this design, reflexive memos serve as the mechanism through which the researcher makes the interpretive process visible and accountable, documenting moments of analytical awareness, positionality risk, and interpretive turning points that shape the emerging framework. Written at the point of occurrence rather than retrospectively, each memo captures the researcher's thinking at a specific analytical moment and records the decision or adjustment that followed.

Three memos are presented here as representative samples of the reflexive practice maintained across the analytical process. Written in first person consistent with reflexive memo convention, each memo documents a moment of genuine interpretive difficulty or awareness, the nature of the positionality risk identified, and the adjustment made in response. Selected because each corresponds to an analytical decision recorded in the decision log (Appendix D5) or to an analytical claim that required positionality management before it could be advanced with confidence, the three memos together illustrate the range of reflexive work undertaken during analysis.

Memo 001 — P07 analytical stance

Date: Early analytical stage, prior to first-cycle coding of P07 transcript

Trigger: Having completed first-cycle coding of the five contemporary interviews conducted to that point, sitting down to code P07's transcript produced almost no matches against the existing framework. P07 described a world in which IA and finance were not differentiated functions, in which the RAS did not appear to exist in recognisable form, and in which the governance arrangements every contemporary participant had described as foundational were simply absent.

Positionality risk identified: The instinct to read P07 as an outlier reflected my own professional formation in contemporary financial-services governance, where the three-lines model and the IA-ERM relationship feel like architectural features of any financial institution rather than recent developments. Applied to P07's account, this contemporary lens was producing analytical noise rather than analytical insight.

Adjustment made: Suspending first-cycle coding, I returned to the transcript with a different question: what does this account tell me about governance before the contemporary framework existed? This reframing produced the longitudinal-contrast function documented in Decision 002 (Appendix D5), with P07's account read for historical texture rather than coded against the primary seven-theme structure. The absence of contemporary features became the analytical point rather than an analytical problem. A residual risk still remains that the historical contrast is framed implicitly as a deficit narrative, positioning contemporary specialisation as an improvement on what preceded it, a normative reading the data do not unambiguously support.

Memo 002 — Advisory boundary: moment of recognition

Date: v5 codebook development, during second-cycle coding of P08 transcript

Trigger: Having coded the advisory boundary material across P05, P09, P10, and P04 and arrived at what felt like a clean convergence on the conditional principle, I reached P08's advisory passage (L562-572) and found myself initially reaching for the conditional code (3.1.3) on the grounds that P08's concern about IA being "conflicted" was the same concern motivating the conditional position.

Positionality risk identified: My familiarity with the conditional position from professional experience was predisposing me to hear P08's strict-separation position as a stricter variant of the conditional position rather than as a categorically distinct stance. P08 is not saying "guide but don't own" — P08 is saying "don't engage at all." Collapsing these to the same code would have suppressed a genuine analytical divergence.

Adjustment made: After a third close reading, I concluded that the strict-separation position was categorically distinct. Code 3.1.2 (Strict-separation) was introduced during

v5 codebook development in direct response, and the three-position taxonomy was formalised at the same time, consistent with Decision 003 of the decision log (Appendix D5). Theme 3 in the findings chapter presents the divergence as analytically informative rather than suppressing it. However, a residual risk remains that my description of P08's strict-separation position as structurally motivated reflects a lingering professional bias toward the conditional position as the more analytically sophisticated stance, and the reader may judge from the findings chapter whether that framing is neutral.

Memo 003 — Trust mechanism: emergence of the "prior to substantive work" formulation

Date: v7 codebook development, during aggregate dimension construction

Trigger: Having coded the advisory boundary material across P05, P09, P10, and P04 and arrived at what felt like a clean convergence on the conditional principle, I reached P08's advisory passage (L562-572) and found myself initially reaching for the conditional code (3.1.3) on the grounds that P08's concern about IA being "conflicted" was the same concern motivating the conditional position.

Positionality risk identified: My familiarity with the conditional position from professional experience was predisposing me to hear P08's strict-separation position as a stricter variant of the conditional position rather than as a categorically distinct stance. P08 is not saying "guide but don't own" — P08 is saying "don't engage at all." Collapsing these to the same code would have suppressed a genuine analytical divergence.

Adjustment made: After a third close reading, I concluded that the strict-separation position was categorically distinct. Code 3.1.2 (Strict-separation) was introduced during v5 codebook development in direct response, and the three-position taxonomy was formalised at the same time, consistent with Decision 003 of the decision log (Appendix D5). Theme 3 in the findings chapter presents the divergence as analytically informative rather than suppressing it. A residual risk remains that the "prior to substantive work" formulation is a researcher interpretation of P05's account rather than a verbatim claim, and a more cautious reading of the same data would treat trust erosion as a specific

failure mode within governance challenge rather than as a framework-level mechanism prior to and enabling substantive work.

Appendix E — Researcher Positionality

E1. Researcher Positionality Statement

Consistent with the reflexive commitment described in Section 3.8 of the methodology chapter and the requirements identified by Berger (2015) and Bourke (2014), this positionality statement was written before the commencement of primary data analysis. Its purpose is to declare, in as specific terms as honesty permits, the researcher's relationship to the field being studied, the assumptions brought to the research, and the professional and personal stakes that may have shaped interpretation. As Bourke (2014) argues, positionality is not a one-time disclosure but an ongoing analytical consideration; this statement should therefore be read alongside the reflexive memos (Appendix D6) which document specific moments of positionality management during analysis.

Professional background and intersection with the research topic

From within the field this study examines, I came to this research. My professional background is in financial services, with experience across risk and compliance roles in the Australian market, and at the time of fieldwork I was working in a governance-adjacent capacity that brought me into regular contact with the kinds of IA-ERM relationships the research investigates. The research topic is therefore not abstract for me: Internal Audit's contribution to Enterprise Risk Management is something I have observed across institutions, formed views about, and in some cases directly experienced, long before I formulated a research question. Berger (2015) argues that the researcher's positional location shapes both what they can see and what they are likely to miss; that argument has been a daily companion throughout this study.

Prior assumptions

Entering the research, my assumptions were not neutral and it would be dishonest to claim otherwise. I expected independence to be central to participants' accounts, which proved correct. I expected capability gaps to surface as a recurring concern, which also proved correct. What I did not anticipate with any precision was the strength of the trust mechanism as an analytical finding, the degree to which the conditional advisory principle would be articulated so consistently across participants from different vantage points, or the extent to which the audit committee chair would emerge as the single most determinative actor in the enabling context dimension. Where the data confirmed my expectations, I tried to remain alert to the risk of confirmation bias, treating confirmatory accounts as evidence to be tested rather than as validation of a prior view. Where the data surprised me, I tried to take the surprise seriously rather than fitting it back into prior framing, and the three moments of analytical recognition documented in Appendix D6 each involved a prior assumption being suspended in favour of a closer reading of the data.

Relationships with participants

Several of the participants in this study are professional contacts of mine, known through industry forums, prior employment contexts, or professional networks operating across the Australian financial services sector. One is a current colleague at my employer at the time of fieldwork, and is my line manager. The existing relationships were simultaneously an enabler and a risk. As an enabler, they produced access that would have been difficult to secure from a researcher without industry standing, and they produced candour, with participants speaking more freely than they might have to a stranger. As a risk, they created the possibility that shared professional assumptions would go unprobed during interviews, that I would assume I understood what a participant meant rather than asking them to articulate it, and that participants would tailor their accounts to what they perceived I wanted to hear.

To manage these risks, the interview guide was designed around open-ended prompts that did not presuppose particular answers, and I made deliberate efforts during interviews to ask participants to elaborate on points where my professional shorthand might have led me to nod when I should have probed. The reflexive journal records several

instances where, on reviewing a transcript, I identified moments where I had assumed shared meaning rather than testing it; in those cases the transcript was returned to with particular care during coding to assess whether my interpretation matched what the participant had actually said rather than what I had assumed they meant.

Stake in the outcome

Honest acknowledgement of my stake in the outcome of this research is necessary, and it is not a small stake. Internal Audit is part of my professional world. A finding that the function contributes meaningfully to ERM governance reflects well on a field I have invested in professionally; a finding that the contribution is performative or marginal would reflect badly on it. I have tried to follow the data rather than the implication, and the framework I have advanced contains findings that are both flattering to the profession (the trust mechanism, the systemic conception of independence, the information quality contribution) and findings that are not (the capability gap, the failure modes of paper closure and challenge fatigue, the conditions under which IA degenerates into wallpaper). The discussion chapter engages directly with Power's (2009) critique of risk management as ritual rather than dismissing it, because the data partly support that critique and intellectual honesty required the engagement. Whether I have succeeded in following the data rather than my professional sympathies is a question the examiner is better placed to answer than I am; I have tried to make the interpretive trail visible enough that the question can be asked.